

SIMD Engines Comparative Reference

Konstantinos Margaritis
markos@freevec.org

June 9, 2014

Chapter 1

Introduction

Chapter 2

SIMD Engines Overview

2.1 Intel SSE family

2.1.1 Intel AVX, AVX2 & AVX512

2.2 ARM NEON (armv7 & armv8)

2.2.1 ARMv6 SIMD

2.3 PowerPC AltiVec/VMX and derivatives

2.3.1 VMX128

2.3.2 VSX

2.4 Others

2.4.1 Cell SPU

2.5 Detection

2.5.1 Compile-time Detection

2.5.2 Runtime Detection

Chapter 3

Data types

3.1 SSE and AVX

3.2 ARM NEON

3.3 PowerPC Altivec

Chapter 4

Comparing Functionality

4.1 Addition, modulo/wrap-around

Modulo/wrap-around addition is the most common amongst CPUs, where in case of an overflow, the value is wrapped around to the minimum value. In case of 8-bit unsigned bytes (unsigned char), adding 1 to 255 using modulo/wrap-around addition would produce the result 0.

4.1.1 8-bit Integers

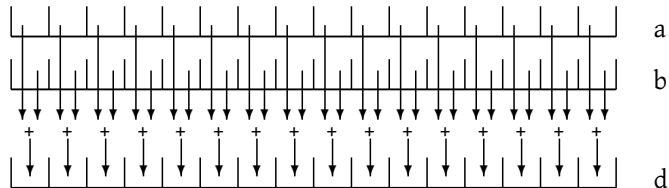


Figure 4.1: Addition of 16 8-bit integer elements in 128-bit vectors

C Intrinsics

SIMD	Code	Input	Output
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>d = _mm_add_pi8(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_add_epi8(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_add_epi8(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vadd_u8(a, b)</code>	<code>uint8x8_t a, b;</code>	<code>uint8x8_t d;</code>
NEON(128-bit)	<code>d = vaddq_u8(a, b)</code>	<code>uint8x16_t a, b;</code>	<code>uint8x16_t d;</code>
Altivec/VMX	<code>d = vec_add(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
VSX ¹	<code>d = vec_add(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
8-bit signed			
NEON(64-bit)	<code>d = vadd_s8(a, b)</code>	<code>int8x8_t a, b;</code>	<code>int8x8_t d;</code>
NEON(128-bit)	<code>d = vaddq_s8(a, b)</code>	<code>int8x16_t a, b;</code>	<code>int8x16_t d;</code>
Altivec/VMX	<code>d = vec_add(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>
VSX ²	<code>d = vec_add(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>

¹the VMX instruction will be used

²the VMX instruction will be used

Notes

Assembly

SIMD	Code	Description	Notes
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>paddb mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddb xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddb xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddb ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vadd.i8 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vaddq.i8 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>add Vd.8B, Vn.8B, Vm.8B</code>	Adds Vn.8B, Vm.8B, result in Vd.8B	d1,d2,d3 are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>add Vd.16B, Vn.16B, Vm.16B</code>	Adds Vn.16B, Vm.16B, result in Vd.16B	Vn.16B, Vm.16B, Vd.16B are 128-bit NEON registers
Altivec/VMX	<code>vaddubm vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are VMX 128-bit integer registers
VSX	<code>vaddubm vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
8-bit signed			
NEON(64-bit)	<code>vadd.s8 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vaddq.s8 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```

__m128i a = _mm_set_epi8(0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15);
__m128i b = _mm_set_epi8(15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0);
__m128i d = _mm_add_epi8(a, b);

```

NEON (128-bit)

```
uint8x16_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
uint8x16_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
uint8x16_t d = vaddq_u8(a, b);
```

Altivec/VMX

```
vector uint8_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
vector uint8_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
vector uint8_t d = vec_add(a, b);
```

4.1.2 16-bit integers

TODO: insert diagram for 16-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
16-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	d = <code>_mm_add_pi16(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	d = <code>_mm_add_epi16(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	d = <code>_mm256_add_epi16(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	d = <code>vadd_u16(a, b)</code>	<code>uint16x4_t a, b;</code>	<code>uint16x4_t d;</code>
NEON(128-bit)	d = <code>vaddq_u16(a, b)</code>	<code>uint16x8_t a, b;</code>	<code>uint16x8_t d;</code>
Altivec/VMX	d = <code>vec_add(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
VSX ³	d = <code>vec_add(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
16-bit signed			
NEON(64-bit)	d = <code>vadd_s16(a, b)</code>	<code>int16x4_t a, b;</code>	<code>int16x4_t d;</code>
NEON(128-bit)	d = <code>vaddq_s16(a, b)</code>	<code>int16x8_t a, b;</code>	<code>int16x8_t d;</code>
Altivec/VMX	d = <code>vec_add(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>
VSX ⁴	d = <code>vec_add(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>

³the VMX instruction will be used

⁴the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>paddw mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddw xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddw xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddw ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vadd.i16 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vaddq.i16 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>add Vd.4H, Vn.4H, Vm.4H</code>	Adds Vn.4H, Vm.4H, result in Vd.4H	Vn.4H, Vm.4H, Vd.4H are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>add Vd.8H, Vn.8H, Vm.8H</code>	Adds Vn.8H, Vm.8H, result in Vd.8H	Vn.8H, Vm.8H, Vd.8H are 128-bit NEON registers
Altivec/VMX	<code>vadduhm vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vadduhm vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
8-bit signed			
NEON(64-bit)	<code>vadd.s16 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vaddq.s16 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi16(0, 1, 2, 3, 4, 5, 6, 7);
__m128i b = _mm_set_epi16(7, 6, 5, 4, 3, 2, 1, 0);
__m128i d = _mm_add_epi16(a, b);
```

NEON (128-bit)

```
uint16x8_t a = {0, 1, 2, 3, 4, 5, 6, 7};
uint16x8_t b = {7, 6, 5, 4, 3, 2, 1, 0};
uint16x8_t d = vaddq_u16(a, b);
```

Altivec/VMX

```
vector uint16_t a = {0, 1, 2, 3, 4, 5, 6, 7};  
vector uint16_t b = {7, 6, 5, 4, 3, 2, 1, 0};  
vector uint16_t d = vec_add(a, b);
```

4.1.3 32-bit integers

TODO: insert diagram for 32-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	d = <code>_mm_add_pi32(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	d = <code>_mm_add_epi32(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	d = <code>_mm256_add_epi32(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	d = <code>vadd_u32(a, b)</code>	<code>uint32x2_t a, b;</code>	<code>uint32x2_t d;</code>
NEON(128-bit)	d = <code>vaddq_u32(a, b)</code>	<code>uint32x4_t a, b;</code>	<code>uint32x4_t d;</code>
Altivec/VMX	d = <code>vec_add(a, b)</code>	<code>vector uint32_t a, b;</code>	<code>vector uint32_t d;</code>
VSX ⁵	d = <code>vec_add(a, b)</code>	<code>vector uint32_t a, b;</code>	<code>vector uint32_t d;</code>
8-bit signed			
NEON(64-bit)	d = <code>vadd_s32(a, b)</code>	<code>int32x2_t a, b;</code>	<code>int32x2_t d;</code>
NEON(128-bit)	d = <code>vaddq_s32(a, b)</code>	<code>int32x4_t a, b;</code>	<code>int32x4_t d;</code>
Altivec/VMX	d = <code>vec_add(a, b)</code>	<code>vector int32_t a, b;</code>	<code>vector int32_t d;</code>
VSX ⁶	d = <code>vec_add(a, b)</code>	<code>vector int32_t a, b;</code>	<code>vector int32_t d;</code>

⁵the VMX instruction will be used

⁶the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>padd mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>padd xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpadd xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpadd ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vadd.i32 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vaddq.i32 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>add Vd.2S, Vn.2S, Vm.2S</code>	Adds Vn.2S, Vm.2S, result in Vd.2S	Vn.2S, Vm.2S, Vd.2S are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>add Vd.4S, Vn.4S, Vm.4S</code>	Adds Vn.4S, Vm.4S, result in Vd.4S	Vn.4S, Vm.4S, Vd.4S are 128-bit NEON registers
Altivec/VMX	<code>vadduwm vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vadduwm vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
8-bit signed			
NEON(64-bit)	<code>vadd.s32 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vaddq.s32 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi32(0, 1, 2, 3);
__m128i b = _mm_set_epi32(3, 2, 1, 0);
__m128i d = _mm_add_epi32(a, b);
```

NEON (128-bit)

```
uint32x4_t a = {0, 1, 2, 3};
uint32x4_t b = {3, 2, 1, 0};
uint32x4_t d = vaddq_u32(a, b);
```

Altivec/VMX

```
vector uint32_t a = {0, 1, 2, 3};  
vector uint32_t b = {3, 2, 1, 0};  
vector uint32_t d = vec_add(a, b);
```

4.1.4 64-bit integers

TODO: insert diagram for 64-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	d = <code>_mm_add_si64(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	d = <code>_mm_add_epi64(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	d = <code>_mm256_add_epi64(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	d = <code>vadd_u64(a, b)</code>	<code>uint64x1_t a, b;</code>	<code>uint64x1_t d;</code>
NEON(128-bit)	d = <code>vaddq_u64(a, b)</code>	<code>uint64x2_t a, b;</code>	<code>uint64x2_t d;</code>
Altivec/VMX ⁷	N/A	-	-
VSX	d = <code>vec_add(a, b)</code>	<code>vector uint64_t a, b;</code>	<code>vector uint64_t d;</code>
8-bit signed			
NEON(64-bit)	d = <code>vadd_s64(a, b)</code>	<code>int64x1_t a, b;</code>	<code>int64x1_t d;</code>
NEON(128-bit)	d = <code>vaddq_s64(a, b)</code>	<code>int64x2_t a, b;</code>	<code>int64x2_t d;</code>
VSX	d = <code>vec_add(a, b)</code>	<code>vector int64_t a, b;</code>	<code>vector int64_t d;</code>

Assembly

SIMD	Code	Description	Notes
8-bit unsigned/signed			
MMX(64-bit)	<code>paddq mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddq xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddq xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddq ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vadd.i64 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vaddq.i64 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>add Dd, Dn, Dm</code>	Adds Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>add Vd.2D, Vn.2D, Vm.2D</code>	Adds Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers
Altivec/VMX	N/A	-	-
VSX	<code>vadduqm vD, vA, vB</code>	Adds vA, vB, result in vD	VSX extension
8-bit signed			
NEON(64-bit)	<code>vadd.s64 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vaddq.s64 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers

⁷Altivec/VMX does not support 64-bit integer arithmetic

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi64(0, 1);  
__m128i b = _mm_set_epi64(3, 2);  
__m128i d = _mm_add_epi64(a, b);
```

AVX (256-bit)

```
__m256i a = _mm256_set_epi64(0, 1, 2, 3);  
__m256i b = _mm256_set_epi64(3, 2, 1, 0);  
__m256i d = _mm256_add_epi64(a, b);
```

NEON (128-bit)

```
uint64x2_t a = {0, 1};  
uint64x2_t b = {3, 2};  
uint64x2_t d = vaddq_u64(a, b);
```

VSX

```
vector uint64_t a = {0, 1};  
vector uint64_t b = {1, 0};  
vector uint64_t d = vec_add(a, b);
```

4.1.5 32-bit floats

TODO: insert diagram for 32-bit float vector addition

C Intrinsics

SIMD	Code	Input	Output
MMX(64-bit) ⁸	N/A	-	-
SSE(64-bit)	<code>d = _mm_add_ss(a, b)</code>	<code>__m128 a, b;</code>	<code>__m128 d;</code>
SSE(128-bit)	<code>d = _mm_add_ps(a, b)</code>	<code>__m128 a, b;</code>	<code>__m128 d;</code>
AVX(256-bit)	<code>d = _mm256_add_ps(a, b)</code>	<code>__m256 a, b;</code>	<code>__m256 d;</code>
NEON(64-bit)	<code>d = vadd_f32(a, b)</code>	<code>float32x2_t a, b;</code>	<code>float32x2_t d;</code>
NEON(128-bit)	<code>d = vaddq_f32(a, b)</code>	<code>float32x4_t a, b;</code>	<code>float32x4_t d;</code>
Altivec/VMX	<code>d = vec_add(a, b)</code>	<code>vector float a, b;</code>	<code>vector float d;</code>
VSX ⁹	<code>d = vec_add(a, b)</code>	<code>vector float a, b;</code>	<code>vector float d;</code>

⁸MMX does not support float arithmetic

⁹the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
MMX(64-bit)	N/A	-	-
SSE(64-bit)	<code>addss xmm1, xmm2</code>	Adds <code>xmm1, xmm2</code> , result in <code>xmm1</code>	<code>xmm1, xmm2</code> are SSE 128-bit double registers, but the operation takes place only at the lower 64-bits
SSE(128-bit)	<code>addps xmm1, xmm2</code>	Adds <code>xmm1, xmm2</code> , result in <code>xmm1</code>	<code>xmm1, xmm2</code> are SSE 128-bit float registers
AVX(128-bit)	<code>addps xmm1, xmm2, xmm3</code>	Adds <code>xmm2, xmm3</code> , result <code>xmm1</code>	<code>xmm1, xmm2, xmm3</code> are AVX 128-bit registers
AVX(256-bit)	<code>addps ymm1, ymm2, ymm3</code>	Adds <code>ymm2, ymm3</code> , result in <code>ymm1</code>	<code>ymm1, ymm2, ymm3</code> are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vadd.f32 d1, d2, d3</code>	Adds <code>d2, d3</code> , result in <code>d1</code>	<code>d1, d2, d3</code> are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vaddq.f32 q1, q2, q3</code>	Adds <code>q2, q3</code> , result in <code>q1</code>	<code>q1, q2, q3</code> are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>fadd Vd.2S, Vn.2S, Vm.2S</code>	Adds <code>Vn.2S, Vm.2S</code> , result in <code>Vd.2S</code>	<code>Vn.2S, Vm.2S, Vd.2S</code> are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>fadd Vd.4S, Vn.4S, Vm.4S</code>	Adds <code>Vn.4S, Vm.4S</code> , result in <code>Vd.4S</code>	<code>Vn.4S, Vm.4S, Vd.4S</code> are 128-bit NEON registers
Altivec/VMX	<code>vaddfp vD, vA, vB</code>	Adds <code>vA, vB</code> , result in <code>vD</code>	<code>vA, vB, vD</code> are normal VMX 128-bit registers
VSX	<code>vaddfp vD, vA, vB</code>	Adds <code>vA, vB</code> , result in <code>vD</code>	VMX instruction
VSX	<code>xsaddsp vD, vA, vB</code>	Adds <code>vA, vB</code> , result in <code>vD</code>	<code>vA, vB, vD</code> are VSX registers holding double values, but the operation takes place only at the lower 64-bits (FIXME: better description)

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```

__m128 a = _mm_set_ps(1.0, 2.0, 3.0, 4.0);
__m128 b = _mm_set_ps(4.0, 3.0, 2.0, 1.0);
__m128 d = _mm_add_ps(a, b);

```

AVX (256-bit)

```
__m256 a = _mm256_set_ps(1.0, 2.0, 3.0, 4.0);  
__m256 b = _mm256_set_ps(4.0, 3.0, 2.0, 1.0);  
__m256 d = _mm256_add_ps(a, b);
```

NEON (128-bit)

```
float32x4_t a = {0, 1, 2, 3};  
float32x4_t b = {3, 2, 1, 0};  
float32x4_t d = vaddq_f32(a, b);
```

AltiVec/VMX

```
vector float a = {0, 1, 2, 3};  
vector float b = {3, 2, 1, 0};  
vector float d = vec_add(a, b);
```

4.1.6 64-bit doubles

TODO: insert diagram for 64-bit float vector addition

C Intrinsics

SIMD	Code	Input	Output
MMX(64-bit) ¹⁰	N/A	-	-
SSE2(64-bit)	d = <code>_mm_add_sd(a, b)</code>	<code>__m128d a, b;</code>	<code>__m128d d;</code>
SSE2(128-bit)	d = <code>_mm_add_pd(a, b)</code>	<code>__m128d a, b;</code>	<code>__m128d d;</code>
AVX(256-bit)	d = <code>_mm256_add_pd(a, b)</code>	<code>__m256d a, b;</code>	<code>__m256d d;</code>
NEON(64-bit) ¹¹	d = <code>vadd_f64(a, b)</code>	<code>float64x1_t a, b;</code>	<code>float64x1_t d;</code>
NEON(128-bit)	d = <code>vaddq_f64(a, b)</code>	<code>float64x2_t a, b;</code>	<code>float64x2_t d;</code>
Altivec/VMX ¹²	N/A	-	-
VSX	d = <code>vec_add(a, b)</code>	<code>vector double a, b;</code>	<code>vector double d;</code>

¹⁰MMX does not support double arithmetic

¹¹double arithmetic is supported only by armv8 NEON

¹²Altivec/VMX does not support double arithmetic

Assembly

SIMD	Code	Description	Notes
MMX(64-bit)	N/A	–	–
SSE2(64-bit)	<code>addsd xmm1, xmm2</code>	Adds <code>xmm1, xmm2</code> , result in <code>xmm1</code>	<code>xmm1, xmm2</code> are SSE 128-bit double registers, but the operation takes place only at the lower 64-bits
SSE2(128-bit)	<code>addpd xmm1, xmm2</code>	Adds <code>xmm1, xmm2</code> , result in <code>xmm1</code>	<code>xmm1, xmm2</code> are SSE 128-bit float registers
AVX(128-bit)	<code>addpd xmm1, xmm2, xmm3</code>	Adds <code>xmm2, xmm3</code> , result <code>xmm1</code>	<code>xmm1, xmm2, xmm3</code> are AVX 128-bit registers
AVX(256-bit)	<code>addpd ymm1, ymm2, ymm3</code>	Adds <code>ymm2, ymm3</code> , result in <code>ymm1</code>	<code>ymm1, ymm2, ymm3</code> are AVX 256-bit registers
NEON(armv7, 64-bit)	N/A	–	–
NEON(armv7, 128-bit)	N/A	–	–
NEON(armv8, 64-bit)	<code>fadd Dd, Dn, Dm</code>	Adds <code>Dn, Dm</code> , result in <code>Dd</code>	<code>Dn, Dm, Dd</code> are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>fadd Vd.2D, Vn.2D, Vm.2D</code>	Adds <code>Vn.2D, Vm.2D</code> , result in <code>Vd.2D</code>	<code>Vn.2D, Vm.2D, Vd.2D</code> are 128-bit NEON registers
Altivec/VMX	N/A	–	–
VSX	<code>xsaddpd vD, vA, vB</code>	Adds <code>vA, vB</code> , result in <code>vD</code>	<code>vA, vB, vD</code> are VSX registers holding double values, but the operation takes place only at the lower 64-bits
VSX	<code>xvaddpd vD, vA, vB</code>	Adds <code>vA, vB</code> , result in <code>vD</code>	<code>vA, vB, vD</code> are VSX registers holding double values

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```

__m128d a = _mm_set_pd(1.0, 2.0, 3.0, 4.0);
__m128d b = _mm_set_pd(4.0, 3.0, 2.0, 1.0);
__m128d d = _mm_add_pd(a, b);

```

AVX (256-bit)

```
__m256d a = _mm256_set_pd(1.0, 2.0, 3.0, 4.0);  
__m256d b = _mm256_set_pd(4.0, 3.0, 2.0, 1.0);  
__m256d d = _mm256_add_pd(a, b);
```

NEON (128-bit)

```
float64x2_t a = {0, 1};  
float64x2_t b = {3, 2};  
float64x2_t d = vaddq_f64(a, b);
```

VSX

```
vector double a = {0, 1, 2, 3};  
vector double b = {3, 2, 1, 0};  
vector double d = vec_add(a, b);
```

4.2 Addition, saturated

Saturated addition in contrast to modulo addition, keeps the result in a fixed range between a minimum and maximum value. If it's greater than the maximum value, then it's "clamped" to the maximum value, if it's below the minimum then it's set to the minimum value. In case of 8-bit unsigned bytes (unsigned char), adding 10 to 255 using saturated addition would produce the result 255.

4.2.1 8-bit Integers

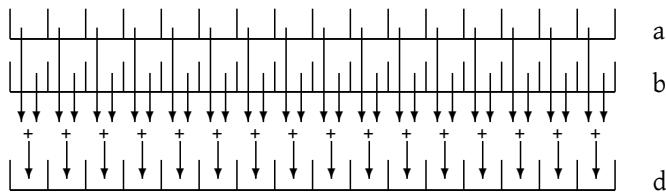


Figure 4.2: Addition of 16 8-bit integer elements in 128-bit vectors

C Intrinsics

SIMD	Code	Input	Output
8-bit unsigned			
MMX(64-bit)	<code>d = _mm_adds_pu8(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_adds_epu8(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_add_epu8(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vaddsq_u8(a, b)</code>	<code>uint8x8_t a, b;</code>	<code>uint8x8_t d;</code>
NEON(128-bit)	<code>d = vaddsq_u8(a, b)</code>	<code>uint8x16_t a, b;</code>	<code>uint8x16_t d;</code>
Altivec/VMX	<code>d = vec_adds(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
VSX ¹³	<code>d = vec_adds(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
8-bit signed			
MMX(64-bit)	<code>d = _mm_adds_pi8(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_adds_epi8(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_adds_epi8(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vqadd_s8(a, b)</code>	<code>int8x8_t a, b;</code>	<code>int8x8_t d;</code>
NEON(128-bit)	<code>d = vqaddq_s8(a, b)</code>	<code>int8x16_t a, b;</code>	<code>int8x16_t d;</code>
Altivec/VMX	<code>d = vec_adds(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>
VSX ¹⁴	<code>d = vec_adds(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>

¹³the VMX instruction will be used

¹⁴the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
8-bit unsigned			
MMX(64-bit)	<code>paddusb mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddusb xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddusb xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddusb ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vqadd.i8 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqaddq.i8 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqadd Vd.8B, Vn.8B, Vm.8B</code>	Adds Vn.8B, Vm.8B, result in Vd.8B	d1,d2,d3 are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqadd Vd.16B, Vn.16B, Vm.16B</code>	Adds Vn.16B, Vm.16B, result in Vd.16B	Vn.16B, Vm.16B, Vd.16B are 128-bit NEON registers
Altivec/VMX	<code>vaddubs vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vaddubs vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
8-bit signed			
MMX(64-bit)	<code>paddsb mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddsb xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddsb xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddsb ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vqadd.s8 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqaddq.s8 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqadd Vd.8B, Vn.8B, Vm.8B</code>	Adds Vn.8B, Vm.8B, result in Vd.8B	d1,d2,d3 are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqadd Vd.16B, Vn.16B, Vm.16B</code>	Adds Vn.16B, Vm.16B, result in Vd.16B	Vn.16B, Vm.16B, Vd.16B are 128-bit NEON registers
Altivec/VMX	<code>vaddsb vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vaddsb vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epu8(0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15);  
__m128i b = _mm_set_epu8(15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0);  
__m128i d = _mm_adds_epu8(a, b);
```

NEON (128-bit)

```
int8x16_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
int8x16_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
int8x16_t d = vqaddq_s8(a, b);
```

AltiVec/VMX

```
vector uint8_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
vector uint8_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
vector uint8_t d = vec_adds(a, b);
```

4.2.2 16-bit integers

TODO: insert diagram for 16-bit saturated vector addition

C Intrinsics

SIMD	Code	Input	Output
16-bit unsigned			
MMX(64-bit)	<code>d = _mm_adds_pu16(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_adds_epu16(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_adds_epu16(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vqadd_u16(a, b)</code>	<code>uint16x4_t a, b;</code>	<code>uint16x4_t d;</code>
NEON(128-bit)	<code>d = vqaddq_u16(a, b)</code>	<code>uint16x8_t a, b;</code>	<code>uint16x8_t d;</code>
AltiVec/VMX	<code>d = vec_adds(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
VSX ¹⁵	<code>d = vec_adds(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
16-bit signed			
MMX(64-bit)	<code>d = _mm_adds_pi16(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_adds_epi16(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_adds_epi16(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vqadd_s16(a, b)</code>	<code>int16x4_t a, b;</code>	<code>int16x4_t d;</code>
NEON(128-bit)	<code>d = vqaddq_s16(a, b)</code>	<code>int16x8_t a, b;</code>	<code>int16x8_t d;</code>
AltiVec/VMX	<code>d = vec_adds(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>
VSX ¹⁶	<code>d = vec_adds(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>

¹⁵the VMX instruction will be used

¹⁶the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
16-bit unsigned			
MMX(64-bit)	<code>paddsw mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddsw xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddsw xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddsw ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vqadd.i16 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqaddq.i16 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqadd Vd.4H, Vn.4H, Vm.4H</code>	Adds Vn.4H, Vm.4H, result in Vd.4H	Vn.4H, Vm.4H, Vd.4H are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqadd Vd.8H, Vn.8H, Vm.8H</code>	Adds Vn.8H, Vm.8H, result in Vd.8H	Vn.8H, Vm.8H, Vd.8H are 128-bit NEON registers
Altivec/VMX	<code>vadduhm vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vadduhm vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
16-bit signed			
MMX(64-bit)	<code>paddsw mm1, mm2</code>	Adds mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>paddsw xmm1, xmm2</code>	Adds xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpaddsw xmm1, xmm2, xmm3</code>	Adds xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpaddsw ymm1, ymm2, ymm3</code>	Adds ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(64-bit)	<code>vqadd.s16 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vqaddq.s16 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqadd Vd.4H, Vn.4H, Vm.4H</code>	Adds Vn.4H, Vm.4H, result in Vd.4H	Vn.4H, Vm.4H, Vd.4H are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqadd Vd.8H, Vn.8H, Vm.8H</code>	Adds Vn.8H, Vm.8H, result in Vd.8H	Vn.8H, Vm.8H, Vd.8H are 128-bit NEON registers
Altivec/VMX	<code>vaddshM vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vaddshM vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi16(0, 1, 2, 3, 4, 5, 6, 7);  
__m128i b = _mm_set_epi16(7, 6, 5, 4, 3, 2, 1, 0);  
__m128i d = _mm_adds_epi16(a, b);
```

NEON (128-bit)

```
int16x8_t a = {0, 1, 2, 3, 4, 5, 6, 7};  
int16x8_t b = {7, 6, 5, 4, 3, 2, 1, 0};  
int16x8_t d = vqaddq_s16(a, b);
```

AltiVec/VMX

```
vector uint16_t a = {0, 1, 2, 3, 4, 5, 6, 7};  
vector uint16_t b = {7, 6, 5, 4, 3, 2, 1, 0};  
vector uint16_t d = vec_adds(a, b);
```


4.2.3 32-bit integers

TODO: insert diagram for 32-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
32-bit unsigned			
MMX(64-bit) ¹⁷	N/A	-	-
SSE2(128-bit) ¹⁸	N/A	-	-
AVX(256-bit) ¹⁹	N/A	-	-
NEON(64-bit)	d = vqadd_u32(a, b)	uint32x2_t a, b;	uint32x2_t d;
NEON(128-bit)	d = vqaddq_u32(a, b)	uint32x4_t a, b;	uint32x4_t d;
Altivec/VMX	d = vec_adds(a, b)	vector uint32_t a, b;	vector uint32_t d;
VSX ²⁰	d = vec_adds(a, b)	vector uint32_t a, b;	vector uint32_t d;
32-bit signed			
MMX(64-bit)	N/A	-	-
SSE2(128-bit)	N/A	-	-
AVX(256-bit)	N/A	-	-
NEON(64-bit)	d = vqadd_s32(a, b)	int32x2_t a, b;	int32x2_t d;
NEON(128-bit)	d = vqaddq_s32(a, b)	int32x4_t a, b;	int32x4_t d;
Altivec/VMX	d = vec_adds(a, b)	vector int32_t a, b;	vector int32_t d;
VSX	d = vec_adds(a, b)	vector int32_t a, b;	vector int32_t d;

¹⁷MMX has no support for saturated addition on 32-bit integers

¹⁸SSE has no support for saturated addition on 32-bit integers

¹⁹AVX has no support for saturated addition on 32-bit integers

²⁰the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
32-bit unsigned			
MMX(64-bit)	N/A	–	–
SSE2(128-bit)	N/A	–	–
AVX(128-bit)	N/A	–	–
AVX(256-bit)	N/A	–	–
NEON(64-bit)	<code>vqadd.i32 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vqaddq.i32 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqadd Dd, Dn, Dm</code>	Adds Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqadd Vd.2D, Vn.2D, Vm.2D</code>	Adds Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers
Altivec/VMX	<code>vadduws vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vadduws vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
32-bit signed			
MMX(64-bit)	N/A	–	–
SSE2(128-bit)	N/A	–	–
AVX(128-bit)	N/A	–	–
AVX(256-bit)	N/A	–	–
NEON(64-bit)	<code>vqadd.s32 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vqaddq.s32 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqadd Dd, Dn, Dm</code>	Adds Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqadd Vd.2D, Vn.2D, Vm.2D</code>	Adds Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers
Altivec/VMX	<code>vaddsws vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vaddsws vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

NEON (128-bit)

```
uint32x4_t a = {0, 1, 2, 3};
uint32x4_t b = {3, 2, 1, 0};
uint32x4_t d = vaddq_u32(a, b);
```

Altivec/VMX

```
vector uint32_t a = {0, 1, 2, 3};  
vector uint32_t b = {3, 2, 1, 0};  
vector uint32_t d = vec_adds(a, b);
```

4.2.4 64-bit integers

At the moment the only architecture that supports 64-bit saturated addition is NEON, in both ARMv7 and ARMv8 implementations.

TODO: insert diagram for 64-bit vector saturated addition

C Intrinsics

SIMD	Code	Input	Output
64-bit unsigned			
NEON(armv7, 64-bit)	<code>d = vqadd_u64(a, b)</code>	<code>uint64x1_t a, b;</code>	<code>uint64x1_t d;</code>
NEON(armv7, 128-bit)	<code>d = vqaddq_u64(a, b)</code>	<code>uint64x2_t a, b;</code>	<code>uint64x2_t d;</code>
64-bit signed			
NEON(64-bit)	<code>d = vqadd_s64(a, b)</code>	<code>int64x1_t a, b;</code>	<code>int64x1_t d;</code>
NEON(128-bit)	<code>d = vqaddq_s64(a, b)</code>	<code>int64x2_t a, b;</code>	<code>int64x2_t d;</code>

Assembly

SIMD	Code	Description	Notes
64-bit unsigned			
NEON(armv7, 64-bit)	<code>vqadd.i64 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqaddq.i64 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqadd Dd, Dn, Dm</code>	Adds Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqadd Vd.2D, Vn.2D, Vm.2D</code>	Adds Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers
64-bit signed			
NEON(64-bit)	<code>vqadd.s64 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vqaddq.s64 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqadd Dd, Dn, Dm</code>	Adds Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqadd Vd.2D, Vn.2D, Vm.2D</code>	Adds Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

NEON (128-bit)

```
uint64x2_t a = {0, 1};  
uint64x2_t b = {3, 2};  
uint64x2_t d = vaddq_u64(a, b);
```

4.3 Subtraction, modulo/wrap-around

FIXME: Modulo/wrap-around addition is the most common amongst CPUs, where in case of an overflow, the value is wrapped around to the minimum value. In case of 8-bit unsigned bytes (unsigned char), adding 1 to 255 using modulo/wrap-around addition would produce the result 0.

4.3.1 8-bit Integers

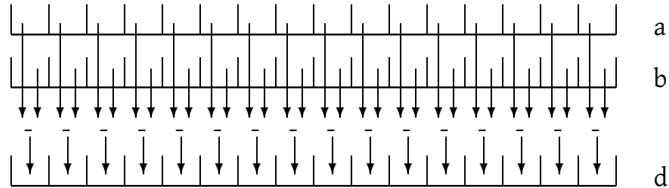


Figure 4.3: Addition of 16 8-bit integer elements in 128-bit vectors

C Intrinsics

SIMD	Code	Input	Output
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>d = _mm_sub_pi8(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_sub_epi8(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_sub_epi8(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vsub_u8(a, b)</code>	<code>uint8x8_t a, b;</code>	<code>uint8x8_t d;</code>
NEON(128-bit)	<code>d = vsubq_u8(a, b)</code>	<code>uint8x16_t a, b;</code>	<code>uint8x16_t d;</code>
Altivec/VMX	<code>d = vec_sub(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
VSX ²¹	<code>d = vec_sub(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
8-bit signed			
NEON(64-bit)	<code>d = vsub_s8(a, b)</code>	<code>int8x8_t a, b;</code>	<code>int8x8_t d;</code>
NEON(128-bit)	<code>d = vsubq_s8(a, b)</code>	<code>int8x16_t a, b;</code>	<code>int8x16_t d;</code>
Altivec/VMX	<code>d = vec_sub(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>
VSX ²²	<code>d = vec_sub(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>

²¹the VMX instruction will be used

²²the VMX instruction will be used

Notes**Assembly**

SIMD	Code	Description	Notes
8-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>psubb mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubb xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubb xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubb ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(64-bit)	<code>vsub.i8 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.i8 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sub Vd.8B, Vn.8B, Vm.8B</code>	Subtracts Vn.8B, Vm.8B, result in Vd.8B	d1,d2,d3 are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sub Vd.16B, Vn.16B, Vm.16B</code>	Subtracts Vn.16B, Vm.16B, result in Vd.16B	Vn.16B, Vm.16B, Vd.16B are 128-bit NEON registers
Altivec/VMX	<code>vsububm vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are VMX 128-bit integer registers
VSX	<code>vsububm vD, vA, vB</code>	Subtracts vA, vB, result in vD	VMX instruction
8-bit signed			
NEON(64-bit)	<code>vsub.s8 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.s8 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers

Notes**Examples**

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```

__m128i a = _mm_set_epi8(0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15);
__m128i b = _mm_set_epi8(15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0);
__m128i d = _mm_sub_epi8(a, b);

```

NEON (128-bit)

```
uint8x16_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
uint8x16_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
uint8x16_t d = vsubq_u8(a, b);
```

Altivec/VMX

```
vector uint8_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
vector uint8_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
vector uint8_t d = vec_sub(a, b);
```


4.3.2 16-bit integers

TODO: insert diagram for 16-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
16-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	d = <code>_mm_sub_pi16(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	d = <code>_mm_sub_epi16(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	d = <code>_mm256_sub_epi16(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	d = <code>vsub_u16(a, b)</code>	<code>uint16x4_t a, b;</code>	<code>uint16x4_t d;</code>
NEON(128-bit)	d = <code>vsubq_u16(a, b)</code>	<code>uint16x8_t a, b;</code>	<code>uint16x8_t d;</code>
Altivec/VMX	d = <code>vec_sub(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
VSX ²³	d = <code>vec_sub(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
16-bit signed			
NEON(64-bit)	d = <code>vsub_s16(a, b)</code>	<code>int16x4_t a, b;</code>	<code>int16x4_t d;</code>
NEON(128-bit)	d = <code>vsubq_s16(a, b)</code>	<code>int16x8_t a, b;</code>	<code>int16x8_t d;</code>
Altivec/VMX	d = <code>vec_sub(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>
VSX ²⁴	d = <code>vec_sub(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>

²³the VMX instruction will be used

²⁴the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
16-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>psubw mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubw xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubw xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubw ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(64-bit)	<code>vsub.i16 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1, d2, d3 are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.i16 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1, q2, q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sub Vd.4H, Vn.4H, Vm.4H</code>	Subtracts Vn.4H, Vm.4H, result in Vd.4H	Vn.4H, Vm.4H, Vd.4H are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sub Vd.8H, Vn.8H, Vm.8H</code>	Subtracts Vn.8H, Vm.8H, result in Vd.8H	Vn.8H, Vm.8H, Vd.8H are 128-bit NEON registers
Altivec/VMX	<code>vsubuhm vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubuhm vD, vA, vB</code>	Subtracts vA, vB, result in vD	VMX instruction
16-bit signed			
NEON(64-bit)	<code>vsub.s16 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1, d2, d3 are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.s16 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1, q2, q3 are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi16(0, 1, 2, 3, 4, 5, 6, 7);
__m128i b = _mm_set_epi16(7, 6, 5, 4, 3, 2, 1, 0);
__m128i d = _mm_sub_epi16(a, b);
```

NEON (128-bit)

```
uint16x8_t a = {0, 1, 2, 3, 4, 5, 6, 7};
uint16x8_t b = {7, 6, 5, 4, 3, 2, 1, 0};
uint16x8_t d = vsubq_u16(a, b);
```

Altivec/VMX

```
vector uint16_t a = {0, 1, 2, 3, 4, 5, 6, 7};  
vector uint16_t b = {7, 6, 5, 4, 3, 2, 1, 0};  
vector uint16_t d = vec_sub(a, b);
```

4.3.3 32-bit integers

TODO: insert diagram for 32-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
32-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	d = <code>_mm_sub_pi32(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	d = <code>_mm_sub_epi32(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	d = <code>_mm256_sub_epi32(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	d = <code>vsub_u32(a, b)</code>	<code>uint32x2_t a, b;</code>	<code>uint32x2_t d;</code>
NEON(128-bit)	d = <code>vsubq_u32(a, b)</code>	<code>uint32x4_t a, b;</code>	<code>uint32x4_t d;</code>
Altivec/VMX	d = <code>vec_sub(a, b)</code>	<code>vector uint32_t a, b;</code>	<code>vector uint32_t d;</code>
VSX ²⁵	d = <code>vec_sub(a, b)</code>	<code>vector uint32_t a, b;</code>	<code>vector uint32_t d;</code>
32-bit signed			
NEON(64-bit)	d = <code>vsub_s32(a, b)</code>	<code>int32x2_t a, b;</code>	<code>int32x2_t d;</code>
NEON(128-bit)	d = <code>vsubq_s32(a, b)</code>	<code>int32x4_t a, b;</code>	<code>int32x4_t d;</code>
Altivec/VMX	d = <code>vec_sub(a, b)</code>	<code>vector int32_t a, b;</code>	<code>vector int32_t d;</code>
VSX ²⁶	d = <code>vec_sub(a, b)</code>	<code>vector int32_t a, b;</code>	<code>vector int32_t d;</code>

²⁵the VMX instruction will be used

²⁶the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
32-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>psubd mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubd xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubd xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubd ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(64-bit)	<code>vsub.i32 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1, d2, d3 are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.i32 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1, q2, q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sub Vd.2S, Vn.2S, Vm.2S</code>	Subtracts Vn.2S, Vm.2S, result in Vd.2S	Vn.2S, Vm.2S, Vd.2S are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sub Vd.4S, Vn.4S, Vm.4S</code>	Subtracts Vn.4S, Vm.4S, result in Vd.4S	Vn.4S, Vm.4S, Vd.4S are 128-bit NEON registers
Altivec/VMX	<code>vsubuwm vD, vA, vB</code>	Subtracts vB from vA, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubuwm vD, vA, vB</code>	Subtracts vB from vA, result in vD	VSX instruction
32-bit signed			
NEON(64-bit)	<code>vsub.s32 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1, d2, d3 are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.s32 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1, q2, q3 are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi32(0, 1, 2, 3);
__m128i b = _mm_set_epi32(3, 2, 1, 0);
__m128i d = _mm_sub_epi32(a, b);
```

NEON (128-bit)

```
uint32x4_t a = {0, 1, 2, 3};
uint32x4_t b = {3, 2, 1, 0};
uint32x4_t d = vsubq_u32(a, b);
```

Altivec/VMX

```
vector uint32_t a = {0, 1, 2, 3};  
vector uint32_t b = {3, 2, 1, 0};  
vector uint32_t d = vec_sub(a, b);
```

4.3.4 64-bit integers

TODO: insert diagram for 64-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
64-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>d = _mm_sub_si64(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_sub_epi64(a, b)</code>	<code>_m128i a, b;</code>	<code>_m128i d;</code>
AVX(256-bit)	<code>d = _mm256_sub_epi64(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vsub_u64(a, b)</code>	<code>uint64x1_t a, b;</code>	<code>uint64x1_t d;</code>
NEON(128-bit)	<code>d = vsubq_u64(a, b)</code>	<code>uint64x2_t a, b;</code>	<code>uint64x2_t d;</code>
AltiVec/VMX ²⁷	N/A	-	-
VSX	<code>d = vec_sub(a, b)</code>	<code>vector uint64_t a, b;</code>	<code>vector uint64_t d;</code>
64-bit signed			
NEON(64-bit)	<code>d = vsub_u64(a, b)</code>	<code>uint64x1_t a, b;</code>	<code>uint64x1_t d;</code>
NEON(128-bit)	<code>d = vsubq_u64(a, b)</code>	<code>uint64x2_t a, b;</code>	<code>uint64x2_t d;</code>
VSX	<code>d = vec_sub(a, b)</code>	<code>vector int64_t a, b;</code>	<code>vector int64_t d;</code>

Assembly

SIMD	Code	Description	Notes
64-bit unsigned/signed (unless specified otherwise)			
MMX(64-bit)	<code>psubq mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubq xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubq xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result in xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubq ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vsub.i64 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1, d2, d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vsubq.i64 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1, q2, q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sub Dd, Dn, Dm</code>	Subtracts Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sub Vd.2D, Vn.2D, Vm.2D</code>	Subtracts Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers
AltiVec/VMX	N/A	-	-
VSX	<code>vsubuqm vD, vA, vB</code>	Subtracts vA, vB, result in vD	VSX extension
64-bit signed			
NEON(armv7, 64-bit)	<code>vsub.s64 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1, d2, d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vsubq.s64 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1, q2, q3 are 128-bit NEON registers

²⁷ AltiVec/VMX does not support 64-bit integer arithmetic

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi64(0, 1);  
__m128i b = _mm_set_epi64(3, 2);  
__m128i d = _mm_sub_epi64(a, b);
```

AVX (256-bit)

```
__m256i a = _mm256_set_epi64(0, 1, 2, 3);  
__m256i b = _mm256_set_epi64(3, 2, 1, 0);  
__m256i d = _mm256_sub_epi64(a, b);
```

VSX

```
vector uint64_t a = {0, 1};  
vector uint64_t b = {1, 0};  
vector uint64_t d = vec_sub(a, b);
```


4.3.5 32-bit floats

TODO: insert diagram for 32-bit float vector subtraction

C Intrinsics

SIMD	Code	Input	Output
MMX(64-bit) ²⁸	N/A	-	-
SSE(64-bit)	<code>d = _mm_sub_ss(a, b)</code>	<code>__m128 a, b;</code>	<code>__m128 d;</code>
SSE(128-bit)	<code>d = _mm_sub_ps(a, b)</code>	<code>__m128 a, b;</code>	<code>__m128 d;</code>
AVX(256-bit)	<code>d = _mm256_sub_ps(a, b)</code>	<code>__m256 a, b;</code>	<code>__m256 d;</code>
NEON(64-bit)	<code>d = vsub_f32(a, b)</code>	<code>float32x2_t a, b;</code>	<code>float32x2_t d;</code>
NEON(128-bit)	<code>d = vsubq_f32(a, b)</code>	<code>float32x4_t a, b;</code>	<code>float32x4_t d;</code>
Altivec/VMX	<code>d = vec_sub(a, b)</code>	<code>vector float a, b;</code>	<code>vector float d;</code>
VSX ²⁹	<code>d = vec_sub(a, b)</code>	<code>vector float a, b;</code>	<code>vector float d;</code>

²⁸MMX does not support float arithmetic

²⁹the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
MMX(64-bit)	N/A	–	–
SSE(64-bit)	<code>subss xmm1, xmm2</code>	Subtracts <code>xmm1</code> , <code>xmm2</code> , result in <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are SSE 128-bit double registers, but the operation takes place only at the lower 64-bits
SSE(128-bit)	<code>subps xmm1, xmm2</code>	Subtracts <code>xmm1</code> , <code>xmm2</code> , result in <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are SSE 128-bit float registers
AVX(128-bit)	<code>subps xmm1, xmm2, xmm3</code>	Subtracts <code>xmm2</code> , <code>xmm3</code> , result <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> , <code>xmm3</code> are AVX 128-bit registers
AVX(256-bit)	<code>subps ymm1, ymm2, ymm3</code>	Subtracts <code>ymm2</code> , <code>ymm3</code> , result in <code>ymm1</code>	<code>ymm1</code> , <code>ymm2</code> , <code>ymm3</code> are AVX 256-bit registers
NEON(64-bit)	<code>vsub.f32 d1, d2, d3</code>	Subtracts <code>d2</code> , <code>d3</code> , result in <code>d1</code>	<code>d1</code> , <code>d2</code> , <code>d3</code> are 64-bit NEON registers
NEON(128-bit)	<code>vsubq.f32 q1, q2, q3</code>	Subtracts <code>q2</code> , <code>q3</code> , result in <code>q1</code>	<code>q1</code> , <code>q2</code> , <code>q3</code> are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>fsub Vd.2S, Vn.2S, Vm.2S</code>	Subtracts <code>Vn.2S</code> , <code>Vm.2S</code> , result in <code>Vd.2S</code>	<code>Vn.2S</code> , <code>Vm.2S</code> , <code>Vd.2S</code> are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>fsub Vd.4S, Vn.4S, Vm.4S</code>	Subtracts <code>Vn.4S</code> , <code>Vm.4S</code> , result in <code>Vd.4S</code>	<code>Vn.4S</code> , <code>Vm.4S</code> , <code>Vd.4S</code> are 128-bit NEON registers
Altivec/VMX	<code>vsubfp vD, vA, vB</code>	Subtracts <code>vA</code> , <code>vB</code> , result in <code>vD</code>	<code>vA</code> , <code>vB</code> , <code>vD</code> are normal VMX 128-bit registers
VSX	<code>vsubfp vD, vA, vB</code>	Subtracts <code>vA</code> , <code>vB</code> , result in <code>vD</code>	VMX instruction
VSX	<code>xssubsp vD, vA, vB</code>	Adds <code>vA</code> , <code>vB</code> , result in <code>vD</code>	<code>vA</code> , <code>vB</code> , <code>vD</code> are VSX registers holding double values, but the operation takes place only at the lower 64-bits (FIXME: better description)

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```

__m128 a = _mm_set_ps(1.0, 2.0, 3.0, 4.0);
__m128 b = _mm_set_ps(4.0, 3.0, 2.0, 1.0);
__m128 d = _mm_sub_ps(a, b);

```

AVX (256-bit)

```
__m256 a = _mm256_set_ps(1.0, 2.0, 3.0, 4.0);  
__m256 b = _mm256_set_ps(4.0, 3.0, 2.0, 1.0);  
__m256 d = _mm256_sub_ps(a, b);
```

NEON (128-bit)

```
float32x4_t a = {0, 1, 2, 3};  
float32x4_t b = {3, 2, 1, 0};  
float32x4_t d = vsubq_f32(a, b);
```

AltiVec/VMX

```
vector float a = {0, 1, 2, 3};  
vector float b = {3, 2, 1, 0};  
vector float d = vec_sub(a, b);
```

4.3.6 64-bit doubles

TODO: insert diagram for 64-bit float vector subtraction

C Intrinsics

SIMD	Code	Input	Output
MMX(64-bit) ³⁰	N/A	-	-
SSE2(64-bit)	d = <code>_mm_sub_sd(a, b)</code>	<code>__m128d a, b;</code>	<code>__m128d d;</code>
SSE2(128-bit)	d = <code>_mm_sub_pd(a, b)</code>	<code>__m128d a, b;</code>	<code>__m128d d;</code>
AVX(256-bit)	d = <code>_mm256_sub_pd(a, b)</code>	<code>__m256d a, b;</code>	<code>__m256d d;</code>
NEON(64-bit) ³¹	d = <code>vadd_f64(a, b)</code>	<code>float64x1_t a, b;</code>	<code>float64x1_t d;</code>
NEON(128-bit)	d = <code>vaddq_f64(a, b)</code>	<code>float64x2_t a, b;</code>	<code>float64x2_t d;</code>
Altivec/VMX ³²	N/A	-	-
VSX	d = <code>vec_sub(a, b)</code>	<code>vector double a, b;</code>	<code>vector double d;</code>

³⁰MMX does not support double arithmetic

³¹double arithmetic is supported only by armv8 NEON

³²Altivec/VMX does not support double arithmetic

Assembly

SIMD	Code	Description	Notes
MMX(64-bit)	N/A	–	–
SSE2(64-bit)	<code>subsd xmm1, xmm2</code>	Subtracts <code>xmm1</code> , <code>xmm2</code> , result in <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are SSE 128-bit double registers, but the operation takes place only at the lower 64-bits
SSE2(128-bit)	<code>subpd xmm1, xmm2</code>	Subtracts <code>xmm1</code> , <code>xmm2</code> , result in <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are SSE 128-bit float registers
AVX(128-bit)	<code>subpd xmm1, xmm2, xmm3</code>	Subtracts <code>xmm2</code> , <code>xmm3</code> , result <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> , <code>xmm3</code> are AVX 128-bit registers
AVX(256-bit)	<code>subpd ymm1, ymm2, ymm3</code>	Subtracts <code>ymm2</code> , <code>ymm3</code> , result in <code>ymm1</code>	<code>ymm1</code> , <code>ymm2</code> , <code>ymm3</code> are AVX 256-bit registers
NEON(armv7, 64-bit)	N/A	–	–
NEON(armv7, 128-bit)	N/A	–	–
NEON(armv8, 64-bit)	<code>fadd Dd, Dn, Dm</code>	Subtracts <code>Dn</code> , <code>Dm</code> , result in <code>Dd</code>	<code>Dn</code> , <code>Dm</code> , <code>Dd</code> are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>fadd Vd.2D, Vn.2D, Vm.2D</code>	Subtracts <code>Vn.2D</code> , <code>Vm.2D</code> , result in <code>Vd.2D</code>	<code>Vn.2D</code> , <code>Vm.2D</code> , <code>Vd.2D</code> are 128-bit NEON registers
Altivec/VMX	N/A	–	–
VSX	<code>xssubdp vD, vA, vB</code>	Subtracts <code>vA</code> , <code>vB</code> , result in <code>vD</code>	<code>vA</code> , <code>vB</code> , <code>vD</code> are VSX registers holding double values, but the operation takes place only at the lower 64-bits
VSX	<code>xvsubdp vD, vA, vB</code>	Subtracts <code>vA</code> , <code>vB</code> , result in <code>vD</code>	<code>vA</code> , <code>vB</code> , <code>vD</code> are VSX registers holding double values

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```

__m128d a = _mm_set_pd(1.0, 2.0, 3.0, 4.0);
__m128d b = _mm_set_pd(4.0, 3.0, 2.0, 1.0);
__m128d d = _mm_sub_pd(a, b);

```

AVX (256-bit)

```
__m256d a = _mm256_set_pd(1.0, 2.0, 3.0, 4.0);  
__m256d b = _mm256_set_pd(4.0, 3.0, 2.0, 1.0);  
__m256d d = _mm256_sub_pd(a, b);
```

NEON (128-bit)

```
float64x2_t a = {0, 1};  
float64x2_t b = {3, 2};  
float64x2_t d = vsubq_f64(a, b);
```

VSX

```
vector double a = {0, 1, 2, 3};  
vector double b = {3, 2, 1, 0};  
vector double d = vec_sub(a, b);
```

4.4 Subtraction, saturated

Saturated subtraction in contrast to modulo addition, keeps the result in a fixed range between a minimum and maximum value. If it's greater than the maximum value, then it's "clamped" to the maximum value, if it's below the minimum then it's set to the minimum value. In case of 8-bit unsigned bytes (unsigned char), adding 10 to 255 using saturated addition would produce the result 255.

4.4.1 8-bit Integers

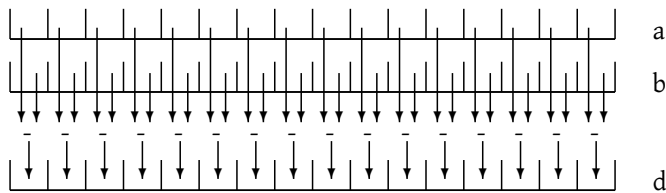


Figure 4.4: Addition of 16 8-bit integer elements in 128-bit vectors

C Intrinsics

SIMD	Code	Input	Output
8-bit unsigned			
MMX(64-bit)	<code>d = _mm_subs_pu8(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_subs_epu8(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_subs_epu8(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vsubs_u8(a, b)</code>	<code>uint8x8_t a, b;</code>	<code>uint8x8_t d;</code>
NEON(128-bit)	<code>d = vsubsq_u8(a, b)</code>	<code>uint8x16_t a, b;</code>	<code>uint8x16_t d;</code>
Altivec/VMX	<code>d = vec_subs(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
VSX ³³	<code>d = vec_subs(a, b)</code>	<code>vector uint8_t a, b;</code>	<code>vector uint8_t d;</code>
8-bit signed			
MMX(64-bit)	<code>d = _mm_subs_pi8(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_subs_epi8(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_subs_epi8(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vqsub_s8(a, b)</code>	<code>int8x8_t a, b;</code>	<code>int8x8_t d;</code>
NEON(128-bit)	<code>d = vqsubq_s8(a, b)</code>	<code>int8x16_t a, b;</code>	<code>int8x16_t d;</code>
Altivec/VMX	<code>d = vec_subs(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>
VSX ³⁴	<code>d = vec_subs(a, b)</code>	<code>vector int8_t a, b;</code>	<code>vector int8_t d;</code>

³³the VMX instruction will be used

³⁴the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
8-bit unsigned			
MMX(64-bit)	<code>psubusb mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubusb xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubusb xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubusb ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vqsub.i8 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqsubq.i8 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqsub Vd.8B, Vn.8B, Vm.8B</code>	Subtracts Vn.8B, Vm.8B, result in Vd.8B	d1,d2,d3 are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqsub Vd.16B, Vn.16B, Vm.16B</code>	Subtracts Vn.16B, Vm.16B, result in Vd.16B	Vn.16B, Vm.16B, Vd.16B are 128-bit NEON registers
Altivec/VMX	<code>vsububs vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsububs vD, vA, vB</code>	Subtracts vA, vB, result in vD	VMX instruction
8-bit signed			
MMX(64-bit)	<code>psubsb mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubsb xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubsb xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubsb ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(64-bit)	<code>vqsub.s8 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vqsubq.s8 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqsub Vd.8B, Vn.8B, Vm.8B</code>	Subtracts Vn.8B, Vm.8B, result in Vd.8B	d1,d2,d3 are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqsub Vd.16B, Vn.16B, Vm.16B</code>	Subtracts Vn.16B, Vm.16B, result in Vd.16B	Vn.16B, Vm.16B, Vd.16B are 128-bit NEON registers
Altivec/VMX	<code>vsubsb vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubsb vD, vA, vB</code>	Subtracts vA, vB, result in vD	VMX instruction

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epu8(0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15);  
__m128i b = _mm_set_epu8(15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0);  
__m128i d = _mm_subs_epu8(a, b);
```

NEON (128-bit)

```
int8x16_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
int8x16_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
int8x16_t d = vqsubq_s8(a, b);
```

AltiVec/VMX

```
vector_uint8_t a = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15};  
vector_uint8_t b = {15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0};  
vector_uint8_t d = vec_subs(a, b);
```

4.4.2 16-bit integers

TODO: insert diagram for 16-bit saturated vector addition

C Intrinsics

SIMD	Code	Input	Output
16-bit unsigned			
MMX(64-bit)	<code>d = _mm_subs_pu16(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_subs_epu16(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_subs_epu16(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vqsub_u16(a, b)</code>	<code>uint16x4_t a, b;</code>	<code>uint16x4_t d;</code>
NEON(128-bit)	<code>d = vqsubq_u16(a, b)</code>	<code>uint16x8_t a, b;</code>	<code>uint16x8_t d;</code>
Altivec/VMX	<code>d = vec_subs(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
VSX ³⁵	<code>d = vec_adds(a, b)</code>	<code>vector uint16_t a, b;</code>	<code>vector uint16_t d;</code>
16-bit signed			
MMX(64-bit)	<code>d = _mm_subs_pi16(a, b)</code>	<code>__m64 a, b;</code>	<code>__m64 d;</code>
SSE2(128-bit)	<code>d = _mm_subs_epi16(a, b)</code>	<code>__m128i a, b;</code>	<code>__m128i d;</code>
AVX(256-bit)	<code>d = _mm256_subs_epi16(a, b)</code>	<code>__m256i a, b;</code>	<code>__m256i d;</code>
NEON(64-bit)	<code>d = vqsub_s16(a, b)</code>	<code>int16x4_t a, b;</code>	<code>int16x4_t d;</code>
NEON(128-bit)	<code>d = vqsubq_s16(a, b)</code>	<code>int16x8_t a, b;</code>	<code>int16x8_t d;</code>
Altivec/VMX	<code>d = vec_subs(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>
VSX ³⁶	<code>d = vec_subs(a, b)</code>	<code>vector int16_t a, b;</code>	<code>vector int16_t d;</code>

³⁵the VMX instruction will be used

³⁶the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
16-bit unsigned			
MMX(64-bit)	<code>psubsw mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubsw xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubsw xmm1, xmm2, xmm3</code>	Subtractss xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubsw ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vqsub.i16 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqsubq.i16 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqsub Vd.4H, Vn.4H, Vm.4H</code>	Subtracts Vn.4H, Vm.4H, result in Vd.4H	Vn.4H, Vm.4H, Vd.4H are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqsub Vd.8H, Vn.8H, Vm.8H</code>	Subtracts Vn.8H, Vm.8H, result in Vd.8H	Vn.8H, Vm.8H, Vd.8H are 128-bit NEON registers
Altivec/VMX	<code>vsubuhm vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubuhm vD, vA, vB</code>	Subtracts vA, vB, result in vD	VMX instruction
16-bit signed			
MMX(64-bit)	<code>psubsw mm1, mm2</code>	Subtracts mm1, mm2, result in mm1	mm1, mm2 are 64-bit MMX registers
SSE2(128-bit)	<code>psubsw xmm1, xmm2</code>	Subtracts xmm1, xmm2, result in xmm1	xmm1, xmm2 are SSE 128-bit registers
AVX(128-bit)	<code>vpsubsw xmm1, xmm2, xmm3</code>	Subtracts xmm2, xmm3, result xmm1	xmm1, xmm2, xmm3 are AVX 128-bit registers
AVX(256-bit)	<code>vpsubsw ymm1, ymm2, ymm3</code>	Subtracts ymm2, ymm3, result in ymm1	ymm1, ymm2, ymm3 are AVX 256-bit registers
NEON(armv7, 64-bit)	<code>vqsub.s16 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqsubq.s16 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqsub Vd.4H, Vn.4H, Vm.4H</code>	Subtracts Vn.4H, Vm.4H, result in Vd.4H	Vn.4H, Vm.4H, Vd.4H are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqsub Vd.8H, Vn.8H, Vm.8H</code>	Subtracts Vn.8H, Vm.8H, result in Vd.8H	Vn.8H, Vm.8H, Vd.8H are 128-bit NEON registers
Altivec/VMX	<code>vsubshm vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubshm vD, vA, vB</code>	Subtracts vA, vB, result in vD	VMX instruction

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

SSE (128-bit)

```
__m128i a = _mm_set_epi16(0, 1, 2, 3, 4, 5, 6, 7);  
__m128i b = _mm_set_epi16(7, 6, 5, 4, 3, 2, 1, 0);  
__m128i d = _mm_subs_epi16(a, b);
```

NEON (128-bit)

```
int16x8_t a = {0, 1, 2, 3, 4, 5, 6, 7};  
int16x8_t b = {7, 6, 5, 4, 3, 2, 1, 0};  
int16x8_t d = vqsubq_s16(a, b);
```

AltiVec/VMX

```
vector uint16_t a = {0, 1, 2, 3, 4, 5, 6, 7};  
vector uint16_t b = {7, 6, 5, 4, 3, 2, 1, 0};  
vector uint16_t d = vec_subs(a, b);
```

4.4.3 32-bit integers

TODO: insert diagram for 32-bit integer saturated subtraction

C Intrinsics

SIMD	Code	Input	Output
32-bit unsigned			
MMX(64-bit) ³⁷	N/A	-	-
SSE2(128-bit) ³⁸	N/A	-	-
AVX(256-bit) ³⁹	N/A	-	-
NEON(64-bit)	d = vqsub_u32(a, b)	uint32x2_t a, b;	uint32x2_t d;
NEON(128-bit)	d = vqsubq_u32(a, b)	uint32x4_t a, b;	uint32x4_t d;
Altivec/VMX	d = vec_subs(a, b)	vector uint32_t a, b;	vector uint32_t d;
VSX ⁴⁰	d = vec_subs(a, b)	vector uint32_t a, b;	vector uint32_t d;
32-bit signed			
MMX(64-bit)	N/A	-	-
SSE2(128-bit)	N/A	-	-
AVX(256-bit)	N/A	-	-
NEON(64-bit)	d = vqsub_s32(a, b)	int32x2_t a, b;	int32x2_t d;
NEON(128-bit)	d = vqsubq_s32(a, b)	int32x4_t a, b;	int32x4_t d;
Altivec/VMX	d = vec_subs(a, b)	vector int32_t a, b;	vector int32_t d;
VSX	d = vec_subs(a, b)	vector int32_t a, b;	vector int32_t d;

³⁷MMX has no support for saturated addition on 32-bit integers

³⁸SSE has no support for saturated addition on 32-bit integers

³⁹AVX has no support for saturated addition on 32-bit integers

⁴⁰the VMX instruction will be used

Assembly

SIMD	Code	Description	Notes
32-bit unsigned			
MMX(64-bit)	N/A	-	-
SSE2(128-bit)	N/A	-	-
AVX(128-bit)	N/A	-	-
AVX(256-bit)	N/A	-	-
NEON(armv7, 64-bit)	<code>vqsub.i32 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqsubq.i32 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqsub Vd.2S, Vn.2S, Vm.2S</code>	Subtracts Vn.2S, Vm.2S, result in Vd.2S	Vn.2S, Vm.2S, Vd.2S are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqsub Vd.4S, Vn.4S, Vm.4S</code>	Subtracts Vn.4S, Vm.4S, result in Vd.4S	Vn.4S, Vm.4S, Vd.4S are 128-bit NEON registers
Altivec/VMX	<code>vsubuvs vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubuvs vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction
32-bit signed			
MMX(64-bit)	N/A	-	-
SSE2(128-bit)	N/A	-	-
AVX(128-bit)	N/A	-	-
AVX(256-bit)	N/A	-	-
NEON(armv7, 64-bit)	<code>vqsub.s32 d1, d2, d3</code>	Adds d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqsubq.s32 q1, q2, q3</code>	Adds q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqsub Vd.2S, Vn.2S, Vm.2S</code>	Subtracts Vn.2S, Vm.2S, result in Vd.2S	Vn.2S, Vm.2S, Vd.2S are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqsub Vd.4S, Vn.4S, Vm.4S</code>	Subtracts Vn.4S, Vm.4S, result in Vd.4S	Vn.4S, Vm.4S, Vd.4S are 128-bit NEON registers
Altivec/VMX	<code>vsubsws vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubsws vD, vA, vB</code>	Adds vA, vB, result in vD	VMX instruction

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

NEON (128-bit)

```
uint32x4_t a = {0, 1, 2, 3};
```

```
uint32x4_t b = {3, 2, 1, 0};  
uint32x4_t d = vqsubq_u32(a, b);
```

Altivec/VMX

```
vector uint32_t a = {0, 1, 2, 3};  
vector uint32_t b = {3, 2, 1, 0};  
vector uint32_t d = vec_subs(a, b);
```

4.4.4 64-bit integers

At the moment the only architecture that supports 64-bit saturated subtraction is NEON, in both ARMv7 and ARMv8 implementations.

TODO: insert diagram for 64-bit vector saturated subtraction

C Intrinsics

SIMD	Code	Input	Output
64-bit unsigned			
NEON(armv7, 64-bit)	<code>d = vqsub_u64(a, b)</code>	<code>uint64x1_t a, b;</code>	<code>uint64x1_t d;</code>
NEON(armv7, 128-bit)	<code>d = vqsubq_u64(a, b)</code>	<code>uint64x2_t a, b;</code>	<code>uint64x2_t d;</code>
64-bit signed			
NEON(64-bit)	<code>d = vqsub_s64(a, b)</code>	<code>int64x1_t a, b;</code>	<code>int64x1_t d;</code>
NEON(128-bit)	<code>d = vqsubq_s64(a, b)</code>	<code>int64x2_t a, b;</code>	<code>int64x2_t d;</code>

Assembly

SIMD	Code	Description	Notes
64-bit unsigned			
NEON(armv7, 64-bit)	<code>vqsub.i64 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(armv7, 128-bit)	<code>vqsubq.i64 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>uqsub Dd, Dn, Dm</code>	Subtracts Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>uqsub Vd.2D, Vn.2D, Vm.2D</code>	Subtracts Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers
64-bit signed			
NEON(64-bit)	<code>vqsub.s64 d1, d2, d3</code>	Subtracts d2, d3, result in d1	d1,d2,d3 are 64-bit NEON registers
NEON(128-bit)	<code>vqsubq.s64 q1, q2, q3</code>	Subtracts q2, q3, result in q1	q1,q2,q3 are 128-bit NEON registers
NEON(armv8, 64-bit)	<code>sqsub Dd, Dn, Dm</code>	Subtracts Dn, Dm, result in Dd	Dn, Dm, Dd are 64-bit NEON registers
NEON(armv8, 128-bit)	<code>sqsub Vd.2D, Vn.2D, Vm.2D</code>	Subtracts Vn.2D, Vm.2D, result in Vd.2D	Vn.2D, Vm.2D, Vd.2D are 128-bit NEON registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

NEON (128-bit)

```
uint64x2_t a = {0, 1};  
uint64x2_t b = {3, 2};  
uint64x2_t d = vsubq_u64(a, b);
```

4.5 Special addition/subtraction instructions

4.5.1 Addition/subtraction with carry (32-bit integers, unsigned)

Altivec/VSX also supports addition/subtraction with carry, but only for 32-bit/64-bit unsigned integers.

TODO: insert diagram for 32-bit vector addition/subtraction with carry

C Ininsics

SIMD	Code	Input	Output
Altivec/VMX/VSX	<code>d = vec_addc(a, b)</code>	vector uint32_t a, b;	vector uint32_t d;
Altivec/VMX/VSX	<code>d = vec_subc(a, b)</code>	vector uint32_t a, b;	vector uint32_t d;

Assembly

SIMD	Code	Description	Notes
Altivec/VMX/VSX	<code>vaddcuw vD, vA, vB</code>	Adds vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
Altivec/VMX/VSX	<code>vsubcuw vD, vA, vB</code>	Subtracts vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

Altivec/VMX/VSX

```
vector uint32_t a = {0, 1, 2, 3};
vector uint32_t b = {3, 2, 1, 0};
vector uint32_t d = vec_addc(a, b);
vector uint32_t e = vec_subc(a, b);
```

4.5.2 Addition/subtraction with carry (64-bit integers, unsigned)

On par with Altivec/VMX, 64-bit integer addition/subtraction with carry is supported only by its extended version VSX.

TODO: insert diagram for 64-bit vector addition

C Intrinsics

SIMD	Code	Input	Output
VSX	<code>d = vec_addc(a, b)</code>	vector uint64_t a, b;	vector uint64_t d;
VSX	<code>d = vec_subc(a, b)</code>	vector uint64_t a, b;	vector uint64_t d;

Assembly

SIMD	Code	Description	Notes
VSX	<code>vaddcuq vD, vA, vB</code>	Adds with carry vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers
VSX	<code>vsubcuq vD, vA, vB</code>	Subtracts with carry vA, vB, result in vD	vA, vB, vD are normal VMX 128-bit registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

VSX

```
vector uint64_t a = {0, 1};
vector uint64_t b = {3, 2};
vector uint64_t d = vec_addc(a, b);
vector uint64_t e = vec_subc(a, b);
```

4.5.3 Odd/Even Addition/subtraction (32-bit floats/64-bit doubles)

Odd/Even Addition/Subtraction is an SSE3-specific instruction (with the respective C intrinsic) that adds odd-numbered elements and subtracts even-numbered elements in the vector.

TODO: insert diagram for 32-bit float vector odd/even addition/subtraction

C Intrinsics

SIMD	Code	Input	Output
SSE3(128-bit)	<code>d = _mm_addsub_ps(a, b)</code>	<code>__m128 a, b;</code>	<code>__m128 d;</code>
SSE3(128-bit)	<code>d = _mm_addsub_pd(a, b)</code>	<code>__m128d a, b;</code>	<code>__m128d d;</code>

Assembly

SIMD	Code	Description	Notes
SSE3(128-bit)	<code>addsubps xmm1, xmm2/m128</code>	Subtracts <code>xmm1</code> , <code>xmm2</code> , result in <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are SSE 128-bit registers
SSE3(128-bit)	<code>addsubpd xmm1, xmm2/m128</code>	Subtracts <code>xmm1</code> , <code>xmm2</code> , result in <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are SSE 128-bit registers
AVX(128-bit)	<code>vaddsubps xmm1, xmm2, xmm3/m128</code>	Subtractss <code>xmm2</code> , <code>xmm3</code> , result <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are AVX 128-bit registers, <code>xmm3</code> can be an AVX
AVX(256-bit)	<code>vaddsubps ymm1, ymm2, ymm3/m128</code>	Subtracts <code>ymm2</code> , <code>ymm3</code> , result in <code>ymm1</code>	<code>ymm1</code> , <code>ymm2</code> , <code>ymm3</code> are AVX 256-bit registers
AVX(128-bit)	<code>vaddsubpd xmm1, xmm2, xmm3/m128</code>	Subtractss <code>xmm2</code> , <code>xmm3</code> , result <code>xmm1</code>	<code>xmm1</code> , <code>xmm2</code> are AVX 128-bit registers, <code>xmm3</code> can be an AVX
AVX(256-bit)	<code>vaddsubpd ymm1, ymm2, ymm3/m128</code>	Subtracts <code>ymm2</code> , <code>ymm3</code> , result in <code>ymm1</code>	<code>ymm1</code> , <code>ymm2</code> , <code>ymm3</code> are AVX 256-bit registers

Notes

Examples

Here, we present some coding examples using the C intrinsics for some of the SIMD engines.

VSX

```
vector uint64_t a = {0, 1};  
vector uint64_t b = {3, 2};  
vector uint64_t d = vec_addc(a, b);
```

4.6 In-vector sum reduction

Many times one must do an in-vector sum reduction, that is, an addition of all elements in a vector, and the normal addition instructions won't do, and more sophisticated instructions are called for. Most SIMD engines provide for at least one of such instructions for at least one datatype, but a 1:1 mapping is not possible. However, since the current document is a Functionality Reference Guide (with emphasis on Functionality) we will provide not only the instruction where available, but also ways to duplicate the functionality in other SIMD engines where a single instruction is missing.

4.6.1 Integer Sum across Partial 1/4 saturated

TODO: insert diagram for 32-bit vector partial 1/4 sum reduction

4.6.2 Sum across Partial 1/2 saturated

4.6.3 Integer Sum across saturated

4.6.4 Floating point sum reduction

4.7 Multiplication Scalar x Vector

4.8 Multiplication of Vectors

4.9 Multiplication with Accumulation

4.10 Polynomial Multiplication

4.11 Division

4.12 Square Root

4.13 Packing/Unpacking

4.14 Shuffling/Permutations/Combinations

4.15 Transpose/Zip/Unzip

4.16 Rounding

4.17 Vector Comparison

4.18 Minimum/Maximum

4.19 Absolute value

4.20 Bit Selection

4.21 Bit Shifting

4.22 Negation/OR/XOR

4.23 Reversing

4.24 Bit counting

4.25 Getting/Setting single elements

Index

`_mm256_add_epi16`, 13
`_mm256_add_epi32`, 16
`_mm256_add_epi64`, 19
`_mm256_add_epi8`, 10
`_mm256_add_ps`, 24
`_mm256_add_ps`, 21
`_mm256_adds_epi16`, 30
`_mm256_adds_epi8`, 27
`_mm256_adds_epu16`, 30
`_mm256_adds_epu8`, 27
`_mm256_sub_epi64`, 47
`_mm256_sub_epi16`, 41
`_mm256_sub_epi32`, 44
`_mm256_sub_epi8`, 38
`_mm256_sub_ps`, 52
`_mm256_sub_ps`, 49
`_mm256_subs_epi16`, 58
`_mm256_subs_epi8`, 55
`_mm256_subs_epu16`, 58
`_mm256_subs_epu8`, 55
`_mm_add_epi16`, 13
`_mm_add_epi32`, 16
`_mm_add_epi64`, 19
`_mm_add_epi8`, 10
`_mm_add_pd`, 24
`_mm_add_pi16`, 13
`_mm_add_pi32`, 16
`_mm_add_pi64`, 19
`_mm_add_pi8`, 10
`_mm_add_ps`, 21
`_mm_add_sd`, 24
`_mm_add_ss`, 21
`_mm_adds_epi16`, 30
`_mm_adds_epu16`, 30
`_mm_adds_pi16`, 30
`_mm_adds_pi8`, 27
`_mm_adds_pu16`, 30
`_mm_adds_pu8`, 27
`_mm_addsub_ps`, 68
`_mm_addsub_ps`, 68
`_mm_adds_epi8`, 27
`_mm_adds_epu8`, 27
`_mm_sub_epi16`, 41
`_mm_sub_epi32`, 44
`_mm_sub_epi64`, 47
`_mm_sub_epi8`, 38
`_mm_sub_pd`, 52
`_mm_sub_pi16`, 41
`_mm_sub_pi32`, 44
`_mm_sub_pi64`, 47
`_mm_sub_pi8`, 38
`_mm_sub_ps`, 49
`_mm_sub_sd`, 52
`_mm_sub_ss`, 49
`_mm_subs_epi16`, 58
`_mm_subs_epi8`, 55
`_mm_subs_epu16`, 58
`_mm_subs_epu8`, 55

`_mm_subs_pi16`, 58
`_mm_subs_pi8`, 55
`_mm_subs_pu16`, 58
`_mm_subs_pu8`, 55

`add`, 11, 14, 17

Addition

- `modulo`, 10
- `odd/Even`, 68
- `saturated`, 27
- `with carry`, 66, 67
- `wrap-around`, 10

`addpd`, 25
`addps`, 22
`addsd`, 25
`addss`, 22
`addsubpd`, 68
`addsubps`, 68

`fadd`, 22

`paddb`, 11
`paddd`, 17
`paddq`, 19
`paddsb`, 28
`paddsw`, 31
`paddw`, 31
`paddusb`, 28
`paddw`, 14
`psubb`, 39
`psubd`, 45
`psubq`, 47
`psubsb`, 56
`psubsw`, 59
`psubw`, 59
`psubusb`, 56
`psubw`, 42

`sqadd`, 28, 31, 34, 36

`sqsub`, 56, 59, 62, 64

`sub`, 42
`subpd`, 53
`subps`, 50
`subsd`, 53
`subss`, 50

Subtraction

- `modulo`, 38
- `odd/Even`, 68
- `saturated`, 55
- `with carry`, 66, 67
- `wrap-around`, 38

`uqadd`, 28, 31, 34, 36
`uqsub`, 56, 59, 62, 64

`vadd.f32`, 22
`vadd.i16`, 14
`vadd.i32`, 17
`vadd.i64`, 19
`vadd.i8`, 11
`vadd.s16`, 14
`vadd.s32`, 17
`vadd.s64`, 19
`vadd.s8`, 11
`vadd.f32`, 21
`vadd.s16`, 13
`vadd.s32`, 16
`vadd.s64`, 19
`vadd.s8`, 10
`vadd.u16`, 13
`vadd.u32`, 16
`vadd.u64`, 19
`vadd.u8`, 10
`vaddcuq`, 67
`vaddcuw`, 66
`vadddp`, 25
`vaddfp`, 22
`vaddq.f32`, 22

vaddq.i16, 14
 vaddq.i32, 17
 vaddq.i64, 19
 vaddq.i8, 11
 vaddq.s16, 14
 vaddq.s32, 17
 vaddq.s64, 19
 vaddq.s8, 11
 vaddq_f32, 21
 vaddq_s16, 13
 vaddq_s32, 16
 vaddq_s64, 19
 vaddq_s8, 10
 vaddq_u16, 13
 vaddq_u32, 16
 vaddq_u64, 19
 vaddq_u8, 10
 vaddsbs, 28
 vaddubs, 28
 vaddshm, 31
 vaddsubpd, 68
 vaddsubps, 68
 vaddsws, 34
 vaddubm, 11
 vadduhm, 14, 31
 vadduqm, 19
 vadduwm, 17
 vadduws, 34
 vec_add, 10, 13, 16, 19, 21, 24
 vec_addc, 67
 vec_addc, 66
 vec_adds, 27, 30, 33
 vec_sub, 38, 41, 44, 47, 49, 52
 vec_subc, 67
 vec_subs, 55, 58, 61
 vec_sum4s, 70
 vpaddb, 11
 vpadd, 17
 vpaddq, 19
 vpaddsb, 28
 vpaddsw, 31
 vpaddw, 31
 vpaddusb, 28
 vpaddw, 14
 vpsubb, 39
 vpsubd, 45
 vpsubq, 47
 vpsubsb, 56
 vpsubsw, 59
 vpsubw, 59
 vpsubusb, 56
 vpsubw, 42
 vqadd.i16, 31
 vqadd.i32, 34
 vqadd.i64, 36
 vqadd.i8, 28
 vqadd.s16, 31
 vqadd.s32, 34
 vqadd.s64, 36
 vqadd.s8, 28
 vqadd_s16, 30
 vqadd_s32, 33
 vqadd_s64, 36, 64
 vqadd_s8, 27
 vqadd_u16, 30
 vqadd_u32, 33
 vqadd_u64, 36, 64
 vqadd_u8, 27
 vqaddq.i16, 31
 vqaddq.i32, 34
 vqaddq.i64, 36
 vqaddq.i8, 28
 vqaddq.s16, 31
 vqaddq.s32, 34
 vqaddq.s64, 36
 vqaddq.s8, 28
 vqaddq_s16, 30
 vqaddq_s32, 33

vqaddq.s64, 36, 64
vqaddq.s8, 27
vqaddq.u16, 30
vqaddq.u32, 33
vqaddq.u64, 36, 64
vqaddq.u8, 27
vqsub.i16, 59
vqadd.i32, 62
vqsub.i64, 64
vqsub.i8, 56
vqsub.s16, 59
vqadd.s32, 62
vqsub.s64, 64
vqsub.s8, 56
vqsub.s16, 58
vqsub.s32, 61
vqsub.s8, 55
vqsub.u16, 58
vqsub.u32, 61
vqsub.u8, 55
vqsubq.i16, 59
vqaddq.i32, 62
vqsubq.i64, 64
vqsubq.i8, 56
vqsubq.s16, 59
vqaddq.s32, 62
vqsubq.s64, 64
vqsubq.s8, 56
vqsubq.s16, 58
vqsubq.s32, 61
vqsubq.s8, 55
vqsubq.u16, 58
vqsubq.u32, 61
vqsubq.u8, 55
vsub.f32, 50
vsub.i16, 42
vsub.i32, 45
vsub.i8, 39
vsub.s16, 42
vsub.s32, 45
vsub.s8, 39
vsub.f32, 49
vsub.s16, 41
vsub.s32, 44
vsub.s8, 38
vsub.u16, 41
vsub.u32, 44
vsub.u8, 38
vsubcuq, 67
vsubcuw, 66
vsubfp, 50
vsubq.f32, 50
vsubq.i16, 42
vsubq.i32, 45
vsubq.i8, 39
vsubq.s16, 42
vsubq.s32, 45
vsubq.s8, 39
vsubq.f32, 49
vsubq.s16, 41
vsubq.s32, 44
vsubq.s8, 38
vsubq.u16, 41
vsubq.u32, 44
vsubq.u8, 38
vsubsubs, 56
vsububs, 56
vsubshm, 59
vaddsws, 62
vsububm, 39
vsubuhm, 42, 59
vsubuqm, 47
vsubuwm, 45
vadduws, 62

xssubdp, 53
xssubsp, 50
xvsubdp, 53