

# IBM Research Report

## QPX Architecture: Quad Processing eXtension to the Power ISA™

**Thomas Fox, Michael Gschwind, Jaime Moreno**

(foxy, mkg, jhmoreno@us.ibm.com)

IBM Research Division

Thomas J. Watson Research Center

P.O. Box 218

Yorktown Heights, NY 10598



Research Division

Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

### **Contributors**

The following individuals contributed to an extensive degree to this architectural specification:

John Gunnels

Alexandre Eichenberger

Daniel Prener

Bruce Fleischer

### **Acknowledgements**

The following individuals played a significant role in the design of the Quad floating-point Processing Unit of the BQC compute chip in Blue Gene/Q:

Charles Wait

Thomas Roewer

Bernard Brezzo

# Chapter 1. Quad-Vector Floating-Point Facility Overview

This document defines the Blue Gene/Q Quad-Processing eXtension (QPX) to IBM's Power Instruction Set Architecture. Refer to IBM's Power ISA™ AS architecture document for descriptions of the base Power instruction set, the storage model, and related facilities available to the application programmer.

The computational model of the QPX architecture is a vector Single Instruction Multiple Data (SIMD) model with four execution slots and a register file containing 32 registers of 256 bits. Each of the 32 registers can be envisioned as containing four elements of 64 bits, whereby each of the execution slots operates on one vector element.

## 1.1 Notation

The following notation is specific to and used throughout the QPX Architecture document.

QRT, QRA, QRB, and QRC refer to Quad Floating-Point Registers, which are 256-bit vector registers containing four elements with 64 bits per element. The vector elements are numbered 0,1,2,3, with element 0 comprising bits 0:63, element 1 comprising bits 64:127, element 2 comprising bits 128:191, and element 3 comprising bits 192:255.

QRT<sup>x</sup> refers to element x of vector register QRT.


## Chapter 2. Quad-Vector Floating-Point Facility Registers

### 2.1 Quad-Vector Floating-Point Registers

Implementations of this architecture provide 32 Quad-vector floating-Point Registers (QPRs), named QPR0 through QPR31. The QPX instruction formats provide 5-bit fields for specifying the QPRs to be used in the execution of the instruction.

Scalar floating-point computational instructions, defined in the Power ISA, operate on element 0 QPRs, which serve as both the scalar FPRs for scalar instructions and the element 0 QPRs for vector instructions.

The figure below shows the Quad floating-point registers.



**Figure 1. Quad Floating-Point Registers**

### 2.2 Floating-Point Status and Control Register

The Floating-Point Exception Summary bits (32:34) and the Floating-Point Exception bits (35:44 and 53:55) of the FPSCR are never updated by QPX instructions, neither implicitly nor explicitly. The remaining status bits (45:51) are never updated by QPX instructions.

The Floating-Point Exception Enable bits (56:60) are ignored by all QPX instructions, which execute as if these bits were disabled. The Floating-Point Non-IEEE Mode (NI) bit (61) and the Floating-Point Rounding Control (RN) bits (62:63) of the FPSCR affect the operations on all four vector elements for QPX instructions.

### 2.3 Store Exception Enable Registers

Certain QPX store instructions provide a novel mechanism for the detection and indication of numerically exceptional conditions at the store interface.

A Store Indicate NaN Exception occurs when the source operand of a Store with Indicate instruction contains a NaN value. The Store Nan Exception Enable (SNEE) register enables the indication of such an exception. If an enabled Store Indicate NaN Exception occurs, the Auxiliary Processor bit of the Exception Syndrome Register is set (ESR[AP] = '1').

A Store Indicate Infinity Exception occurs when the source operand contains an Infinity value during a Store with Indicate instruction. The Store Infinity Exception Enable (SIEE) register enables the indication of such an exception. If an enabled Store Indicate Infinity Exception occurs, the Auxiliary Processor bit of the Exception Syndrome Register is set (ESR[AP] = '1').

The precedence of simultaneously occurring indication exceptions and memory fault exceptions is implementation defined.

#### Implementation Note

In the QPU for BGQ, the following bits in the AXUCR0 Special Purpose Register contain the SNEE and SIEE state on a per thread basis:

```
axucr0(20) : Thread 0 SNEE
axucr0(21) : Thread 0 SIEE
axucr0(22) : Thread 1 SNEE
axucr0(23) : Thread 1 SIEE
axucr0(24) : Thread 2 SNEE
axucr0(25) : Thread 2 SIEE
axucr0(26) : Thread 3 SNEE
axucr0(27) : Thread 3 SIEE
```

## Chapter 3. Scalar Instructions

Scalar floating-point load instructions, defined in the Power ISA, cause a replication of the source data across all elements of the target register.

Scalar floating-point move, arithmetic, rounding and conversion, compare, and select instructions, defined in the Power ISA, are executed in execution slot 0. Source operands for these instructions are read from element 0 QPRs, while target results are written to element 0 QPRs. Target elements 1, 2, and 3 are left in an undefined state.

## Chapter 4. Quad-Vector Floating-Point Facility Instructions

### 4.1 Quad-Vector Floating-Point Load Instructions

#### Quad-Vector Load Floating-point Single indexed X-form

qvlfsvx      QRT,RA,RB      (X=0)  
 qvlfsvxa    QRT,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 519 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

if RA = 0 then b ← 0
else            b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MVAL ← MEM(EA, 16)
QRT0 ← DOUBLE(MVAL0:31)
QRT1 ← DOUBLE(MVAL32:63)
QRT2 ← DOUBLE(MVAL64:95)
QRT3 ← DOUBLE(MVAL96:127)
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The 16 bytes in storage addressed by the 16-byte-aligned EA are interpreted as four single-precision vector elements, converted to double-precision format, and placed into register QRT.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

#### Quad-Vector Load Floating-point Single with Update indexed X-form

qvlfsvux      QRT,RA,RB      (X=0)  
 qvlfsvuxa    QRT,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 551 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF0
MVAL ← MEM(EA, 16)
QRT0 ← DOUBLE(MVAL0:31)
QRT1 ← DOUBLE(MVAL32:63)
QRT2 ← DOUBLE(MVAL64:95)
QRT3 ← DOUBLE(MVAL96:127)
RA ← EA
    
```

Let the effective address (EA) be the sum (RA)+(RB).

The 16 bytes in storage addressed by the 16-byte-aligned EA are interpreted as four single-precision vector elements, converted to double-precision format, and placed into register QRT.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Load Floating-point Double  
indexEd X-form**

qvlfdx      QRT,RA,RB      (X=0)  
 qvlfdxa    QRT,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 583 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

if RA = 0 then b ← 0
else            b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFE0
QRT ← MEM(EA, 32)
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The 32 bytes in storage addressed by the 32-byte-aligned EA are interpreted as four double-precision vector elements, and placed into register QRT.

If the X bit is set, and the address is not aligned on a 32-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Load Floating-point Double  
with Update indexEd X-form**

qvlfdux      QRT,RA,RB      (X=0)  
 qvlfduxa    QRT,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 615 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFE0
QRT ← MEM(EA, 32)
RA ← EA
    
```

Let the effective address (EA) be the sum (RA)+(RB).

The 32 bytes in storage addressed by the 32-byte-aligned EA are interpreted as four double-precision vector elements, and placed into register QRT.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 32-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None



**Quad-Vector Load Floating-point Complex Single indexed X-form**

qvlfcSX     QRT,RA,RB     (X=0)  
 qvlfcSxa     QRT,RA,RB     (X=1)

|   |    |     |    |    |   |    |
|---|----|-----|----|----|---|----|
| 0 | 31 | QRT | RA | RB | 7 | X  |
|   | 6  | 11  | 16 | 21 |   | 31 |

```

if RA = 0 then b ← 0
else            b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF8
MVAL ← MEM(EA, 8)
QRT0 ← DOUBLE(MVAL0:31)
QRT1 ← DOUBLE(MVAL32:63)
QRT2 ← DOUBLE(MVAL0:31)
QRT3 ← DOUBLE(MVAL32:63)
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).  
 The 8 bytes in storage addressed by the 8-byte-aligned EA are interpreted as two single-precision vector elements, converted to double-precision format, and replicated into register QRT.

If the X bit is set, and the address is not aligned on an 8-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Load Floating-point Complex Double indexed X-form**

qvlfcDx     QRT,RA,RB     (X=0)  
 qvlfcDxa     QRT,RA,RB     (X=1)

|   |    |     |    |    |    |    |
|---|----|-----|----|----|----|----|
| 0 | 31 | QRT | RA | RB | 71 | X  |
|   | 6  | 11  | 16 | 21 |    | 31 |

```

if RA = 0 then b ← 0
else            b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MVAL ← MEM(EA, 16)
QRT0 ← MVAL0
QRT1 ← MVAL1
QRT2 ← MVAL0
QRT3 ← MVAL1
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).  
 The 16 bytes in storage addressed by the 16-byte-aligned EA are interpreted as two double-precision vector elements, and replicated into register QRT.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Load Floating-point Complex Single with Update indexed X-form**

qvlfc<sub>sux</sub> QRT,RA,RB (X=0)  
 qvlfc<sub>suxa</sub> QRT,RA,RB (X=1)

|    |     |    |    |    |    |
|----|-----|----|----|----|----|
| 31 | QRT | RA | RB | 39 | X  |
| 0  | 6   | 11 | 16 | 21 | 31 |

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF8  
 MVAL ← MEM(EA, 8)  
 QRT<sup>0</sup> ← DOUBLE(MVAL<sub>0:31</sub>)  
 QRT<sup>1</sup> ← DOUBLE(MVAL<sub>32:63</sub>)  
 QRT<sup>2</sup> ← DOUBLE(MVAL<sub>0:31</sub>)  
 QRT<sup>3</sup> ← DOUBLE(MVAL<sub>32:63</sub>)  
 RA ← EA

Let the effective address (EA) be the sum (RA)+(RB).

The 8 bytes in storage addressed by the 8-byte-aligned EA are interpreted as two single-precision vector elements, converted to double-precision format, and replicated into register QRT.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 8-byte boundary, an exception is raised.

**Special Registers Altered:**

None

**Quad-Vector Load Floating-point Complex Double with Update indexed X-form**

qvlfc<sub>dux</sub> QRT,RA,RB (X=0)  
 qvlfc<sub>dux</sub>a QRT,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 103 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF0  
 MVAL ← MEM(EA, 16)  
 QRT<sup>0</sup> ← MVAL<sup>0</sup>  
 QRT<sup>1</sup> ← MVAL<sup>1</sup>  
 QRT<sup>2</sup> ← MVAL<sup>0</sup>  
 QRT<sup>3</sup> ← MVAL<sup>1</sup>  
 RA ← EA

Let the effective address (EA) be the sum (RA)+(RB).

The 16 bytes in storage addressed by the 16-byte-aligned EA are interpreted as two double-precision vector elements, and replicated into register QRT.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**

None



**Quad-Vector Load Permute Control Left  
Double indexed X-form**

qvlpcldx QRT,RA,RB

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 582 | /  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

if RA = 0 then b ← 0
else      b ← (RA)
EA ← b + (RB)
AA = EA & 0b11000
QRT0 ← 0x400 || (AA) 58:60 || 490
QRT1 ← 0x400 || (AA+ 8) 58:60 || 490
QRT2 ← 0x400 || (AA+16) 58:60 || 490
QRT3 ← 0x400 || (AA+24) 58:60 || 490
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

A quad-vector (32 bytes) describing a dynamic double-precision data alignment to be performed using the quad-vector permute instruction *qvfpperm* is generated based on the address EA.

The instruction may raise a memory translation exception if EA is not a valid address.

The behavior of this instruction is boundedly undefined when the address does not correspond to at least the natural alignment of an IEEE double precision floating point number.

**Special Registers Altered:**

None

**Programming Note**

This instruction allows the implementation of a software based alignment sequence for double-precision floating-point quad-vectors

```

qvlpcldx  qalign, ra, rb
qvlfdx   qmem1, ra, rb
qvlfdx   qmem2, ra, rb
qvfpperm qaligned, qmem1, qmem2, qalign
    
```

**Quad-Vector Load Permute Control Left  
Single indexed X-form**

qvlpclsx QRT,RA,RB

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRT | RA | RB | 518 | /  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

if RA = 0 then b ← 0
else      b ← (RA)
EA ← b + (RB)
AA = (EA * 2) & 0b11000
QRT0 ← 0x400 || (AA) 58:60 || 490
QRT1 ← 0x400 || (AA+ 8) 58:60 || 490
QRT2 ← 0x400 || (AA+16) 58:60 || 490
QRT3 ← 0x400 || (AA+24) 58:60 || 490
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

A quad-vector (32 bytes) describing a dynamic single-precision data alignment to be performed using the quad-vector permute instruction *qvfpperm* is generated based on the address EA.

The instruction may raise a memory translation exception if EA is not a valid address.

The behavior of this instruction is boundedly undefined when the address does not correspond to at least the natural alignment of an IEEE double precision floating point number.

**Special Registers Altered:**

None

**Programming Note**

This instruction allows the implementation of a software based alignment sequence for single-precision floating-point quad-vectors

```

qvlpclsx  qalign, ra, rb
qvlfsux   qmem1, ra, rb
qvlfsux   qmem2, ra, rb
qvfpperm  qaligned, qmem1, qmem2, qalign
    
```

## QPX Architecture

### Quad-Vector Load Permute Control Right Double indexed X-form

qvlpcrdx QRT,RA,RB

|   |    |     |    |    |    |    |
|---|----|-----|----|----|----|----|
| 0 | 31 | QRT | RA | RB | 70 | /  |
|   | 6  |     | 11 | 16 | 21 | 31 |

```

if RA = 0 then b ← 0
else          b ← (RA)
EA ← b + (RB)
AA = (32 - (EA & 0b11000))
QRT0 ← 0x400 || (AA) 58:60 || 490
QRT1 ← 0x400 || (AA+ 8) 58:60 || 490
QRT2 ← 0x400 || (AA+16) 58:60 || 490
QRT3 ← 0x400 || (AA+24) 58:60 || 490

```

Let the effective address (EA) be the sum (RA|0)+(RB).

A quad-vector (32 bytes) describing a dynamic data alignment to be performed using the quad-vector permute instruction *qvperm* is generated based on the address EA.

The instruction may raise a memory translation exception if EA is not a valid address.

The behavior of this instruction is boundedly undefined when the address does not correspond to at least the natural alignment of an IEEE double precision floating point number.

#### Special Registers Altered:

None

#### Programming Note

This instruction allows the implementation of a software based alignment sequence for double-precision floating-point quad-vectors.

### Quad-Vector Load Permute Control Right Single indexed X-form

qvlpcrsx QRT,RA,RB

|   |    |     |    |    |    |    |
|---|----|-----|----|----|----|----|
| 0 | 31 | QRT | RA | RB | 6  | /  |
|   | 6  |     | 11 | 16 | 21 | 31 |

```

if RA = 0 then b ← 0
else          b ← (RA)
EA ← b + (RB)
AA = (32 - ((EA * 2) & 0b11000))
QRT0 ← 0x400 || (AA) 58:60 || 490
QRT1 ← 0x400 || (AA+ 8) 58:60 || 490
QRT2 ← 0x400 || (AA+16) 58:60 || 490
QRT3 ← 0x400 || (AA+24) 58:60 || 490

```

Let the effective address (EA) be the sum (RA|0)+(RB).

A quad-vector (32 bytes) describing a dynamic data alignment to be performed using the quad-vector permute instruction *qvperm* is generated based on the address EA.

The instruction may raise a memory translation exception if EA is not a valid address.

The behavior of this instruction is boundedly undefined when the address does not correspond to at least the natural alignment of an IEEE double precision floating point number.

#### Special Registers Altered:

None

#### Programming Note

This instruction allows the implementation of a software based alignment sequence for single-precision floating-point quad-vectors.

## 4.2 Quad-Vector Floating-Point Store Instructions

### Quad-Vector Store Floating-point Single indexed X-form

qvstfsx      QRS,RA,RB      (X=0)  
 qvstfsxa    QRS,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 647 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

if RA = 0 then b ← 0
else
    b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← SINGLE(QRS0) || SINGLE(QRS1) ||
              SINGLE(QRS2) || SINGLE(QRS3)
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The four vector elements of register QRS are converted to single-precision format and stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

### Quad-Vector Store Floating-point Single with Update indexed X-form

qvstfsux      QRS,RA,RB      (X=0)  
 qvstfsuxa    QRS,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 679 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← SINGLE(QRS0) || SINGLE(QRS1) ||
              SINGLE(QRS2) || SINGLE(QRS3)
RA ← EA
    
```

Let the effective address (EA) be the sum (RA)+(RB).

The four vector elements of register QRS are converted to single-precision format and stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Store Floating-point Single indexed and Indicate X-form**

qvstfsxi QRS,RA,RB (X=0)  
 qvstfsxia QRS,RA,RB (X=1)

|   |    |     |    |    |     |    |
|---|----|-----|----|----|-----|----|
| 0 | 31 | QRS | RA | RB | 645 | X  |
|   | 6  | 11  | 16 | 21 |     | 31 |

```

if RA = 0 then b ← 0
else b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← SINGLE(QRS0) || SINGLE(QRS1) ||
                SINGLE(QRS2) || SINGLE(QRS3)

if (SNEE = 1) then
    if (isNaN (QRS0) OR
        isNaN (QRS1) OR
        isNaN (QRS2) OR
        isNaN (QRS3)) then
        ESR[AP] ← 1

if (SIEE = 1) then
    if (isInf (QRS0) OR
        isInf (QRS1) OR
        isInf (QRS2) OR
        isInf (QRS3)) then
        ESR[AP] ← 1
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The four vector elements of register QRS are converted to single-precision format and stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 ESR[AP]

**Quad-Vector Store Floating-point Single with Update indexed and Indicate X-form**

qvstfsuxi QRS,RA,RB (X=0)  
 qvstfsuxia QRS,RA,RB (X=1)

|   |    |     |    |    |     |    |
|---|----|-----|----|----|-----|----|
| 0 | 31 | QRS | RA | RB | 677 | X  |
|   | 6  | 11  | 16 | 21 |     | 31 |

```

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← SINGLE(QRS0) || SINGLE(QRS1) ||
                SINGLE(QRS2) || SINGLE(QRS3)
RA ← EA

if (SNEE = 1) then
    if (isNaN (QRS0) OR
        isNaN (QRS1) OR
        isNaN (QRS2) OR
        isNaN (QRS3)) then
        ESR[AP] ← 1

if (SIEE = 1) then
    if (isInf (QRS0) OR
        isInf (QRS1) OR
        isInf (QRS2) OR
        isInf (QRS3)) then
        ESR[AP] ← 1
    
```

Let the effective address (EA) be the sum (RA)+(RB).

The four vector elements of register QRS are converted to single-precision format and stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 ESR[AP]

**Quad-Vector Store Floating-point Double indexed X-form**

qvstfdx      QRS,RA,RB      (X=0)  
 qvstfdxa    QRS,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 711 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```
if RA = 0 then b ← 0
else          b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFE0
MEM(EA, 32) ← (QRS)
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The contents of register QRS are stored into the 32 bytes in storage addressed by the 32-byte-aligned EA.

If the X bit is set, and the address is not aligned on a 32-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Store Floating-point Double with Update indexed X-form**

qvstfdx      QRS,RA,RB      (X=0)  
 qvstfduxa    QRS,RA,RB      (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 743 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```
EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFE0
MEM(EA, 32) ← (QRS)
RA ← EA
```

Let the effective address (EA) be the sum (RA)+(RB).

The contents of register QRS are stored into the 32 bytes in storage addressed by the 32-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 32-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None



### Quad-Vector Store Floating-point Double indexed and Indicate X-form

qvstfdxi QRS,RA,RB (X=0)  
 qvstfdxia QRS,RA,RB (X=1)

|   |    |     |    |    |     |    |
|---|----|-----|----|----|-----|----|
| 0 | 31 | QRS | RA | RB | 709 | X  |
|   | 6  | 11  | 16 | 21 |     | 31 |

```

if RA = 0 then b ← 0
else          b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFE0
MEM(EA, 32) ← (QRS)
    
```

```

if (SNEE = 1) then
    if (isNaN (QRS0) OR
        isNaN (QRS1) OR
        isNaN (QRS2) OR
        isNaN (QRS3)) then
        ESR[AP] ← 1
    
```

```

if (SIEE = 1) then
    if (isInf (QRS0) OR
        isInf (QRS1) OR
        isInf (QRS2) OR
        isInf (QRS3)) then
        ESR[AP] ← 1
    
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The contents of register QRS are stored into the 32 bytes in storage addressed by the 32-byte-aligned EA.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on a 32-byte boundary, an exception is raised.

#### Special Registers Altered:

ESR[AP]

### Quad-Vector Store Floating-point Double with Update indexed and Indicate X-form

qvstfduxi QRS,RA,RB (X=0)  
 qvstfduxia QRS,RA,RB (X=1)

|   |    |     |    |    |     |    |
|---|----|-----|----|----|-----|----|
| 0 | 31 | QRS | RA | RB | 741 | X  |
|   | 6  | 11  | 16 | 21 |     | 31 |

```

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFE0
MEM(EA, 32) ← (QRS)
RA ← EA
    
```

```

if (SNEE = 1) then
    if (isNaN (QRS0) OR
        isNaN (QRS1) OR
        isNaN (QRS2) OR
        isNaN (QRS3)) then
        ESR[AP] ← 1
    
```

```

if (SIEE = 1) then
    if (isInf (QRS0) OR
        isInf (QRS1) OR
        isInf (QRS2) OR
        isInf (QRS3)) then
        ESR[AP] ← 1
    
```

Let the effective address (EA) be the sum (RA)+(RB).

The contents of register QRS are stored into the 32 bytes in storage addressed by the 32-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on a 32-byte boundary, an exception is raised.

#### Special Registers Altered:

ESR[AP]

**Quad-Vector Store Floating-point Complex Single indexed X-form**

qvstfcsx    QRS,RA,RB                    (X=0)  
 qvstfcsxa    QRS,RA,RB                    (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 135 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```
if RA = 0 then b ← 0
else      b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF8
MEM(EA, 8) ← SINGLE(QRS0) || SINGLE(QRS1)
```

Let the effective address (EA) be the sum (RA|0)+(RB).

Vector elements 0 and 1 of register QRS are converted to single-precision format and stored into the 8 bytes in storage addressed by the 8-byte-aligned EA.

If the X bit is set, and the address is not aligned on an 8-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Store Floating-point Complex Double indexed X-form**

qvstfcdx    QRS,RA,RB                    (X=0)  
 qvstfcdxa    QRS,RA,RB                    (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 199 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```
if RA = 0 then b ← 0
else      b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← QRS0 || QRS1
```

Let the effective address (EA) be the sum (RA|0)+(RB).

Vector elements 0 and 1 of register QRS are stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Store Floating-point Complex Single indexed and Indicate X-form**

qvstfcsxi QRS,RA,RB (X=0)  
 qvstfcsxia QRS,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 133 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```
if RA = 0 then b ← 0
else b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF8
MEM(EA, 8) ← SINGLE(QRS0) || SINGLE(QRS1)
```

```
if (SNEE = 1) then
    if (isNaN (QRS0) OR
        isNaN (QRS1)) then
        ESR[AP] ← 1
```

```
if (SIEE = 1) then
    if (isInf (QRS0) OR
        isInf (QRS1)) then
        ESR[AP] ← 1
```

Let the effective address (EA) be the sum (RA|0)+(RB).

Vector elements 0 and 1 of register QRS are converted to single-precision format and stored into the 8 bytes in storage addressed by the 8-byte-aligned EA.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on an 8-byte boundary, an exception is raised.

**Special Registers Altered:**  
 ESR[AP]

**Quad-Vector Store Floating-point Complex Double indexed and Indicate X-form**

qvstfcdxi QRS,RA,RB (X=0)  
 qvstfcdxia QRS,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 197 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```
if RA = 0 then b ← 0
else b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← QRS0 || QRS1
```

```
if (SNEE = 1) then
    if (isNaN (QRS0) OR
        isNaN (QRS1)) then
        ESR[AP] ← 1
```

```
if (SIEE = 1) then
    if (isInf (QRS0) OR
        isInf (QRS1)) then
        ESR[AP] ← 1
```

Let the effective address (EA) be the sum (RA|0)+(RB).

Vector elements 0 and 1 of register QRS are stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 ESR[AP]

**Quad-Vector Store Floating-point Complex Single with Update indexed X-form**

qvstfcsux QRS,RA,RB (X=0)  
 qvstfcsuxa QRS,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 167 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

$EA \leftarrow ((RA) + (RB)) \& 0xFFFFFFFFFFFFFFF8$   
 $MEM(EA, 8) \leftarrow SINGLE(QRS^0) \parallel SINGLE(QRS^1)$   
 $RA \leftarrow EA$

Let the effective address (EA) be the sum (RA)+(RB).

Vector elements 0 and 1 of register QRS are converted to single-precision format and stored into the 8 bytes in storage addressed by the 8-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on an 8-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Store Floating-point Complex Double with Update indexed X-form**

qvstfcdx QRS,RA,RB (X=0)  
 qvstfcdxa QRS,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 231 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

$EA \leftarrow ((RA) + (RB)) \& 0xFFFFFFFFFFFFFFF0$   
 $MEM(EA, 16) \leftarrow QRS^0 \parallel QRS^1$   
 $RA \leftarrow EA$

Let the effective address (EA) be the sum (RA)+(RB).

Vector elements 0 and 1 of register QRS are stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**  
 None

**Quad-Vector Store Floating-point Complex Single with Update indexed and Indicate X-form**

qvstfcsuxi QRS,RA,RB (X=0)  
 qvstfcsuxia QRS,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 165 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF8  
 MEM(EA, 8) ← SINGLE(QRS<sup>0</sup>) || SINGLE(QRS<sup>1</sup>)  
 RA ← EA

if (SNEE = 1) then  
 if (isNaN (QRS<sup>0</sup>) OR  
     isNaN (QRS<sup>1</sup>)) then  
     ESR[AP] ← 1

if (SIEE = 1) then  
 if (isInf (QRS<sup>0</sup>) OR  
     isInf (QRS<sup>1</sup>)) then  
     ESR[AP] ← 1

Let the effective address (EA) be the sum (RA)+(RB).

Vector elements 0 and 1 of register QRS are converted to single-precision format and stored into the 8 bytes in storage addressed by the 8-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on an 8-byte boundary, an exception is raised.

**Special Registers Altered:**

ESR[AP]

**Quad-Vector Store Floating-point Complex Double with Update indexed and Indicate X-form**

qvstfcduxi QRS,RA,RB (X=0)  
 qvstfcduxia QRS,RA,RB (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 229 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

EA ← ((RA) + (RB)) & 0xFFFFFFFFFFFFFFF0  
 MEM(EA, 16) ← QRS<sup>0</sup> || QRS<sup>1</sup>  
 RA ← EA

if (SNEE = 1) then  
 if (isNaN (QRS<sup>0</sup>) OR  
     isNaN (QRS<sup>1</sup>)) then  
     ESR[AP] ← 1

if (SIEE = 1) then  
 if (isInf (QRS<sup>0</sup>) OR  
     isInf (QRS<sup>1</sup>)) then  
     ESR[AP] ← 1

Let the effective address (EA) be the sum (RA)+(RB).

Vector elements 0 and 1 of register QRS are stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

If any vector element being stored is a NaN (or Infinity), and the corresponding Store NaN (or Infinity) Exception is enabled, then the Auxiliary Processor bit of the Exception Syndrome Register (ESR[AP]) is set.

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**

ESR[AP]

**Quad-Vector Store Floating-point as Integer Word indexed X-form**

qvstfiwx    QRS,RA,RB                    (X=0)  
 qvstfiwxa    QRS,RA,RB                    (X=1)

|    |     |    |    |     |    |
|----|-----|----|----|-----|----|
| 31 | QRS | RA | RB | 967 | X  |
| 0  | 6   | 11 | 16 | 21  | 31 |

```

if RA = 0 then b ← 0
else            b ← (RA)
EA ← (b + (RB)) & 0xFFFFFFFFFFFFFFF0
MEM(EA, 16) ← QRS032:63 || QRS132:63 ||
              QRS232:63 || QRS332:63
  
```

Let the effective address (EA) be the sum (RA|0)+(RB).

The least significant 32 bits of each vector element of register QRS are stored into the 16 bytes in storage addressed by the 16-byte-aligned EA.

If the contents of register QRS were produced, either directly or indirectly, by a *Load Floating-Point Single* instruction, a single-precision *Arithmetic* instruction, or *frsp*, then the value stored is undefined. (The contents of register QRS are produced directly by such an instruction if QRS is the target register for the instruction. The contents of register QRS are produced indirectly by such an instruction if QRS is the final target register of a sequence of one or more *Floating-Point Move* instructions, with the input to the sequence having been produced directly by such an instruction.)

If the X bit is set, and the address is not aligned on a 16-byte boundary, an exception is raised.

**Special Registers Altered:**

None

### 4.3 Quad-Vector Floating-Point Move Instructions

**Quad-Vector Floating-point Move Register X-form**

qvfmr QRT,QRB

|   |   |     |     |     |    |    |
|---|---|-----|-----|-----|----|----|
| 0 | 4 | QRT | /// | QRB | 72 | /  |
|   | 6 |     | 11  | 16  | 21 | 31 |

For each vector element, the contents of register QRB are placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point NEGate X-form**

qvfneg QRT,QRB

|   |   |     |     |     |    |    |
|---|---|-----|-----|-----|----|----|
| 0 | 4 | QRT | /// | QRB | 40 | /  |
|   | 6 |     | 11  | 16  | 21 | 31 |

For each vector element, the contents of register QRB, with bit 0 inverted, are placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point ABSolute value X-form**

qvfabs QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 264 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the contents of register QRB, with bit 0 set to zero, are placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Negative ABSolute value X-form**

qvfnabs QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 136 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the contents of register QRB, with bit 0 set to one, are placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point CoPy SiGN X-form**

qvfcpsgn QRT,QRA,QRB

|   |   |     |     |     |    |    |
|---|---|-----|-----|-----|----|----|
| 0 | 4 | QRT | QRA | QRB | 8  | /  |
|   | 6 |     | 11  | 16  | 21 | 31 |

- QRT<sub>0</sub> ← QRA<sub>0</sub>
- QRT<sub>1:63</sub> ← QRB<sub>1:63</sub>
- QRT<sub>64</sub> ← QRA<sub>64</sub>
- QRT<sub>65:127</sub> ← QRB<sub>65:127</sub>
- QRT<sub>128</sub> ← QRA<sub>128</sub>
- QRT<sub>129:191</sub> ← QRB<sub>129:191</sub>
- QRT<sub>192</sub> ← QRA<sub>192</sub>
- QRT<sub>193:255</sub> ← QRB<sub>193:255</sub>

For each vector element, the contents of register QRB, with bit 0 set to the value of bit 0 of register QRA, are placed into register QRT.

**Special Registers Altered:**  
None

## 4.4 Quad-Vector Floating-Point Arithmetic Instructions

### 4.4.1 Quad-Vector Floating-Point Elementary Arithmetic Instructions

#### *Quad-Vector Floating-point ADD [Single] A-form*

qvfadd QRT,QRA,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | /// | 21 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvfadds QRT,QRA,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | QRA | QRB | /// | 21 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

For each vector element, the floating-point operand in register QRA is added to the floating-point operand in register QRB.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits (G, R, and X) enter into the computation.

If a carry occurs, the sum's significand is shifted right one bit position and the exponent is increased by one.

**Special Registers Altered:**

None

#### *Quad-Vector Floating-point SUBtract [Single] A-form*

qvsub QRT,QRA,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | /// | 20 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvsubs QRT,QRA,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | QRA | QRB | /// | 20 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

For each vector element, the floating-point operand in register QRB is subtracted from the floating-point operand in register QRA.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

The execution of the Floating Subtract instruction is identical to that of Floating Add, except that the contents of QRB participate in the operation with the sign bit (bit 0) inverted.

**Special Registers Altered:**

None



## QPX Architecture

---

### Quad-Vector Floating-point MULtipl [Single] A-form

qvfmul QRT,QRA,QRC

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 4 | QRT | QRA | /// | QRC | 25 | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

qvfmuls QRT,QRA,QRC

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 0 | QRT | QRA | /// | QRC | 25 | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

For each vector element, the floating-point operand in register QRA is multiplied by the floating-point operand in register QRC.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

Floating-point multiplication is based on exponent addition and multiplication of the significands.

#### Special Registers Altered:

None

### Quad-Vector Floating-point Reciprocal Estimate [Single] A-form

qvfre QRT,QRB

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 4 | QRT | /// | QRB | /// | 24 | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

qvfres QRT,QRB

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 0 | QRT | /// | QRB | /// | 24 | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

For each vector element, an estimate of the reciprocal of the floating-point operand in register QRB is placed into register QRT. The estimate placed into register QRT is correct to a precision of one part in 16384 of the reciprocal of (QRB), i.e.,

$$\text{ABS}\left(\frac{\text{estimate} - 1/x}{1/x}\right) \leq \frac{1}{16384}$$

where x is the initial value in QRB.

Operation with various special values of the operand is summarized below.

| Operand   | Result    |
|-----------|-----------|
| $-\infty$ | -0        |
| -0        | $-\infty$ |
| +0        | $+\infty$ |
| $+\infty$ | +0        |
| SNaN      | QNaN      |
| QNaN      | QNaN      |

The results of executing this instruction may vary between implementations.

#### Special Registers Altered:

None

**Quad-Vector Floating-point Reciprocal Square Root Estimate [Single] A-form**

qvfrsqрте QRT,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | /// | QRB | /// | 26 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvfrsqрtes QRT,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | /// | QRB | /// | 26 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

For each vector element, an estimate of the reciprocal of the square root of the floating-point operand in register QRB is placed into register QRT. The estimate placed into register QRT is correct to a precision of one part in 16384 of the reciprocal of the square root of (QRB), i.e.,

$$ABS\left(\frac{\text{estimate} - 1/(\sqrt{x})}{1/(\sqrt{x})}\right) \leq \frac{1}{16384}$$

where x is the initial value in QRB.

Operation with various special values of the operand is summarized below.

| Operand   | Result    |
|-----------|-----------|
| $-\infty$ | QNaN      |
| $< 0$     | QNaN      |
| -0        | $-\infty$ |
| +0        | $+\infty$ |
| $+\infty$ | +0        |
| SNaN      | QNaN      |
| QNaN      | QNaN      |

The results of executing this instruction may vary between implementations.

**Special Registers Altered:**

None

## 4.4.2 Quad-Vector Floating-Point Multiply-Add Instructions

### Quad-Vector Floating-point Multiply-ADD [Single] A-form

qvfmmadd QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |    |
|---|---|-----|-----|-----|-----|----|----|
| 0 | 4 | QRT | QRA | QRB | QRC | 29 | /  |
|   |   | 6   | 11  | 16  | 21  | 26 | 31 |

qvfmmadds QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |    |
|---|---|-----|-----|-----|-----|----|----|
| 0 | 0 | QRT | QRA | QRB | QRC | 29 | /  |
|   |   | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow [(\text{QRA}^0) \times (\text{QRC}^0)] + (\text{QRB}^0) \\ \text{QRT}^1 &\leftarrow [(\text{QRA}^1) \times (\text{QRC}^1)] + (\text{QRB}^1) \\ \text{QRT}^2 &\leftarrow [(\text{QRA}^2) \times (\text{QRC}^2)] + (\text{QRB}^2) \\ \text{QRT}^3 &\leftarrow [(\text{QRA}^3) \times (\text{QRC}^3)] + (\text{QRB}^3) \end{aligned}$$

are performed.

For each vector element, the floating-point operand in register QRA is multiplied by the floating-point operand in register QRC. The floating-point operand in register QRB is added to this intermediate result.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

#### Special Registers Altered:

None

### Quad-Vector Floating-point Multiply-SUBtract [Single] A-form

qvfmsub QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |    |
|---|---|-----|-----|-----|-----|----|----|
| 0 | 4 | QRT | QRA | QRB | QRC | 28 | /  |
|   |   | 6   | 11  | 16  | 21  | 26 | 31 |

qvfmsubs QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |    |
|---|---|-----|-----|-----|-----|----|----|
| 0 | 0 | QRT | QRA | QRB | QRC | 28 | /  |
|   |   | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow [(\text{QRA}^0) \times (\text{QRC}^0)] - (\text{QRB}^0) \\ \text{QRT}^1 &\leftarrow [(\text{QRA}^1) \times (\text{QRC}^1)] - (\text{QRB}^1) \\ \text{QRT}^2 &\leftarrow [(\text{QRA}^2) \times (\text{QRC}^2)] - (\text{QRB}^2) \\ \text{QRT}^3 &\leftarrow [(\text{QRA}^3) \times (\text{QRC}^3)] - (\text{QRB}^3) \end{aligned}$$

are performed.

For each vector element, the floating-point operand in register QRA is multiplied by the floating-point operand in register QRC. The floating-point operand in register QRB is subtracted from this intermediate result.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

#### Special Registers Altered:

None

**Quad-Vector Floating-point Negative Multiply-ADD [Single] A-form**

qvfnmadd QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | QRC | 31 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvfnmadds QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | QRA | QRB | QRC | 31 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow - ( [(QRA^0) \times (QRC^0)] + (QRB^0) ) \\ \text{QRT}^1 &\leftarrow - ( [(QRA^1) \times (QRC^1)] + (QRB^1) ) \\ \text{QRT}^2 &\leftarrow - ( [(QRA^2) \times (QRC^2)] + (QRB^2) ) \\ \text{QRT}^3 &\leftarrow - ( [(QRA^3) \times (QRC^3)] + (QRB^3) ) \end{aligned}$$

are performed.

For each vector element, the floating-point operand in register QRA is multiplied by the floating-point operand in register QRC. The floating-point operand in register QRB is added to this intermediate result.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, then negated and placed into register QRT.

This instruction produces the same result as would be obtained by using the *qvfmadd* instruction and then negating the result, with the following exceptions.

QNaNs propagate with no effect on their “sign” bit. QNaNs that are generated as the result of a disabled Invalid Operation Exception have a “sign” bit of 0.

SNaNs that are converted to QNaNs as the result of a disabled Invalid Operation Exception retain the “sign” bit of the SNaN.

**Special Registers Altered:**

None

**Quad-Vector Floating-point Negative Multiply-SUBtract [Single] A-form**

qvfnmsub QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | QRC | 30 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvfnmsubs QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | QRA | QRB | QRC | 30 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow - ( [(QRA^0) \times (QRC^0)] - (QRB^0) ) \\ \text{QRT}^1 &\leftarrow - ( [(QRA^1) \times (QRC^1)] - (QRB^1) ) \\ \text{QRT}^2 &\leftarrow - ( [(QRA^2) \times (QRC^2)] - (QRB^2) ) \\ \text{QRT}^3 &\leftarrow - ( [(QRA^3) \times (QRC^3)] - (QRB^3) ) \end{aligned}$$

are performed.

For each vector element, the floating-point operand in register QRA is multiplied by the floating-point operand in register QRC. The floating-point operand in register QRB is subtracted from this intermediate result.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, then negated and placed into register QRT.

This instruction produces the same result as would be obtained by using the *qvfnmsub* instruction and then negating the result, with the following exceptions.

QNaNs propagate with no effect on their “sign” bit. QNaNs that are generated as the result of a disabled Invalid Operation Exception have a “sign” bit of 0.

SNaNs that are converted to QNaNs as the result of a disabled Invalid Operation Exception retain the “sign” bit of the SNaN.

**Special Registers Altered:**

None

## QPX Architecture

### Quad-Vector Floating-point cross (X) Multiply-ADD [Single] A-form

qvfxmadd QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 4 | QRT | QRA | QRB | QRC | 9  | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

qvfxmadds QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 0 | QRT | QRA | QRB | QRC | 9  | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow [(\text{QRA}^0) \times (\text{QRC}^0)] + (\text{QRB}^0) \\ \text{QRT}^1 &\leftarrow [(\text{QRA}^0) \times (\text{QRC}^1)] + (\text{QRB}^1) \\ \text{QRT}^2 &\leftarrow [(\text{QRA}^2) \times (\text{QRC}^2)] + (\text{QRB}^2) \\ \text{QRT}^3 &\leftarrow [(\text{QRA}^2) \times (\text{QRC}^3)] + (\text{QRB}^3) \end{aligned}$$

are performed.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

#### Special Registers Altered:

None

#### Programming Note

This instruction is typically used in cross-product multiplication, and in conjunction with qvfxnpxmadd.

### Quad-Vector Floating-point double-cross complex (XXNP) Multiply-ADD [Single] A-form

qvfxnpxmadd QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 4 | QRT | QRA | QRB | QRC | 11 | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

qvfxnpxmadds QRT,QRA,QRC,QRB

|   |   |     |     |     |     |    |   |
|---|---|-----|-----|-----|-----|----|---|
| 0 | 0 | QRT | QRA | QRB | QRC | 11 | / |
|   | 6 | 11  | 16  | 21  | 26  | 31 |   |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow - ( [(\text{QRA}^1) \times (\text{QRC}^1)] - (\text{QRB}^0) ) \\ \text{QRT}^1 &\leftarrow [(\text{QRA}^0) \times (\text{QRC}^1)] + (\text{QRB}^1) \\ \text{QRT}^2 &\leftarrow - ( [(\text{QRA}^3) \times (\text{QRC}^3)] - (\text{QRB}^2) ) \\ \text{QRT}^3 &\leftarrow [(\text{QRA}^2) \times (\text{QRC}^3)] + (\text{QRB}^3) \end{aligned}$$

are performed.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR. For vector elements 0 and 2, the rounded result is negated and placed into register QRT. For vector elements 1 and 3, the rounded result is placed into register QRT.

#### Special Registers Altered:

None

#### Programming Note

This instruction is typically used in cross-product multiplication of complex numbers, in conjunction with qvfxmul or qvfxmadd.

$$\begin{bmatrix} M_0 + N_0j & M_1 + N_1j \\ M_2 + N_2j & M_3 + N_3j \end{bmatrix} \begin{bmatrix} P_0 + Q_0j & P_1 + Q_1j \\ P_2 + Q_2j & P_3 + Q_3j \end{bmatrix}$$

Consecutive Memory Locations:  $M_0 \ N_0 \ M_1 \ N_1 \ M_2 \ N_2 \ M_3 \ N_3$   
Separate from above, but consecutive in memory:  $P_0 \ Q_0 \ P_1 \ Q_1 \ P_2 \ Q_2 \ P_3 \ Q_3$

A 2x2 matrix MN times a 2x2 matrix PQ produces a resultant 2x2 matrix R

$$\begin{aligned} \text{Entry Row}i\text{Column}j \text{ of the Resultant Matrix R} \\ = (M_i + N_0j)(P_j + Q_0j) + (M_1 + N_1j)(P_1 + Q_1j) \\ = M_iP_j + M_0Q_0j + N_0P_jj - N_0Q_0 + M_1P_1 + M_1Q_1j + N_1P_1j - N_1Q_1 \end{aligned}$$

|       |                      |                      |                      |                      |
|-------|----------------------|----------------------|----------------------|----------------------|
|       | Element <sup>0</sup> | Element <sup>1</sup> | Element <sup>2</sup> | Element <sup>3</sup> |
| QPR20 | M <sub>0</sub>       | N <sub>0</sub>       | M <sub>1</sub>       | N <sub>1</sub>       |

|       |                      |                      |                      |                      |
|-------|----------------------|----------------------|----------------------|----------------------|
|       | Element <sup>0</sup> | Element <sup>1</sup> | Element <sup>2</sup> | Element <sup>3</sup> |
| QPR21 | P <sub>0</sub>       | Q <sub>0</sub>       | P <sub>1</sub>       | Q <sub>1</sub>       |

qvfxmul QPR22, QPR20, QPR21 (notice A=QPR20 and C=QPR21) yields:

|       |                               |                               |                               |                               |
|-------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
|       | Element <sup>0</sup>          | Element <sup>1</sup>          | Element <sup>2</sup>          | Element <sup>3</sup>          |
| QPR22 | M <sub>0</sub> P <sub>0</sub> | M <sub>0</sub> Q <sub>0</sub> | M <sub>1</sub> P <sub>1</sub> | M <sub>1</sub> Q <sub>1</sub> |

qvfxnpxmadd QPR23, QPR21, QPR20, QPR22 (notice A=QPR21 and C=QPR20) yields:

|       |   |   |   |   |
|-------|---|---|---|---|
|       | Element <sup>0</sup>  | Element <sup>1</sup>  | Element <sup>2</sup>  | Element <sup>3</sup>  |
| QPR23 | M <sub>0</sub> P <sub>0</sub> - N <sub>0</sub> Q <sub>0</sub> | M <sub>0</sub> Q <sub>0</sub> + N <sub>0</sub> P <sub>0</sub> | M <sub>1</sub> P <sub>1</sub> - N <sub>1</sub> Q <sub>1</sub> | M <sub>1</sub> Q <sub>1</sub> + N <sub>1</sub> P <sub>1</sub> |

Now need to add Element<sup>0</sup>+Element<sup>2</sup> and Element<sup>1</sup>+Element<sup>3</sup>

**Quad-Vector Floating-point double-cross conjugate (XXCPN) Multiply-ADD [Single] A-form**

qvfxcpnmadd QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | QRC | 3  | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvfxcpnmadds QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | QRA | QRB | QRC | 3  | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} QRT^0 &\leftarrow [(QRA^1) \times (QRC^1)] + (QRB^0) \\ QRT^1 &\leftarrow - ( [(QRA^0) \times (QRC^1)] - (QRB^1) ) \\ QRT^2 &\leftarrow [(QRA^3) \times (QRC^3)] + (QRB^2) \\ QRT^3 &\leftarrow - ( [(QRA^2) \times (QRC^3)] - (QRB^3) ) \end{aligned}$$

are performed.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR. For vector elements 0 and 2, the rounded result is placed into register QRT. For vector elements 1 and 3, the rounded result is negated and placed into register QRT.

**Special Registers Altered:**

None

**Quad-Vector Floating-point double-cross (XX) Multiply-ADD [Single] A-form**

qvfxmadd QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | QRC | 1  | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

qvfxmadds QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 0 | QRT | QRA | QRB | QRC | 1  | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} QRT^0 &\leftarrow [(QRA^1) \times (QRC^1)] + (QRB^0) \\ QRT^1 &\leftarrow [(QRA^0) \times (QRC^1)] + (QRB^1) \\ QRT^2 &\leftarrow [(QRA^3) \times (QRC^3)] + (QRB^2) \\ QRT^3 &\leftarrow [(QRA^2) \times (QRC^3)] + (QRB^3) \end{aligned}$$

are performed.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

**Special Registers Altered:**

None

## QPX Architecture

---

### **Quad-Vector Floating-point cross (X) MULTiPLY [Single] A-form**

qvfxmul QRT,QRA,QRC

|   |   |     |     |     |     |    |    |
|---|---|-----|-----|-----|-----|----|----|
| 0 | 4 | QRT | QRA | /// | QRC | 17 | /  |
|   |   | 6   | 11  | 16  | 21  | 26 | 31 |

qvfxmuls QRT,QRA,QRC

|   |   |     |     |     |     |    |    |
|---|---|-----|-----|-----|-----|----|----|
| 0 | 0 | QRT | QRA | /// | QRC | 17 | /  |
|   |   | 6   | 11  | 16  | 21  | 26 | 31 |

The operations

$$\begin{aligned} \text{QRT}^0 &\leftarrow (\text{QRA}^0) \times (\text{QRC}^0) \\ \text{QRT}^1 &\leftarrow (\text{QRA}^1) \times (\text{QRC}^1) \\ \text{QRT}^2 &\leftarrow (\text{QRA}^2) \times (\text{QRC}^2) \\ \text{QRT}^3 &\leftarrow (\text{QRA}^3) \times (\text{QRC}^3) \end{aligned}$$

are performed.

For each vector element, the resultant significand may require normalization or denormalization, depending on the values of the two most significant bits of the resultant significand and on the value of the resultant exponent. The result is rounded to the target precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

#### **Special Registers Altered:**

None

## 4.5 Quad-Vector Floating-Point Rounding and Conversion Instructions

### 4.5.1 Quad-Vector Floating-Point Rounding Instruction

**Quad-Vector Floating-point Round to Single-Precision** *X-form*

qvfrsp      QRT,QRB

|   |     |     |     |    |    |
|---|-----|-----|-----|----|----|
| 4 | QRT | /// | QRB | 12 | /  |
| 0 | 6   | 11  | 16  | 21 | 31 |

For each vector element, the floating-point operand in register QRB is rounded to single-precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

**Special Registers Altered:**

None



## 4.5.2 Quad-Vector Floating-Point Convert To/From Integer Instructions

### *Quad-Vector Floating-point Convert To Integer Doubleword* X-form

qvftcid QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 814 | /  |
|   | 6 |     | 11  |     | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer under control of the Floating-Point Rounding Control field (RN) of the FPSCR.

For each vector element, if the rounded floating-point integer is greater than  $2^{63} - 1$ , then QRT is set to `0x7FFF_FFFF_FFFF_FFFF`.

For each vector element, if the rounded floating-point integer is less than  $-2^{63}$ , then QRT is set to `0x8000_0000_0000_0000`.

Otherwise, for each vector element, QRT is set to the value of the rounded floating-point integer converted to 64-bit signed-integer format.

**Special Registers Altered:**

None

### *Quad-Vector Floating-point Convert To Integer Doubleword Unsigned* X-form

qvftcidu QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 942 | /  |
|   | 6 |     | 11  |     | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer under control of the Floating-Point Rounding Control field (RN) of the FPSCR.

For each vector element, if the rounded floating-point integer is greater than  $2^{64} - 1$ , then QRT is set to `0xFFFF_FFFF_FFFF_FFFF`.

For each vector element, if the rounded floating-point integer is less than 0.0, then QRT is set to `0x0000_0000_0000_0000`.

Otherwise, for each vector element, QRT is set to the value of the rounded floating-point integer converted to 64-bit unsigned-integer format.

**Special Registers Altered:**

None

### *Quad-Vector Floating-point Convert To Integer Doubleword with round toward Zero* X-form

qvftcidz QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 815 | /  |
|   | 6 |     | 11  |     | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward Zero.

For each vector element, if the rounded floating-point integer is greater than  $2^{63} - 1$ , then QRT is set to `0x7FFF_FFFF_FFFF_FFFF`.

For each vector element, if the rounded floating-point integer is less than  $-2^{63}$ , then QRT is set to `0x8000_0000_0000_0000`.

Otherwise, for each vector element, QRT is set to the value of the rounded floating-point integer converted to 64-bit signed-integer format.

**Special Registers Altered:**

None

### *Quad-Vector Floating-point Convert To Integer Doubleword Unsigned with round toward Zero* X-form

qvftciduz QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 943 | /  |
|   | 6 |     | 11  |     | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward Zero.

For each vector element, if the rounded floating-point integer is greater than  $2^{64} - 1$ , then QRT is set to `0xFFFF_FFFF_FFFF_FFFF`.

For each vector element, if the rounded floating-point integer is less than 0.0, then QRT is set to `0x0000_0000_0000_0000`.

Otherwise, for each vector element, QRT is set to the value of the rounded floating-point integer converted to 64-bit unsigned-integer format.

**Special Registers Altered:**

None

**Quad-Vector Floating-point Convert To Integer Word X-form**

qvftiw QRT,QRB

|   |   |     |     |     |    |    |
|---|---|-----|-----|-----|----|----|
| 0 | 4 | QRT | /// | QRB | 14 | /  |
|   | 6 |     | 11  | 16  | 21 | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer under control of the Floating-Point Rounding Control field (RN) of the FPSCR.

For each vector element, if the rounded floating-point integer is greater than  $2^{31} - 1$ , then QRT<sub>32:63</sub> is set to 0x7FFF\_FFFF.

For each vector element, if the rounded floating-point integer is less than  $-2^{31}$ , then QRT<sub>32:63</sub> is set to 0x8000\_0000.

Otherwise, for each vector element, QRT<sub>32:63</sub> is set to the value of the rounded floating-point integer converted to 32-bit signed-integer format.

QRT<sub>0:31</sub> of each vector element is undefined.

**Special Registers Altered:**

None

**Implementation Note**

In the QPU of BGQ, for each vector element, QRT<sub>0:31</sub> ← 0x7FF80000

**Quad-Vector Floating-point Convert To Integer Word Unsigned X-form**

qvftiwu QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 142 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer under control of the Floating-Point Rounding Control field (RN) of the FPSCR.

For each vector element, if the rounded floating-point integer is greater than  $2^{32} - 1$ , then QRT<sub>32:63</sub> is set to 0xFFFF\_FFFF.

For each vector element, if the rounded floating-point integer is less than 0.0, then QRT<sub>32:63</sub> is set to 0x0000\_0000.

Otherwise, for each vector element, QRT<sub>32:63</sub> is set to the value of the rounded floating-point integer converted to 32-bit unsigned-integer format.

QRT<sub>0:31</sub> of each vector element is undefined.

**Special Registers Altered:**

None

**Implementation Note**

In the QPU of BGQ, for each vector element, QRT<sub>0:31</sub> ← 0x7FF80000

**Quad-Vector Floating-point Convert To Integer Word with round toward Zero X-form**

qvftiwz QRT,QRB

|   |   |     |     |     |    |   |
|---|---|-----|-----|-----|----|---|
|   | 4 | QRT | /// | QRB | 15 | / |
| 0 | 6 | 11  | 16  | 21  | 31 |   |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward Zero.

For each vector element, if the rounded floating-point integer is greater than  $2^{31} - 1$ , then QRT<sub>32:63</sub> is set to 0x7FFF\_FFFF.

For each vector element, if the rounded floating-point integer is less than  $-2^{31}$ , then QRT<sub>32:63</sub> is set to 0x8000\_0000.

Otherwise, for each vector element, QRT<sub>32:63</sub> is set to the value of the rounded floating-point integer converted to 32-bit signed-integer format.

QRT<sub>0:31</sub> of each vector element is undefined.

**Special Registers Altered:**

None

**Implementation Note**

In the QPU of BGQ, for each vector element, QRT<sub>0:31</sub> ← 0x7FF80000

**Quad-Vector Floating-point Convert To Integer Word Unsigned with round toward Zero X-form**

qvftiwuz QRT,QRB

|   |   |     |     |     |     |   |
|---|---|-----|-----|-----|-----|---|
|   | 4 | QRT | /// | QRB | 143 | / |
| 0 | 6 | 11  | 16  | 21  | 31  |   |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward Zero.

For each vector element, if the rounded floating-point integer is greater than  $2^{32} - 1$ , then QRT<sub>32:63</sub> is set to 0xFFFF\_FFFF.

For each vector element, if the rounded floating-point integer is less than 0.0, then QRT<sub>32:63</sub> is set to 0x0000\_0000.

Otherwise, for each vector element, QRT<sub>32:63</sub> is set to the value of the rounded floating-point integer converted to 32-bit unsigned-integer format.

QRT<sub>0:31</sub> of each vector element is undefined.

**Special Registers Altered:**

None

**Implementation Note**

In the QPU of BGQ, for each vector element, QRT<sub>0:31</sub> ← 0x7FF80000

**Quad-Vector Floating-point Convert From Integer Doubleword X-form**

qvfcfid QRT,QRB

|   |     |     |     |     |    |
|---|-----|-----|-----|-----|----|
| 4 | QRT | /// | QRB | 846 | /  |
| 0 | 6   | 11  | 16  | 21  | 31 |

For each vector element, the 64-bit signed fixed-point operand in register QRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to double-precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Convert From Integer Doubleword Unsigned X-form**

qvfcfidu QRT,QRB

|   |     |     |     |     |    |
|---|-----|-----|-----|-----|----|
| 4 | QRT | /// | QRB | 974 | /  |
| 0 | 6   | 11  | 16  | 21  | 31 |

For each vector element, the 64-bit unsigned fixed-point operand in register QRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to double-precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Convert From Integer Doubleword Single X-form**

qvfcfids QRT,QRB

|   |     |     |     |     |    |
|---|-----|-----|-----|-----|----|
| 0 | QRT | /// | QRB | 846 | /  |
| 0 | 6   | 11  | 16  | 21  | 31 |

For each vector element, the 64-bit signed fixed-point operand in register QRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to single-precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Convert From Integer Doubleword Unsigned Single X-form**

qvfcfidus QRT,QRB

|   |     |     |     |     |    |
|---|-----|-----|-----|-----|----|
| 0 | QRT | /// | QRB | 974 | /  |
| 0 | 6   | 11  | 16  | 21  | 31 |

For each vector element, the 64-bit unsigned fixed-point operand in register QRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to single-precision under control of the Floating-Point Rounding Control field (RN) of the FPSCR, and placed into register QRT.

**Special Registers Altered:**  
None

### 4.5.3 Quad-Vector Floating-Point Round to Integer Instructions

**Quad-Vector Floating-point Round to Integer Nearest** *X-form*

qvfrin QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 392 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer as follows, with the result placed into register QRT. If the sign of the operand is positive, (QRB) + 0.5 is truncated to a floating-point integer, otherwise (QRB) - 0.5 is truncated to a floating-point integer.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Round to Integer Plus** *X-form*

qvfrip QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 456 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward +Infinity, and the result is placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Round to Integer toward Zero** *X-form*

qvfriz QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 424 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward Zero, and the result is placed into register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Floating-point Round to Integer Minus** *X-form*

qvfrim QRT,QRB

|   |   |     |     |     |     |    |
|---|---|-----|-----|-----|-----|----|
| 0 | 4 | QRT | /// | QRB | 488 | /  |
|   | 6 |     | 11  | 16  | 21  | 31 |

For each vector element, the floating-point operand in register QRB is rounded to a floating-point integer using the rounding mode Round toward -Infinity, and the result is placed into register QRT.

**Special Registers Altered:**  
None

## 4.6 Quad-Vector Floating-Point Compare Instructions

### Quad-Vector Floating-point TeST for NAN X-form

qvftstnan QRT,QRA,QRB

|   |   |     |     |     |    |    |
|---|---|-----|-----|-----|----|----|
| 0 | 4 | QRT | QRA | QRB | 64 | /  |
|   | 6 | 11  | 16  | 21  |    | 31 |

```

if isNaN(QRA0) OR isNaN(QRB0)
  then QRT0 ← 0x3FF0_0000_0000_0000
  else QRT0 ← 0xBFF0_0000_0000_0000
if isNaN(QRA1) OR isNaN(QRB1)
  then QRT1 ← 0x3FF0_0000_0000_0000
  else QRT1 ← 0xBFF0_0000_0000_0000
if isNaN(QRA2) OR isNaN(QRB2)
  then QRT2 ← 0x3FF0_0000_0000_0000
  else QRT2 ← 0xBFF0_0000_0000_0000
if isNaN(QRA3) OR isNaN(QRB3)
  then QRT3 ← 0x3FF0_0000_0000_0000
  else QRT3 ← 0xBFF0_0000_0000_0000
  
```

Each vector element is compared for the specified condition, and the result is encoded. The Boolean value TRUE is encoded as 1.0. The Boolean value of FALSE is encoded as -1.0.

**Special Registers Altered:**

None

### Quad-Vector Floating-point CoMPare Greater Than X-form

qvfcmpgt QRT,QRA,QRB

|   |   |     |     |     |    |    |
|---|---|-----|-----|-----|----|----|
| 0 | 4 | QRT | QRA | QRB | 32 | /  |
|   | 6 | 11  | 16  | 21  |    | 31 |

```

if (QRA0) > (QRB0)
  then QRT0 ← 0x3FF0_0000_0000_0000
  else QRT0 ← 0xBFF0_0000_0000_0000
if (QRA1) > (QRB1)
  then QRT1 ← 0x3FF0_0000_0000_0000
  else QRT1 ← 0xBFF0_0000_0000_0000
if (QRA2) > (QRB2)
  then QRT2 ← 0x3FF0_0000_0000_0000
  else QRT2 ← 0xBFF0_0000_0000_0000
if (QRA3) > (QRB3)
  then QRT3 ← 0x3FF0_0000_0000_0000
  else QRT3 ← 0xBFF0_0000_0000_0000
  
```

Each vector element is compared for the specified condition, and the result is encoded. The Boolean value TRUE is encoded as 1.0. The Boolean value of FALSE is encoded as -1.0.

When one of the operands is a NaN, the value -1.0 (FALSE) is returned.

**Special Registers Altered:**

None

### Quad-Vector Floating-point CoMPare Less Than X-form

qvfcmlpt QRT,QRA,QRB

|   |   |     |     |     |    |   |
|---|---|-----|-----|-----|----|---|
|   | 4 | QRT | QRA | QRB | 96 | / |
| 0 | 6 | 11  | 16  | 21  | 31 |   |

```

if (QRA0) < (QRB0)
  then QRT0 ← 0x3FF0_0000_0000_0000
  else QRT0 ← 0xBFF0_0000_0000_0000
if (QRA1) < (QRB1)
  then QRT1 ← 0x3FF0_0000_0000_0000
  else QRT1 ← 0xBFF0_0000_0000_0000
if (QRA2) < (QRB2)
  then QRT2 ← 0x3FF0_0000_0000_0000
  else QRT2 ← 0xBFF0_0000_0000_0000
if (QRA3) < (QRB3)
  then QRT3 ← 0x3FF0_0000_0000_0000
  else QRT3 ← 0xBFF0_0000_0000_0000
  
```

Each vector element is compared for the specified condition, and the result is encoded. The Boolean value TRUE is encoded as 1.0. The Boolean value of FALSE is encoded as -1.0.

When one of the operands is a NaN, the value -1.0 (FALSE) is returned.

**Special Registers Altered:**  
None

### Quad-Vector Floating-point CoMPare Equal X-form

qvfcmeq QRT,QRA,QRB

|   |   |     |     |     |    |   |
|---|---|-----|-----|-----|----|---|
|   | 4 | QRT | QRA | QRB | 0  | / |
| 0 | 6 | 11  | 16  | 21  | 31 |   |

```

if (QRA0) = (QRB0)
  then QRT0 ← 0x3FF0_0000_0000_0000
  else QRT0 ← 0xBFF0_0000_0000_0000
if (QRA1) = (QRB1)
  then QRT1 ← 0x3FF0_0000_0000_0000
  else QRT1 ← 0xBFF0_0000_0000_0000
if (QRA2) = (QRB2)
  then QRT2 ← 0x3FF0_0000_0000_0000
  else QRT2 ← 0xBFF0_0000_0000_0000
if (QRA3) = (QRB3)
  then QRT3 ← 0x3FF0_0000_0000_0000
  else QRT3 ← 0xBFF0_0000_0000_0000
  
```

Each vector element is compared for the specified condition, and the result is encoded. The Boolean value TRUE is encoded as 1.0. The Boolean value of FALSE is encoded as -1.0.

When one of the operands is a NaN, the value -1.0 (FALSE) is returned.

**Special Registers Altered:**  
None

## 4.7 Quad Floating-Point Select Instruction

### *Quad-Vector Floating-point SElectA-form*

qvfsel      QRT,QRA,QRC,QRB

|   |     |     |     |     |    |    |
|---|-----|-----|-----|-----|----|----|
| 4 | QRT | QRA | QRB | QRC | 23 | /  |
| 0 | 6   | 11  | 16  | 21  | 26 | 31 |

```

if (QRA0) ≥ 0.0
  then QRT0 ← (QRC0)
  else QRT0 ← (QRB0)
if (QRA1) ≥ 0.0
  then QRT1 ← (QRC1)
  else QRT1 ← (QRB1)
if (QRA2) ≥ 0.0
  then QRT2 ← (QRC2)
  else QRT2 ← (QRB2)
if (QRA3) ≥ 0.0
  then QRT3 ← (QRC3)
  else QRT3 ← (QRB3)

```

For each vector element, the floating-point operand in register QRA is compared to the value zero. If the operand is greater than or equal to zero, register QRT is set to the contents of register QRC. If the operand is less than zero or is a NaN, register QRT is set to the contents of register QRB. The comparison ignores the sign of zero (i.e., regards +0 as equal to -0).

**Special Registers Altered:**

None



## 4.8 Quad-Vector Alignment and Formatting Instructions

### Quad-Vector ALIGN Immediate

Z23-form

qvaligni QRT,QRA,QRB,VD

|   |   |     |     |     |    |   |    |
|---|---|-----|-----|-----|----|---|----|
| 0 | 4 | QRT | QRA | QRB | VD | 5 | /  |
|   | 6 | 11  | 16  | 21  | 23 |   | 31 |

```

if VD = 00 then
    QRT ← (QRA)
else if VD = 01 then
    QRT ← (QRA1) || (QRA2) || (QRA3) || (QRB0)
else if VD = 10 then
    QRT ← (QRA2) || (QRA3) || (QRB0) || (QRB1)
else if VD = 11 then
    QRT ← (QRA3) || (QRB0) || (QRB1) || (QRB2)
    
```

The contents of registers QRA and QRB are concatenated, and a quad-vector is extracted starting at the vector element specified by field VD. The resulting quad-vector is placed into register QRT.

**Special Registers Altered:**

None

### Quad-Vector Floating-point PERMute A-form

qvfperm QRT,QRA,QRB,QRC

|   |   |     |     |     |     |   |    |
|---|---|-----|-----|-----|-----|---|----|
| 0 | 4 | QRT | QRA | QRB | QRC | 6 | /  |
|   | 6 | 11  | 16  | 21  | 26  |   | 31 |

For each vector element,

```

if QRC1:11 = 0x400 then
    case QRC12:14
        QRT ← (QRA0) when 000
        QRT ← (QRA1) when 001
        QRT ← (QRA2) when 010
        QRT ← (QRA3) when 011
        QRT ← (QRB0) when 100
        QRT ← (QRB1) when 101
        QRT ← (QRB2) when 110
        QRT ← (QRB3) when 111
    
```

else

QRT ← Undefined

The contents of registers QRA and QRB are concatenated. A quad-vector is composed from vector elements extracted from the concatenated registers, as specified by the contents of register QRC.

**Special Registers Altered:**

None

**Quad-Vector Element SPLAT Immediate  
Z23-form**

qvesplati QRT,QRA,VD

|   |     |     |    |       |    |   |
|---|-----|-----|----|-------|----|---|
| 4 | QRT | QRA | // | VD    | 37 | / |
| 0 | 6   | 11  | 16 | 21 23 | 31 |   |

```

if VD = 00 then
    QRT ← (QRA0) || (QRA0) || (QRA0) || (QRA0)
else if VD = 01 then
    QRT ← (QRA1) || (QRA1) || (QRA1) || (QRA1)
else if VD = 10 then
    QRT ← (QRA2) || (QRA2) || (QRA2) || (QRA2)
else if VD = 11 then
    QRT ← (QRA3) || (QRA3) || (QRA3) || (QRA3)
    
```

The vector element from register QRA, specified by field VD, is placed into each vector element of register QRT.

**Special Registers Altered:**  
None

**Quad-Vector Generate Permute Control  
Immediate  
Z23-form**

qvgpci QRT,GPC

|   |     |     |     |    |
|---|-----|-----|-----|----|
| 4 | QRT | GPC | 133 | /  |
| 0 | 6   | 11  | 23  | 31 |

```

QRT0 ← 0x400 || GPC0:2 || 490
QRT1 ← 0x400 || GPC3:5 || 490
QRT2 ← 0x400 || GPC6:8 || 490
QRT3 ← 0x400 || GPC9:11 || 490
    
```

Register QRT is loaded with the 12-bit immediate field GPC, dispersed across its four elements, to serve as control for a QVFPERM instruction.

**Special Registers Altered:**  
None

## 4.9 Floating-Point Boolean Instruction

### Quad-Vector Floating-point boolean LOGICAL *X-form*

qvflogical QRT,QRA,QRB,TT

|   |     |     |     |    |    |    |
|---|-----|-----|-----|----|----|----|
| 4 | QRT | QRA | QRB | TT | 4  | /  |
| 0 | 6   | 11  | 16  | 21 | 25 | 31 |

For each vector element,

```

if [(QRA) < 0.0 OR isNaN(QRA)] AND
  [(QRB) < 0.0 OR isNaN(QRB)] then
  if TT0 = 1 then QRT ← 0x3FF0_0000_0000_0000
  else QRT ← 0xBFF0_0000_0000_0000
if [(QRA) ≥ 0.0] AND
  [(QRB) < 0.0 OR isNaN(QRB)] then
  if TT1 = 1 then QRT ← 0x3FF0_0000_0000_0000
  else QRT ← 0xBFF0_0000_0000_0000
if [(QRA) < 0.0 OR isNaN(QRA)] AND
  [(QRB) ≥ 0.0] then
  if TT2 = 1 then QRT ← 0x3FF0_0000_0000_0000
  else QRT ← 0xBFF0_0000_0000_0000
if [(QRA) ≥ 0.0] AND
  [(QRB) ≥ 0.0] then
  if TT3 = 1 then QRT ← 0x3FF0_0000_0000_0000
  else QRT ← 0xBFF0_0000_0000_0000
    
```

The floating-point operands in registers QRA and QRB are treated as boolean values of TRUE if greater than or equal to +/- 0.0, and as FALSE if less than 0.0 or a NaN. Immediate field TT is used in conjunction with these values to create a logical operation.

#### Programming Note

Some common logical operations can be accessed via pseudo mnemonics, expressed in the table below.

| Extended Mnemonic     | Equivalent                | Function                   |
|-----------------------|---------------------------|----------------------------|
| qvfcir QRT            | qvflogical QRT,QRT,QRT,0  | clear (set as FALSE)       |
| qvfcand QRT,QRA,QRB   | qvflogical QRT,QRA,QRB,1  | and                        |
| qvfcandc QRT,QRA,QRB  | qvflogical QRT,QRA,QRB,4  | and complement B           |
| qvfcftb QRT,QRA       | qvflogical QRT,QRA,QRA,5  | convert to float-boolean A |
| qvfcxor QRT,QRA,QRB   | qvflogical QRT,QRA,QRB,6  | xor                        |
| qvfcor QRT,QRA,QRB    | qvflogical QRT,QRA,QRB,7  | or                         |
| qvfcnor QRT,QRA,QRB   | qvflogical QRT,QRA,QRB,8  | nor                        |
| qvfcfequ QRT,QRA,QRB  | qvflogical QRT,QRA,QRB,9  | Boolean equivalent (XNOR)  |
| qvfcfnor QRT,QRA      | qvflogical QRT,QRA,QRA,10 | not                        |
| qvfcforc QRT,QRA,QRB  | qvflogical QRT,QRA,QRB,13 | or complement B            |
| qvfcfnand QRT,QRA,QRB | qvflogical QRT,QRA,QRB,14 | nand                       |
| qvfcset QRT           | qvflogical QRT,QRT,QRT,15 | set (set as TRUE)          |