

# xCORE-200 DSP Library

This API reference manual describes the XMOS fixed-point digital signal processing software library. The library implements a suite of common signal processing functions for use on XMOS xCORE-200 multi-core microcontrollers.

---

## Required tools and libraries

- xTIMEcomposer Tools Version 14.0.1 or later

## Required hardware

Only XMOS xCORE-200 based multicore microcontrollers are supported with this library. The previous generation XS1 based multicore microcontrollers are not supported.

The xCORE-200 has a single cycle 32x32->64 bit multiply/accumulate unit, single cycle double-word load and store, dual issue instruction execution, and other instruction set enhancements. These features make xCORE-200 an efficient platform for executing digital signal processing algorithms.

## Prerequisites

This document assumes familiarity with the XMOS xCORE architecture, the XMOS tool chain, the 'C' programming language, and digital signal processing concepts.

## Software version and dependencies

This document pertains to version 1.0.0 of this library. It is known to work on version 14.1.1 of the xTIMEcomposer tools suite, it may work on other versions.

The library does not have any dependencies (i.e. it does not rely on any other libraries).

## Related application notes

The following application notes use this library:

- AN00209 - xCORE-200 DSP Library

# 1 Overview

## 1.1 Introduction

This API reference manual describes the XMOS xCORE-200 fixed-point digital signal processing firmware library. The library implements a suite of common signal processing functions for use on XMOS xCORE-200 multicore microcontrollers.

## 1.2 Library Organization

The library is divided into function collections with each collection covering a specific digital signal processing algorithm category. The API and implementation for each category are provided by a single 'C' header file and implementation file.

Category	Source Files	Functions
Fixed point	lib_dsp_qformat	Q16 through Q31 formats, fixed and floating point conversions
Filters	lib_dsp_filters	FIR, biquad, cascaded biquad, and convolution
Adaptive	lib_dsp_adaptive	LMS and NLMS Adaptive filters
Scalar math	lib_dsp_math	Multiply, square root, reciprocal, inverse square root
Vector math	lib_dsp_vector	Scalar/vector add/subtract/multiply, dot product
Matrix math	lib_dsp_matrix	Scalar/matrix add/subtract/multiply, inverse and transpose
Statistics	lib_dsp_statistics	Vector mean, sum-of-squares, root-mean-square, variance

Table 1: DSP library organization :class: narrow

## 2 Fixed-Point Format

### 2.1 Q Format Introduction

The library functions support 32 bit input and output data, with internal 64 bit accumulator. The output data can be scaled to any of the supported Q Formats (Q16 through Q31). Further details about Q Format numbers is available here : [https://en.wikipedia.org/wiki/Q\\_\(number\\_format\)](https://en.wikipedia.org/wiki/Q_(number_format)).

### 2.2 The 'q\_format' Parameter

All XMOS DSP library functions that incorporate a multiply operation accept a parameter called q\_format. This parameter can naively be used to specify the fixed point format for all operands and results (if applicable) where the formats are the same for all parameters. For example:

```
result_q28 = lib_dsp_math_multiply( input1_q28, input2_q28, 28 );
```

The 'q\_format' parameter, being used after one or more sequences of multiply and/or multiply-accumulate, is used to right-shift the 64-bit accumulator before truncating the value back to a 32-bit integer (i.e. the 32-bit fixed-point result). Therefore the 'q\_format' parameter can be used to perform the proper fixed-point adjustment for any combination of input operand fixed-point format and desired fixed-point result format.

The output fixed-point fraction bit count is equal to the sum of the two input fraction bit counts minus the desired result fraction bit count:

```
q_format = input1 fraction bit count + input2 fraction bit count - result fraction bit count
```

For example:

```
// q_format_parameter = 31 = 30 + 29 - 28
result_q28 = lib_dsp_math_multiply( input1_q30, input2_q29, 31 );

// q_format_parameter = 27 = 28 + 29 - 30
result_q30 = lib_dsp_math_multiply( input1_q28, input2_q29, 27 );
```

### 3 Filter Functions: Finite Impulse Response (FIR) Filter

<b>Function</b>	<b>lib_dsp_filters_fir</b>
<b>Description</b>	<p>This function implements a Finite Impulse Response (FIR) filter. The function operates on a single sample of input and output data (i.e. each call to the function processes one sample). The FIR filter algorithm is based upon a sequence of multiply-accumulate (MAC) operations. Each filter coefficient <math>h[i]</math> is multiplied by a state variable which equals a previous input sample <math>x[i]</math>, or <math>y[n]=x[n]*h[0]+x[n-1]*h[1]+x[n-2]*h[2]+x[n-N+1]*h[N-1]</math>. The parameter <code>filter_coeffs</code> points to a coefficient array of size <math>N = \text{num\_taps}</math>. The filter coefficients are stored in forward order (e.g. <math>h[0], h[1], h[N-1]</math>). The following example shows a five-tap (4th order) FIR filter with samples and coefficients represented in Q28 fixed-point format.</p> <pre>int filter_coeff[5] = { Q28(0.5),Q(-0.5),Q28(0.0),Q28(-0.5),Q28(0.5) }; int filter_state[4] = { 0, 0, 0, 0 }; int result = lib_dsp_fir( sample, filter_coeff, filter_state, 5, 28 );</pre> <p>The FIR algorithm involves multiplication between 32-bit filter coefficients and 32-bit state data producing a 64-bit result for each coefficient and state data pair. Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered.</p>
<b>Type</b>	<pre>int lib_dsp_filters_fir(int input_sample,                    const int filter_coeffs[],                    int state_data[],                    int tap_count,                    int q_format)</pre>
<b>Parameters</b>	<p><code>input_sample</code>      The new sample to be processed.</p> <p><code>filter_coeffs</code>      Pointer to FIR coefficients array arranged as <math>[b_0, b_1, b_2, b_{N-1}]</math>.</p> <p><code>state_data</code>          Pointer to filter state data array of length <math>N</math>. Must be initialized at startup to all zeros.</p> <p><code>tap_count</code>          Filter tap count (<math>N = \text{tap\_count} = \text{filter order} + 1</math>).</p> <p><code>q_format</code>            Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The resulting filter output sample.

## 4 Filter Functions: Interpolating FIR Filter

<b>Function</b>	<b>lib_dsp_filters_interpolate</b>
<b>Description</b>	<p>This function implements an interpolating FIR filter. The function operates on a single input sample and outputs a set of samples representing the interpolated data, whose sample count is equal to <code>interp_factor</code>. (i.e. and each call to the function processes one sample and results in <code>interp_factor</code> output samples).</p> <p>The FIR filter algorithm is based upon a sequence of multiply-accumulate (MAC) operations. Each filter coefficient <code>h[i]</code> is multiplied by a state variable which equals a previous input sample <code>x[i]</code>, or <math>y[n]=x[n]*h[0]+x[n-1]*h[1]+x[n-2]*h[2]+x[n-N+1]*h[N-1]</math></p> <p><code>filter_coeffs</code> points to a coefficient array of size <code>N = num_taps</code>. The filter coefficients are stored in forward order (e.g. <code>h[0]</code>, <code>h[1]</code>, <code>h[N-1]</code>).</p> <p>The FIR algorithm involves multiplication between 32-bit filter coefficients and 32-bit state data producing a 64-bit result for each coefficient and state data pair. Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered.</p>
<b>Type</b>	<pre>void lib_dsp_filters_interpolate(int input_sample,                            const int filter_coeffs[],                            int state_data[],                            int tap_count,                            int interp_factor,                            int output_samples[],                            int q_format)</pre>

*Continued on next page*

Parameters	
	<p><b>input_sample</b> The new sample to be processed.</p>
	<p><b>filter_coeffs</b> Pointer to FIR coefficients array arranged as:  <math>h_M, h_{(1L+M)}, h_{(2L+M)}, h_{((N-1)L+M)}, h_1, h_{(1L+1)}, h_{(2L+1)}, h_{((N-1)L+1)}, h_0, h_{(1L+0)}, h_{(2L+0)}, h_{((N-1)L+0)}</math>,                      where <math>M = N-1</math></p>
	<p><b>state_data</b> Pointer to filter state data array of length N. Must be initialized at startup to all zeros.</p>
	<p><b>tap_count</b> Filter tap count (<math>N = \text{tap\_count} = \text{filter order} + 1</math>).</p>
	<p><b>interp_factor</b> The interpolation factor/index (i.e. the up-sampling ratio). The interpolation factor/index can range from 2 to 16.</p>
	<p><b>output_samples</b> The resulting interpolated samples.</p>
	<p><b>q_format</b> Fixed point format (i.e. number of fractional bits).</p>

## 5 Filter Functions: Decimating FIR Filter

Function	<code>lib_dsp_filters_decimate</code>
<b>Description</b>	<p>This function implements an decimating FIR filter.</p> <p>The function operates on a single set of input samples whose count is equal to the decimation factor. (i.e. and each call to the function processes <code>decim_factor</code> samples and results in one sample).</p> <p>The FIR filter algorithm is based upon a sequence of multiply-accumulate (MAC) operations. Each filter coefficient <code>h[i]</code> is multiplied by a state variable which equals a previous input sample <code>x[i]</code>, or <math>y[n]=x[n]*h[0]+x[n-1]*h[1]+x[n-2]*h[2]+x[n-N+1]*h[N-1]</math></p> <p><code>filter_coeffs</code> points to a coefficient array of size <code>N = num_taps</code>. The filter coefficients are stored in forward order (e.g. <code>h[0]</code>, <code>h[1]</code>, <code>h[N-1]</code>).</p> <p>The FIR algorithm involves multiplication between 32-bit filter coefficients and 32-bit state data producing a 64-bit result for each coefficient and state data pair. Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered.</p>
<b>Type</b>	<pre>int lib_dsp_filters_decimate(int input_samples[],                         const int filter_coeffs[],                         int state_data[],                         int tap_count,                         int decim_factor,                         int q_format)</pre>
<b>Parameters</b>	<p><code>input_samples</code> The new samples to be decimated.</p> <p><code>filter_coeffs</code> Pointer to FIR coefficients array arranged as <code>[b0,b1,b2,bN-1]</code>.</p> <p><code>state_data</code> Pointer to filter state data array of length <code>N</code>. Must be initialized at startup to all zeros.</p> <p><code>tap_count</code> Filter tap count (<code>N = tap_count = filter order + 1</code>).</p> <p><code>decim_factor</code> The decimation factor/index (i.e. the down-sampling ratio).</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The resulting decimated sample.

## 6 Filter Functions: Bi-Quadratic (BiQuad) IIR Filter

Function	lib_dsp_filters_biquad
<b>Description</b>	<p>This function implements a second order IIR (direct form I). The function operates on a single sample of input and output data (i.e. and each call to the function processes one sample). The IIR filter algorithm is based upon a sequence of multiply-accumulate (MAC) operations. Each filter coefficient <math>b[i]</math> is multiplied by a state variable which equals a previous input sample <math>x[i]</math>, or <math>y[i]=x[n]*b[0]+x[n-1]*b[1]+x[n-2]*b2+x[n-1]*a[1]+x[n-2]*a[2]</math>. The filter coefficients are stored in forward order (e.g. <math>b_0, b_1, b_2, a_1, a_2</math>). Example showing a single Biquad filter with samples and coefficients represented in Q28 fixed-point format:</p> <pre>int filter_coeff[5] = { Q28(+0.5), Q(-0.1), Q28(-0.5), Q28(-0.1), Q28     ↪ (0.1) }; int filter_state[4] = { 0, 0, 0, 0 }; int result = lib_dsp_biquad( sample, filter_coeff, filter_state, 28 );</pre> <p>The IIR algorithm involves multiplication between 32-bit filter coefficients and 32-bit state data producing a 64-bit result for each coefficient and state data pair. Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered.</p>
<b>Type</b>	<pre>int lib_dsp_filters_biquad(int input_sample,                       const int filter_coeffs[],                       int state_data[],                       int q_format)</pre>
<b>Parameters</b>	<p><b>input_sample</b> The new sample to be processed.</p> <p><b>filter_coeffs</b> Pointer to biquad coefficients array arranged as <math>[b_0, b_1, b_2, a_1, a_2]</math>.</p> <p><b>state_data</b> Pointer to filter state data array (initialized at startup to zeros). The length of the state data array is 4.</p> <p><b>q_format</b> Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The resulting filter output sample.



## 7 Filter Functions: Cascaded BiQuad Filter

Function	lib_dsp_filters_biquads
<b>Description</b>	<p>This function implements a cascaded direct form I BiQuad filter. The function operates on a single sample of input and output data (i.e. and each call to the function processes one sample).</p> <p>The IIR filter algorithm is based upon a sequence of multiply-accumulate (MAC) operations. Each filter coefficient <math>b[i]</math> is multiplied by a state variable which equals a previous input sample <math>x[i]</math>, or <math>y[n]=x[n]*b[0]+x[n-1]*b[1]+x[n-2]*b2+x[n-1]*a[1]+x[n-2]*a[2]</math></p> <p>The filter coefficients are stored in forward order (e.g. section1:b0,b1,b2,a1,a2,sectionN:b0,b1,b2,a1,a2).</p> <p>Example showing a 4x cascaded Biquad filter with samples and coefficients represented in Q28 fixed-point format:</p> <pre>int filter_coeff[20] = { Q28(+0.5), Q(-0.1), Q28(-0.5), Q28(-0.1), Q28(0.1),                       Q28(+0.5), Q(-0.1), Q28(-0.5), Q28(-0.1), Q28(0.1),                       Q28(+0.5), Q(-0.1), Q28(-0.5), Q28(-0.1), Q28(0.1),                       Q28(+0.5), Q(-0.1), Q28(-0.5), Q28(-0.1), Q28(0.1) }; int filter_state[16] = { 0,0,0,0, 0,0,0,0, 0,0,0,0, 0,0,0,0 }; int result = lib_dsp_cascaded_biquad( sample, filter_coeff, filter_state, 4, 28 );</pre> <p>The IIR algorithm involves multiplication between 32-bit filter coefficients and 32-bit state data producing a 64-bit result for each coefficient and state data pair. Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered.</p>
<b>Type</b>	<pre>int lib_dsp_filters_biquads(int input_sample,                        const int filter_coeffs[],                        int state_data[],                        int num_sections,                        int q_format)</pre>
<b>Parameters</b>	<p><b>input_sample</b> The new sample to be processed.</p> <p><b>filter_coeffs</b> Pointer to biquad coefficients array for all BiQuad sections. Arranged as [section1:b0,b1,b2,a1,a2, . . . sectionN:b0,b1,b2,a1,a2].</p> <p><b>state_data</b> Pointer to filter state data array (initialized at startup to zeros). The length of the state data array is <math>\text{num\_sections} * 4</math>.</p> <p><b>num_sections</b> Number of BiQuad sections.</p> <p><b>q_format</b> Fixed point format (i.e. number of fractional bits).</p>

*Continued on next page*

<b>Returns</b>	The resulting filter output sample.
----------------	-------------------------------------

## 8 Adaptive Filter Functions: LMS Adaptive Filter

Function	lib_dsp_adaptive_lms
<b>Description</b>	<p>This function implements a least-mean-squares adaptive FIR filter. LMS filters are a class of adaptive filters that adjust filter coefficients in order to create a transfer function that minimizes the error between the input and reference signals. FIR coefficients are adjusted on a per sample basis by an amount calculated from the given step size and the instantaneous error.</p> <p>The function operates on a single sample of input and output data (i.e. and each call to the function processes one sample and each call results in changes to the FIR coefficients). The general LMS algorithm, on a per sample basis, is to:</p> <ol style="list-style-type: none"> <li>1) Apply the transfer function: <math>output = FIR( input )</math></li> <li>2) Compute the instantaneous error value: <math>error = reference - output</math></li> <li>3) Compute current coefficient adjustment delta: <math>delta = mu * error</math></li> <li>4) Adjust transfer function coefficients:  <math>FIR\_COEFFS[n] = FIR\_COEFFS[n] + FIR\_STATE[n] * delta</math></li> </ol> <p>Example of a 100-tap LMS filter with samples and coefficients represented in Q28 fixed-point format:</p> <pre>int filter_coeff[100] = { ... not shown for brevity }; int filter_state[100] = { 0, 0, 0, 0, ... not shown for brevity };  int output_sample = lib_dsp_adaptive_lms (     input_sample, reference_sample, &amp;error_sample,     filter_coeff_array, filter_state_array, 100, Q28(0.01), 28 );</pre> <p>The LMS filter algorithm involves multiplication between two 32-bit values and 64-bit accumulation as a result of using an FIR as well as coefficient step size calculations). Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered for both FIR operations as well as for coefficient step size calculation and FIR coefficient adjustment.</p>
<b>Type</b>	<pre>int lib_dsp_adaptive_lms(int input_sample,                     int reference_sample,                     int error_sample[],                     int filter_coeffs[],                     int state_data[],                     int tap_count,                     int step_size,                     int q_format)</pre>

*Continued on next page*

<b>Parameters</b>	<p><code>input_sample</code> The new sample to be processed.</p> <p><code>reference_sample</code> Reference sample.</p> <p><code>error_sample</code> Pointer to resulting error sample (error = reference - output)</p> <p><code>filter_coeffs</code> Pointer to FIR coefficients arranged as [b0,b1,b2, ...,bN-1].</p> <p><code>state_data</code> Pointer to FIR filter state data array of length N. Must be initialized at startup to all zeros.</p> <p><code>tap_count</code> Filter tap count where <math>N = \text{tap\_count} = \text{filter order} + 1</math>.</p> <p><code>step_size</code> Coefficient adjustment step size, controls rate of convergence.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The resulting filter output sample.

## 9 Adaptive Filter Functions: Normalized LMS Filter

Function	lib_dsp_adaptive_nlms
<b>Description</b>	<p>This function implements a normalized LMS FIR filter. LMS filters are a class of adaptive filters that adjust filter coefficients in order to create the a transfer function that minimizes the error between the input and reference signals. FIR coefficients are adjusted on a per sample basis by an amount calculated from the given step size and the instantaneous error. The function operates on a single sample of input and output data (i.e. and each call to the function processes one sample and each call results in changes to the FIR coefficients).</p> <p>The general NLMS algorithm, on a per sample basis, is to:</p> <ol style="list-style-type: none"> <li>1) Apply the transfer function: <math>output = FIR( input )</math></li> <li>2) Compute the instantaneous error value: <math>error = reference - output</math></li> <li>3) Normalize the error using the instantaneous power computed by:  <math>E = x[n]^2 + \dots + x[n-N+1]^2</math></li> <li>4) Update error value: <math>error = (reference - output) / E</math></li> <li>5) Compute current coefficient adjustment delta: <math>delta = mu * error</math></li> <li>6) Adjust transfer function coefficients:  <math>FIR\_COEFFS[n] = FIR\_COEFFS[n] + FIR\_STATE[n] * delta</math></li> </ol> <p>Example of a 100-tap NLMS filter with samples and coefficients represented in Q28 fixed-point format:</p> <pre>int filter_coeff[100] = { ... not shown for brevity }; int filter_state[100] = { 0, 0, 0, 0, ... not shown for brevity };  int output_sample = lib_dsp_adaptive_nlms (     input_sample, reference_sample, &amp;error_sample,     filter_coeff_array, filter_state_array, 100, Q28(0.01), 28 );</pre> <p>The LMS filter algorithm involves multiplication between two 32-bit values and 64-bit accumulation as a result of using an FIR as well as coefficient step size calculations). Multiplication results are accumulated in 64-bit accumulator with the final result shifted to the required fixed-point format. Therefore overflow behavior of the 32-bit multiply operation and truncation behavior from final shifting of the accumulated multiplication results must be considered for both FIR operations as well as for coefficient step size calculation and FIR coefficient adjustment. Computing the coefficient adjustment involves taking the reciprocal of the instantaneous power computed by <math>E = x[n]^2 + x[n-1]^2 + \dots + x[n-N+1]^2</math>. The reciprocal is subject to overflow since the instantaneous power may be less than one.</p>
<b>Type</b>	<pre>int lib_dsp_adaptive_nlms(int input_sample,                     int reference_sample,                     int error_sample[],                     int filter_coeffs[],                     int state_data[],                     int tap_count,                     int step_size,                     int q_format)</pre>

*Continued on next page*

<b>Parameters</b>	<p><code>input_sample</code> The new sample to be processed.</p> <p><code>reference_sample</code> Reference sample.</p> <p><code>error_sample</code> Pointer to resulting error sample (error = reference - output)</p> <p><code>filter_coeffs</code> Pointer to FIR coefficients arranged as [b0,b1,b2, ...,bN-1].</p> <p><code>state_data</code> Pointer to FIR filter state data array of length N. Must be initialized at startup to all zeros.</p> <p><code>tap_count</code> Filter tap count where <math>N = \text{tap\_count} = \text{filter order} + 1</math>.</p> <p><code>step_size</code> Coefficient adjustment step size, controls rate of convergence.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The resulting filter output sample.

## 10 Scalar Math Functions: Multiply

<b>Function</b>	<b>lib_dsp_math_multiply</b>
<b>Description</b>	<p>Scalar multiplication.</p> <p>This function multiplies two scalar values and produces a result according to fixed-point format specified by the <code>q_format</code> parameter.</p> <p>The two operands are multiplied to produce a 64-bit result which is tested for overflow, clamped at the minimum/maximum value given the fixed-point format if overflow occurs, and finally shifted right by <code>q_format</code> bits.</p> <p>Algorithm:</p> <ol style="list-style-type: none"> <li>1) <math>Y = X1 * X2</math></li> <li>2) <math>Y = \min(\max(Q\_FORMAT\_MIN, Y), Q\_FORMAT\_MAX, Y)</math></li> <li>3) <math>Y = Y \gg q\_format</math></li> </ol> <p>Example:</p> <pre>int result; result = lib_dsp_math_multiply( Q28(-0.33), sample, 28 );</pre> <p>While saturation is employed after multiplication an overflow condition when preparing the final result must still be considered when specifying a Q-format whose fixed-point numerical range do not accomodate the final result of multiplication and saturation (if applied).</p>
<b>Type</b>	<pre>int lib_dsp_math_multiply(int input1_value,                     int input2_value,                     int q_format)</pre>
<b>Parameters</b>	<pre>input1_value    Multiply operand #1. input2_value    Multiply operand #2. q_format        Fixed point format (i.e. number of fractional bits).</pre>
<b>Returns</b>	<code>input1_value * input2_value.</code>

## 11 Scalar Math Functions: Reciprocal

<b>Function</b>	<b>lib_dsp_math_reciprocal</b>
<b>Description</b>	<p>Scalar reciprocal.</p> <p>This function computes the reciprocal of the input value using an iterative approximation method as follows:</p> <ol style="list-style-type: none"> <li>1) result = 1.0</li> <li>2) result = result + result * (1 - input * result)</li> <li>3) Repeat step #2 until desired precision is achieved</li> </ol> <p>Example:</p> <pre>int result; result = lib_dsp_math_reciprocal( sample, 28 );</pre>
<b>Type</b>	<p>int</p> <pre>lib_dsp_math_reciprocal(int input_value, int q_format)</pre>
<b>Parameters</b>	<p>input_value      Input value for computation.</p> <p>q_format      Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The reciprocal of the input value.



## 12 Scalar Math Functions: Inverse Square Root

<b>Function</b>	<b>lib_dsp_math_invsqrt</b>
<b>Description</b>	<p>Scalar inverse square root.                  This function computes the reciprocal of the square root of the input value using an iterative approximation method as follows:</p> <ol style="list-style-type: none"> <li>1) <math>result = 1.0</math></li> <li>2) <math>result = result + result * (1 - input * result^2) / 2</math></li> <li>3) Repeat step #2 until desired precision is achieved</li> </ol> <p>Example:</p> <pre>int result; result = lib_dsp_math_invsqrt( sample, 28 );</pre>
<b>Type</b>	<p>int                  lib_dsp_math_invsqrt(int input_value, int q_format)</p>
<b>Parameters</b>	<p>input_value      Input value for computation.</p> <p>q_format      Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The inverse square root of the input value.

## 13 Scalar Math Functions: Square Root

<b>Function</b>	<b>lib_dsp_math_squareroot</b>
<b>Description</b>	<p>Scalar square root.                  This function computes the square root of the input value using the following steps:</p> <pre>int result; result = lib_dsp_math_invsqrtroot( input ) result = lib_dsp_math_reciprocal( result )</pre> <p>Example:</p> <pre>int result; result = lib_dsp_math_squareroot( sample, 28 );</pre>
<b>Type</b>	<pre>int lib_dsp_math_squareroot(int input_value, int q_format)</pre>
<b>Parameters</b>	<pre>input_value</pre> <p>Input value for computation.</p> <pre>q_format</pre> <p>Fixed point format (i.e. number of fractional bits).</p>
<b>Returns</b>	The square root of the input value.

## 14 Vector Math Functions: Minimum Value

<b>Function</b>	<b>lib_dsp_vector_minimum</b>
<b>Description</b>	<p>Vector Minimum.                      Locate the vector's first occurring minimum value, returning the index of the first occurring minimum value.                      Example:</p> <pre>int samples[256]; int result = lib_dsp_vector_minimum( samples, 256 );</pre>
<b>Type</b>	<pre>int lib_dsp_vector_minimum(const int input_vector[],                        int vector_length)</pre>
<b>Parameters</b>	<p><b>input_vector</b>                      Pointer to source data array.</p> <p><b>vector_length</b>                      Length of the output and input arrays.</p>
<b>Returns</b>	Array index where first minimum value occurs.

## 15 Vector Math Functions: Maximum Value

<b>Function</b>	<b>lib_dsp_vector_maximum</b>
<b>Description</b>	Vector Minimum. Locate the vector's first occurring maximum value, returning the index of the first occurring maximum value. Example: <pre>int samples[256]; int result = lib_dsp_vector_maximum( samples, 256 );</pre>
<b>Type</b>	int lib_dsp_vector_maximum(const int input_vector[], int vector_length)
<b>Parameters</b>	input_vector Pointer to source data array.  vector_length Length of the output and input arrays.
<b>Returns</b>	Array index where first maximum value occurs.

## 16 Vector Math Functions: Element Negation

<b>Function</b>	<b>lib_dsp_vector_negate</b>
<b>Description</b>	<p>Vector negation: <math>R[i] = -X[i]</math>.</p> <p>This function computes the negative value for each input element and sets the corresponding result element to its negative value.</p> <p>Each negated element is computed by twos-compliment negation therefore the minimum negative fixed-point value can not be negated to generate its corresponding maximum positive fixed-point value. For example: -Q28(-8.0) will not result in a fixed-point value representing +8.0.</p> <p>Example:</p> <pre>int samples[256]; int result[256]; lib_dsp_vector_negate( samples, result, 256 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_negate(const int input_vector_X[],                     int result_vector_R[],                     int vector_length)</pre>
<b>Parameters</b>	<p><b>input_vector_X</b> Pointer/reference to source data.</p> <p><b>result_vector_R</b> Pointer to the resulting data array.</p> <p><b>vector_length</b> Length of the input and output vectors.</p>

## 17 Vector Math Functions: Element Absolute Value

<b>Function</b>	<b>lib_dsp_vector_abs</b>
<b>Description</b>	<p>Vector absolute value: <math>R[i] =  X[i] </math>.                      Set each element of the result vector to the absolute value of the corresponding input vector element.                      Example:</p> <pre>int samples[256]; int result[256]; lib_dsp_vector_abs( samples, result, 256 );</pre> <p>If an element is less than zero it is negated to compute its absolute value. Negation is computed via twos-compliment negation therefore the minimum negative fixed-point value can not be negated to generate its corresponding maximum positive fixed-point value. For example: -Q28(-8.0) will not result in a fixed-point value representing +8.0.</p>
<b>Type</b>	<pre>void lib_dsp_vector_abs(const int input_vector_X[],                   int result_vector_R[],                   int vector_length)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code>                      Pointer/reference to source data.</p> <p><code>result_vector_R</code>                      Pointer to the resulting data array.</p> <p><code>vector_length</code>                      Length of the input and output vectors.</p>

## 18 Vector Math Functions: Scalar Addition

<b>Function</b>	<b>lib_dsp_vector_adds</b>
<b>Description</b>	<p>Vector / scalar addition: <math>R[i] = X[i] + A</math>.</p> <p>This function adds a scalar value to each vector element. 32-bit addition is used to compute the scalar plus vector element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_scalar_A = Q28( 0.333 ); int result_vector_R[256]; lib_dsp_vector_adds( input_vector_X, scalar_value_A, result_vector_R,     ↪ 256 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_adds(const int input_vector_X[],     int input_scalar_A,     int result_vector_R[],     int vector_length)</pre>
<b>Parameters</b>	<p><b>input_vector_X</b> Pointer/reference to source data array X</p> <p><b>input_scalar_A</b> Scalar value to add to each input element</p> <p><b>result_vector_R</b> Pointer to the resulting data array</p> <p><b>vector_length</b> Length of the input and output vectors</p>

## 19 Vector Math Functions: Scalar Multiplication

<b>Function</b>	<b>lib_dsp_vector_muls</b>
<b>Description</b>	<p>Vector / scalar multiplication: <math>R[i] = X[i] * A</math>.                      32-bit addition is used to compute the scalar plus vector element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.                      Example:</p> <pre>int input_vector_X[256]; int input_scalar_A = Q28( 0.333 ); int result_vector_R[256]; lib_dsp_vector_adds( input_vector_X, scalar_value_A, result_vector_R,                     ↪ 256 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_muls(const int input_vector_X[],                    int input_scalar_A,                    int result_vector_R[],                    int vector_length,                    int q_format)</pre>
<b>Parameters</b>	<p><b>input_vector_X</b>                      Pointer/reference to source data array X.</p> <p><b>input_scalar_A</b>                      Scalar value to multiply each element by.</p> <p><b>result_vector_R</b>                      Pointer to the resulting data array.</p> <p><b>vector_length</b>                      Length of the input and output vectors.</p> <p><b>q_format</b>    Fixed point format, the number of bits making up fractional part.</p>



## 20 Vector Math Functions: Vector Addition

<b>Function</b>	<b>lib_dsp_vector_addv</b>
<b>Description</b>	<p>Vector / vector addition: <math>R[i] = X[i] + Y[i]</math>.                      32-bit addition is used to compute the scalar plus vector element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.                      Example:</p> <pre> int input_vector_X[256]; int input_vector_Y[256]; int result_vector_R[256]; lib_dsp_vector_addv( input_vector_X, input_vector_Y, result_vector_R,                     ↪ 256 );                     </pre>
<b>Type</b>	<pre> void lib_dsp_vector_addv(const int input_vector_X[],                    const int input_vector_Y[],                    int result_vector_R[],                    int vector_length)                     </pre>
<b>Parameters</b>	<p><b>input_vector_X</b>                      Pointer to source data array X.</p> <p><b>input_vector_Y</b>                      Pointer to source data array Y.</p> <p><b>result_vector_R</b>                      Pointer to the resulting data array.</p> <p><b>vector_length</b>                      Length of the input and output vectors.</p>

## 21 Vector Math Functions: Vector Subtraction

<b>Function</b>	<b>lib_dsp_vector_subv</b>
<b>Description</b>	<p>Vector / vector subtraction: <math>R[i] = X[i] - Y[i]</math>.                      32-bit subtraction is used to compute the scaler plus vector element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.                      Example:</p> <pre>                     int input_vector_X[256];                     int input_vector_Y[256];                     int result_vector_R[256];                     lib_dsp_vector_subv( input_vector_X, input_vector_Y, result_vector_R,                     ↪ 256 );                     </pre>
<b>Type</b>	<pre>                     void                     lib_dsp_vector_subv(const int input_vector_X[],                     const int input_vector_Y[],                     int result_vector_R[],                     int vector_length)                     </pre>
<b>Parameters</b>	<p><b>input_vector_X</b>                      Pointer to source data array X.</p> <p><b>input_vector_Y</b>                      Pointer to source data array Y.</p> <p><b>result_vector_R</b>                      Pointer to the resulting data array.</p> <p><b>vector_length</b>                      Length of the input and output vectors.</p>

## 22 Vector Math Functions: Vector Multiplication

<b>Function</b>	<b>lib_dsp_vector_mulv</b>
<b>Description</b>	<p>Vector / vector multiplication: <math>R[i] = X[i] * Y[i]</math>.                      Elements in each of the input vectors are multiplied together using a 32-bit multiply 64-bit accumulate function therefore fixed-point multiplication and q-format adjustment overflow behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>).</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_vector_Y[256]; int result_vector_R[256]; lib_dsp_vector_mulv( input_vector_X, input_vector_Y, result_vector_R,     ↪ 256, 28 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_mulv(const int input_vector_X[],     const int input_vector_Y[],     int result_vector_R[],     int vector_length,     int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code>                      Pointer to source data array X.</p> <p><code>input_vector_Y</code>                      Pointer to source data array Y.</p> <p><code>result_vector_R</code>                      Pointer to the resulting data array.</p> <p><code>vector_length</code>                      Length of the input and output vectors.</p> <p><code>q_format</code>    Fixed point format (i.e. number of fractional bits).</p>

## 23 Vector Math Functions: Vector multiplication and scalar addition

Function	<code>lib_dsp_vector_mulv_adds</code>
<b>Description</b>	<p>Vector multiplication and scalar addition: <math>R[i] = X[i] * Y[i] + A</math>. Elements in each of the input vectors are multiplied together using a 32-bit multiply 64-bit accumulate function therefore fixed-point multiplication and q-format adjustment overflow behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). 32-bit addition is used to compute the vector element plus scalar value result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_vector_Y[256]; int input_scalar_A = Q28( 0.333 ); int result_vector_R[256]; lib_dsp_vector_mulv_adds( input_vector_X, input_vector_Y, scalar_value_A,     ↪ result_vector_R, 256, 28 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_mulv_adds(const int input_vector_X[],     const int input_vector_Y[],     int input_scalar_A,     int result_vector_R[],     int vector_length,     int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>input_vector_Y</code> Pointer to source data array Y.</p> <p><code>input_scalar_A</code> Scalar value to add to each X*Y result.</p> <p><code>result_vector_R</code> Pointer to the resulting data array.</p> <p><code>vector_length</code> Length of the input and output vectors.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

## 24 Vector Math Functions: Scalar multiplication and vector addition

Function	<code>lib_dsp_vector_muls_addv</code>
<b>Description</b>	<p>Scalar multiplication and vector addition: <math>R[i] = X[i] * A + Y[i]</math>. Each element in the input vectors is multiplied by a scalar using a 32bit multiply 64-bit accumulate function therefore fixed-point multiplication and q-format adjustment overflow behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). 32-bit addition is used to compute the vector element minus vector element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_scalar_A = Q28( 0.333 ); int input_vector_Y[256]; int result_vector_R[256]; lib_dsp_vector_muls_addv( input_vector_X, input_scalar_A, input_vector_Y,     ↪ result_vector_R, 256, 28 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_muls_addv(const int input_vector_X[],     int input_scalar_A,     const int input_vector_Y[],     int result_vector_R[],     int vector_length,     int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>input_scalar_A</code> Scalar value to multiply each element by.</p> <p><code>input_vector_Y</code> Pointer to source data array Y</p> <p><code>result_vector_R</code> Pointer to the resulting data array.</p> <p><code>vector_length</code> Length of the input and output vectors</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

## 25 Vector Math Functions: Scalar multiplication and vector subtraction

<b>Function</b>	<b>lib_dsp_vector_muls_subv</b>
<b>Description</b>	<p>Scalar multiplication and vector subtraction: <math>R[i] = X[i] * A - Y[i]</math>. Each element in the input vectors is multiplied by a scalar using a 32bit multiply 64-bit accumulate function therefore fixed-point multiplication and q-format adjustment overflow behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). 32-bit subtraction is used to compute the vector element minus vector element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_scalar_A = Q28( 0.333 ); int input_vector_Y[256]; int result_vector_R[256]; lib_dsp_vector_muls_subv( input_vector_X, input_scalar_A, input_vector_Y,     ↪ result_vector_R, 256, 28 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_muls_subv(const int input_vector_X[],     int input_scalar_A,     const int input_vector_Y[],     int result_vector_R[],     int vector_length,     int q_format)</pre>
<b>Parameters</b>	<p><code>input_scalar_A</code> Scalar value to multiply each element by.</p> <p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>input_vector_Y</code> Pointer to source data array Y.</p> <p><code>result_vector_R</code> Pointer to the resulting data array.</p> <p><code>vector_length</code> Length of the input and output vectors.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

## 26 Vector Math Functions: Vector multiplication and vector addition

<b>Function</b>	<b>lib_dsp_vector_mulv_addv</b>
<b>Description</b>	<p>Vector multiplication and vector addition: <math>R[i] = X[i] * Y[i] + Z[i]</math>. The elements in the input vectors are multiplied before being summed therefore fixed-point multiplication behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). Due to successive 32-bit additions being accumulated using 64-bit arithmetic overflow during the summation process is unlikely. The final value, being effectively the result of a left-shift by <code>q_format</code> bits will potentially overflow the final fixed-point value depending on the resulting summed value and the chosen Q-format.</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_vector_Y[256]; int input_vector_Z[256]; int result_vector_R[256]; lib_dsp_vector_mulv_subv( input_vector_X, input_vector_Y, input_vector_Z,     ↪ result_vector_R, 256 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_mulv_addv(const int input_vector_X[],     const int input_vector_Y[],     const int input_vector_Z[],     int result_vector_R[],     int vector_length,     int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>input_vector_Y</code> Pointer to source data array Y.</p> <p><code>input_vector_Z</code> Pointer to source data array Z.</p> <p><code>result_vector_R</code> Pointer to the resulting data array.</p> <p><code>vector_length</code> Length of the input and output vectors.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

## 27 Vector Math Functions: Vector multiplication and vector subtraction

<b>Function</b>	<b>lib_dsp_vector_mulv_subv</b>
<b>Description</b>	<p>Vector multiplication and vector addition: <math>R[i] = X[i] * Y[i] - Z[i]</math>. The elements in the input vectors are multiplied before being subtracted therefore fixed-point multiplication behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). Due to successive 32-bit subtractions being accumulated using 64-bit arithmetic overflow during the summation process is unlikely. The final value, being effectively the result of a left-shift by <code>q_format</code> bits will potentially overflow the final fixed-point value depending on the resulting summed value and the chosen Q-format.</p> <p>Example:</p> <pre>int input_vector_X[256]; int input_vector_Y[256]; int input_vector_Z[256]; int result_vector_R[256]; lib_dsp_vector_mulv_subv( input_vector_X, input_vector_Y, input_vector_Z,     ↪ result_vector_R, 256 );</pre>
<b>Type</b>	<pre>void lib_dsp_vector_mulv_subv(const int input_vector_X[],     const int input_vector_Y[],     const int input_vector_Z[],     int result_vector_R[],     int vector_length,     int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>input_vector_Y</code> Pointer to source data array Y.</p> <p><code>input_vector_Z</code> Pointer to source data array Z.</p> <p><code>result_vector_R</code> Pointer to the resulting data array.</p> <p><code>vector_length</code> Length of the input and output vectors.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>



## 28 Matrix Math Functions: Element Negation

<b>Function</b>	<b>lib_dsp_matrix_negate</b>
<b>Description</b>	<p>Matrix negation: <math>R[i][j] = -X[i][j]</math>.                      Each negated element is computed by twos-compliment negation therefore the minimum negative fixed-point value can not be negated to generate its corresponding maximum positive fixed-point value. For example: -Q28(-8.0) will not result in a fixed-point value representing +8.0.                      Example:</p> <pre>int samples[8][32]; int result[8][32]; lib_dsp_matrix_negate( samples, result, 8, 32 );</pre>
<b>Type</b>	<pre>void lib_dsp_matrix_negate(const int input_matrix_X[],                     int result_matrix_R[],                     int row_count,                     int column_count)</pre>
<b>Parameters</b>	<p><code>input_matrix_X</code>                      Pointer/reference to source data.</p> <p><code>result_matrix_R</code>                      Pointer to the resulting 2-dimensional data array.</p> <p><code>row_count</code>    Number of rows in input matrix.</p> <p><code>column_count</code>                      Number of columns in input matrix.</p>

## 29 Matrix Math Functions: Scalar Addition

<b>Function</b>	<b>lib_dsp_matrix_adds</b>
<b>Description</b>	<p>Matrix / scalar addition: <math>R[i][j] = X[i][j] + A</math>.                      32-bit addition is used to compute the scalar plus matrix element result. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.                      Example:</p> <pre>int input_matrix_X[8][32]; int input_scalar_A = Q28( 0.333 ); int result_vector_R[8][32]; lib_dsp_matrix_adds( input_matrix_X, scalar_matrix_A, result_matrix_R,                     ↪ 8, 32 );</pre>
<b>Type</b>	<pre>void lib_dsp_matrix_adds(const int input_matrix_X[],                     int input_scalar_A,                     int result_matrix_R[],                     int row_count,                     int column_count)</pre>
<b>Parameters</b>	<p><b>input_matrix_X</b>                      Pointer/reference to source data.</p> <p><b>input_scalar_A</b>                      Scalar value to add to each input element.</p> <p><b>result_matrix_R</b>                      Pointer to the resulting 2-dimensional data array.</p> <p><b>row_count</b>    Number of rows in input and output matrices.</p> <p><b>column_count</b>                      Number of columns in input and output matrices.</p>

### 30 Matrix Math Functions: Scalar Multiplication

<b>Function</b>	<b>lib_dsp_matrix_muls</b>
<b>Description</b>	<p>Matrix / scalar multiplication: <math>R[i][j] = X[i][j] * A</math>.</p> <p>Each element of the input matrix is multiplied by a scalar value using a 32bit multiply 64-bit accumulate function therefore fixed-point multiplication and q-format adjustment overflow behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>).</p> <p>Example:</p> <pre>int input_matrix_X[8][32]; int input_scalar_A = Q28( 0.333 ); int result_vector_R[8][32]; lib_dsp_matrix_muls( input_matrix_X, scalar_value_A, result_matrix_R, 256, 8, 32, 28 ↪ );</pre>
<b>Type</b>	<pre>void lib_dsp_matrix_muls(const int input_matrix_X[],                     int input_scalar_A,                     int result_matrix_R[],                     int row_count,                     int column_count,                     int q_format)</pre>
<b>Parameters</b>	<p><code>input_matrix_X</code>      Pointer/reference to source data X.</p> <p><code>input_scalar_A</code>      Scalar value to multiply each element by.</p> <p><code>result_matrix_R</code>      Pointer to the resulting 2-dimensional data array.</p> <p><code>row_count</code>      Number of rows in input and output matrices.</p> <p><code>column_count</code>      Number of columns in input and output matrices.</p> <p><code>q_format</code>      Fixed point format (i.e. number of fractional bits).</p>

### 31 Matrix Math Functions: Matrix Addition

<b>Function</b>	<b>lib_dsp_matrix_addm</b>
<b>Description</b>	<p>Matrix / matrix addition: <math>R[i][j] = X[i][j] + Y[i][j]</math>.                      32-bit addition is used to compute the result for each element. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.                      Example:</p> <pre> int input_matrix_X [256]; int input_matrix_Y [256]; int result_matrix_R[256]; lib_dsp_matrix_addm( input_matrix_X, input_matrix_Y, result_matrix_R,                     ↪ 8, 32 );                     </pre>
<b>Type</b>	<pre> void lib_dsp_matrix_addm(const int input_matrix_X[],                     const int input_matrix_Y[],                     int result_matrix_R[],                     int row_count,                     int column_count)                     </pre>
<b>Parameters</b>	<p><b>input_matrix_X</b>                      Pointer to source data array X.</p> <p><b>input_matrix_Y</b>                      Pointer to source data array Y.</p> <p><b>result_matrix_R</b>                      Pointer to the resulting 2-dimensional data array.</p> <p><b>row_count</b>    Number of rows in input and output matrices.</p> <p><b>column_count</b>                      Number of columns in input and output matrices.</p>

## 32 Matrix Math Functions: Matrix Subtraction

<b>Function</b>	<b>lib_dsp_matrix_subm</b>
<b>Description</b>	<p>Matrix / matrix subtraction: <math>R[i][j] = X[i][j] - Y[i][j]</math>.                      32-bit subtraction is used to compute the result for each element. Therefore fixed-point value overflow conditions should be observed. The resulting values are not saturated.</p> <p>Example:</p> <pre>int input_matrix_X [256]; int input_matrix_Y [256]; int result_matrix_R[256]; lib_dsp_matrix_addv( input_matrix_X, input_matrix_Y, result_matrix_R,                     ↪ 8, 32 );</pre>
<b>Type</b>	<pre>void lib_dsp_matrix_subm(const int input_matrix_X[],                    const int input_matrix_Y[],                    int result_matrix_R[],                    int row_count,                    int column_count)</pre>
<b>Parameters</b>	<p><code>input_matrix_X</code>                      Pointer to source data array X.</p> <p><code>input_matrix_Y</code>                      Pointer to source data array Y.</p> <p><code>result_matrix_R</code>                      Pointer to the resulting 2-dimensional data array.</p> <p><code>row_count</code>    Number of rows in input and output matrices.</p> <p><code>column_count</code>                      Number of columns in input and output matrices.</p>

### 33 Matrix Math Functions: Matrix Multiplication

<b>Function</b>	<b>lib_dsp_matrix_mulm</b>
<b>Description</b>	<p>Matrix / matrix multiplication: <math>R[i][j] = X[i][j] * Y[i][j]</math>.                      Elements in each of the input matrices are multiplied together using a 32bit multiply 64-bit accumulate function therefore fixed-point multiplication and q-format adjustment overflow behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>).</p> <p>Example:</p> <pre>int input_matrix_X[8][32]; int input_matrix_Y[8][32]; int result_vector_R[8][32]; lib_dsp_matrix_mulm( input_matrix_X, input_matrix_Y, result_matrix_R, 256, 8, 32, 28 ↪ );</pre>
<b>Type</b>	<pre>void lib_dsp_matrix_mulm(const int input_matrix_X[],                     const int input_matrix_Y[],                     int result_matrix_R[],                     int row_count,                     int column_count,                     int q_format)</pre>
<b>Parameters</b>	<p><code>input_matrix_X</code>                      Pointer to source data array X.</p> <p><code>input_matrix_Y</code>                      Pointer to source data array Y.</p> <p><code>result_matrix_R</code>                      Pointer to the resulting 2-dimensional data array.</p> <p><code>row_count</code>    Number of rows in input and output matrices.</p> <p><code>column_count</code>                      Number of columns in input and output matrices.</p> <p><code>q_format</code>    Fixed point format (i.e. number of fractional bits).</p>

## 34 Statistics Functions: Vector Mean

<b>Function</b>	<b>lib_dsp_vector_mean</b>
<b>Description</b>	<p>Vector mean: <math>R = (X[0] + X[N-1]) / N</math>.</p> <p>This function computes the mean of the values contained within the input vector. Due to successive 32-bit additions being accumulated using 64-bit arithmetic overflow during the summation process is unlikely. The final value, being effectively the result of a left-shift by <code>q_format</code> bits will potentially overflow the final fixed-point value depending on the resulting summed value and the chosen Q-format.</p> <p>Example:</p> <pre>int result; result = lib_dsp_vector_mean( input_vector, 256, 28 );</pre>
<b>Type</b>	<pre>int lib_dsp_vector_mean(const int input_vector_X[],                     int vector_length,                     int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>vector_length</code> Length of the input vector.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

### 35 Statistics Functions: Vector Power (Sum-of-Squares)

<b>Function</b>	<b>lib_dsp_vector_power</b>
<b>Description</b>	<p>Vector power (sum of squares): <math>R = X[0]^2 + X[N-1]^2</math>.</p> <p>This function computes the power (also know as the sum-of-squares) of the values contained within the input vector.</p> <p>Since each element in the vector is squared the behavior for fixed-point multiplication should be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). Due to successive 32-bit additions being accumulated using 64-bit arithmetic overflow during the summation process is unlikely. The final value, being effectively the result of a left-shift by <code>q_format</code> bits will potentially overflow the final fixed-point value depending on the resulting summed value and the chosen Q-format.</p> <p>Example:</p> <pre>int result; result = lib_dsp_vector_power( input_vector, 256, 28 );</pre>
<b>Type</b>	<pre>int lib_dsp_vector_power(const int input_vector_X[],                     int vector_length,                     int q_format)</pre>
<b>Parameters</b>	<pre>input_vector_X     Pointer to source data array X.  vector_length     Length of the input vector.  q_format     Fixed point format (i.e. number of fractional bits).</pre>



## 36 Statistics Functions: Root Mean Square (RMS)

<b>Function</b>	<b>lib_dsp_vector_rms</b>
<b>Description</b>	<p>Vector root mean square: <math>R = ((X[0]^2 + X[N-1]^2) / N) ^ 0.5</math>.</p> <p>This function computes the root-mean-square (RMS) of the values contained within the input vector.</p> <pre> result = 0 for i = 0 to N-1: result += input_vector_X[i] return lib_dsp_math_squareroot( result / vector_length )                     </pre> <p>Since each element in the vector is squared the behavior for fixed-point multiplication should be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). Due to successive 32-bit additions being accumulated using 64-bit arithmetic overflow during the summation process is unlikely. The squareroot of the 'sum-of-squares divided by N uses the function <code>lib_dsp_math_squareroot</code>; see behavior for that function. The final value, being effectively the result of a left-shift by <code>q_format</code> bits will potentially overflow the final fixed-point value depending on the resulting summed value and the chosen Q-format.</p> <p>Example:</p> <pre> int result; result = lib_dsp_vector_rms( input_vector, 256, 28 );                     </pre>
<b>Type</b>	<pre> int lib_dsp_vector_rms(const int input_vector_X[],                   int vector_length,                   int q_format)                     </pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>vector_length</code> Length (N) of the input vector.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

### 37 Statistics Functions: Dot Product

<b>Function</b>	<b>lib_dsp_vector_dotprod</b>
<b>Description</b>	<p>Vector dot product: <math>R = X[0] * Y[0] + X[N-1] * Y[N-1]</math>.</p> <p>This function computes the dot-product of two equal length vectors. The elements in the input vectors are multiplied before being summed therefore fixed-point multiplication behavior must be considered (see behavior for the function <code>lib_dsp_math_multiply</code>). Due to successive 32-bit additions being accumulated using 64-bit arithmetic overflow during the summation process is unlikely. The final value, being effectively the result of a left-shift by <code>q_format</code> bits will potentially overflow the final fixed-point value depending on the resulting summed value and the chosen Q-format.</p> <p>Example:</p> <pre>int result; result = lib_dsp_vector_dotprod( input_vector, 256, 28 );</pre>
<b>Type</b>	<pre>int lib_dsp_vector_dotprod(const int input_vector_X[],                       const int input_vector_Y[],                       int vector_length,                       int q_format)</pre>
<b>Parameters</b>	<p><code>input_vector_X</code> Pointer to source data array X.</p> <p><code>input_vector_Y</code> Pointer to source data array Y.</p> <p><code>vector_length</code> Length of the input vectors.</p> <p><code>q_format</code> Fixed point format (i.e. number of fractional bits).</p>

## APPENDIX A - Known Issues

There are no known issues.

## APPENDIX B - xCORE-200 DSP library change log

### B.1 1.0.0

- Initial version