

Enabling High-Performance DSP Applications with Stratix V Variable-Precision DSP Blocks


The silicon digital signal processing (DSP) architecture of the FPGA can make a big difference when implementing complex signal-processing algorithms. Altera’s Stratix® V FPGAs, with the variable-precision DSP block architecture, are the only programmable devices that efficiently support many different precision levels, including floating-point implementations. Also, with a 64-bit cascade bus and accumulator, the designer does not have to sacrifice precision when the algorithm implementation requires multiple DSP blocks. This unique architecture provides increased system performance, reduced power consumption, and reduced architecture constraints for system algorithm designers.

Introduction

FPGAs are the platform of choice when it comes to implementing high-performance digital signal processing (DSP) systems. As Kevin Morris wrote in a recent article, “...for high-performance algorithmic design, FPGAs are capable of performance, efficiency, and cost-effectiveness orders of magnitude better than alternative solutions like DSP processors...”⁽¹⁾

Every designer implements an algorithm in different ways and with different precision requirements. These requirements vary not only from design to design, but also within each stage of a design, such as finite impulse response (FIR) filters, fast Fourier transforms (FFTs), detection processing, and adaptive algorithms. By mapping the signal-processing precision requirements along a continuum, it is found that different applications fall naturally across the DSP precision spectrum, as shown in Figure 1.

Figure 1. Types of Applications on the DSP Precision Spectrum



FPGAs have traditionally supported 18-bit signal-processing datapaths. However, high-performance signal-processing designs require more than 18-bit precision. Examples of higher precision requirements in traditional DSP functions are FIR filters, FFTs, and also custom signal-processing datapaths. The need for more than 18-bit precision modes sometimes leads to the implementation of floating-point signal processing on LTE channel cards in high-end designs such as advanced military radars and multiple-input multiple-output (MIMO) processing.

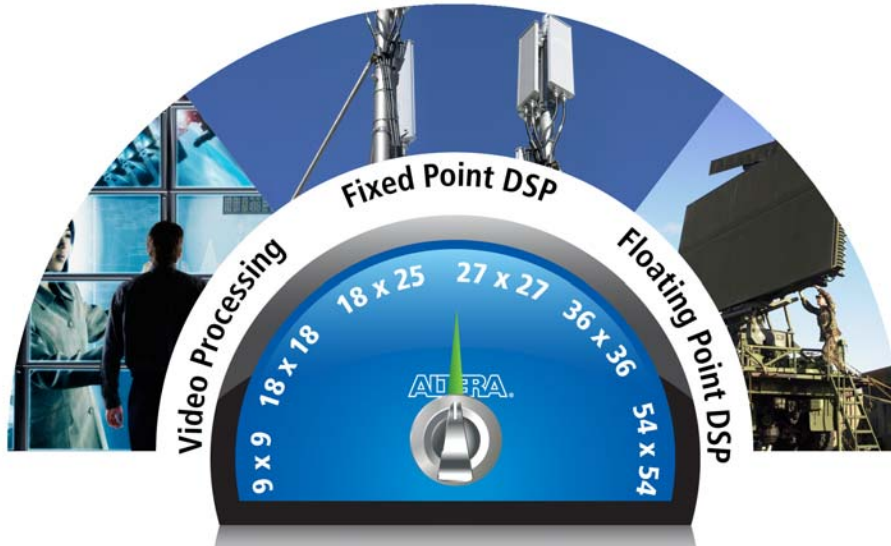
Today's FPGAs use either 18x18 or 18x25 fixed-precision DSP architecture. Other precision modes can be supported, but at the cost of wasting DSP capabilities—for example, using an 18x25 DSP block to implement a 9x9 operation—or by cascading multiple DSP blocks.

As DSP applications move towards higher ranges of datapath precision, the DSP architecture for the next-generation FPGA must not only be flexible enough to accommodate high precision and even floating points, but also able to continue supporting traditional 18-bit precision signal processing in the most area- and power-efficient manner.

Introducing the First Variable-Precision DSP Architecture

With the 28-nm Stratix® V FPGA, Altera has developed a DSP block architecture that uniquely provides native support for higher precision customer needs. These needs are for fixed precision modes that are higher than 18 bits, lower precision modes such as 18x18, 12x12, and 9x9, and double-precision floating-point implementations. All implementations are enabled at high performance and in an area-efficient manner. Conceptually, variable precision can be thought of as a dial, such as the one shown in Figure 2, that designers can use to set the precision of the DSP architecture to better match an application's needs on a block-by-block basis.

Figure 2. Precision Dial for Determining DSP Architecture Precision Mode



At the heart of this new architecture is a variable-precision DSP block that can be configured in either an 18-bit precision mode, as shown in Figure 3, or a high-precision mode, as shown in Figure 4.

Figure 3. 18-bit Precision Mode of Stratix V FPGA

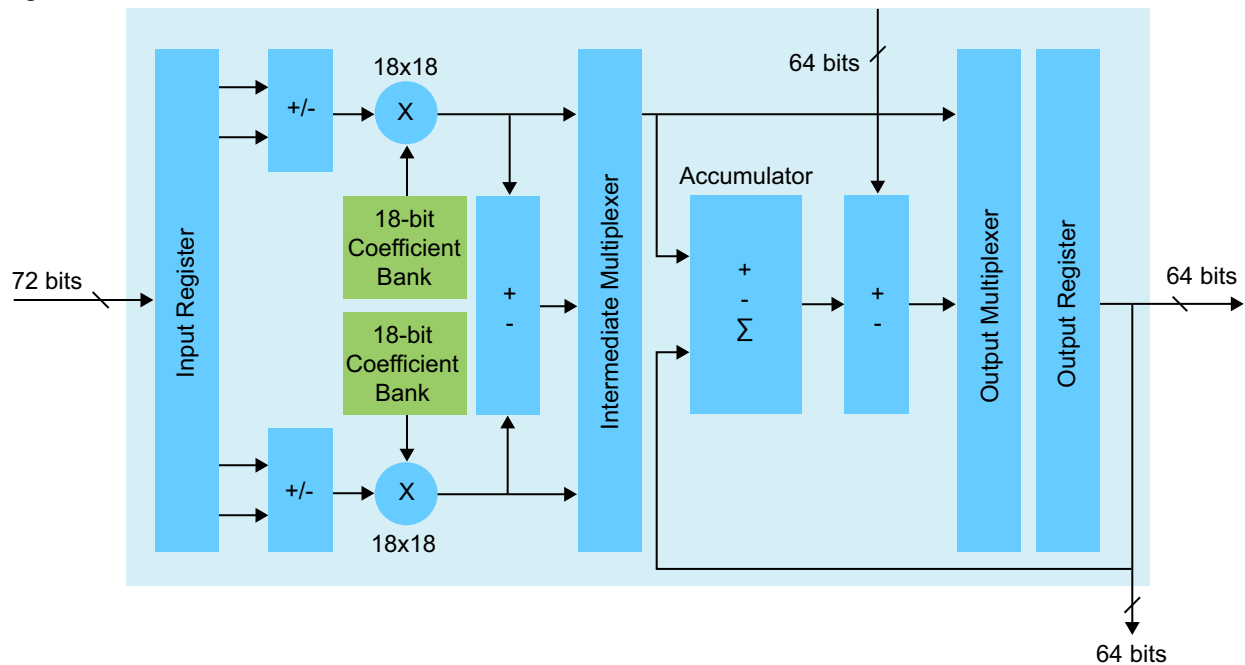
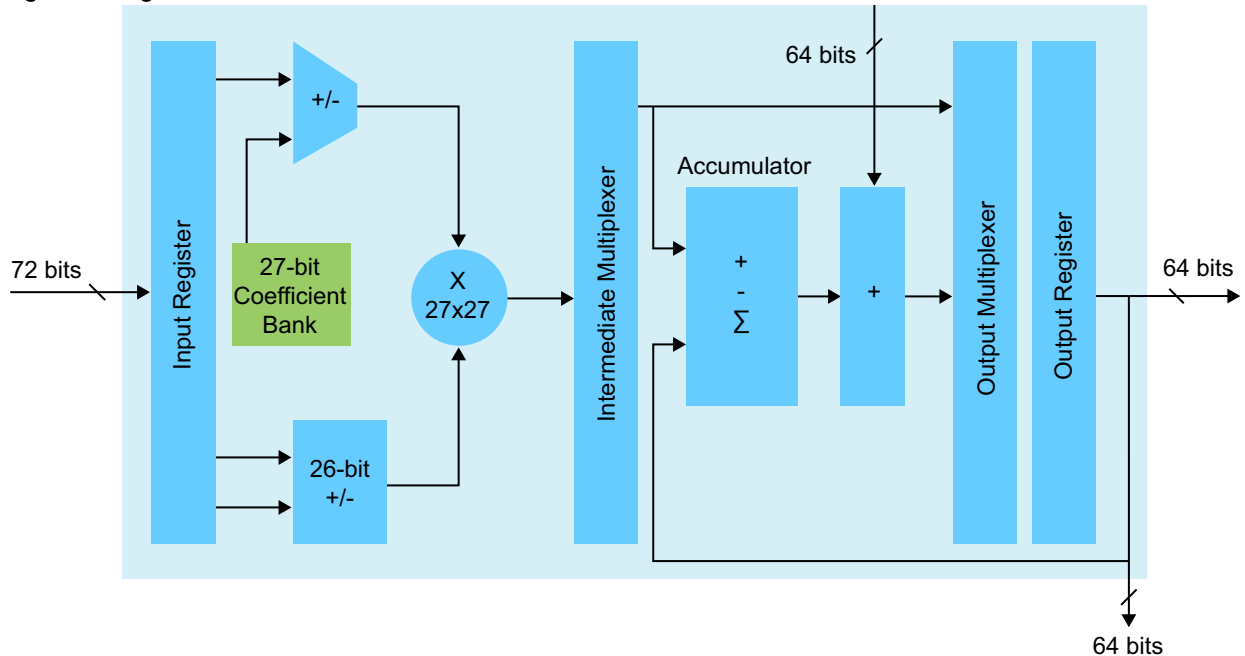


Figure 4. High-Precision Mode of Stratix V FPGA



Each variable-precision DSP block can be configured using a 27x27-bit multiplier. Alternatively, it can be configured using dual 18x18-bit multipliers. In either case, the multiplier stage is preceded by a pre-adder and followed by a 64-bit accumulator and cascaded-adder stage.

The ability to configure a DSP block to either implement a configuration of dual 18-bit multipliers or a single high-precision multiplier is unique. This ability allows the implementation of signal-processing datapaths of different precision modes, either within a single block or by cascading multiple blocks.

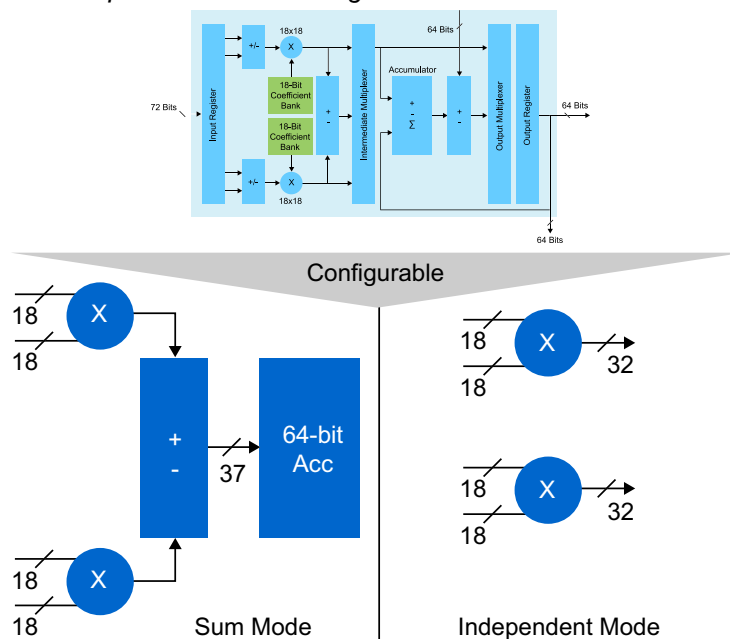
By accommodating varying precision modes without the need for external logic, Stratix V FPGAs can now deliver area-efficient implementations with the highest system f_{MAX} and the lowest system latency. The major new features of the variable-precision DSP block include:

- Native 27x27-bit and 18x18-bit multiplier precision
- New 18x25 complex multiplication mode
- 64-bit accumulator (the largest in the industry)
- Hard pre-adder (available in both 18-bit and 27-bit modes)
- Cascaded output adders for efficient systolic FIR filters
- Internal coefficient register banks
- Enhanced independent multiplier operation
- Efficient support for single- and double-precision floating-point arithmetic
- Inferrability of all modes by the Quartus® II synthesis engine

Variable-Precision DSP Block Configurations

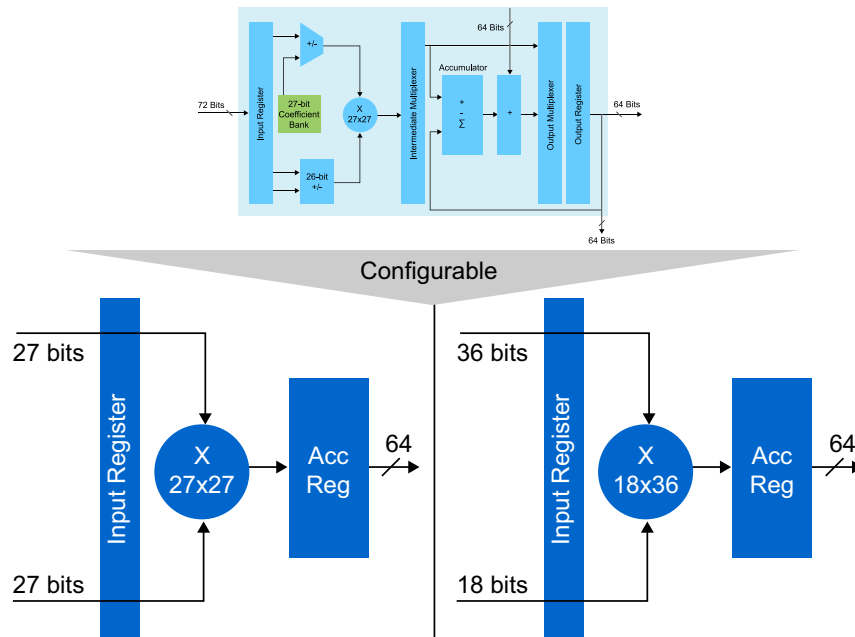
In the 18-bit precision mode, the block can be configured to support dual multipliers in either the sum mode or the independent mode, as shown in Figure 5. The sum mode, which is commonly used to implement FIR filters, is the same mode that is available with the previous generation of Stratix FPGAs, thus backward compatibility with existing customer designs can be maintained. Alternatively, the multipliers can be configured in the independent mode, in which the output of the multipliers need not be summed. In this mode, the output resolution is limited to 32 bits. In both modes, the coefficients can be stored in the internal coefficient store or routed from outside the block, and the pre-adder can either be used or bypassed.

Figure 5. Sum Mode and Independent Mode Configurations of Stratix V FPGAs



In the high-precision mode, a designer can configure each variable-precision DSP block to implement either a 27x27 multiplier or an 18x36 multiplier, as shown in Figure 6. The bit width of the internal coefficient store and the size of the hard pre-adder also increase to accommodate the higher precision.

Figure 6. Configurations for the High-Precision Mode of Stratix V FPGAs



With the increased precision mode, the FPGA's variable-precision block is efficient for implementing single-precision floating-point operations. A single-precision mantissa multiplication can be implemented using only one variable-precision block configured in the high-precision mode. The designer can choose to use the internal coefficient storage or bypass it.

A single variable-precision DSP block can implement different multiplication operations for each precision mode, as shown in Table 1. Each block can be configured differently within the same design.

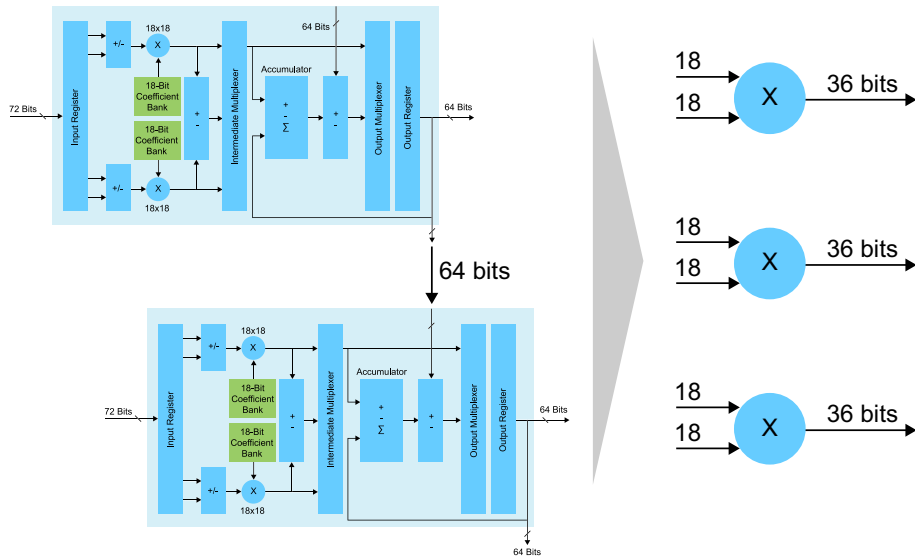
Table 1. Precision Modes Supported by Each Block of Stratix V FPGAs

Precision Mode	Application
3 independent mode 9x9	Low-precision fixed point
2 sum mode 18x18	Medium-precision fixed point
2 independent mode 18x18 with 32-bit resolution	Medium-precision fixed point
1 independent mode 18x25 or 18x36	High-precision fixed point
1 independent mode 27x27	High-precision fixed point or single-precision floating point

Cascading the Variable-Precision DSP Blocks

Each variable-precision DSP block can be cascaded with other blocks to further extend the range of precision modes that it can support. Cascading is performed with a dedicated 64-bit bus, in which precision modes higher than 27x27, complex multipliers, and operations that require more than one variable-precision DSP block can be built without resorting to external logic. In addition, using two variable-precision blocks can implement three independent 18x18 multipliers with a full 36-bit resolution as shown in Figure 7.

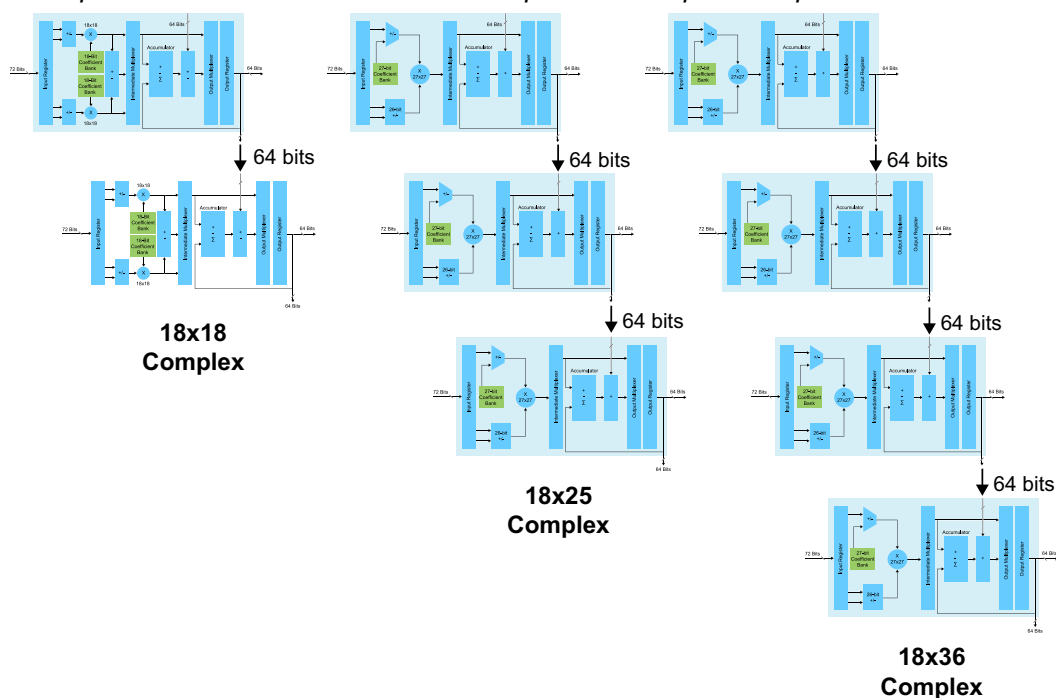
Figure 7. Cascading Two Variable-Precision DSP Blocks of Stratix V FPGA



The FFT algorithm, a principal algorithm in radar, medical imaging, and wireless systems, is commonly implemented using complex multiplication. The variable-precision DSP block optimally supports the FFT algorithm, which increases precision requirements on only one side of the multiplier. The variable-precision DSP block supports the asymmetric multiplier size in several precision modes to accommodate data growth through the FFT stages, resulting in higher dynamic range and a lower noise floor.

Figure 8 shows how complex multipliers of various precision modes can be implemented efficiently by using multiple variable-precision DSP blocks. To implement an 18x18 complex multiplier, two variable-precision DSP blocks configured in the 18-bit precision mode can be cascaded together. An 18x25 complex multiplier would traditionally require four multipliers in the high-precision mode. But with the 27x27 multiplication capability and 26-bit pre-adders, the same complex multiplier can be implemented using three variable-precision DSP blocks. Competing architectures that are limited to the 18x25 precision need four blocks to implement this function, and to implement even higher precision complex multipliers such as 18x36, four variable-precision DSP blocks can be cascaded.

Figure 8. Multiple Variable-Precision DSP Blocks Implement Complex Multipliers in Stratix V FPGAs



For the highest precision applications, floating-point FFTs using a 27×27 complex multiplier can be supported by cascading four variable-precision DSP blocks configured in the high-precision mode. Additionally, for FFT applications with high dynamic-range requirements, only the Altera® FFT MegaCore® function offers an option of single-precision floating-point implementation with resource usage and performance similar to high-precision fixed-point implementations. Table 2 lists the different precision modes that can be supported by cascading these blocks.

Table 2. List of Precision Modes Supported by Cascading Multiple Variable-Precision Blocks of Stratix V FPGAs

Multiplier Mode	DSP Silicon Resources	Application
1 independent mode 36×36	2 variable-precision blocks	Very-high-precision fixed point
1 independent mode 54×54	2 variable-precision blocks	Double-precision floating point
Complex multiply 18×18	2 variable-precision blocks	Resource-optimized FFTs
Complex multiply 18×25	3 variable-precision blocks	Accommodate bit growth in FFTs
Complex multiply 18×36	4 variable-precision blocks	Full-precision large FFT stages
Complex multiply 27×27	4 variable-precision blocks	Single-precision floating-point FFTs

Variable-Precision DSP Blocks—Other Features

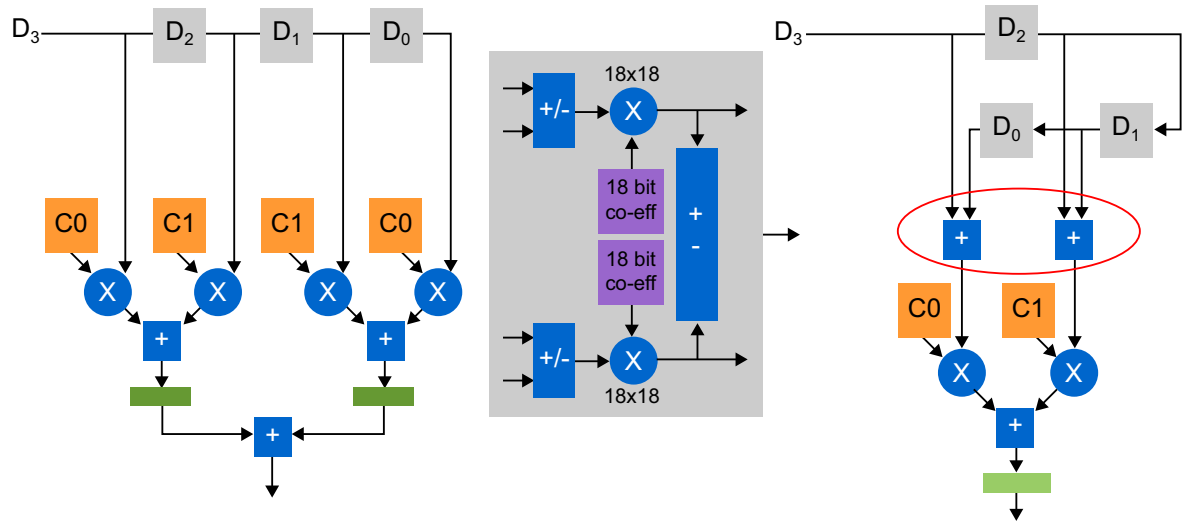
Altera's variable-precision DSP blocks have been carefully designed with features that enable an area- and power-efficient implementation for various commonly implemented DSP functions, including hard pre-adders (or subtractors), accumulators, post-adders (or subtractors) and distributed adders, coefficient storage, and $A \times B + C$ mode.

Hard Pre-Adder (or Subtractor)

Pre-adders are primarily used for symmetric FIR filters. As the data is shifted across the coefficient set, two data samples are multiplied by a common coefficient due to the symmetrical structure. The pre-adder adds the two

samples prior to multiplication, requiring the use of only one multiplier for every two taps rather than two multipliers. Figure 9 shows the architecture of the hard pre-adder or subtractor in the variable-precision DSP block.

Figure 9. Hard Pre-Adder or Subtractor Architecture in the Variable-Precision DSP Block of Stratix V FPGA



Which precision mode is selected determines whether a dual 18-bit or a single 26-bit pre-adder is used. The same pre-adder can also be configured as a subtractor, so the DSP block can be used to perform the sum of squared differences, yet another DSP algorithmic function.

Accumulator

Accumulators are integral to many DSP operations. They are necessary when a single multiplier is used to calculate a series of multiplication and addition operations.

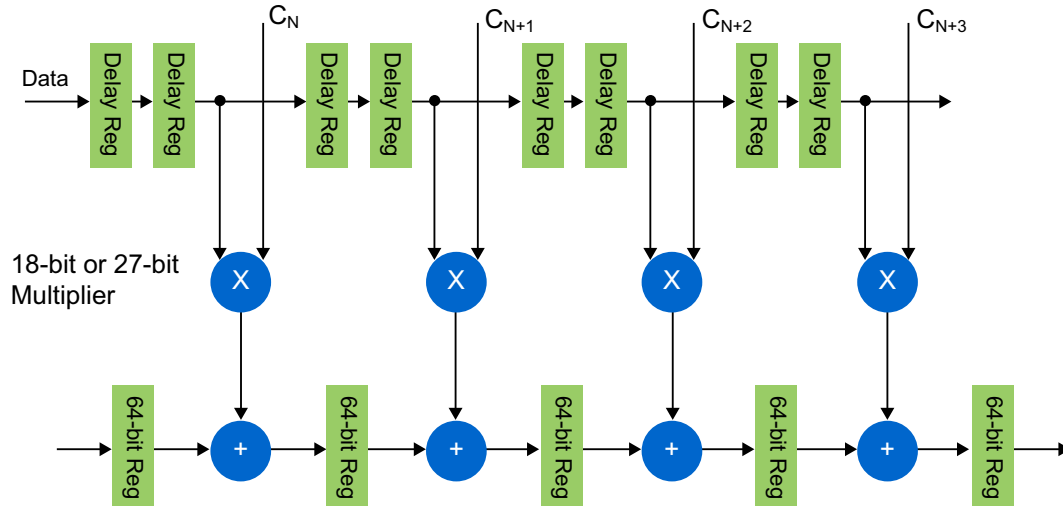
Each variable-precision DSP block has a 64-bit accumulator, which is the largest in the industry. This feature facilitates high-precision DSP computations. The accumulator supports incrementing, decrementing, and pre-loading. Additionally, the ability to perform biased rounding on the final accumulation result allows for bit truncations within the DSP block as needed for further processing stages.

Post-Adder (or Subtractor) and Distributed Adder

Post-adders are used to construct larger adders from smaller adders to perform summation or subtraction operations in complex multiplications. Post-adders are also used to perform sum-of-products operations used in FIR filters.

A popular FIR-filter architecture, also known as systolic architecture, is shown in Figure 10. It uses a distributed output adder, and therefore, programmable logic-based adder circuits are not required regardless of the FIR-filter size. This architecture allows for more efficient and higher f_{MAX} (clock frequency) FIR-filter implementations, and supports both 18-bit and 27-bit FIR filter structures. A 64-bit vertical cascade path is used between variable-precision DSP blocks, which are normally placed in columns in an FPGA.

Figure 10. Distributed Adder or Systolic FIR Architecture




Coefficient Storage

Coefficients are required for most DSP operations such as FIR filters and FFTs. These coefficients are normally stored in distributed memory blocks external to the DSP block. This allows for large numbers of coefficients to be used and for easy updating of adaptive filters, for example. However, as most FIR filters in a hardware implementation are built using a parallel, or partially parallel, structure, the number of coefficients used per multiplier is often small. For these cases, it is advantageous to provide a coefficient storage bank that is dynamically selectable on each clock cycle inside the DSP block.

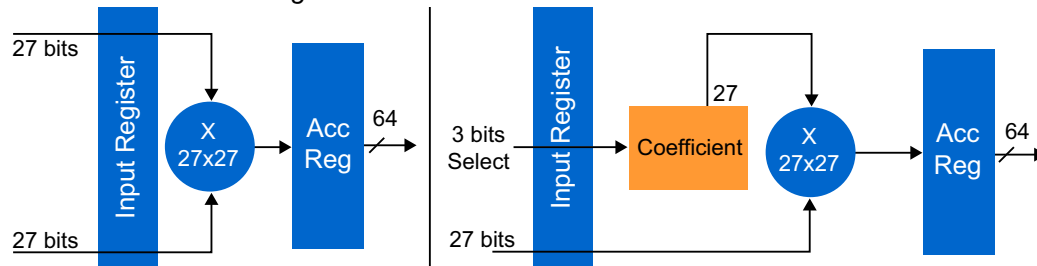
As shown in Figure 11, the variable-precision DSP block integrates an 8-deep, 27-bit wide internal coefficient bank, or two 8-deep, 18-bit wide internal coefficient banks, depending on mode. Besides saving FPGA memory resources, the advantages are reduced power and routing congestions, and easier timing closure of high-performance designs.

Figure 11. Architecture of Variable-Precision DSP Blocks



The internal coefficient storage can be used or bypassed, as shown in Figure 12.


Figure 12. Internal Coefficient Storage in Variable-Precision DSP Blocks of Stratix V FPGAs



$A \times B + C$ Mode

The variable-precision DSP block can support multiplication operations followed by the summation of the product with external inputs, as shown in Figure 13. The final result can be output or accumulated with previous operation results.

Figure 13. Implementing $A \times B + C$ Mode in Stratix V FPGAs



Conclusion

While many data converters commonly fall within the 14-bit to 16-bit range, sampling rate decimation, MIMO techniques, frequency domain transforms, and other processing-gain techniques often increase datapath precision requirements. Many newer DSP applications are becoming more demanding in terms of performance as well as precision. In addition, it is not uncommon to have a range of precision modes that must be supported across a datapath in many of tomorrow's DSP applications.

Stratix V FPGAs, with the variable-precision DSP block architecture, is the only programmable device family that can efficiently support many different precision levels, including floating-point implementations. A variable-precision DSP architecture allows the designer to set the precision mode to match the algorithm instead of the other way around. Also, with a 64-bit cascade bus and accumulator, the designer does not have to sacrifice precision when the algorithm implementation requires multiple DSP blocks.

With the introduction of this DSP architecture, Altera has opened a silicon DSP technology gap against the competition. Today designers can natively, and within a single block, implement a 27×27 multiplier, which is not only useful for high-precision fixed-point DSP applications, but is also very efficient for emerging floating-point DSP applications. This unique architecture provides increased system performance, reduced power consumption, and reduced architecture constraints for system algorithm designers.

Further Information

1. Kevin Morris, "A Perfect DSP Storm: BDTi + High Level Synthesis + FPGA," *FPGA and Programmable Logic Journal*, January 19, 2010:
www.fpgajournal.com/fpgajournal/feature_articles/20100119-storm
2. Stratix V FPGAs: Built for Bandwidth:
www.altera.com/products/devices/stratix-fpgas/stratix-v/stxv-index.jsp
3. Literature: Stratix V Devices:
www.altera.com/products/devices/stratix-fpgas/stratix-v/literature/stv-literature.jsp
4. 28-nm Variable-Precision DSP Block Architecture:
www.altera.com/technology/dsp/variable-precision/dsp-variable-precision.html
5. Webcast: "Enabling High-Performance DSP with Stratix V Variable-Precision DSP Architecture":
www.altera.com/education/webcasts/all/wc-2010-dsp-var-prec-dsp-arch.html

Acknowledgements

- Michael Parker, Senior Technical Marketing Manager, IP and Technology Product Marketing, Altera Corporation
- Suhel Dhanani, Senior Manager, Software, Embedded, and DSP Marketing, Altera Corporation

Document Revision History

Table 1. Document Revision History

Date	Version	Changes
May 2011	1.1	Added details for Stratix V variable-precision DSP blocks.
May 2010	1.0	Initial release.



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2011 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.