Implementing FIR Filters and FFTs with 28-nm Variable-Precision DSP Architecture

WP-01140-1.0 White Paper

Across a range of applications, the two most common functions implemented in FPGA-based high-performance signal processing are finite impulse response (FIR) filters and fast Fourier transforms (FFTs). The FPGA's digital signal processing (DSP) architecture must be optimized to allow the most efficient implementation of these structures as this directly translates into cost and power benefits to the customer. This white paper introduces the DSP architecture of the latest 28-nm Altera® FPGAs and shows how this architecture enables the most efficient implementation of FIR filters and FFTs.

Introduction

FIR filters and FFTs are two of the most common DSP functions implemented in FPGAS. As shown in Figure 1, Altera's internal primary market research shows that these functions dominate implementation when it comes to DSP in FPGAs. This is most likely due to the fact that FPGAs alone are able to meet the throughput and latency requirements for these functions, given the parallel DSP datapath implementation that is feasible with the FPGA architecture.

Figure 1. Types of DSP Functions Implemented in FPGAs

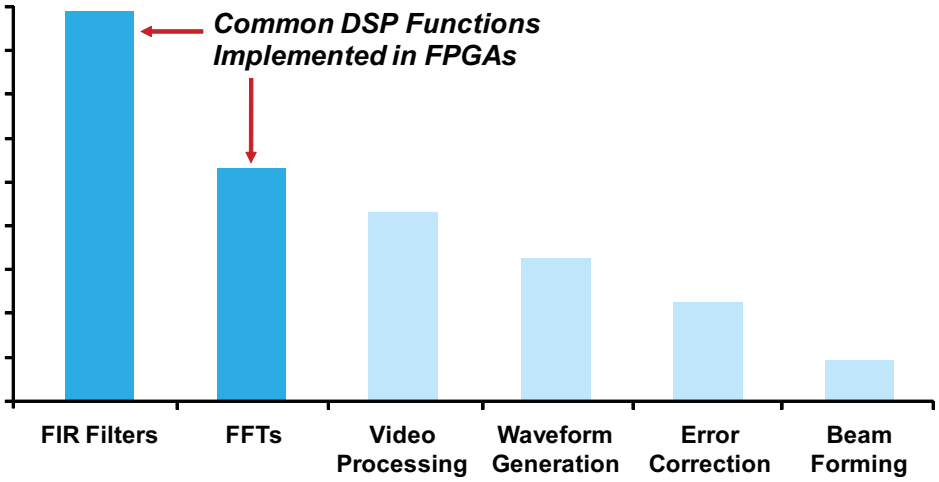

Given the preponderance of FIR and FFT implementation, it is critical that the FPGA DSP architecture be designed to enable this implementation with highest performance and the least resources. At the 28-nm process node, Altera has developed the FPGA industry's first variable-precision DSP architecture in its Stratix® V devices. This architecture enables designs with varying precision and performance requirements to be implemented using the 28-nm silicon fabric with two to three times the implementation efficiency compared to competing, fixed-precision 18x25 DSP architectures.

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

This innovative 28-nm DSP architecture is also designed with key features that enable the efficient implementation of high-performance FIR filters and FFT structures. This white paper details the specific key features that are included for FIR and FFT implementation optimization.

FIR Filter Optimization Features

As shown in Figure 2, Altera's 28-nm DSP architecture includes a host of features for optimizing FIR filter implementations:

- Hard, built-in pre-adders can be used when implementing symmetric filters to cut multiplier usage by half.
- Internal co-efficient register storage allows the designer to store the filter coefficients inside the DSP block, which not only saves registers and memory but allows for faster f_{MAX} because coefficients do not have to be routed from the logic.
- Two levels of adders within a single block, which are important when building "direct-form" FIR filters. Competing DSP blocks have only a single level of post-multiply adder stage, which necessitates external logic-based adders to build an adder tree.
- Output register and cascade path for implementing systolic filters.

Figure 2. DSP Architecture Features for FIR Filter Implementation

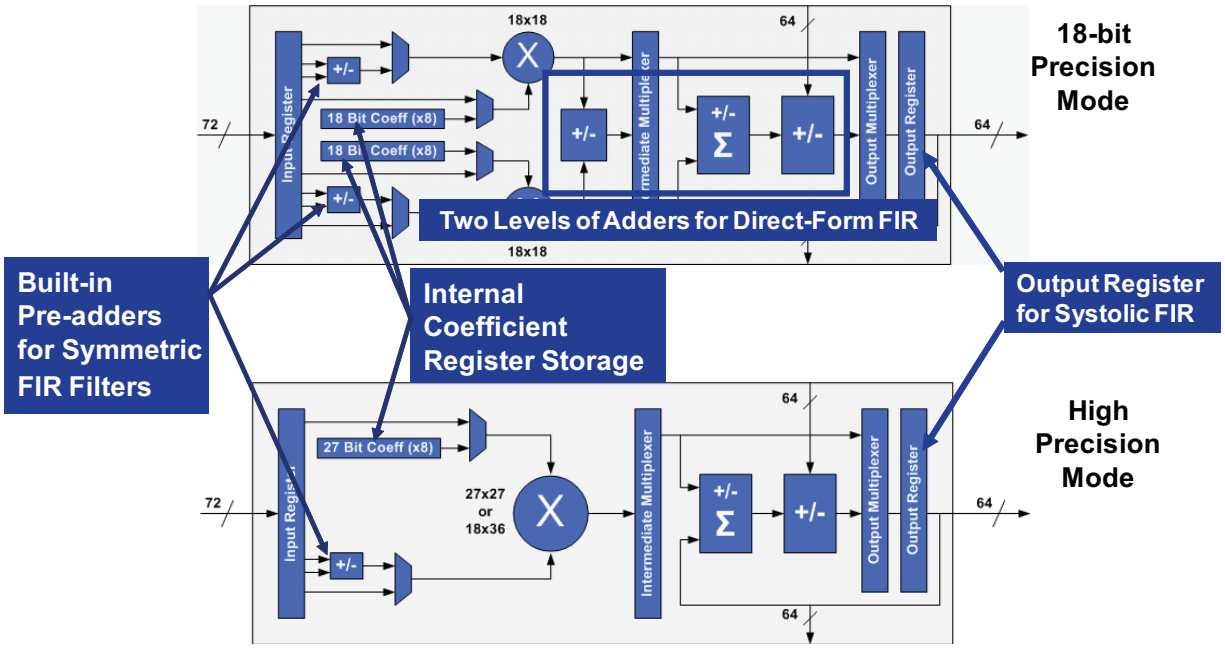

FIR Filter Implementation

The most common forms of FIR filters include the direct-form FIR, the systolic FIR, and the serial FIR.

Figure 3 shows the signal flow graph for a direct-form FIR filter, which requires multipliers, delay elements (registers), and adders for its implementation. The direct-form FIR filter structure maps easily to the variable-precision DSP block, which can implement the delay element, the multiplier and two levels of the adder tree within a single DSP block. For a single-channel, single-rate filter, the input delay element can be built using an internal cascade register.

Figure 3. Direct-Form FIR Filter Implemented Using Altera's 28-nm DSP Block

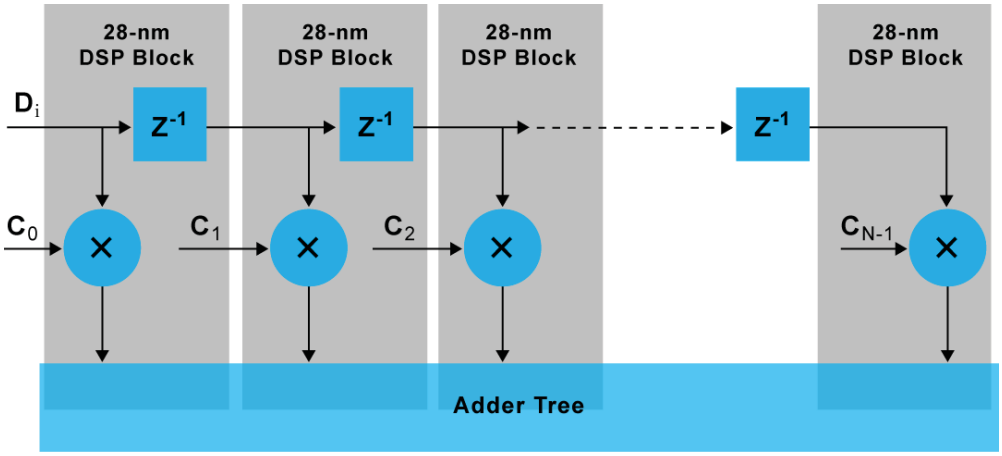

By having two adder stages within a single block (in the 18x18 mode), four multipliers can be added using two DSP blocks and no external logic. This ability to implement two stages of an adder tree within one DSP block is unique to Altera's 28-nm DSP architecture.

Figure 4 shows how multiple variable-precision DSP blocks can be cascaded to create a filter structure. Two DSP blocks create a sum of four multiplicands, the products of which can be summed using an external logic adder tree, if required.

Figure 4. Single-Channel, Single-Rate Direct-Form FIR Filter Implemented Using Altera's 28-nm DSP Blocks

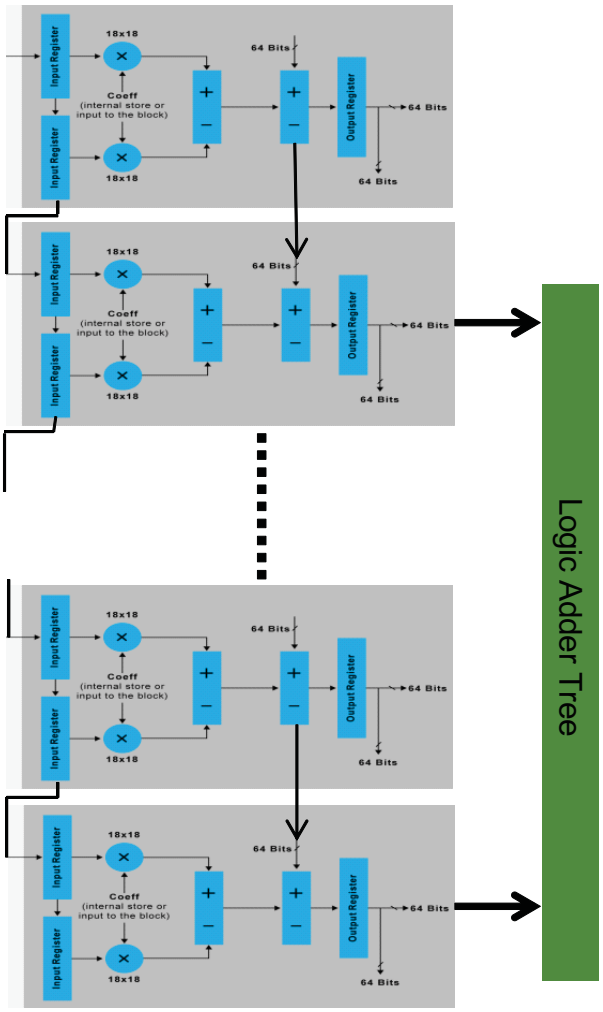

For multi-channel, multi-rate filters, distributed memory can be used to implement input delay elements, rather than input cascaded registers. Symmetric FIR filters (Figure 5) are used widely in various applications because they cut down the number of multipliers by half. By pre-adding the two data samples and then multiplying with the common coefficient, a symmetric FIR filter can replace two multiply operators with one multiplier and a one adder.

Figure 5. Implementing a Symmetric Direct-Form FIR Filter

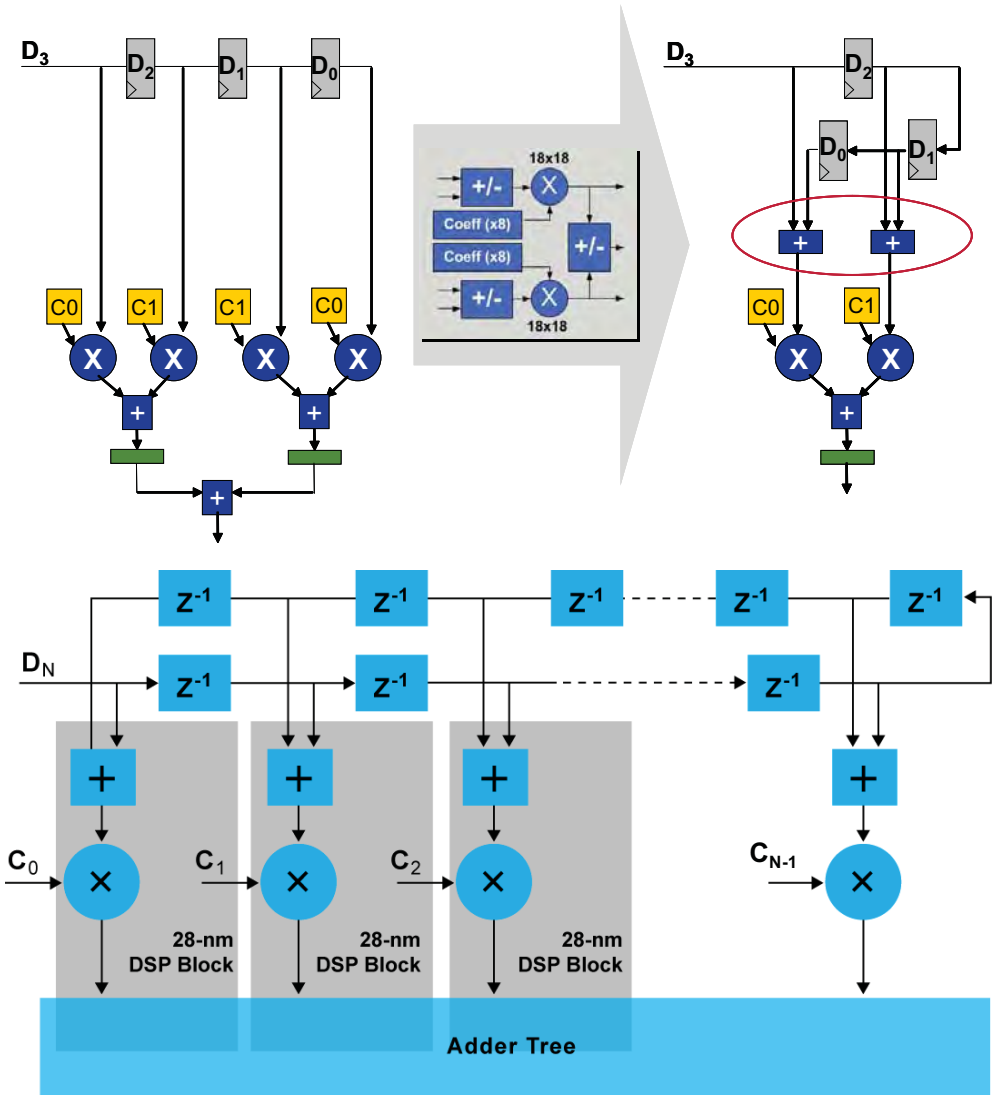

With the variable-precision DSP architecture, the pre-adder is included as a hard feature within the block. In this case, the pre-adder, the multiplier, and part of the adder tree are implemented within the DSP block, while the delay elements are implemented using registers or distributed memory outside the DSP block. Using the hard pre-adder saves half of the multipliers for a same-tap FIR filter.

Another common FIR implementation is the systolic FIR, where the output adder tree is distributed and needs a register delay between each multiply-add. With a systolic FIR, the direct-form FIR's adder tree is replaced by a distributed adder, as shown in Figure 6.

Figure 6. Direct-Form FIR (top) vs. Systolic FIR (bottom)

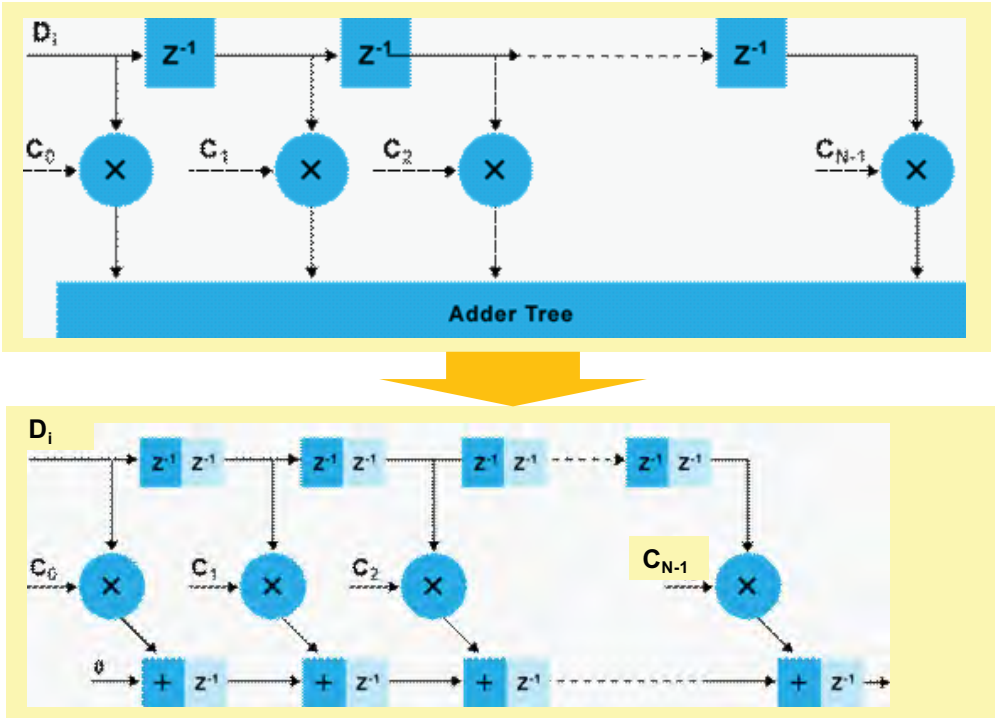

To implement a structure like this, the designer must add in an output systolic register after each adder stage. The variable-precision DSP architecture has such an output register built into the DSP block.

Figure 7 illustrates the configuration of the variable-precision DSP block necessary to implement two stages of the systolic filter per DSP block:

- In the 18-bit systolic mode, the adders are configured as dual 44-bit registered adders, thereby giving 8 bits of overhead when using an 18-bit operation (i.e., 36-bit products) and allowing for a sum of up to 256 multiplier products.
- In the high-precision systolic mode, the designer can configure the single adder with 64 bits. This mode allows the implementation of one stage of the systolic filter per DSP block. A 64-bit adder provides 10 bits of overhead when using 27-bit data with 54-bit products, allowing for a sum of up to 1024 multiplier products.

Figure 7. Registered Adders in Altera's 28-nm DSP Blocks Enable Implementation of Systolic Filters

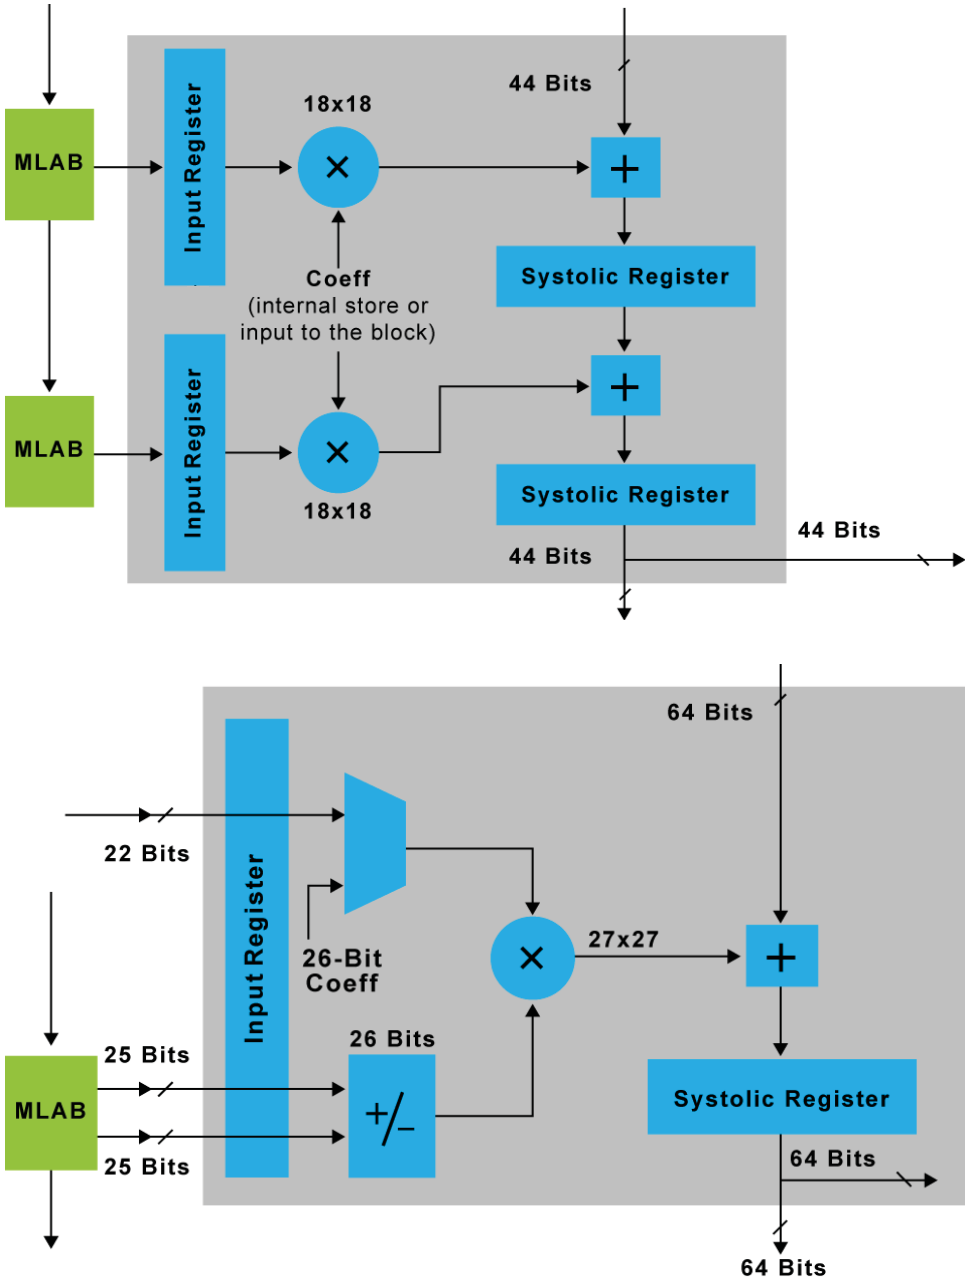

The competing 18x25 fixed-precision architecture where the adder resolution is limited to 48 bits provides only a 5-bit overhead, limiting the designer to only 32 multiplier products without losing precision. This limits the size of a filter that can be implemented without compromising on the overall system precision.

The third type of FIR filter implementation is the serial filter generally implemented in DSP processors with a single multiplier-accumulator (MAC) unit. Here a single multiplier is used in conjunction with an accumulator, as shown in Figure 8. (The pre-adder is optional.) With a 64-bit accumulator and a feedback path from the output register to the accumulator, the designer can implement a serial filter completely within the variable-precision DSP block.

Figure 8. Serial Filters Implemented Using the Accumulator Stage in the Variable-Precision DSP Block

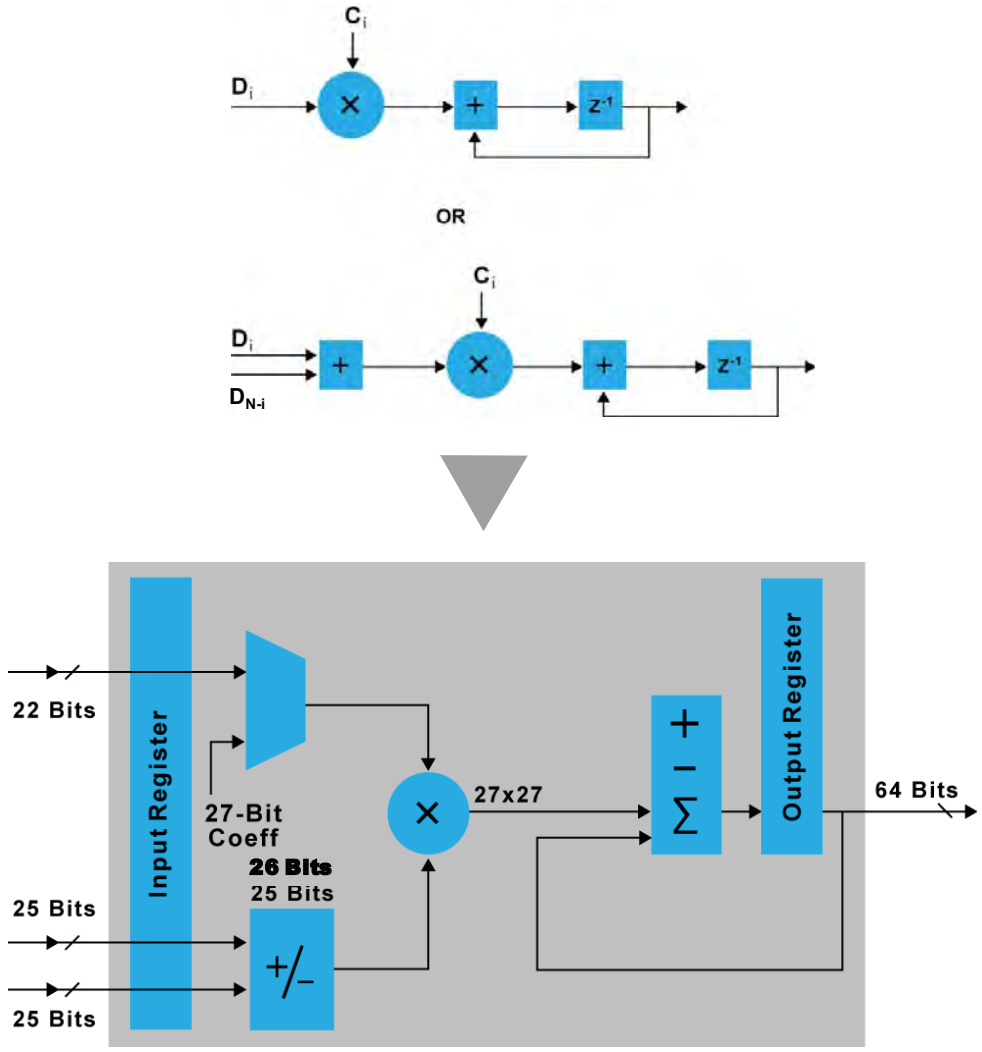

FFT Optimization Features

FFT is a principal algorithm used in military radar, medical imaging, and wireless systems. As shown in Figure 9, Altera's 28-nm DSP architecture includes a host of features for optimizing FIR implementations, including:

- 26-bit pre-adder that allows for optimized implementation of 18x25 complex multiply operation
- Dual 18x18 multipliers that allow the designer to build a real or imaginary product using a single block
- 8x36 mode for supporting even higher data precision
- 64-bit cascade bus that allows blocks to be cascaded without a loss in precision

Figure 9. Features in the Variable-Precision DSP Block Optimize the Implementation of FFT Structures

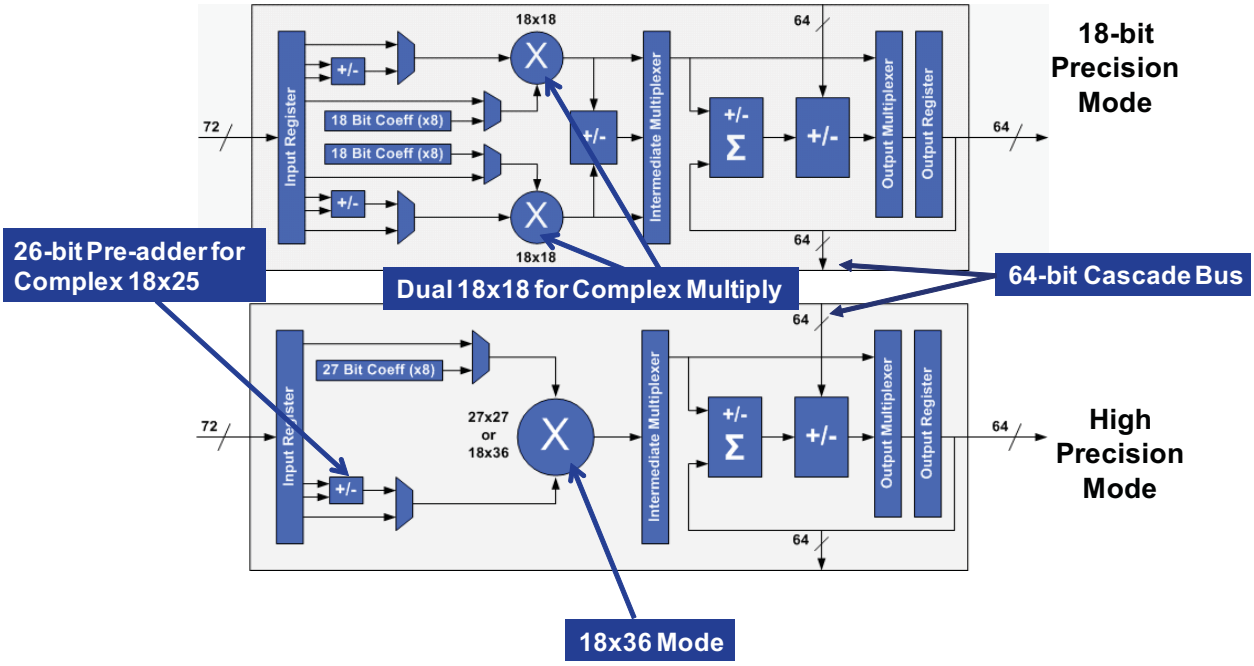

The FFT algorithm has a characteristic of increasing precision requirements on only one side of the multiplier. Data precision growth through the FFT stages is required for higher dynamic range and lower noise floor characteristics, both of which are crucial in advanced systems such as radars and sensitive medical diagnostic systems.

Figure 10 shows how the implementation of data precision growth, needs asymmetric multipliers with increasing precision. The Stratix V FPGA's variable-precision DSP block is the only FPGA DSP architecture that optimally supports the FFT algorithm by allowing the efficient implementation of asymmetric complex multipliers of increasing precision without the need for external logic.

Figure 10. Complex Multipliers with Increasing Data Width at Each Stage, Coefficient Width Stays the Same

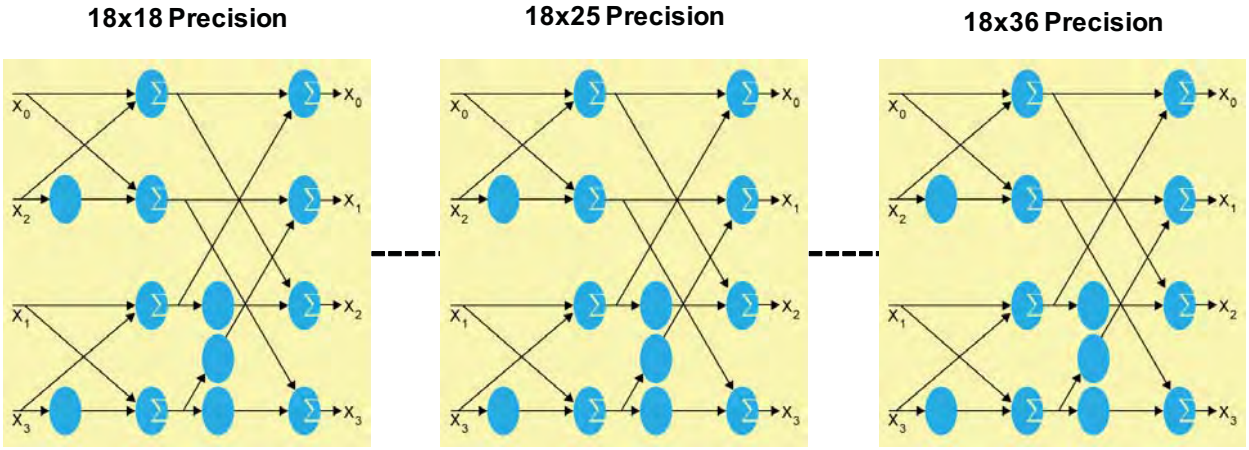

Using the dual 18x18 multipliers in a single DSP block as well as the 64-bit wide cascade chain, allows a designer to implement an 18x18 bit complex multiply using only two DSP blocks, as shown in Figure 11. Competing 18x25, fixed-precision DSP blocks would require twice the number of DSP blocks (i.e., four) when implementing the same algorithm.

Figure 11. Two Variable-Precision Blocks Used to Implement an 18x18 Complex Multiply Operation

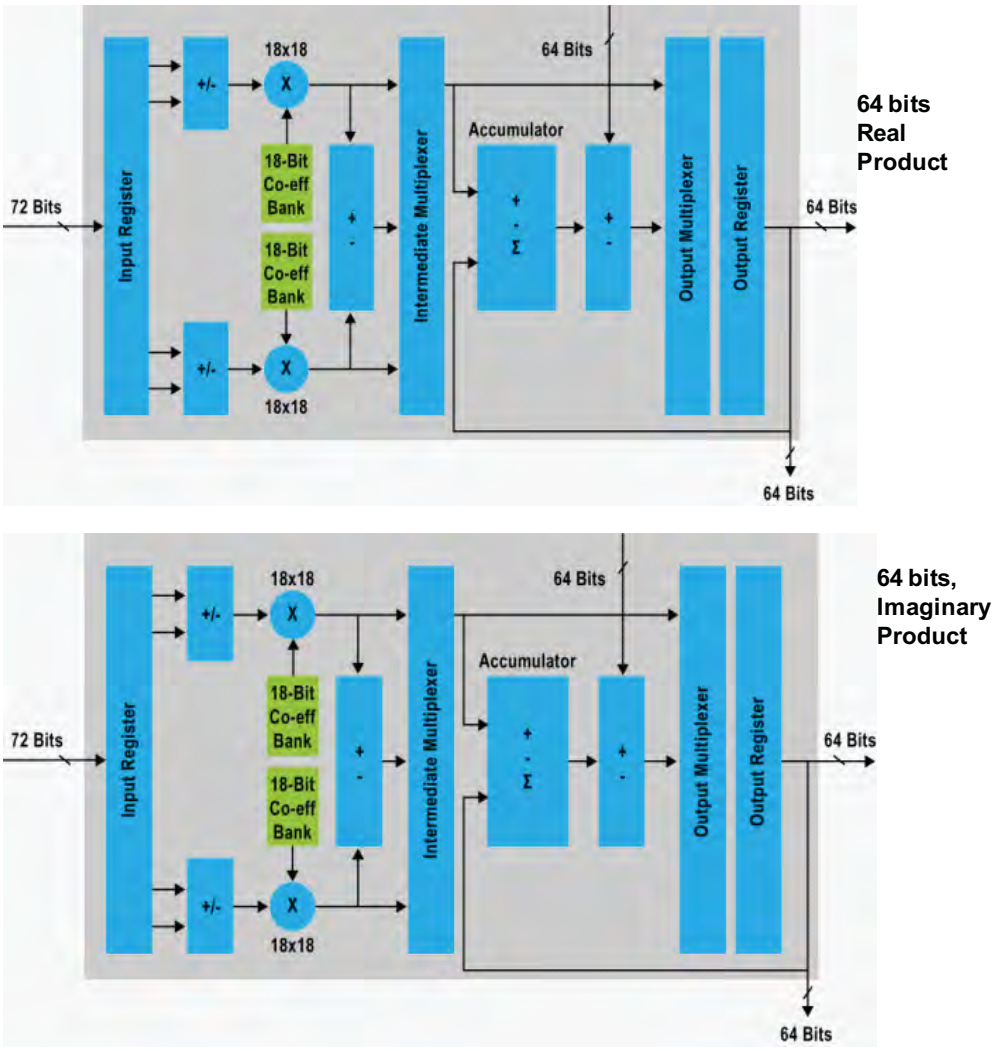

Another popular complex multiply operation is the 18x25 complex multiply. This algorithm can take four competing 18x25, fixed-precision DSP blocks when implemented as shown in Figure 12.

Figure 12. 18x25 Complex Multiply Implementation with Four Competing, Fixed-Precision 18x25 DSP Blocks

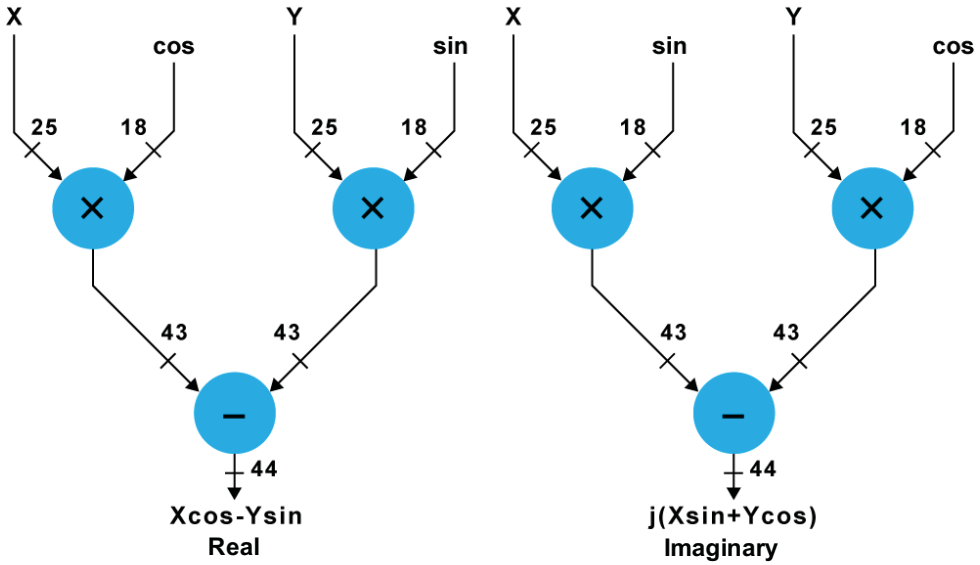

However by using the 26-bit pre-adder and a 27-bit x 27-bit multiplier available exclusively in each 28-nm Stratix V variable-precision DSP block, the same algorithm can be structured as shown in Figure 13.

Figure 13. 18x25 Complex Multiply Implementation with Only Three Variable-Precision DSP Blocks

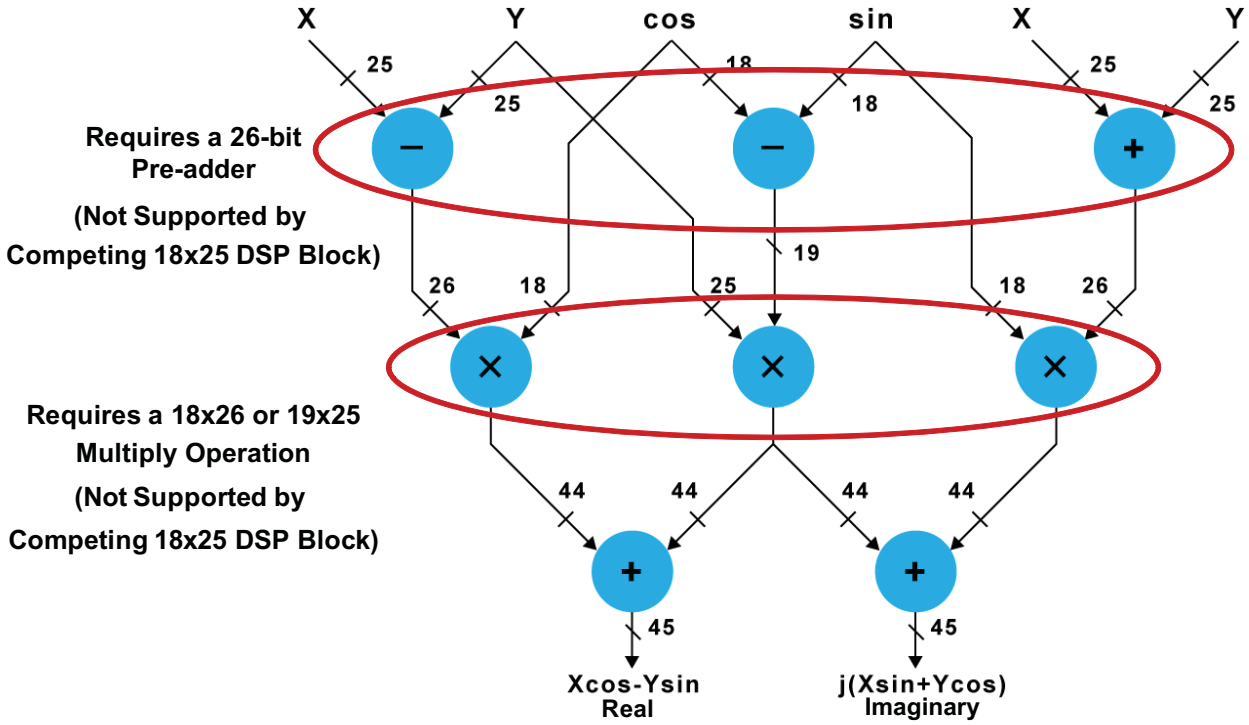

Structured this way, a 18x25 complex multiply operation takes three variable precision block compared to four competing fixed-precision 18x25 DSP blocks. This structure offers a substantial savings in both power and cost when the design requires hundreds if not thousands of DSP blocks.

When dealing with any data precision greater than 25 bits, the variable-precision block can be configured to implement a multiplier up to 18x36 wide. This configuration allows the implementation of higher precision complex multipliers such as 18x36 using only four variable-precision DSP blocks, as shown in Figure 14. Complex multiply operations with increasing bit growth can be accommodated by cascading the DSP blocks without having to resort to logic.

Figure 14. Four Variable-Precision DSP Blocks Cascaded to Implement an 18x36 Complex Multiplier

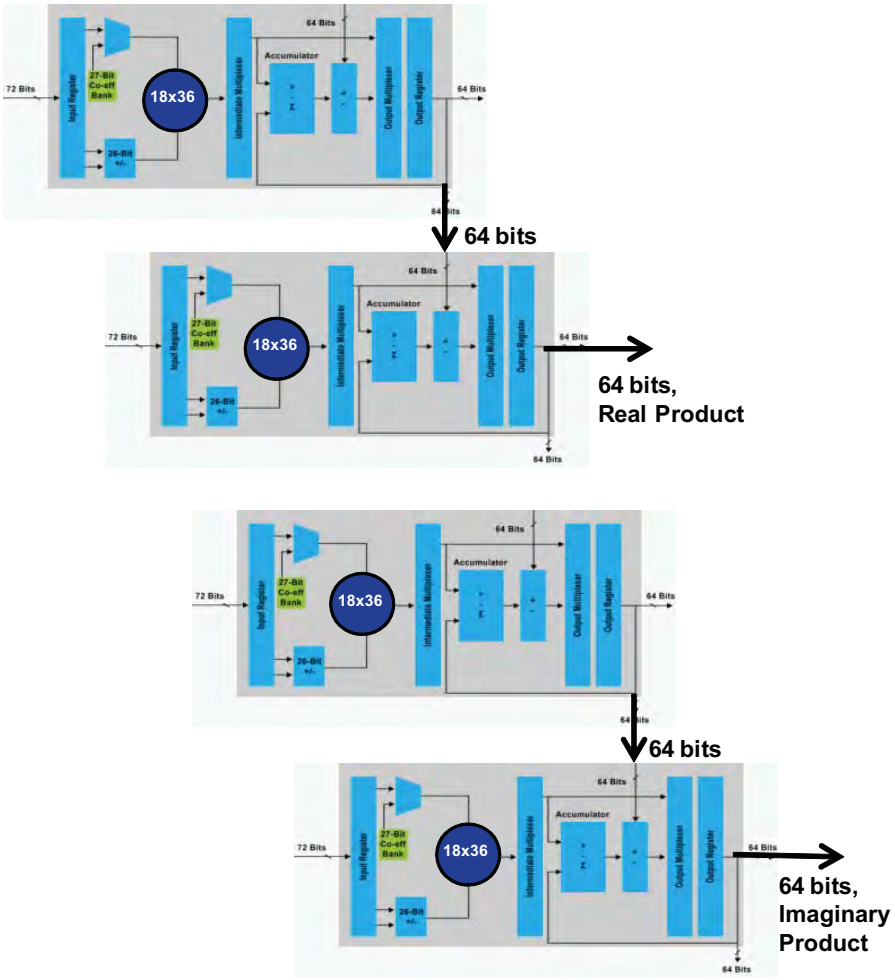

It would take eight competing DSP blocks to implement the same 18x36 bit complex multiply functionality. In short, the variable-precision DSP block includes features that allow for up to twice the silicon efficiency when implementing FFT algorithms.

Summary

The 28-nm variable-precision DSP architecture is carefully designed to provide the best implementation efficiency for FIR and FFT structures, two of the most commonly used DSP structures across a range of systems. Using the unique features of Altera's 28-nm DSP architectures, the system designer can implement these structures in many cases with half the DSP resources compared with the competing 28-nm DSP architecture. This implementation efficiency advantage translates directly to a power and cost advantage for the overall system.

Further Information

- Altera's Total 28-nm DSP Portfolio: Fastest Path to Highest Performance Signal Processing:
  www.altera.com/b/28-nm-dsp-portfolio.html
- Webcast: "Enabling High-Precision, High-Performance DSP With Variable-Precision DSP Architecture":
  www.altera.com/education/webcasts/all/wc-2010-dsp-var-prec-dsp-arch.html
- White Paper: *Enabling High-Precision DSP Applications with the FPGA Industry's First Variable-Precision Architecture*:
  www.altera.com/literature/wp/wp-01131-stxv-dsp-architecture.pdf

Acknowledgements

- Suhel Dhanani, Senior Manager, DSP Marketing, Altera Corporation
- Michael Parker, Senior Technical Marketing Manager, DSP Marketing, Altera Corporation

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
September 2010	1.0	Initial release.