



Embedded Signal Processing Capabilities of the LatticeECP3 sysDSP Block

A Lattice Semiconductor White Paper

February 2009

Lattice Semiconductor
5555 Northeast Moore Ct.
Hillsboro, Oregon 97124 USA
Telephone: (503) 268-8000
www.latticesemi.com

Introduction

Field programmable devices are continually being adopted in new market segments, where they are being implemented as mainstream logic devices. These new market segments are increasingly driving competing FPGA vendors to incorporate a wider variety of functionality and flexibility within their devices. Embedded digital signal processing (DSP) is one such function, addressing a wide gamut of market segments. In order to meet increasing market demands, these processing elements and their supporting hardware platforms must be able to provide increased calculation throughput without the expense of additional latency. For example, in 3G and 4G wireless applications, baseband and remote radio head (RRH) cards are required to handle both multiple protocols and increased throughput in order to support higher cellular data rates, even while maintaining a high Signal to Noise Ratio.

To address these emerging needs, Lattice Semiconductor has continued its tradition of providing high performance DSP capabilities in its most recent low-cost, SERDES-capable LatticeECP3 FPGA family. Features such as a dual slice architecture, the ability to cascade/chain DSP slices and blocks and an enhanced instruction set establish the LatticeECP3 family as a compelling alternative for signal processing applications such as FIR filtering and FFT/iFFT implementations.

Enhanced DSP Capabilities

The success of the LatticeECP2M family in addressing the 2G/3G wireless arena, coupled with discussions with customers, has served to drive the architectural improvements made in the LatticeECP3 DSP block. Several key points have been addressed. First and foremost, backward compatibility has been maintained with the current LatticeECP2/M DSP. There is also finer control of the DSP functions, which increases overall DSP block performance. The following sections of this white paper will outline each targeted feature, and how the LatticeECP3 family provides superb overall performance for many signal processing applications.

Figure 1 below provides a high level diagram of the major data flow differences between the LatticeECP2/M and LatticeECP3 DSP block.

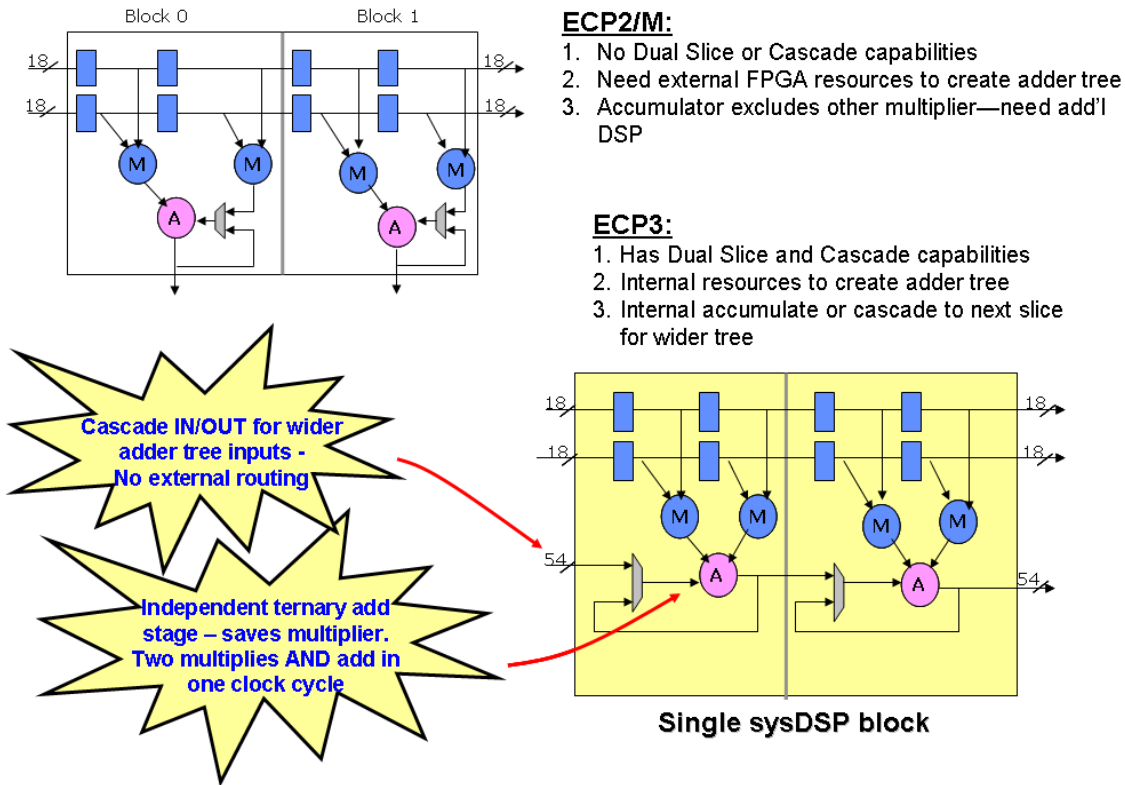



Figure 1 - LatticeECP2/M vs. LatticeECP3

Dual Slice Architecture

The first performance enhancement introduced in the LatticeECP3 DSP block is that the block has been divided into two identical slices. This feature has been added for several reasons: to enable increased performance within the DSP block, provide finer control capability and to allow independent ALU operation. Additionally, within each slice bypassable pipeline registers are available, permitting the designer to remove propagation latencies where necessary to meet rigid timing requirements. Figure 2 provides a block level view of the

slice architecture of the ECP3 DSP block.



- ◆ **Enhanced 3rd Generation sysDSP Architecture**
 - Fully cascadable blocks
 - Backwards compatible with ECP2M sysDSP block
 - Dual-slice architecture
 - » *higher performance, finer control*
 - » *Independent ALU operation*
- ◆ **Programmable Multipliers**
 - Two 18x18, four 9x9, one 36x36 for double precision / floating point
 - 36x36 across two adjacent slices
 - 18x36 MAC & 18x18 MMAC modes
- ◆ **54-bit Cascadable ALU**
 - Rounding & truncation

Figure 2 - Dual Slice Architecture

While each slice is logically separated, what is added is the capability to chain DSP slices with no routing penalties. One advantage of chained slices is that they permit wider multiplication and accumulator operations, with the partial products and sums from slice 0 to be propagated to slice 1, where the final product or sum will be computed. As indicated, this is all done within the DSP block without using external resources. If the chaining is not required, the slices can be configured to operate independently, with Input and output signals that are dedicated to each slice.

Cascadability

For many signal processing applications where large FIR filters or FFTs are employed, it may be necessary to create large signal processing functions. To accommodate this need, it is necessary to have DSP blocks cascaded together. The LatticeECP3 addresses the need for high performance signal processing functions by connecting the accumulator output of one block directly to its adjacent DSP block input. Figure 3 is a high-level diagram of the accumulator cascade I/O busses contained within each DSP block.

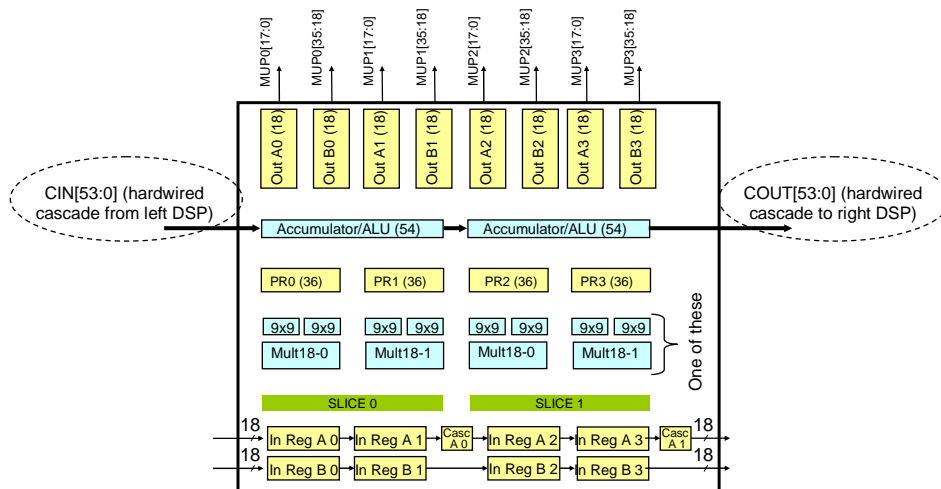


Figure 3 - DSP Cascade Diagram

This cascading is accomplished within the LatticeECP3 family by tiling the enhanced DSP blocks across specific rows, with the accumulator within each block being connected directly to the DSP accumulator adjacent to it. As with chaining, this interconnect does not use regular FPGA routing resources, which improves performance in terms of routing latency. Additionally, there is no external pipeline register required. Cascading can be very useful for some common DSP applications, such as creating larger adder trees in FIR filter applications. Without the cascade feature, as the adder tree grows, more levels of soft routing resources would be required, as well as pipeline registers to maintain the performance requirements. Larger cascade adder trees in LatticeECP3 FPGAs maintain increased performance when compared to LatticeECP2/M. Now, with cascade capability, up to four-stage adder trees can be generated, with the final adder placed at the end of the tree without any external FPGA logic. Some other useful DSP operations that are supported with the new cascade feature are rounding, barrel shifting and creating 36x36 multipliers.

Part of the cascade feature in the LatticeECP3 DSP block is the addition of a wide mux at the inputs to the accumulators found within each slice. The output of the mux feeds into the ternary-adder based ALU circuit. This can be seen in Figure 2, as the grayed out element to the left of the pink accumulator. Among other advantages, this new mux enables the introduction of a new instruction, MMAC: multiple-multiply and accumulates. This new instruction permits the sum of a slice accumulator to be multiplexed back into the accumulator as an addend for the next MAC operation.

Another improvement within the LatticeECP3 is its accumulator. The LatticeECP3 has a 54-bit data word accumulator, which doubles the MMAC instructions per cycle within a slice, as well as now supporting two 18x18 accumulations per cycle. The accumulator has been enhanced to contain a ternary adder. Combining the features of the wide mux with the multiple input adder permits improved summing, rounding and shifting operations.

Finally, as indicated previously, the third generation LatticeECP3 DSP is completely backward compatible with LatticeECP2 DSP but with added performance. Many of the existing opcodes are directly implemented, while others have been added or modified to ensure all legacy DSP instructions are supported within the new hardware platform.

High Performance Modes

Table 1 below provides a summary of the major functional similarities and differences between the LatticeECP2/M and LatticeECP3 DSP blocks. With the new LatticeECP3 features, advanced signal processing can be economically realized to meet advanced system performance requirements.

DSP Modes	ECP3	ECP2	Implication
MULT	Yes	Yes	
MAC	Yes	Yes	
MULTADDSUB (MultAdd2)	Yes	Yes	
MULTADDSUBSUM (MultAdd4)	Yes	Yes	
MMAC	Yes	No	<ul style="list-style-type: none"> • Re-uses logic to double accumulator performance • Single slice support for multiple MACs • Key for filtering and other applications where running sums must be maintained
Adder Tree	Yes	No	<ul style="list-style-type: none"> • Long (wide) addends can be summed w/o using generic FPGA resources... up to 30% increase in performance • Allows accumulator to be at the end of the adder without a performance penalty • Improves performance of barrel shifting/rounding
Wide Mux	Yes	No	<ul style="list-style-type: none"> • Two extra bits permit quadrupling of math functions (One extra bit for each doubling of data value). • Increased performance by using same slice for extra math operations.
SLICE	Yes	No	<ul style="list-style-type: none"> • Separates clock and control signals, providing more granular control per slice. • Improved performance/clock speed. • Cascading implemented across slices, no external routing resources needed.

 - New modes for ECP3

Table 1 – High Performance Modes

Summary

The world of signal processing is a dynamic and demanding one. Continuous improvement is expected, both from designers and vendors. Architectural improvements such as dual slice architecture allow for finer block control and higher clock speeds, while other enhancements, such as cascadability and a wider ALU/input mux, allow the LatticeECP3 to support more efficient implementations of processing functions, such as adder trees, by not requiring the use of any external logic. These are just some of the many improvements made to the LatticeECP3 family, demonstrating once again that Lattice is committed to delivering the best possible solutions. By combining high performance, low power and low cost, the LatticeECP3 FPGA family is the next chapter in the Lattice story of unparalleled overall value.