

ConnX SSP16 Soft Stream Processor

16-Way SIMD DPU for Multi-Standard Baseband PHY

FEATURES

- High-performance 16-way SIMD arithmetic processor
- Supports 3-bit, 8-bit, and 16-bit scalar data types and 8-bit vector data types that use 10-bit internal representation per element providing two guard bits
- Extensible interfaces with custom designed Port, Queue and Lookup interfaces
- Further tailor instruction set to meet the application needs
- 3-slot VLIW for efficient parallel load/store and compute operations
- Dual 128-bit load/store unit supports configurations with up to 4MB addressable region
- Architecture optimized for 8-bit and 10-bit vector operations
- C operator overloading
- 160-bit wide, 8-entry vector register file with support for signed 10-bit x 16 and 8-bit x 16 vector types
- 128-bit wide, 4 entry alignment buffer register file
- Optional acceleration units:
 - Soft demapper
 - Viterbi
 - PRBS, convolutional encoding

BENEFITS

- Optimized for small size and low power
- High performance on vector arithmetic and bitwise operations
- Multi-standard (LTE and HSPA+) soft bit processing
- Sixteen 10-bit or 8-bit vectors can be operated on in one cycle
- Ease of development with C code programming model
- Easy integration into existing hardware systems as well as multi-core systems
- High I/O throughput with the ability to create custom interfaces to hardwired custom coprocessors / RTL blocks

Small Size, Low Power, 16-Way SIMD DPU

The ConnX SSP16 (Soft Stream Processor) is a high-performance Dataplane Processor Unit (DPU) designed for use in SOC designs for next-generation communication baseband PHY-level systems such as those found in LTE and HSPA+ cellular radios and multi-standard broadcast receivers. It is specifically optimized for processing streams of soft bits, which are 4- to 8-bit representations of transmitted bits. Soft bits are generated by the demodulator in the receive chain.

The high compute requirements of HARQ pre-processing and header decoding within multi-standard wireless baseband applications require new and innovative architectures with a high degree of parallelism and efficient I/Os. The ConnX SSP16 meets these needs by combining a 16-way SIMD, 3-slot VLIW processing pipeline optimized for 10-bit and 8-bit processing (10-bit supports the required precision for multiple operations on 8-bit data).

The dual 128-bit wide data path allows 16-way loading and operations for higher performance. The ConnX SSP16 also supports specialized functions such as the transpose memory module and the Viterbi accelerator module.

The ConnX SSP16 is an integral part of the ConnX Atlas LTE Reference Architecture along with the ConnX BBE16 DSP core and the ConnX BSP3 and ConnX Turbo16 DPUs.

Instruction Set

The ConnX SSP16 is built on the baseline Xtensa RISC architecture, which implements a rich set of generic instructions optimized for efficient embedded processing. The power of the ConnX SSP16 comes from a comprehensive set of instructions for SIMD arithmetic operations operating on 10b or 8b vectors. The SIMD arithmetic operations include add, subtract, compare as well as logic operations such as shift, bitwise OR, bitwise AND, and bitwise negate. In total there are over 150 instructions using up to two slots.

Because ConnX SSP16 is a high performance 16-way SIMD computation engine, efficient loading of vectors into the register arrays is critical. Along with the dual 128-bit load/store support,

a full set of instructions is available that work with the alignment register bank to deliver full bandwidth loads and stores for aligned data accesses. With up to 16 vectors stored in the register array, flexibility and vector manipulation is offered within the instruction set with a comprehensive set of vector select instructions.

Specific acceleration has been added for high processing required algorithms for the baseband modem. These include the optional accelerators for LTE and HSPA+ Viterbi (giving 32-state update per cycle), soft demapper, PRBS, and convolutional encoding.

Programming Model

ConnX SSP16 was designed to get maximum performance from standard C code. Tensilica's advanced compiler automatically vectorizes code and performs automatic data alignment.

Looping delays are kept to a minimum and zero overhead looping is used. General arithmetic and logical operations can be automatically mapped to the SIMD engine by using operator overloading to give optimal performance. C intrinsics can be used for algorithms that can not be automatically vectorized.

Extensibility

The ConnX SSP16 supports custom defined Port (GPIOs) and Queue (FIFO) interfaces. The Ports can be defined to match the interfaces of existing RTL hardware blocks. Buffered communication between two ConnX SSP16s or between a ConnX SSP16 and an RTL block can be automatically implemented using queue interfaces and are fully supported in programming and modeling tools

Local memories can be connected directly to the ConnX SSP16 using the lookup interface, bypassing the processor memory bus.

This efficiently implements functions that require storage of multiple intermediate datasets. The ConnX SSP16 also can be modified and extended by defining new instructions, registers, and execution units to augment the existing instruction set.

Multi-core Integration

ConnX SSP16 is a building block within the Atlas Reference Architecture. As a result the integration and support of multi-core systems is a key factor in the definition of ConnX SSP16 DPU. Connectivity to other cores can be supported by a memory mapped AMBA AXI interface for a shared bus architecture. Also the Tensilica PIF interface can be used as a shared bus or a point-to-point connection scheme offering better performance on high data rate data paths as well as lower power. Control and synchronization can be done via PIF as well as queue and port interfaces. These multi-core systems are all supported in the Tensilica development tool chain, with multi-core system simulation, profiling and debug.

Toolchain

A complete set of tools are available to support the ConnX SSP16. A comprehensive instruction set simulator (ISS) is included as part of the Xplorer™ IDE, which allows developers to quickly simulate and evaluate performance. The fast, functional TurboXim™ simulator option achieves speeds that are 40 to 80 times faster than the ISS for efficient software development and functional verification. System C (XTSC) and C-based (XTMP) system modeling at the transaction level and pin level can aid in full-chip simulations.

The toolset includes a high-performance C/C++ compiler with automatic vectorization to support the VLIW pipeline in the SBP16. This comprehensive tool set also includes the linker, assembler, debugger, profiler, an energy estimation tool and graphical visualization tools. All major back-end EDA flows are supported.

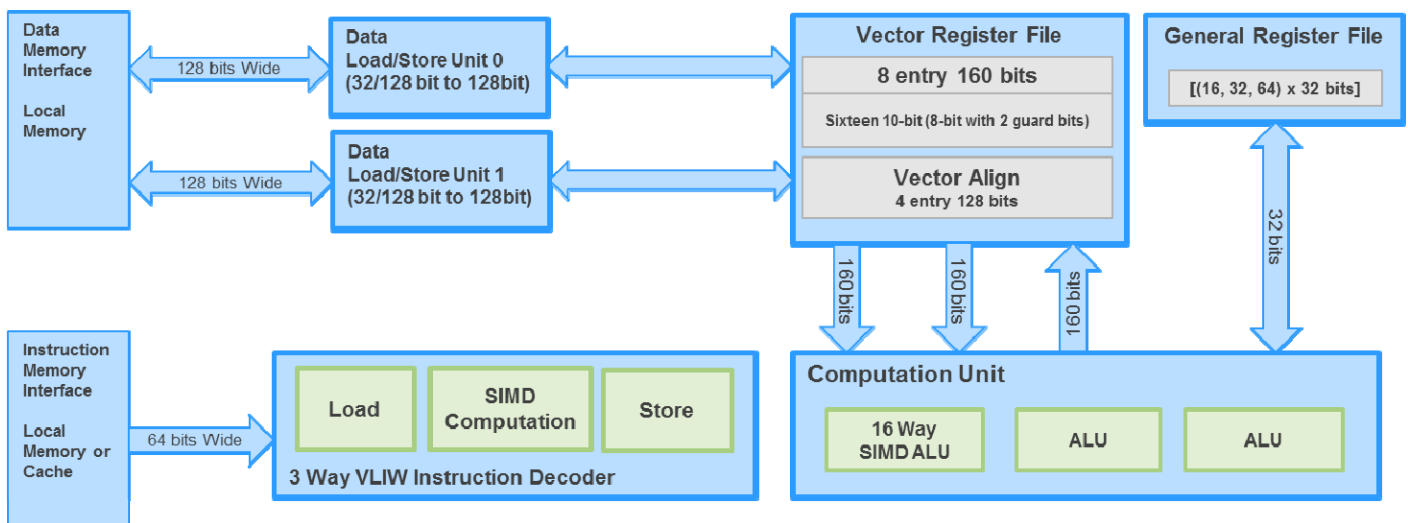


Figure 1. A simplified block structure of the ConnX SSP16 DPU



Tensilica, Inc.
 3255-6 Scott Blvd., Santa Clara, CA 95054-3013, USA
 Tel: 1-408-986-8000 Fax: 408-986-8919 Website: www.tensilica.com

©June 2012, Tensilica and Xtensa are registered trademarks of Tensilica, Inc. The Tensilica logo, Xenergy, Xplorer and TurboXim are trademarks of Tensilica, Inc. All other trademarks are the property of their respective owners.