cadence®

Tensilica Vision DSP Family

High-performance, low-energy vision/AI and image processing

The Cadence® Tensilica® Vision digital signal processor (DSP) family is designed for demanding embedded vision and artificial intelligence (AI) applications in the mobile, automotive, surveillance, augmented reality (AR) / virtual reality (VR), drone, and wearable markets. The Vision P6 DSP, the Vision Q6 DSP, and the Vision Q7 DSP are our three embedded vision-specific products that establish a new standard in high-performance, low-energy digital signal processing. Plus, since all our DSPs are built on the highly successful Cadence Tensilica Xtensa® processor, the Vision DSP family shares the same development environment.

Overview

For sufficient pixel-processing throughput, the Vision DSP family architecture incorporates advanced VLIW/SIMD support for the industry's highest number of ALU and MAC operations per processor cycle, as well as the industry's widest and most flexible memory bus.

Specialized instructions also allow the Vision DSP family to efficiently speed up pixel processing. Various architecture enhancements boost the performance while keeping the energy consumption low. The Vision DSP family provides unprecedented flexibility in system implementations at power-consumption levels that significantly reduce the need for hardware accelerators. The DSPs also offer an integrated DMA engine, interface for instruction memory, instruction cache, and two AXI interfaces. They offer industry's widest data memory bus of 1024-bit. In the Vision P6, Q6, and Q7 DSPs, we also offer an optional vector floating-point unit.

		Vision P6 DSP	Vision Q6 DSP	Vision Q7 DSP
Use Case		**Vision and AI up to 256MAC/sec**	**Vision and AI up to 384MAC/sec**	**Vision and AI up to 786MAC/sec**
MACs	8x8	256		512
	8x16	128		
	16x16	64		128
VFPU	16b half precision	32-way SIMD (optional)		2x 32-way SIMD
	32b single precision	16-way SIMD (optional)		2x 16-way SIMD
MAX SIMD Width		64-way 8-bit		
SuperGather		Yes		
Coefficient Decompression - Saves memory bandwidth		Yes		
AXI Interface		3 AXIs 2-128b Bus (1 master, 1 slave) shared by Instruction and Data 1-128b or 1-256 Bus for iDMA (master)	5 AXIs 2-128b Bus for Instruction 2-128b Bus for Data (2 masters, 2 slaves) 1-128b or 1-256 Bus for iDMA (master)	**3 AXIs** **2-128b Bus (1 master, 1 slave) shared by Instruction and Data** **1-128b or 1-256 Bus for iDMA (master)**

Vision DSP Family Features and Benefits

- Provides high-performance, energy-efficient vision and AI embedded DSPs
- Implements 64-way 8-bit SIMD with multiple VLIW slots
- Vision P6 DSP achieves up to 1.1GHz on 16nm process technology
- Vision Q6 and Q7 DSP achieve up to 1.5GHz on 16nm process technology
- Only DSPs in the industry to offer 1024-bit memory bus for transferring high-resolution data associated with today's imaging systems
- Provides a complete subsystem using the integrated DMA that allows the system to transfer high-resolution data directly into the local memory of the DSP, thus hiding the data access latency associated with accessing data from an external DRAM
- Implements the Tensilica SuperGather™ enhanced memory interface to quickly and efficiently read/write non-contiguous locations from local memory
- Features an instruction set customized for better code density, fewer cycles, and lower power
- The Vision DSP family delivers scalable, optimized performance with low energy for computer-vision and pixel-processing applications that span a large range of data types from 8b to 32b, such as face detection, object detection, lens distortion correction, and many more advanced vision applications
- Optional vector floating-point unit (VFPU) in the Vision DSP family offers flexibility to provide high-precision math at a minimal area penalty

Vision Q7 DSP Features and Benefits

The Vision Q7 DSP delivers up to 1.82 tera operations per second (TOPS), 1.7X higher TOPS compared to the Vision Q6 DSP in the same area. To address the increasing computational requirements for embedded vision and AI applications, the sixth-generation Vision Q7 DSP provides up to 2X greater AI and floating-point performance compare to, the Vision Q6 DSP. The Vision Q7 DSP is specifically optimized for simultaneous localization and mapping (SLAM), a technique **commonly used in the robotics, drone, mobile, and automotive markets to automatically construct or update a map of an unknown environment, and in the AR/VR market for inside-out tracking**.

- Very long instruction word (VLIW) SIMD architecture delivers up to 1.7X higher TOPS compared to the Vision Q6 DSP in the same area.
- An enhanced instruction set supporting **8/16/32-bit data types and optional VFPU support for single and half precision** enable up to 2X faster performance on SLAM kernels compared to the Vision Q6 DSP
- Delivers up to 2X improvement in floating-point operations per mm² (FLOPS/mm²) for both half-precision (FP16) and single precision (FP32) compared to the Vision Q6 DSP
- Up to 2X greater AI performance in the same area compared to the Vision Q6 DSP results in up to 2X improvement in GMAC/mm² compared to the Vision Q6 DSP

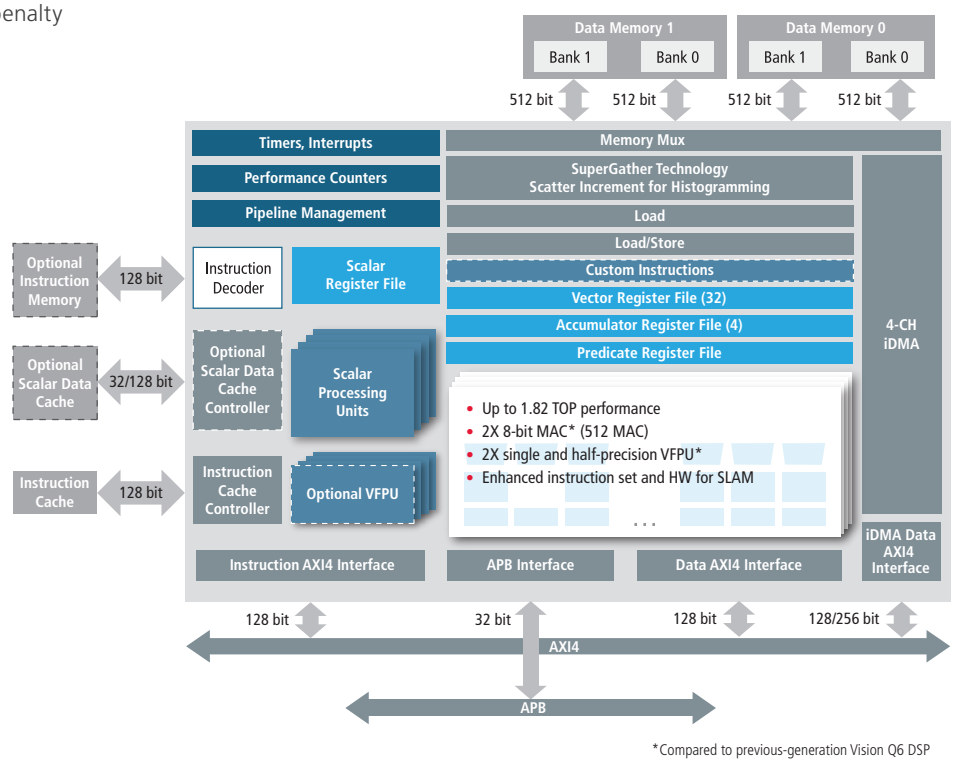

Figure 1: Vision Q7 DSP block diagram

Vision Q6 DSP Features and Benefits

A deeper, 13-stage processor pipeline and system architecture designed for use with large local memories enable the Vision Q6 DSP to achieve 1.5GHz peak frequency and 1GHz typical frequency at 16nm, in the same floorplan area as the Vision P6 DSP. As a result, designers using the Vision Q6 DSP can develop high-performance products that meet increasing vision and AI demands and power-efficiency needs.

- An enhanced DSP instruction set results in up to 20 percent fewer cycles than the Vision P6 DSP for embedded vision applications/ kernels such as Optical Flow, Transpose, and warpAffine, and for commonly used filters such as Median and Sobel
- 2X system data bandwidth with separate master/slave AXI interfaces for data/instructions and multi-channel DMA alleviates memory bandwidth challenges in vision and AI applications, and also reduces latency and overhead associated with task switching and DMA setup
- Backwards compatibility with the Vision P6 DSP, so customers can preserve their software investment for an easy migration
- Optional vector floating-point unit (VFPU) also supports half precision (FP16)

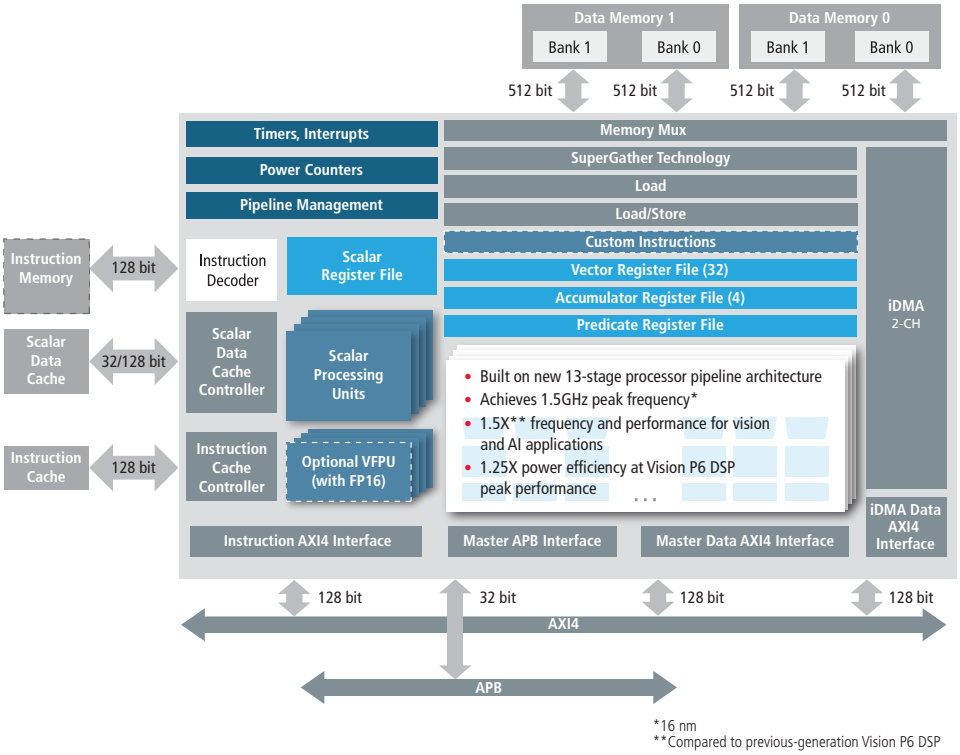

Figure 2: Vision Q6 DSP block diagram

Vision P6 DSP Features and Benefits

- Processes 9728 bits per cycle
- Offers 256 MAC
- Enhanced instruction set
- Smart instruction slotting
- Supports SuperGather enhanced memory interface
- Optional VFPU with single-precision (32-bit) floating-point and half-precision (16-bit) floating-point support offers flexibility to port GPU code and high-precision math at a minimal area penalty

Library Support

- Vision P6, Q6, and Q7 DSPs offer highly optimized OpenCV-based library functions, which can accelerate applications development
- Vision P6 and Q6 DSPs offer OpenVX-based support to enable fast application development in a heterogeneous environment
- Vision P6, Q6, and Q7 DSPs offer highly optimized NN library

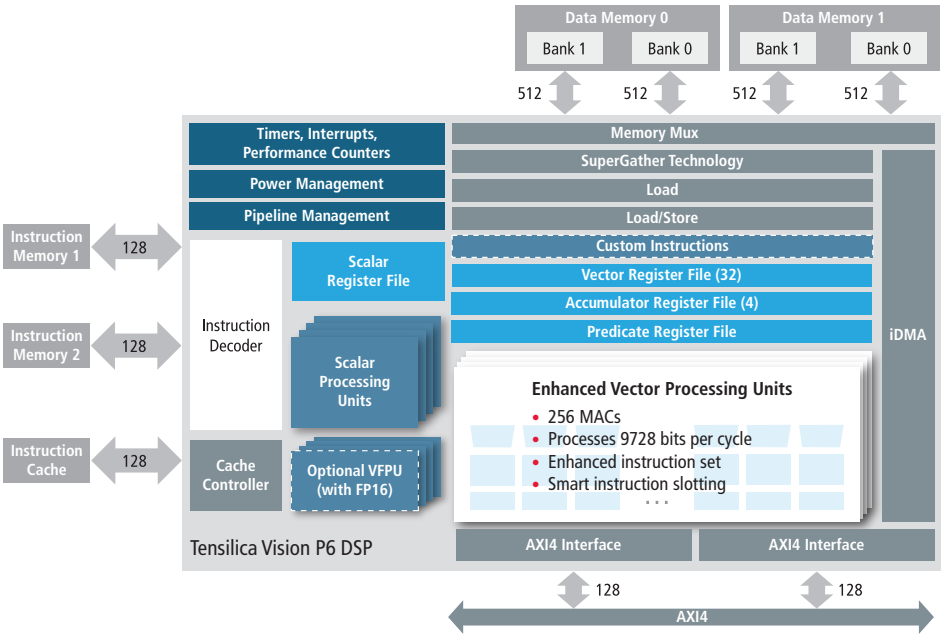

Figure 3: Vision P6 DSP block diagram

Toolchain

The Vision DSPs are delivered with a complete set of software tools:

- A high-performance C/C++ compiler with automatic bundling and vectorization supports the VLIW and SIMD capabilities.
- Linker, assembler, debugger, profiler, and graphical visualization tools are included
- A comprehensive instruction set simulator (ISS) allows you to quickly simulate and evaluate performance
- When working with large systems or lengthy test vectors, the fast, functional TurboXim simulator achieves speeds that are 40X to 80X faster than the ISS for efficient software development and functional verification
- Tensilica Xtensa Modeling Protocol (XTMP) for system modeling in C and Xtensa SystemC (XTSC) for system modeling in SystemC® provide for full-chip simulations. The pin-level XTSC model offers co-simulation of the SystemC model at the pin level for fast, cycle-accurate system simulations
- Vision P6, Q6, and Q7 DSPs come with support for the Tensilica Neural Network Compiler, which maps any neural network trained with a framework such as Caffe, TensorFlow, and TensorFlow Lite into executable and highly optimized fixed-point code for target DSPs, leveraging a comprehensive set of hand-optimized neural network library functions. Vision P6 and Q6 DSPs also support the Android Neural Network API allowing the customers to take advantage of the higher performance offered on these DSPs on the Android platforms.
- All major back-end EDA flows are supported

Cadence Services and Support

- Cadence Tensilica application engineers can answer your technical questions and provide technical assistance and custom training.
- Cadence-certified instructors teach a series of courses on Tensilica IP and bring their real-world experience into the classroom.
- Internet Learning Series (iLS) online courses allow you the flexibility of training at your own computer via the Internet.
- The Cadence Tensilica IP support site gives you 24x7 online access to a knowledgebase of the latest solutions, technical documentation, software downloads, and more at ip.cadence.com/support.

cadence®

Cadence software, hardware and semiconductor IP enable electronic systems and semiconductor companies to create the innovative end products that are transforming the way people live, work, and play. The company's System Design Enablement strategy helps customers develop differentiated products—from chips to boards to systems. www.cadence.com