

# ARC EV7xFS Processor IP

## Highlights

- Adds hardware safety features to EV7x vision processor with minimal area and power impact
- Integrates 32-bit scalar core, 512-bit vector DSP, and optional DNN accelerator to deliver up to 35 TOPS performance
- IEEE 754-compliant vector floating point unit option offers single precision or half precision operations and advanced math functions
- Supports ISO 26262 automotive safety standard
- ASIL B, C, or D Ready hardware and/or Hybrid mode
- Software test libraries complement integrated hardware safety features to achieve ASIL B compliance
- Supported by MetaWare EV Development Toolkit for Safety with ASIL D Ready compiler and graph mapping tool
- Extensive safety documentation eases certification process

## Target Applications

- Safety-critical automotive systems using AI and deep learning
- Self-driving cars
- Radar/lidar systems
- Advanced Driver Assistance Systems (ADAS)
- Industrial applications requiring enhanced reliability

## Overview

The ASIL B or D Ready DesignWare EV7xFS Embedded Vision Processors enable automotive system-on-chip (SoC) designers to accelerate Advanced Driver Assistance Systems (ADAS) and autonomous vehicle application development and ISO 26262 certification for systems using vision processing, artificial intelligence (AI) and deep learning. To provide greater flexibility to automotive design teams and address evolving requirements, the EV7xFS offers an Automotive Safety Integrity Level (ASIL) B/D Ready “hybrid” option that enables users to select ASIL configurations levels that trade-off performance vs safety levels post-silicon, in software. The EV7xFS processors include state-of-the-art hardware safety features including diagnostic error injection, windowed watchdog timers, error classification, and software diagnostic tests as well as safety monitors and lockstep capabilities for safety-critical modules. Comprehensive safety documentation, including safety manuals, FMEDA and DFMEA reports accelerate SoC-level functional safety assessments. These features enable designers to achieve high levels of fault coverage as required for ASIL certifications without a significant effect on performance, power or area compared to the non-ASIL Ready EV7x processor.

The DesignWare EV7xFS Embedded Vision Processors are fully programmable and combine the flexibility of software solutions with the high performance and low power consumption of dedicated hardware. The EV7xFS family integrates up to four high-performance 32-bit scalar cores and 512-bit vector DSPs, and an optimized deep neural network (DNN) acceleration engine for fast and accurate object detection, classification and scene segmentation.

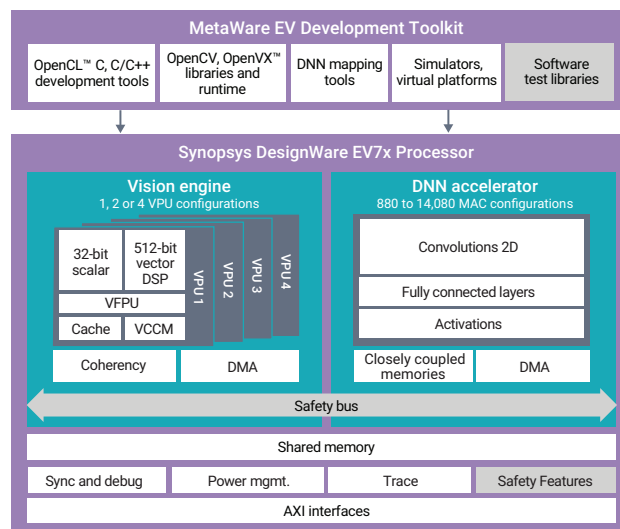


Figure 1: DesignWare EV7xFS block diagram

## EV7x Safety Features

- ASIL B, C, or D Ready, or hybrid ASIL B/D
- Integrated safety-critical hardware features
- ECC memories
- Diagnostic error injection
- Error classification
- Error checking on core registers and safety-critical registers
- Windowed watchdog timer for each core
- Software diagnostic tests
- Lockstep capabilities for safety-critical modules
- Optional dedicated Safety Island monitors and executes safety escalations and diagnostics within the SoC and protects system bring-up
- Safety documentation: FMEDA reports and safety manuals speeds functional safety assessments

## EV Vision Engine

The EV7xFS vision engine features a quad issue super vector architecture with a 32-bit, high-performance scalar pipeline, and a 512-bit wide SIMD Vector DSP (VDSP) that are optimally balanced to achieve excellent performance with low power consumption. The vector processing unit executes one scalar and three vector instructions (128-bit instruction bundle) per cycle. The vision engine is supported by configurable instruction and data cache for scalar operations and vector closely-coupled memory (VCCM) – a single cycle access RAM – for vector processing.

The EV7xFS Processors support 8KB to 128KB of data closely coupled memory (DCCM), and 32KB to 256KB of vector memory (VCCM) for the VDSP. The CCM is implemented as a separate memory space and can be accessed every clock cycle.

DMA for the VPU(s) is provided by a Streaming Transfer Unit (STU) which is tightly coupled with the VCCM. The STU brings in external data for vector DSP processing, supports 1D and 2D (image) transfers and IO Coherency, and has 4x 128b connections to the AXI bus. The coherency unit, the optional clustered shared memory and hardware support for synchronization (semaphores, interrupt dispatches, small messages in a mailbox) allows the multiple VPUs of the EV72 and the EV74 to be integrated seamlessly.

The EV7xFS includes an optional IEEE-754 compliant Vector Floating Unit that supports both full (32-bit) and half (16-bit) floating point operations. This high performance VFPU implementation uses pipelined, high-performance floating point components to achieve up to 512 GFLOPS (4 VPUs, 1 GHz, half-precision). The VFPU also supports an extensive set of math functions including:  $\text{div}$ ,  $\sqrt{x}$ ,  $1/\sqrt{x}$ ,  $\sin(x)$ ,  $\cos(x)$ ,  $\log_2(x)$ ,  $2^x$ ,  $e^x$ ,  $\text{atan2}(x)$

## DNN Accelerator

The optional embedded deep neural network accelerator adds scalable deep learning and AI capabilities to the EV7x family. The DNN accelerator is optimized for CNNs and batched or convolutional Recurrent Neural Networks (RNNs) or LSTMs and includes advanced hardware features to support the latest pruning, compression and layer merging techniques to increase performance and minimize bandwidth. The DNN can be configured from 880 multiply-accumulators (MACs) up to 14,080 MAC versions. Most of the MACs are used for 2D convolutions while a portion is dedicated to 1D convolutions needed for fully connected layers. The DNN datapath supports 8- and 12-bit data precision. The DNN accelerator supports flexible activation functions, including ReLU, PReLU, ReLU6, tanh and sigmoid. The EV7x supports all CNNs including popular networks such as MobileNet, GoogLeNet, ResNet, Yolo, Faster R-DNN, and ICNet. Designers can run CNN graphs originally trained for 32-bit floating point hardware on the EV7x's DNN accelerator using 8- or 12-bit resolution significantly reducing the power and area of their designs while maintaining high levels of detection accuracy.

In addition to supporting CNNs, the DNN supports batched LSTMs (long short-term memories) for applications that require time-based results, such as predicting the location of a pedestrian based on their observed path and speed.

The DNN accelerator is supported by a high-performance DMA for transferring image data from external memory into the internal closely coupled memories.

## Cluster Shared Memory

A low-latency shared data memory is included in the processor to support information passing and coordination between the multiple CPUs and the DNN processing element cores. This memory is used as a software-managed scratch pad and is configurable from 0 to 8MB. To allow for larger sizes, the memory is internally multi-banked, but this is invisible to the software. It includes arbitration to support concurrent access from the CPU cores and/or the DNN processing elements. The shared subsystem data memory is optional.

## Real-Time Trace

The DesignWare® ARC® Real-Time Trace (RTT) unit is a hardware helps trace executed instructions or program flow and data. ARC RTT generates Nexus 5001 class 3-compliant trace messages. The RTT system can be set up in many different configurations which need to be specified as build-time configurations by including the trace generator in the core and the RTT module at build time. ARC RTT can support on- and off-chip memory setups to suit your application tracing needs.

## Comprehensive Software Environment

The ARC MetaWare EV for Safety Development Toolkit is a complete set of tools that provides everything needed to program the EV7xFS processors.

The MetaWare EV For Safety Toolkit includes the MetaWare C/C++ Compiler, MetaWare Debugger and Instruction Set Simulator and adds an OpenCL C compiler for writing vision kernels for the vector DSP. The DNN mapping tool provides a complete CNN software environment to support Caffe and Tensorflow neural network frameworks with additional frameworks supported via the industry-backed ONNX interchange format. DNN graph training is done off-line, typically on a server farm, and the resulting graph is programmed into the object detection engine by the user with the DNN graph mapping tool.

| Components                     | Description   |
|--------------------------------|---|
| MetaWare EV Compiler           | Single ASIL D Ready compiler supports C/C++ and OpenCL C programming                  |
| MetaWare EV Debugger and nSIM  | Debugger and simulator for debugging, profiling and optimizing                        |
| EV Runtime and Libraries       | OpenVX and OpenVX SC frameworks and libraries eases vision graph development          |
| DNN Software Development Kit   | ASIL D Ready DNN graph mapping tools automates mapping to the DNN acceleration engine |
| EV Virtualizer Development Kit | EV Virtualizer Development Kit  |

Table 1: MetaWare EV Development Toolkit for Safety components

## About MIPS:

MIPS by GlobalFoundries delivers software to silicon with RISC-V for building physical AI platforms. MIPS delivers software-hardware co-design, optimized AI, and custom ASSP design and manufacturing. Together with ARC, MIPS delivers the open, standards-based processor IP portfolio for embedded applications. Physical AI is built on MIPS.

For more information, visit [www.mips.com/arc](http://www.mips.com/arc).