

XCORE[®].AI

THE FAST, FLEXIBLE, ECONOMICAL CROSSOVER PROCESSOR FOR THE AIOT

A crossover processor for the artificial intelligence of things (AIoT), helping you to get to market fast, with products that stand out from the competition.

AIoT product proliferation creates an increasing demand for a flexible, performant processing core that offers short time-to-market. Demands around privacy and user-experience are pushing tasks from cloud to edge (on-device) processing and challenging the capabilities of traditional processors.

With xcore.ai, real-time inferencing, decisioning at the edge, signal processing, control and communications are wrapped up in a single chip. Product Designers no longer need to rely on a costly applications processor or a microcontroller with additional components.

xcore.ai combines AI acceleration, powerful DSP, connectivity and general-purpose compute in a low eBOM solution - ideal for developers working on smart products that enhance our everyday living.

FEATURE HIGHLIGHTS

The device has two "tiles." Each tile is a self-contained processor with 512 kByte single cycle SRAM. The tile has a scalar unit (up to 1600MIPS), a vector unit (up to 25,600 MMACS), and a floating point unit (up to 800 MFLOPS); 1 Mbyte tightly coupled SRAM, 3200 MIPS, 1600 MFLOPS, and 51,200 MMACCS across the tiles. The device has three integrated PHYs: a high-speed USB, a MIPI D-PHY receiver, and LPDDR1.

FLEXIBILITY

Supports 1.8V and 3.3V on each bank of I/O, allowing the exact mix of peripheral interfaces and connections for your system. With USB, MIPI RX and a software driven I/O structure supporting I2C, I2S, S/PDIF, SPI, UART, PWM, and many more, xcore.ai offers a single chip solution (replacing the need for multi-components).

ON-DEVICE INTELLIGENCE


Convolutional (CNN) and Dense (DNN) Neural Network acceleration supports 32-bit, 16-bit, 8-bit and 1-bit vector operations. Instructions loading a vector from memory, multiplication and accumulation occur in a single clock cycle (achieving 51.2 GMACC/s peak, and 30GMACC/s sustained 8-bit vector arithmetic).

TOOL CHAIN

Industry standard tools enable accelerated applications. xcore.ai is fully programmable in C and C++, with optimised libraries for rapid prototyping and development. FreeRTOS support and deep learning framework TensorFlow-Lite ensures developers will enjoy an easy familiarity with the programming model.

COMPUTE

With up to 3200MIPS of compute, the fast and predictable xcore.ai can handle the most challenging edge-AI operations. With up to 8 FIR taps (32'bit) per clock cycle and 1 million 512 FFTs per second, it's ideal for voice, imaging and sensor processing applications.



APPLICATIONS



SMART SPEAKER
AUDIO VISUAL
APPLIANCES
LIGHTING/SECURITY



HEALTH
FITNESS
CARE & DIAGNOSTICS
MONITORING



TRAFFIC & PARKING
ENVIRONMENTAL
PUBLIC SAFETY
MONITORING



SAFETY
TRACKING
MAINTENANCE
ENERGY MANAGEMENT



ASSET TRACKING
PREDICTIVE CARE
PEOPLE TRACKING
AUTONOMOUS (L1)

DEVELOPMENT TOOLS

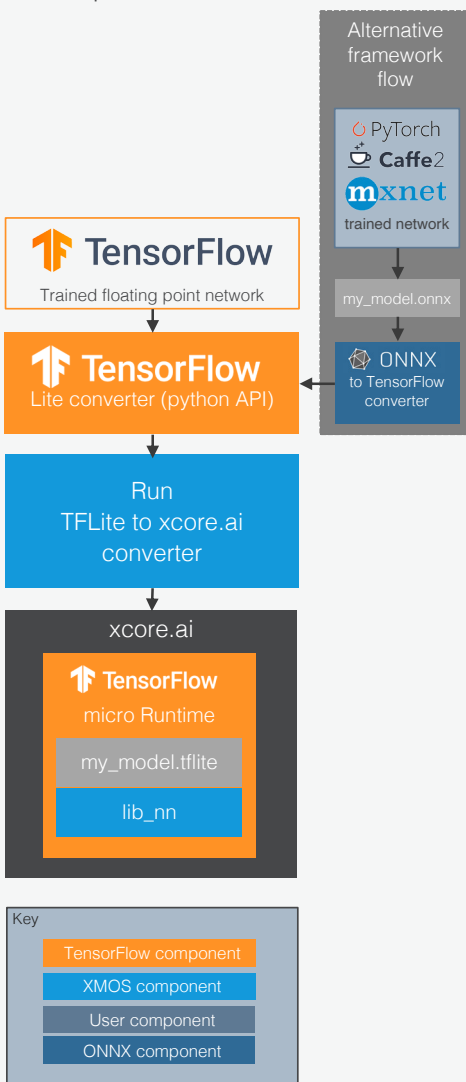
Standard programming tools:

- LLVM compiler suite
- GNU debugger
- Make for building

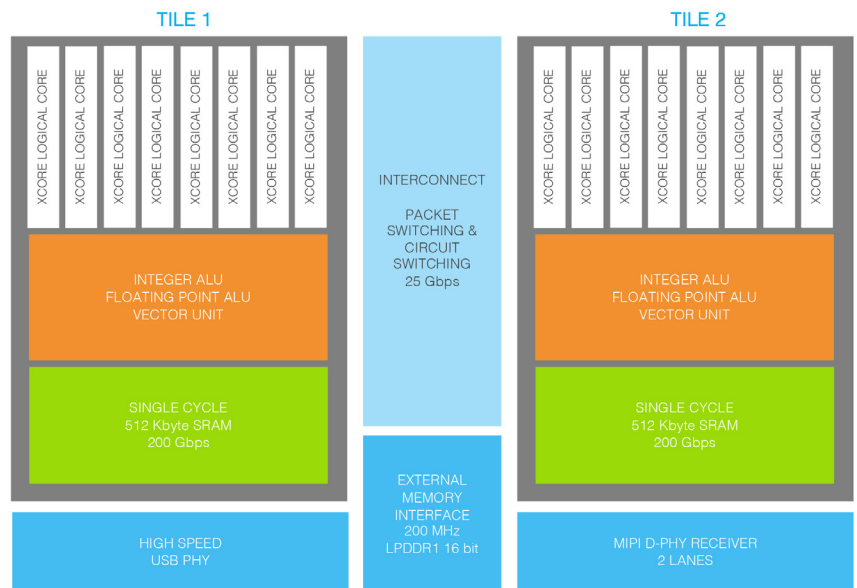
Bespoke tools include:

- Mapper to link code and eliminate unused code
- Global optimiser
- Programmer for Flash, OTP and volume programming

TensorFlow Lite for microcontrollers enables high performance code for immediate compilation / modification to deliver optimised models.



BLOCK DIAGRAM



CROSSOVER PROCESSOR SPECIFICATION

PACKAGE	7x7mm, 60-pin QFN 0.4mm pitch 14x14mm, 265-pin BGA 0.8mm pitch
COMPUTE	16 real-time logical cores across 2 tiles
MEMORY	2 x 512KB on-chip SRAM, extendable through an LPDDR1 interface or external flash
CONNECTIVITY	Up to 128 GPIO defined in software; port sampling rates up to 100MHz with respect to external clock; selectable I/O voltage 1V8-3V3 (package dependent)
AI	Vector unit with 256bit wide registers; operates in 32/16/8bit integer mode. Binarised neural network (XNOR) support
DSP	32x32 bit complex vector arithmetic, 32x32 -> 64bit MAC instructions, and single precision floating point
PERIPHERALS	Built-in physical layers for high-speed USB PHY; Dual-lane MIPI D-PHY receiver; and PLL for application clock generation
SECURITY	8 kByte internal OTP



For further information, please visit xmos.ai/xcore-ai/

Copyright © 2022 XMOS Ltd. All rights reserved