

Fibonacci ML SoC

General description


Hierarchical scalability is the foundation principle of CSEM's Fibonacci machine-learning (ML) system-on-chip (SoC). Like the Fibonacci number series, combining each element by the sum of the previous ones, the SoC can dynamically increase its computational performance by adding accelerator resources based on the application's needs. Its heterogeneous architecture features a low power time-series ML accelerator (FETA), two clusters of highly parallelized neural processing units (NPU), energy-optimized on-chip memories, a flexible RISC-V microcontroller core, and a rich set of peripherals for easy system integration. Trained models can be deployed through CSEM's ML compiler, supporting all common formats (e.g. ONNX).

NPU Clusters

- Optimized for spatial neural networks (e.g. CNNs, ResNets, MobileNets)
- Sparsity exploitation
- Peak MAC performance: 960 GOPS

FETA Cluster

- Optimized for temporal neural networks (e.g. RNNs like LSTM or GRU)
- Smart temporal feature extraction engine



Target applications

- Multi-modal concurrent data analysis from different sensor types (e.g. audio-visual sensor fusion)
- Multi-stage evaluation: hierarchical execution with increasing complexity to reduce average power consumption.
- Low power edge processing, down to uW power budgets
- Spatial and time-series signal analysis

Features

- General purpose RISC-V core (RV32IMC)
- Standard communication peripherals: UART, I2C, SPI (x2), Octo-SPI, DCMI, I2S
- JTAG debugging interface
- Up to 4 MB of on-chip SRAM + 0.5MB of MRAM
- Multi neural network execution
- Selective execution and early exit
- Dynamic precision scaling
- Dynamic power switching
- Bank and block level power gating
- Flexible DMA engines (x2)
- Power consumption: 100 uW - 500 mW

