



Data Pre/Post-processing Engine

MangoBoost's patent-pending Data Pre/Post-processing Engine (MB-DP) offers data pre/post-processing and orchestrates the data communication between PCIe devices and CPU (e.g., SSD, GPU, MangoDPU).

Key Features

- **Configurable architecture** using standard AXI interfaces and Partial Reconfiguration, which supports **easy integration with custom IPs**
- User space C and Python based APIs for **light-weight software integration**
- **Fully compatible** with the existing kernel drivers of the other devices (e.g., NVMe SSD, NVIDIA GPU)
- Supports AMD Alveo-series FPGAs
- Written in C-based HLS (AMD Vivado)
- SystemC version available soon (tested using Siemens Catapult HLS) for portability to other FPGA/ASIC targets

Application Use Cases

- AI inference server (e.g., Triton)
- AI training server
- RADAR processing

Case Study: AI Inference

- **Setup:** MB-DP on AMD U280 FPGA with two Nvidia A40 GPUs, on host Intel Xeon Gold 6338 2.0 GHz CPUs
- Installed Nvidia Triton inference server with Resnet-50 and client over a single 100-GbE network
- MB-DP runs at 250MHz, utilizing 151K LUTs, 763K Regs, 1973 BRAMs, 11 URAMs
- Achieved **1.64x higher inference throughput** and **1.33x higher GPU utilization** versus software baseline (i.e., Nvidia Dali) with **88% CPU savings**

