Dracula’s Capabilities

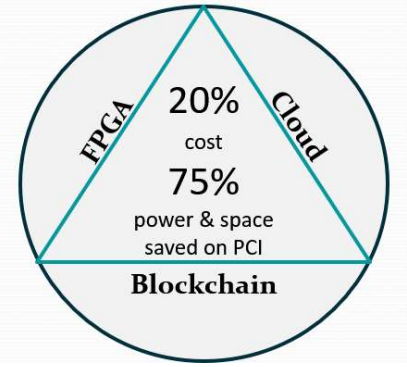

General

Ypologist’s Dracula is a massively parallel novel architecture targeting Deep Learning and Neural network applications. Dracula's many-core FPGA and ASIC implementations will deliver excellent cost and power-efficient alternatives to accelerators such as Nvidia’s GPCPU or Intel’s MIC.

Integration with Cloud Computing

The structure of the accelerator consists of a linear array of execution/processing cells with two global loops, a reduction loop and a scan loop closed over the map array, working under the control of a programmable sequencer. The accelerator is connected to a host system responsible for the complex part of the computation.

Dracula seamless software integration with AI Heterogeneous Computing (HC) clouds is achieved by Ypologist optimized AI libraries such as TensorFlow. Further, Dracula FPGAs can be used as reconfigurable and virtualized accelerators in current Azure and AWS EC2 F1clouds. ASIC implementations further reduce the total cost of cloud operations by reducing the use of space, electric power and cooling requirements in data centers.

Performance highlights

The main application domains, for which libraries will be developed, are linear algebra, cryptocurrencies, automotive and bio-informatics. The following are the performance characteristics for each of these applications:

- **Linear Algebra** 2 to 3X improvement in energy savings and 3X increase the ratio performance/peak performance compared with off-the-shelf solutions. In most Big-Data applications power consumption per task is 20%.
- **Cryptocurrency** applications request high MH/sec/Watt (MH stands for mega hashes). Measured performance 337% against NVIDIA GPU.
- **Automotive stereo** vision, where tight real-time performance and temperature requirements are mandatory, Dracula's 1024-cell system consumes less than 6W at 100° C.
- **Bio-Informatics** applications are accelerated 6X when implemented in a 28nm FPGA when compared with a 22nm 4-core Intel CPU with SSE. Simulations predict a 30X improvement. The energy efficiency is improved by 20X for the Dracula PGA version and 300X for the ASIC version.