GRAPHCORE

POPLAR OVERVIEW

Poplar™ is a scalable graph programming framework targeting Intelligent Processing Unit (IPU) accelerated servers and IPU accelerated server clusters, designed to meet the growing needs of both advanced research teams and commercial deployment in the enterprise. It's not a new language, it's a C++ framework which abstracts the graph-based machine learning development process from the underlying graph processing IPU hardware.

Poplar includes a comprehensive, open source set of Poplar graph libraries for machine learning. In essence, this means existing user applications written in standard machine learning frameworks, like Tensorflow and MXNet, will work out of the box on an IPU. It will also be a natural basis for future machine intelligence programming paradigms which extend beyond tensor-centric deep learning. Poplar has a full set of debugging and analysis tools to help tune performance and a C++ and Python interface for application development if required.

The Poplar graph compiler has been built from the ground up for translating the standard operations used by machine learning frameworks into highly optimized application code for the IPU. The graph compiler builds up an intermediate representation of the computational graph to be scheduled and deployed across one or many IPU devices. The compiler can display this computational graph, so an application written at the level of a machine learning framework reveals an image of the computational graph which runs on the IPU.

The Poplar framework is designed to be extensible; the IPU will accelerate today's deep learning applications, but the combination of Poplar and IPU provides access to the full richness of the computational graph abstraction for future innovation.

Poplar™ Highlights:
Optimised graph processing framework

- Supports the powerful computational graph abstraction that underpins efficient machine learning development
- Abstracts the graph-based machine learning development process from the underlying graph processing IPU hardware systems
- Includes:
    - Seamless interface to multiple machine learning frameworks such as TensorFlow
    - C++ interface for application development
        - Future support for Python interface
    - Comprehensive set of open source libraries and primitives which allows the Poplar framework to be easily integrated into existing applications
    - Full set of debugging and analysis tools to aid application development and performance tuning

GRAPHCORE

POPLAR OVERVIEW

- Advanced graph compiler that optimizes hardware utilization on an IPU graph processor
- Runtime graph engine – for IPU graph processor deployment
- Hardware abstraction layer for future proof IPU support
- Drivers and comprehensive IPU system support

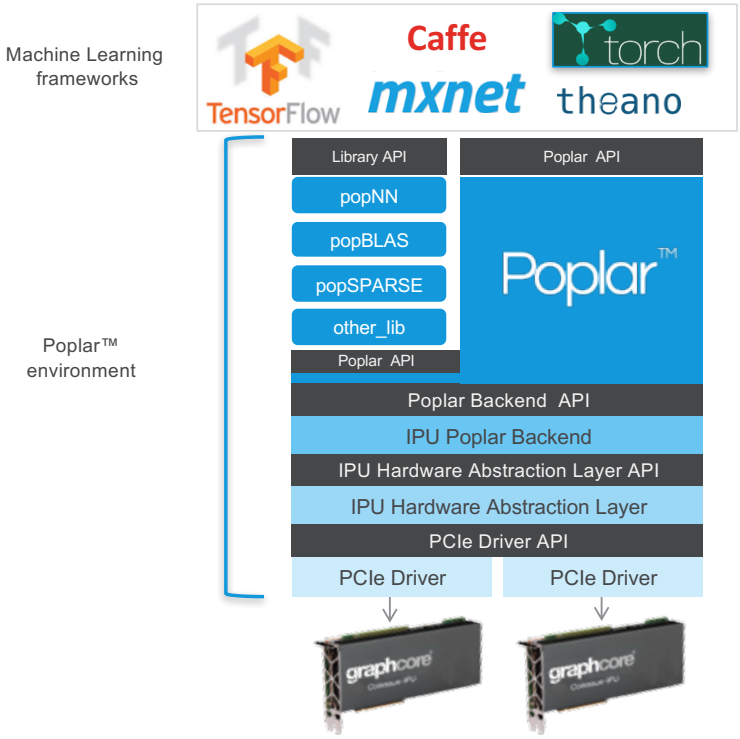

Machine learning framework independent

- Seamless interface to current and future machine learning development tools
- Allows machine learning projects to be developed using a variety of industry standard machine learning frameworks and tools, targeting single or multiple IPU processors
- Supports easy scaling across multiple IPU processors
- Optimised support for **TensorFlow**
- Future support for: **MxNET**, **PyTorch**, **Caffe-2**, **CNTK**, and many others...

Allows much faster innovation in machine learning systems

- Supports current and next generation machine learning systems
  - Graph abstraction model allows full flexibility to support the development of innovative new machine learning models and deep learning network topologies
- Fully supports model parallelism, collaborative models and generative adversarial model structures

GRAPHCORE

POPLAR OVERVIEW

Poplar graph libraries

- C++ and Python user interfaces
- No new languages to learn
- Comprehensive set of open sourced Poplar libraries and primitives including
    - ***popDNN***: deep neural network libraries
    - ***popBLAS***: linear algebra libraries
    - ***popSparse***: sparse matrix multiply libraries
    - ***popRandom***: random number and noise shaping libraries
    - *Future support for*
        - ***popFFT***: Fast Fourier Transform libraries
        - ***popRobotics***: SLAM, trajectory planning, autonomous car and robotics primitives
- Fully supports the ability to develop your own libraries and primitives
    - Modify and extend open sourced Poplar libraries
    - All libraries developed using Poplar framework with source code included

Scalable hardware deployments and virtualization support

- Single IPU accelerated servers
- Chassis based multiple IPU accelerated servers
- Clusters of IPU accelerated servers
- Embedded IPU accelerated devices
- IPU accelerated server virtualization including support for Docker

Comprehensive debug tools

- Optimised machine learning debug suite
- Supercomputer-class tools for comprehensive machine learning application debug
- Methodology designed from the bottom up and integrated fully into IPU architecture
- Powerful visualization tools

Documentation and support

- Easy to use system and familiar graph methodology
- No new language to learn
- Comprehensive documentation delivered through website
- Online tutorials, code examples and application notes

Drivers and support tools

- User space PCIe driver allowing host access to IPU card and IPU applications
- Virtual to physical memory mapping and comprehensive IPU-DMA support
- Supports multiple IPU card servers and systems
- Supports Docker containers and IPU accelerated server virtualization