

Applications Engineer Job Description:

OVERVIEW

The use of machine learning to reimagine software applications and service development is exploding. Companies from every corner of the industry -- the biggest cloud service providers to corporate industrials to financial services to healthcare to retailers -- are exploring new ways of building products and services using data-centric learning models in place of traditional explicit programming. The drive to deliver more timely and more accurate results is compelling an ever greater need for specialized computing power. GPUs have been hailed as the solution to those computing needs but the industry is actively searching for a better, more efficient solution. Graphcore has that solution.

Graphcore's Intelligence Processing Units, or IPUs, are specifically designed for artificial intelligence and compute-dense graph applications. These are not GPUs, Graphics Processing Units, but rather graph processors especially adept at the kinds of computations used in understanding relationships within a sea of data. Architecturally, IPUs look nothing like GPUs. They offer performance, latency, and power efficiency advantages a GPU will be unable to match. Graphcore offers support for popular industry ML frameworks and a full tool suite for developers to innovate both within and outside those frameworks.

As an Application Engineer, you will work to port and optimize machine learning and artificial intelligence applications using Graphcore's Poplar™ software and IPU processors, enabling breakthroughs in this rapidly moving field. You will create Application Notes and blog content, and work closely with Graphcore's field teams, customers and partners to help them in understanding and getting the most from our Intelligence Processing Unit (IPU) technology. You will also act as a senior technical figure within our product support organization, debugging customer issues and providing concise summaries and recommended fixes to our core engineering teams.

In your work you will support some of the world's top machine learning innovators at deep learning research groups, at academic institutions, at innovative machine learning start-ups, at leading automotive companies, and at some of the world's largest cloud and internet companies. You will need to develop a deep understanding of the IPU architecture and the associated Poplar™ tools and become familiar with leading machine learning frameworks.

RESPONSIBILITIES

- Port and optimize machine learning applications and create new high performance library implementations using Graphcore's Poplar™ tools
- Become a recognized expert on Graphcore's IPU technology and Poplar™ tools and deliver compelling training to our customers and partners
- Field & resolve challenging/complex customer support issues
- Shepherd critical customer issues and provide timely advance warning of critical issues that need attention
- Work with the Product Management and Engineering to ensure a good flow of customer and market feedback that can be incorporated into future products

REQUIREMENTS

- 4+ years of related experience writing and optimizing applications, scientific libraries or compilers for high performance computing, digital signal processors, GPUs or FPGAs
- Bachelors in Engineering, Computer Science, Mathematics, Physics or similar field
- Deep experience with C++ and in-depth knowledge of computer architectures, high performance programming and parallel programming
- Ability to multitask effectively in a fast paced environment
- Action-oriented with strong analytical and problem solving skills
- Strong time-management and organization skills for coordinating and prioritizing multiple projects and initiatives
- Keen interest to learn about the exciting new field of AI

DIFFERENTIATORS

- Masters or PhD in related computationally intensive science or engineering field
- Experience working with modern deep learning software architecture and frameworks including: Tensorflow, MxNet, Caffe, Caffe2, Torch, and/or PyTorch
- Excellent communication & presentations skills and comfortable in a customer-facing environment
- Experience working with PCIe form-factor accelerators such as GPUs, DSPs or FPGAs
- Experience in the AI or Machine Learning space

This position will be based at our office in our Bristol, UK or Palo Alto, USA and could require frequent travel to customer & partner sites.

Graphcore is an equal opportunity workplace and we would like to hear from all qualified applicants. If you have a disability or any special needs that we might need to accommodate, please just let us know.