GRAPHCORE

SuperTest – helping Graphcore accelerate the future of AI

Machine intelligence processor maker Graphcore (Bristol, UK) develops massively-parallel computing platforms that let innovators create next-generation AI products. As part of its compiler development tool chain, the company is using SuperTest to test and validate compilers for its unique Intelligence Processing Unit architecture.

There are several factors that compel the need for exhaustive compiler testing. Firstly, if you are building a compiler for a proprietary target architecture, where you can't draw on the wealth of experience already acquired by the developer community. Secondly, if your compiler must perform complex application-dependent optimizations, reordering or modifying code to make maximum use of target architecture features and instructions.

For Graphcore, a UK-based chip-maker that develops silicon platforms for accelerating Artificial Intelligence (AI) applications, both these conditions apply. The instruction set and memory architecture of its massively parallel Intelligence Processing Unit (IPU), which integrates over 1200 processor cores on a single piece of silicon, is uniquely designed to accelerate today's AI applications while also allowing engineers to create new categories of AI models and products. The back-end of its software stack, called Poplar®, employs multi-layered optimizations to distribute tasks between the cores and ensure they run at optimum speed and efficiency.

To promote exploration and innovation in the AI application space, the company is also committed to allowing application developers to program its IPU at the C++ level, which makes it extremely important to offer them a highly robust compiler that can handle anything they throw at it.

"We have developed a new silicon platform that is totally different in terms of its structure to anything that exists today and that can outperform existing solutions and open up new applications for AI. To open the door to these new applications, we are committed to providing users with a highly robust programming environment that gives them the freedom to explore using standard programming languages such as C++, which is why we chose SuperTest to test and validate our compilers," says Graphcore's Vice-President Software, Matt Fyles.

In terms of user confidence and convenience, there are also other reasons why the C++ compiler for Graphcore's IPU needs to be rigorously tested and validated. For example, in addition to allowing programming at the C++ level, the Poplar software development stack accepts input from highly abstracted deep-learning AI frameworks, which adds additional layers of translation between the abstracted model and the target hardware machine code.

"A major part of our solution is the software platform and how it maps applications onto the IPU," says Matt. "At the top level we have AI frameworks such as TensorFlow and PyTorch, which provide users with an abstraction layer that allows them to build neural networks without having to program the underlying hardware. Our Poplar software stack takes the output of these frameworks and compiles it down to individual hardware cores in the IPU via an intermediate language, in our case C++. When it comes to functionally testing the application, users need to have every confidence it's not the compiler that's causing problems."

"SuperTest's strengths lie in its breadth of coverage of language features and what it allows you to test, making it particularly useful when you are bringing up a compiler from scratch with a custom back-end. Its ability to do that and its general utility in that space is why we come back to using it again and again. It's almost an industry-standard," says Matt.

For more information on Graphcore's Intelligence Processing Unit (IPU) technology visit www.graphcore.ai

GRAPHCORE

Graphcore has created a completely new processor, the Intelligence Processing Unit (IPU), specifically designed for machine intelligence. The IPU's unique architecture means developers can run current machine learning models orders of magnitude faster. More importantly, it lets AI researchers undertake entirely new types of work, not possible using current technologies, to drive the next great breakthroughs in general machine intelligence.

We believe our IPU technology will become the worldwide standard for machine intelligence compute. The performance of Graphcore's IPU is going to be transformative across all industries and sectors whether you are a medical researcher, roboticist or building autonomous cars.

Solid Sands is based in Amsterdam, the Netherlands. Our mission is to put quality into C. We do that by improving the quality of C and C++ compilers, libraries and analysis tools, and their safe and secure use, with the best possible test and validation suite. With SuperTest, Solid Sands serves its customers to achieve the software quality level required by the ISO language and functional safety standards. With our history in compiler development, our knowledge of past, current and upcoming versions of the C and C++ standards, new analysis and optimizations techniques and new use cases, Solid Sands stays at the fore-front of tools testing and validation.

SOLID SANDS
from Amsterdam is the one-stop shop
for C and C++ compiler and library testing,
validation and safety services.

Postbus 7897 | 1008 AB AMSTERDAM | The Netherlands | www.solidsands.nl