EVEREST — DESIGN ENVIRONMENT FOR EXTREME-SCALE BIG DATA ANALYTICS ON HETEROGENEOUS PLATFORMS

The EVEREST project released its open-source System Development Kit

The Horizon 2020-funded project EVEREST – a consortium including ten partners from industry, industrial research laboratories, and academia across six European countries – announces the first public release of its open-source System Development Kit (EVEREST SDK). The EVEREST SDK integrates tools to simplify the deployment of applications on heterogeneous high-performance cloud infrastructure, especially for environments offering FPGA acceleration.

The European data market has been growing substantially and recently has seen an unprecedented boost through the advent of generative AI. Data processing is having a profound impact on industry, society, and the environment. The EVEREST project gathered application developers, compiler experts, runtime specialists, platform engineers, and security experts with the common goal of creating an integrated set of tools to co-develop and securely deploy high-performance big data analytics applications on heterogeneous FPGA-based platforms. While these systems can improve energy efficiency and performance, fixed-function accelerators introduce programmability and scalability issues.

"*The EVEREST SDK, publicly released on 18 May this year, provides a powerful and easy-to-use design environment to deploy big data applications on heterogeneous platforms using FPGA acceleration. Starting from the ability to express functional and non-functional requirements, the EVEREST SDK supports the end-users on their journey to accelerate compute kernels and move large amounts of data under defined security and privacy constraints*," explains the project coordinator Christoph Hagleitner from IBM Research Europe, Zurich.*

The development of the EVEREST SDK was driven by real-world use cases provided by the application-oriented partners in the EVEREST consortium. The use cases involve diverse applications, including improved forecasting of renewable energy production based on accelerated and AI-enhanced weather prediction, air-quality prediction around large industrial sites, and traffic prediction based on AI models combined with real-time data from smart-city environments. A preview of the newly released SDK was presented as a tutorial at the 32nd IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM) on 5 May, where interested users were able to gain hands-on experience with the development and deployment tools offered by the SDK. "*The EVEREST SDK demonstrated the ability to join forces among several top institutions, companies, and research centres in Europe. It provides a unique set of features to create and optimise complex FPGA architectures – with an emphasis on reusability and interoperability,*" comments the scientific coordinator Christian Pilato from Politecnico di Milano.*

The EVEREST project was launched on 1 October 2020. The project's initial focus was to define and elaborate the requirements for the design environment posed by the three use cases. Using constant feedback from application partners, the scientific community, and industry experts, the EVEREST consortium created the EVEREST SDK (https://github.com/everest-h2020/everest-sdk), which is now publicly available. The EVEREST project brought together partners from Switzerland, Italy, Germany, France, the Czech Republic and Slovakia, with IBM Research Europe GmbH from Switzerland acting as the project coordinator. For more information, see https://everest-h2020.eu.

Contacts
Christoph Hagleitner, IBM Zurich, hle@zurich.ibm.com
Christian Pilato, Politecnico di Milano, christian.pilato@polimi.it