

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

cFDevOps: Workshop on DevOps Support for Cloud FPGA platforms @ FPL 2021

## Climbing EVEREST: A design environment for extreme-scale big data analytics on heterogeneous platforms

#### **CHRISTIAN PILATO**

Politecnico di Milano, Scientific Coordinator

christian.pilato@polimi.it

#### **MICHELE PAOLINO**

Virtual Open System, Virtualization Leader

m.paolino@virtualopensystems.com

# **EVEREST: Using cloudFPGA for Big Data Analytics**

H2020 project funded under the call – "Big Data technologies and extreme-scale analytics" [Kick-off on Oct 1, 2020][http://www.everest-h2020.eu]

Big focus on FPGA acceleration in data centers and related issues → cloudFPGA

#### Key idea:

a coordinated action with the appropriate technology areas (e.g., AI, analytics, software engineering, HPC, Cloud technologies, IoT and edge/fog/ubiquitous computing) → FPGA acceleration in (federated) data centers

 system engineering/tools to contribute to the co-design of federated/distributed systems → EVEREST system development kit

adaptive memory management architectures for collecting, managing and exploiting virtualization data security

hardware acceleration

high-level synthesis

runtime management

domain-specific extensions



standardized interconnection methods

## **Application Concepts**

#### Three use cases provided by the application partners

- Looking for hardware acceleration (intense data computation) with efficient and secure data management (distributed data sources)
- Possibility of AI/ML-based decision making
- Combination of the tasks in different pipelines







# **EVEREST cloudFPGA: Key Features**

### Network-attached solution composed of:

- Interface logic already designed (cF Shell) to support system integration
  - TPC/UDP communication is managed transparently to the user logic
- User logic (ROLE) that can be easily designed and customized) with traditional HLS tools

### Application code running on host

- Low-level libraries for host-FPGA communication
- Possibility to use a cluster of FPGAs
- IDE for allocation and management of resources





# **Climbing EVEREST: Obstacles on the Road**

Programmability: Application developers have often limited hardware skills and limited knowledge of the target platform

- How to specify the application functionality to get the best results?
- How to design the hardware accelerator and the memory subsystem not only to optimize the performance but also to avoid bottleneck
- **Portability:** Designing a FPGA system is hard, but **designing an application for many systems** is even harder
- How to specify a platform-agnostic functionality?
- How to match such functionality with the actual hardware?
- How to deal with dynamic changes?





# **EVEREST SDK: System Development Kit**

**Coordinated design environment** composed of **four major phases**:

- **1.** Application specification (data ▷ application and requirements)
- **Architecture abstraction** (target system ▷ arch. description) 2.
- **Programming environment** (app+arch+reqs desc. ▷ hw/sw bin.) 3.
- **4.** Execution monitoring and management (hw/sw bin. ▷ execution)



EVEREST @ cFDevOps 2021 6

# **EVEREST Programming Environment**

(ML) frameworks

- **1. Compilation Environment:** analyzes application and creates all "variants" based on architecture abstraction and application/data requirements
  - Unified IR framework (MLIR)
  - Hardware acceleration and High-level synthesis (Bambu, Vivado/Vitis HLS)
  - Integration of non-functional properties with domain-specific extensions



Standard IR format and

exchange files

## **Hardware Compilation Flow**

Automated **DSL-to-bitstream generation** for accelerating selected application kernels with specialized memory architectures





8 EVEREST @ cFDevOps 2021

## **From DSL to Bitstream**





## **Creation of Parallel Architectures**









## **Preliminary Evaluation**

- Xilinx Zynq UltraScale+ MPSoC ZCU106 board
  - CFD simulation of 50,000 elements
- Preliminary comparison with embedded ARM





K. Friebel, et al. "From Domain-Specific Languages to Memory-Optimized Accelerators for Fluid Dynamics" HPCFPGA'21



# Next Step: Let's Put Memory First

We are building a **compilation flow** based on **LLVM MLIR** for **automatic specialization**:

- MLIR Input From DSL descriptions of the system functionality
- **Data Organization** Determine which data resides off chip (also based on user/compiler annotations)
- **Layout** Reorganize communication to exploit local memories (cache/PLM)
- Communication Configure prefetcher to hide transfer latency
- Local Partitioning Determine multi-bank PLM architecture (Mnemosyne<sup>1</sup>)
- **HLS** Generate computation part (interfacing with existing HLS tools, e.g., open-source Vitis HLS frontend)
- HDL Output Automated code generation and system-level integration based on the target platform



Domain-Specific Memory Templates" LATTE'21



# **EVEREST Runtime Environment**

- 2. Runtime Environment: implements the <u>selection of "variants"</u> and the <u>hardware configuration</u> based on the <u>system status</u>
  - Dynamic adaptation and autotuning (mARGOt)
  - Two-level runtime for (1) virtualization of hardware resources regardless their distribution and the low-level details of the platforms; (2) implement functional decisions (mARGOt, HyperLoom)



**Runtime API** 

Autotuning API

Hiding communication latency (e.g., prefetching)



Seamless execution when varying the system configuration (resources, nodes, data, etc.)



# **EVEREST Runtime Virtualization**

The **EVEREST runtime virtualization** is based on specific host and guest extensions:

- Host
  - The hypervisor of reference is KVM, and the Virtualized Infrastructure Manager is OpenStack
  - Libvirt will be extended to recognize harwdare accelerators and allocate them to the VMs. It will be used as driver for OpenStack Nova

#### Guests

- Will be extended with multi-node and monitoring/profiling applications that communicate with the host
- Guests will benefit from optimized communication mechanisms based on shared memory



EVEREST

## **EVEREST Runtime Virtualization**

### The EVEREST runtime **key virtualization techniques** will be **FPGA Virtualization** and **API remoting**:

**FPGA Virtualization**: we consider both network-based and host-attached approaches

- Network based: using cloudFPGA mentioned before
- Host attached: SR-IOV and mediated pass-through used to allow guests direct access to the FPGA partitions
  - Single Root Input-Output Virtualization (SR-IOV) enables the creation of multiple virtual functions
  - Each of them will be mapped to a VM (and a portion of the FPGA)





# **EVEREST Runtime Virtualization (ii)**

The EVEREST runtime **key virtualization techniques** will be **FPGA Virtualization** and **API remoting**:

## **API remoting**

- Virtualization at a library level will be implemented through API remoting
- API calls in the guests will be redirected to the host
- A zero-copy shared memory mechanism is in place between the host and guests OSes





## **EVEREST Runtime Virtualization at the edge**

The EVEREST runtime key virtualization focuses on **both edge and cloud sides**:

- Arm, RISC-V, and x86 supported for the cloud side
- Specific virtualization solutions based on system resource partitioning are developed to cope with security and safety requirements of edge systems
- VOSySmonitor and VOSySmonitoRV by Virtual Open Systems are based respectively on Arm TrustZone and RISC-V PMP zones





## **EVEREST Runtime functional decision**

#### The EVEREST FPGA systems include a **monitoring and decision infrastructure** for **dynamic autotuning** based on workload conditions



 Application variants (either software or hardware) are generated at design time (compilation and hardware synthesis), and selected at run time based on the actual available hardware resources



## Conclusions

**EVEREST** is a new H2020 project that aims at **simplyfing the use of FPGA** for the acceleration of **Big Data applications** 

- Data-centric approach focusing on domain-specific extensions, high-level synthesis, and dynamic adaptivity
- Three application use cases: weather-based renewable energy prediction, airquality monitoring of industrial sites, and intelligent traffic management
- **CloudFPGA** will be used as primary target
- Looking for **interoperability** with existing solutions
- EVEREST SDK will be released as open-source to the community

The main goal is to simplify the description of complex Big Data applications and improve the programmability of distributed FPGA-based systems







This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957269