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

DATE22 Workshop

Climbing EVEREST: dEsign enVironmEnt foR Extreme-Scale big data analyTics on heterogeneous platforms

GIANLUCA PALERMO

Politecnico di Milano

EVEREST: Big Data Processing on FPGA

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 associated issues

Key idea:

- a coordinated action with the appropriate technology areas (e.g., AI, analytics, software engineering, HPC, Cloud technologies) ➔ **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

domain-specific extensions

data security

architectures for collecting, managing and exploiting data

hardware acceleration

high-level synthesis

runtime management

virtualization

autotuning and dunamic adaptation

EVEREST

EVEREST Partners

IBM Reseach Lab, Zurich (Switzerland)
Project coordination, prototype of the target system

Politecnico di Milano (Italy)
Scientific management, high-level synthesis, flexible memory managers, dynamic autotuning

Università della Svizzera italiana (Switzerland)
Data security requirements and protection techniques

TU Dresden (Germany)
Domain-specific extensions, code optimizations and variants generation

Virtual Open Systems (France)
Virtualization techniques, runtime extensions to manage heterogeneous resources

IT4Innovations (Czech Republic)
Exploitation leaders, HPC infrastructure, workflow libraries

Centro Internazionale di Monitoraggio Ambientale (Italy)
Weather prediction models

Duferco Energia (Italy)
Application for prediction of renewable energies

Numtech (France)
Application for monitoring the air quality of industrial sites

Sygic A/S (Slovakia)
Application for intelligent transportation in smart cities

Data Driven World

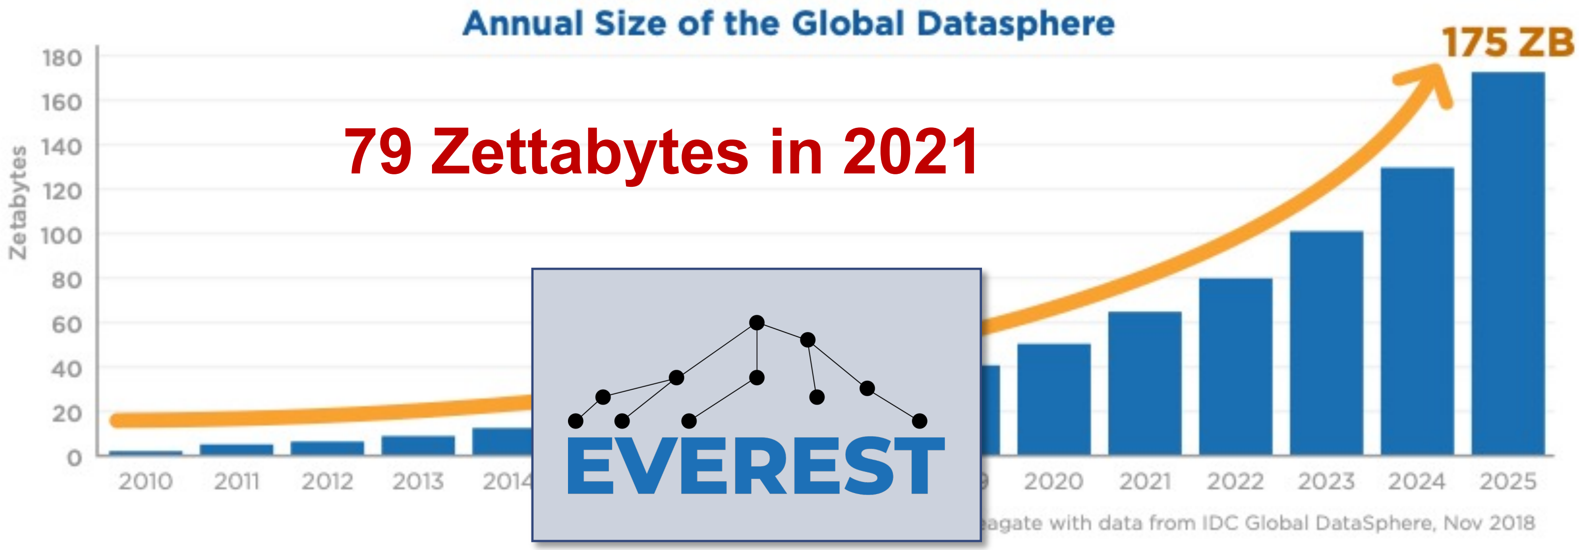

DARK DATA

NOT STORABLE DATA

Outline

- Everest Use Cases
- Everest Target Architectures
- Everest System Development Kit

Everest Use Cases: 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 tasks in different pipelines (creation of "scientific services")

Traffic modelling for intelligent transportation

Air-quality monitoring in industrial sites

Weather-based prediction of renewable energy production

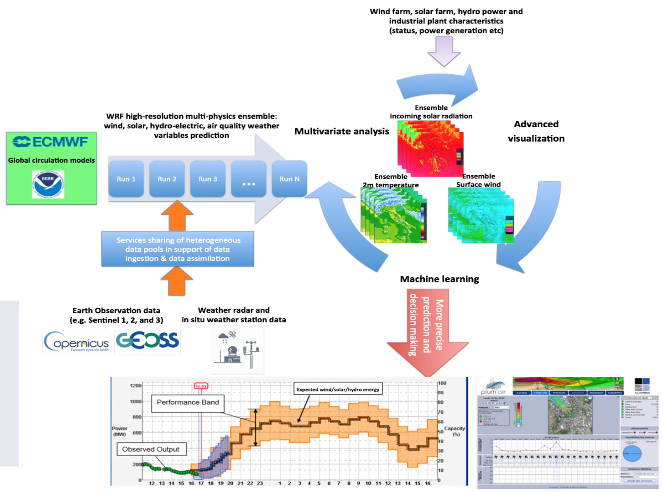

EVEREST

The WRF Model

- First step of two use cases…

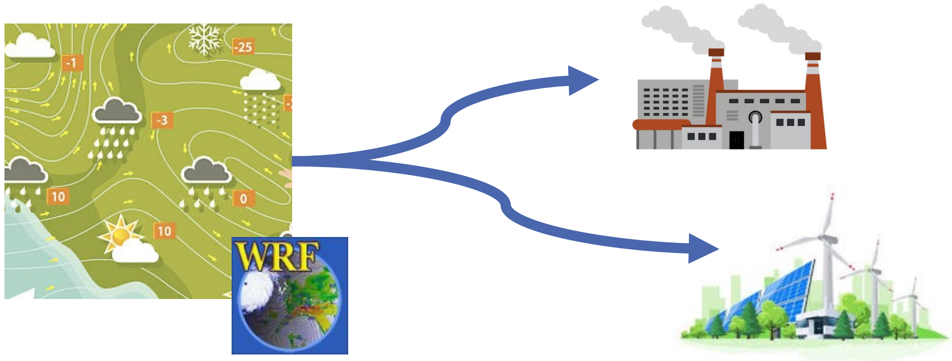

- WRF is an collaborative open source model supported principally by the US National Center for Atmospheric Research (NCAR), the US National Oceanic and Atmospheric Administration and the US National Center for Environmental Prediction – NCEP)

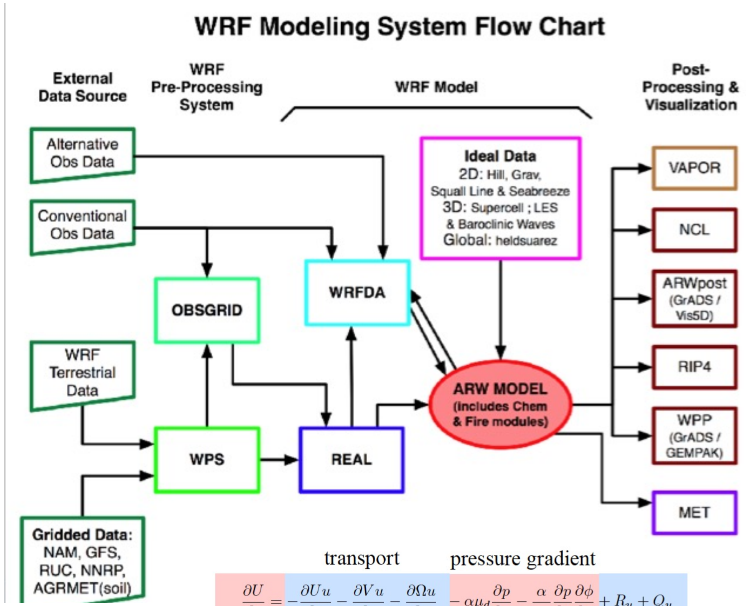

transport — pressure gradient

$$\frac{\partial U}{\partial t} = -\frac{\partial U u}{\partial x} - \frac{\partial V u}{\partial y} - \frac{\partial \Omega u}{\partial \eta} - \alpha\mu_d \frac{\partial p}{\partial x} - \frac{\alpha}{\alpha_d}\frac{\partial p}{\partial \eta}\frac{\partial \phi}{\partial x} + R_u + Q_u$$

$$\frac{\partial V}{\partial t} = -\frac{\partial U v}{\partial x} - \frac{\partial V v}{\partial y} - \frac{\partial \Omega v}{\partial \eta} - \alpha\mu_d \frac{\partial p}{\partial y} - \frac{\alpha}{\alpha_d}\frac{\partial p}{\partial \eta}\frac{\partial \phi}{\partial y} + R_v + Q_v$$

$$\frac{\partial W}{\partial t} = -\frac{\partial U w}{\partial x} - \frac{\partial V w}{\partial y} - \frac{\partial \Omega w}{\partial \eta} - g\left(\mu_d - \frac{\alpha}{\alpha_d}\frac{\partial p}{\partial \eta}\right) + R_w + Q_w$$

$$\frac{\partial \mu_d}{\partial t} = -\frac{\partial U}{\partial x} - \frac{\partial V}{\partial y} - \frac{\partial \Omega}{\partial \eta}$$

$$\frac{\partial \Theta}{\partial t} = -\frac{\partial U \theta}{\partial x} - \frac{\partial V \theta}{\partial y} - \frac{\partial \Omega \theta}{\partial \eta} + R_\theta + Q_\theta$$

$$\frac{\partial \mu_d q_j}{\partial t} = -\frac{\partial U q_j}{\partial x} - \frac{\partial V q_j}{\partial y} - \frac{\partial \Omega q_j}{\partial \eta} + R_{q_j} + Q_{q_j}$$

$$\frac{\partial \phi}{\partial t} = -u\frac{\partial \phi}{\partial x} - v\frac{\partial \phi}{\partial y} - \omega\frac{\partial \phi}{\partial \eta} + gw$$

↑ numerical filters, physics, projection terms ($+R_u+Q_u$, $+R_v+Q_v$, $+R_w+Q_w$, $+R_\theta+Q_\theta$, $+R_{q_j}+Q_{q_j}$)

← geopotential eqn term ($+gw$)

Diagnostic relations: $\frac{\partial \phi}{\partial \eta} = -\alpha_d \mu_d$, $p = \left(\frac{R_d \Theta_m}{p_o \mu_d \alpha_d}\right)^{\gamma}$, $\Theta_m = \Theta\left(1 + \frac{R_v}{R_d} q_v\right)$

The WRF Model

- First step of two use cases…

CHALLANGES

- **Automatic deployment of several instances of the model in HPC resources with multiple boundaries conditions**
- **Efficient High Resolution output evaluating the usage of FPGA resources**

- WRF is an collab… supported principal… for Atmospheric Research (NCAR), the US National Oceanic and Atmospheric Administration and the US National Center for Environmental Prediction – NCEP)

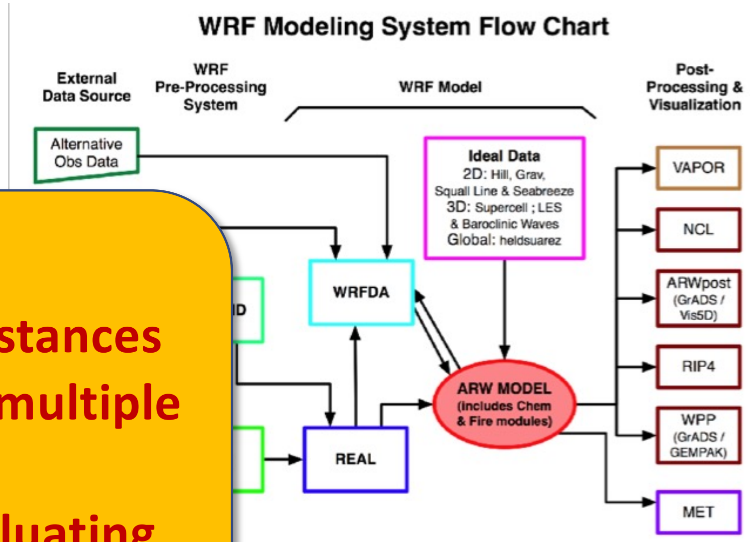

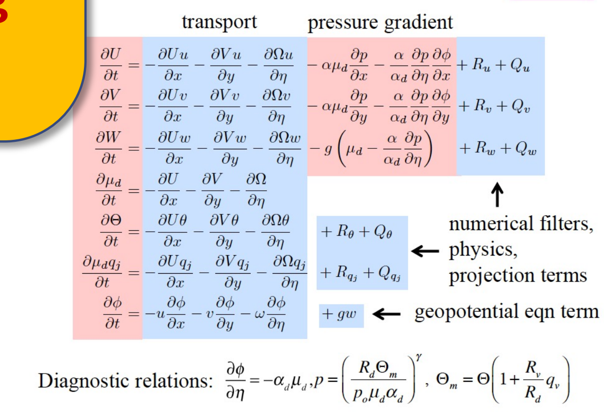

EVERLAST

Air-quality use case: General context

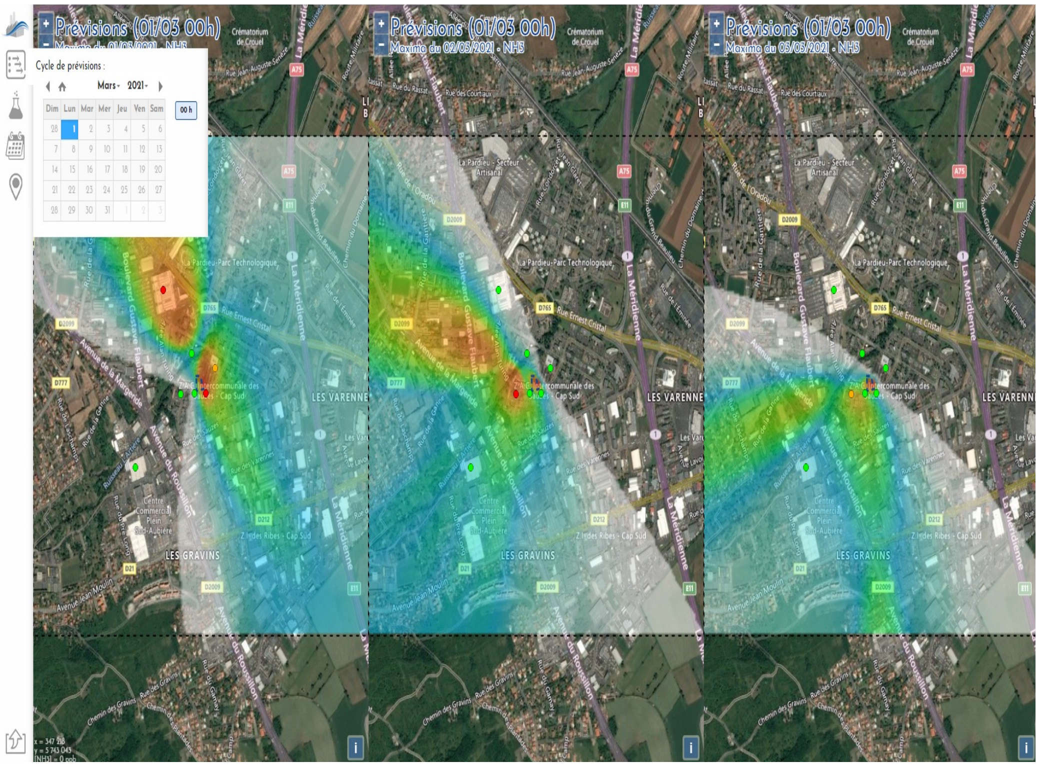

- **Platform to predict the air quality of industrial sites**
- If predicted impacts exceeds some concentration levels, the industrial site may adapt its production, activate some emission reduction process, …
  - Action = Financial cost
  - Objectives :
    - Minimize false predicted peaks
    - Optimize real predicted peaks
- One key part is the meteorological forecast part
- Site specific predictions

Air-quality use case: Workflows and Challenges

NUMTECH

1. WRF Deterministic weather forecast

Global observation

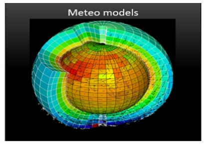

WRF execution

Improve speed to produce forecast

2. Ensemble weather aggregation

N x deterministic weather forecast

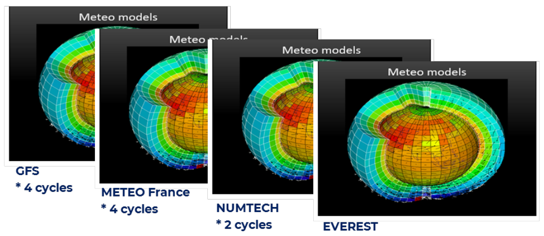

\+

Local weather observation on-site

One aggregated weather forecast forced by observation

Improve quality of local weather forecast

3. Air-quality dispersion forecast

Landuse, Topography

Industrial site data

Emission forecast

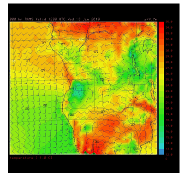

Local weather forecast

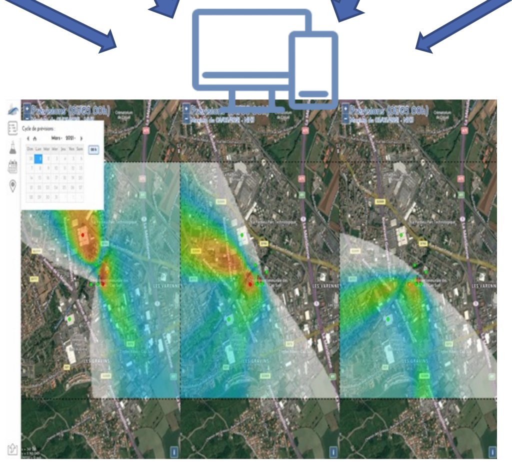

Air-quality forecast

Improve speed to produce air-quality forecast and its quality

EVEREST

Renewable Energy use case: Context and Challenge

- **Effort to improve the Wind Power Forecast accuracy**
- Different challenges due to intermittency of the wind power generation:
  - Transmission System Operator (in Italy TERNA) to ensure the balance of grid (very short term horizon: 1s to 1h)
  - Traders to forecast the power to sell on energy market, intraday or day ahead (short term horizon: 1h to 24h)
  - Wind farm owners to schedule their maintenance programs (long term horizon)

Wind Power Global Capacity and Annual Additions, 2008-2018

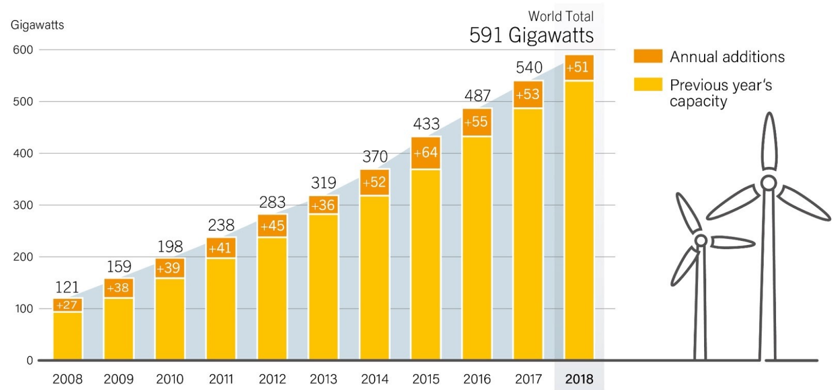

EVEREST

Renewable Energy use case: Context and Challenge

- Different challenges due to intermittency of the wind power generation:
  - Transmissio[…]e of grid (very short t[…]
  - Traders to f[…]y ahead (short term […]
  - Wind farm o[…] horizon)
  - → **great eff[…]**

CHALLENGES:

1. **Integration of the the data coming from sensors available on wind turbine for improving weather forecast model**
2. **Extend the wind speed estimation at 80m instead of the 10m default in WRF to improve prediction capabilities**
3. **Adopt AI models to learn and predict power generated by the specific site/turbine given historical data**
4. **… making use of Heterogeneous Resources when needed**

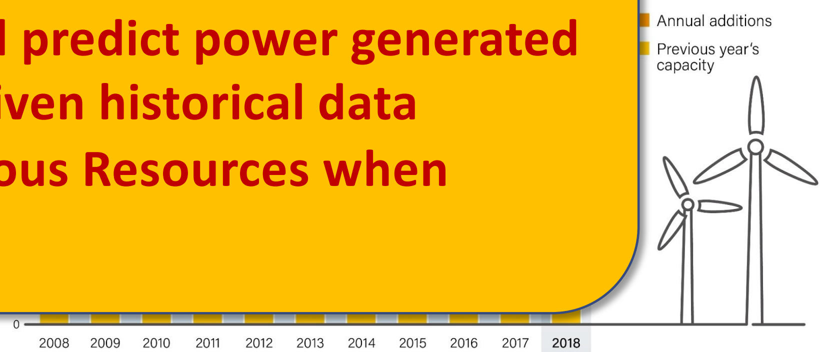

EVERIST

Advance Traffic Modeling for Smart Cities use case

- Mobility platform for supporting cities with advanced traffic modelling

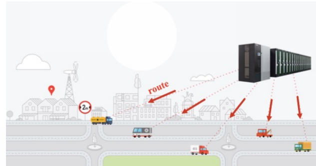

- **Data sources**
  - Historical and real-time **Floating Car Data** (FCD)
    - e.g. GPS position, timestamp and speed
  - **Origin-destination matrix** (ODM) defining city
  - **Road network graph** including road restrictions;
  - Historical **weather data** (temperature, precipitation)
- **Traffic services**
  - **What-if analysis** for given scenarios, e.g. road closure;
  - **Intelligent routing** for large amount of vehicles
  - **Traffic prediction** for major road elements of cities

Advance Traffic Modeling for Smart Cities use case

- Mobility platform for supporting cities with advanced traffic modelling
- **Data source**
  - Historical a
    - e.g. GPS
  - **Origin-des**
  - **Road netw**
  - Historical **w**
- **Traffic servi**
  - **What-if ana**
  - **Intelligent routing** for large amount of vehicles
  - **Traffic prediction** for major road elements of cities

CHALLENGES:

1. **Boost the FCD data collection and processing to compute 3D traffic models**
2. **Develop and efficiently deploy a traffic simulator on a multinode architecture to predict traffic behaviour**
3. **Adopt AI techniques to learn the traffic patterns resulting into a traffic prediction service.**
4. **… making use of Heterogeneous Resources**

EVEREST

EVERST Target System: Brief Overview

Datacenters and **Supercomputer** infrastructures

Network-attached and **PCIe-attached** FPGSd architectures:

- Off-the-shelf FPGA devices
- User logic can be easily designed and customized with traditional HLS tools

FPGA as a Co-Processor

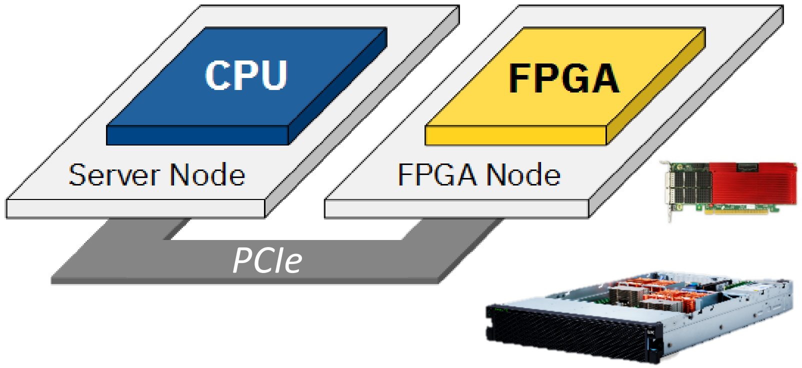

FPGA as a Peer-Processor

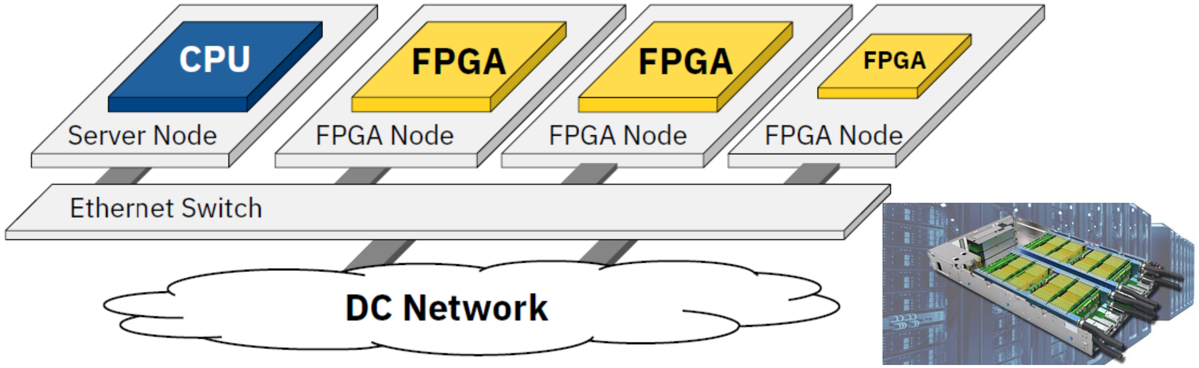

EVEREST

cloudFPGA

FPGA as a 1st class citizen within a DataCenter (DC):

- FPGA disaggregated from the server nodes
- FPGAs connected directly to the DC network for its access and to communicate with CPUs and other FPGAs
- FPGAs are den densely packed into DC chassis and racks and distributed across the DC

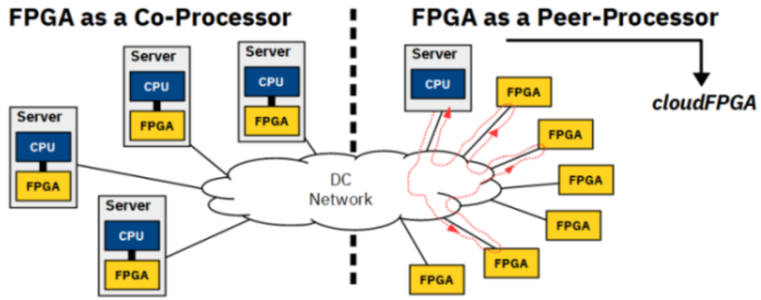

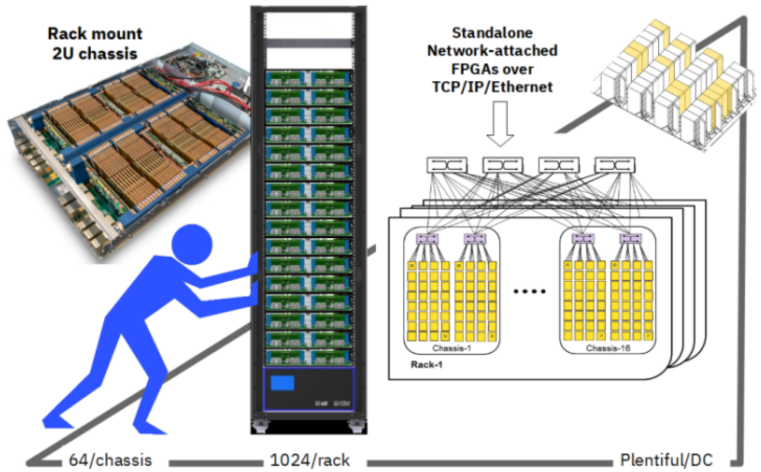

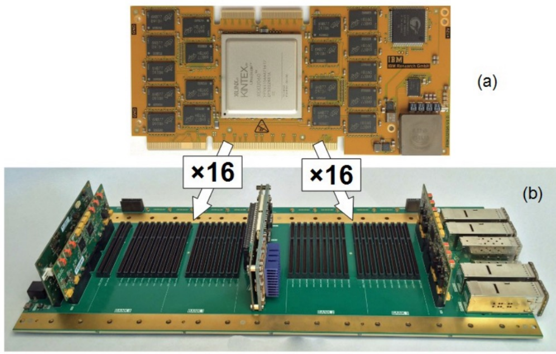

Figure 1: (a) The disaggregated FPGA and (b) the carrier board.

Xilinx Kintex UltraScale XCKU060 FPGA with 2x8GB of DDR4 memory

F. Abel, et al. "An FPGA Platform for Hyperscalers", HOTI

B. Ringlein, et al. “A Case for Function-as-a-Service with Disaggregated FPGAs,” CLOUD

EVEREST

cloudFPGA: Key Features

IBM

- 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
  - FPGA accessible through the network
  - Low-level libraries for host-FPGA communication
- Possibility to instanciate a cluster of FPGAs
- IDE and custom tools for allocation and management of resources

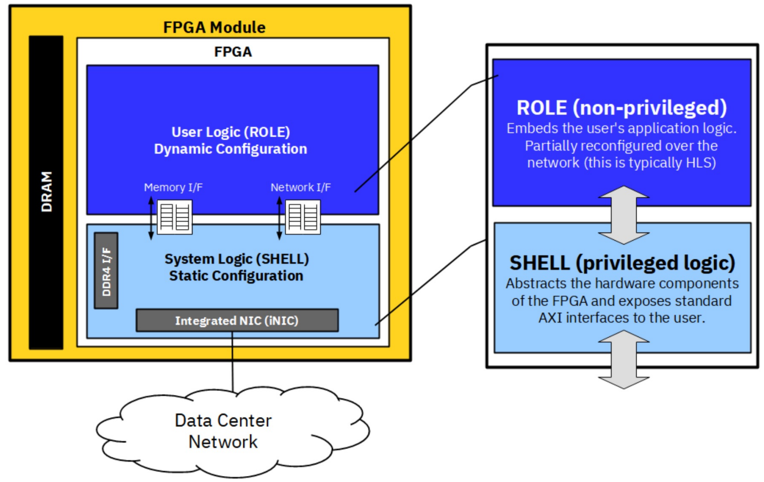

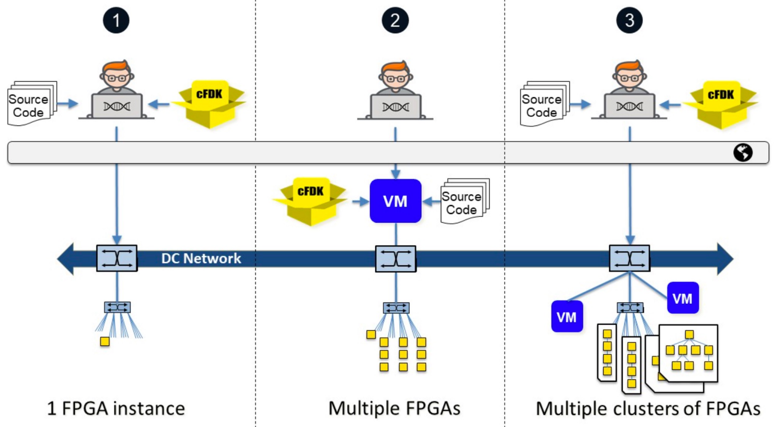

https://github.com/cloudFPGA

EVEREST

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 respect resource constraints

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?

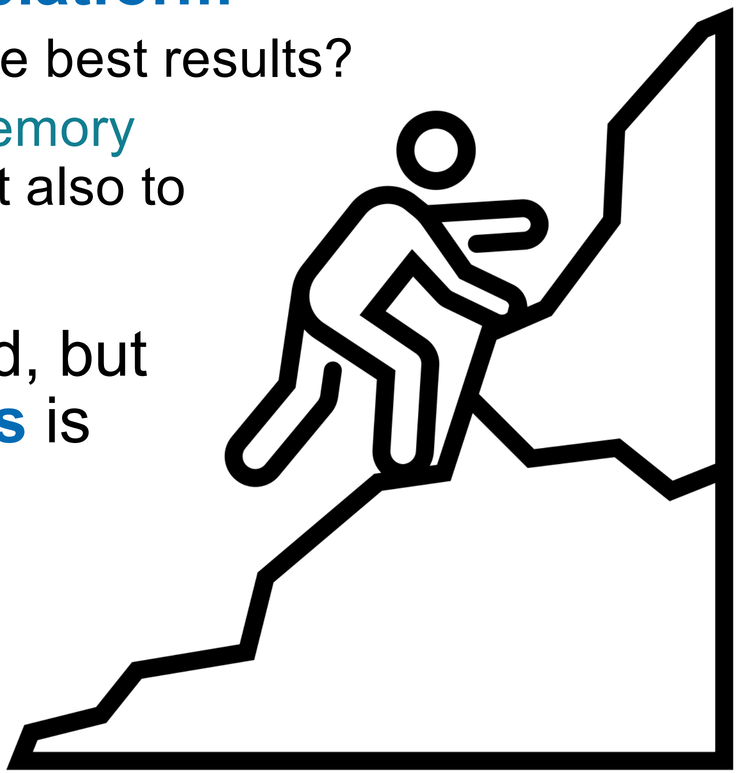

EVEREST SDK: SYSTEM Development Kit

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

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

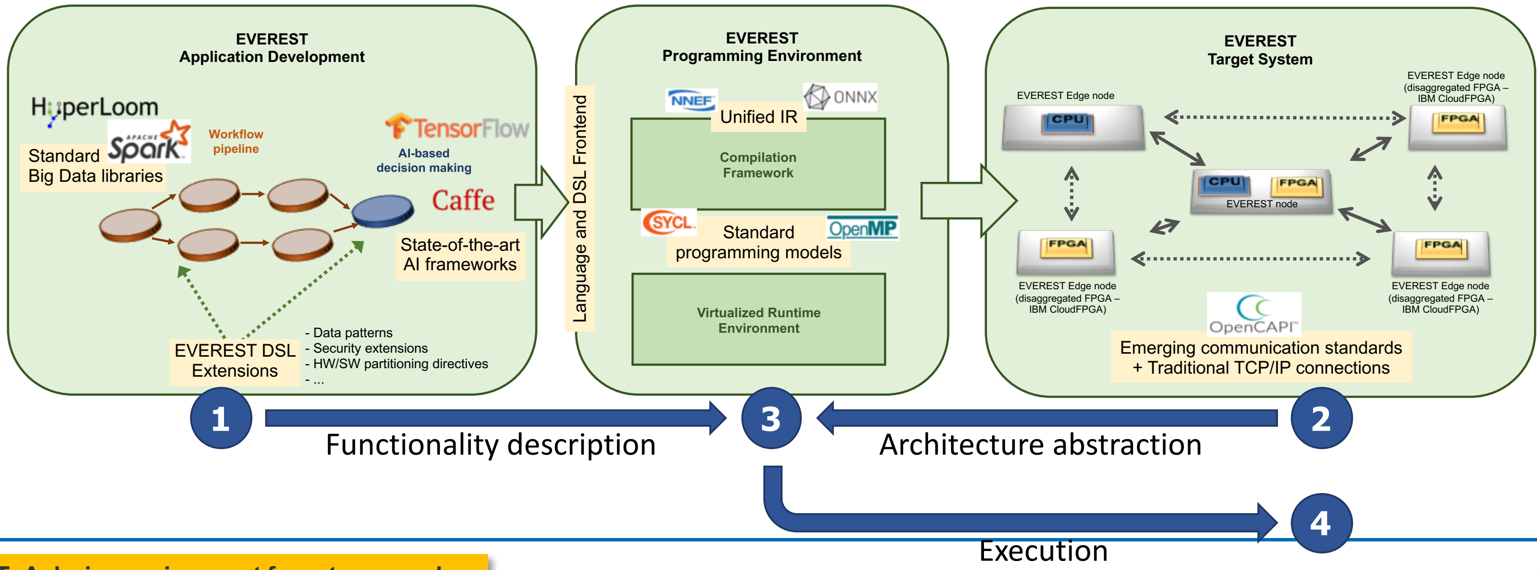

"EVEREST: A design environment for extreme-scale big data analytics on heterogeneous platforms" DATE

EVEREST

EVEREST Programming Environment

Compilation Environment: analyzes application and creates all "variants" based on *architecture abstraction* and *application/data requirements*

- Integration of **functional description** and **non-functional properties** with DSL extensions
- **Unified IR framework** (**MLIR**) to support different inputs
  - Support for custom DSL and existing ML frameworks
- Interoperability of **HLS tools** (**Bambu, VivadoVitis HLS**)
- "Intelligent" **memory managers** to coordinate data transfers (**Mnemosyne**)
- **Multitarget system integration** for different FPGAs (**Olympus**, **cFDK**)

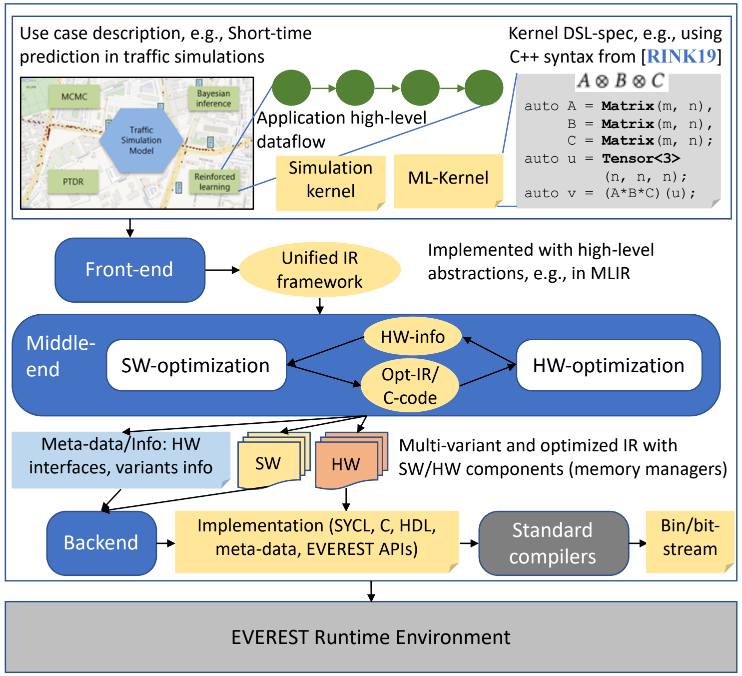

Creation of bitstreams and binaries to be deployed onto the target platform hiding its complexity

F. Ferrandi et al., "Bambu: an Open-Source Research Framework for the High-Level Synthesis of Complex Applications" DAC

C. Pilato et al. "System-Level Optimization of Accelerator Local Memory for Heterogeneous Systems-on-Chip" TCAD

C. Friebel et al. "From Domain-Specific Languages to Memory-Optimized Accelerators for Fluid Dynamics." CLUSTER

EVEREST Runtime Environment

Runtime Environment: implements the dynamic *selection of "variants"* and the *hardware configuration* based on the *system status*

- **Two-level runtime** to support

  (1) **virtualization** of hardware resources regardless their distribution and the low-level details of the platforms;

  (2) Application-level **functional decisions**

    (a) mARGOt for **dynamic adaptation and autotuning**

    (b) HyperLoom/Hyperqueue for **multinode management** (scheduling and resource allocation)

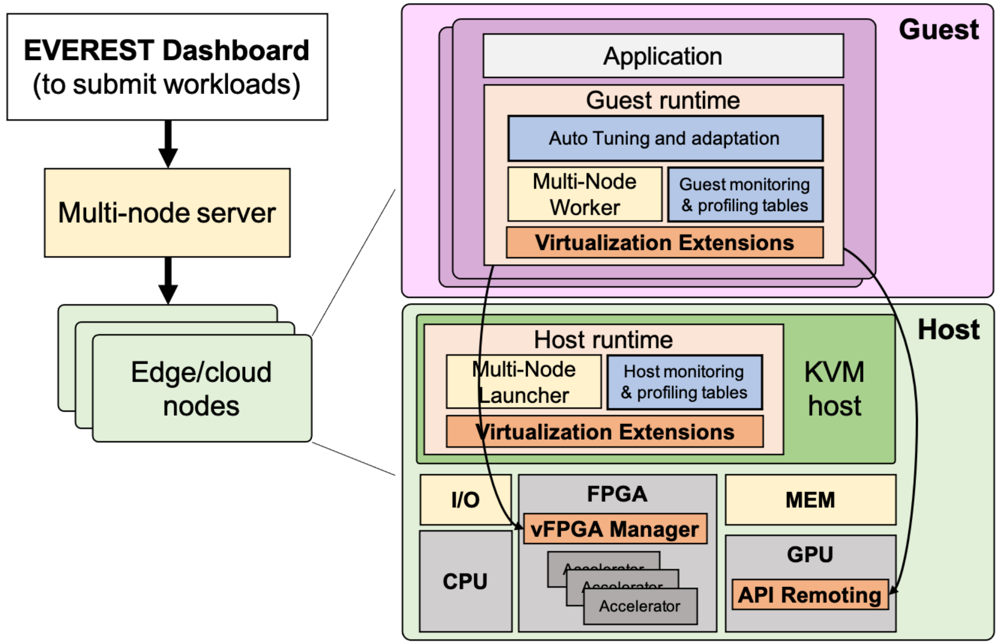

Automatic deployment and configuration on the (cloud) FPGA resources

D. Gadioli, et al. "mARGOt: A Dynamic Autotuning Framework for Self-Aware Approximate Computing" TCOM

V. Cima et al."HyperLoom: A platform for defining and executing scientific pipelines in distributed environments." PARMA-DITAM

EVEREST

Conclusions

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

- **Three application use cases**: *weather-based renewable energy prediction*, *air-quality monitoring of industrial sites*, and *intelligent traffic management*
- **Data-centric approach** focusing on the combination of **domain-specific and data-oriented extensions**, **high-level synthesis**, and **dynamic adaptivity**
- **CloudFPGA** and **Alveo-based clusters** will be used as acceleration platforms
- 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**

!!! Many piecies available, the global picture is not composed yet !!!
Stay tuned on http://www.everest-h2020.eu

IBM

POLITECNICO
MILANO 1863

cima
RESEARCH
FOUNDATION

IT4I

NUMTECH

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