

This issue's round-up of news and results from EU-funded projects sees an update on the VINEYARD project, new data centre architectures and I/O software stack of extreme-scale systems. We also hear about a novel toolchain to assess GPU performance and power in a vendor-agnostic way.

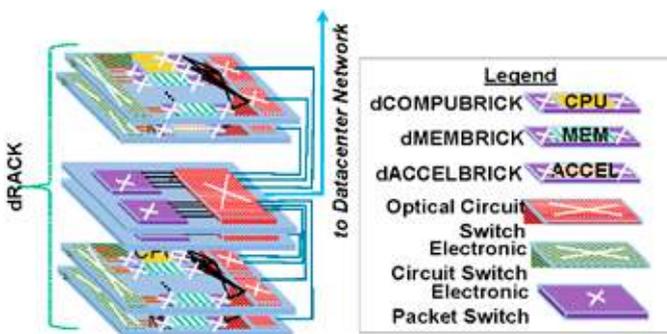
Innovation Europe

DISAGGREGATED RECURSIVE DATACENTRE-IN-A-BOX



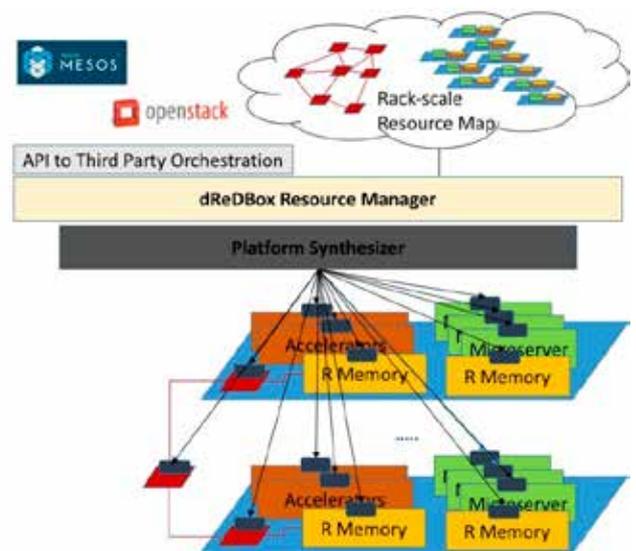
dReDBox, a three-year project funded by the EC H2020 programme aspires to change the way in which data centres are built, proposing the deployment of pooled, disaggregated resources rather than monolithic and tightly integrated components. The goal is to improve utilization, scalability, reliability, and therefore power efficiency in cloud and edge data centres.

Contemporary cloud data centres use individual server units for processing, memory, acceleration, and storage. This arrangement incurs significant underutilization of resources due to the inherent inability of the system to closely match the amount of available resources within a server (or set of servers) to the variable user IT requirements. dReDBox aspires to remedy these limitations by moving to a pooled computing approach that enables the flexible pooling of disaggregated resources to match the runtime application requirements.



1: The dRedBox architecture

The project approach (Figure 1) is to disaggregate the traditional tray and overall rack architecture by introducing a set of bricks per tray, i.e. dedicated modules of specific functionality (compute, memory, acceleration, etc.). The main processing block is the Compute Brick, which can host multiple Virtual Machines (VMs). Memory Bricks collectively serve as a large pool of memory that can be (dynamically) redistributed to VMs. Accelerator Bricks host dedicated hardware modules to boost application performance. Intra-tray communication between bricks is realized via electrical circuit crossbar switches, while an optical network will provide inter-tray connectivity. Therefore, inter-brick communication will be implemented through multiple communication networks of different speeds depending on the brick locations.



2: Orchestration tools high-level architecture

From the software management perspective (Figure 2), the Linux Kernel Virtual Machine (KVM) hypervisor will allow authorized tenants to reserve system resources to efficiently execute applications on hardware. dReDBox revisits the memory management

architecture at the VMM-level, implementing physical/logical attachments of remote memory to VMs.

Collectively, dRedBox will deliver a prototype of a fully-fledged, vertically-integrated data centre-in-a-box to demonstrate dRedBox in three use-cases: security, network analytics, and telecoms. Contact: katrinisk@ie.ibm.com

NAME: Disaggregated Recursive Datacentre-in-a-Box (dRedBox)

START/END DATE: 01/01/2016 – 31/12/2018

KEYWORDS: Advanced computing, Embedded computing, Computing for servers, data centres

PARTNERS: IBM Research Ltd (Ireland), University of Thessaly (Greece), University of Bristol (UK), Barcelona Supercomputing Center (Spain), Sintecs BV (Netherlands), Foundation for Research and Technology - Hellas (Greece), Telefonica Investigation y Desarrollo SA (Spain), Kinesense Ltd (Ireland), Naudit High Performance, Computing and Networking SL (Spain), Virtual Open Systems SARL (France), Polatis Ltd (UK)

BUDGET: €6.45M

WEBSITE: www.dredbox.eu

The dRedBox project has received funding from the European Union's Horizon 2020 Programme under grant agreement no. 687632.

VERSATILE INTEGRATED ACCELERATOR-BASED HETEROGENEOUS DATA CENTRES



VINEYARD's goal is to develop the technology and the ecosystem that will enable the efficient and seamless integration of the hardware acceleration into data centre applications. The deployment of energy-efficient hardware accelerators will permit the significant improvement in the performance of cloud computing applications and the reduction of energy consumption in data centres.

VINEYARD is developing an integrated framework for energy-efficient data centres based on programmable hardware accelerators. It is working towards a high-level programming framework that allows end-users to seamlessly utilize these accelerators in heterogeneous computing systems by using typical data centre cluster frameworks (i.e. Spark). The VINEYARD framework and the required system software hides the programming complexity of the heterogeneous computing system based on hardware accelerators. This programming framework also allows the hardware accelerators to be swapped in and out of the heterogeneous infrastructure so as to offer both efficient energy use and flexibility. To allow the efficient utilization of the accelerators by several applications, a novel VM appliance model for provisioning of data to shared accelerators has been developed. The enhanced VINEYARD middleware augments the functionality of the

resource manager, by enabling more informed allocation of tasks to accelerators.

VINEYARD is also developing two types of novel energy-efficient servers integrating two kinds of hardware accelerator: programmable dataflow-based accelerators and FPGA-based accelerators. The servers coupled with dataflow-based accelerators are suitable for cloud computing applications that can be represented in dataflow graphs while the latter will be used for accelerating applications that need close communication between the processor and the hardware accelerators.

VINEYARD also fosters the establishment of an ecosystem that will empower open innovation based on hardware accelerators as data centre plugins, thereby facilitating innovative enterprises (large industries, SMEs, and creative startups) to develop novel solutions using VINEYARD's leading edge developments. The ecosystem will bring together existing communities from all relevant stakeholders including providers of hardware intellectual-property (IP) technologies, data centre developers, data centre operators and more. This ecosystem will allow the promotion of open pluggable custom hardware accelerators (i.e. a hardware accelerator Application Store) that can be used in data centres in the same way that software libraries are currently being utilized.

VINEYARD plans to demonstrate the advantages of its approach in three real use-cases:

- a) a bioinformatics application for high-performance brain simulations
- b) two critical financial applications
- c) a big-data analysis application

The VINEYARD framework will also be integrated with widely-used cluster computing frameworks, like Spark, that will allow the utilization of the accelerators for typical cloud computing applications like machine learning and graph computations.

