A Software Stack for Operation Beyond Conservative Scaling Boundaries

Panos Koutsovasilis* | Konstantinos Tovletoglou† | Manolis Maroudas* | Christos Kalogirou* | Charalambos Chaliost | Christos Antonopoulos*
Spyros Lalis* | Nikolaos Bellas* | Georgios Karakonstantis† | Hans Vandierendonck† | Dimitrios Nikolopoulos†

* University of Thessaly, Greece | † Queen's University of Belfast, United Kingdom

Horizon 2020
European Union funding
for Research & Innovation

uniserver
Universal Micro-Server EcoSystem

1. Introduction

UniServer aims to prepare the next generation of micro-servers that can handle the upcoming immerse data flood of Internet of Things (IoT) by combining ***Edge/Fog Computing*** and ***intrinsic Heterogeneity*** by shaving the rather conservative manufacturer operation margins to a *just-right* point.

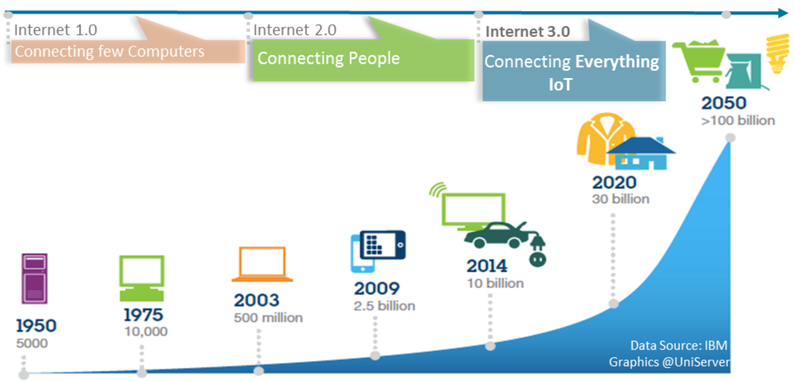

Evolution of internet of things (Data source: IBM)

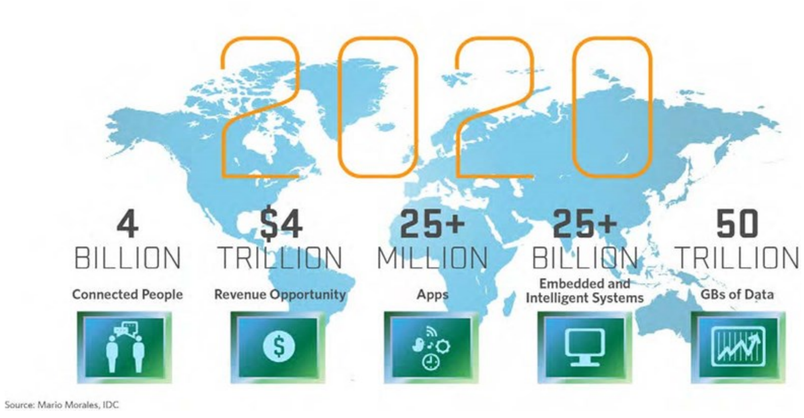

50 Billion connected devices by 2020

Advent of ***Edge/Fog Computing***

- Extremely Energy Efficient
- Low deployment cost
- Small form factor
- Error Resilient—Robust
- Secure

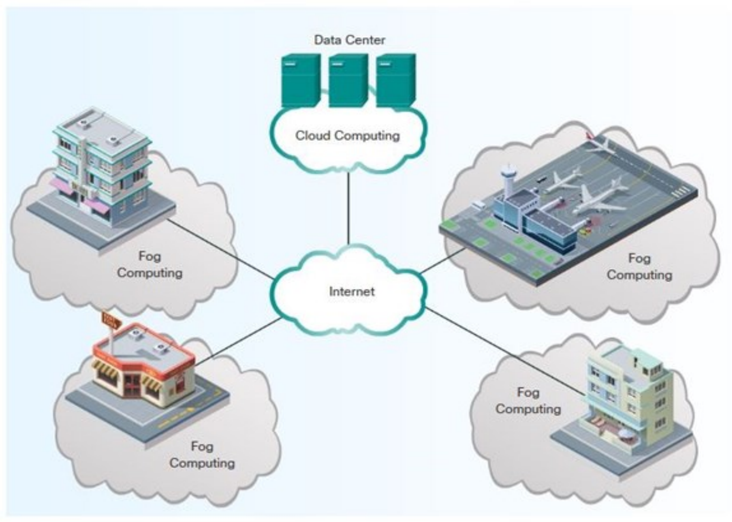

Fog computing

Exceeding the energy and performance scaling boundaries by exploiting ***intrinsic Heterogeneity*** can really help overcome

- Power Constrains
- Density

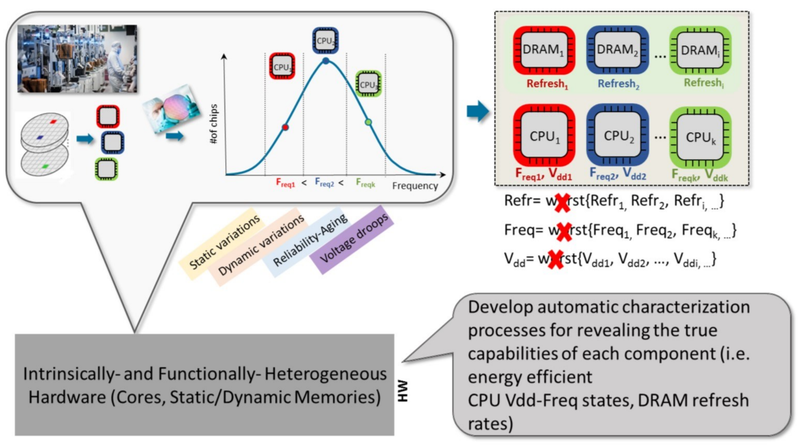

2. Uniserver Ecosystem

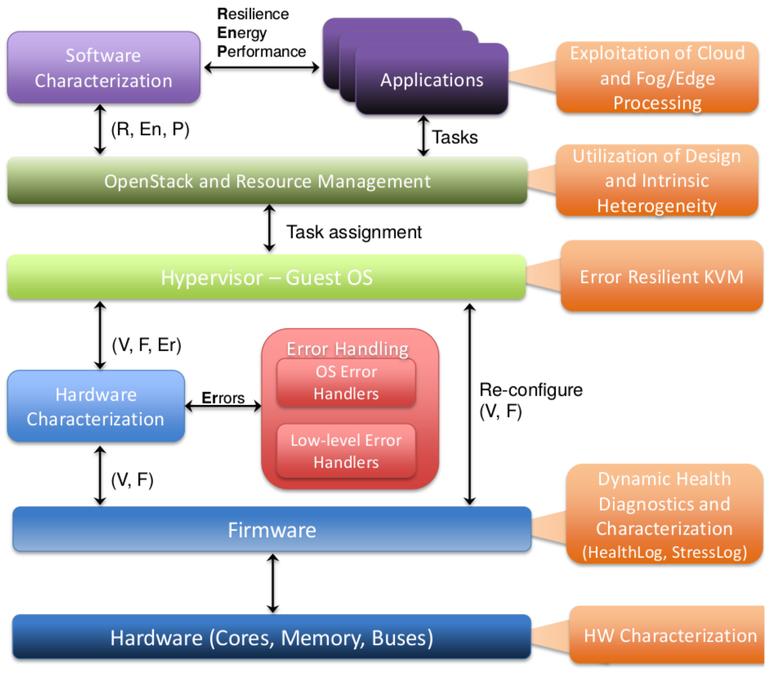

Uniserver ecosystem

The proposed cross-layer approach plans to contribute to the following layers:

- ***at the circuit, micro-architecture and architecture layer*** by revealing the possible extended operating points .
- ***at the firmware layer, with low-level handlers*** by monitoring and controlling the operating status of the underlying hardware components.
- ***at the software layer*** by enabling virtualization and programmability ensuring full utilization of the operational margins, and featuring extensions for resource management.

3. System Software Stack

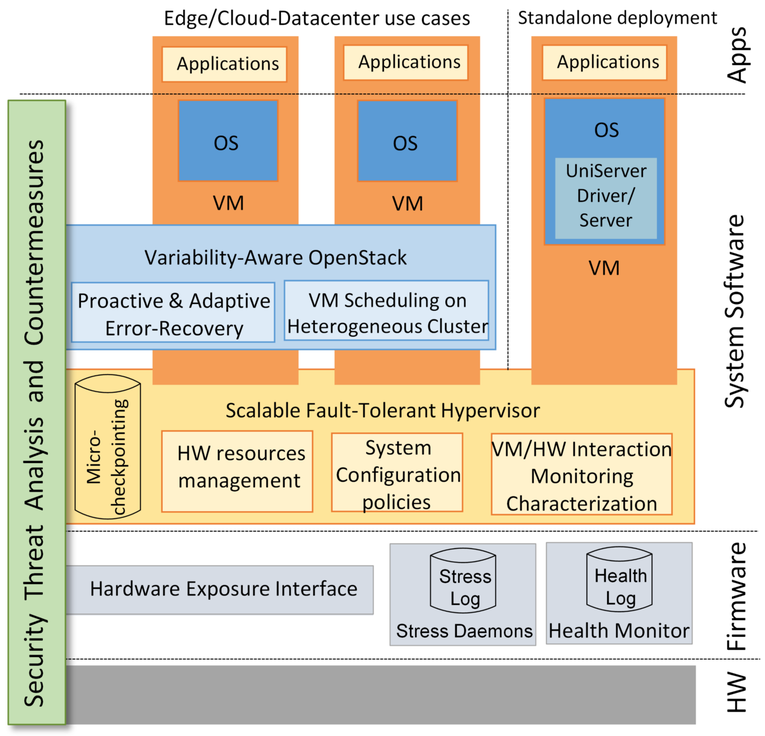

Uniserver software stack

Extended operating points may dynamically change depending on the workload, variations of environmental conditions and thus the system should be able to decide the right configuration parameters

- ***Hypervisor*** will enable the creation of an appropriate execution environment for Virtual Machines by manipulating the power/performance/reliability trade-offs in an educated and safe manner. Also ***Hypervisor*** will mask and manage transparently the errors from higher software layers.
- ***Openstack*** will include support for monitoring VMs and determining their dynamically changing charactering and their resource utilization. This information will be used to form new scheduling policies, as well as to assess the susceptibility of VMs to catastrophic errors.

4. Initial Experimentation

The added safety voltage margins increase energy consumption and force operation at higher voltage or lower frequency. They may also result in lower yield or field returns if a part operates at higher power than its specification allows. The scale of pessimism is also observed on recently measured ARM processors.

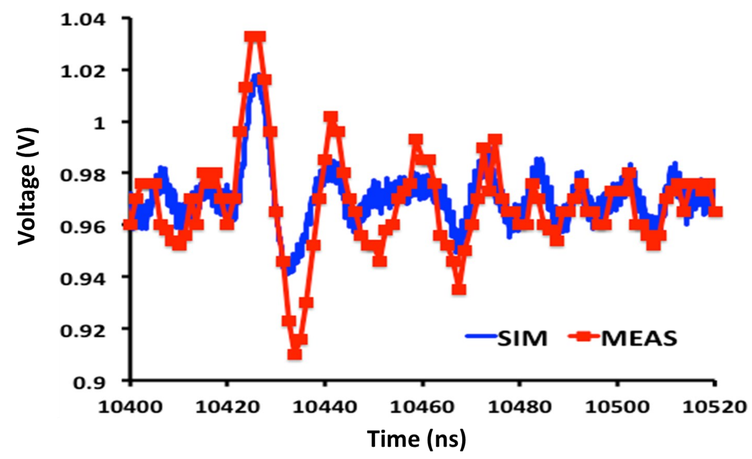

Branch Misprediction causes voltage overshoot

Measured droops reveal 15% voltage droops for a dual-core ARM A57 cluster operating at 1.1GHz.

Each DRAM row needs to be periodically ***refreshed*** to restore charge but this introduces ***energy*** and ***performance*** overhead.

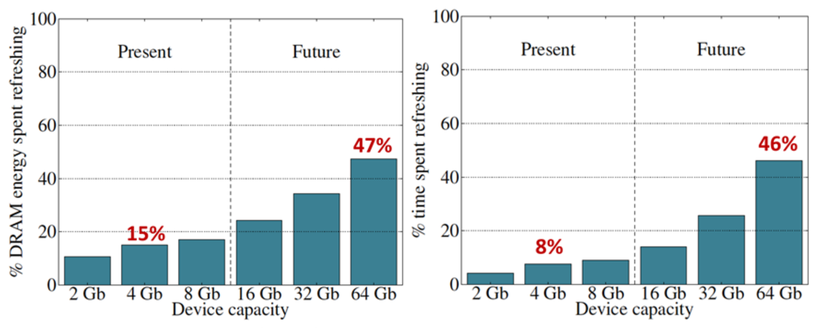

energy overhead ***performance*** overhead

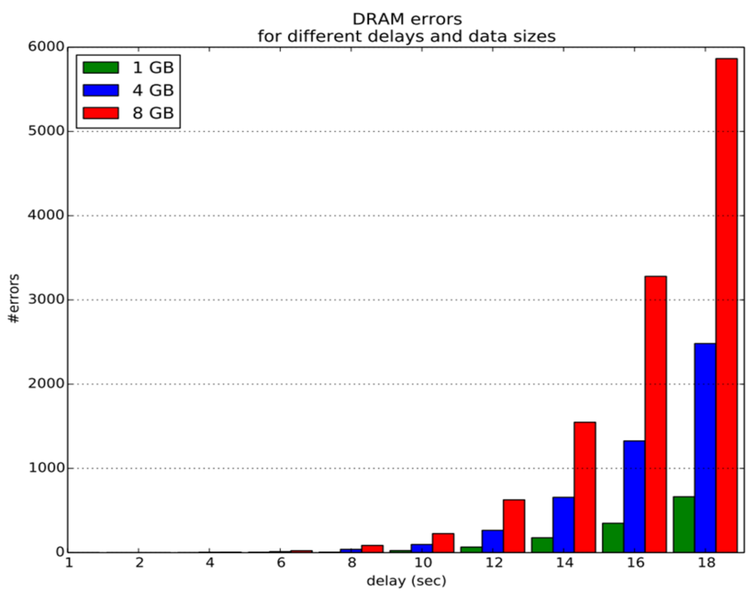

Disabled refresh rate of DRAMs

By changing the refresh interval of Dram we have observed that

- The first bit-flip occurs after an interval of more than 2 sec.
- The pessimistic worst-case timing of 64ms for retention time of each cell is useful only in extreme cases.
- There are large margins for reduction of the overheads.

5. Acknowledgement

The research leading to these results has received funding from the European Community: (a) Through the H2020 programme for research and technical development, grant agreement no. 688540 (UniServer), and (b) partially through the 7th Framework Programme (FP7/2007-2013), grant agreement no. FP7-323872 (SCoRPiO).