LINKSFOUNDATION.COM

FONDAZIONE links
PASSION FOR INNOVATION

Data-Centric Applications Management

Future Perspective

Alberto Scionti, Ph.D

ISC-2021 – HPC-IODC Workshop, July 2nd, 2021

Before starting …

Few words about me

- Master degree in computer engineering received in 2007 from Politecnico di Torino
- Ph.D degree in computer engineering received in 2011 from Politecnico di Torino:
  - Visiting researcher at UPC (Spain)
  - Visiting researcher at TU-Delft (The Netherldands)
- Post-doc position at Università degli Studi di Siena
- Background:
  - Digital circuit and memory design (mainly from a Test Engineer perspective)
  - CPU microarchitecture design (under EU funded projects' framework): data-flow based concept
  - Cloud computing, High-performance Computing and ... recently Quantum Computing application development
- Current position:
  - Senior researcher at LINKS Foundation (no-profit research center located in Turin – ITALY)

ACROSS

EuroHPC Project

- HPC BIG-DATA ARTIFICIAL INTELLIGENCE CROSS-STACK PLATFORM TOWARDS EXASCALE
- **EuroHPC-IA Call: H2020-JTI-EuroHPC-2019-1 -** Innovating and Widening the HPC use and skills base
- **Topic: EuroHPC-02-2019 - Type of action: IA:** HPC and data centric environments and application platforms
- Total cost: 8.815.845,00 Euros, EU contribution: 3.999.115,88 Euros
- Duration: 36 month, 13 partners
- Coordinator: Olivier Terzo (LINKS Foundation)

ACROSS
HPC Big DAta Artificial intelligence cross Stack PlatfoRm TOwardS ExaScale

ACROSS OBJECTIVES

EuroHPC Project

Overall

- O1.1 – Foundation and co-design of energy efficient HPC/BD/AI cross-stack platform
- O1.2 – Validation of ACROSS platform through industrial applications and fostering societal applications adoption
- O1.3 – Ensuring ACROSS platform interoperability, adaptability and security

Pilot

- O2.1 – Enhancing effectiveness in designing key aeronautical components (engine components, combustors, turbines) by adopting new workflows, multiscale/multiphysics unsteady approach, and artificial intelligence
- O2.2 – Enhancing global numerical weather prediction by means of hardware acceleration, low-latency exploitation of climate simulations, and enabling HPDA on large datasets
- O2.3 – Improving capability of performing large-scale carbon geologic sequestration simulations; enable direct subsurface flow simulations on processed seismic data; develop cross-stack workflows for subsurface simulation/analysis

Technological

- O3.1 – Co-designing HW/SW integration in strict collaboration with pilots, for creating exascale-ready services and ensuring compliance with the future EPI initiative
- O3.2 – Building a platform supporting the execution of large-scale HPC/Big Data/AI workflows with seamless exploitation of hardware heterogeneity
- O3.3 – **Efficient execution of mixed HPC/BD/DL workflows through hardware accelerators**, improving overall energy efficiency
- O3.4 – Optimizing workflow execution through the development of innovative orchestration and workflow management strategies
- O3.5 – Enhancing **performance for data stream processing, using in-situ/in-transit techniques**, in support of extreme-scale applications

MOTIVATIONS

Why improving data processing

- Data-centric Applications
  - Modern (scientific) workflows include machine-learning/deep-learning computations and data analytics tasks
  - In-situ/In-transit processing operations and visualization
- HPC machines require new levels of energy efficiency to scale towards Exascale and beyond
- von Neumann bottleneck:
  - I/O bandwidth is large but not infinite: moving back and forth large amount of data must be done carefully
  - There is (still) a large performance gap between memory elements close to the processing elements (e.g., cache hierarchy on CPU) and the large storage system (i.e., main memory, local disks, etc.):
    - Less-capacity communication channels connect main memory and the host processor
    - For instance, a DDR4 DIMM has an internal bandwidth of thousands of GB/s. vs an external bandwidth of few tens of GB/s
  - Poor spatial and temporal locality and unpredictable data access for modern data-centric applications that limit the benefit of HW solutions for latency-hiding (e.g., cache hierarchy)

DAMARIS

Scalable Middleware for Asynchronous I/O and In-Situ Processing On HPC Systems

- **Idea**: one dedicated I/O core per multicore node
- **Originality**: shared memory, asynchronous processing
- **Implementation**: software library
- **Experimented on**: Titan, Jaguar, Kraken, Blue Waters, Pangea, PRACE Hawk

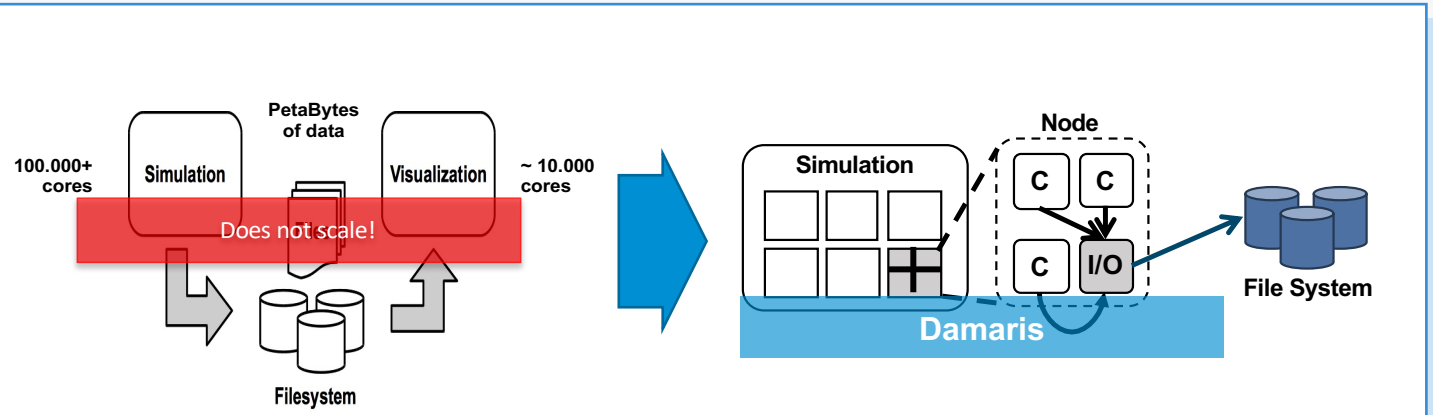

https://project.inria.fr/damaris/

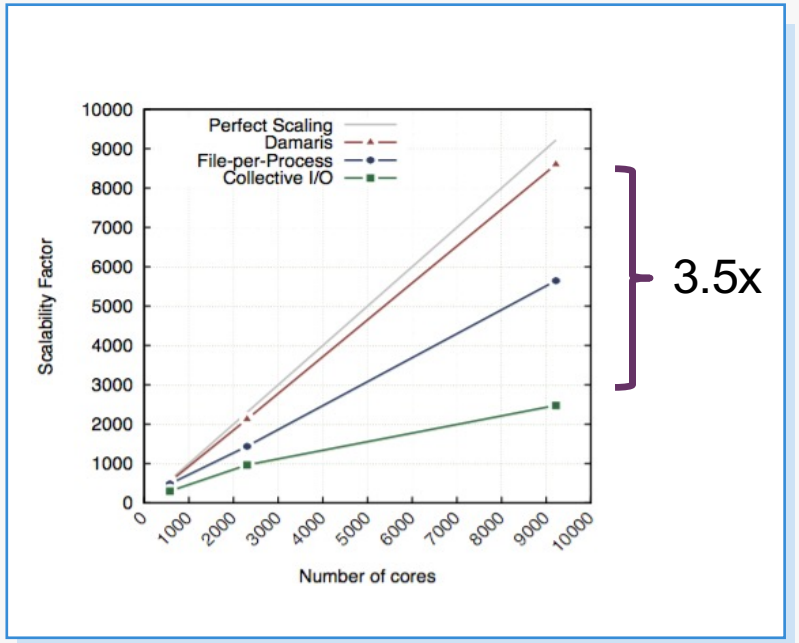

BURST BUFFERS

Scaling up I/O performance

- Improving I/O performance on HPC systems by using (dedicated) intermediate storage resources
- Do not sacrifice cores for managing I/O
- Large use of fast storage devices (SSDs)
- Further extension of the hierarchy by introducing NVM (i.e., NVRAM)

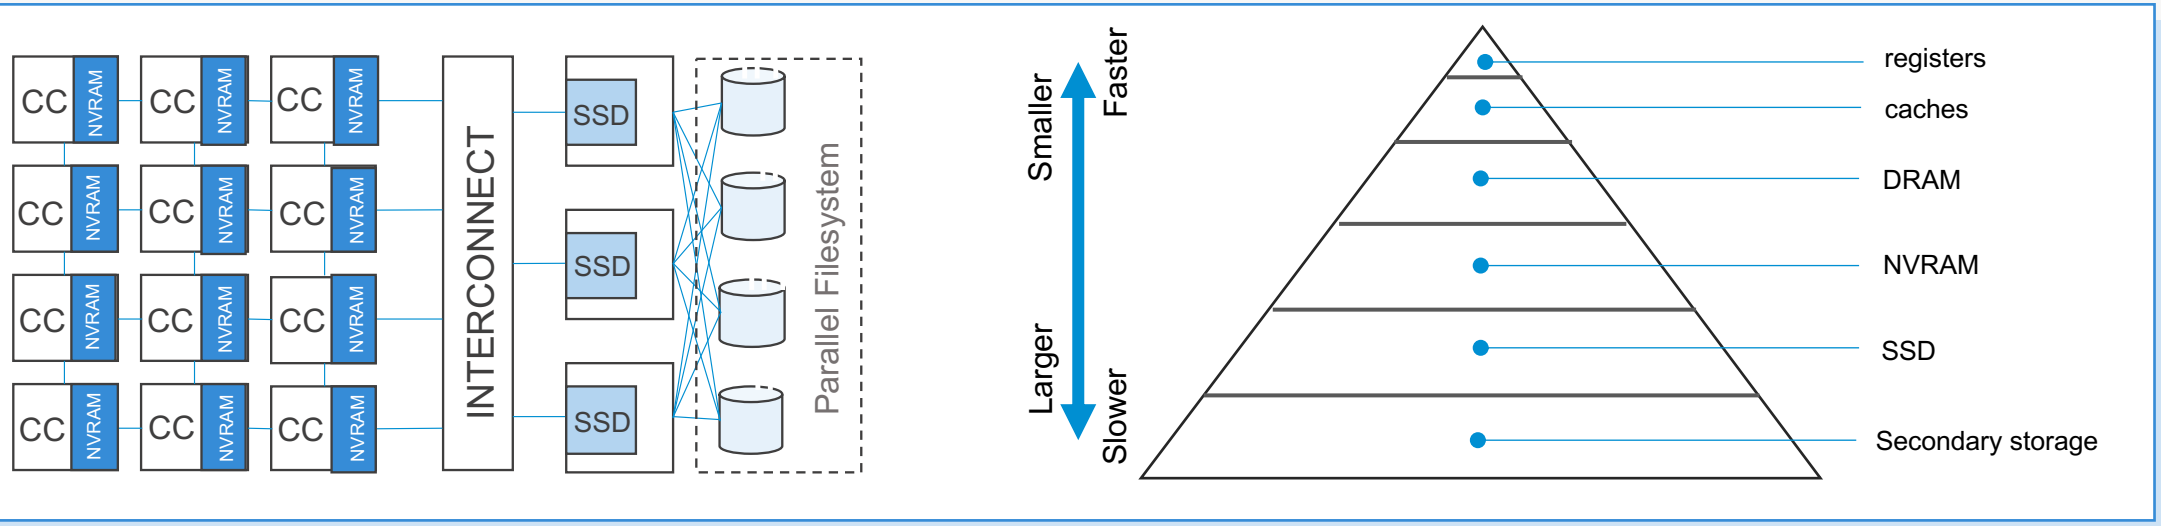

PROCESS IN-MEMORY

Is it the next step for Data-Centric Application Acceleration?

- Offloading data-centric (intensive) tasks from CPU cores to the memory device:
  - Exploiting the larger (internal to the device) bandwidth
  - Performing massive operation (e.g., MAC) with higher energy-efficiency
- Leveraging technology integration (3D stacking – improved memory bandwidth) for augmenting the standard main memory modules with processing capability
- Leveraging on NVM technology for a tradeoff between memory capacity, I/O bandwidth and data retention

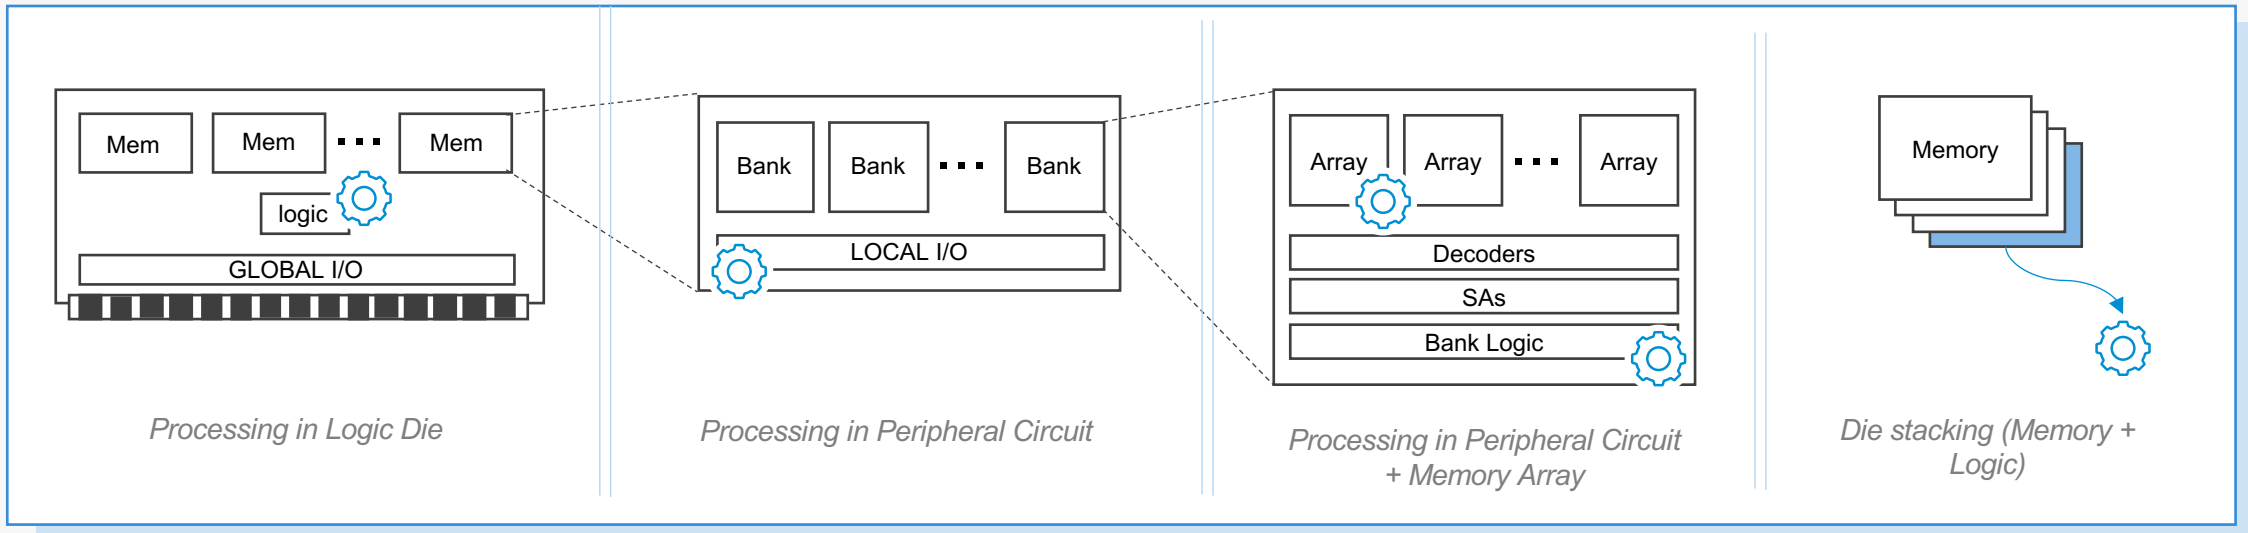

PROCESS IN-MEMORY

Is it the next step for Data-Centric Application Acceleration?

Technology readiness

- The idea is not entirely new: first experiments in the 1970s
- Several academic and industrial prototypes have been proposed:
  - E.g., Micron Automata Processor (based on DRAM devices), Micron In-Memory-Intelligence (IMI), Upmem DRAM Processing Unit (DPU), and recently Samsung Function-In-Memory DRAM (a programmable processing unit is added to each bank of a DRAM chip)
- Different technologies:
  - SRAM (very fast but volatile), ReRAM, STT-RAM, PCM
- Processing operations:
  - Pattern matching, filtering, massive MAC, support to CNNs/DNNs

Challenges

- Creation of a PIM ecosystem: software development and simulation environments, application vehicles, commercial products
- Carefully designing the interface between the PIM accelerators and the remainder of the system
- Management solutions: providing support at the BIOS, OS and system software (programming libraries, programming and execution models, etc.)
- Security (processing layer exposes a larger surface of attack)

FONDAZIONE links

PASSION FOR INNOVATION

CONTACTS

Alberto Scionti

Ph.D – Senior Researcher (Advanced Computing and Applications)

p. +39 011 22762905

e. alberto.scionti@linksfoundation.com

LINKS FOUNDATION

Via Pier Carlo Boggio 61 | 10138 Torino

P. +39 011 22 76 150

LINKSFOUNDATION.COM