MAESTRO
DATA ORCHESTRATION

Middleware for Memory and Data-Awareness in Workflows

Run time

3-year project, started in September 2018

Current status

Description of the requirements of partners' applications and workflows in order to design the first specification of the Maestro middleware API

Next steps

Finalize the specifications of the Maestro middleware architecture and design the demonstrators for the ECMWF and SIRIUS use cases

Partners

JÜLICH Forschungszentrum — JÜLICH SUPERCOMPUTING CENTRE

cea

ETH zürich

CSCS

appentra make code parallel

CRAY

ECMWF

SEAGATE

Motivation

- HPC and HPDA workloads are more and more I/O-intensive
- Performance bottlenecks are usually in the memory and storage systems
- Reducing and minimising data movement is very hard in general
- The HPC software stack was designed in a different era, to solve a different problem
- Few abstractions exist that capture data semantics of applications, so reasoning about data movement and memory in software is impossible
- Few useful models of memory systems and data movement exist, so estimation of cost of data movement is hard
- The memory-storage hierarchy is becoming more heterogeneous and complex, so a unified API and automatic promotion are needed

Maestro consortium is building a middleware library that characterises application data, reasons about how to load and store that data, assesses the cost of moving it and automates data movement across diverse memory systems

Approach and Methodology

- Co-design: ascertain data movement and access requirements of target applications
- Develop new data-aware abstractions:
  - Used in any level of software (compiler, runtime, application)
  - Relevant for any type of data (array, file, unspecified)
- Design a middleware and library that enables:
  - Modelling of memory hierarchy
  - Reasoning based on cost of moving data objects
  - Automatic movement and promotion of data in memories
  - Powerful data transformations and optimisation
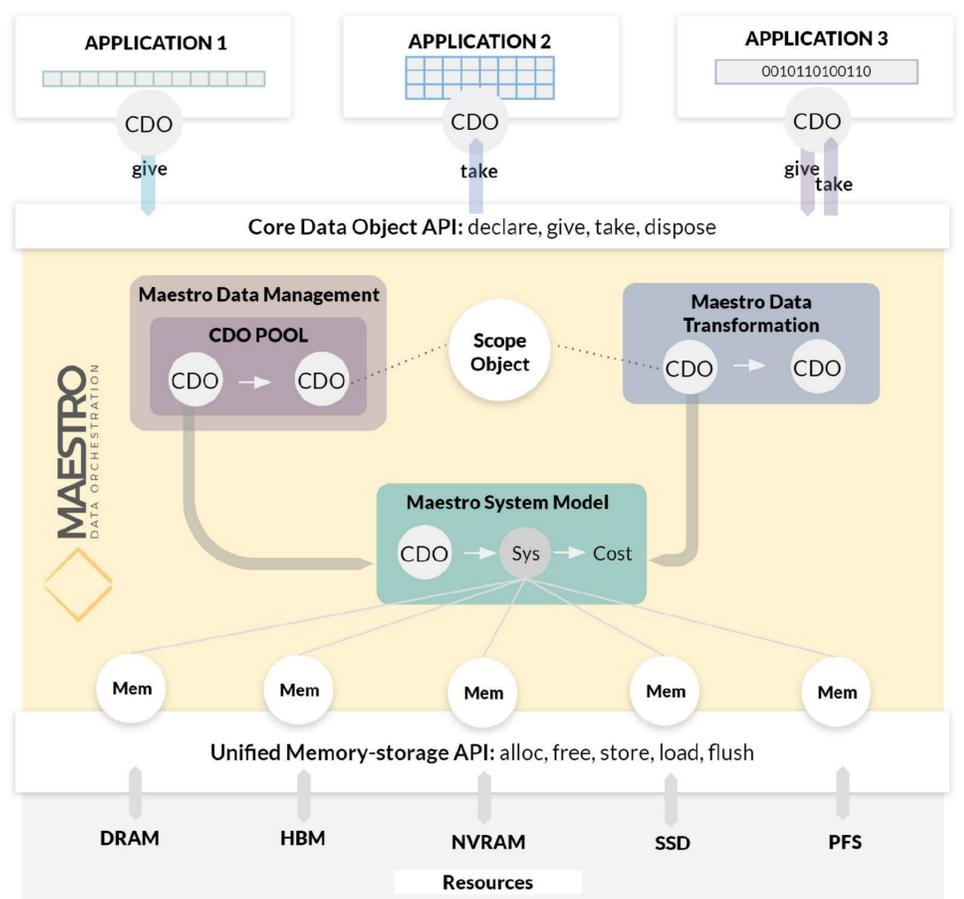- Explore data-based performance portability of Maestro applications

Figure:
Design of the Maestro middleware. The CDO (Core Data Object) is at the heart of Maestro's design. It is used to encapsulate data and meta-data.

maestro-data.eu @maestrodata /company/maestrodata