



This project has received funding from the European Union's Horizon 2020 research and innovation program through grant agreement 801101.



MAESTRO
DATA ORCHESTRATION

Maestro Project Introduction

1st European Communities Workshop on Exascale Computing

Dirk Pleiter

Jülich Supercomputing Centre

Today's Shortcomings: Lacking Data Awareness

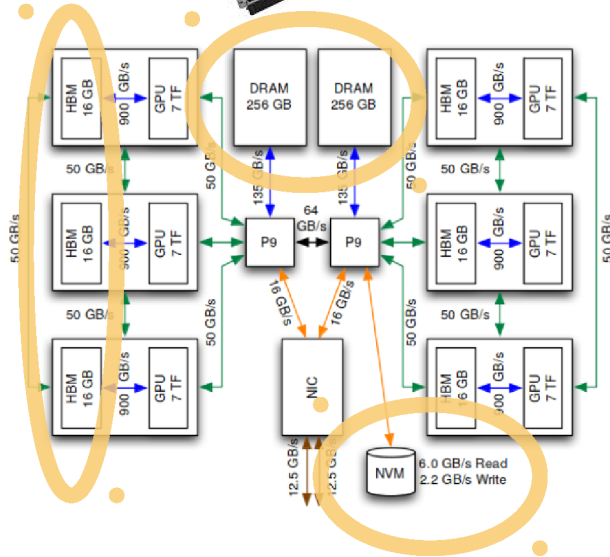
- Software stacks focussing on data processing
 - Optimised for filling of processing pipelines
 - Provide means for leveraging parallelism
- Lacking focus on basic data handling
 - Lacking functionality for controlling data handling
 - Lacking (unified) semantics for guiding data transport

Today's Shortcomings: Lacking Memory Awareness


- Missing information about available memory/storage hardware and its characteristics
 - Lacking ability for making data transport decisions
 - Missing information leads to hardware-neutral decisions
- Challenging variety of data access methods
 - Example storage class memory:
Block store, file system, object storage
- This becomes more critical with deeper memory and storage hierarchies

Hardware Architecture Examples: Summit

Compute nodes



External storage



Co-Design Applications

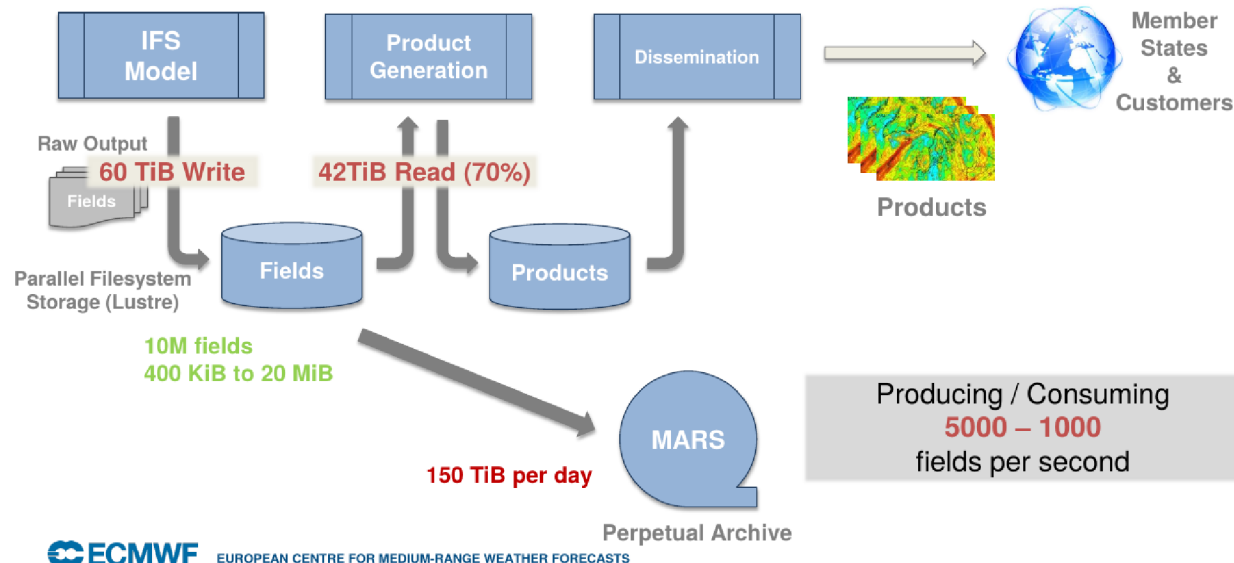
- IFS numerical weather prediction system (ECMWF)
 - Complex data processing and simulation system with multiple data producers and consumers
- Computational Fluid Dynamics plus in-situ analysis (CEA)
 - Pipeline coupling multiple simulations plus data post-processing
- Electronic structure calculation library SIRIUS (CSCS)
 - Simulations involving GPU acceleration
- Global Earth Modelling system TerrSysMP (JSC)
 - Coupled simulations

Co-Design Approach

- Usage scenarios
 - Description of a characteristic set of components and manner of use that is meaningful for application domain expert
- Use cases
 - Specific actions, functions or instances extracted from usage scenarios
- Requirements
 - capture anything that should influence the design or implementation of the Maestro middleware



Example: Weather Prediction Workflow



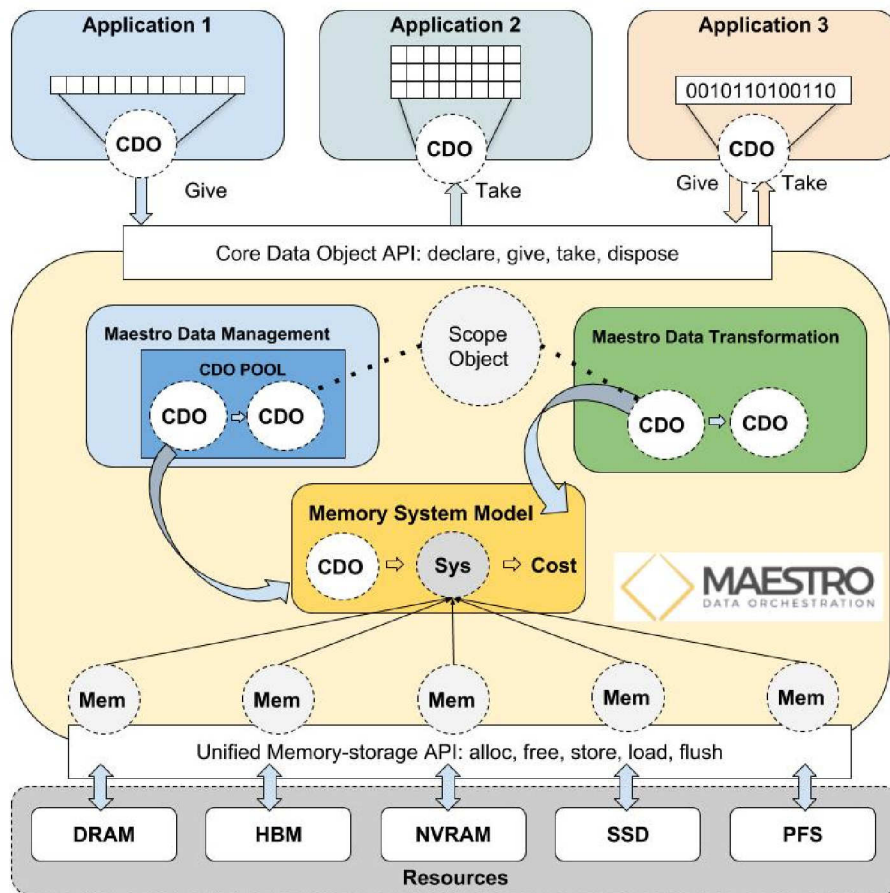
Today's bottlenecks

- Data movement between forecast stages and product generation
- Irregular archiving of output from research workflows

Maestro Solution Concepts

- Object-based approach to encapsulate data with application and Maestro related metadata
 - Core data objects
- Data movement decision based on workflow annotations and real-time I/O monitoring
 - Data object producer/consumer model

Architecture Overview



Project Schedule

- Requirements definition completed by August 2019
- Core design fully specified by April 2020
- Start system software and application demonstration in autumn 2019
- Project completion in August 2021

Summary and Outlook

- Today's HPC (and HPDA) solutions lack data and memory awareness
- Maestro will develop a data and memory aware middleware
 - Abstractions based on data objects
 - Memory-aware data transport and placement in middleware
- Open for providing early access to technology



MAESTRO

DATA ORCHESTRATION