

I/O Workload Simulator - a tool for future benchmarking and procurement of HPC systems

H2020 – NEXTGenIO project – 1st Scalability Day

Antonino Bonanni, Tiago Quintino, Simon Smart


Forecast Department (FD) – Development Section

Antonino.Bonanni@ecmwf.int, Tiago.Quintino@ecmwf.int, Simon.Smart@ecmwf.int

I/O Workload Simulator IOWS

- **Objectives**
 - Benchmarking
 - Assist HPC systems procurement
- **I/O Workload Simulator**
 - *A software tool capable of generating a workload model that is representative of a real-life HPC system workload*

I/O Workload Simulator IOWS



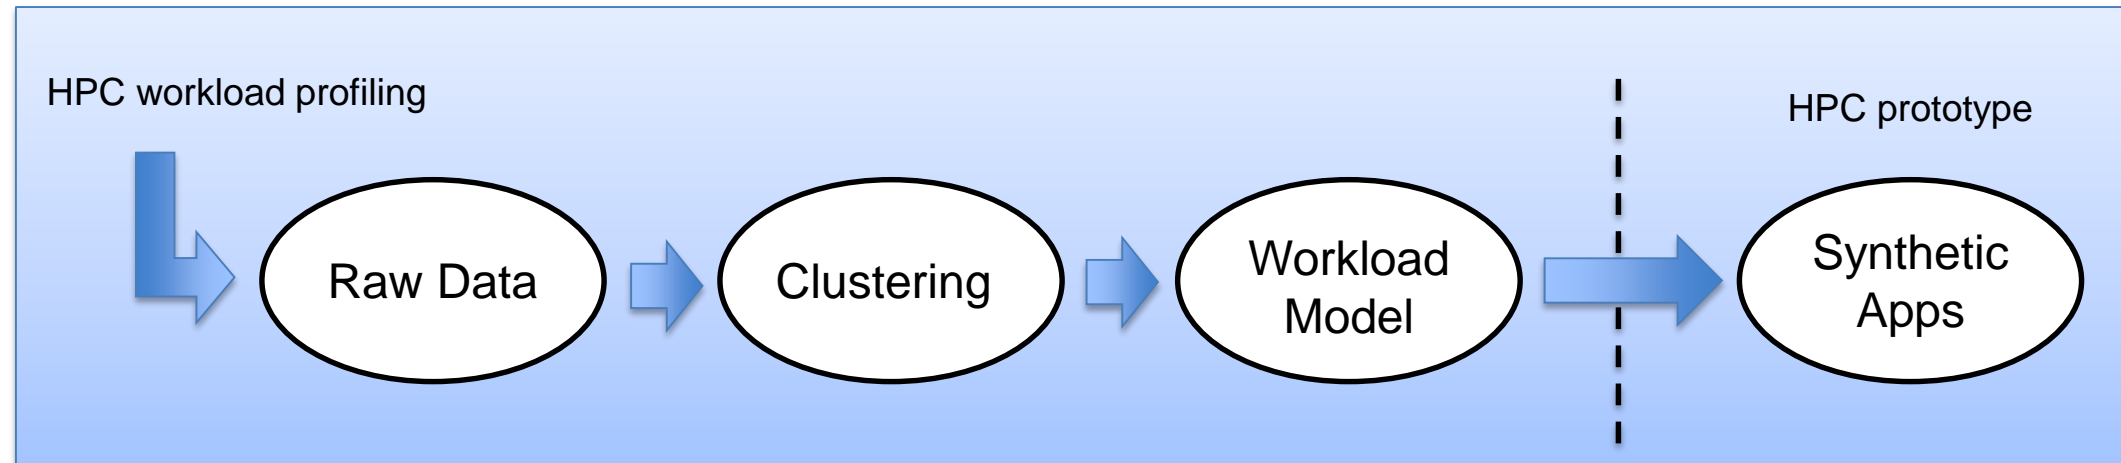
But not just I/O...

I/O Workload Simulator IOWS

- Use for benchmarking new HPC systems more thoroughly, including:
 - **I/O systems**
 - **Compute**
 - **Networking**
 - **Scheduling**

IOWS architecture

- **Data gathering.** The HPC systems are instrumented
- **Data clustering.** The collected data are clustered into “classes” of similar jobs
- **Workload model/schedule.** The workload is mapped onto a schedule of “idealised tasks”
- **Model execution.** A real/scaled schedule of synthetic applications is generated
- **Model simulation.** The execution of the meta-tasks “*might*” also be “simulated” in the abstract



Raw Data collection/handling

- **Job Metadata** collected from scheduler logs
 - Job submit statistics
 - HPC usage patterns
- **Time-dependent CPU, memory, Network, I/O metrics** from profiling tools
 - Profiling tools deployed at ECMWF, EPCC, ARCTUR
 - Provide more in depth understanding of the usage of HPC resources by each job
- **Data “gaps”** handled by regression algorithms
 - Measurements won’t be available system-wide and missing data are expected
 - Artificial Neural Network (efficient for large parameter spaces)

IOWS Modelling/Clustering

- **Job Characterization through:**


- Discretization of time signals
- Search for similar patterns

- **Clustering**

- **K-Means** Algorithm
- Self Organizing Maps (**SOM**)



- **Still to explore..**

- DBSCAN
- RBM/Deep-Learning
- PCA



K-Means

- Simple/scalable
- Clusters constrained to be convex subsets...



Example: colour mapping by SOM. Similar clusters end up close to each other (right pic)

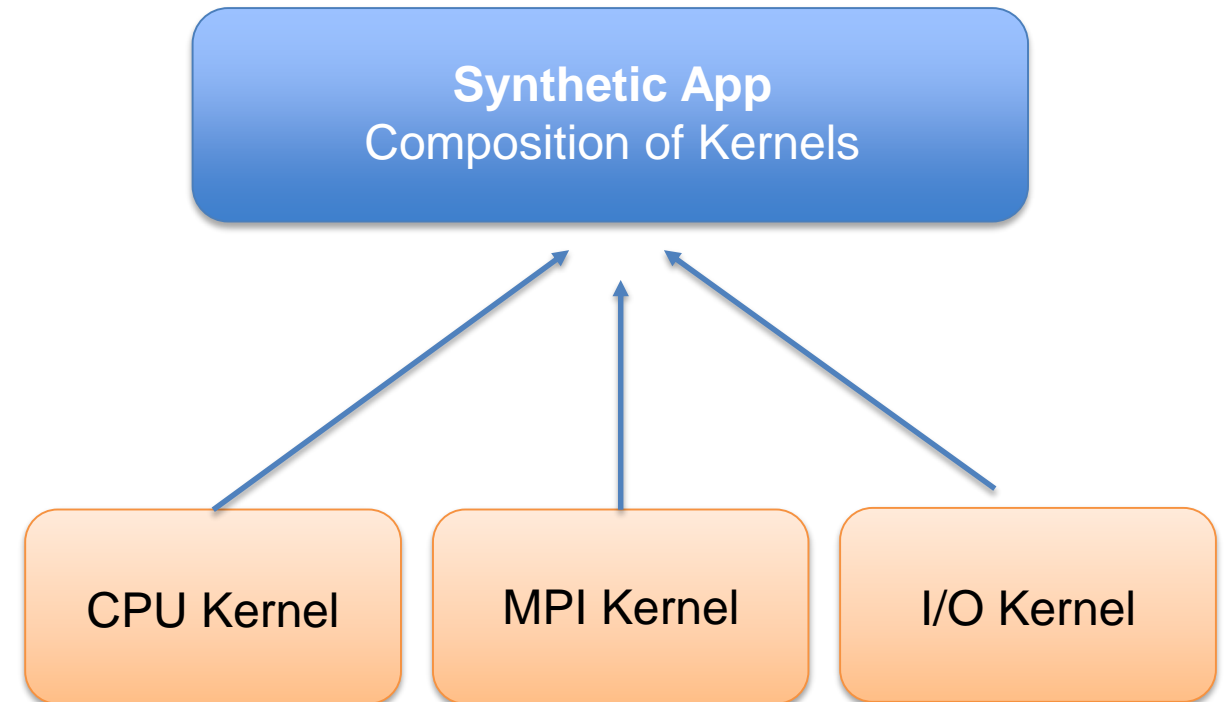
Synthetic Apps

- **Synthetic Apps**

- Deployed on HPC prototypes
- Light and portable apps
- Composition of kernels
 - CPU intensive kernels
 - I/O intensive kernels
 - Network intensive kernels (MPI)

- **Benchmarking:**

- RAPS (capability oriented)
- IOWS (capacity benchmarking)



Coming Next

- **Workload Executor**

- Deploys meta-tasks on an real HPC system
- Scaled if deployed on prototypes
- Full scale if deployed on operational hardware (testing and validation)

- **Workload Simulator**

- Artificially evaluates meta-tasks
- Extrapolated scaling factors, Discrete Event Simulations
- HPC architecture details needed for model setup

Many Thanks!
Questions?