


HyperLoom

S. Böhm¹, V. Cima¹, J. Martinovič¹, T. Ashby², T. Vander Aa², V. Chupakhin³
¹IT4Innovations, VSB-TU Ostrava, Czechia; ²IMEC, Belgium; ³Janssen Pharmaceutica NV, Belgium

A Platform for Defining and Executing Scientific Pipelines in Large-Scale Distributed Environments

GOAL: Efficient end-to-end data processing in large-scale distributed environments made simple



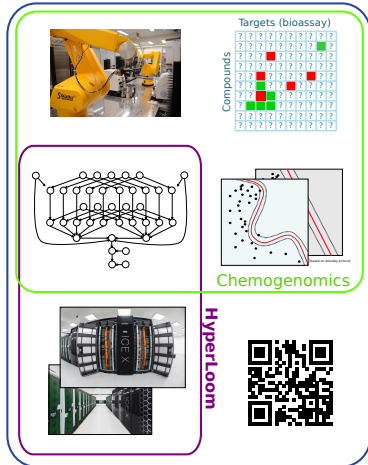
Motivation: Novel Drug Discovery

Chemogenomics

Modelling the interaction of the chemical compounds with biological organisms

Challenge: Running complex machine learning-based computational pipelines in HPC environments


Main objective: Easy to use framework for running large pipelines on HPC systems



Workflow

STEP 1: Setup

- node allocation (e.g. PBS job)
- HyperLoom deployment




STEP 2: Pipeline definition

- pipeline definition and submission

```
import loom.client as lc
COUNT = 100
tasks = [lc.Task.run("hostname") for i in range(COUNT)]
array = lc.Task.merge(tasks)
client = lc.Client("localhost", 9010)
future = client.submit_one(array)
#COMPUTATION
result = future.gather()
```

STEP 3: Pipeline execution

- task scheduling, distribution & execution



STEP 4: Result collection

- waiting for the submitted tasks
- collection of the results
- client connection termination

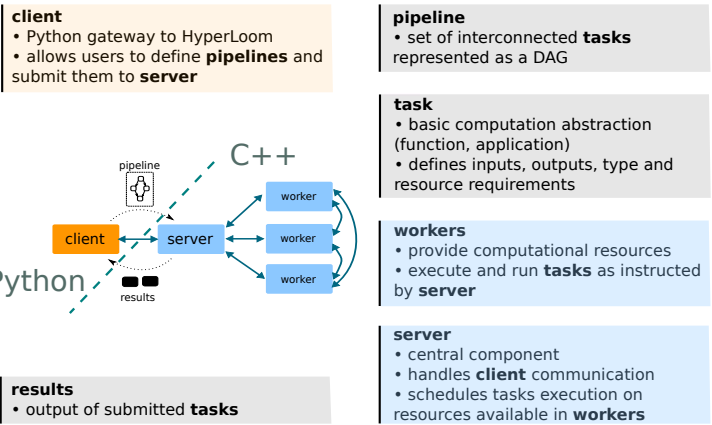
Features

- In-memory data processing reducing filesystem load
- Direct worker-to-worker data transfer reducing server overhead
- Support for execution of third party applications
- Support for short and long tasks (from tens of milliseconds to hours)
- Data-location aware scheduling algorithm reducing inter-node network traffic
- C++ core with a Python client enabling high performance through a simple API
- High scalability and native HPC support

Challenges

- Pipeline size** - pipeline can contain millions of tasks
- Pipeline shape** - pipeline can take shape of any directed acyclic graph
- Pipeline heterogeneity** - pipeline can contain various types of tasks
- Task execution time** - task execution time varies from milliseconds to hours and is difficult to be predicted
- Task output size** - size of the output generated by a task is not known before task finishes
- Cluster size** - HPC clusters contain thousands of computational nodes
- Cluster heterogeneity** - computational nodes may provide different resources and their capacities

Architecture




Performance

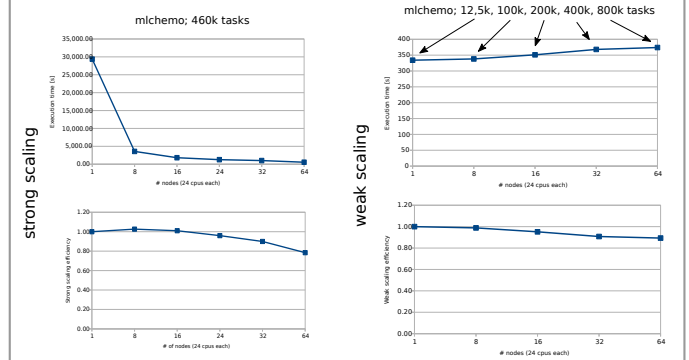
Test Scenarios

- 50kh**
 - test generating significant scheduling overhead
 - homogenous pipeline
 - ~50k short running independent tasks
- gridcat**
 - test generating significant network traffic
 - ~5k tasks generating significant amount of data
 - non-trivial pipeline shape
- mlchemo**
 - test representing complex ML pipeline derived from a real chemogenomics use case
 - highly varying task execution time
 - non-trivial pipeline shape

Comparing HyperLoom to Disk/Distributed¹



HyperLoom Scaling



¹Rocklin, Matthew. "Disk: Parallel computation with blocked algorithms and task scheduling." Proceedings of the 14th Python in Science Conference, 2015.



This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under Grant Agreement no. 671555. This work was supported by The Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project „IT4Innovations National Supercomputing Center – LM2015070“.