Predicting Execution Time of Bags-of-Tasks in Real Distributed Environments

Nathan Tran

Supervisor: Javid Taheri

School of Information Technologies
FACULTY OF ENGINEERING & INFORMATION TECHNOLOGIES

MOTIVATION

- Cluster schedulers need to know in advance the execution time of a task on a machine in order to generate efficient task assignments to machines
- Simulation and emulation tools are often employed to test scheduling algorithms
- However, there often exist discrepancies between test and real world distributed environments due to fluctuations in performance of computing resources and the network
- Thus, there is a need to develop model that can take a scheduling decision and predict how long it would actually take if deployed in a real environment

AIM

- Develop a statistical model for predicting makespans given a scheduling decision and a computing cluster
- Focus on compute-intensive bag-of-tasks (BoT) applications, thus transfer times are considered negligible

METHODS

Platform Setup

- HTCondor
  - An open-source high-throughput computing framework
  - Specialises in workload management for compute-intensive applications
- VMware vSphere
  - Cloud computing virtualisation operating system
  - Used to set up and manage our cluster of homogeneous virtual machines

Image-Resizing Program

- The program that tasks run on execute machines is an image-resizing application
- Written in java due to good Java support by HTCondor

Dataset

- Consists of 300 input images of different sizes

Data Collection and Storage

- Ran a total of 40 experiments and 12000 tasks
- Collect metrics such as:
  - Execution time
  - Task interval
  - Queue time
- Metrics are stored in an instance of InfluxDB, an open-source time series database

Regression Model

- The prediction engine performs regression analysis of collected metrics to generate makespan predictions
- 5-fold cross-validation implemented in order to assess the model

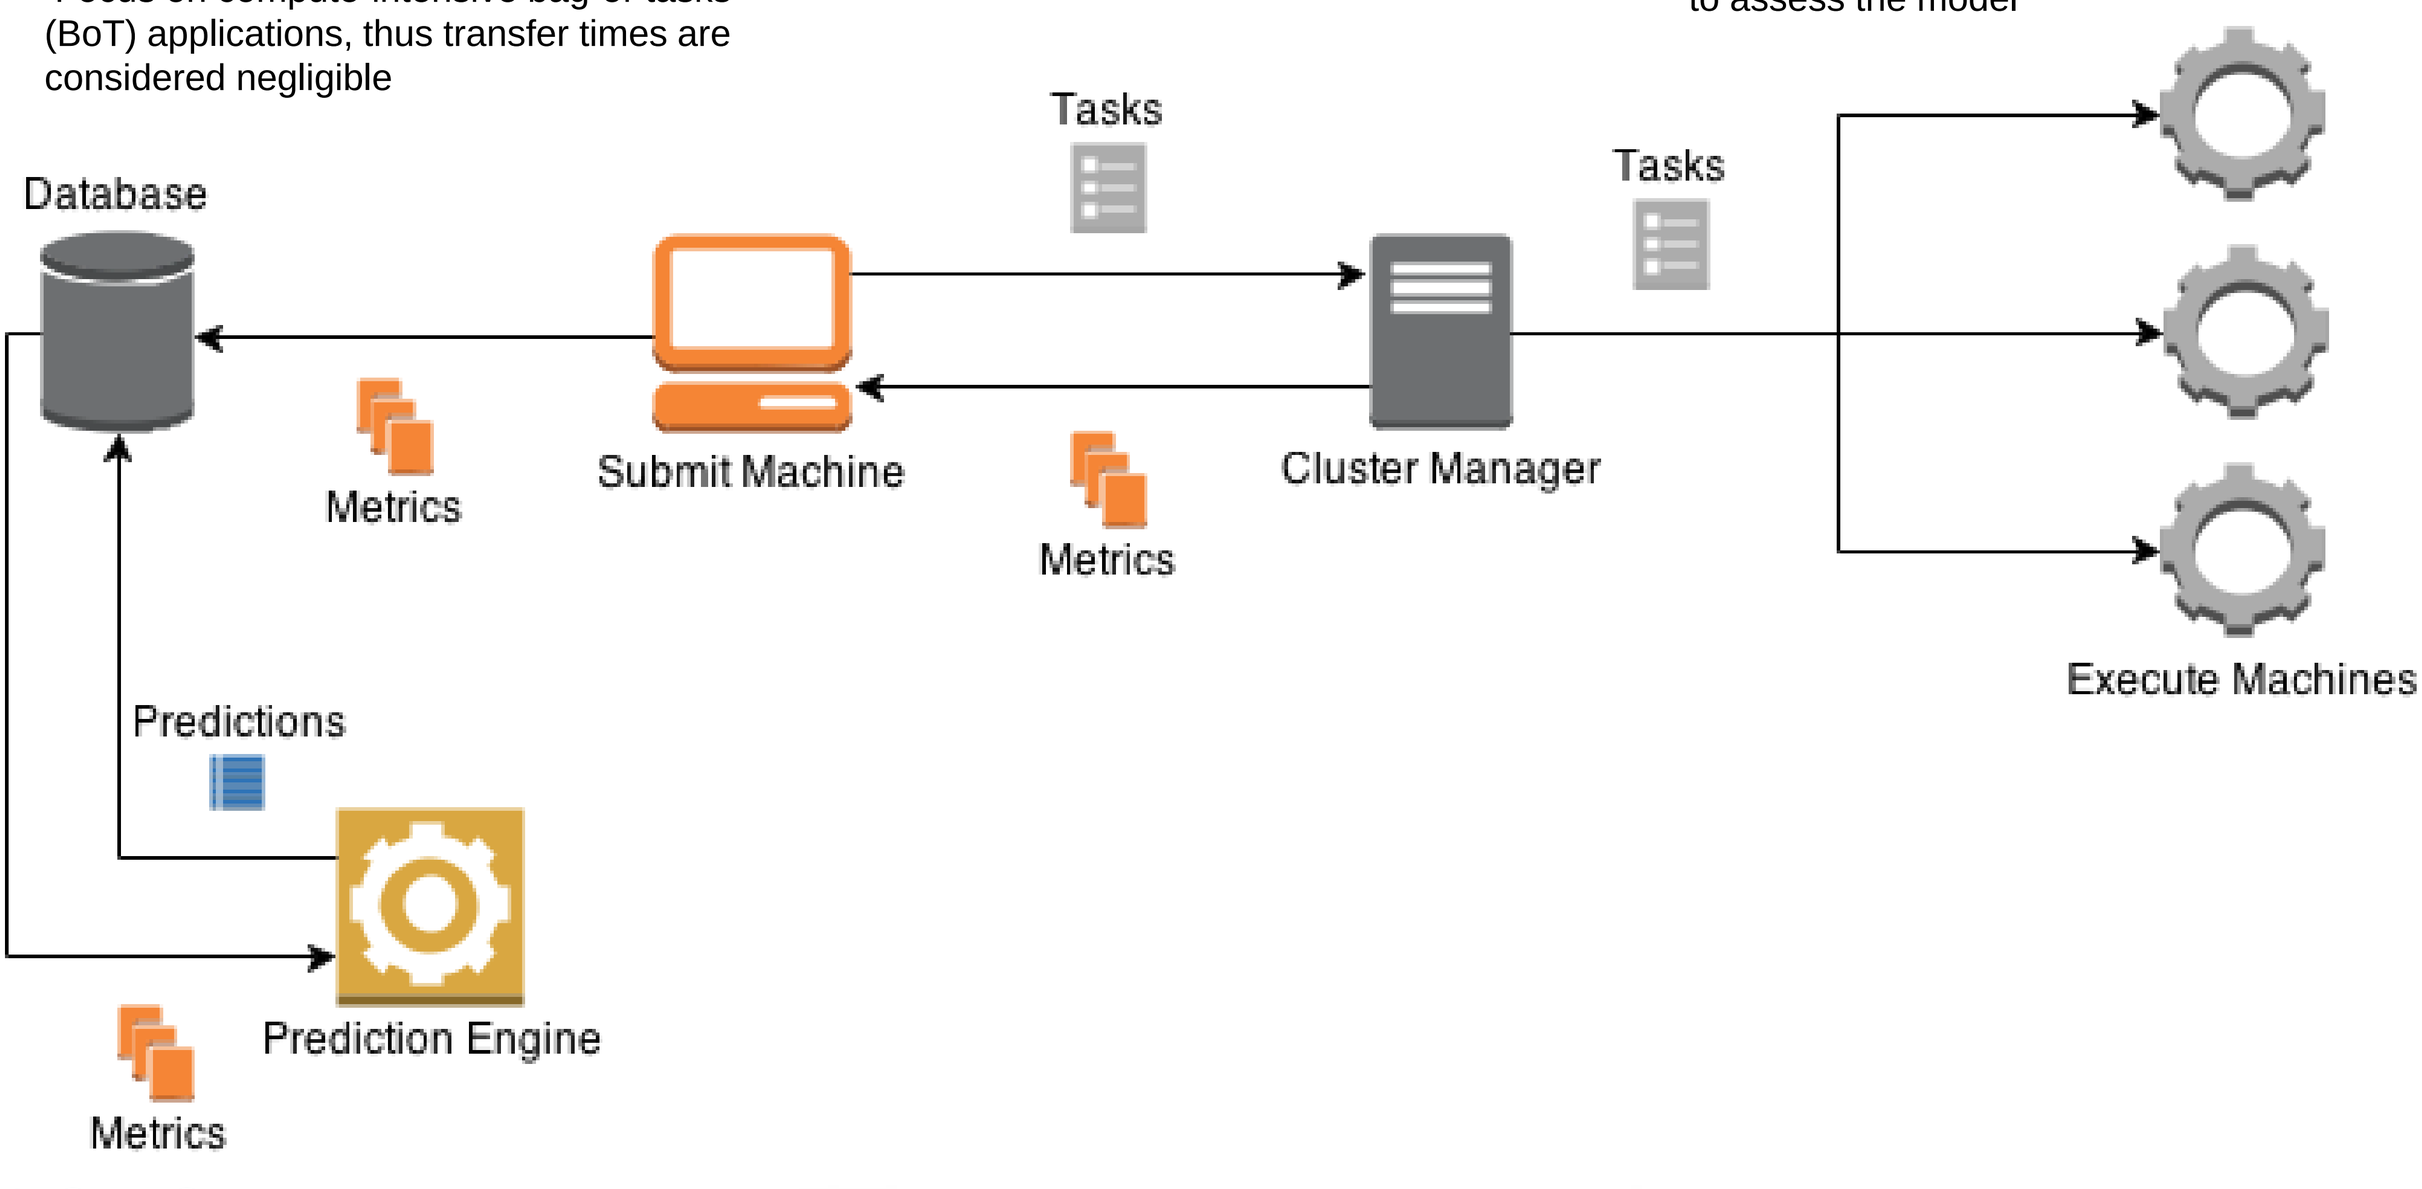

RESULTS

- The prediction engine generates results for linear regression, second-degree, third-degree, and fourth-degree polynomial regression models
- The models' accuracies are measured using *mean absolute percentage error (MAPE)*

degree	MAPE	stddev
linear	3.3847%	0.387%
2nd	3.4706%	0.384%
3rd	3.5608%	0.386%
4th	3.3826%	0.403%

- Low percentage errors for all models
- Low standard deviations
- No significant differences in terms of MAPE and standard deviation between the models
- Thus, a linear regression model suffices for most cases

Example output of linear regression between file size and execution time

CONCLUSION

- An accurate makespan prediction model for real distributed environments

FUTURE WORK

- Take into account transfer times for both input data and task
- Test with data-intensive BOT applications
- Improve the model to work with workflows
- Test on clusters of heterogeneous resources
- Test in environments with multiple users
- Test in environments where resources are geographically spread across multiple administrative domains