Popularity Prediction Tool for ATLAS Distributed Data Management

T Beermann, P Maettig, **G A Stewart**, M Lassnig, V Garonne, M Barisits, R Vigne, C Serfon, L Goossens, A Nairz and A Molfetas

Introduction

- ATLAS Distribution Data Management responsible for 150PB of data
- Constantly growing (see plot)
- Jobs running at over 100 computing sites
- **How do we best place data at the sites?**

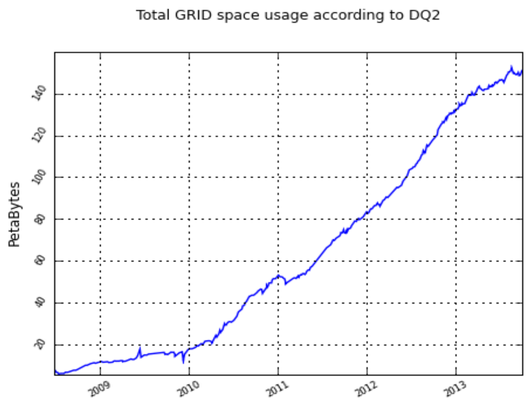

- Use Neural Networks to predict data popularity
- Ensure that the *right* data is in the *right* place, before the users run jobs
- Workload Simulation shows how much waiting time can be reduced

Popularity Prediction

- Try to find patterns in historic dataset popularity to predict future popularity.
- We predict future accesses using an Artificial Neural Network trained on past access to similar data.
- Take the first n weeks of a dataset lifecycle and predict week n+1.
- Different neural networks are used for different data types and physics streams as user behaviour is different.

- To evaluate the performance the method was compared to other means of predictions.

- **Static prediction:** The *n*th value of the input vector will be prediction for week *n+1.*
- **Linear prediction:** The prediction is the *n*th value of the input vector plus the difference of the *n*th value and the *n-1*th value.

- To compare the results the RMS of the difference between the actual and the predicted value was computed.

- Below are some sample results for one specific data type / physics stream combination.

  - **Neural Network:** 24.48
  - **Static:** 82.47
  - **Linear:** 207.7

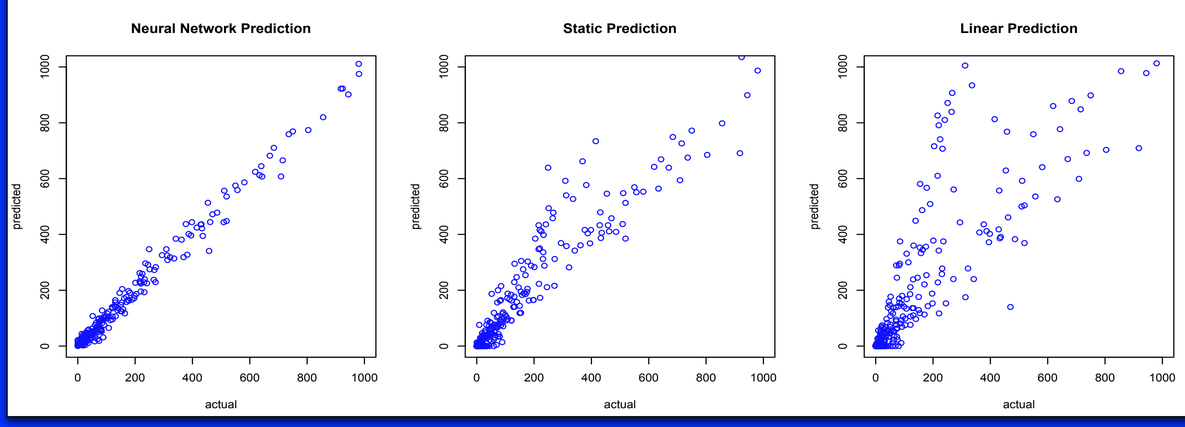

Simulator

- The predictions now can be used to improve data distribution.
- This process is complex because of the many degrees of freedom.
- To be able to evaluate the benefits of a different data distribution a simulator is needed
- The simulator has to be able to rerun the same workload on different data distributions.
- It needs to be able to collect useful metrics to compare the results of different distributions

The simulator consists of three main components:

- **Site Component**: Manages the available disk space, the number of job slots and the waiting queue.
- **Distributed Data Management:** Manages replicas and provides an interface to query replica locations.
- **Workload Management:** Schedules jobs at sites depending on the site's workload and the data locality.

The simulator was tested with different data distribution strategies:

- **Actual distribution:** As extracted from DDM.
- **Actual distribution, plus one extra replica:** Actual + 1 extra replica if dataset will be accesses.
- **Two extra replicas:** Same as above but with 2 replicas.
- **One extra replicas with space constraints:** First all extra replicas for datasets that will not be accessed are removed. Then the space is filled up with extra replicas of the most popular datasets.

The Graph below shows and example of the average job waiting time for the different methods for a test run of one day.

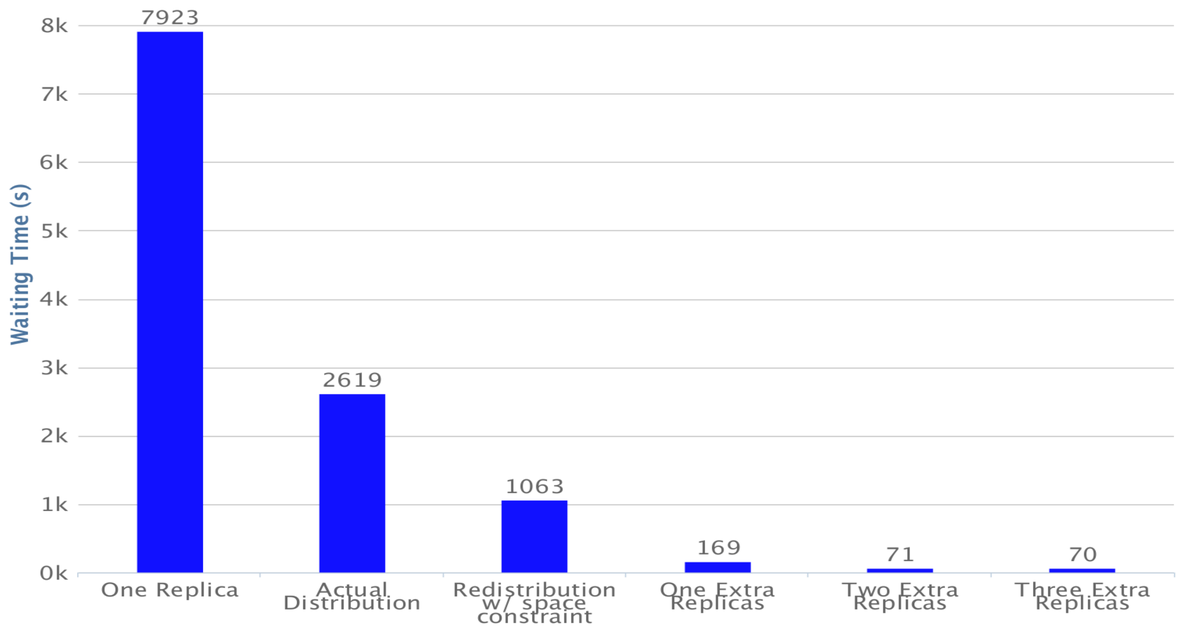

Conclusion

- In the currently running system data distribution is either a manual process or proceeds by creating new replicas reactively.
- We have shown how future access patterns can be predicted using ANNs. With this knowledge a better data distribution can be achieved.
- To evaluate the benefits a simulator was developed and first results were presented.
- These results showed that redistributing the data can lead to substantial reductions in waiting time for users.
- Based on these encouraging results different distribution strategies will be developed and evaluated.

CERN

ATLAS EXPERIMENT http://atlas.ch

BERGISCHE UNIVERSITÄT WUPPERTAL