

Nephology:

Modelling heterogeneity in cloud computing

Derek G. Murray and Steven Hand

Challenges

- **Heterogeneous performance**

Cloud providers use virtualisation to divide a single server into multiple virtual machines. The use of a work-conserving I/O scheduler means that activity in one customer's VM can affect I/O performance in another customer's VM. Therefore, new algorithms are necessary to achieve optimal performance.

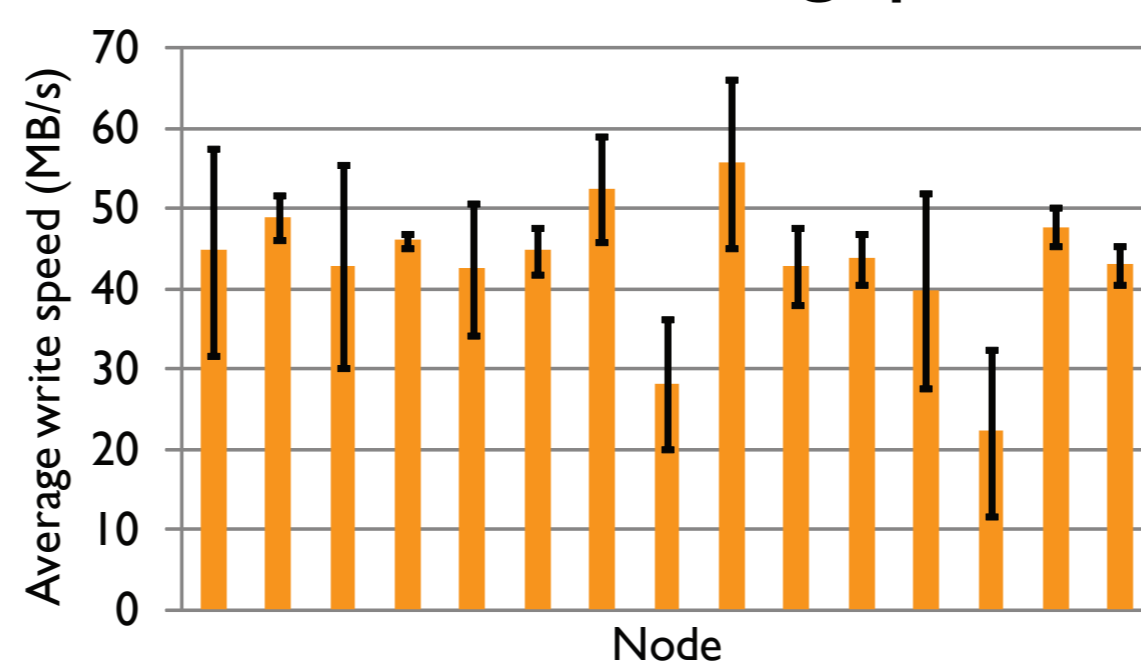
- **Expensive experiments**

The pay-per-hour model (e.g. \$0.10 per CPU-hour on Amazon EC2) makes it possible for cloud users to scale their systems to meet peaks in demand. It also enables researchers to run experiments at a large scale. However, the use of leased resources makes it expensive to conduct repeated or long-running experiments. Therefore, we want to conduct off-line experiments that accurately predict real cloud performance.

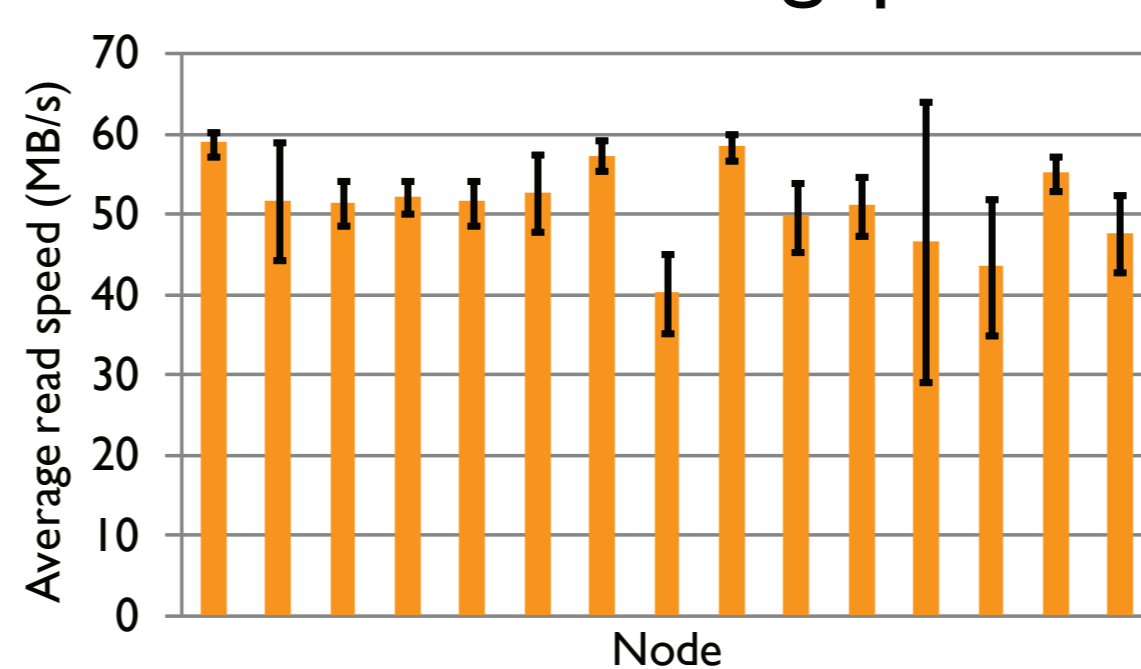
- **Unreproducible experiments**

Since the cloud resources are dynamically heterogeneous, it is not possible to reproduce an experiment under identical conditions. It is also difficult to measure the performance of two different algorithms under the same conditions. Therefore, to improve confidence in cloud measurements, we must capture the dynamic behaviour so that it may be replayed.




Disk write throughput




Disk read throughput



Network bandwidth




Cloud application
e.g. Hadoop, Dryad, etc.



Application Model
e.g. Hidden Markov Model,
Queueing Model



Cloud resources
e.g. Amazon EC2,
Windows Azure



Performance trace
e.g. Network I/O, Disk I/O,
CPU speed

Cloud Simulator

This combines the model and trace to predict end-to-end application performance. It will enable:

- Experimentation with a modified application model (e.g. scheduling)
- Independently reproducible results
- Performance prediction on other cloud platforms
- Cloud resource cost amortisation over multiple experiments

Opportunities

- **Better scheduling algorithms**

We will develop new application schedulers that use detailed performance measurements to make better decisions.

- **Computation dispersal**

In earlier work, we looked at how erasure coding can improve the performance of some parallel algorithms under machine failures and heterogeneity. We will evaluate this technique using our cloud traces.

- **Data placement**

Inter-node bandwidth depends on the cloud network topology (above). By studying real traces, we will devise topology-aware placement algorithms that achieve better performance.



UNIVERSITY OF
CAMBRIDGE

Derek.Murray@cl.cam.ac.uk
<http://www.cl.cam.ac.uk/~dgm36>