Optimizing Computation & Data Movement for Fresher Features

Sarah Wooders

RE STORE

PhD Student @ UC Berkeley Sky Computing (prev. RISELab)



What causes features (or data) to be stale?

1. Computational Bottlenecks

Problem: New data has not be processed to be reflected in feature values

ralf: Optimizing *computation* for fresher features

2. Data Movement Bottlenecks

Problem: Data is slow and expensive to replicate

Skyplane: Optimizing *data movement* for faster, cheaper data replication

ralf: Optimizing Computation for Feature Maintenance



ralf: Optimizing Computation for Feature Maintenance



Computationally Expensive

ralf: Optimizing Computation for Feature Maintenance



Model Performance Degradation

Computationally Expensive

Tradeoffs for Pre-computed Features



Opportunity for More Intelligent Scheduling





Scheduling Policy

We use the cumulative regret of a feature (since it was updated) as a measure to prioritize updating feature value.





Result: Higher Accuracy, Same Cost

Lower prediction error some equivalent feature update budgets across different workloads (anomaly detection, recommendation)



(From current paper in-review. Please email wooders@berkeley.edu if you're interested in seeing the submission!)



Result: Higher Accuracy, Same Cost

Lower prediction error some equivalent feature update budgets across different workloads (anomaly detection, recommendation)



(From current paper in-review. Please email wooders@berkeley.edu if you're interested in seeing the submission!)



For many applications, data movement is the bottleneck (<u>not</u> compute) - e.g. for feature publishing and serving.



Models may need to access data located on different clouds and regions.

Data may need to be delivered to applications on different cloud regions or providers.



Direct transfer is slow

Moving a 70GB dataset between regions (single cloud) or providers (multi-cloud)



1. Slow transfers = stale data

.

ubuntu@ip-172-31-82-174: ~

7.#3

(base) ubuntu@ip-172-31-82-174:-\$ aws s3 cp --recursive s3://skyplane-us-east-1/ s3://exps-paras-skylark-us-east-2/_ copy: s3://skyplane-us-east-1/00300.bin to s3://exps-paras-skylark-us-east-2/_/00300.bin copy: s3://skyplane-us-east-1/00302.bin to s3://exps-paras-skylark-us-east-2/_/00302.bin copy: s3://skyplane-us-east-1/00301.bin to s3://exps-paras-skylark-us-east-2/_/00302.bin copy: s3://skyplane-us-east-1/00301.bin to s3://exps-paras-skylark-us-east-2/_/00301.bin copy: s3://skyplane-us-east-1/00306.bin to s3://exps-paras-skylark-us-east-2/_/00305.bin copy: s3://skyplane-us-east-1/00306.bin to s3://exps-paras-skylark-us-east-2/_/00306.bin copy: s3://skyplane-us-east-1/00304.bin to s3://exps-paras-skylark-us-east-2/_/00304.bin copy: s3://skyplane-us-east-1/00304.bin to s3://exps-paras-skylark-us-east-2/_/00304.bin copy: s3://skyplane-us-east-1/00304.bin to s3://exps-paras-skylark-us-east-2/_/00304.bin copy: s3://skyplane-us-east-1/00304.bin to s3://exps-paras-skylark-us-east-2/_/00304.bin

70GiB dataset at 21MiB/s = 1 hour

Cost to move 70GB dataset

Running **34 instances**

2. High egress fees = \$\$\$





Why do slow transfers matter? Slow transfers = worse accuracy







What is Skyplane?

Problem: Managing data across regions and across clouds is <u>slow</u> and <u>expensive</u>

Skyplane is a system for fast, low-cost transfers between object stores.

skyplane cp {s3,gs,az}://... {s3,gs,az}://...

How does it work?

1. Faster networking:



\$

Control network route (overlay network) Striping over parallel VMs

2. Lower egress cost:

Compression + bandwidth tiering

Up to **113x faster** within single cloud

Up to **267x faster** between clouds

Up to 3.8x cheaper

Open source project!

\$ pip install skyplane

<u>expensive</u>

Skyplane is a system

skyplane cp {s3,gs,az}:/,

How does it work?

- 1. Faster network
 - Control network Striping over pare



skyplane.org

ers between object stores.

Up to **113x faster** within single cloud

Up to **267x faster** between clouds

p to 3.8x cheaper

2. Lower egr Compress

Direct internet path between clouds are often slow



Reasons for slow transfers

Congestion along direct path
Poor peering between providers
Packet loss from the physical layer

17

Direct internet path between clouds are often slow



Overlay routing allows circumventing slow links



Striping transfers large files across parallel VMs



2.7 gbps => 17.6 gbps

Low cost: Compression to reduce egress cost



All together: transfer datasets in minutes with Skyplane



Full benchmarks online at https://skyplane.org/en/latest/benchmark.html

Low cost: up to 3.8x cheaper for compressible data



23

Workload: 228GB Wikipedia dump Source: AWS us-east-1 Destination: AWS us-west-2 # of VMs: 8



All techniques explained in our NSDI 2023 paper

Overlay routing

Longer indirect paths are worthwhile for slow links

of VMs per region

Access throughput beyond NIC, AWS and GCP throttle egress

of parallel TCP connections

Unlike internet, fairness is a provider-level concern due to egress fees

Network tier selection

Hot potato routing up to 40% cheaper than cold potato

Skyplane: Network Overlays for Navigating the Cost–Performance Tradeoff for Inter-Cloud Bulk Transfers

Paras Jain, Sam Kumar, Sarah Wooders, Shishir Patil, Joseph Gonzalez, and Ion Stoica University of California, Berkeley

Abstract

Cloud applications are increasingly distributing data across multiple regions and cloud providers due to privacy regulations, availability of specialized hardware, and to prevent vendor lock-in. Unfortunately, wide-area bulk data transfers are often slow, bottlenecking applications. We demonstrate that it is possible to significantly improve inter-cloud bulk transfer throughput by adapting network overlays to the cloud. setting-that is, by routing data through indirect paths at the application layer. However, directly applying network overlays in this setting results in unaccentable increases in cloud egress prices. We present Skyplane, a system for bulk data transfer between cloud object stores that uses network overlays to ontimally navigate the trade-off between price and performance. Skyplane's planner uses mixed-integer linear programming to determine the optimal overlay path and resource allocation for data transfer subject to user-provided constraints on price or performance. Skyplane's data plane saturates the available bandwidth found by the planner with parallel reads/writes and integrates well with object stores. aligning with the stores' internal layouts. Together, these techniques allow Skyplane to significantly improve object transfer throughput at very low cost overheads.

1 Introduction

Many applications replicate data across multiple regions and datacenters, requiring fast built data transfers across the wide area. These applications include transferring data from one datacenter to others for data analysis (e.g. ETL [6], Geo-Diaributed Analytics [38]), and transferring the results of analysis to other datacenters for production (e.g. moving search indices [25]). There has been extensive prior work in optimizing the transfer time of built data transfers between datacenters, with application requirements on transfer time often modeled as deadines [25, 22, 82, 46]. All mainstream cloud providers offer services for built transfers (e.g. AWS' DataSyne, Aurz AcCopy, and GCP Cloud Transfer Service). Increasingly, applications are built not only on multiple regions within a cloud provider (multi-region), but also across multiple cloud providers (multi-cloud), due to variations in pricing, features, and availability of cloud services. In a recent survey [19], more than 86% of 727 respondents had adopted a multi-cloud strategy across diverse workloads. Thus, support for fast, cross-cloud built knasfers is increasingly important.

Bulk transfer performance is governed by two factors: transfer time (latency) and cost. Thus, we ask: How can we improve the transfer of data objects to support latency-sensitive workloads, while mininging the additional cost of achieving faster data transfers? We study this question in the context of designing and implementing Skyplane, an inter-cloud object transfer system which we plan to open-source.

A seemingly natural approach is to optimize the routing protocols used internally by cload providers to support higherthroughput inter-cload data transfers. Unfortunately, this is no feasible for two reasons. (1) Re-architecing the IP-layer routing the structure of the structure of the structure routing the structure of protocol and structure of the structure of the structure optimize efficient data transfer to other cloads. Indeed, AWS Snowball (14] and Azare Data Box Disk [8], all support data structure frainfor the other of the structure of the structure of the network plants between cloads is altimately spill among achieved with the source cload's cocoration as well.

Skyphane's key observation is that we can instead identify overlay pains—paths that send data via intermediate cloud regions—that are faster than the direct paths. For example, consider transferring data from Azare's Central Canada region to GCP's sails=northeast1. Using direct TCP connections would yield 6.2 GNps. Instead, Skyplane can first transfer data from iostiral Canada to Azare's US 2.2 to asil=northeast1 at 12.4 GNps. for a 2.0x speedup (Fig. 1). Crucially, this can be implemented on top of the cloud provider's verices, without their explicit buy-in.

Thank you!

Contact: wooders@berkeley.edu

ralf: <u>https://github.com/feature-store/ralf</u>

Skyplane:

- \$ pip install skyplane
- \$ skyplane init
- \$ skyplane cp -r s3://... gcs://...

