Low-Latency Geo-Distributed Data Analytics (Iridium)

$\bullet \bullet \bullet$

Qifan Pu, Ganesh Ananthanarayanan, Peter Bodik, Srikanth Kandula, Aditya Akella, Paramvir Bahl, Ion Stoica

SIGCOMM '15 (Aug 17-21)

Nicholas Turner

- Many industries have use cases where they would like to execute a query across data which is distributed over a large geographical area

- Many industries have use cases where they would like to execute a query across data which is distributed over a large geographical area
 - some examples include
 - determining global properties of a service across users median response time
 - whole cluster system log queries

- Many industries have use cases where they would like to execute a query across data which is distributed over a large geographical area

- The common approach to doing this involves copying ALL of the relevant data to a single data center, or to spread the work 'evenly' across datacenters

- Many industries have use cases where they would like to execute a query across data which is distributed over a large geographical area

- The common approach to doing this involves copying ALL of the relevant data to a single data center, or to spread the work 'evenly' across datacenters
 - but these methods can be time inefficient, or have high *query response time* when datacenter qualities are variable
 - This is particularly undesirable when the results of these queries are used directly by operators (analysts), or fuel interactive services

Core Ideas

- Can we instead process the data in a globally distributed way, assigning the heavy lifting of analytics queries intelligently?

Core Ideas

- Can we instead process the data in a globally distributed way, assigning the heavy lifting of analytics queries intelligently?

- Can we shuffle data around preemptively so that future queries are faster?

- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core_* (communication bottlenecks between site and core)

- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core*_ (communication bottlenecks lie between site and



- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core* (communication bottlenecks lie between site and core)
 - IO and CPU tasks take 0 time to execute (or at least that their time cost is negligible compared to communication latencies

- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core* (communication bottlenecks lie between site and core)
 - IO and CPU tasks take 0 time to execute (or at least that their time cost is negligible compared to communication latencies
 - Computation is split across three sites with the following characteristics

	Site-1	Site-2	Site-3
Input Data (MB), /	300	240	240
Intermediate Data (MB), ${\cal S}$	150	120	120
Uplink (MB/s), <i>U</i>	10	10	10
Downlink (MB/s), D	1	10	10

- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core* (communication bottlenecks lie between site and core)
 - IO and CPU tasks take 0 time to execute (or at least that their time cost is negligible compared to communication latencies
 - Computation is split across three sites with the following characteristics

	Site-1	Site-2	Site-3
Input Data (MB), /	300	240	240
Intermediate Data (MB), S	150	120	120
Uplink (MB/s), <i>U</i>	10	10	10
Downlink (MB/s), D	(1)	10	10

- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core* (communication bottlenecks lie between site and core)
 - IO and CPU tasks take 0 time to execute (or at least that their time cost is negligible compared to communication latencies
 - Computation is split across three sites with the following characteristics
 - Sites shuffle data in an 'all-to-all' fashion always sending a percentage of their data proportional to the number of reduce tasks on the destination site

	Site-1	Site-2	Site-3
Input Data (MB), /	300	240	240
Intermediate Data (MB), ${\cal S}$	150	120	120
Uplink (MB/s), <i>U</i>	10	10	10
Downlink (MB/s), D	1	10	10



- Consider a MapReduce query where
 - nodes communicate via a *congestion-free core* (communication bottlenecks lie between site and core)
 - IO and CPU tasks take 0 time to execute (or at least that their time cost is negligible compared to communication latencies
 - Computation is split across three sites with the following characteristics
 - Sites shuffle data in an 'all-to-all' fashion always sending a percentage of their data proportional to the number of reduce tasks on the destination site

	Site-1	Site-2	Site-3
Input Data (MB), /	300	240	240
Intermediate Data (MB), ${\cal S}$	150	120	120
Uplink (MB/s), <i>U</i>	10	10	10
Downlink (MB/s), D	$\begin{pmatrix} 1 \end{pmatrix}$	10	10



- Assuming the model above (and that tasks are infinitesimally divisible), you can find an optimal solution for a single reduce task by the following a linear program

$$S = \sum_i S_i$$
 $T_i^U(r_i) = rac{(1-r_i)S_i}{U_i}$ $T_i^D(r_i) = rac{r_i(S-S_i)}{D_i}$

$$\begin{array}{ll} \min & z \\ \text{s.t.} & \forall i:r_i \geq 0 \\ & \sum_i r_i = 1 \\ \forall i:T_i^U(r_i) \leq z, \quad T_i^D(r_i) \leq z \end{array}$$

- Assuming the model above (and that tasks are infinitesimally divisible), you can find an optimal solution for a single reduce task by the following a linear program
- In practice, they relax the divisibility assumption above by using a Mixed Integer Program instead (which takes a bit more time)

$$S = \sum_i S_i$$
 $T_i^U(r_i) = rac{(1-r_i)S_i}{U_i}$ $T_i^D(r_i) = rac{r_i(S-S_i)}{D_i}$

$$\begin{array}{ll} \min & z \\ \text{s.t.} & \forall i:r_i \geq 0 \\ & \sum_i r_i = 1 \\ \forall i:T_i^U(r_i) \leq z, \quad T_i^D(r_i) \leq z \end{array}$$

- Assuming the model above (and that tasks are infinitesimally divisible), you can find an optimal solution for a single reduce task by the following a linear program
- In practice, they relax the divisibility assumption above by using a Mixed Integer Program instead (which takes a bit more time)
- They also approximate the solution for multiple linked tasks (DAGs) greedily

$$S = \sum_i S_i$$
 $T_i^U(r_i) = rac{(1-r_i)S_i}{U_i}$
 $T_i^D(r_i) = rac{r_i(S-S_i)}{D_i}$

$$\begin{array}{ll} \min & z \\ \text{s.t.} & \forall i: r_i \geq 0 \\ & \sum_i r_i = 1 \\ \forall i: T_i^U(r_i) \leq z, \quad T_i^D(r_i) \leq z \end{array}$$

- The terms in the last equations depended heavily upon S_i (which in turn depend on I_i), if we knew which queries were coming, we could optimize these terms to find better solutions (though the theoretical formulation isn't tractable here).

$$S = \sum_i S_i$$
 $T_i^U(r_i) = rac{(1-r_i)S_i}{U_i}$ $T_i^D(r_i) = rac{r_i(S-S_i)}{D_i}$

- The terms in the last equations depended heavily upon S_i, if we knew which queries were coming, we could optimize these terms to find better solutions (though the theoretical formulation isn't tractable here).



- The conceptual lesson from the last example is that data shouldn't reside on 'bottleneck' sites when possible

- The conceptual lesson from the last example is that data shouldn't reside on 'bottleneck' sites when possible

- Luckily, the LP is efficient, and it tells us which sites are bottleneck nodes

$$\begin{array}{ll} \min & z \\ \text{s.t.} & \forall i: r_i \geq 0 \\ & \sum_i r_i = 1 \\ & \forall i: T_i^U(r_i) \leq z, \quad T_i^D(r_i) \leq z \end{array}$$

- The conceptual lesson from the last example is that data shouldn't reside on 'bottleneck' sites when possible

- Luckily, the LP is efficient, and it tells us which sites are bottleneck nodes

- So, although we can't find an exact solution, we can take a greedy approach by considering data movements in small increments - frequently running the LP to see the effects on the eventual query response time

Technical Design - Multiple Datasets

- None of this is useful in real data centers unless you're able to prioritize which datasets to optimize at a given point in time

Technical Design - Multiple Datasets

- None of this is useful in real data centers unless you're able to prioritize which datasets to optimize at a given point in time
- Fortunately, we can do this by iterating over all datasets, and consolidating the query frequency of each query for a particular dataset

```
1: procedure AllocateMoves(List(Dataset) D)
2:
       for each Dataset d in D do
3:
           Move d.m \leftarrow FINDMOVE(d)
           lag \leftarrow \sum_{q \in d.Queries} q.lag / d.Queries.Count
4:
           d.value \leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] /
5:
   lag
           d.score \leftarrow \frac{d.value}{d.m.cost}
6:
7:
       for each Dataset d in D.SortedByDesc(d.score) do
8:
           if d.m.bottleneck.canMove() then
9:
               execute d.m
```

Technical Design - Other Features

- Can impose a tunable **WAN-usage Budget** by first running a WAN-optimized strategy, and then only moving data if it fits under some multiplicative factor of the usage of that strategy (notably, 1.3 factor yields 90% of the performance gains)

- Using *query/data contention* they only move a data packet if it won't upset currently running queries too much (also estimated using the same greedy methods)

Implementation

- Written as a modified version of Spark (~MapReduce successor) using Hadoop Distributed File System
 - Measure bandwidth of each site and intermediate data size by empirical estimation from data throughput and task flows

- 8 Globally distributed EC2 instances
 - in Tokyo, Singapore, Sydney, Frankfurt, Ireland, Sao Paulo, Virginia, and California
 - also use 30 instances within one region with tuned bandwidth distributions

- 8 Globally distributed EC2 instances
 - in Tokyo, Singapore, Sydney, Frankfurt, Ireland, Sao Paulo, Virginia, and California
 - also use 30 instances within one region with tuned bandwidth distributions
- Run query workload benchmarks from 4 different applications
 - Conviva Video Analytics queries from a video delivery and monitoring company
 - Microsoft Bing Edge Distribution Edge server streaming queries
 - TPC-DS Benchmark Decision support queries modelled after Amazon
 - **AMPLab Big-Data** combination of Hive and Spark queries using the same schema throughout

- 8 Globally distributed EC2 instances
 - in Tokyo, Singapore, Sydney, Frankfurt, Ireland, Sao Paulo, Virginia, and California
 - also use 30 instances within one region with tuned bandwidth distributions
- Run query workload benchmarks from 4 different applications
 - Conviva Video Analytics queries from a video delivery and monitoring company
 - Microsoft Bing Edge Distribution Edge server streaming queries
 - TPC-DS Benchmark Decision support queries modelled after Amazon
 - **AMPLab Big-Data** combination of Hive and Spark queries using the same schema throughout
- Also perform a trace-driven simulation of Facebook's Hadoop cluster

- 8 Globally distributed EC2 instances
 - in Tokyo, Singapore, Sydney, Frankfurt, Ireland, Sao Paulo, Virginia, and California
 - also use 30 instances within one region with tuned bandwidth distributions
- Run query workload benchmarks from 4 different applications
 - Conviva Video Analytics queries from a video delivery and monitoring company
 - Microsoft Bing Edge Distribution Edge server streaming queries
 - TPC-DS Benchmark Decision support queries modelled after Amazon
 - **AMPLab Big-Data** combination of Hive and Spark queries using the same schema throughout
- Also perform a trace-driven simulation of Facebook's Hadoop cluster
- Compare against two baseline strategies:
 - leaving data 'in-place' & centralization strategy (with very high in-DC bandwidth)
 - both use Spark's standard scheduling routines



Figure 5: EC2 Results across eight worldwide regions (a): Tokyo, Singapore, Sydney, Frankfurt, Ireland, Sao Paulo, Virginia (US) and California (US). The figure on the right (b) is on a larger 30-site setup. Iridium is $3 \times -19 \times$ better compared to the two baselines.

	lridium vs. In-place	lridium vs. Centralized
Core	26%	32%
Core + Query Lag	41%	46%
Core + Query Lag		
+ Contention	59%	74%
Core + Contention	45%	53%

Table 2: Progression of Iridium's gains as additional features of considering query lag and contention between query/data movements are added to the basic heuristic. (Facebook workload)





Figure 7: Iridium's improvements (and % queries), bucketed by various query characteristics: (a) intermediate/input data ratio, (b) dataset access count, (c) query size (# tasks), and (d) cross-site skew in intermediate data.

Lag Metric	Vs. In-place	Vs. Centralized
Iridium (Avg.)	59%	74%
Iridium (Median)	56%	75%
Iridium (Earliest)	38%	42%
Iridium (Latest)	24%	40%
Oracle	66%	81%

Table 3: Effectiveness of estimating query lag. Iridium's approach of using the average lag outperforms other options and crucially, has gains of $\sim 90\%$ of an oracle that has full knowledge about query arrivals.

class Move double cost $\langle \mathbf{QueryID}, \mathbf{double} \rangle$ timeReduction Site bottleneck1: procedure ALLOCATEMOVES(List $\langle \mathbf{Dataset} \rangle D$)2: for each Dataset d in D do3: Move d.m \leftarrow FINDMOVE(d)4: lag $\leftarrow \sum_{q \in d.Queries} q.lag / d.Queries.Count$ 5: d.value $\leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] / lag$ 6: d.score $\leftarrow \frac{d.value}{d.m.cost}$ 7: for each Dataset d in D.SortedByDesc(d.score) do8: if d.m.bottleneck.canMove() then9: execute d.m	
$\begin{array}{c} \operatorname{double \ cost} & \langle \operatorname{QueryID, \ double} \rangle \ \operatorname{timeReduction} \\ \operatorname{Site \ bottleneck} \\ 1: \ \operatorname{procedure \ ALLOCATEMOVES(List\langle \operatorname{Dataset} \rangle \ D)} \\ 2: \ \ \operatorname{for \ each \ Dataset} \ d \ in \ D \ do \\ 3: \ \ \operatorname{Move \ d.m} \leftarrow \operatorname{FINDMOVE(d)} \\ 4: \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ $	class Move
$\begin{array}{ll} \langle \mathbf{QueryID, double} \rangle \mbox{ timeReduction} \\ \mathbf{Site \ bottleneck} \\ 1: \ \mathbf{procedure \ ALLOCATEMOVES(List\langle \mathbf{Dataset} \rangle \ D)} \\ 2: \ \ \mathbf{for \ each \ Dataset \ d \ in \ D \ do} \\ 3: \ \ \mathbf{Move \ d.m \leftarrow FINDMOVE(d)} \\ 4: \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ $	double cost
Sitebottleneck1:procedureALLOCATEMOVES(List (Dataset) D)2:for each Dataset d in D do3:Move d.m \leftarrow FINDMOVE(d)4:lag $\leftarrow \sum_{q \in d.Queries} q.lag / d.Queries.Count$ 5:d.value $\leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] / lag6:d.score \leftarrow \frac{d.value}{d.m.cost}7:for each Dataset d in D.SortedByDesc(d.score) do8:if d.m.bottleneck.canMove() then9:execute d.m$	$\langle \mathbf{QueryID, double} \rangle$ timeReduction
1: procedure ALLOCATEMOVES(List(Dataset) D) 2: for each Dataset d in D do 3: Move d.m \leftarrow FINDMOVE(d) 4: lag $\leftarrow \sum_{q \in d.Queries}$ q.lag / d.Queries.Count 5: d.value $\leftarrow \sum_{q \in d.Queries}$ d.m.timeReduction[q] / lag 6: d.score $\leftarrow \frac{d.value}{d.m.cost}$ 7: for each Dataset d in D.SortedByDesc(d.score) do 8: if d.m.bottleneck.canMove() then 9: execute d.m	Site bottleneck
2:for each Dataset d in D do3:Move d.m \leftarrow FINDMOVE(d)4:lag $\leftarrow \sum_{q \in d.Queries} q.lag / d.Queries.Count$ 5:d.value $\leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] / lag6:d.score \leftarrow \frac{d.value}{d.m.cost}7:for each Dataset d in D.SortedByDesc(d.score) do8:if d.m.bottleneck.canMove() then9:execute d.m$	1: procedure AllocateMoves(List(Dataset) D)
3: Move $d.m \leftarrow FINDMOVE(d)$ 4: $lag \leftarrow \sum_{q \in d.Queries} q.lag / d.Queries.Count$ 5: $d.value \leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] / lag$ 6: $d.score \leftarrow \frac{d.value}{d.m.cost}$ 7: for each Dataset d in D.SortedByDesc(d.score) do 8: if d.m.bottleneck.canMove() then 9: execute d.m	2: for each Dataset d in D do
4: $\operatorname{lag} \leftarrow \sum_{q \in d.Queries} q.\operatorname{lag} / d.\operatorname{Queries.Count}$ 5: $\operatorname{d.value} \leftarrow \sum_{q \in d.Queries} \operatorname{d.m.timeReduction}[q] / \operatorname{lag}$ 6: $\operatorname{d.score} \leftarrow \frac{\operatorname{d.value}}{\operatorname{d.m.cost}}$ 7: for each Dataset d in <i>D</i> .SortedByDesc(d.score) do 8: if d.m.bottleneck.canMove() then 9: execute d.m	3: Move $d.m \leftarrow FINDMOVE(d)$
5: d.value $\leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] / lag 6: d.score \leftarrow \frac{d.value}{d.m.cost}7: for each Dataset d in D.SortedByDesc(d.score) do8: if d.m.bottleneck.canMove() then9: execute d.m$	4: $\operatorname{lag} \leftarrow \sum_{q \in d. Queries} q. \operatorname{lag} / d. \operatorname{Queries. Count}$
$\begin{array}{ll} & \underset{\text{d.score}}{\text{lag}} \\ 6: & \underset{\text{d.score}}{\text{d.score}} \leftarrow \frac{\text{d.value}}{\text{d.m.cost}} \\ 7: & \underset{\text{for each Dataset d in } D.\text{SortedByDesc(d.score) do} \\ 8: & \underset{\text{if d.m.bottleneck.canMove() then}}{\text{9:}} \\ 9: & \underset{\text{execute d.m}}{\text{execute d.m}} \end{array}$	5: d.value $\leftarrow \sum_{q \in d.Queries} d.m.timeReduction[q] /$
6: $d.score \leftarrow \frac{d.value}{d.m.cost}$ 7: for each Dataset d in <i>D</i> .SortedByDesc(d.score) do 8: if d.m.bottleneck.canMove() then 9: execute d.m	lag
 7: for each Dataset d in D.SortedByDesc(d.score) do 8: if d.m.bottleneck.canMove() then 9: execute d.m 	6: $d.score \leftarrow \frac{d.value}{d.m.cost}$
8: if d.m.bottleneck.canMove() then 9: execute d.m	7: for each Dataset d in D.SortedByDesc(d.score) do
9: execute d.m	8: if d.m.bottleneck.canMove() then
A STATE STAT	9: execute d.m



Figure 8: WAN Bandwidth Usage knob, B. MinBW is the scheme that optimizes for WAN bandwidth usage. Even with same WAN usage as MinBW (B = 1), Iridium's gains in query response time are significantly higher. MinBW slows down queries against the in-place baseline.

Strengths

- Provide significant performance gains in a zone which was seemingly poorly optimized beforehand
- Avoids blindly optimizing one metric, and allows the user to tune their preferences (WAN knob)

- Very comprehensive and relevant evaluations, which effectively probe how the system works
- Intuition models were effective ways to communicate their ideas

Places for Improvement

 Greedy strategy seems to have clear negative effects on their current performance in some places

- Haven't properly modelled tasks which take a significant amount of time

- Assumes very large storage and compute capacities, which limit their relevant applications

Low-Latency Geo-Distributed Data Analytics

$\bullet \bullet \bullet$

Qifan Pu, Ganesh Ananthanarayanan, Peter Bodik, Srikanth Kandula, Aditya Akella, Paramvir Bahl, Ion Stoica

SIGCOMM '15 (Aug 17-21)

Nicholas Turner