

Chain:

Operator Scheduling For Memory Minimization In Data Stream Systems.

Brian Babcock, Shivnath Babu,
Mayur Datar, and Rajeev Motwani.

Presented by:
Kareem Elgebaly

ACM SIGMOD Int. Conf. on Management of Data, pages 253-264, 2003.

Outline:

Introduction.
Intuition.
Proof.
Experiments.
Conclusion.
Assessment.

2

Introduction:

Problem:

Data stream data rates are not only
fast but also irregular.

(2 orders of magnitude)

3

Introduction:

Solutions so far:

Drop tuples: (e.g. Load Shedding (Aurora,STREAM))
Loss of data leads to inaccuracies.

Overflow on disk:
Disastrous performance degradation.

4

Introduction:

Answer:

Reduce memory needed for queuing.

How ?

Through better scheduling.

Why ?

No penalties in performance or accuracy.

5

Introduction:

Devise a scheduler that discriminates among operators according to their memory impact.

Fast Operators:

Expected to have a very fast run-time.

Selective Operators:

Operators that consume a lot of records.

6

Intuition:

Two dimensional problem:

Fast + Selective \longrightarrow High Priority

Slow + Unselective \longrightarrow Low Priority

7

Intuition:

How many tuples per unit time does the Op consume?

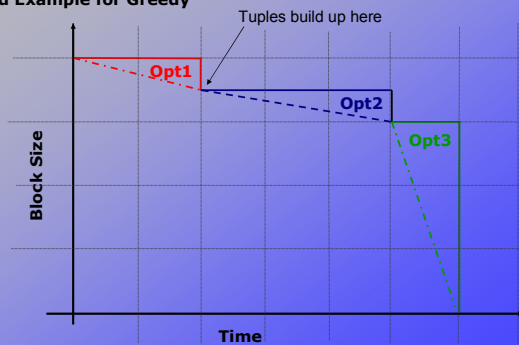
Greedy evaluation:

Priority \propto (Selectivity / Time)

8

Intuition:

Bad Example for Greedy



9

Intuition:

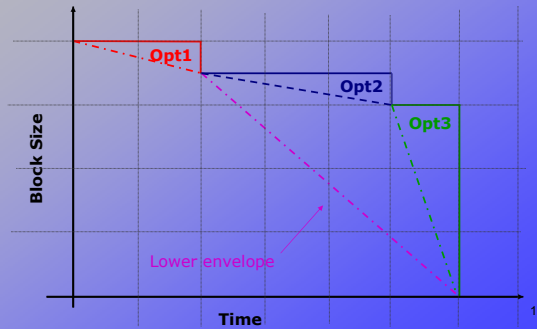
Answer is not straight forward...

A very good operator that takes results from a bad operator will never get scheduled.

(Local Minima)

10

Intuition:



11

Intuition:

Chain evaluation:

Priority α Lower Envelope Slope

12

Proof:

Claim:

Memory needed by Chain scheduling is within constant factor of optimal offline algorithm.

(Clairvoyant)

Proof sketch:

1. Greedy scheduling is optimal for convex progress charts
(since) Best operators are immediately available
2. Lower envelope is convex
3. Lower envelope closely approximates actual progress chart


13

Proof:

Claim:

Lower envelope closely approximates actual progress chart

1. At most one block in the middle of each lower envelope segment
(Due to) tie-breaking rule
2. (Lower envelope + 1) gives upper bound on actual memory usage
3. Additive error of 1 block per progress chart



14

Experiments:

Setup:

Data Sets:


- 1) Synthetic Data Set.
- 2) Real Data Set.

Queries:

- 1) Single Queries.
- 2) Multi Queries.
- 3) Join Queries.


15

Experiments:




16

Experiments:



17

Experiments:



18

Conclusions:

	Pros.	Cons.
FIFO	No starvation	Poor performance
Round-Robin	No starvation	Poor performance
Greedy	Good performance	Starvation Stuck in local maxima
Chain	Near optimal performance	Starvation

19

Conclusions:

Chain is orthogonal to traditional memory requirements minimization techniques. Hence you are not trading Chain's benefits with anything, you are getting it for free.

Chain is an algorithm that guaranties certain performance standard without introducing any extra over heads.

Good !!!

20

Assessment:

Limitation:

Chain is no better than FIFO in normal rates.
Chain is most useful when rates are irregular.

(plus) no experimentation done to compare performance in such cases.

What if:

The SDMS was implementing an Early Selection optimization technique ?

Would Chain make sense?

Would it be any better than greedy?

21

Assessment:

TODO:

More QOS guarantees. Like (low response time)
tuples may wait for an unacceptable long time before it gets scheduled.

Chain doesn't take into account:

- 1) Parallelism
- 2) Shared sub plans (shared queues)

22

Questions?

23