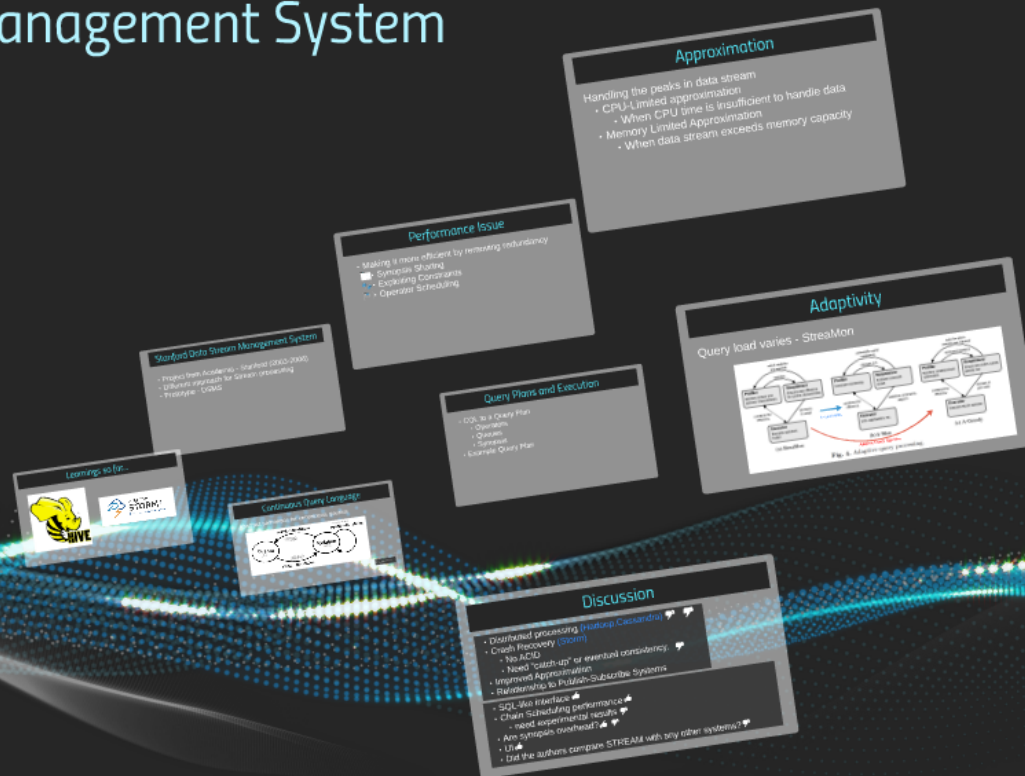


STREAM

Stanford Data Stream Management System

Pranav Muktali



STREAM

Stanford Data Stream Management System

Pranav Muktali

Learning to fly

- Project from Accelerate - Stanford (2008-2010)
- Ultimate approach for serverless computing
- Prototype: COWS

Stanford Data Stream Management System

- Project from Accelerate - Stanford (2008-2010)
- Ultimate approach for serverless computing
- Prototype: COWS

Performance Issue

- Making it more efficient by removing redundancy
- Synopsis Sharing
- Evolving Control plans
- Operator scheduling

Approximation

Handling the peaks in data stream

- CPU-Limited approximation
- When CPU time is insufficient to handle data
- Memory Limited Approximation
- When data stream exceeds memory capacity

Adaptivity


Query load varies - StreamMon

Query Plans and Execution


- Q1: Is a Query Plan
- Operators
- Control
- Example Query Plan

Discussion

- Distributed processing (Dariusz Casagrande)
- Crash Recovery (Lorenz)
- Is ACID
- Need "catch-up" or eventual consistency
- Implement Approximation
- Relationships to Publish-Subscribe Systems
- SQL-like Interface
- Chain Scheduling performance
- Need experimental results
- Are synthetic overhead?
- UI
- Did the authors compare STREAM with any other systems?



LAST YEAR WE
RECOGNIZED THAT OUR
PROCESSES WERE FAR
TOO COMPLEX



LET THE CLOUDS MAKE YOUR LIFE EASIER

Learnings so far...



Hive

- Abstraction over Hadoop.
- Supports SQL-like queries.
- Runs Map-Reduces jobs from declarative statements.

Here comes continuous stream of data!

Apache Storm

- Stream processing system
- Spouts and Bolt
- A record is a tuple of $\langle \text{key}, \text{value(s)} \rangle$

What about to SQL-like query?

Stanford Data Stream Management System


- Project from Academia - Stanford (2003-2006)
- Different approach for Stream processing
- Prototype - DSMS

Data Stream Management System

- CQL: Continuous Query Language
- CQL Under the hood
- Optimize
- Handle high volume
- User Interface to query and visualize

Continuous Query Language

Abstract semantics for continuous queries



Examples

```
Select (stream1) From S [Rows Unbounded] Where S.A > 10  
Select * From S1 [Rows 1000], S2 [Range 2 Minutes]  
Where S1.A = S2.A And S1.A > 10
```

L L C A


- Unbound bag (multiset) of pairs $\langle s, t \rangle$
- s is a tuple in stream S
- t is a timestamp
- t denotes logical time of arrival of s
- Leslie Lamport: happened-before

ETALU

- R is a time-varying bag of tuple
- $R(t)$ denotes bag of tuples at time t

Continuous Query Language

Abstract semantics for continuous queries



Examples

```
Select (stream1) From S [Rows Unbounded] Where S.A > 10  
Select * From S1 [Rows 1000], S2 [Range 2 Minutes]  
Where S1.A = S2.A And S1.A > 10
```

queries

Uses SQL construct to expressions

relation-to-relation



- tuple-based sliding window
- time-based sliding window
- partitioned sliding window

stream-to-relation

tuple-based

- Input:
 - int $N > 0$ at t
- Output:
 - $R(t)$, N tuples with timestamp $\leq t$
- Usage
 - [Rows N]
 - [Rows Unbounded]

time-based

- Input:
 - timestamp w
- Output:
 - $R(t)$, tuples between $t-w$ and t
- [Range w]
- [Now] ($w=0$)

partitioned

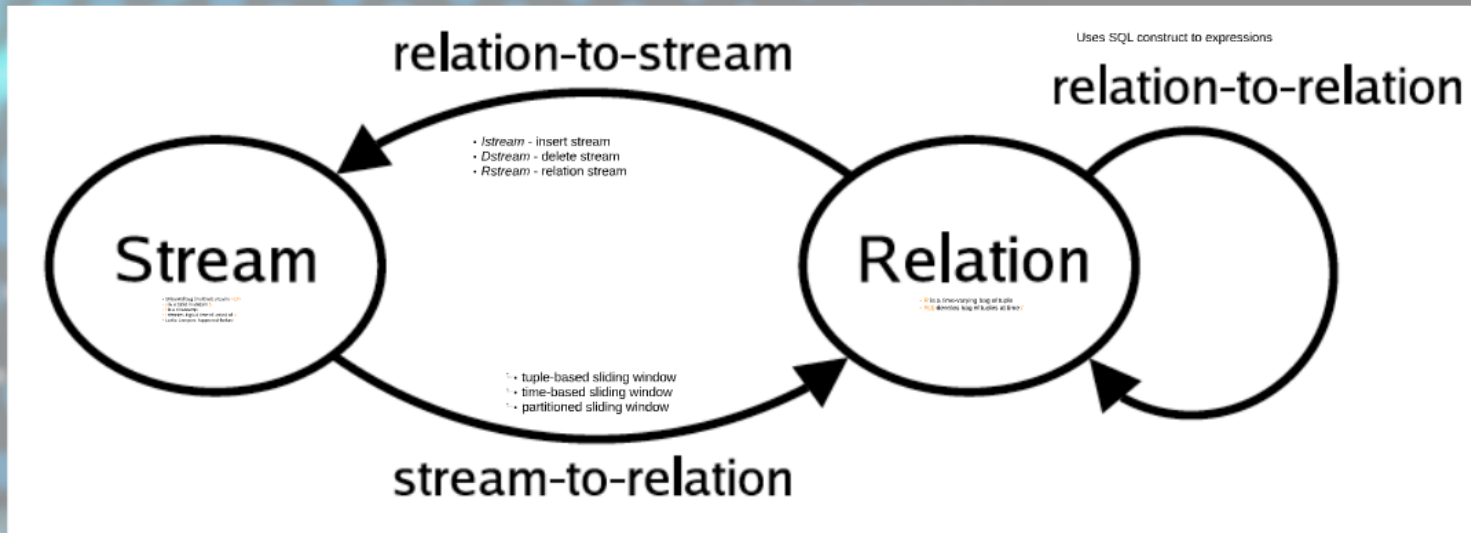
- Input:
 - int N and set $\{A_1, \dots, A_k\}$
- Output:
 - $R(t)$, union of substreams by attr of size N
- [Partition By A_1, \dots, A_k Rows N]

relation-to-stream

- *Istream* - insert stream
- *Dstream* - delete stream
- *Rstream* - relation stream

Continuous Query Language

Abstract semantics for continuous queries



Examples

```
Select (stream1) From S [Rows Unbounded] Where S.A > 10  
Select * From S1 [Rows 1000], S2 [Range 2 Minutes]  
Where S1.A = S2.A And S1.A > 10
```

Examples

Select Istream(*) From S [Rows Unbounded] Where S.A > 10

r-to-s

s-to-r

r-to-r

Select * From S1 [Rows 1000], S2 [Range 2 Minutes]
Where S1.A = S2.A And S1.A > 10

Istream(*)

r-to-s



[Rows Unbounded]

s-to-r

10000] S2 [Range 2



S.A > 10

r-to-r

Examples

Select Istream(*) From S [Rows Unbounded] Where S.A > 10

r-to-s

s-to-r

r-to-r

Select * From S1 [Rows 1000], S2 [Range 2 Minutes]
Where S1.A = S2.A And S1.A > 10

Query Plans and Execution

- CQL to a Query Plan
 - Operators
 - Queues
 - Synopsis
- Example Query Plan


Name	Operator Type	Description
<code>select</code>	relation-to-relation	Filters elements based on predicate(s)
<code>project</code>	relation-to-relation	Duplicate-preserving projection
<code>binary-join</code>	relation-to-relation	Joins two input relations
<code>mjoin</code>	relation-to-relation	Multiway join from [22]
<code>union</code>	relation-to-relation	Bag union
<code>except</code>	relation-to-relation	Bag difference
<code>intersect</code>	relation-to-relation	Bag intersection
<code>antijoin</code>	relation-to-relation	Antijoin of two input relations
<code>aggregate</code>	relation-to-relation	Performs grouping and aggregation
<code>deduplicate</code>	relation-to-relation	Performs duplicate elimination
<code>seq-window</code>	stream-to-relation	Implements time-based, tuple-based, and partitioned windows
<code>i-stream</code>	relation-to-stream	Implements <i>Istream</i> semantics
<code>d-stream</code>	relation-to-stream	Implements <i>Dstream</i> semantics
<code>r-stream</code>	relation-to-stream	Implements <i>Rstream</i> semantics

Table 1. Operators used in STREAM query plans.

Query Plans and Execution

- CQL to a Query Plan
 - Operators
 - Queues
 - Synopsis
- Example Query Plan


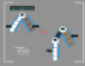

Example



Select * From S1 [Rows 1000], S2 [Range 2 Minutes]
Where S1.A = S2.A And S1.A > 10



Performance Issue

- Making it more efficient by removing redundancy
 -  • Synopsis Sharing
 -  • Exploiting Constraints
 -  • Operator Scheduling

Synopsis Sharing

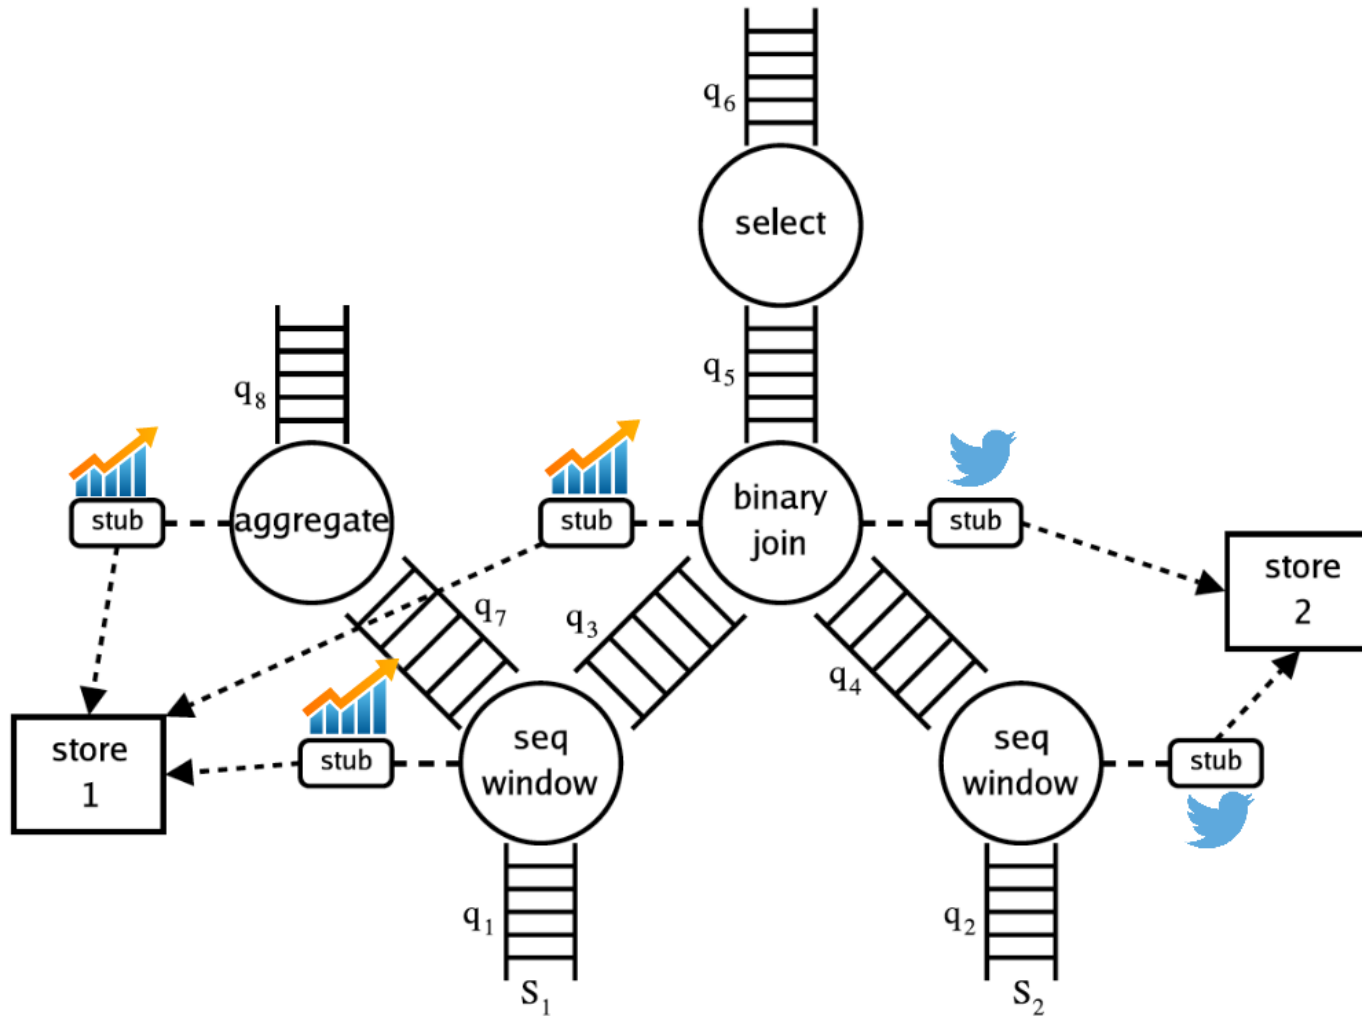
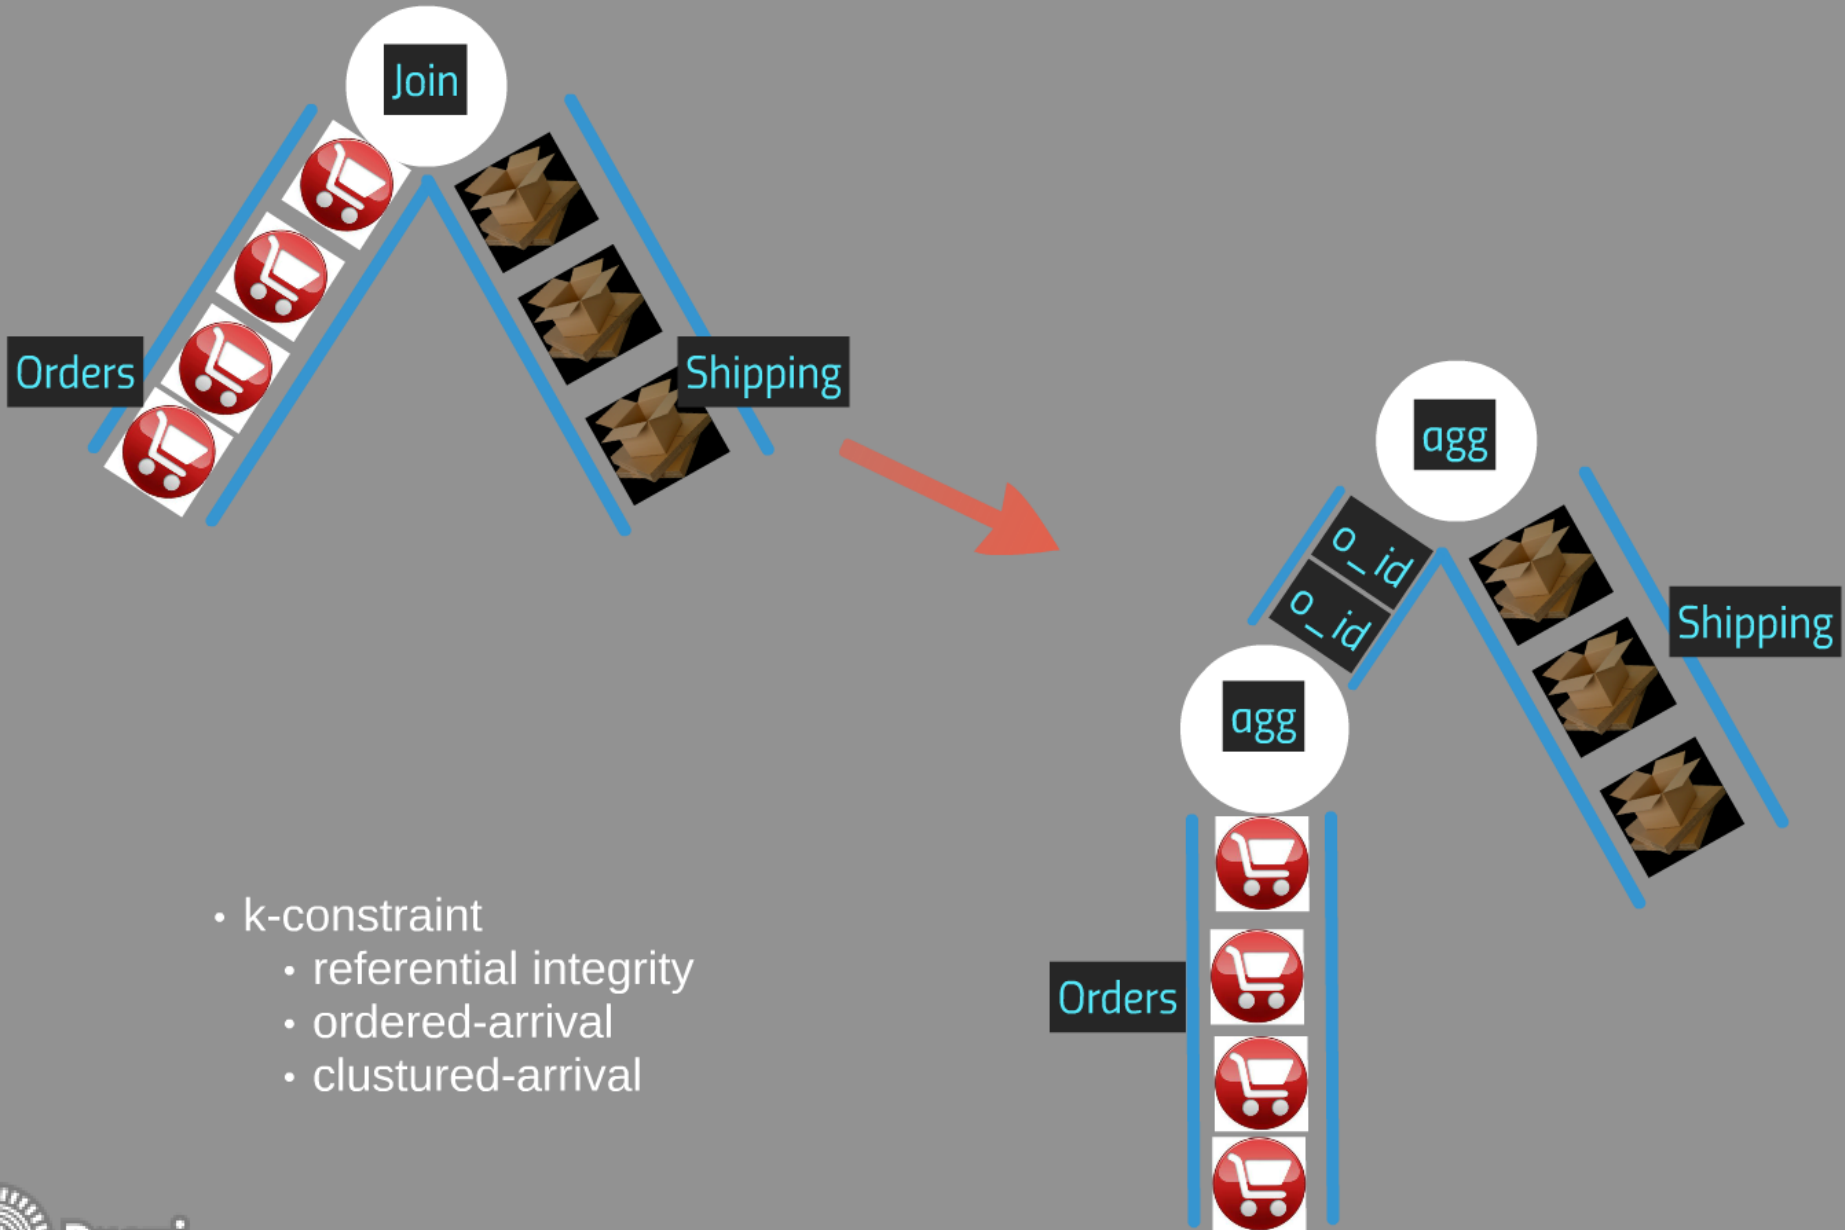


Fig. 3. A query plan illustrating synopsis sharing.

Exploiting Constraints



Operator Scheduling

Possible schedulers

- FIFO



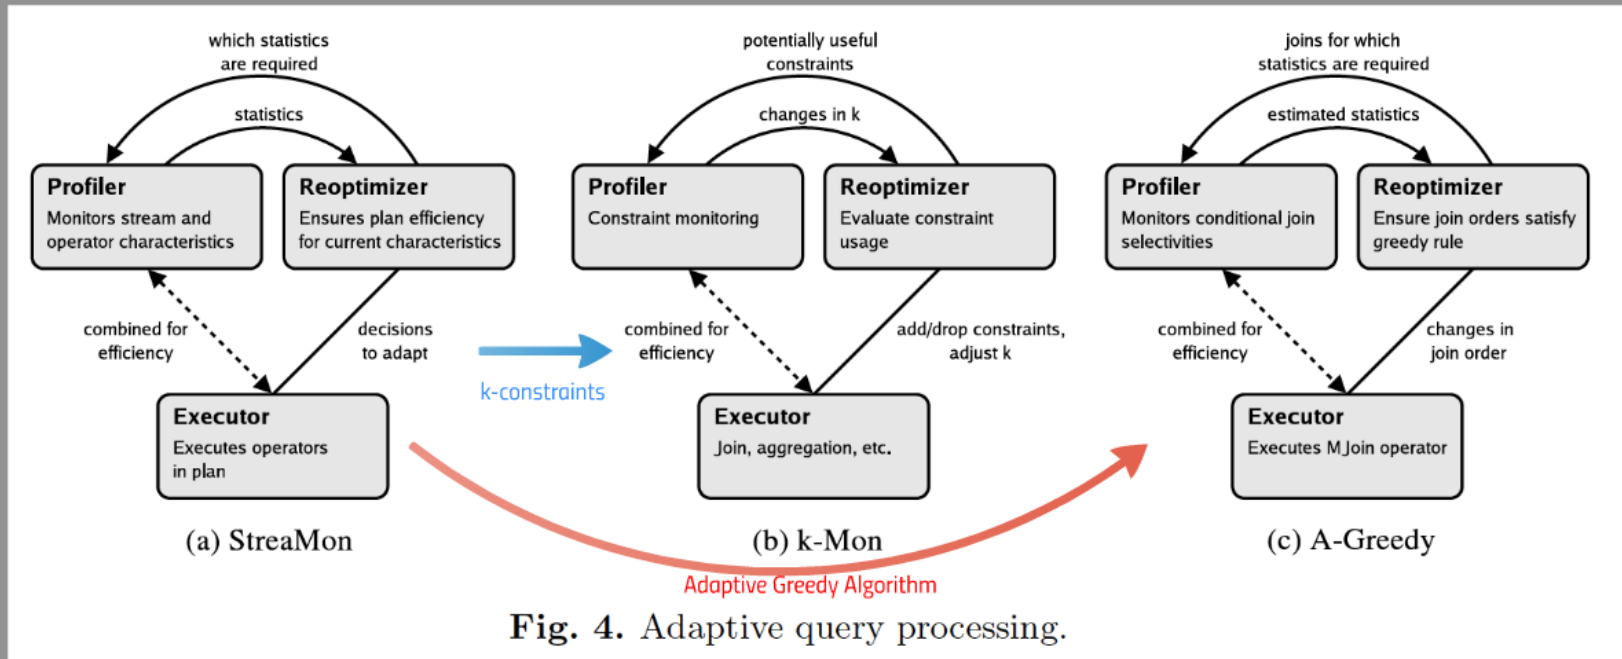
- Greedy

Introducing Chain Scheduling

- Greedy chain creation
- FIFO operations in chain

Adaptivity

Query load varies - StreamMon



Approximation

Handling the peaks in data stream

- CPU-Limited approximation
 - When CPU time is insufficient to handle data
- Memory Limited Approximation
 - When data stream exceeds memory capacity

Discussion

- Distributed processing ([Hadoop](#), [Cassandra](#)) 👎 👎
 - Crash Recovery ([Storm](#))
 - No ACID
 - Need "catch-up" or eventual consistency. 👎
 - Improved Approximation
 - Relationship to Publish-Subscribe Systems
-
- SQL-like interface 👍
 - Chain Scheduling performance 👍
 - need experimental results 👎
 - Are synopsis overhead? 👍 👎
 - UI 👍
 - Did the authors compare STREAM with any other systems? 👎

STREAM

Stanford Data Stream Management System

Pranav Muktali

Learning to fly

- Project from Accelerator - Stanford (2000-2005)
- Ultimate approach for server processing
- Prototype: C++

Stanford Data Stream Management System

- Project from Accelerator - Stanford (2000-2005)
- Ultimate approach for server processing
- Prototype: C++

Performance Issue

- Making it more efficient by removing redundancy
- Synopsis Sharing
- Evolving Control plans
- Operator scheduling

Approximation

Handling the peaks in data stream

- CPU-Limited approximation
- When CPU time is insufficient to handle data
- Memory Limited Approximation
- When data stream exceeds memory capacity

Adaptivity

Query load varies - StreamMon

Query Plans and Execution

- Q1: Is a Query Plan
- Operators
- Streams
- Example Query Plan

Discussion

- Distributed processing (Dariusz Casagrande)
- Crash Recovery (Lorenz)
- Is ACID
- Need "catch-up" or eventual consistency
- Implement Approximation
- Relationships to Publish-Subscribe Systems
- SQL-like Interface
- Chain Scheduling performance
- Need experimental results
- Are synthetic overhead?
- UI
- Did the authors compare STREAM with any other systems?