Load Shedding in a Data Stream Manager

Nesime Tatbul, Uğur Çetintemel, Stan Zdonik Brown University

Mitch Cherniack Brandeis University Michael Stonebraker M.I.T.

The Overload Problem



"Load Shedding": eliminating excess load by <u>dropping data</u>

Aurora Data Stream Management System

Aurora Query Network



- Stream Query Algebra (SQuAl)
- Quality of Service

- Run-time Operation:
 - Operator Scheduling
 - Storage Management
 - Performance Monitoring
 - Load Shedding

Load Shedding by Inserting Drops





Quality of Service





Assumptions

Processor is the main scarce resource
Operators that produce new attribute values are not considered (e.g., aggregate, map)
Load Shedder operates independently from

Scheduler

Problem Statement

N: query networkC : processing capacityI : set of input streams

when Load(N(I)) > C, transform N to N' such that
 Load(N'(I)) < C</pre>

Utility(N(I))-Utility(N'(I)) is minimized

key questions:

when to shed load?
where to shed load?
how much load to shed?
which tuples to drop?
VLDB 2003 Berlin

Talk Outline

Introduction
Technical Overview
Experimental Results
Related Work
Conclusions and Future Work

Algorithmic Overview

evaluate the Query Network Load

IF Load > Capacity
 insert Drops to the Query Network

ELSE IF Load < Capacity AND Drops in the Query Network remove Drops from the Query Network

read stats, network modify network

System Catalog

Queries Statistics

Drop Insertion Plans

look up

Load Evaluation

"Load Coefficients" for each input stream



 $\sum_{i=1}^{m} L_i \times R_i$

(CPU cycles per time unit)

Load Shedding Road Map (LSRM)



Constructing the LSRM

- 1. identify potential "drop locations"
- 2. sort the drop locations
- 3. apply the drop locations in order:
 - *insert* one unit of drop
 - if semantic drop, determine the filter predicate
 - create an LSRM entry

Drop Locations



drops before splits cause more utility loss

Best Drop Location

Goal: maximize cycle gain, minimize utility loss



From Values to % Delivery "when, where, how much?" have the same answer for Random and Semantic Load Shedding Trick: translation between QoS functions Assumption: drop the lowest utility values first



Determining the Filter Predicate



Talk Outline

Introduction
Technical Overview
Experimental Results
Related Work
Conclusions and Future Work

Experimental Study

Setup:

- streams: uniformly distributed integer values
- networks: a mix of filters and unions
- QoS: value-based QoS, range utilities chosen from a Zipf distribution
- Metrics used:
 - percent loss in total average utility on
 loss-tolerance QoS (*Tuple Utility Loss*)
 value-based QoS (*Value Utility Loss*)

Random-LS beats Admission Control



Semantic-LS beats Admission Control



Semantic-LS beats Random-LS



Performance for Networks with Sharing



Talk Outline

Introduction
Technical Overview
Experimental Results
Related Work
Conclusions and Future Work

Relevance to Existing Work

- Congestion control in networks
- Multimedia streaming
- Approximate query answering
- Data stream processing
 - sampling, shedding on aggregates STREAM [MW+03, BDM03]
 - approximate join processing [DGR03]
 - adjusting rates to windowed joins [KNV03]

Talk Outline

Introduction
Technical Overview
Experimental Results
Related Work
Conclusions and Future Work

Highlights

An end-to-end solution for detecting and resolving overload in data stream systems
 Utility loss minimization for networks of queries which share resources and processing

 → semantic utility as well as quantity-based utility

 Static drop insertion plans, dynamic instantiation

Future Directions

Handling complex operators

joins and aggregates

Other resource limitations

memory - windowed operators
bandwidth - Aurora*
power - at sensor level

More Information

Aurora Web Page:

http://www.cs.brown.edu/research/aurora

Email:

tatbul@cs.brown.edu

Demo