

Semantics of Data Streams and Operators

David Maier, Jin Li, Kristin Tufte,
Vassilis Papadimos
Portland State University
Pete Tucker
Whitworth College

16 February 2005

Stream Semantics

1

What do you say in an invited talk?

- Retrospective
- Tutorial
- All the stuff you've been mulling over but haven't quite figured out yet
 - *Maybe someone here will get it better*

Stream representation, reconstruction functions,
monotonicity, punctuation, disorder, windows

16 February 2005

Stream Semantics

2

What does a stream mean?

- A recent paper define a data streams as
"bags of append-only of ordered tuples"
or with the addition of timestamps, as
"an unbounded appended-only bags of elements"
- Lots of questions
 - an ordered bag?
 - what constitutes a duplicate?
 - can a stream be finite?
 - is it being updated?

Stream Denotation and Representation

- Denotation (what it means)
- Representation (encoding)
- Can have multiple representations for the same denotation

Example: Sequence of finite relation states

$r1(A\ B)$	$r2(A\ B)$	$r3(A\ B)$...
1 2	1 2	1 3	
1 3	1 3	2 4	
	2 4	2 5	

- **serial**

$\langle 1\ 2 \rangle, \langle 1\ 3 \rangle, \text{EOR}, \langle 1\ 2 \rangle, \langle 1\ 3 \rangle, \langle 2\ 4 \rangle, \text{EOR},$
 $\langle 1\ 3 \rangle, \langle 2\ 4 \rangle, \langle 2\ 5 \rangle, \text{EOR}, \dots$


- **tagged**

$(\langle 1\ 2 \rangle:1), (\langle 1\ 3 \rangle:2), (\langle 1\ 3 \rangle:1), (\langle 1\ 2 \rangle:2),$
 $(\langle 1\ 3 \rangle:3), (\langle 2\ 4 \rangle:2), (\langle 2\ 4 \rangle:3), (\langle 2\ 5 \rangle:3),$
...

- **delta**

$+\langle 1\ 2 \rangle, +\langle 1\ 3 \rangle, \text{EOR}, +\langle 2\ 4 \rangle, \text{EOR},$
 $-\langle 1\ 2 \rangle, +\langle 2\ 5 \rangle, \text{EOR}, \dots$

- **windowed**

$\langle 1\ 2 \rangle \langle 1\ 3 \rangle \langle 2\ 4 \rangle \langle 2\ 5 \rangle \dots$


Different characteristics

- Space efficiency
- States can be interleaved with tagged (but hard to know when a state is complete)
- Operator implementation, e.g., $\sigma_{A+1=B}$
 - Can apply elementwise on serial, tagged, delta
 - Not on windowed


16 February 2005

Stream Semantics

7

Example: Sensor readings

- Reg. intervals: 23, 23, 22, 21, 20, 22, 22, 23, ...
- Change only: $\langle 23 \ 5s \rangle$, $\langle 22 \ 15s \rangle$, $\langle 21 \ 20s \rangle$, $\langle 20 \ 25s \rangle$, $\langle 22, \ 30s \rangle$, $\langle 23, \ 40s \rangle$, ...
Hard to tell no data vs. stall



16 February 2005



Stream Semantics

8

Can interpret same representation different ways

- Stream of stock quotes

<ticker, time, price>



16 February 2005

Stream Semantics

9

Reconstitution functions

- Want to treat streams incrementally
- Use a *reconstitution function* to give successive approximations of stream denotation from successive prefixes

$S \in \text{strm}(\mathbf{T})$ [finite or infinite]

P a prefix of S has type $\text{seq}(\mathbf{T})$

\mathbf{D} is domain of interpretation

$\text{reconst}(P) = d \in \mathbf{D}$

[1, 5, 6, 5, 7, 2, ...]

16 February 2005

Stream Semantics

10

Example: Insert (bag)

- $\text{ins}([]) = \{\}$
- $\text{ins}(P:i) = \text{insert}(i, \text{ins}(P))$

$\text{ins}([1, 5, 6, 5]) = \{\{1, 5, 5, 6\}\}$

Example: Insert unique

- $\text{ins-u}([]) = \emptyset$
- $\text{ins-u}(P:i) = \text{if } i \in \text{ins-u}(P) \text{ then } \text{ins-u}(P) \text{ else } \text{insert}(i, \text{ins-u}(P))$

$\text{ins-u}([1, 5, 6, 5]) = \{1, 5, 6\}$

Example: Insert-collect

- $\text{ins-c}([]) = []$
- $\text{ins-c}(P:i) = \text{ins-c}(P):\text{ins-u}(P:i)$

$\text{ins}([1, 5, 6, 5]) =$
 $[\{1\}, \{1,5\}, \{1,5,6\}, \{1,5,6\}]$


Remarks

- These examples are all incremental, but reconst function need not be
- Element types of input and output structures need not be the same
 \mathbf{T} vs. $\pm\mathbf{T}$
- if $\text{reconst}(P) = d$, then write
 $\text{const}(d) = P$
(not really a function)


Faithfulness of stream versions of operators

- When is a stream operator a reasonable analogue of its domain counterpart?


For every prefix P_1 of input stream S



Monotone/Non-blocking/Reconstitution

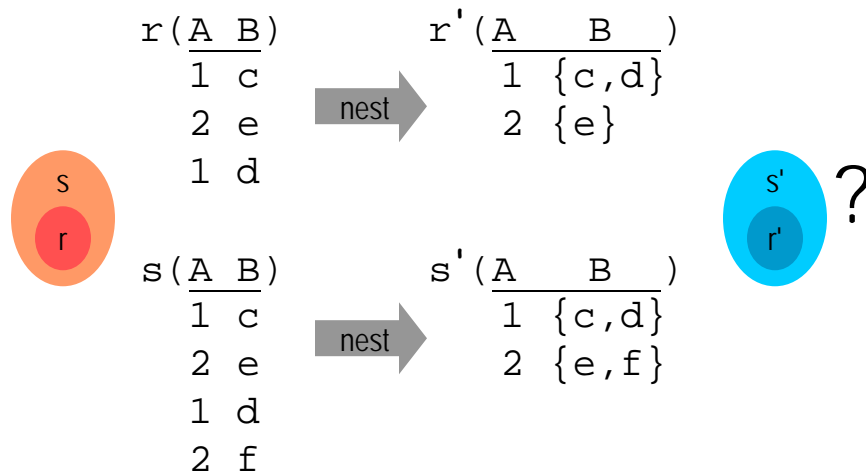


Monotone relational ops carry over easily to stream counterparts (under ins)



$$cr = \text{const}(r)$$

What about hierarchical data?



16 February 2005

Stream Semantics

17

Depends on the ordering

- s' *superset* r' : not monotone
- s' *subsumes* r' : monotone

Can we get



Not if ins is the reconstitution function

What works?

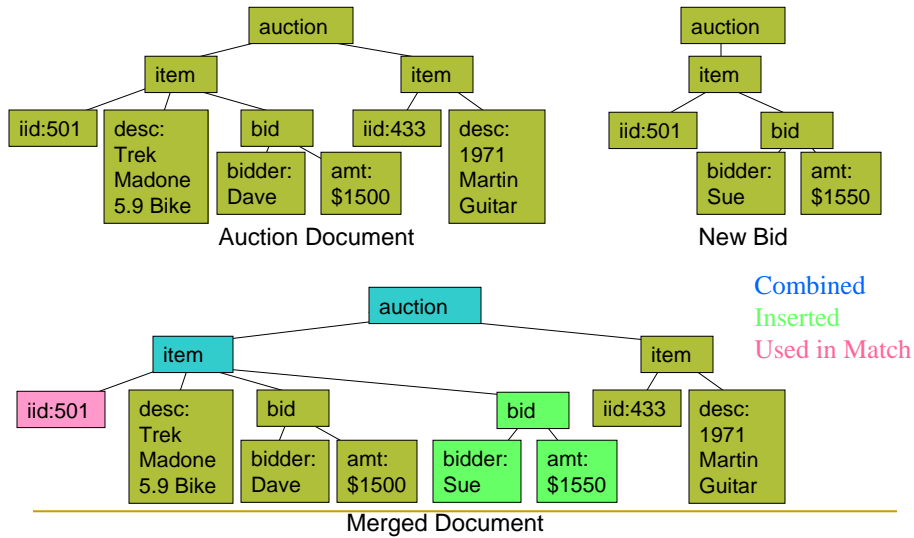
- insert-replace - treat A as key
 $\text{ins-r}(r', \langle 2, \{e,f\} \rangle) = s'$
- merge (recursive union on XML)

16 February 2005

Stream Semantics

18

Merge is like Deep Union (Penn)

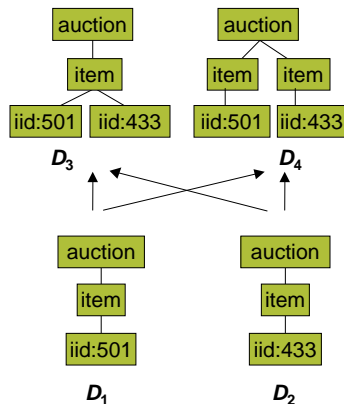


16 February 2005

Stream Semantics

19

Issue: Non-unique upper bounds



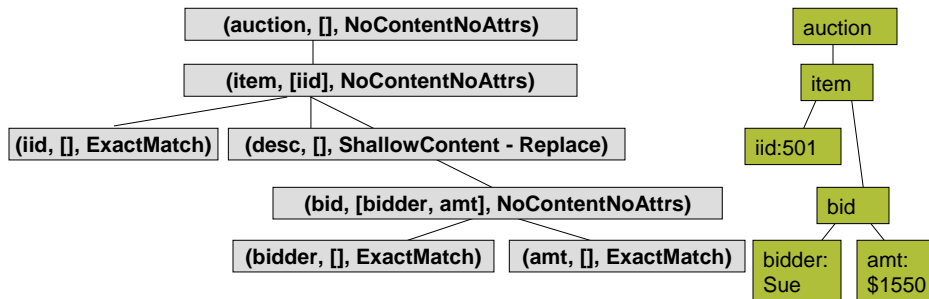
- Would like Merge to be LUB in subsumption order
- Keys in Merge Template eliminate ambiguity
- Know D_4 is correct result if we know iid is a key for item

16 February 2005

Stream Semantics

20

Merge Template removes ambiguity



- Merge Template is an XML document consisting of a tree of Element Merge Templates (EMT)
- EMT is a triplet containing: <name, local key, content-combine function>

16 February 2005

Stream Semantics

21

Additional stream semantics

May have useful prior knowledge about stream content

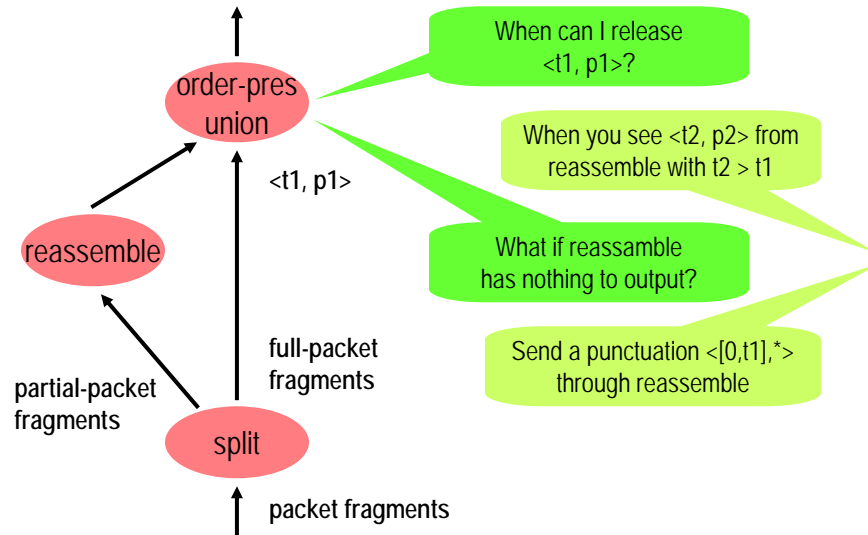
- $|\text{reconst}(P)| < k$
- B-sorted on attribute A with slack k [Aurora]
- Limits on remainder of stream (punctuation)
Everything matching pattern p has already been seen

16 February 2005

Stream Semantics

22

Example from Gigascope



16 February 2005

Stream Semantics

23

Disorder

Many reasons for physical stream to be disordered

- ❑ different data paths
- ❑ combining sources
- ❑ more than one order
start, end time of a netflow record

16 February 2005

Stream Semantics

24

Example of disorder

- flow start, ordered by 8th packet observation time




Figure 6: Band Disorder

16 February 2005

25

How to deal with disorder?

Allow for a maximum displacement: *slack*

Slack of k_2 catches all points; slack of k_1 introduces a little inaccuracy, but improves latency.




Figure 6: Band Disorder

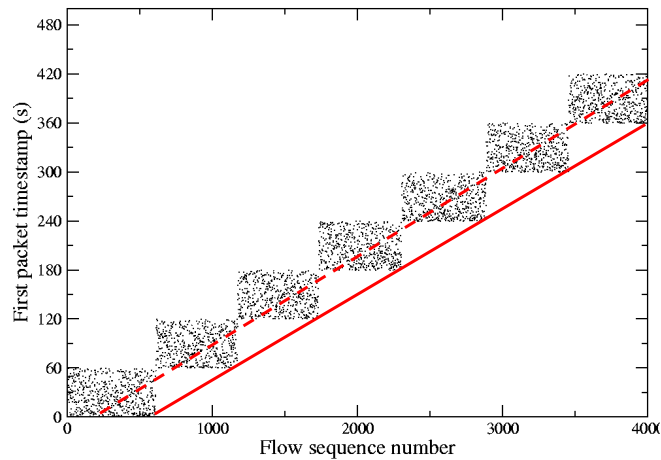
16 February 2005

26

What about this?

Netflow records from a router; start time by output order

Much inaccuracy with anything other than maximal slack.

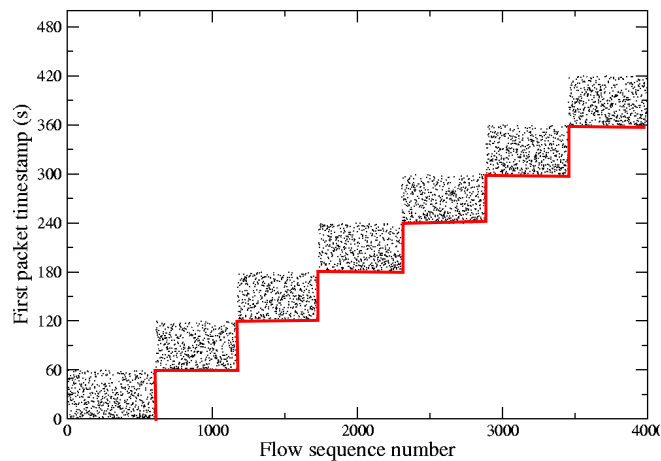


16 February 2005

Problem: constant slack

Punctuation can give "variable slack"

Data source likely knows block boundaries.



16 February 2005

Shield your eyes: Performance graph




Figure 8: Latency vs. Accuracy Band-Disorder
(average error percentage)




Figure 9: Latency vs. Accuracy Block-Sorted-Disorder
(average error percentage)

16 February 2005

Stream Semantics

29

Windowed operators

- General idea easy to get; details are dicey
- Useful for
 - Blocking operators (e.g., aggregate)
 - Unbounded stateful operators (e.g., join)
- Divide stream into finite substreams, execute operation over each, say max




16 February 2005

Stream Semantics

30

Can overlap: Sliding window



- Often described operationally on physical order
- Causes problems when stream is disordered


Specify windows independent of physical presentation

Consider a window specification

[Range 30, Slide 10] on A

Use explicit window IDs


{0, 1, 2, ... }



Define extent of each window

Gives stream items for each window ID w


$$\text{extent}(w) = \{i \mid w \cdot 10 \leq i.A \leq w \cdot 10 + 30\}$$



Useful to have dual function

Gives window IDs for each item i

$$\text{wids}(i) = \{w \mid \lceil i.A/10 \rceil - 1 < w < \lceil (i.A+30)/10 \rceil - 1\}$$



Can be basis for implementation

Tag data items with window IDs; treat ID as a regular attribute

Need a means, such as punctuation, to know when a window is complete




Figure 8 (b): Execution Time: WID versus Buffering – Zoom-in

Conclusion: What would make me like you

- Explain to me what I'm doing with reconstitution functions - trying to reason about the infinite with finite fragments?
- Understand the monotone/non-blocking/reconstitution connection better.
- Automate the use of punctuation in stream query processing - for example, when can we prove space bounds?

- Find an effective way to describe disorder;
effective =
 - can determine for a stream
 - can synthesize
 - can calculate latency vs. accuracy trade-offs
- Use properties of window specifications to determine properties of windowed operator implementations.

Questions

Niagara West:

<http://www.cse.ogi.edu/dot/niagara>

Acknowledgements:

- Niagara Midwest (Wisconsin)
- Tim Sheard, Leo Fegaras, Ted Johnson
- Abilene Observatory
- NSF grant IIS 00-86002