S-Store: A Streaming NewSQL System for Big Velocity Applications

U. Cetintemel, J. Du, T. Kraska, S. Madden, D. Maier, J. Meehan, A. Pavlo, M. Stonebraker, E.Sutherland, N. Tatbul, K. Tufte, H.Wang, S. Zdonik

Motivation and Contributions

State of the art for big-velocity applications:

- Stream processing: No/weak transactional guarantees
- OLTP systems: No support for data-driven processing

Features for transaction processing over big-velocity data

+ Main-memory OLTP foundation (SQL-based and ACID)
+ Data-driven processing
+ Fault-tolerance: command logging & upstream backup
+ Uniform state management

Added constructs: Streams, Windows, Triggers, Workflows

Streaming Transactions

- Streaming Transaction = stored procedure + batch of input tuples
- Workflow is a DAG of SPs
- Due to **data and processing dependencies**:
  - Transaction Executions (TEs) can't be arbitrarily ordered
  - **Window state** can only be accessed by TEs of a given SP
- Presence of **shared tables** and OLTP transactions require additional isolation rules

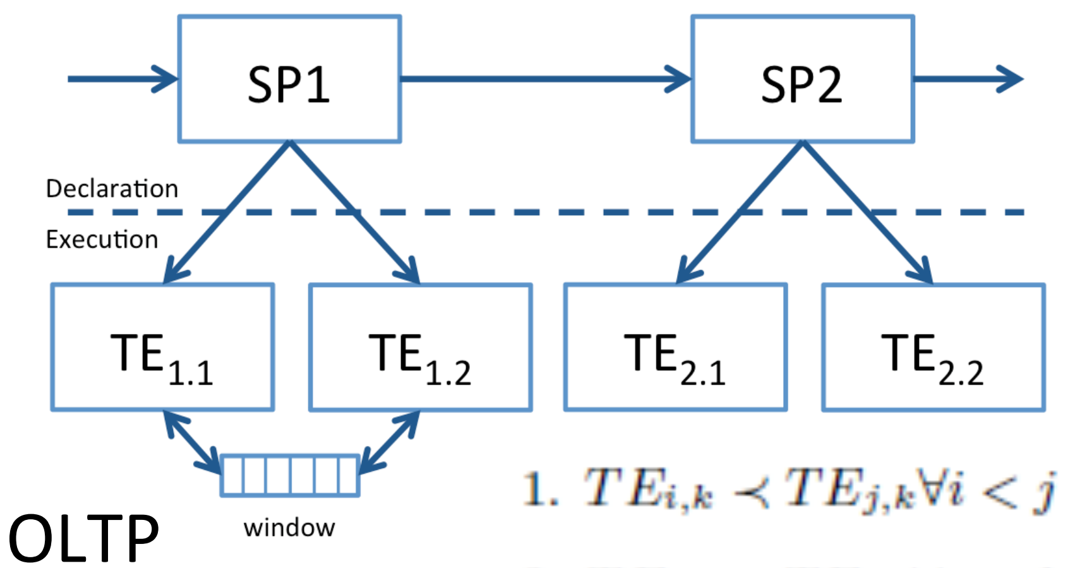

1. $TE_{i,k} \prec TE_{j,k} \forall i < j$
2. $TE_{i,j} \prec TE_{i,k} \forall j < k$

Application: Canadian Dreamboat

The Canadian pop star popularity contest!

- Viewers submit votes via text (one vote per viewer)
- Leaderboard maintained for:
  - Top three contestants
  - Bottom three contestants
  - "Trending" three (in last 100 votes)
- Every 1,000 votes, the lowest contestant is removed
- Viewers who voted for a removed contestant may resubmit their vote for a remaining candidate
- Voting continues until a single winner is decided

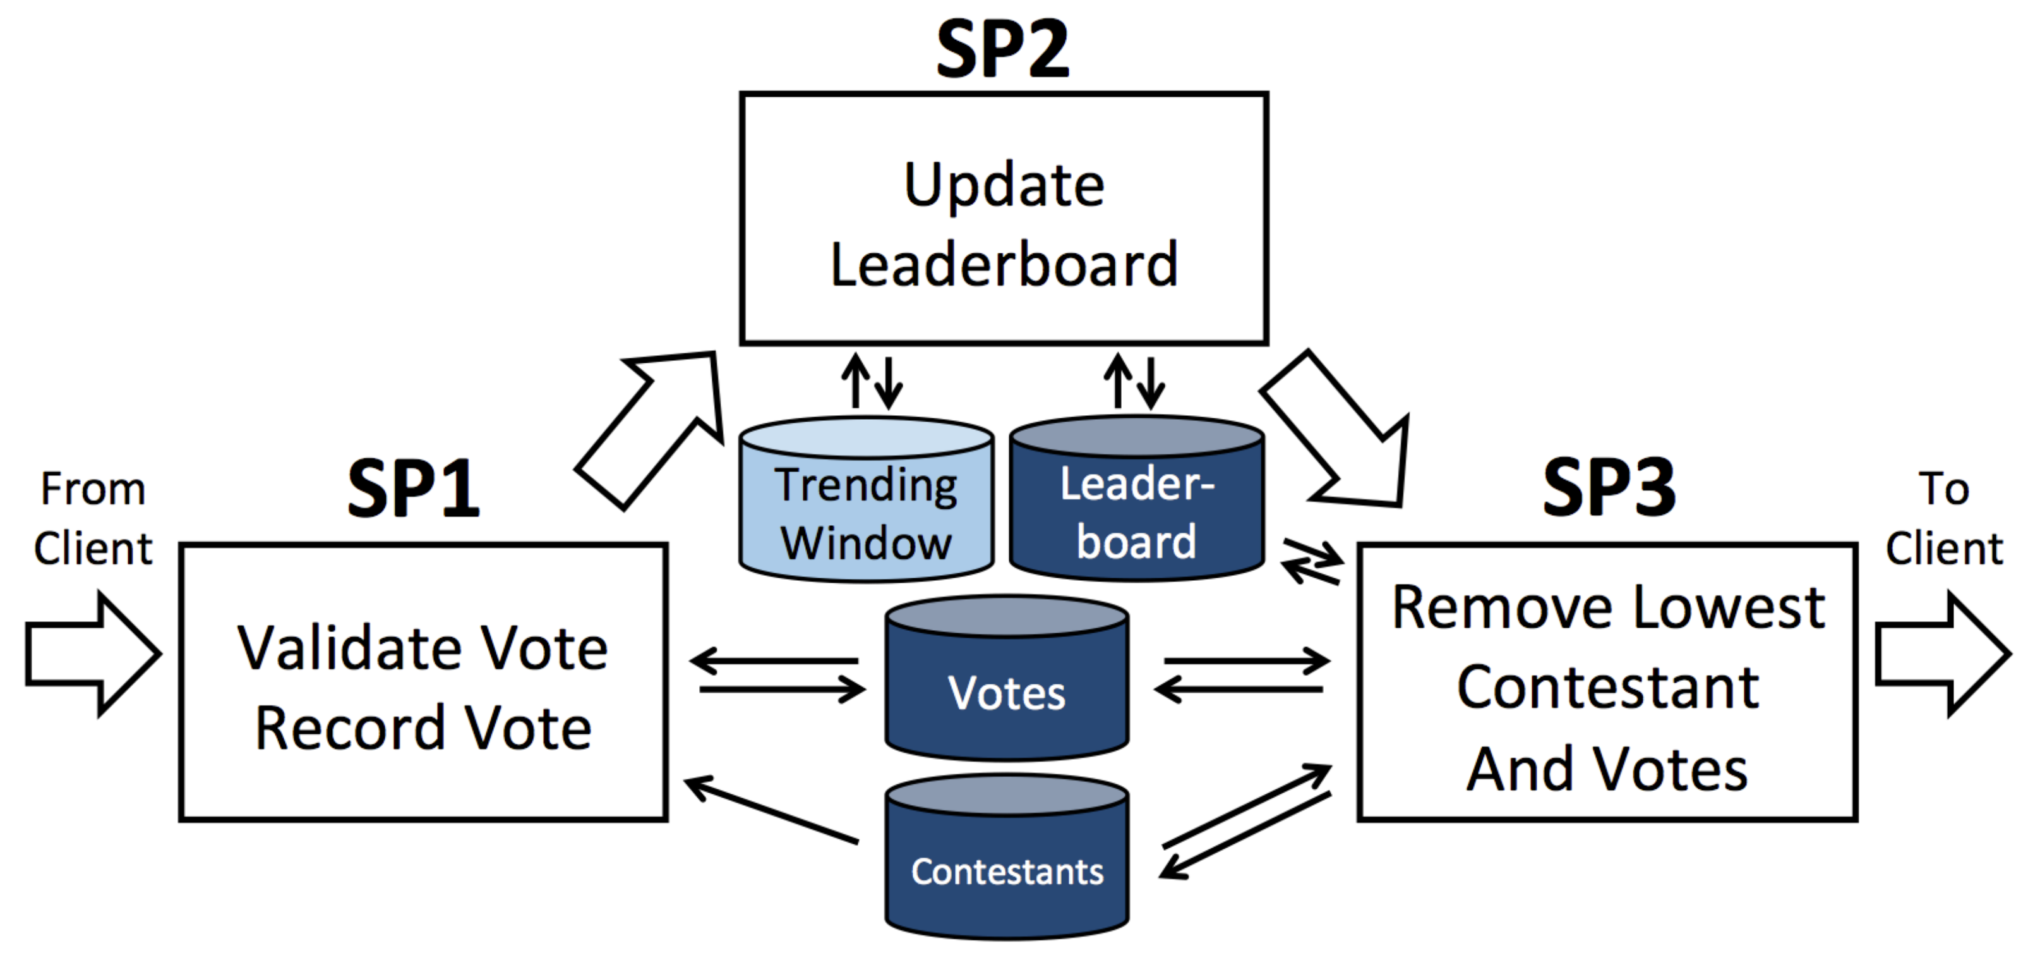

Workflow separated into three stored procedures

- Each SP comprises an ACID transaction
- State stored in tables, accessed by multiple SPs

Voting System Comparison

H-Store: Transactions, no workflow definitions

- SP execution can be arbitrarily ordered
- Additional votes submitted before the deletion occurs can lead to incorrect candidate deletion

Place	Contestant	Votes
12	Céline Dion	64
11	Nelly Furtado	160
10	Michael Bublé	261
9	Jim Bryson	364
8	Micah Barnes	494
7	Ashley Leggat	634
6	Emily Haines	815
5	Adam Gontier	1023
4	Avril Lavigne	1272
3	Leonard Cohen	1605
2	Justin Bieber	2076
1	Jann Arden	2232

Final expected results

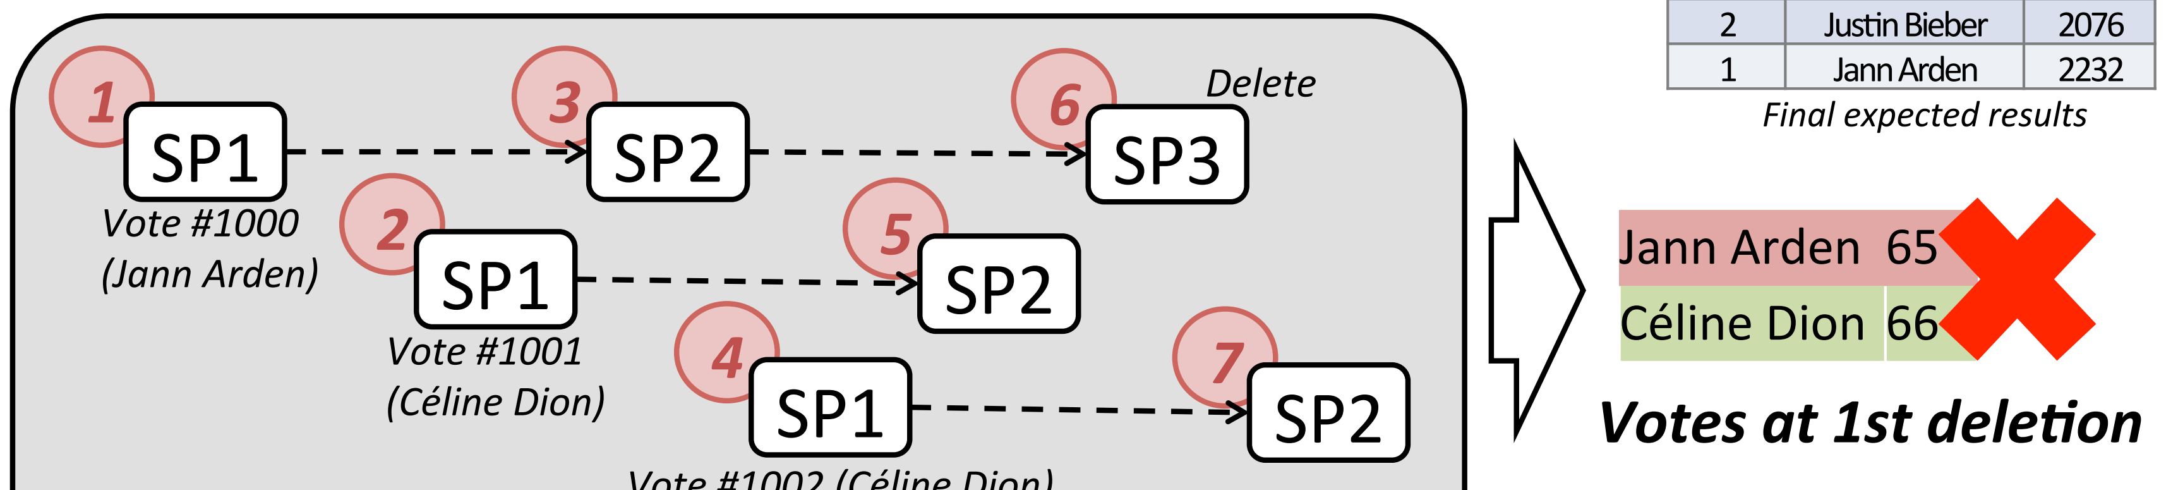

Votes at 1st deletion

S-Store: Transactions with streaming workflows

- Workflows are serially processed when state is shared across stored procedures
- Workflows prevent incorrect ordering of transactions

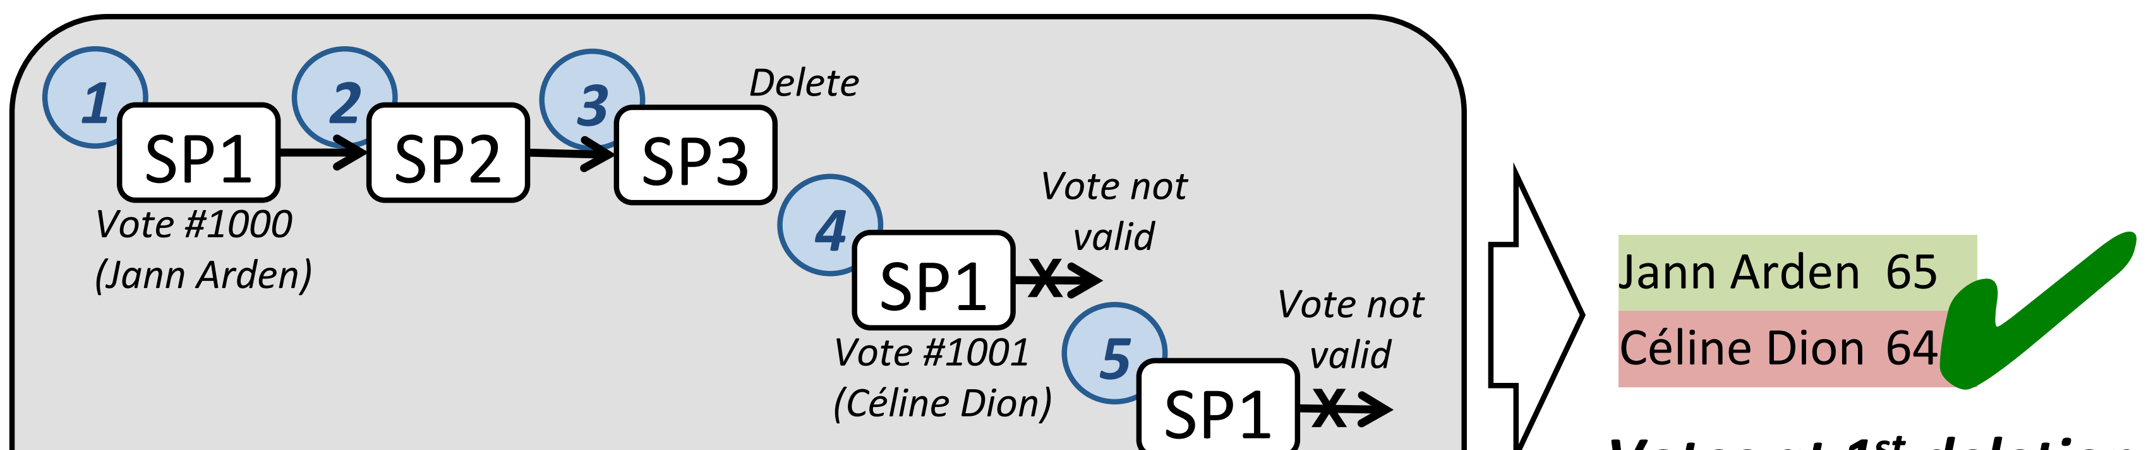

Votes at 1st deletion

Application: Bicycle Sharing

A website and mobile app based on bike sharing programs

- Full integration of streaming data and OLTP workloads
- Customers check out, ride, & check in bikes
- Checkout and check in are OLTP transactions
- Bike locations arrive as streaming data
- Discounts are used to encourage bike returns to stations lacking bikes (calculated dynamically using location data)

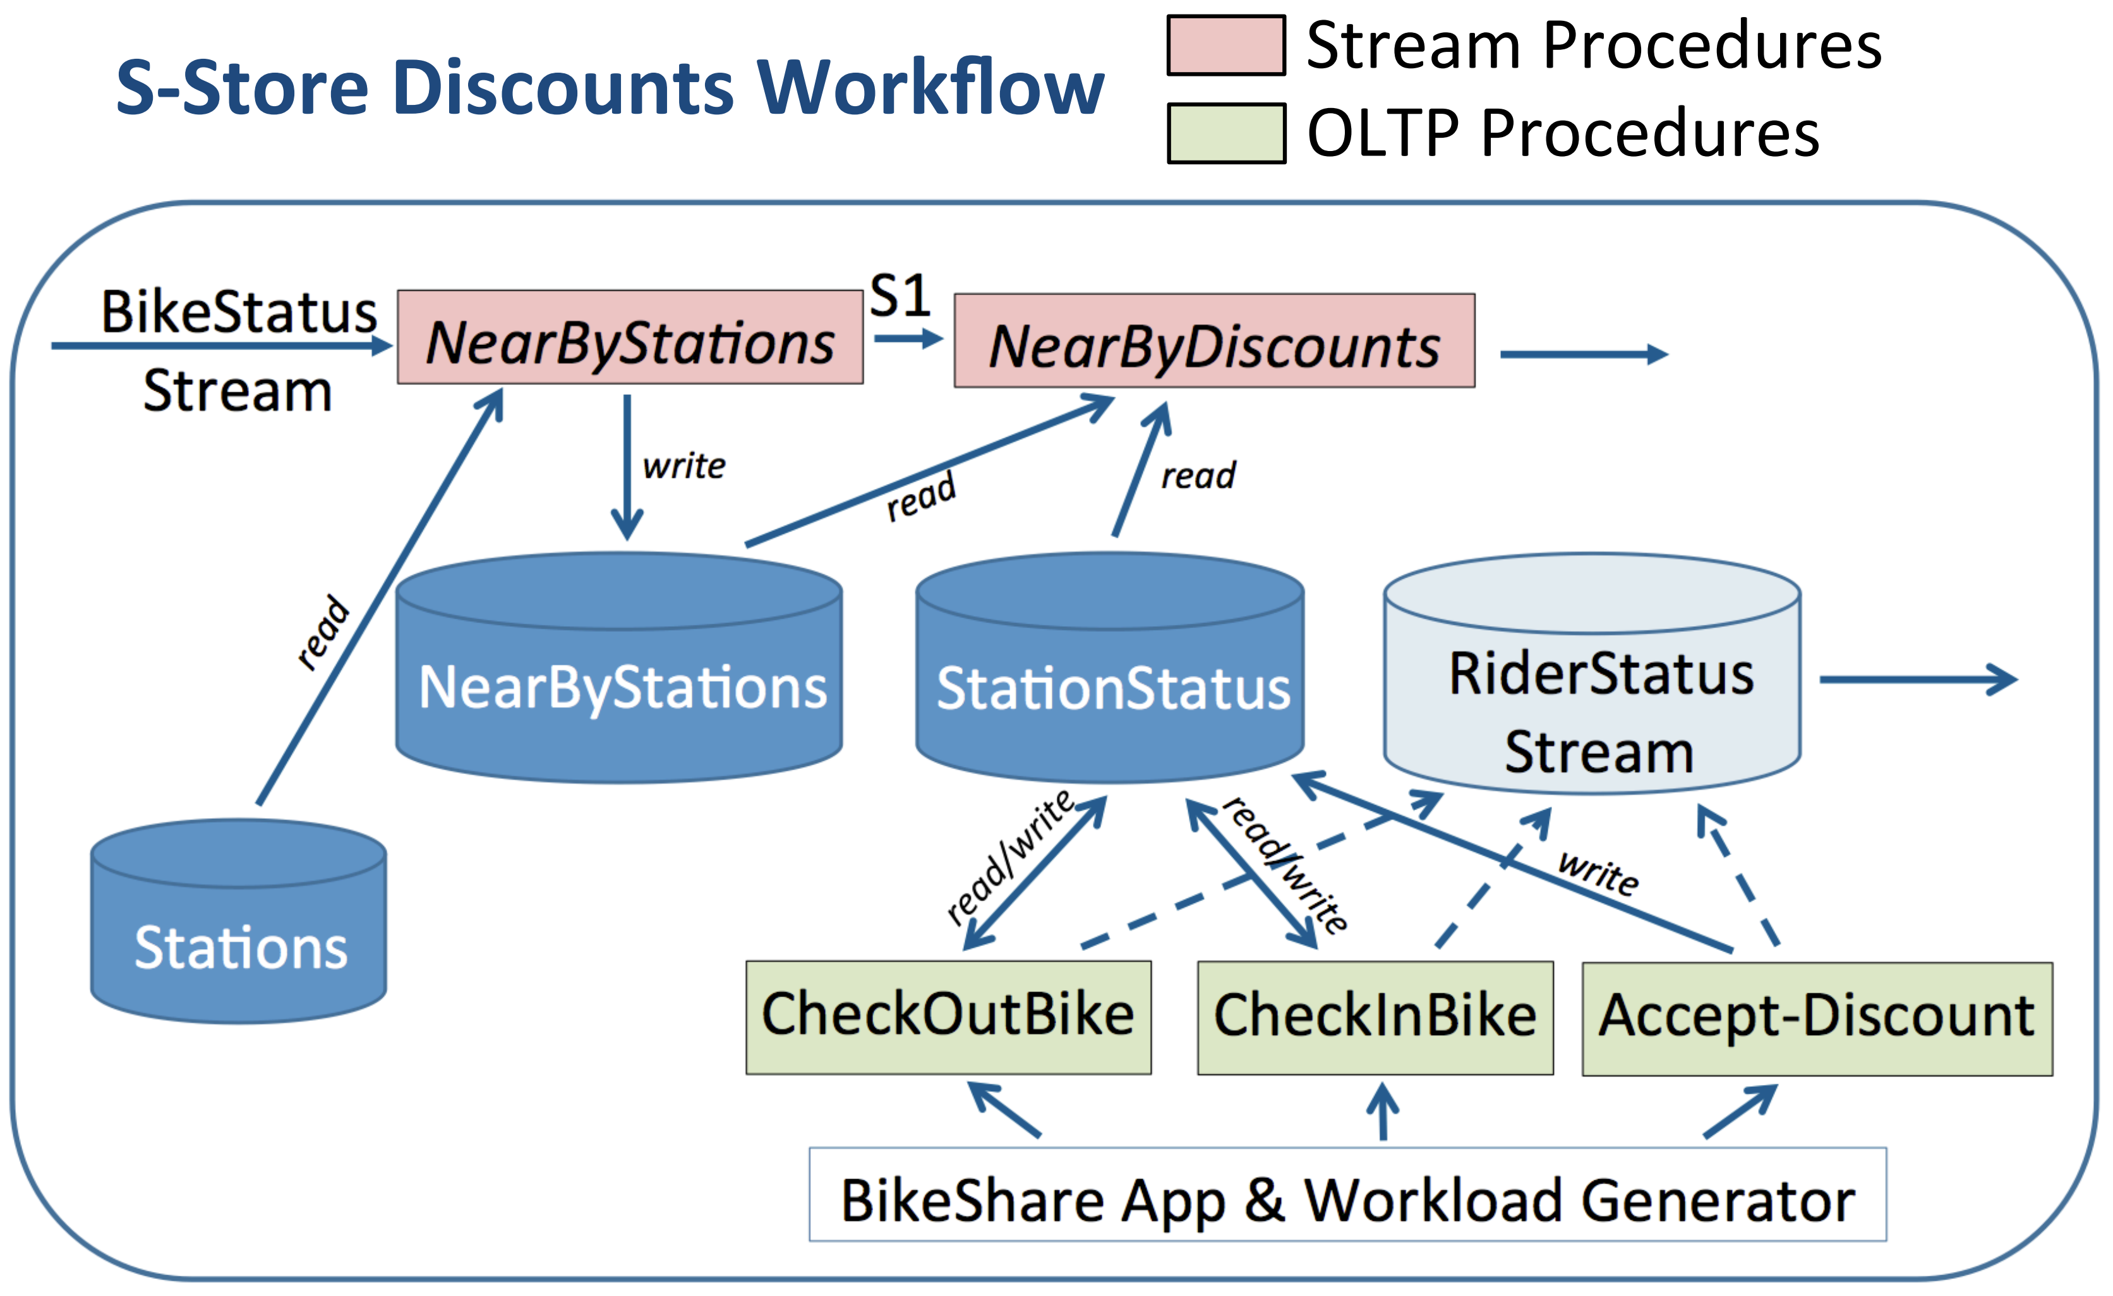

S-Store Discounts Workflow

BikeShare System Demonstration

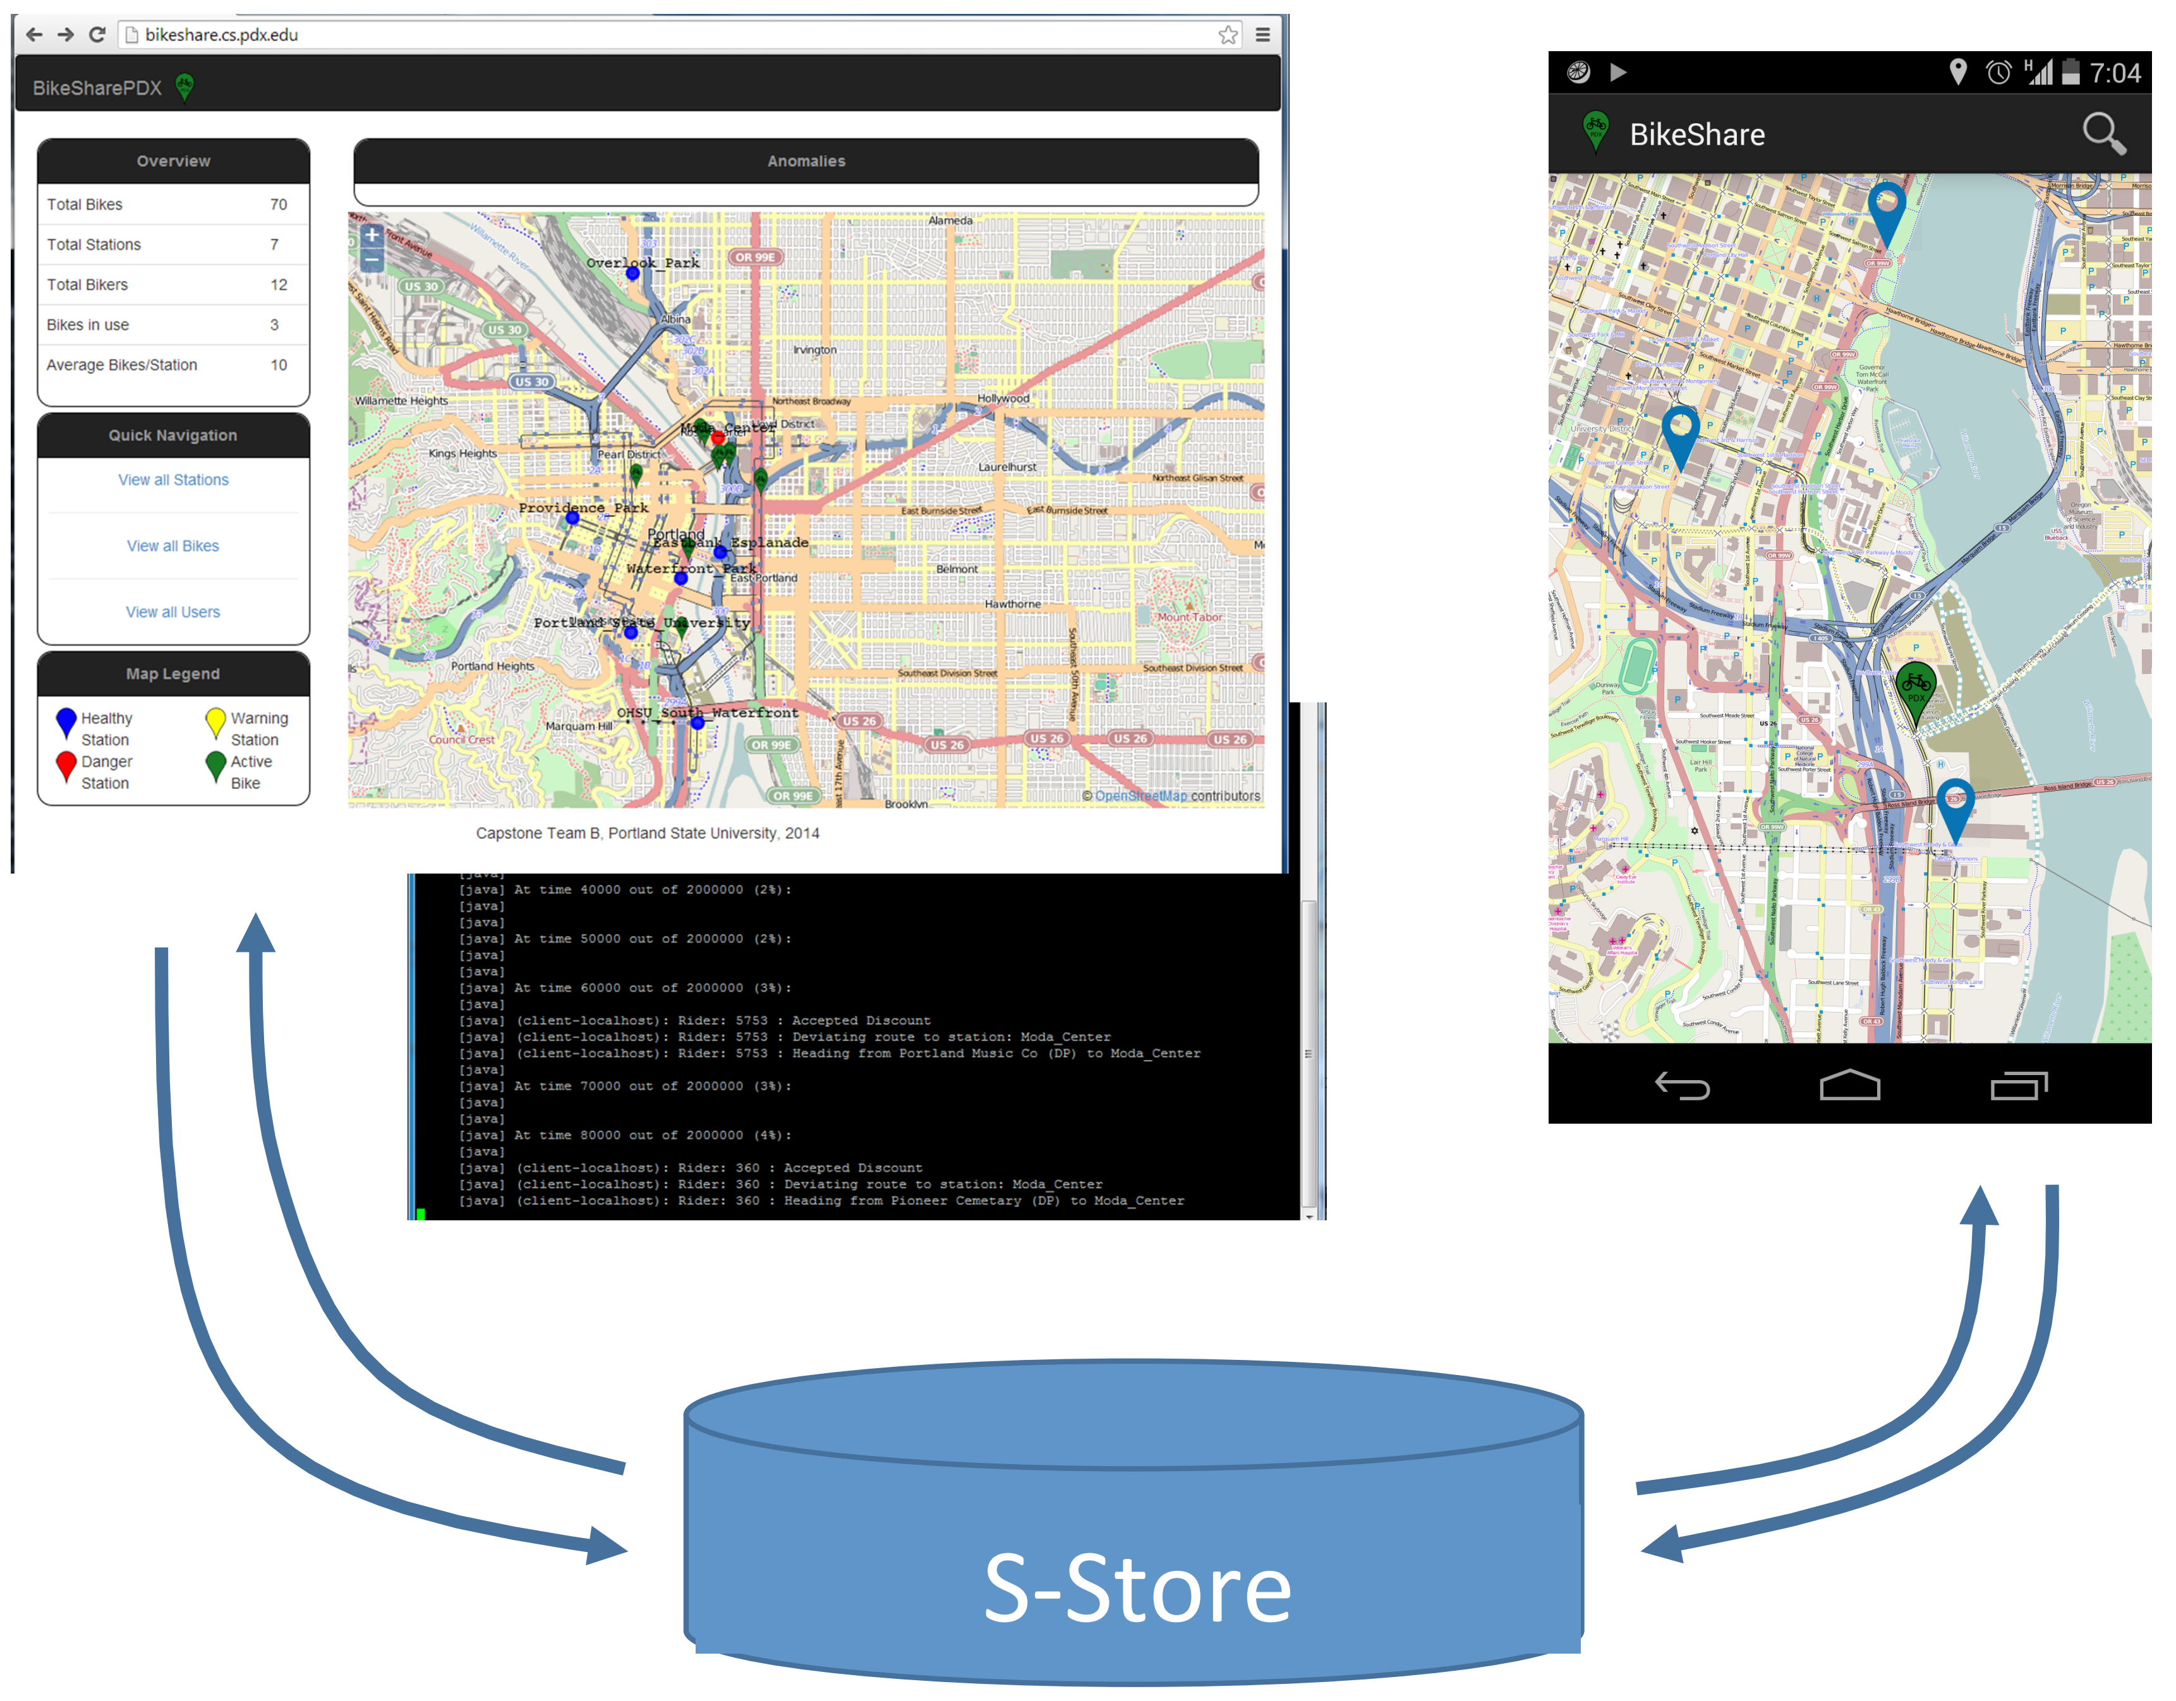

Overview	
Total Bikes	70
Total Stations	7
Total Bikers	12
Bikes in use	3
Average Bikes/Station	10

```
[java] At time 40000 out of 2000000 (2%):
[java]
[java]
[java] At time 50000 out of 2000000 (2%):
[java]
[java]
[java] At time 60000 out of 2000000 (3%):
[java]
[java] (client-localhost): Rider: 5753 : Accepted Discount
[java] (client-localhost): Rider: 5753 : Deviating route to station: Moda_Center
[java] (client-localhost): Rider: 5753 : Heading from Portland Music Co (DP) to Moda_Center
[java]
[java] At time 70000 out of 2000000 (3%):
[java]
[java]
[java] At time 80000 out of 2000000 (4%):
[java]
[java] (client-localhost): Rider: 360 : Accepted Discount
[java] (client-localhost): Rider: 360 : Deviating route to station: Moda_Center
[java] (client-localhost): Rider: 360 : Heading from Pioneer Cemetary (DP) to Moda_Center
```

BikeShare Development Team

PSU Capstone Team B: Chris Giossi, Jonathan Hong, Sara Kane, Craig Meinschein, Colleen Murphy, David Ramage, Clayton Ward

Backend Development: Hong Quach