

In-Network Execution of Monitoring Queries in Sensor Networks

Xiaoyan Yang, Hock-Beng Lim, M. Tamer Ozsu, Kian-Lee Tan

Presented by: Herodotos Herodotou
March 25, 2008

Outline

- Problem Setting
- Challenges
- Query Semantics
- Naïve Approaches
- Two-Phase Self-Join Approach
- Continuous TPSJ
- Experiments
- Discussion

03/25/2008

Herodotos Herodotou

2 / 22

Problem Setting

- Sensor Networks
 - Consist of many small sensor nodes with sensing, data processing and wireless communication capabilities
- Monitoring Queries
 - Track correlation among sensor data within a time window to detect events of interests
- In-network processing of window self-join queries

03/25/2008

Herodotos Herodotou

3 / 22

Challenges

- Sensor nodes are resource-constraint
 - Limited battery energy
 - Limited processing power
 - Limited storage
 - Limited communication bandwidth
- Ability to handle node failures
- Ability to handle communication failures

03/25/2008

Herodotos Herodotou

4 / 22

Query Semantics

Query Template:

```
SELECT S1.AT1, S2.AT2
FROM Sensor AS S1, Sensor AS S2 ← Same Table (self-join)
WHERE p1(S1.AT3) ← Selection Predicate
AND p2(S1.att1, S2.att1) ← Join Predicate
AND window(S1.ts, S2.ts, W) ← Time window of size W
```

Notes:

- Sensor relation is sorted by the reading's timestamp
- AT_i represents subset of attributes from Sensor table
- p₁ is of the form att_i opt Constant
- p₂ is of the form att_i opt att_j
- Opt = { <, >, ≤, ≥, ≠, = }

03/25/2008

Herodotos Herodotou

5 / 22

Query Example

Volcano Monitoring Application:

```
SELECT P1.pressure, P1.time, P2.pressure, P2.time
FROM Pressure AS P1, Pressure AS P2
WHERE P1.pressure > δ
AND P2.pressure > P1.pressure
AND P2.time > P1.time AND P2.time - P1.time < h
```

Query Explanation:

Learn whether the pressure detected has crossed a certain threshold, δ , and is continuously increasing within some period of time, h .

03/25/2008

Herodotos Herodotou

6 / 22

Naïve Approach 1: Flooding

- Approach:
 - Propagate messages containing local sensor data to neighboring nodes
- Pros:
 - In-network query execution
- Cons:
 - Overwhelm limited network bandwidth
 - Need large amount of local storage
 - Incur a lot of computational power and memory

03/25/2008

Herodotos Herodotou

7 / 22

Naïve Approach 2: Centralization

- Approach:
 - All nodes periodically sense and transmit all data back to base station
- Pros:
 - Base station is more powerful
=> can handle self-join queries
- Cons:
 - High data transmission cost
 - Base station could still become bottleneck

03/25/2008

Herodotos Herodotou

8 / 22

Proposed Approach: TPSJ

- TPSJ: Two-Phase Self-Join
 - Phase 1:
 - Preliminary filtering
 - Find candidates that might be in final result
 - Phase 2:
 - Perform window join
 - Candidates found in phase 1 are used to do further filtering within the network

03/25/2008

Herodotos Herodotou

9 / 22

Preprocessing

- Query decomposition into 2 queries:

Q_1^* :
`SELECT S.AT1 INTO R1
 FROM Sensor AS S
 WHERE $p_1(S.AT_3)$`

Q_2^* :
`SELECT S.AT2
 FROM R1, Sensor AS S
 WHERE $p_2(R_1.att_i, S.att_j)$
 AND window($R_1.ts, S.ts, W$)`

Selection query:

- Finds tuples that satisfy selection predicate
- Stores results in temporary relation R_1

Join query:

- Between *Sensor* and intermediate table R_1

03/25/2008

Herodotos Herodotou

10 / 22

Preprocessing Example

- Volcano Monitoring Application:

Q_1^* :
`SELECT P.pressure, P.time
 INTO R1
 FROM Pressure AS P
 WHERE P.pressure > δ`

Q_2^* :
`SELECT P.pressure, P.time
 FROM R1, Pressure AS P
 WHERE P.pressure > R1.pressure
 AND window($R_1.time, P.time, h$)`

Selection query:

- Finds tuples with pressure higher than a certain threshold, δ
- Stores results in temporary relation R_1

Join query:

- Between *Pressure* and intermediate table R_1


03/25/2008

Herodotos Herodotou

11 / 22

TPSJ: Phase 1

- Goal
 - Find candidates that might contribute to final result
- Tasks
 - Execute query Q_1^*
 - If tuple satisfies p_1 , forward it to base station
- Example
 - $\delta = 500$




03/25/2008

Herodotos Herodotou

12 / 22

TPSJ: Phase 2

- **Goal**
 - Find matching tuples for candidates from Ph.1
 - Complete event detection in one time window
- **Tasks**
 - Construct R_1
 - Execute query Q_2^*
- **Example:**
 - $h = 10$



03/25/2008

Herodotos Herodotou

13 / 22

Basic Optimizations

- Join predicate only involves $\{<, >, \leq, \geq\}$
 - Sort R_1
 - Send smallest/largest value instead of R_1
- Rewrite Q_2^* to a simple selection query


```
SELECT S.AT2
FROM Sensor AS S
WHERE p2(r.att1, S.atth)
AND window(W)
```

03/25/2008

Herodotos Herodotou

14 / 22

Continuous TPSJ

- **Goal:**
 - Detect target events over a long period of time
- **Challenges:**
 - Some tuples may trigger a new window self-join within current processing window
 - Avoid unnecessary/duplicate work
- **Solution:**
 - New algorithm for efficient window triggering

03/25/2008

Herodotos Herodotou

15 / 22

Triggering of window self-joins

- **Rule A:**
 - One window self-join per sampling interval
 - **Rule B:**
 - Delay Query Triggering As Much As Possible
 - **Rule C:**
 - Hidden Query
- ↓
- **Theorem:**
 - Number of phase 2 queries injected is minimal for self-join query involving only $\{<, >, \leq, \geq\}$


03/25/2008

Herodotos Herodotou

16 / 22

Experiments for TPSJ (1 window)

- **Network Topology**
 - Total Transmissions




03/25/2008

Herodotos Herodotou

17 / 22

Experiments for TPSJ (1 window)

- **Window Size**
 - Total Transmissions




03/25/2008

Herodotos Herodotou

18 / 22

Experiments for TPSJ (1 window)

- Self-Join Selectivity
 - Total Transmissions




03/25/2008

Herodotos Herodotou

19 / 22

Experiments for Continuous TPSJ

- Network Topology
 - Total Transmissions




03/25/2008

Herodotos Herodotou

20 / 22

Experiments for Continuous TPSJ

- Window Size
 - Savings



03/25/2008

Herodotos Herodotou

21 / 22

Discussion

- Unrealistic Assumptions
 - Node Synchronization
 - Same sampling rate
 - Same transmission time from node to base station (no network delay, no distance issues)
- Scalability - Large number of tuples satisfying the selection predicate =>
 - How to disseminate efficiently?
 - How to store a number of queries in a memory-limited sensor node?
- Redundant tuple transfer from nodes to base station and back
- Insufficient optimization for queries with (in) equality constraints
- No Recursive Generalization for Phase 2 Optimizations
- Experimental Procedure: 8*8 nodes is very limited

03/25/2008

Herodotos Herodotou

22 / 22