



# Optimization of Query Streams Using Semantic Prefetching

---

Ivan T. Bowman  
School of Computer Science  
University of Waterloo

Kenneth Salem  
School of Computer Science  
University of Waterloo

*Presented By : Ahmed A. Soror*



## 1. Behind The Scenes


---

- Query Streams
  - A request stream of open, fetch, close.
  - Single stream/connection.
- Query Patterns

Batches



Nesting



Data Structure  
Correlation





## 1. Behind the Scenes

---

- Optimization
  - Rewriting Semantically related query.
  - Predictively execute new queries.
  - Reduce communication and system interface layers latency and overhead.
- Why not manually ..?
  - Dynamically changing semantics.
  - Manual tuning costs.
  - Decoupling implementation from application logic.



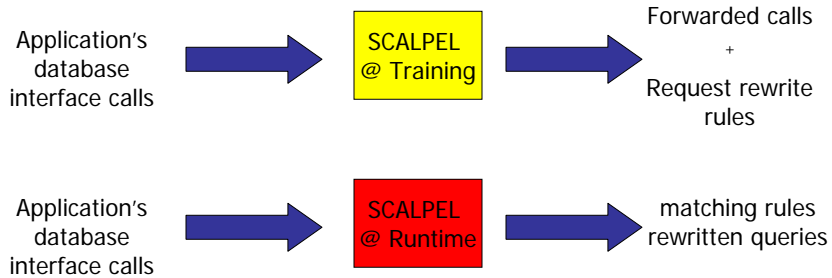
## 1. Behind the Scenes

---

- Semantic Prefetching
  - Execute predicted semantic (context related) queries before actual call.
- Updates after prefetches
  - Same connection : invalidate prefetches
  - Other connections : serializable isolation level


## 2. Introducing .. SCALPEL

- A system for detecting and optimizing patterns of repeated requests within a query stream.



## 3. Training Scalpel

- Context Detection
  - Monitor application requests stream.
  - Track evolving request context.
- Pattern detection
  - Detect correlation between queries and their context
- Pattern optimizer
  - Rewrite cost efficient patterns





### 3. Training Scalpel (Example)

```
function get_openinvoices( cust_id, currency )
// this is Q1
c1 = open( SELECT id, curr, transdate
FROM ar
WHERE customer_id = :cust_id
AND ar.curr = :currency
AND NOT ar.amount = ar.paid
ORDER BY id )
while( r1 = fetch(c1) )
...
if( currency != defaultcurrency )
rate = get_exchangerate( r1.curr,
r1.transdate );
end

function get_exchangerate( curr, transdate )
// this is Q2
r2 = fetch( SELECT exchangerate FROM exchangerate
WHERE curr = :curr
AND transdate = :transdate )
return r2.exchangerate;
end
```

Initial application side Query




```
SELECT id, curr, exchangerate, ...
FROM ar LEFT JOIN exchangerate er
ON ar.curr = er.curr
AND ar.transdate=er.transdate
WHERE customer_id = :cust_id
AND ar.curr = :currency
AND NOT ar.amount = ar.paid
ORDER BY id
```

Optimized Combined Query




### 3. Training Scalpel (Example)

Trace
OPEN(Q <sub>1</sub> (cust1,CDN\$)) → c1
FETCH(c1) → (3305,CDN\$,30/10/03)
OPEN(Q <sub>2</sub> (CDN\$,30/10/03)) → c2
FETCH(c2) → 1.4304
CLOSE(c2)
FETCH(c1) → (3307,CDN\$,30/10/03)
OPEN(Q <sub>2</sub> (CDN\$,30/10/03)) → c2
FETCH(c2) → 1.4304
CLOSE(c2)
FETCH(c1) → (3308,CDN\$,31/10/03)
OPEN(Q <sub>2</sub> (CDN\$,31/10/03)) → c2
FETCH(c2) → 1.4522
CLOSE(c2)
CLOSE(c1)



Query Context	Parameter Context
-	-
Q <sub>1</sub>	(cust1,CDN\$,--,-)
Q <sub>1</sub>	(cust1,CDN\$,3305,CDN\$,30/10/03)
Q <sub>1</sub> ,Q <sub>2</sub>	(cust1,CDN\$,3305,CDN\$,30/10/03),(CDN\$,30/10/03,-)
Q <sub>1</sub> ,Q <sub>2</sub>	(cust1,CDN\$,3305,CDN\$,30/10/03),(CDN\$,30/10/03,1.4304)
Q <sub>1</sub>	(cust1,CDN\$,3305,CDN\$,30/10/03)
Q <sub>1</sub>	(cust1,CDN\$,3307,CDN\$,30/10/03)
Q <sub>1</sub> ,Q <sub>2</sub>	(cust1,CDN\$,3307,CDN\$,30/10/03),(CDN\$,30/10/03,-)
Q <sub>1</sub> ,Q <sub>2</sub>	(cust1,CDN\$,3307,CDN\$,30/10/03),(CDN\$,30/10/03,1.4304)
Q <sub>1</sub>	(cust1,CDN\$,3307,CDN\$,30/10/03)
Q <sub>1</sub>	(cust1,CDN\$,3308,CDN\$,31/10/03)
Q <sub>1</sub> ,Q <sub>2</sub>	(cust1,CDN\$,3308,CDN\$,31/10/03),(CDN\$,31/10/03,-)
Q <sub>1</sub> ,Q <sub>2</sub>	(cust1,CDN\$,3308,CDN\$,31/10/03),(CDN\$,31/10/03,1.4522)
Q <sub>1</sub>	(cust1,CDN\$,3308,CDN\$,31/10/03)



[Q<sub>1</sub>],Q<sub>2</sub>

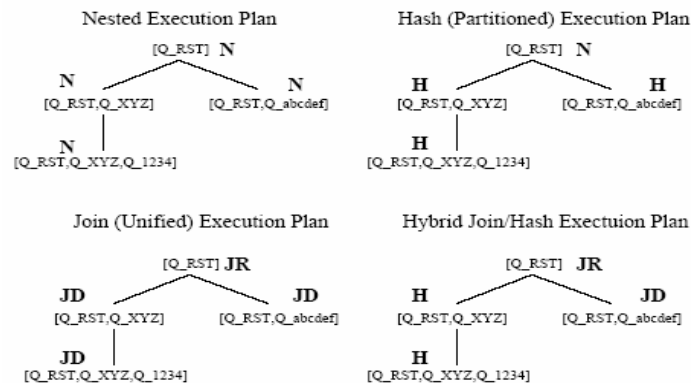


## 3. Pattern Optimizer

- Plan generator
  - Contexts identified may be related to on another.
  - Enumerate all possible execution alternatives
    - Nested Execution (as-is)
    - Partitioned Execution (rewritten inner query hash joined with outer query at client)
    - Unified execution (join inner query with its context)
- Ranking Module
  - Response time cost based ranking



## 3. Pattern Optimizer (Plan generator)



### 3. Pattern Optimizer (Ranking Module)

$$\text{CONTEXT-COST}(C) = \text{OPENS}(C) \times \text{COST}(Q')$$

$$\text{CONTEXT-COST}(C) = P_0(C) \text{OPENS}(C_0) [\text{COST}(Q') + |Q'| H_{\text{add}}] + H_{\text{find}} \text{OPENS}(C)$$

$$\text{OPENS}(C) = \begin{cases} 1 & : C = [Q] \\ P(C) |Q_0| \text{OPENS}(C_0) & : C = [C_0, Q_0] \end{cases}$$

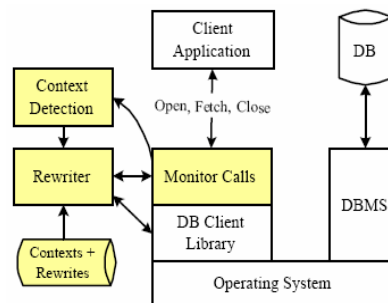
$$\text{COST}(Q) = U_0 + \text{SERVER-COST}(Q) + \text{COMM-COST}(|Q|, \text{BYTES}(Q))$$

$$P(C) = P_0(C) P_1(C)$$

Quantity	Source	Description
SERVER-COST(Q)	RDBMS	Server costs for Q in seconds
Q	RDBMS	Rows returned by Q
BYTES(Q)	RDBMS	Average row length for Q
COMM-COST(N, B)	Scalpel	Communication latency for N rows of B bytes
U <sub>0</sub>	Scalpel	Overhead of a single request
H <sub>add</sub>	Scalpel	Cost of adding to hash table
H <sub>find</sub>	Scalpel	Cost of finding in hash table
P <sub>0</sub> , P <sub>1</sub>	Scalpel	Selectivity of client predicates

### 4. Running Scalpel

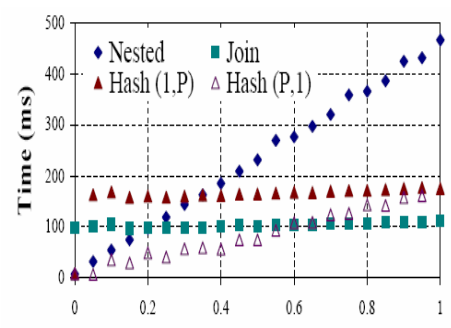
- On context detection Rewriter issues new optimized query.
- Further, Scalpel intercepts predicted calls and replies with prefetched results



# 5. Experiments

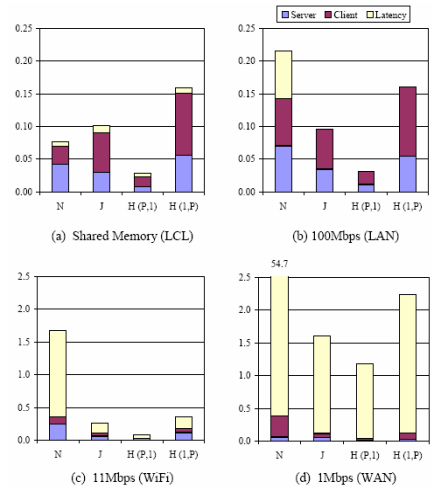
## Effect of client predicate selectivity

- For small selectivity, nested is most efficient.
- Hash strongly depends on  $P_0$  selectivity.



# Experiments

- Execution Costs
  - Nested execution is optimal for local connection.
  - Server costs are 50% lower with joined variants.
  - Nested outperformed in deployments with higher network latency.





## Concluding remarks

---

- Tolerated overhead.
- Optimizing stream instead of individual requests.
- Providing more flexibility to the optimizer.
- Multi-query optimization.
- Is Nesting + stored procedure competitive..?
- Different cost models (resource consumption)
- Why scalpel..?
- Other partitioning, join techniques
- Why log the most recently fetched query results..?
- How long does it take to train it..?



Thank You