

TAG

April 29, 2004

I. Basics

Program the ensemble

SQL-like query language

In network aggregation

II. Query Language

```
SELECT AVG(volume), room FROM sensors
```

```
WHERE floor = 6
```

```
GROUP BY room
```

```
HAVING AVG(volume) > t
```

```
EPOCH DURATION 30s
```

Output is one row for normal queries (columns are AVG(volume), room); or a table for groups with one row per group.

HAVING filters the groups (not the readings)

EPOCH is how time between query executions -- output is thus a stream of values

Queries run until cancelled

III. Aggregation

Aggregation functions:

- o messages contain state records (same for all nodes); e.g. average = <sum, count>
- o define three functions i, f, e
- o i(x) is the initializer; e.g. i(x) -> <x, 1> for average
- o $z = f(x, y)$ is the merge procedure; where x, y, z are state records
 - $\langle sz = sx + sy, cz = cx + cy \rangle$
- o e is the evaluator; computes the final version from a state record

- $e(z) = sz/cz$
- o queries use names for values; a local catalog maps names to local ids

Taxonomy:

- o duplicate sensitive; MAX no, AVG yes
- o Exemplary vs. Summary; MAX is E, AVG is S
- o Monotonic: MAX yes, AVG no
- o Partial state
 - distributive: constant size for each group and same as final value ($e(z) = z$); e.g. MAX
 - algebraic: constant size, but need an evaluator to get final result; e.g. AVG
 - holistic: state records proportional to total network size; e.g. MEDIAN
 - unique: holistic, but only keep unique values, not all values
 - content-sensitive: size depends on data attributes; e.g. histogram

Basics:

- o distribute the query
- o collect output for each epoch
- o tree based routing
- o sub-intervals based on depth in the tree; epoch duration = $d * \text{interval}$, where d is total depth
- o aggregate state records for each group, but can batch messages
- o can push down the HAVING predicates if the aggregation is monotonic
- o may need to evict group values for space reasons (too many groups)
 - evicted groups can be pushed up the tree and fully aggregated later (trading bandwidth and power for temporary space)

IV. Other Issues

Deal with disconnections and loss

- o use potential parents to recover more quickly
- o error due to losses can be very large for COUNT, but is generally small (e.g. 10% for AVG)
- o cache old values and reuse them if needed; more effective use of space than retransmission
- o use multiple paths (via broadcast), can reduce the variance (but faults are not really independent)

Even power usage (assuming aggregation)

Long idle times to save power

Overall, up 20x savings in messages via aggregation (for max), 10x for AVG

Optimizations:

- o snooping can reduce traffic, e.g. MAX
- o hypothesis testing