3-29-2009

Recursive Computation of Regions and Connectivity in Networks

Mengmeng Liu

University of Pennsylvania, mengmeng@seas.upenn.edu

Nicholas E. Taylor

University of Pennsylvania, netaylor@cis.upenn.edu

Wenchao Zhou

University of Pennsylvania, wenchaoz@seas.upenn.edu

Zachary G. Ives

University of Pennsylvania, zives@cis.upenn.edu

Boon Thau Loo

University of Pennsylvania, boonloo@central.cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_papers

Recommended Citation

Mengmeng Liu, Nicholas E. Taylor, Wenchao Zhou, Zachary G. Ives, and Boon Thau Loo, "Recursive Computation of Regions and Connectivity in Networks", . March 2009.

Copyright 2009 IEEE. Reprinted from:

Mengmeng Liu; Taylor, N.E.; Wenchao Zhou; Ives, Z.G.; Boon Thau Loo, "Recursive Computation of Regions and Connectivity in Networks," Data Engineering, 2009. ICDE '09. IEEE 25th International Conference on , vol., no., pp.1108-1119, March 29 2009-April 2 2009 URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4812481&isnumber=4812372>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/417
For more information, please contact repository@pobox.upenn.edu.

Recursive Computation of Regions and Connectivity in Networks

Abstract

In recent years, the data management community has begun to consider situations in which data access is closely tied to network routing and distributed acquisition: examples include, sensor networks that execute queries about reachable nodes or contiguous regions, declarative networks that maintain information about shortest paths and reachable endpoints, and distributed and peer-to-peer stream systems that detect associations (e.g., transitive relationships) among data at the distributed sources. In each case, the fundamental operation is to maintain a view over dynamic network state. This view is typically distributed, recursive, and may contain aggregation, e.g., describing transitive connectivity, shortest paths, least costly paths, or region membership. Surprisingly, solutions to computing such views are often domain-specific, expensive, and incomplete. In this paper, we recast the problem as one of incremental recursive view maintenance in the presence of distributed streams of updates to tuples: new stream data becomes insert operations and tuple expirations become deletions. We develop a set of techniques that maintain compact information about tuple derivability or data provenance. We complement this with techniques to reduce communication: aggregate selections to prune irrelevant aggregation tuples, provenance-aware operators that can determine when tuples are no longer derivable and remove them from their state, and shipping operators that greatly reduce the tuple and provenance information being propagated while still maintaining correct answers. We validate our work in a distributed setting with sensor and network router queries, showing significant gains in communication overhead without sacrificing performance.

Keywords

query processing, sensor fusion, data associations, data management, data provenance, declarative networks, distributed acquisition, distributed systems, incremental recursive view maintenance, network routing, peer-to-peer stream systems, recursive computation, sensor networks

Comments

Copyright 2009 IEEE. Reprinted from:

Mengmeng Liu; Taylor, N.E.; Wenchao Zhou; Ives, Z.G.; Boon Thau Loo, "Recursive Computation of Regions and Connectivity in Networks," Data Engineering, 2009. ICDE '09. IEEE 25th International Conference on , vol., no., pp.1108-1119, March 29 2009-April 2 2009 URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4812481&isnumber=4812372>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Recursive Computation of Regions and Connectivity in Networks

Mengmeng Liu, Nicholas E. Taylor, Wenchao Zhou, Zachary G. Ives, Boon Thau Loo

Computer and Information Science Department, University of Pennsylvania
Philadelphia, PA, U.S.A.

{mengmeng, netaylor, wenchaoz, zives, boonloo}@cis.upenn.edu

Abstract—In recent years, the data management community has begun to consider situations in which data access is closely tied to network routing and distributed acquisition: examples include, sensor networks that execute queries about reachable nodes or contiguous regions, declarative networks that maintain information about shortest paths and reachable endpoints, and distributed and peer-to-peer stream systems that detect associations (e.g., transitive relationships) among data at the distributed sources. In each case, the fundamental operation is to maintain a *view* over dynamic network state. This view is typically distributed, recursive, and may contain aggregation, e.g., describing transitive connectivity, shortest paths, least costly paths, or region membership.

Surprisingly, solutions to computing such views are often domain-specific, expensive, and incomplete. In this paper, we recast the problem as one of *incremental recursive view maintenance* in the presence of distributed streams of updates to tuples: new stream data becomes insert operations and tuple expirations become deletions. We develop a set of techniques that maintain compact information about tuple derivability or *data provenance*. We complement this with techniques to reduce communication: aggregate selections to prune irrelevant aggregation tuples, provenance-aware operators that can determine when tuples are no longer derivable and remove them from their state, and shipping operators that greatly reduce the tuple and provenance information being propagated while still maintaining correct answers. We validate our work in a distributed setting with sensor and network router queries, showing significant gains in communication overhead without sacrificing performance.

I. INTRODUCTION

As data management systems are handling increasingly distributed and dynamic data, the line between a network and a query processor is blurring. In a plethora of emerging applications, data originates at a variety of nodes and is being frequently updated: routing tables in a peer-to-peer overlay network [1] or in a declarative networking system [2], [3], sensors embedded in an environment [4], [5], monitors within various clusters at geographically distributed hosting sites [6], [7], data producers in large-scale distributed scientific data integration [8]. It is often natural to express distributed data acquisition, integration, and processing for these settings using declarative queries — and in some cases to compute and incrementally maintain the results of these queries, e.g., in the form of a routing table, an activity log, or a status display.

The queries that are of interest in this domain are frequently quite different from the OLAP or OLTP queries that exemplify centralized DBMS query processing. We consider two main settings.

Declarative networking. In declarative networking [9], [3], an extended variant of datalog has been used to manage the state in routing tables — and thus to control how network messages are forwarded through the network. Perhaps the central task in this work is to compute paths available through multi-hop connectivity, based on information in neighboring routers tables. It has been shown that recursive path queries, used to determine reachability and cost, can express conventional and new network protocols in a declarative way.

Sensor networks. Declarative, database-style query systems have also been shown to be effective in the sensor realm [4], [5], primarily for aggregation-style queries. Outside the database community, a variety of macroprogramming languages [10], [11] have been proposed as alternatives, which include features like region and path computations. In the long run, we argue that the declarative query approach is superior because of data independence and optimization. However, the query languages and runtime systems must be extended to match the functionality of macroprogramming, particularly with respect to computing regions and paths.

Section II provides a number of detailed use cases and declarative queries for regions and paths in these two domains. The use cases are heavily reliant on *recursive* computations, which must be performed over distributed data that is being frequently updated in “stream” fashion (e.g., sensor state and router links are dynamic properties that must be constantly refreshed). The majority of past work on recursive queries [12], [13] has focused on recursion in the context of centralized deductive databases, and some aspects of that work have ultimately been incorporated into the SQL-99 standard and today’s commercial databases. However, recursion is relatively uncommon in traditional database applications, and hence little work has been done to extend this work to a distributed setting. We argue that the advent of declarative querying over networks has made recursion of fundamental interest: it is at the core of the main query abstractions we need in a network, namely regions, reachability, shortest paths, and transitive associations.

To this point, only specializations of recursive queries have been studied in networks. In the sensor domain, algorithms have been proposed for computing regions and neighborhoods [10], [11], [14], but these are limited to situations in which data comes from physically contiguous devices, and computation is relatively simple. In the declarative networking domain, a semantics has been defined [3] that closely matches

router behavior, but it is not formalized, and hence the solution does not generalize. Furthermore, little consideration has been given to the problem of *incremental computation* of results in response to data arrival, expiration, and deletion.

In this paper, we show how to compute and incrementally maintain recursive views over data streams, in support of networked applications. In contrast to previous maintenance strategies for recursive views [15], our approach emphasizes *minimizing the propagation of state* — both across the network (which is vital to reduce communication overhead) and inside the query plan (which reduces computational cost). Our methods generalize to sensors, declarative networking, and data stream processing. We make the following contributions:

- We develop a novel, compact *absorption provenance*, which enables us to directly detect when view tuples are no longer derivable and should be removed.
- We propose a *MinShip* operator that reduces the number of times that tuples annotated with provenance need to be propagated across the network and in the query.
- We generalize *aggregate selection* to handle streams of insertions and deletions, in order to reduce the propagation of tuples that do not contribute to the answer.
- We evaluate our schemes within a distributed query processor, and experimentally validate their performance in real distributed settings, with realistic Internet topologies and simulated sensor data.

Section II presents use cases for declarative recursive views. In Section III we discuss the distributed query processing settings we address. Sections IV through VI discuss our main contributions: absorption provenance, the *MinShip* operator, and our extended version of aggregate selection. We present experimental validation in Section VII, describe related work in Section VIII, and wrap up and discuss future work in Section IX.

II. DISTRIBUTED RECURSIVE VIEW USE CASES

We motivate our work with several examples that frame network monitoring functionalities as distributed recursive views. This is not intended to be an exhaustive coverage of the possibilities of our techniques, but rather an illustration of the ease with which distributed recursive queries can be used.

Throughout the paper, we assume a model in which logical relations describe state horizontally partitioned across many nodes, as in declarative networking [9]. In our examples, we shall assume the existence of a relation *link(src, dst)*, which represents all router link state in the network. Such state is partitioned according to some key attribute; unless otherwise specified, we adopt the convention that a relation is partitioned based on the value of its first attribute (*src*), which may (depending on the setting) directly specify an IP address at which the data is located, or a logical address like a DNS name or a key in a content-addressable network [1].

Network Reachability. The textbook example of a recursive query is graph transitive closure, which can be used to compute network reachability. Assume the query processor at node *X* has access to *X*'s routing table. Let a tuple *link(X, Y)* denote the presence of a link between node *X* and its neighbor *Y*.

Then the following query computes all pairs of nodes that can reach each other.

```
with recursive reachable(src,dst) as
( select src,dst
  from link
  union
  select link.src, reachable.dst
  from link, reachable
  where link.dst = reachable.src)
```

The techniques of this paper are agnostic as to the query language; we could express all queries in datalog, as in [9]. However, since SQL has a more familiar syntax, we present our examples using SQL-99's recursive query syntax¹. The SQL query (view) above takes base data from the *link* table, then recursively joins *link* with its current contents to generate a transitive closure of links. Note that since all tables are originally partitioned based on the *src*, computing the view requires a distributed join that sends *link* tuples to nodes based on their *dst* attributes, who join with *reachable.src*.

There are many potential enhancements to this query, e.g., to compute reachable pairs within a radius, or to find cycles.

Network Shortest Path. We next consider how to compute the shortest path between each pair of nodes, in terms of the hop count (number of links) between the nodes:

```
with recursive path(src,dst,vec,length) as
( select src,dst,src || '.' || dst,1 from link
  union
  select link.src,path.dst,link.src || '.' || vec,
    length+1
  from link, path where link.dst = path.src)

create view minHops(src,dst,length) as
(select src,dst,min(length) from path
 group by src,dst)

create view shortestPath(src,dst,vec,length) as
(select P.src,P.dst,vec,P.length
 from path P, minHops H where P.src = H.src
 and P.dst = H.dst and P.length = H.length)
```

This represents the composition of three views. The *path* recursive view is similar to the previous *reachable* query, with additional computation of the path length, as well as the path itself. The other (non-recursive) views *minHops* and *shortestPath* determine the length of the shortest path, and the set of paths with that length, respectively.

Network Highest-Bandwidth Path. We can similarly define the *highest bandwidth* path: instead of counting the number of links, we instead set a path's bandwidth to be the minimum bandwidth along any link; and then find, for any pair of endpoints, the path with maximum bandwidth.

Sensing Contiguous Regions. In addition to querying the graph topology itself, distributed recursive queries can be used to detect regions of neighboring nodes that have correlated activity. One example is a *horizon* query, where a node computes a property of nodes within a bounded number of hops of itself. A second example (which we show and experimentally evaluate in Section VII) starts with a series of reference nodes, and computes contiguous regions of triggered sensors near these nodes. This is useful in sensor networks, e.g., in order to determine the average temperature of a fire.

¹We assume SQL UNIONS with set semantics, and that a query executes until it reaches fixpoint. Not all SQL implementations support these features.

Fig. 1. Basic architecture: query processing nodes are placed in a number of sub-networks. Each collects state information about its sub-network, and the nodes share state to compute distributed recursive views such as shortest paths across the network.

Other Example Queries. The *routing resilience* query counts the number of paths (alternate routes) between any two nodes. Another class of queries examines multicast or aggregation trees constructed within the network. A query could compute the height of each subtree and store this height at the subtree root. Alternatively, we might query for the *imbalance* in the tree – the difference in height between the lowest and highest leaf node. Finally, a query could identify all the nodes at each level of the tree (referred to as the “same generation” query in the datalog literature).

III. EXECUTION MODEL AND APPROACH

We consider techniques applicable to a broad variety of networked environments, and we make few assumptions about our execution environment. We assume that our networked query processor executes across a number of distributed nodes in a network; in addition, we allow for the possibility that there exist other legacy nodes that may not run the query processor (as indicated in Figure 1). In this flexible architecture, the query processing nodes will serve as proxy nodes storing state information (connectivity, sensor status, etc) about devices on their sub-networks: IP routers, overlay nodes, sensors, devices, etc.

Individual sub-networks may have a variety of types of link-layers (wired IP, wireless IP with a single base station, multi-hop wireless/mesh, or tree-structured sensor networks). They may even represent different autonomous systems on the Internet backbone, or different locations within a multi-site organization. Through polling, notifications, or snooping, our distributed query processing nodes can acquire detailed information about these sub-networks. The query processing nodes each maintain a *horizontal partition* of one or more views about the overall network state: cross-sub-network shortest paths, regions that may span physically neighboring sub-networks (e.g., a fire in a multi-story building), etc. During operation, the nodes may exchange state with one another, either (1) to partition state across the nodes according to keys or ranges, or (2) to perform computation of joins or recursive queries.

Importantly, in a volatile environment such as a network, both sensed state and connectivity will frequently change. Hence a major task will be to *maintain* the state of the views, as base data (sensor readings, individual links) are added or deleted, as distributed state ages beyond a *time-to-live* and gets expired, and as the effects of deletions or expirations get propagated to derived data.

A. Query Execution Model

In networks, query execution is a distributed, continuous stream computation, over a set of horizontally partitioned base relations that are updated constantly. We assume that all communication among nodes is carried out using a reliable in-order delivery mechanism. We also assume that our goal is to compute and update *set* relations, not bag relations: we stop computing recursive results when we reach a fixpoint.

In our model, inputs to a query are streams of insertions or deletions over the base data. Hence, we process more general *update streams* rather than tuple streams. *Sliding windows*, commonly used in stream processing, can be used to process *soft-state* [16] data, where the time-based window size essentially specifies the useful lifetime of base tuples. Thus, a base tuple that results from an insertion may receive an associated timeout, after which the tuple gets deleted. When this happens, the derived tuples that depend on the base tuples have to be deleted as well. Due to the needs of network state management, we consider timeouts or windows to be specified over base data only, not derived tuples.

B. Motivation for New Distributed Recursive Techniques

To illustrate the need for our approach, we consider an example. Assume our goal is to maintain, at every node, the set of all nodes reachable from this node. Refer to Figure 2, which shows a network consisting of three nodes and four links (visualized in Figure 3). Each node “knows” its direct neighbors: we represent these in the *link* table, consisting of four entries $link(A, B)$, $link(B, C)$, $link(C, A)$, and $link(C, B)$. As in our previous examples, the *link* table is partitioned such that all values with source *src* are stored on node *src*. In our simple example, there is a direct correspondence between *src* value and location, although one could decouple each location from its physical encoding by using logical addresses (e.g., doing hash-based partitioning).

Now we define a materialized view $reachable(src, dst)$, which is also partitioned so tuples with source *src* are stored on node *src*. This query computes the transitive closure over the *link* table, and was shown in the **Network Reachability** example of Section II. Unlike in traditional recursive query execution (e.g., for datalog), here computing the transitive closure requires a good deal of communications traffic: *link* data must be shipped to the node corresponding to its *dst* attribute in order to join with *reachable* tuples²; and the output of this join may need to be shipped to a new location depending on what its *src* is. Consider the execution plan shown in Figure 4. This plan is disseminated to all nodes, from which it continuously generates and updates partitions of the reachability relation. The left *DistributedScan* represents the table scan required for the base case, which fetches the contents of *link* and sends them to the *Fixpoint* operator. In the recursive case, the *Fixpoint* invokes the right subtree of the query plan: it sends its current contents to a *FixPointReceiver*, where they are joined via a *PipelinedHashJoin* with a copy of *link* — whose contents have been re-partitioned and shipped to the nodes corresponding to the *dst* attribute. The output

²Or vice-versa, depending on the query plan.

is shipped to the fixpoint via the *MinShip* (tuple shipping) operator, which in the simplest case simply sends data to a receiving node.

Computing the View Instance. Figure 2 steps through the execution of *reachable*, showing state after each computation step in semi-naïve evaluation (equivalent to steps in stratified execution), as well as communication (the “*at* → *to*” columns). We defer discussion of the column marked *pv*.

The base-case contents of *reachable* are computed directly from *link*, as specified in the first “branch” of the view definition (See **Network Reachability** query in Section II). The recursive query block joins all *link* tuples with those currently in *reachable*. Since the tables are distributed by their first attribute, all *link* tuples must first be shipped to nodes corresponding to their *dst* attribute, where they are joined with *reachable* tuples with matching *srcs*. Finally, the resulting *reachable* tuples must be shipped to the nodes corresponding to their *src* attributes. For instance, in step 1, *reachable*(*C*, *B*) is computed by joining *link*(*C*, *A*) and *reachable*(*A*, *B*) as computed from step 0. That requires first shipping *link*(*C*, *A*) to node *A*, performing the join to generate *reachable*(*C*, *B*), and sending the resulting tuple to node *C*. In our figure, we indicate the communication for the resulting *reachable* table in the third column as *A* → *C*.

Since we are following set-semantics execution, duplicate removal will eliminate tuples with identical values; but this only occurs *after* they are created and sent to the appropriate node. For instance, consider *reachable*(*C*, *C*), which is first computed in step 1 and sent to node *C*. During step 2, node *A* re-derives this same tuple; however, it must send this result to node *C* before the duplication can be detected, and the tuple eliminated. In total, 16 tuples (4 initial *link* tuples, and 12 *reachable* tuples) are shipped during the recursive computation. In the final step, a fixpoint is reached when no new tuples are derived. Observe that since we have a fully-connected network, the final resulting *reachable* table at every node contains the set of all node pairs in the network with the first attribute matching the node’s address.

Incremental Deletion (Standard Approach). Now consider the case when *link*(*C*, *B*) expires (hence is deleted). Commonly used schemes for maintaining non-recursive views, such as counting tuple derivations, do not apply to this recursive view. Instead, one might employ the standard algorithm for recursive view maintenance, DRed [15]. DRed works by first *over-deleting* tuples conservatively and then *re-deriving* tuples that may have alternative derivations. Figure 5 shows the DRed over-deletion phase (steps 0-4), followed by the rederivation phase (steps 5-8). In the over-deletion phase, it first deletes *reachable*(*C*, *B*) based on the initial deletion of *link*(*C*, *B*). This in turn leads to the deletion of all *reachable* tuples with *src* = *C* (step 1), then those with *src* = *B* (step 2) and *src* = *A* (step 3). The *reachable* table is empty in step 4. DRed will ultimately re-derive *every reachable* tuple, as shown in steps 5-8. Overall, DRed requires shipping a total of 16 tuples, equivalent to computing the entire *reachable* view from scratch, despite having just a single deletion.

In the above example, DRed is prohibitively expensive:

deleting a single link resulted in the deletions of all *reachable* tuples; yet, it is clear that nodes *A*, *B*, and *C* are still connected after *link*(*C*, *B*) is deleted. One source of deletions in network settings is tuple expirations; given the fact that large-scale network tends to be highly dynamic, tuples will need to expire frequently, thus triggering frequent re-computation and exacerbating the overhead. Perhaps surprisingly, our example illustrates the *common case* behavior for network state queries: most networks are well-connected with bi-directional connectivity along several redundant paths. DRed will over-delete such paths, and then re-derive data.

We have ignored a further issue that DRed must wait until all deletions have been processed before it can start rederiving. (This requires distributed synchronization, which may be expensive.)

C. Our Approach

We now propose a solution that eliminates the need for recomputation, and that also avoids global synchronization. The major challenge with distributed incremental view maintenance lies in handling deletions of tuples. In general, we must either buffer base tuples, then recompute the majority of the query (as in our example); or we must maintain *state* at intermediate nodes, which enables them to propagate the appropriate updates when a base tuple is removed. We adopt the latter approach, developing a scheme that:

- Maintains a concise form of *data provenance* — bookkeeping about the derivations and derivability of tuples — such that it is easy to determine whether a view tuple should be removed when a base tuple is removed. (Section IV.)
- Propagates *provenance information* from one node to another only when necessary to ensure correctness — thus reducing network and computation costs. (Section V.)
- Propagates *tuples* through distributed aggregate computations only when necessary for correctness — also reducing network and computation costs. (Section VI.)

We describe these features in the next three sections, with the query plan of Figure 4 as the central example.

IV. PROVENANCE FOR EFFICIENT DELETIONS

In order to support view maintenance when a base tuple is deleted, we must be able to test whether a derived tuple is *still derivable*. Rather than over-delete and re-derive (as with DRed), we instead propose to keep around metadata about derivations, i.e., *provenance* [17], also called *lineage* [18].

Provenance alternatives. Different proposed forms of provenance capture different amounts of information. Lineage in [18] encodes the set of tuples from which a view tuple was derived — but this is not sufficiently expressive to distinguish what happens if a base tuple is removed. Alternatives include why-provenance [17], which encodes *sets of source tuples* that produced the answer; and the semiring polynomial provenance representation of [8], [19], whose implementation we term *relative provenance* here. In physical form, the latter encodes a derivation graph capturing which tuples are created as *immediate consequents* of others. The graph can be traversed after a deletion to determine whether a tuple is still derivable from base data [8]. Either of these latter two forms of provenance

Algorithm 1 Fixpoint operator

Fixpoint(B^Δ, R^Δ)

Inputs: Input base stream B^Δ , recursive stream R^Δ

Output: Output stream U'^Δ

```

1: Init hash map  $P: U(\bar{x}) \rightarrow$  provenance expressions over  $U(\bar{x})$ 
2: if there is a aggregate selection option then
3:   Get the grouping key  $uk$ , number of aggregate functions  $n$  and
   aggregate functions  $agg_1, \dots, agg_n$ 
4:    $B'^\Delta := AggSel(B^\Delta, uk, n, agg_1, \dots, agg_n)$ 
5:    $B^\Delta := B'^\Delta$ 
6:    $R'^\Delta := AggSel(R^\Delta, uk, n, agg_1, \dots, agg_n)$ 
7:    $R^\Delta := R'^\Delta$ 
8: end if
9: while not EndOfStream( $B^\Delta$ ) and not EndOfStream( $R^\Delta$ ) do
10:   Read an update  $u$  from  $B^\Delta$  or  $R^\Delta$ 
11:   if  $u.type = \text{INS}$  then
12:     if  $P$  does not contain  $u.tuple$  then
13:        $P[u.tuple] := u.pv$ 
14:       Add  $u.tuple$  to the view
15:       Output  $u$  to the next operator
16:     else
17:        $oldPv := P[u.tuple]$ 
18:        $P[u.tuple] := P[u.tuple] \vee u.pv$ 
19:        $deltaPv := P[u.tuple] \wedge \neg oldPv$ 
20:       if  $oldPv \neq P[u.tuple]$  then
21:          $u'.tuple := u.tuple$ 
22:          $u'.type := \text{INS}$ 
23:          $u'.pv := deltaPv$ 
24:         Output  $u'$  to the next operator
25:       end if
26:     end if
27:   else if  $u$  is from  $B^\Delta$  then
28:     for each  $t$  in  $P$  do
29:        $oldPv := P[t]$ 
30:        $P[t] := restrict(P[t], \neg u.pv)$ 
31:       if  $P[t]$  indicates no derivability then
32:         Remove  $t$  from  $P$ 
33:       end if
34:     end for
35:   end if
36: end while

```

$\sigma_\theta(R)$:	If tuple t in R satisfies θ , annotate t with $\mathbf{P}(t)$
$R_1 \bowtie R_2$:	For each tuple t_1 in R_1 and tuple t_2 in R_2 , annotate the output tuple $t_1 \bowtie t_2$ with $\mathbf{P}(t_1) \wedge \mathbf{P}(t_2)$.
$R_1 \cup R_2$:	For each tuple t output by $R_1 \cup R_2$, annotate t with $\mathbf{P}(t_1) \vee \mathbf{P}(t_2)$, where $\mathbf{P}(t_1)$ is false iff t does not exist in R_1 ; similarly for $\mathbf{P}(t_2)$, R_2
$\Pi_A(R)$:	Given tuples $t_1, t_2, \dots, t_n$ that project to the same tuple t' , annotate t' with $\mathbf{P}(t_1) \vee \mathbf{P}(t_2) \vee \dots \vee \mathbf{P}(t_n)$

Fig. 6. Relational algebra rules for composition of provenance expressions. of provenance described in the theoretical paper [19]).

Our key innovation with respect to provenance is to develop a physical representation in which we can exploit *Boolean absorption* to minimize the provenance expressions: absorption is based on the law $a \wedge (a \vee b) \equiv a \vee (a \wedge b) \equiv a$, and it eliminates terms and variables from a Boolean expression that are not necessary to preserve equivalence. We term this model *absorption provenance*. It describes in a minimal way exactly which tuples, in which combinations of join and union, are essential to the existence of a tuple in the view. The benefit of a compact provenance annotation is reduced network traffic. Even better, we can use absorption provenance to help maintain a view after a base tuple has been deleted: we assign the value **false** to the provenance variable for each deleted base tuple, then substitute this value into all provenance annotations of tuples in the view. If applying absorption to the tuple's provenance results in the value **false**, we remove the tuple. Otherwise, it remains derivable.

Absorption provenance in the example of Figure 2. Ab-

sorption provenance adds a bit of overhead to normal query computation: the fixpoint operator must propagate a tuple through to the recursive step whenever it receives a *new derivation* (even of an existing tuple), not simply when it receives a new tuple. Refer back to the *reachable* query example of Figure 2. The *pv* column shows the absorption provenance for every tuple during the initial view computation, with respect to the input *link* tuples annotated p_1, p_2, p_3 , and p_4 ; we see that an additional 4 tuples (beyond the previous set-oriented execution model) is shipped during query evaluation, as a result of computing absorption provenance. For instance, *reachable*(B, B) is derived in both strata 1 and 2. They have different provenance that cannot be absorbed, hence we must track both derivations.

Absorption provenance shows its value in handling deletions. When *link*(C, B) is deleted, the only step required with absorption provenance is to zero out p_4 in the provenance expressions of all *reachable* tuples. In this example, zeroing out this derivation only requires two message transmissions, and it does not result in the removal of any tuples from the view. (In the worst case it is still possible that deletions may need to be propagated to all nodes in the network.)

A. Implementing Absorption Provenance

There are multiple alternatives when attempting to encode an absorption provenance expression. Each expression can, of course, be normalized to a sum-of-products expression, since in the end there are possibly multiple derivations of the same tuple, and each derivation is formed by a conjunctive rule (or a conjunction of tuples that resulted from conjunctive rules). From there we could implement absorption logic that is invoked every time the provenance expression changes. We choose an alternative — and often more compact — encoding for absorption provenance: the *binary decision diagram* [20] (BDD), a compact encoding of a Boolean expression in a DAG. A BDD (specifically, a *reduced ordered BDD*) represents each Boolean expression in a canonical way, which automatically eliminates redundancy by merging isomorphic subgraphs and removing isomorphic children: this process automatically applies absorption. Since BDDs are frequently used in circuit synthesis applications and formal verification, many highly optimized libraries are available [21]. Such libraries, e.g., [21], provide abstract BDD types as well as Boolean operators to perform on them: pairs of BDDs can be ANDed or ORed; individual BDDs can be negated; and variables within BDDs can be set or cleared. We exploit such capabilities in our provenance-aware stateful query operators.

Now we describe in detail the implementation of absorption provenance within the Fixpoint operator. We defer a discussion of how aggregation state management works to Section VI.

B. Fixpoint Operator

The key operator for supporting recursion is the Fixpoint operator, which first calls a **base case** query to produce results, then repeatedly invokes a **recursive case** query. It repeatedly unions together the results of the base case and each recursive step, and terminates when no new results have been derived. We define the fixpoint in a recursive query as follows: we reach

a fixpoint when we can no longer derive any new results that affect the absorption provenance of any tuple in the result.

Unlike traditional semi-naïve evaluation, our fixpoint operator does not block or require computations in synchronous rounds (or iterations), a prohibitively expensive operation in distributed settings. We achieve this with the use of pipelined semi-naïve evaluation [9], where tuples are handled in the order in which they arrive via the network (assuming a FIFO channel), and are only combined with tuples that arrived previously.

Pseudocode for this operator is shown in Algorithm 1. The fixpoint operator receives insertions from either the base (B^Δ) or recursive (R^Δ) streams. It maintains a hash table P containing the absorption provenance of each tuple that it has received, which remains derivable. Note that in our algorithms, each tuple now contains three fields, *type* which indicates whether it is an INS or DEL tuple, *tuple* which records its raw tuple values, and *pv* which stores its annotated provenance.

Initially (Lines 2–8), we apply any portions of an aggregation operation that might have been “pushed into” the fixpoint — this uses a technique called *aggregate selection* discussed in Section VI. Now, upon receipt of an insertion operation u (Lines 11–26), the fixpoint operator first determines whether the tuple has already been encountered (perhaps with a different provenance). If u is new, it is simply stored in $P[u.\text{tuple}]$ as the first possible derivation; otherwise we merge it with the existing absorption provenance in $P[u.\text{tuple}]$. We save the resulting *difference* in deltaPv . If the provenance has indeed changed despite absorption, u gets propagated to the next operator, annotated with provenance deltaPv .

Deletions are handled in a straightforward fashion (Lines 27–35), given our implementation of absorption provenance. In our scheme deletions on the recursive stream are directly caused by deletions on the base stream. Hence, we only need to focus on deletion tuples generated from the base (B^Δ) stream. When we receive a deletion operation u , for each tuple t in the table P , we zero out the associated provenance of tuple u ($u.\text{pv}$) from the provenance expression of each t ($P[t]$), computed by BDD operation “restrict” [21] shown in Line 30. If the result is a provenance expression returning **false** (zero), a deletion operation on t is propagated to the next operator after removing its entry from P .

C. Join Operator

The *PipelinedHashJoin* must not only maintain two hash tables for its input relations (as is the norm), but also a hash table from each tuple to its current absorption provenance. It maintains this provenance state in a manner similar to the *Fixpoint*; due to space constraints we refer the reader to the extended technical report [22] for pseudocode. As insertions are received, provenance is updated for the associated tuple. The *difference* between the tuple’s existing and new provenance is computed; then the tuple is added to the appropriate hash table (if it does not already exist), and probed against the opposite relation. Deletion happens similarly, except that a tuple is removed from the join hash table *only if* its provenance becomes **false** (i.e., it is no longer derivable).

V. MINIMIZING PROPAGATION OF TUPLE PROVENANCE

With provenance, each time a given operator receives a new *derivation* of a tuple, it must typically propagate that tuple and derivation, in much the same fashion as it would a completely new tuple. If a tuple is derivable in many ways, it will be processed many times, just as a tuple might be propagated multiple times in a bag relation (versus a set). This increases the amount of work done in query processing, as well as the amount of state shipped across the network. Even worse, in the general case, a recursive query may produce an infinite number of possible derivations.

Fortunately, absorption helps in the last case. If a new tuple derivation is received whose provenance is completely absorbed, we do not need to propagate any information forward. We will reach a fixpoint when we can no longer derive any new results that affect the absorption provenance of any tuple in the result.

However, we must take additional steps to reduce the amount of state shipped by our distributed query processor nodes. Our goal is to reduce the number of *derivations* (provenance annotations) we propagate through the query plan and the network, while still maintaining the ability to handle deletions. Here we define a special stateful *MinShip* operator. *MinShip* replaces a conventional *Ship* operator, but maintains provenance information about the tuples produced by incoming updates. It always propagates the *rst* derivation of every tuple it receives, but simply buffers all subsequent derivations of the same tuple — merely updating their absorption provenance. By absorption, the stored provenance expression absorbs multiple derivations into a simpler expression.

Now if the original tuple derivation is deleted, *MinShip* responds by propagating forward any alternate derivations it has buffered — then it propagates that deletion operation. Additionally, depending on our preferences about state propagation, we can require the *MinShip* operator to propagate all of its buffered state periodically, e.g., when the buffer exceeds a capacity or a time threshold. By changing the batching interval or conditions, we can adjust how many alternate derivations are propagated through the query plan — a smaller interval will propagate more state, and a larger interval will propagate less state. In the extreme case, we can set the interval to infinity, resulting in what we term *lazy provenance propagation*. In the lazy case, alternate derivations of a tuple will only be propagated when they affect downstream results; this significantly reduces the cost of insertions. (In some cases it may slightly increase the cost of deletion propagation.)

MinShip’s internal state management again resembles that of the *Fixpoint* operator. Pseudocode is given in [22].

VI. MINIMIZING PROPAGATION OF STATE

Our third challenge is to minimize the amount of state (in terms of unique tuples, not just alternate derivations of the same tuple) that gets propagated from one node to the next. Given that aggregation is commonplace in network-based queries (as in most queries of Section II), we need a way to also suppress tuples that have no bearing on the output aggregate values. We adapt a technique called *aggregate selection* [23] to a streaming model, with a *windowed* aggregation

(group-by) operation [24]. We consider MIN, MAX, COUNT, and SUM functions³. In essence, the aggregate computation is split between a partial-aggregate operation that is used internally by stateful operators like the *Fixpoint* and *MinShip* to prune irrelevant state, and a final aggregation computation is done at the end over the partial aggregates' outputs. Our main contributions are to support revision (particularly deletion) of results within a windowed aggregation model, and to combine aggregate selection with minimal provenance shipping.

Our aggregate selection (*AggSel* for short) module (Algorithm 2) can be embedded within any operator that creates and ships state. (In our system, both *Fixpoint* and *MinShip* have calls to this module.) The module takes as input a stream U^Δ , a grouping key $\bar{u}k$, the number of aggregate functions n , and a set of aggregate functions $agg_1, agg_2, \dots, agg_n$. The module maintains a hash table H indexed on the grouping key $\bar{u}k$, which records all the buffered tuples met so far based on its grouping key values — this is necessary to support tuple deletion. A corresponding hash table P maps from each tuple to their absorption provenance. Another hash table B is maintained to record the value associated with each aggregate attribute agg_i , for the grouping key $\bar{u}k$. *AggSel* finally outputs a stream U'^Δ of the update tuples.

Each time *AggSel* receives a stream insertion (Lines 6–30), it inserts this tuple into the internal map H from group-by key $\bar{u}k$ to source tuple set. (If a tuple with the same value already exists in the set, then it simply updates the provenance P for the tuple.) Next, if the insertion affects the result of any aggregate attribute associated with $\bar{u}k$ — it changes the MIN or MAX value, or it revises the COUNT or SUM — the aggregation selection module will then propagate a *deletion* operation on the old aggregate value. After checking all the aggregate functions, if at least one of the aggregate values is affected, then it propagates this input insertion tuple as an *insertion*; if none of them is affected, it propagates nothing (see the loop starting at Line 15). Meanwhile, the module applies the change to its internal state.

Upon encountering a stream deletion or an expiration (Lines 30–56), *AggSel* checks whether the deletion has any affect on the derivability of the deleted tuple (Lines 31–33), and then whether any aggregate value associated with the group-by key $\bar{u}k$ is affected. If an aggregate value is modified (i.e., this deletion tuple at least partly determines the aggregate value), then *AggSel* traverses through the current version of buffered tuple table, computes the updated aggregate value, and propagates an *insertion* of the tuple with the new aggregate value. If any of the aggregate values is affected, then it propagates a *deletion*. Meanwhile, the module applies the change to its internal state.

VII. EXPERIMENTAL EVALUATION

We have developed a Java-based distributed query processor that implements all operators as described in Sections IV–VI. Our implementation utilizes the FreePastry 2.0.03 [26] DHT for data distribution, and JavaBDD v1.0b2 [21] as the BDD library for absorption provenance maintenance. Our

³AVERAGE can be derived from SUM and COUNT, as in [25].

Algorithm 2 Aggregate selection sub-module

AggSel($U^\Delta, \bar{u}k, n, agg_1, agg_2, \dots, agg_n$)

Inputs: Input stream U^Δ , grouping keys $\bar{u}k$, number of aggregate functions n , aggregate function $agg_1, agg_2, \dots, agg_n$.

Output: Stream U'^Δ .

```

1: Init hash map  $H: U(\bar{x})[\bar{u}k] \rightarrow \{U(\bar{x})\}$ 
2: Init hash map  $P: U(\bar{x}) \rightarrow$  provenance expressions over  $U(\bar{x})$ 
3: Init hash map  $B: U(\bar{x})[\bar{u}k] \rightarrow [1..n] * \{U(\bar{x})\}$ 
4: while not EndOfStream( $U^\Delta$ ) do
5:   Read an update  $u$  from  $U^\Delta$ 
6:   if  $u.type = \text{INS}$  then
7:     if  $H$  does not contain  $u.tuple$  then
8:        $H[u.tuple[\bar{u}k]] := u.tuple$ 
9:       Set  $P[u.tuple]$  to the provenance of  $u$ 
10:    else
11:      Add the provenance of  $u$  to  $P[u.tuple]$ 
12:    end if
13:    if  $oldPv \neq P[u.tuple]$  then
14:       $changed := \text{false}$ 
15:      for  $i = 1$  to  $n$  do
16:        if  $B$  does not contain  $u.tuple[\bar{u}k].i$  then
17:           $B[u.tuple[\bar{u}k]].i := u.tuple$ 
18:           $changed := \text{true}$ 
19:        else if  $u.tuple$  is better than  $B[u.tuple[\bar{u}k]].i$  for  $agg_i$  then
20:           $u'.tuple := B[u.tuple[\bar{u}k]].i$ 
21:           $u'.type := \text{DEL}$ 
22:          Set provenance of  $u'$  to  $P[B[u.tuple[\bar{u}k]].i]$ 
23:          Output  $u'$ 
24:           $B[u.tuple[\bar{u}k]].i := u.tuple$ 
25:           $changed := \text{true}$ 
26:        end if
27:      end for
28:      if  $changed$  then Output  $u$ 
29:    end if
30:  else if  $H$  contains  $u.tuple$  then
31:     $oldPv := P[u.tuple]$ 
32:    Remove the provenance of  $u$  from  $P[u.tuple]$ 
33:    if  $P[u.tuple]$  indicates no derivability then
34:      Remove  $u.tuple$  from  $P$ 
35:      Remove  $u.tuple[\bar{u}k]$  from  $H$ 
36:    end if
37:    if  $oldPv \neq P[u.tuple]$  then
38:       $changed := \text{false}$ 
39:      for  $i = 1$  to  $n$  do
40:        if  $B[u.tuple[\bar{u}k]].i = u.tuple$  then
41:           $changed := \text{true}$ 
42:          Remove  $u.tuple$  from  $B[u.tuple[\bar{u}k]].i$ 
43:          for each tuple  $t$  in  $H[u.tuple[\bar{u}k]]$  do
44:            if  $B[u.tuple[\bar{u}k]].i = \text{null}$  or  $t$  is better than
               $B[u.tuple[\bar{u}k]].i$  for  $agg_i$  then
45:               $B[u.tuple[\bar{u}k]].i := t$ 
46:            end if
47:          end for
48:           $u'.tuple := B[u.tuple[\bar{u}k]].i$ 
49:           $u'.type = \text{INS}$ 
50:          Set provenance of  $u'$  to  $P[B[u.tuple[\bar{u}k]].i]$ 
51:          Output  $u'$ 
52:        end if
53:      end for
54:      if  $changed$  then Output  $u$ 
55:    end if
56:  end if
57: end while

```

experiments are carried out on two clusters: a 16-node cluster consisting of quad-core Intel Xeon 2.4GHz PCs with 4GB RAM running Linux 2.6.23, and an 8-node cluster consisting of dual-core Pentium D 2.8GHz PCs with 2GB RAM running Linux 2.6.20. The machines are internally connected within each cluster via a high-speed Gigabit network, and the clusters are interconnected via a 100Mbps network shared with the rest of campus traffic. Our default setting involves 12 nodes from the first cluster; when we scale up, we first use all 16 nodes from this cluster, then add 8 more nodes from the second

cluster to reach 24 nodes. All experimental results are averaged across 10 runs with 95% confidence intervals included.

A. Experimental Setup

We studied two query workloads taken from our use cases:

Workload 1: Declarative networks. Our query workloads consist of the *reachable query* and the *shortest-path query* (Section II). As input to these queries, we use simulated Internet topologies generated by GT-ITM [27], a package that is widely used to model Internet topologies. By default we use GT-ITM to create “transit-stub” topologies consisting of eight nodes per stub, three stubs per transit node, and four nodes per transit domain. In this setup, there are 100 nodes in the network, and approximately 200 bidirectional links (hence 400 *link* tuples in our case). Each input *link* tuple contains *src* and *dst* attributes, as well as an additional *latency* cost attribute. Latencies between transit nodes are set to 50 ms, the latency between a transit and a stub node is 10 ms, and the latency between any two nodes in the same stub is 2 ms. To emulate network connectivity changes, we add and delete *link* tuples during query execution.

Workload 2: Sensor networks. Our second workload consists of region-based sensor queries executed over a simulated 100m by 100m grid of sensors, where the sensors report data to their local query processing node. We include 5 “seed” groups, each initialized to contain a single device. Our recursive view “activeRegion” finds contiguous (within k meters, where by default $k=20$) triggered nodes and adds them to the group — or removes them if they are no longer triggered. Based on that, we can compute the the largest such active region.

```
with recursive activeRegion(regionid,sensorid) as
( select M.regionid, S.sensorid
  from sensor S, coordSensor M, isTriggered T
  where M.sensorid = S.sensorid
        and S.sensorid = T.sensorid
  union
    select A.regionid, S2.sensorid
    from sensor S1, sensor S2, activeRegion A,
         isTriggered T
    where distance(S1.coord, S2.coord) < k
          and S1.sensorid = A.sensorid and
          S1.sensorid = T.sensorid )

create view regionSizes(regionid,size) as
(select regionid, count(sensorid)
  from activeRegion
  group by regionid)

create view largestRegion(size) as
(select max(size) from regionSizes)

create view largestRegions(regionid) as
(select R.regionid
  from regionSizes R, largestRegion L
  where R.size = L.size)
```

Initially all the seed sensors are triggered. Also we trigger half of the sensors in the network to study the effects of insertions, and then randomly remove them to study the effects of deletions. Note that while the input topology simulates a grid-based sensor topology, the queries are executed over our real distributed query processor implementation.

Our evaluation metrics are as follows:

- **Per-tuple provenance overhead (B):** the space taken by the provenance annotations on a per-tuple basis.

- **Communication overhead (MB):** the total size of communication messages processed by each distributed node for executing a distributed query to completion.
- **Per-node state within operators (MB):** the total overhead of the state maintained inside operators on each distributed node.
- **Convergence time (s):** the time taken for a distributed query to finish execution on all distributed nodes.

B. Incremental View Maintenance with Provenance

Our first set of experiments focuses on measuring the overhead of incremental view maintenance. Using the *reachable query* as a starting point, we compare three different schemes: the traditional *DRed* recursive view maintenance strategy, *relative provenance* [8] where each tuple is annotated with information describing derivation “edges” from other tuples, and our proposed *absorption provenance*. We also consider two schemes for propagating provenance: an *eager* strategy (propagate state from *MinShip* once a second) and a *lazy* one (propagate state only when necessary).

Insertions-only workload: We first measure the overhead of maintaining provenance, versus normal set-oriented execution. Figure 7 shows the performance of the *reachable query*, where the Y-axis shows our four evaluation metrics, and the X-axis shows the fraction of links inserted, in an incremental fashion, up to the maximum of 400 link tuples required to create the 100-node GT-ITM topology. Given an insertion-only workload, *DRed* has the best overall performance, since no provenance needs to be computed or maintained. Relative provenance encodes more information than absorption provenance, resulting in larger tuple annotations, more communication, and more operator state. Relative provenance with eager propagation (*Relative Eager*) did not converge within 5 minutes for insertion ratios of 0.75 or higher; hence, we only show lazy propagation (*Relative Lazy*) for the remaining graphs. Eager propagation with absorption provenance (*Absorption Eager*) also is costly due to the overhead of sending every new derivation of a tuple. Lazy propagation of absorption provenance (*Absorption Lazy*) is clearly the most efficient of the provenance schemes.

Insertions-followed-by-deletions workload: Our next set of experiments separately measures the overhead of deletions: here provenance becomes useful, whereas in the insertion case it was merely an overhead. (One can estimate the performance over a mixed workload by considering the relative distribution of insertions vs. deletions and looking at the overheads on each component.) Given the same 100-node topology, after inserting all the *link* tuples as above, we then delete *link* tuples in sequence. Each deletion occurs in isolation and we measure the time the query results take to converge after every deletion is injected. Figure 8 shows that *DRed* is prohibitively expensive for deletions when compared to our absorption provenance schemes: it is an order of magnitude more expensive in both communication overhead and execution time. Relative provenance wins versus *DRed* in communication cost and convergence time because it does not over-delete and re-derive. However, its performance is far

worse than absorption provenance, and it also incurs more per-tuple overhead and operator state. Relative provenance relies on graph traversal operations to determine derivability from base tuples (see [8]), and thus is expensive in a distributed setting. In contrast, absorption provenance directly encodes whether a derived tuple is dependent on a base tuple. Overall, absorption provenance is the most efficient method in deletion handling, and consequently ships fewer tuples than the other methods. Taking both insertions and deletions into account, *Absorption Lazy* has the best mix of performance.

Region-based sensor query: The *region* query is computed over a different topology from the *reachable* case, and it exhibits slightly different update characteristics. Still, as we see in Figure 9, which measures performance with the insertion workload described earlier in the experimental setup, performance follows similar patterns. (The overhead is lower across each of the four metrics, since the network is smaller here and neighbors are within closer proximity.) Under deletion workloads, the trends shown by the *region* query also closely mirror that of the *reachable* query and those graphs are shown in [22]. Since the queries exhibit similar performance, we focus on the *reachable* query for our remaining experiments.

C. Scalability

Next we consider how our absorption provenance schemes scale, with respect to inputs and to query processing nodes.

Scaling Data. We increase the number of input link tuples, by increasing the average number of transit nodes in the GT-ITM generated topology. We considered two network topologies: each node in the *dense* topology has four links (as in our default setting) on average, whereas the *sparse* setting has two. Figure 10 shows the insertion-only workload.⁴ We observe that the dense network is more costly to evaluate than the sparse network: there are far more derivations. Here, lazy propagation is essential: *Eager Dense* did not complete after 5 minutes on a 800-link network, whereas *Lazy Dense* finished in under 5 seconds.

Increasing Query Processing Nodes. Next, we increase the number of query processing nodes to up to 24 machines, while keeping the input dataset constant. Figure 11 shows the results. Per-tuple provenance overhead increases, then eventually levels off, as the number of nodes increases: each node will now process fewer tuples, and the opportunities of absorption and buffering are reduced. More query processors lead to a reduction in query execution latency, per-node communication overhead, and per-node operator state. The increase of latency between 16 and 24 nodes is due to the lower-bandwidth connection between our two subnets. In all cases, DRed incurs higher communication overhead and takes longer to complete than our approach.

D. Multi-aggregate Selection

Figure 12 shows the effectiveness of aggregate selections over the dense and sparse topology of 100 nodes. We experiment with two extensions of the *shortest path* query presented

⁴We further experimented with deleting an additional 20% of the links. Observations were similar and we omit graphs due to space constraints.

in Section II: *Multi AggSel* computes two aggregates (one for shortest path and the other for cheapest cost path); *Single AggSel* minimizes only based on the cheapest cost path. We observe that aggregate selections are most effective in dense topologies, and *Multi AggSel* costs only half as much as *Single AggSel* due to aggressive pruning of the two aggregates simultaneously. Without the use of aggregate selections, all queries are prohibitively expensive, and do not complete within 5 minutes for dense topologies.

E. Summary of Results

We summarize our experimental results with reference to the contributions of this paper as outlined in Section III-C.

- *Absorption provenance* (Section IV) incurs some overhead during insertions and consumes increased memory, compared to traditional schemes such as DRed. That increase is offset by huge improvements in communication overhead and execution times when deletions are part of the workload. Moreover, our concise representation of data provenance is far more efficient than an encoding of relative provenance. Most applications (both for declarative networking and sensor monitoring) include time-based expiration for state, and hence require frequent deletion processing.
- Our second technique, *lazy propagation* of derivations (Section V) using the *MinShip* operator, reduces traffic when there are multiple possible derivations. Lazy propagation results in significant communication cost savings. Given the dense network topology with 800 links and many alternative routes, lazy propagation resulted in 5-second running times, versus 5 minutes for eager propagation in the same network.
- Our third technique of multiple *aggregate selections* results in minimal propagation of tuples during query evaluation (Section VI). A dense network produces several alternative routes, and aggregate selections are especially effective in this setting, resulting in at least an order of magnitude reduction in communication cost and execution times. While the benefits of aggregate selections have been explored previously in centralized settings, our main contribution here was the extension to a stream model, including support for deletions, and validating that similar benefits are observed in a distributed recursive stream query processor.

VIII. RELATED WORK

Stream query processing has been popular in the recent database literature, encompassing sensor network query systems [4], [5] as well as Internet-based distributed stream management systems [28], [29], [30]. To the best of our knowledge, none of these systems support recursive queries. Distributed recursive queries have been proposed as a mechanism for managing state in declarative networks. Our work formalizes aspects of soft-state management and significantly improves the ability to maintain recursive views. Our distributed recursive view maintenance techniques are applicable to other networked environments, particularly programming abstractions for region-based computations in sensor networks [10], [11].

Provenance (also called lineage) has often been studied to help “explain” why a tuple exists [17] or to assign a ranking or score [8], [31]. Lineage was studied in [18] as a

means of maintaining data warehouse data. Our absorption provenance semiring in [19] (which provides a theoretical provenance framework, but does not consider implementabil-

Fig. 12. Aggregate selections performance on *shortestPath* and *cheapestCostPath* query

ity). We specialized it for maintenance of derived data in recursive settings. Our approach improves over the counting algorithm [15] which does not support recursion. We have experimentally demonstrated benefits versus DRed [15] and maintenance based on relative provenance [8] (both of which were developed for non-distributed query settings).

IX. CONCLUSIONS AND FUTURE WORK

We have proposed novel techniques for distributed recursive stream view maintenance. Our work is driven by emerging applications in declarative networking and sensor monitoring, where distributed recursive queries are increasingly important. We demonstrated that existing recursive query processing techniques such as DRed [15] are not well-suited for the distributed environment. We then showed how absorption provenance could be used to encode tuple derivability in a compact fashion, then incorporated into provenance-aware operators that are bandwidth efficient and avoid propagating unnecessary information, while maintaining correct answers.

Our work is proceeding along several fronts. Since our experimental results have demonstrated the effectiveness of techniques, we are working towards deploying our system in both declarative networking and sensor network domains. We intend not only to support efficient distributed view maintenance, but also to utilize the provenance information to enforce decentralized trust policies, and perform real-time network diagnostics and forensic analysis. We also hope to explore opportunities for adaptive cost-based optimizations based on the query workload, network density, network connectivity, rate of network change, etc.

ACKNOWLEDGMENTS

This work was funded by NSF grants CNS-0721845, CNS-0721541, IIS-0812270, IIS-0447972, and IIS-0713267, and a DARPA Computer Science Study Group grant. We thank the anonymous reviewers for their suggestions.

REFERENCES

- [1] H. Balakrishnan, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica, "Looking Up Data in P2P Systems," *Communications of the ACM*, Vol. 46, No. 2, Feb. 2003.
- [2] D. Chu, L. Popa, A. Tavakoli, J. M. Hellerstein, P. Levis, S. Shenker, and I. Stoica, "The design and implementation of a declarative sensor network system," in *SenSys*, New York, NY, USA, 2007.
- [3] B. T. Loo, J. M. Hellerstein, I. Stoica, and R. Ramakrishnan, "Declarative routing: extensible routing with declarative queries," in *SIGCOMM*, 2005.
- [4] A. J. Demers, J. Gehrke, R. Rajaraman, A. Trigoni, and Y. Yao, "The Cougar project: a work-in-progress report," *SIGMOD Record*, vol. 32(3), 2003.
- [5] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong, "Design of an acquisitional query processor for sensor networks," in *SIGMOD*, 2003.
- [6] D. Narayanan, A. Donnelly, R. Mortier, and A. Rowstron, "Delay aware querying with Seaweed," in *VLDB*, 2006.
- [7] R. Huebsch, J. M. Hellerstein, N. Lanham, B. T. Loo, S. Shenker, and I. Stoica, "Querying the Internet with PIER," in *VLDB*, 2003.
- [8] T. J. Green, G. Karvounarakis, Z. G. Ives, and V. Tannen, "Update exchange with mappings and provenance," in *VLDB*, 2007, amended version available as Univ. of Pennsylvania report MS-CIS-07-26.
- [9] B. T. Loo, T. Condie, M. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica, "Declarative Networking: Language, Execution and Optimization," in *Proc. SIGMOD*, June 2006.
- [10] M. Welsh and G. Mainland, "Programming sensor networks using abstract regions," in *NSDI*, March 2004.
- [11] K. Whitehouse, C. Sharp, E. Brewer, and D. Culler, "Hood: a neighborhood abstraction for sensor networks," in *MASN*, 2004.
- [12] I. Balbin and K. Ramamohanarao, "A generalization of the differential approach to recursive query evaluation," *J. Log. Program.*, vol. 4, no. 3, 1987.
- [13] F. Bancilhon, D. Maier, Y. Sagiv, and J. D. Ullman, "Magic sets and other strange ways to implement logic programs," in *PODS*, 1986.
- [14] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: a scalable and robust communication paradigm for sensor networks," in *MobiCom*, 2000.
- [15] A. Gupta, I. S. Mumick, and V. S. Subrahmanian, "Maintaining views incrementally," in *SIGMOD*, 1993.
- [16] S. Raman and S. McCanne, "A model, analysis, and protocol framework for soft state-based communication," in *Proceedings of ACM SIGCOMM Conference on Data Communication*, 1999, pp. 15–25.
- [17] P. Buneman, S. Khanna, and W. C. Tan, "Why and where: A characterization of data provenance," in *ICDT*, 2001.
- [18] Y. Cui, "Lineage tracing in data warehouses," Ph.D. dissertation, Stanford University, 2001.
- [19] T. J. Green, G. Karvounarakis, and V. Tannen, "Provenance semirings," in *PODS*, 2007.
- [20] R. Bryant, "Graph-based algorithms for boolean function manipulation," *IEEE Transactions on Computers*, vol. 35, no. 8, pp. 677–691, 1986.
- [21] John Whaley, "Javabdd library," <http://javabdd.sourceforge.net>.
- [22] M. Liu, N. E. Taylor, W. Zhou, Z. G. Ives, and B. T. Loo, "Recursive computation of regions and connectivity in networks," University of Pennsylvania, Tech. Rep. MS-CIS-08-32, 2008.
- [23] S. Sudarshan and R. Ramakrishnan, "Aggregation and Relevance in Deductive Databases," in *Proceedings of VLDB Conference*, 1991.
- [24] R. Motwani, J. Widom, A. Arasu, B. Babcock, S. Babu, M. Datar, G. Manku, C. Olston, J. Rosenstein, and R. Varna, "Query processing, resource management, and approximation in a data stream management system," in *CIDR*, 2003.
- [25] S. Chaudhuri and K. Shim, "Including group-by in query optimization," in *VLDB*, 1994.
- [26] A. Rowstron and P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems," in *Middleware*, Nov. 2001, pp. 329–350.
- [27] GT-ITM, "Modelling topology of large networks," <http://www.cc.gatech.edu/projects/gtintm/>.
- [28] D. J. Abadi, D. Carney, U. Cetintemel, M. Cherniack, C. Convey, S. Lee, M. Stonebraker, N. Tatbul, and S. Zdonik, "Aurora: a new model and architecture for data stream management," *VLDB J.*, vol. 12(2), August 2003.
- [29] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. A. Shah, "TelegraphCQ: Continuous dataflow processing for an uncertain world," in *CIDR*, 2003.
- [30] A. Arasu, S. Babu, and J. Widom, "The CQL continuous query language: semantic foundations and query execution," *VLDB J.*, vol. 15, no. 2, 2006.
- [31] O. Benjelloun, A. D. Sarma, A. Y. Halevy, and J. Widom, "ULDBs: Databases with uncertainty and lineage," in *VLDB*, 2006.