Improved Decremental Algorithms for Maintaining Transitive Closure and All-Pairs Shortest Paths ⋆

Surender Baswana [a,1] Sandeep Sen [a,*,2] Ramesh Hariharan [b]

[a] *Dept. of Comp. Sc. and Engg., Indian Institute of Technology Delhi,Hauz Khas, New Delhi, India*

[b] *Dept. of Comp. Sc. and Automation,Indian Institute of Science, Bangalore, India*

Abstract

This paper presents improved algorithms for the following problem : Given an unweighted digraph $G(V, E)$ and an on-line sequence of shortest-path/reachability queries interspersed with edge-deletions, develop a data-structure that can answer each query on-line in optimal time, and can be updated efficiently after each edge-deletion.

The central idea underlying our algorithms is a scheme for implicitly storing all-pairs reachability/shortest-path information, and an efficient way to maintain this information.

Our algorithms are randomized and have one-sided inverse polynomial error for query.

Key words: BFS tree, dynamic, graph, transitive closure, shortest paths

⋆ Preliminary version of this paper appeared in 34th *ACM Symposium on Theory of Computing (STOC)*, May 19-21, 2002, Montreal,Quebec, Canada.

* Corresponding author.

Email addresses: sbaswana@cse.iitd.ernet.in (Surender Baswana), ssen@cse.iitd.ernet.in (Sandeep Sen), ramesh@csa.iisc.ernet.in (Ramesh Hariharan).

[1] Work was supported in part by a fellowship from Infosys Technologies Ltd., Bangalore.

[2] Work was supported in part by an IBM UPP award

1 Introduction

The following two problems are among the most fundamental algorithmic problems in computer science :

- **Transitive Closure :** Process a given directed graph $G(V, E)$ so that any query of the form, "*Is there a path from u to v in the graph?*", can be answered efficiently.
- **All-Pairs Shortest Paths Problem :** Process a given directed unweighted graph $G(V, E)$ so that any query of the form, "*Report the shortest path (or distance) from vertex u to a vertex v?*", can be answered efficiently.

There exist classical algorithms that take $O(mn)$ time for the two problems so that any query can be answered in $O(1)$ time. There also exist algorithms based on fast matrix multiplication that achieve sub-cubic time for the two problems. In particular, using the fastest known matrix multiplication algorithms (Coppersmith and Winograd [2]), the best bound on computing transitive closure is $O(n^{2.376})$, whereas for the all-pairs shortest paths problem, it is $O(n^{2.575})$ [15].

There are many applications in communication networks, incremental parser generation [10] and relational databases augmented with transitive closure [11,7] that require efficient solutions of the above problems for a dynamic graph. In a dynamic graph problem, an initial graph is given, followed by an on-line sequence of queries interspersed with updates which can be insertion or deletion of edges. We have to carry out the updates and answer the queries on-line in an efficient manner. Each query has to be answered with respect to the present state of the graph, i.e., incorporating all updates preceding the query. One trivial solution for solving such a dynamic graph problem is that we run the static graph algorithm after every update. The goal of a dynamic graph algorithm is to update the solution efficiently after the dynamic changes, rather than having to re-compute it from scratch each time.

For every static graph problem, there exists its dynamic counterpart. We can classify dynamic graph problems according to the types of updates allowed. A problem is said to be *fully dynamic* if the update operations include both insertions and deletions of edges. A problem is called *partially dynamic* if only one type of update, either insertion or deletions, is allowed. If only insertions are allowed, the problem is called *incremental*; if only deletions are allowed, it is called *decremental*.

We present efficient decremental algorithms for maintaining APSP and transitive closure in an unweighted digraph. The algorithms are Monte-Carlo with one sided error (the probability of error being inverse polynomial).

1.1 Previous work and our contribution

Transitive Closure :

Poutre and Leeuwen [14] gave a decremental algorithm for maintaining transitive closure with $O(m)$ amortized update time per edge deletion and answering each query in $O(1)$ time. Demetrescu and Italiano [4] gave a decremental algorithm for the problem that requires $O(n^3/m)$ amortized update time which is better for dense graphs. For an initial complete graph, the algorithm achieves $O(n)$ amortized update time per edge deletion [3], but for sparse graphs, the update time can be even $O(n^2)$. It can be seen that a combination of these two algorithms yields an upper bound of $O(n^{\frac{3}{2}})$ on the update time while keeping $O(1)$ query time. Henzinger and King [9] gave a decremental randomized algorithm that achieves amortized update time $(n \log^2 n)$ but at the expense of increased query time of $O(\frac{n}{\log n})$. The query has one sided inverse polynomial error, i.e., the answer to any query may be incorrect in the sense that it may not report a path when there exists one. However, the probability of an error can be made smaller than $\frac{1}{n^c}$ for arbitrary $c > 0$. Subsequently we will use w.h.p. to denote probability exceeding $1 - \frac{1}{n^c}$ for any $c > 0$.

In this paper, we present an efficient algorithm that achieves $O(1)$ query time w.h.p. and requires $O(n \log^2 n + \frac{n^2}{\sqrt{m}} \log n)$ amortized update time per edge-deletion. Our algorithm achieves an improvement in the update time compared to the deterministic algorithms while ensuring $O(1)$ query time. By suitably combining with [14], our algorithm achieves an improved upper bound of $O(n^{\frac{4}{3}} \sqrt[3]{\log n})$ on amortized update time per edge deletion for the problem of maintaining transitive closure with optimal query time.

All-pairs shortest paths :

Demetrescu and Italiano [5] gave a decremental algorithm for maintaining all-pairs shortest paths under deletion of edges that achieves $O(\frac{n^3}{m} \log^3 n)$ amortized update time while achieving optimal query time w.h.p. Their algorithm improves the previous $O(n^2)$ update time of [12] for dense graphs.

We present two decremental algorithms for the all-pairs shortest path problem. For distance reporting problem, we give a simpler combinatorial algorithm that requires $O(\frac{n^3}{m} \log^2 n)$ amortized update time while achieving $O(1)$ query time. For shortest path reporting problem, we use the idea of filtering search in a novel way to design an algorithm that achieves $O(\min(n^{\frac{3}{2}} \sqrt{\log n}, \frac{n^3}{m} \log^2 n))$ update time while achieving optimal query time. Hence we reduce the upper bound on the worst case amortized update time for the problem of maintaining all-pairs shortest paths under deletion of edges by a factor of $O(\sqrt{\frac{n}{\log n}})$.

Maintaining All-Pairs Reachability		
Query	Previous Update Time	Improved Update Time
Is v reachable from u ?	$O\left(n^{\frac{3}{2}}\right)$	$O\left(n^{\frac{4}{3}}\sqrt[3]{\log n}\right)$
report a path from u to v		
Maintaining All-Pairs $(1+\epsilon)$-Approximate Shortest Paths		
Query	Previous Update Time	Improved Update Time
report $(1+\epsilon)$-approx. distance	*none*	$O\left(\frac{n\log n}{\epsilon^2}+\frac{n^2\sqrt{\log n}}{\epsilon\sqrt{m}}\right)$
report $(1+\epsilon)$-approx. shortest path		
Maintaining All-Pairs Shortest Paths		
Query	Previous Update Time	Improved Update Time
report the distance from u to v	$O\left(\frac{n^3}{m}\log^3 n\right)$	$O\left(\frac{n^3}{m}\log^2 n\right)$
report the shortest path from u to v	$O(\frac{n^3}{m}\log^3 n)$	$O\left(\min\left(n^{\frac{3}{2}}\sqrt{\log n}, \frac{n^3}{m}\log^2 n\right)\right)$

Table 1
All the update bounds (old and new) are amortized. Throughout this paper, our query times are optimal but have one sided inverse polynomial error.

Next, we present an efficient decremental algorithm that offers a trade-off between update time and approximation factor of the shortest path. For maintaining $(1+\epsilon)$-approximate all-pairs shortest paths, our algorithm achieves $O(\frac{n\log n}{\epsilon^2}+\frac{n^2\sqrt{\log n}}{\epsilon\sqrt{m}})$ amortized update time per edge-deletion for arbitrarily small $\epsilon > 0$. We summarize our results in Table 1.

It may be noted that our algorithms are simple to implement and do not make use of any sophisticated matrix multiplication algorithms. We show that the space requirement of our data structures is $O(n^2)$.

2 Overview of our algorithms

2.1 Notations

Given an unweighted directed graph $G(V, E)$, $v \in V$, and a positive integer d, we define the following notations to be used frequently in this paper.

- *out_tree*(v) : a complete breadth-first-search (BFS) tree rooted at $v \in V$ in the graph $G(V, E)$.
- *in_tree*(v) : a complete BFS tree rooted at $v \in V$ in the graph formed by reversing all the edge-directions in G.
- *out_tree*(v, d) : a BFS tree of depth d rooted at the vertex v in the graph $G(V, E)$.
- *in_tree*(v, d) : a BFS tree of depth d rooted at the vertex v in the graph formed by reversing all the edge-directions in G.
- (*in_tree*, *out_tree*)(v, d) : a pair of *in_tree*(v, d) and *out_tree*(v, d).
- $W(S, d)$: $\{$(*in_tree*, *out_tree*)(v, d) $|v \in S\}$, for a given $S \subset V$.
- S_r : A random set of vertices formed by picking each vertex randomly independently with probability $\frac{\ln n}{r}$. The expected size of the set S_r is $\frac{n \ln n}{r}$.

It can be seen that the *out_tree*(v, d) stores all the vertices that lie within distance d from v. Also note that each vertex in a BFS tree is assigned to a level which is also its distance from the root. The level information of a vertex in a BFS tree rooted at v can be used to report the distance of the vertex from v. Hence, in addition to maintaining reachability information from v, the *out_tree*(v, d) can also be used for maintaining the shortest paths to all the vertices lying within distance d from v. In the same way, the *in_tree*(v, d) can be used for maintaining reachability information as well as the shortest path to v from all the vertices for whom v lies within distance d.

While edges are being deleted in a graph, the levels of vertices in a BFS tree may fall. Thus maintaining a BFS tree of depth d, after an edge deletion, involves finding the set of vertices whose level has fallen, eliminating those from this set whose level has fallen beyond d, and assigning the rest of the vertices (falling to level $\leq d$) to their new levels.

The best known algorithm for maintaining an *in_tree* or an *out_tree* is given by Henzinger and King [9] along the lines of [6].

Lemma 1 *[9] Given an unweighted digraph $G(V, E)$, a vertex $v \in V$ and a positive integer d, an* in_tree*(v, d) (or an* out_tree*(v, d)) can be maintained in $O(d)$ amortized update time per edge deletion.*

2.2 Main Idea

The simplest way of maintaining all-pairs reachability (shortest paths) in an unweighted digraph is to maintain an *out_tree* up to depth n from each vertex. It follows from the Lemma 1 that the amortized update time required by this approach (maintaining n *out_trees* of depth n) will be $O(n^2)$ per edge deletion.

In order to develop efficient decremental algorithm for transitive closure (and all-pairs shortest paths), we explore alternate schemes of maintaining all-pairs reachability (shortest paths). In this endeavor, we present schemes that maintains reachability (and exact/approximate shortest path) information *implicitly*.

Consider a pair of vertices $u, w \in V$ such that there is a path from vertex u to w, and let v be any intermediate vertex on the path. It can be seen that the vertex u belongs to *in_tree* rooted at v, and the vertex w belongs to *out_tree* rooted at v. Thus combined together the *in_tree*, *out_tree* rooted at v stores the path from u to w implicitly (the complete path and its length can be retrieved by querying *in_tree*(v) and *out_tree*(v)). In other words, the (*in_tree*(v),*out_tree*(v))-pair acts as a *witness of reachability* from u to w. Analogously if the vertex v lies on the shortest path from u to w, the (*in_tree*(v),*out_tree*(v)) pair acts as *a witness of shortest path* from u to w.

The above mentioned scheme of keeping reachability (shortest path) information implicitly suggests that in order to maintain all-pairs reachability (shortest paths) under deletion of edges, it suffices to maintain a witness (if one exists) of reachability (shortest path) for each vertex-pair (u, w). In the following section, we design efficient algorithms for maintaining witnesses of all-pairs reachability and shortest paths corresponding to paths of length in an interval $[d/2, d]$ for a given $d \leq n$. These algorithms form the basis for developing efficient decremental algorithms for the main problems as follows.

Maintaining transitive closure : The scheme of maintaining reachability implicitly (by keeping witnesses) proves to be efficient for maintaining all-pairs reachability corresponding to long paths. On the other hand, the scheme of maintaining reachability explicitly (by keeping *out_tree* from every vertex) proves to be efficient for maintaining reachability corresponding to short paths. We combine the two strategies together to achieve improved update time of $O(n \log^2 n + \frac{n^2}{\sqrt{m}}\sqrt{\log n})$ per edge deletion.

Maintaining all-pairs approximate shortest paths : Analogous to witness of reachability and shortest paths, we propose the terminology of witness of *approximate shortest paths*. It turns out that our data structure for maintaining transitive closure can be suitably adapted to maintain all-pairs

$(1+\epsilon)$-approximate shortest paths [3] in unweighted digraphs.

Maintaining all-pairs exact shortest paths : For the problem of maintaining all-pairs shortest distances, the strategy of maintaining witness of shortest path for every vertex-pair leads to achieving $O(\frac{n^3}{m}\log^2 n)$ update time. For the case when query is to report the shortest path, we use idea of filtering search [1] to reduce the update time further to $O(\min(n^{\frac{3}{2}}\sqrt{\log n}, \frac{n^3}{m}\log^2 n))$ which is bounded by $O(n^{\frac{3}{2}}\sqrt{\log n})$ for all graphs.

3 Maintaining witnesses of reachability/shortest-paths for all-pairs of vertices separated by distance $[\frac{\mathbf{d}}{\mathbf{2}}, \mathbf{d}]$

In this section we design efficient algorithms for maintaining reachability and shortest paths for all the vertex-pairs (u, w) such that there is a path from u to w of length $\in [\frac{d}{2}, d]$. For the reachability problem, the algorithm will maintain a witness of reachability for every vertex-pair (u, w) if there is path of length $\in [\frac{d}{2}, d]$ from u to w. For the all-pairs shortest path problem, the algorithm will maintain a witness of the shortest path for every vertex-pair (u, w) among all paths from u to w of length $\in [\frac{d}{2}, d]$.

Let $S \subset V$, and let $W(S, d)$ be the set of pairs of *in_tree* and *out_tree* of depth d built on each vertex of the set S (see the notations defined in section 2.1). We begin with the design of efficient algorithms for maintaining all-pairs reachability and shortest paths with tree-pairs from the set $W(S, d)$ as witnesses, i.e., the set of paths under consideration is restricted to the paths captured by the tree-pairs of the set $W(S, d)$ only. The algorithm can be used for maintaining reachability (shortest paths) for all-pairs of vertices separated by distance $\leq [\frac{d}{2}, d]$, if we choose $S = V$. It turns out that the update time of the algorithm is directly proportional to the size of the witness set. Subsequently, to improve the update time, we reduce the number of witnesses, i.e. $|S|$ using random sampling. This observation was exploited by Henzinger and King [9] for designing a decremental algorithm to maintain all-pairs reachability with $O(n\log^2 n)$ update time at the expense of $O(\frac{n}{\log n})$ query time. We extend this approach to its full potential for maintaining all-pairs shortest paths and transitive closure with optimal query time.

3.1 Maintaining all-pairs reachability with respect to a witness set

Let $W(S, d)$=$\{$*in_tree*(v, d), *out_tree*$(v, d)|v \in S\}$ be the set of witnesses kept in a list, also denoted by $W(S, d)$. There is a path from u to w of length $\leq d$

[3] Path with length at most $(1+\epsilon)$ times the length of the shortest path

passing through some $v \in S$ if u and w lie respectively in *in_tree*(v,d) and *out_tree*(v,d). Thus it takes $O(1)$ time to test whether (*in_tree*, *out_tree*)(v,d) is a witness of reachability from u to w. Finding out if there is any witness in the set $W(S,d)$ for reachability from u to w will require querying each (*in_tree*, *out_tree*)$(v,d) \in W(S,d)$. Thus the query time will be $O(|S|)$. In order to achieve $O(1)$ query time, we maintain a witness matrix M. For each pair (u,w), at every stage $M[u,w]$ points to some vertex $v \in S$ if (*in_tree*, *out_tree*)$(v,d) \in W(S,d)$ is a witness of reachability from u to w.

We initialize the matrix M as follows : The matrix M has all entries pointing to **null** initially. For every (*in_tree*, *out_tree*)$(v,d) \in W(S,d)$, we do the following. For each $u \in$ *in_tree*(v,d) and for every $w \in$ *out_tree*(v,d), we update $M[u,w]$ to point to v if it was pointing to **null** previously. Thus it requires $O(n^2)$ time per tree-pair from the set $W(S,d)$ to initialize the witness matrix M.

Notice that as the edges are being deleted, a pair (*in_tree*, *out_tree*)$(v,d) \in W(S,d)$ may cease to be witness of reachability for a vertex pair (u,w). This happens when either u ceases to belong to *in_tree*(v,d) or the vertex w ceases to belong to *out_tree*(v,d). Thus we have to update the entries of matrix appropriately after an edge deletion.

To maintain validity of entries of matrix M after an edge deletion, we process each (*in_tree*, *out_tree*)$(v,d) \in W(S,d)$ (starting from the first tree-pair in the list $W(S,d)$) as follows. We update *in_tree*(v,d) and *out_tree*(v,d) for the edge deletion, and compute the sets of vertices X and Y that cease to belong to *in_tree*(v,d) and *out_tree*(v,d) respectively. This step takes $O(|X| + |Y|)$ time in addition to $O(d)$ amortized time for updating the two BFS trees at v. For every $u \in X$, we find the set s_u of vertices such that $M[u,w]$ is v for each $w \in s_u$. By inspecting the row $M[u,_]$ it requires $O(n)$ time per $u \in X$ to find s_u. It can be seen that (*in_tree*, *out_tree*)$(v,d) \in W(S,d)$ has ceased to be a witness of reachability from u to each $w \in s_u$, and so a new witness is to be searched for. We search the list $W(S,d)$ starting from the tree-pair following (*in_tree*, *out_tree*)(v,d), and find a witness of reachability from u to w (spending $O(1)$ time per tree-pair in $W(S,d)$). If we reach the end of the list $W(S,d)$, we update $M[u,w]$ to point to **null**, otherwise we update $M[u,w]$ to point to the root of the new witness found. In a similar way, we process the set Y of vertices that cease to belong to *out_tree*(v,d).

The following invariant is maintained throughout the series of edge deletions.
I^τ : *$M[u,w]$ points to the first tree-pair in the list $W(S,d)$ that is a witness of reachability from u to w.*

The invariant holds just after the initialization of M. It can be easily verified that the procedure described above preserves the invariant after each edge deletion. By induction on the number of edge deletions, we can conclude that

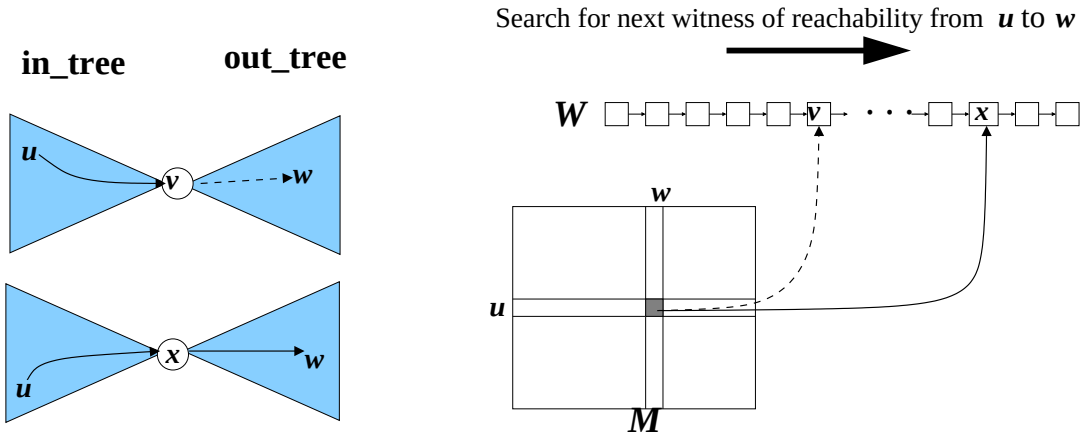

Fig. 1. The search for a new witness of reachability from u to w when $(in_tree, out_tree)(v, d)$ ceases to be the witness.

the invariant I^τ always holds.

Cost Analysis : As mentioned above, the total cost of initializing the witness matrix M is $O(n^2|S|)$. We now assess the total cost over all edge deletions. It takes $O(md)$ cost for maintaining the *in_tree* and the *out_tree* of depth d rooted at a vertex over the entire sequence of edge deletions, thus amounting to $O(md|S|)$ total cost for all tree-pairs from $W(S, d)$. Also for every $(in_tree, out_tree)(v, d) \in W(S, d)$, we incur $O(n)$ cost per vertex when the vertex ceases to belong to either the $in_tree(v, d)$ or the $out_tree(v, d)$, amounting to a total of $O(n^2)$ cost per tree-pair of the set $W(S, d)$. To assess the total cost for searching for a new witness in $W(S, d)$, note that we consider a tree-pair from the list $W(S, d)$ exactly once for being a witness of reachability for a pair (u, w). This is because a tree-pair $(in_tree, out_tree)(v, d)$, that ceases to be a witness of reachability from u to w, can never become a witness of reachability from u to w in future. Hence a total of $O(n^2|S|)$ cost will be spent in searching for a new witness over all vertex-pairs. Summing up the cost of all operations, we conclude that the total cost incurred in maintaining all-pairs reachability over any sequence of edge deletions is $O(md|S| + n^2|S|)$.

Lemma 2 *Given a digraph $G(V, E)$, a set of tree-pairs $W(S, d)$ for a set $S \subset V$ and a positive integer d, the reachability information with respect to the witnesses $W(S, d)$ can be maintained under edge-deletions for all-pairs of vertices separated by distance $\leq d$ in $O((md + n^2)|S|)$ total update time and $O(1)$ query time.*

3.2 Maintaining all-pairs shortest paths with respect to a witness set

The algorithm we describe is analogous to the algorithm described in the previous subsection. Let $W(S, d) = \{in_tree(v, d), out_tree(v, d) | v \in S\}$ be the

set of witnesses kept in a list, also denoted by $W(S,d)$. We maintain a witness matrix M such that for each $u, w \in V$, $M[u,w]$ points to the first tree-pair in the list $W(S,d)$ which is a witness of the shortest path of length $\leq d$ among all paths from u to w passing through any vertex $v \in S$. The matrix M is initialized as follows. The matrix M has all entries pointing to **null** initially. For every $(in_tree, out_tree)(v,d) \in W(S,d)$, we do the following. For each $u \in in_tree(v,d)$ and for every $w \in out_tree(v,d)$, we update $M[u,w]$ to point to v if it is pointing to **null** or if the tree-pair associated with the existing $M[u,w]$ is a witness of a path from u to w with longer length than the tree-pair $(in_tree, out_tree)(v,d)$. It follows easily that it requires $O(n^2)$ time per element of the set $W(S,d)$ to initialize the matrix.

Notice that as the edges are being deleted, a tree-pair $(in_tree, out_tree)(v,d) \in W(S,d)$ may cease to be a witness of the shortest path for a vertex pair (u,w). This happens when either u increases its level in $in_tree(v,d)$ or the vertex w increases its level in $out_tree(v,d)$. Thus we have to update the entries of matrix appropriately after an edge deletion.

To maintain validity of entries of M after an edge deletion, we process each $(in_tree, out_tree)(v,d) \in W(S,d)$ (starting from the first tree-pair in the list $W(S,d)$) as follows. We update $in_tree(v,d)$ and $out_tree(v,d)$ for the edge deletion and compute the sets of vertices X and Y whose level has increased in the $in_tree(v,d)$ and $out_tree(v,d)$ respectively. This step takes $O(|X| + |Y|)$ time in addition to $O(d)$ amortized time for updating the the tree-pair $in_tree, out_tree(v,d)$. It can be seen that for each $u \in X, w \in V$ (likewise $u \in V, w \in Y$), the length of the shortest path passing through v has increased and so we need to search for a new witness if $M[u,w]$ points to v. We process each such pair (u,w) as follows :

Let r be the shortest path length from u to w prior to the recent edge deletion. We scan the list $W(S,d)$ starting from the successor of $(in_tree, out_tree)(v,d)$ in search for a witness of path from u to w with length r. If we find one, we stop; otherwise there is no path from u to w of length $\leq r$ in any tree-pair from the set $W(S,d)$. In this case, we perform another scan starting from the head of the list $W(S,d)$ in search of a witness of path-length $r+1$. We increment r until we find one witness or r increases beyond d. In the latter case, we conclude that there is no path from u to w of length $\leq d$ that passes through any $x \in S$.

In a similar way we process the set Y of vertices whose distance from v has increased.

The following invariant is maintained throughout the series of edge deletions.
I^{APSP} : *At every stage $M[u,w]$ points to the first tree-pair in the list $W(S,d)$ that is a witness of the shortest path from u to w among all paths from u to w captured by tree-pairs of the set $W(S,d)$.*

From the initialization of M, it follows that the invariant holds in the beginning. A simple inductive proof on number of edge deletions can be used to show that the invariant I^{APSP} always holds.

Cost Analysis : The total cost of initializing the matrix M and maintaining *in_tree* and *out_tree* of depth d for every tree-pair from $W(S,d)$, is $O(n^2|S| + md|S|)$. Now we shall assess the total cost for maintaining $M[u,w]$ other than the cost incurred in scanning the list $W(S,d)$ (in search of witnesses). For every $v \in S$, after an edge deletion we compute the set of vertex-pairs whose distance in $(in_tree, out_tree)(v,d)$ has increased. Since we do not process a vertex-pair (u,w) whenever either of them falls beyond distance d from v, therefore, a vertex-pair will be reported in this set at most $2d$ times for a tree-pair from the set $W(S,d)$. It follows that a total of $O(n^2 d)$ time will be spent per tree-pair in computing such sets over the entire sequence of edge deletions. Also note that for a particular distance $r \leq d$ and a vertex-pair (u,w), a tree-pair from the set $W(S,d)$ is considered at most once for being a witness of path-length r (from u to w). This is because a tree-pair that ceases to be a witness of path-length r from u to w can never become a witness of path-length r from u to w in future. Thus for a pair of vertices, the total cost incurred in searching the list $W(S,d)$ will be $O(d|S|)$. Summing up the cost of all operations, we can conclude that the total cost incurred in maintaining all-pairs shortest-paths with respect to the witness set $W(S,d)$ is $O(n^2|S| + md|S| + n^2 d|S|)$, i.e., $O(n^2 d|S|)$.

Lemma 3 *Given a digraph $G(V,E)$, a set of tree-pairs $W(S,d)$ for a set $S \subset V$ and a positive integer d, the shortest-path information with respect to the witnesses $W(S,d)$ can be maintained under edge-deletions for all-pairs of vertices separated by distance $\leq d$ in $O(n^2 d|S|)$ total update time and $O(1)$ query time.*

3.3 Maintaining all-pairs reachability and shortest paths for paths of length $\in [\frac{d}{2}, d]$

It follows from Lemma 2 (and 3) that choosing witness set $W(V,d)$, we can maintain all pairs reachability (shortest-paths) corresponding to paths of length $\leq d$. Notice that the update cost achieved is proportional to the size of witness set. Therefore, to improve the update cost, a relevant question is : *Can we maintain all-pairs reachability (shortest paths) corresponding to paths of length $\leq d$ using $o(n)$ size witness set?*. It can be observed that for maintaining reachability (shortest paths) from u to w, it would suffice if at every stage there is at-least one vertex from the set S such that there is a path (the shortest path) of length $\leq d$ from u to w passing through it. We use random sampling to build S (of sub-linear size) ensuring that with very high probability at least one witness of reachability (shortest path) for each

vertex-pair is there in the set $W(S,d)$.

Consider two vertices u and w separated by a path (shortest path) p_{uw} of length $\in [\frac{d}{2}, d]$. There are $\Omega(d)$ vertices lying on the path p_{uw}, and at least one of them will be present in a random sample of $\frac{2n}{d}$ vertices with probability $\geq 1 - \frac{1}{e}$. The success probability bound can be made arbitrarily close to 1, as mentioned in the following Lemma.

Lemma 4 *[8] Given a path p_{uw} of length l from u to w, if we sample $c\frac{n}{l}\ln n$ vertices (for any $c > 0$), then with probability $1 - \frac{1}{n^c}$, at least one of the vertices will be picked from the path p_{uw}.*

Therefore, for maintaining all-pairs reachability (shortest paths) corresponding to paths of length $\in [\frac{d}{2}, d]$, it would suffice to maintain a witness-matrix with respect to a set $W(S_d, d)$, where S_d is a set formed by picking each vertex of the graph independently with probability $\frac{\ln n}{d}$. Expected size of the set S_d is $\frac{n \ln n}{d}$. This observation can be combined with Lemmas 2 and 3 to yield the following result.

Theorem 5 *Given a graph $G(V,E)$ and $d \leq n$, all-pairs reachability (or shortest paths) corresponding to paths of length $\in [\frac{d}{2}, d]$, can be maintained with optimal query time w.h.p. by maintaining a witness matrix M with respect to the set $W(S_d, d)$. It requires $O(mn \ln n + \frac{n^3 \ln n}{d})$ total update time for maintaining all-pairs reachability, whereas for maintaining all-pairs shortest paths the total update time required is $O(n^3 \ln n)$.*

4 Decremental algorithm for maintaining transitive closure

We maintain all-pairs reachability corresponding to *short* and *long* paths separately. We call the paths of lengths $\leq d$ the *short paths*, and the paths of length $> d$ the *long paths*, where d will be fixed shortly. It follows that for a pair of vertices $u, w \in V$, there is a path from u to w if and only if there is a short path or a long path from u to w. Therefore, for maintaining all-pairs reachability under deletion of edges it suffices to solve the following two subproblems.

Maintaining reachability corresponding to short paths : It follows from the description given in section 2 that *out_tree* of depth d rooted at a vertex u maintains the set of vertices reachable within distance d from u. The set of *out_trees* of depth d from every vertex can thus be used for maintaining all-pairs reachability corresponding to short paths, and the total update time required is $O(mnd)$. In future we shall refer to this data structure by S_0^d.

Maintaining reachability corresponding to long paths : Theorem 5 shows that all-pairs reachability corresponding to paths of length $\in [\frac{r}{2}, r]$ can

be maintained by keeping a witness matrix M with respect to the set $W(S_r, r)$. In order to maintain all-pairs reachability for paths of length $\in [d, n]$, we partition the interval $[d, n]$ into $\log \frac{n}{d}$ sub-intervals : $[d, 2d], \cdots, [2^i d, 2^{i+1} d], \cdots, [n/2, n]$. For each sub-interval $[2^i d, 2^{i+1} d]$ starting from $i = 0$, we build and maintain the set $W(S_{2^{i+1}d}, 2^{i+1} d)$. It can be seen that at any stage in the sequence of edge-deletions, if there is a path of length $\geq d$ from a vertex u to a vertex w, the set $\cup_{i=0}^{\log n/d} W(S_{2^i d}, 2^i d)$ has a witness of reachability from u to w. Therefore, in order to maintain all-pairs reachability corresponding to long paths, it suffices to maintain a witness matrix M with respect to the set $\cup_{i=0}^{\log n/d} W(S_{2^i d}, 2^i d)$. Now, using Theorem 5, the total update time for maintaining the witness matrix M with respect to this set will be

$$\sum_{i=0}^{\log \frac{n}{d}} \left(mn \ln n + \frac{n^3 \ln n}{2^i d} \right) = O\left(mn \log^2 n + \frac{n^3 \ln n}{d} \right)$$

We shall denote the above data structure by L_d^n. We maintain all pairs reachability by keeping the data structures S_0^d and L_d^n simultaneously. Deleting an edge will require updating both the data structures, and thus the total update time for maintaining all-pairs reachability (transitive closure) over any sequence of edge-deletions will be

$$O\left(mnd + mn \log^2 n + \frac{n^3 \ln n}{d} \right) = O\left(mn \log^2 n + n^2 \sqrt{m \ln n} \right) \quad \text{for} \quad d = \frac{n\sqrt{\ln n}}{\sqrt{m}}$$

Theorem 6 *Given a digraph $G(V, E)$, all-pairs reachability information can be maintained under deletion of edges by an algorithm that achieves $O(1)$ query time w.h.p. and $O(n \log^2 n + \frac{n^2 \sqrt{\ln n}}{\sqrt{m}})$ amortized update time per edge deletion.*

There are two previous algorithms for maintaining transitive closure under deletion of edges with $O(1)$ query time. The first algorithm due to Poutre and Leeuwen [14] achieves $O(m)$ amortized update time while the second algorithm due to Demetrescu and Italiano [4] achieves $O(\frac{n^3}{m})$ amortized update time. These two algorithms together establish an upper bound of $O(n^{\frac{3}{2}})$ on the update time. By suitably combining our algorithm with [14], we can state the following Corollary.

Corollary 7 *There exists a decremental algorithm for maintaining transitive closure that achieves $O(n^{\frac{4}{3}} \sqrt[3]{\log n})$ amortized update time and optimal query time w.h.p.*

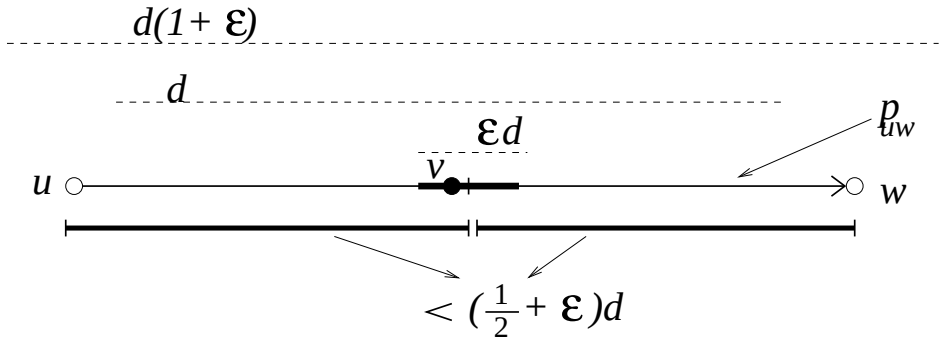

Fig. 2.

5 Decremental algorithm for maintaining all-pairs approximate shortest paths

We introduced the terminologies of witness of reachability and witness of shortest path. In this section, we introduce one more type of witness, namely a witness of *approximate shortest path.*

Consider a pair of vertices $u, w \in V$ such that the vertex u lies in *in_tree* and the vertex w lies in *out_tree* of depth d rooted at a vertex v. The pair (*in_tree*,*in_tree*)(v, d) not only act as a witness of reachability between the two vertices, it also gives an upper bound (thus estimates) of $2d$ on the distance between the two vertices. In other words, if $\delta(u, w)$ is the distance from vertex u to vertex w, the tree-pair (*in_tree*,*in_tree*)(v, d) acts as a witness of $\frac{2d}{\delta(u,w)}$-approximate distance between the two vertices.

We present an efficient decremental algorithm for maintaining all-pairs $(1+\epsilon)$-approximate shortest paths for any $\epsilon > 0$. The algorithm can be visualized as a careful refinement of the algorithm for maintaining all-pairs reachability described in the previous section.

The following Lemma is the basis for maintaining witnesses of $(1 + 2\epsilon)$-approximate distances for all-pairs separated by distance $\in [d, d(1 + \epsilon)]$.

Lemma 8 *Given a sample $S_{\epsilon d} \subset V$ formed by picking each vertex independently with probability $\frac{\ln n}{\epsilon d}$, the set $W(S_{\epsilon d}, (1/2+\epsilon)d)$ has a witness of $(1+2\epsilon)$-approximate distance with probability $1 - 1/n$, for a pair of vertices separated by distance $\in [d, d(1 + \epsilon)]$.*

PROOF. Let p_{uw} be a shortest-path from vertex u to vertex w of length $\in [d, d(1 + \epsilon)]$ (see Figure 2). Let s_{uw} be the set of the vertices lying on the path p_{uw} that are at distance $\leq \epsilon d/2$ from the mid-point of p_{uw}. Clearly $|s_{uw}| = \epsilon d$. It can be observed that an *in_tree* and an *out_tree* of depth $(1/2+\epsilon)d$ rooted at any $v \in s_{uw}$ capture the path p_{uw} completely, and thus establish an upper bound of $(1 + 2\epsilon)d$ on the distance from u to w. Since the actual dis-

tance from u to w lies in the interval $\in [d,(1+\epsilon)d]$, therefore, the tree-pair $(in_tree, out_tree)(v,d) : v \in s_{uw}$ is a witness of $(1+2\epsilon)$-approximate distance from u to w.

A sample $S_{\epsilon d} \subset V$ formed by picking each vertex independently with probability $\frac{\ln n}{\epsilon d}$ has at-least one vertex from the set s_{uw} with probability $\geq 1-1/n$. So the set $W(S_{\epsilon d}, (\frac{1}{2}+\epsilon)d)$ has a witness of $(1+2\epsilon)$-approximate distance for a pair of vertices separated by distance $\in [d, d(1+\epsilon)]$ with probability $> 1-\frac{1}{n}$.

It follows from the Lemma given above that in order to maintain $(1+2\epsilon)$-approximate distances for all-pairs of vertices separated by distance $\in [d,(1+\epsilon)d]$, it suffices if we maintain witness matrix for all-pairs reachability with respect to the witness set $W(S_{\epsilon d}, (\frac{1}{2}+\epsilon)d)$. So we can state the following theorem along the lines of Theorem 5.

Theorem 9 *Given a graph $G(V,E)$ and $d \leq n$, all-pairs $(1+2\epsilon)$-approximate shortest paths for all-pairs of vertices separated by distance $\in [d,(1+\epsilon)d]$, can be maintained under edge-deletions with optimal query time w.h.p. by maintaining a witness matrix M for all-pairs reachability with respect to the set $W(S_{\epsilon d}, \frac{d}{2})$. The total update time required over any sequence of edge-deletions is $O(\frac{mn\ln n}{\epsilon} + \frac{n^3 \ln n}{\epsilon d})$.*

We use the idea of maintaining $(1+\epsilon)$-approximate shortest paths corresponding to *short* and *long* paths separately (also used for maintaining transitive closure in the previous section). We call the paths of lengths $\leq d$ the *short paths*, and the paths of length $> d$ the *long paths*, where d will be determined in the analysis.

Maintaining $(1+\epsilon)$-approximate shortest paths for vertex-pairs separated by long paths:

In order to maintain $(1+\epsilon)$-approximate shortest paths for pair of vertices separated by distance $\in [d,n]$, we partition the interval $[d,n]$ into $\ln_{1+\epsilon}\frac{n}{d}$ sub-intervals : $\{((1+\epsilon)^i d, (1+\epsilon)^{i+1}d) | i \leq \ln_{1+\epsilon}\frac{n}{d}\}$. For each sub-interval $[(1+\epsilon)^i d, (1+\epsilon)^{i+1}d]$ starting from $i=0$, we build and maintain the sets $\mathbf{W}_i = W(S_{\epsilon(1+\epsilon)^i d}, (1+\epsilon)^i(\frac{1}{2}+\epsilon)d)$ and update the matrix M accordingly.

Processing of an edge deletion involves updating each tree-pair, and searching for new witness for each $u,w \in V$ if $M[u,w]$ ceases to be the witness due to the recent edge deletion. Whenever search for witness of reachability from u to w fails in list $\mathbf{W}_i$, we start search in next list $\mathbf{W}_{i+1}$. Thus the following invariant (along the lines of I^τ) will be maintained for each $u,w \in V$.

I^R : At any time if sets $\mathbf{W}_{i_1}, \mathbf{W}_{i_2}, \cdots, \mathbf{W}_{i_k} : i_1 < i_2 \cdots < i_k$, have witnesses of reachability from u to w, then $M[u,w]$ will point to the first witness of reachability (from u to w) in the list $\mathbf{W}_{i_1}$

Let the length of the shortest-path from vertex u to vertex w lies in interval $((1+\epsilon)^j d, (1+\epsilon)^{j+1} d)$. The Lemma 8 ensures that the list $\mathbf{W}_j$ must have a tree-pair which is a witness of $(1+2\epsilon)$-approximate shortest path from u to w, and the invariant I^R implies that $M[u,w]$ will point to this tree-pair. In other words, the reachability-witness matrix M with respect to the set $\cup_{i=0}^{\ln_{1+\epsilon} n/d} \mathbf{W}_i$ maintains witnesses of $(1+2\epsilon)$-approximate shortest path for each pair (u,w) separated by distance $\in [d,n]$. Using Theorem 9, the total update time required for maintaining the witness matrix M over any sequence of edge-deletions will be

$$\sum_{i=0}^{\ln_{1+\epsilon} \frac{n}{d}} \left(\frac{mn \ln n}{\epsilon} + \frac{n^3 \ln n}{2^i \epsilon d} \right) = O\left(\frac{mn \ln n}{\epsilon^2} + \frac{n^3 \ln n}{\epsilon^2 d} \right)$$

So the amortized update time per edge deletion required for maintaining $(1+\epsilon)$-approximate shortest paths for all-pairs of vertices separated by long paths is $O(\frac{n \ln n}{\epsilon^2} + \frac{n^3 \ln n}{\epsilon^2 dm})$.

Maintaining shortest paths for vertex-pairs separated by short paths : We maintain *out_tree* of depth d from each vertex to report exact distance between any pair of vertices separated by short paths. The amortized update time per edge deletion is $O(nd)$

Balancing the update bounds for maintaining $(1+\epsilon)$-approximate shortest path for long and short paths, we get $d = \frac{n\sqrt{\ln n}}{\epsilon\sqrt{m}}$. Hence the amortized update time per edge deletion for maintaining $(1+\epsilon)$-approximate shortest paths will be $O(\frac{n \ln n}{\epsilon^2} + \frac{n^2\sqrt{\ln n}}{\epsilon\sqrt{m}})$.

Theorem 10 *Given an unweighted digraph $G(V,E)$, there exists a data structure for maintaining all-pairs $(1+\epsilon)$ approximate shortest paths/distance with optimal query time w.h.p. and $O(\frac{n \ln n}{\epsilon^2} + \frac{n^2\sqrt{\ln n}}{\epsilon\sqrt{m}})$ amortized update time per edge-deletion.*

King [12] gave a fully dynamic algorithm for maintaining all-pairs $(1+\epsilon)$ approximate shortest paths that require $O(\frac{n^2 \log^3 n}{\epsilon^2})$ update time per edge deletion. For updates consisting of edge deletions only, our algorithm improves the update time by a factor of $O(\frac{\sqrt{m \log^5 n}}{\epsilon})$.

6 Decremental algorithm for maintaining all-pairs distances and all-pairs shortest paths

We maintain all-pairs shortest paths/distances by keeping a witness of shortest path (if exists) for each vertex pair. It follows from Theorem 5 that for main-

taining shortest paths for all-pairs of vertices separated by distance $\in [r/2, r]$, it suffices to maintain a matrix of shortest-path witnesses with respect to the set $W(S_r, r)$. In order to maintain shortest paths for all-pairs of vertices (irrespective of the distance), we build and maintain $W(S_r, r)$ for each $r \in \{1, 2, 4, \cdots, 2^i, \cdots, n\}$. It follows from subsection 3.2 that with respect to the set $\cup_{i=0}^{\log n} W(S_{2^i}, 2^i)$, an entry $M[u, w]$ gets initialized to a witness of the shortest path from u to w for all $u, w \in V$ in total time $O(n^3)$.

Processing of an edge deletion involves updating each tree-pair and searching for new shortest-path witness for each pair $u, w \in V$ if path length from u to w passing through $M[u, w]$ has increased due to the recent edge deletion. While searching for a witness of path length r from u to w, we will confine our search within the list $W(S_{2^i}, 2^i)$ for $2^{i-1} < r \leq 2^i$. Whenever search for witness of path of length 2^i fails in the list $W(S_{2^i}, 2^i)$, we move onto the next list $W(S_{2^{i+1}}, 2^{i+1})$. We proceed in this way for each pair, performing $O(n \log n)$ work per witness list (scanning list of length $\frac{n}{2^i} \ln n$, for 2^i times). Since there are now $\log n$ lists and a total of n^2 pairs, the total update cost will be $O(n^3 \log^2 n)$ over any sequence of edge deletions. We can thus state the following theorem.

Theorem 11 *Given an unweighted digraph $G(V, E)$, there exists a data structure for maintaining all-pairs shortest distances in $O(\frac{n^3 \log^2 n}{m})$ amortized update time per edge-deletion and taking $O(1)$ time to answer a distance query w.h.p.*

Remark : Demetrescu and Italiano [5] designed an $O(\frac{n^3}{m} \log^3 n)$ update time algorithm for maintaining all-pairs shortest paths under deletion of edges. Their algorithms achieves $O(1)$ query time w.h.p. and improves the previous $O(n^2)$ update time of King [12] for dense graphs. Our algorithm improves the update time further by a factor of $O(\log n)$.

For initial complete graph, the new algorithm achieves $O(n \log^2 n)$ amortized update time. But for sparse graphs the update time may be $O(n^2)$. So there is still no decremental algorithm for maintaining all-pairs shortest distances that achieves sub-quadratic update time for all digraphs and takes $O(1)$ time to answer distance query. However, for the case of shortest path reporting problem, we are able to design a decremental algorithm that requires $O(n^{\frac{3}{2}} \sqrt{\log n})$ update time for any digraph and answers any query in optimal time. The improvement is achieved by combining the concept of filtering search [1] with the data structures S_0^d, L_d^n in a novel way as follows.

If the shortest path is of length $l_{uw} \leq d$, we are able to answer the distance query exactly in $O(1)$ time using *out_tree* rooted at u. Otherwise $2^{i-1} d \leq l_{uw} < 2^i d$ holds for some $2 \leq i < \log \frac{n}{d}$. Note that w.h.p. the set $W(S_{2^i d}, 2^i d)$ has a witness of the shortest path from u to w. But finding the witness may take

$O(|S_{2^i d}|) = O(\frac{n}{2^i d} \log n)$ time since it takes $O(1)$ time to find the length of the path from u to w lying in a tree-pair $\in W(S_{2^i d}, 2^i d)$. This forced us to maintain a witness-matrix for achieving $O(1)$ query time for distance reporting problem. However, for a shortest-path reporting query, we must in any case spend $O(l_{uw})$ time to report all the edges on the path. Therefore, if we spend additional $O(l_{uw})$ time to find shortest path witness, we will still be achieving $O(l_{uw})$ query time for path reporting (which is optimal). This is indeed the idea of filtering search [1]. Now choosing $d = \sqrt{n \log n}$ ensures that $|W_{2^i d}| \leq d < l_{uw}$ for all i. Our algorithm employs the data structures S_0^d and L_d^n for $d = \sqrt{n \log n}$ and processes a shortest-path reporting query (say, from u to w) in optimal time as follows :

First inquire if w is present in the out_tree *rooted at u. This operation takes constant time. If the answer is yes, we can report the shortest path using the* out_tree *rooted at u; otherwise inquire if w is reachable from u or not, using the witness matrix M. If $M[u, w]$ is pointing to* **null**, *we report in $O(1)$ time that there is no path from u to w. Otherwise, it can be concluded that the shortest path length $l_{uw} \in [\sqrt{n \log n} + 1, n - 1]$. In this case, at least one of the vertex of the shortest path from u to w must be present w.h.p. in the sampled set of witnesses. We search the entire set of witnesses to find the witness of shortest path. The size of the witness set being $\sqrt{n \log n}$, we can thus find the witness of the shortest path from u to w in $O(\sqrt{n \log n})$ time and then report the shortest path from u to w passing through it in additional $O(l_{uw})$ time. Since $l_{uw} > \sqrt{n \log n}$, the total time taken for processing a shortest path reporting query is $O(l_{uw})$ (and hence optimal).*

The amortized update time per edge deletion for the two data structures S_0^d and L_d^n will be $O(nd + \frac{n^3}{md} \log n)$, which is $O(n^{\frac{3}{2}} \sqrt{\log n})$ for $d = \sqrt{n \log n}$.

Theorem 12 *An unweighted digraph $G(V, E)$ can be preprocessed to build a data structure that answers any on-line shortest path reporting query in optimal time w.h.p. while ensuring $O(\min(n^{\frac{3}{2}} \sqrt{\log n}, \frac{n^3 \log^2 n}{m}))$ amortized update time per edge-deletion.*

7 Space Requirement of our Algorithms

It can be seen that all the algorithms given in this paper employ a witness matrix M and $O(n)$ number of *in_trees* and *out_trees*. The matrix M clearly occupies $\theta(n^2)$ space. Although the earlier scheme given in [12] requires $\theta(m)$ space for maintaining an *in_tree* (or an *out_tree*), it has been improved to $O(n)$ by King and Thorup [13]. Thus the total space requirement for maintaining $O(n)$ number of *in_trees* or *out_trees* is $O(n^2)$. Hence the space requirement of all the algorithms given in this paper is $O(n^2)$.

8 Conclusion and Open Problems

In this paper, we gave improved decremental algorithms for maintaining transitive closure and All-pairs shortest paths with optimal query time. We pose the following problem :

Can a given graph $G(V, E)$ be preprocessed so that the set of vertices lying within distance d from a vertex can be maintained under deletion of edges while taking $o(md)$ total update time?

An answer to the above problem will lead to improved upper bound on the update time for the decremental problem of transitive closure (and all-pairs approximate shortest paths).

9 Acknowledgments

We sincerely thank Camil Demetrescu for providing a very helpful clarification about the paper [4]. We also thank L. Sunil Chandran and L. Shankar Ram for their valuable suggestions and comments on an earlier draft of this paper.

References

[1] B. Chazelle. Filtering search : A new approach to query-answering. *SIAM J. Comput.*, 15:703–724, 1986.

[2] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9:251–280, 1990.

[3] C. Demetrescu. Personal communication.

[4] C. Demetrescu and G. Italiano. Fully dynamic transitive closure : Breaking through the $O(n^2)$ barrier. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 381–389, 2000.

[5] C. Demetrescu and G. Italiano. Fully dynamic all pairs shortest paths with real edge weights. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 381–389, 2001.

[6] S. Even and Y. Shiloach. An on-line edge-deletion problem. *Journal of association for computing machinery*, 28:1–4, 1981.

[7] R. Goldman, N. Shivakumar, S. Venkatasubramanian, and H. Garcia-Molnias. Proximity searches in databases. In *Proceedings of the 24th VLDB Conference*, 1998.

[8] D. Greene and D. Knuth. Mathematics for the analysis of algorithms. *Birkhauser, Boston*, 1982.

[9] M. R. Henzinger and V. King. Fully dynamic bi-connectivity and transitive closure. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 664–672, 1995.

[10] N. Horspool. Incremental generation of lr parsers. In *Technical Repor, University of Victoria, Computer Science Department*, 1988.

[11] H. V. Jagdish. A compressed transitive closure technique for efficient fixed-point query processing. In *Proceedings of the Second International Conference on Expert Database Systems*, 1988.

[12] V. King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 325–334, 1999.

[13] V. King and M. Thorup. A space saving trick for directed dynamic transitive closure and shortest path algorithms. In *Proceedings of 7th Annual International Conference, COCOON*, volume 2108 of *LNCS*, pages 268–277. Springer Verlag, 2001.

[14] H. Poutre and J. van Leeuwen. Maintenance of transitive closure and transitive reduction of a graph. In *Proceedings of Workshop on Graph-Theoretic Concepts in Computer Science*, volume 314 of *LNCS*, pages 106–120. Springer Verlag, 1988.

[15] U. Zwick. All-pairs shortest paths in weighted directed graphs - exact and almost exact algorithms. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 310–319, 1998.