# Average update times for fully-dynamic all-pairs shortest paths ${ }^{\star}$ 

Tobias Friedrich*, Nils Hebbinghaus<br>Max-Planck-Institut für Informatik, Campus E1.4, 66123 Saarbrücken, Germany

## ARTICLE INFO

## Article history:

Received 26 August 2009
Received in revised form 5 February 2011
Accepted 8 February 2011
Available online 31 March 2011

## Keywords:

Dynamic graph algorithms
Shortest paths
Average-case analysis
Random graphs


#### Abstract

We study the fully-dynamic all pairs shortest path problem for graphs with arbitrary nonnegative edge weights. It is known for digraphs that an update of the distance matrix costs $\mathcal{O}\left(n^{2.75}\right.$ polylog(n)) worst-case time (Thorup, 2005 [20]) and $\mathcal{O}\left(n^{2} \log ^{3}(n)\right)$ amortized time (Demetrescu and Italiano, 2004 [4]) where $n$ is the number of vertices. We present the first average-case analysis of the undirected problem. For a random update we show that the expected time per update is bounded by $\mathcal{O}\left(n^{4 / 3+\varepsilon}\right)$ for all $\varepsilon>0$. If the graph is outside the critical window, we prove even smaller bounds.


© 2011 Elsevier B.V. All rights reserved.

## 1. Introduction

Dynamic graph algorithms maintain a certain property (e.g., connectivity information) of a graph that changes (a new edge inserted or an existing edge deleted) dynamically over time. They are used in a variety of contexts, e.g., operating systems, information systems, database systems, network management, assembly planning, VLSI design and graphical applications. An algorithm is called fully-dynamic if both edge weight increases and edge weight decreases are allowed. While a number of fully dynamic algorithms have been obtained for various properties on undirected graphs (see [6]), the design and analysis of fully-dynamic algorithms for directed graphs has turned out to be much harder (e.g., [16,17,15,9]).

In this article, we consider the fully-dynamic all-pairs shortest path problem (APSP) for undirected graphs, which is one of the most fundamental problems in dynamic graph algorithms. The problem has been studied intensively since the late sixties (see [4] and references therein). We are interested in algorithms that maintain a complete distance matrix as edges are inserted or deleted. The static directed APSP problem can be solved in $\mathcal{O}\left(m n+n^{2} \log \log n\right)$ time [14] where $n$ is the number of vertices and $m$ is the number of edges. For dense graphs the problem can be solved in time $\mathcal{O}\left(n^{3} \sqrt{\log \log (n) / \log n}\right)$ [19] in general and in time $\mathcal{O}\left(n^{2.575}\right)$ [22] for integer edge weights. These bounds also give an upper bound on the worst-case update time for a static recomputation from scratch. The first major improvement that is provably faster than this only worked on digraphs with small integer weights. King [10] presented a fully-dynamic APSP algorithm for general directed graphs with positive integer weights less than $C$ that supported updates in $\mathcal{O}\left(n^{2.5} \sqrt{C \log n}\right)$. In the remainder of the paper, we will only consider non-negative real-valued edge weights. Demetrescu and Italiano pursued this problem in a series of papers and showed that it can be solved in $\mathcal{O}\left(n^{2} \log ^{3} n\right)$ amortized time per update [4]. This has been slightly improved to $\mathcal{O}\left(n^{2}\left(\log n+\log ^{2}((m+n) / n)\right)\right)$ amortized time per update by Thorup [21]. In [20], Thorup showed a worst-case update time of $\mathcal{O}\left(n^{2.75}\right.$ polylog $\left.(n)\right)$.

We are interested in expected update times. It is known that the static APSP problem can be solved in average time $\mathcal{O}\left(n^{2} \log n\right)$ for a nonnegatively weighted directed graph [12]. The only known result for the online case is for the undirected, unweighted, decremental, randomized, and approximate version of the APSP problem. Roditty and Zwick [18] showed for

[^0]this setting an expected amortized time of $\mathcal{O}(n \operatorname{polylog}(n))$. For our setting of the problem on undirected graphs with arbitrary non-negative edge weights, there is nothing known about the average-case update times. We analyze a variant of Demetrescu and Italiano's algorithm described in Section 3. Let $\mathcal{R}(p)$ denote the expected runtime of our algorithm for a single random edge update of a random graph $G \in \mathcal{G}(n, p)$. Let $\varepsilon, \varepsilon^{\prime}>0$. For arbitrary $p$, we can show $\mathcal{R}(p)=\mathcal{O}\left(n^{4 / 3+\varepsilon}\right)$. However, for most $p$ we can prove that the runtime is actually much smaller. The above bound is best only at the phase transition around $p n=1$, i.e., when the size of the largest component rapidly grows from $\Theta(\log n)$ to $\Theta(n)$. When the graph is sparser, our algorithm is much faster. In this case, we can show $\mathcal{R}(p)=\mathcal{O}\left(n^{2 / 3+\varepsilon}\right)$ for $p n \leq 1-n^{-1 / 3}$ and $\mathcal{R}(p)=\mathcal{O}\left(n^{\varepsilon}\right)$ for $p n<1 / 2$. Similarly, the algorithm becomes faster when a giant component emerges. We show $\mathcal{R}(p)=\mathcal{O}\left(n^{\varepsilon} / p\right)$ for $p n \geq 1+\varepsilon^{\prime}$. The final result is given in Theorem 11. For inserting random edges into the empty graph until it becomes complete this yields an expected amortized update time of $\mathcal{O}\left(n^{1 / 3+\varepsilon}\right)$.

Additionally to these asymptotic upper bounds on the expected runtime, we also examined the empirical average runtime. Interestingly, this also shows that the update costs are first increasing and later decreasing when more edges are inserted. This corresponds well with the above phase distinction for the asymptotic bounds.

The remainder of this paper is organized as follows. The next section presents all necessary graph theoretical notations. In Section 3 we present our algorithm. In Section 4 we prove a number of random graph properties which are then used in Section 5 to show the asymptotic bounds. The last section presents some empirical results.

## 2. Preliminaries

Demetrescu and Italiano [5] performed several experiments on directed random graphs. We want to bound the expected runtime of random updates on a random graph of a very similar algorithm. We utilize the random graph model $\mathcal{g}(n, p)$ introduced and popularized by Erdős and Rényi [7]. The $\mathcal{q}(n, p)$ model consists of a graph with $n$ vertices in which each edge is chosen independently with probability $p$. We assume that the edge weights are chosen uniformly at random from $[0,1]$. In our model, a random update first chooses two vertices $x$ and $y(x \neq y)$. Then, with (fixed) probability $\delta$, it inserts the edge $(x, y)$ with a random weight $w \in[0,1]$. If the edge was already in the graph, it changes its weight to $w$. Otherwise, with probability $1-\delta$, a random update deletes the edge $(x, y)$ if $(x, y)$ is in the graph (otherwise, it does nothing).

A random update of a graph where each edge is present with the same probability yields a graph where also each edge is present with the same probability. Let $\rho(T)$ be this probability after $T$ random updates. Hence,

$$
\begin{aligned}
& \rho(0)=p \\
& \rho(T+1)= \delta\left[(1-\rho(T))\left(\rho(T)+1 /\binom{n}{2}\right)+(\rho(T))^{2}\right] \\
&+(1-\delta)\left[(1-\rho(T)) \rho(T)+\rho(T)\left(\rho(T)-1 /\binom{n}{2}\right)\right] \\
&= \rho(T)\left(1-1 /\binom{n}{2}\right)+\delta /\binom{n}{2} \\
&= \delta+(\rho(0)-\delta)\left(1-1 /\binom{n}{2}\right)^{T+1} .
\end{aligned}
$$

This implies that $T$ random updates on a random graph $G \in g(n, p)$ lead to a graph with edges present with probability $p^{\prime}=\delta+(p-\delta)\left(1-1 /\binom{n}{2}\right)^{T}$, but not necessarily mutually independent.

Throughout the paper, we use the following notations:

- $G=(V, E)$ is an undirected graph with arbitrary non-negative edge weights.
- $\Delta:=\max _{v \in V} \operatorname{deg}(v)$ is the maximum degree of the graph $G$.
- length of a path is the sum of its edge weights
- $\operatorname{dist}(x, y)$ (distance) is the length of a shortest path from $x$ to $y$.
- diam $(G)$ (diameter) is the greatest (hop-)distance between any two vertices of one component.
- $C(x)$ denotes the component that contains the vertex $x$.
- $C(G)$ denotes the largest component of $G$.
- $w_{x y}$ denotes the weight of an edge $(x, y)$.
- $\pi_{x y}=\left\langle u_{0}, u_{1}, \ldots, u_{k}\right\rangle$ is a path from vertex $x=u_{0}$ to vertex $y=u_{k}$, i.e., a sequence of vertices such that with $\left(u_{i}, u_{i+1}\right) \in E$ for each $0 \leq i<k$ (no repeated edges).
- $w\left(\pi_{x y}\right)=\sum_{i=0}^{k-1} w_{u_{i} u_{i+1}}$ is the weight of a path.
- $\pi_{x y} \circ \pi_{y z}$ denotes the concatenation of two paths $\pi_{x y}$ and $\pi_{y z}$.
- $\ell\left(\pi_{x y}\right)$ denotes the subpath $\pi_{x a}$ of $\pi_{x y}$ such that $\pi_{x y}=\pi_{x a} \circ\langle a, y\rangle$.
- $r\left(\pi_{x y}\right)$ denotes the subpath $\pi_{b y}$ of $\pi_{x y}$ such that $\pi_{x y}=\langle x, b\rangle \circ \pi_{b y}$.

We assume without loss of generality that there is only one shortest path between each pair of vertices in G. Otherwise, ties can be broken as discussed in Section 3.4 of [4].

## 3. Algorithm

We will now describe our algorithm. It is a slight modification of the algorithm of Demetrescu and Italiano [4] as our aim is an average-case analysis of the undirected problem while they were interested in the amortized costs for directed graphs.

The main tool Demetrescu and Italiano [4] very cleverly introduced and applied is the concept of "locally shortest paths". A path $\pi_{x y}$ is locally shortest if every proper subpath is a shortest path or it consists of only a single vertex. The algorithm maintains the following data structures:

- $w_{x y}$ weight of edge $(x, y)$
- $P_{x y}$ priority queue of the locally shortest paths from $x$ to $y$ (priority $w\left(\pi_{x y}\right)$ )
- $P_{x y}^{*}$ set containing the shortest path from $x$ to $y$, if $y$ is reachable from $x$
- $L\left(\pi_{x y}\right)$ set of left-extensions $\left\langle x^{\prime}, x\right\rangle \circ \pi_{x y}$ of $\pi_{x y}$ that are locally shortest paths
- $L^{*}\left(\pi_{x y}\right)$ set of left-extensions $\left\langle x^{\prime}, x\right\rangle \circ \pi_{x y}$ of $\pi_{x y}$ that are shortest paths
- $R\left(\pi_{x y}\right)$ set of right-extensions $\pi_{x y} \circ\left\langle y, y^{\prime}\right\rangle$ of $\pi_{x y}$ that are locally shortest paths
- $R^{*}\left(\pi_{x y}\right)$ set of right-extensions $\pi_{x y} \circ\left\langle y, y^{\prime}\right\rangle$ of $\pi_{x y}$ that are shortest paths

Note that $P_{x y}^{*} \subseteq P_{x y}$ and that every minimum weight path in $P_{x y}$ is also a shortest path. Each path $\pi_{x y} \in P_{x y}$ is stored implicitly with constant space by just storing two pointers to the subpaths $\ell\left(\pi_{x y}\right)$ and $r\left(\pi_{x y}\right)$.

The pseudo-code of our algorithm is given in Fig. 1. The first four phases are equivalent to [5]. We will just describe them briefly. A detailed description can be found in [4]. In the first phase, the algorithm deletes from the data structure all the paths that would stop being locally shortest if we deleted the edge $(u, v)$. In doing so it stores the pairs of the endpoints of the affected paths in the temporary list $A$. In the following phase it adds the edge (if it is an insert or update operation) to the data structures. The third phase initializes the heap $H$ with the minimum weight paths $\pi_{x y}$ for all $(x, y) \in A$. In the fourth phase the algorithm repeatedly extracts the cheapest path $\pi_{x y}$ from $H$. The first extracted path for each pair ( $x, y$ ) must be a shortest path. If this is the case, the path is stored in the data structures. To propagate this information, also its left- and right-extensions are updated and added to $H$ to find all further extensions.

The amortized number of new locally shortest paths can be $\Omega\left(n^{3}\right)$ per update. To allow a better amortized performance, Demetrescu and Italiano [4] had to delay the update of the data structure in a very clever way. Their data structure can contain paths in $P_{x y}$ which are not locally shortest any more. We avoid this with the fifth phase. There, all locally shortest paths which stopped being locally shortest because one of their two subpaths stopped being shortest path are detected and deleted.

We analyze the expected time for the algorithm to insert a randomly chosen edge $e$ in the $\operatorname{graph} G \in \mathcal{G}(n, p)$ and maintain the sets of shortest path and locally shortest path. The weights of $e$ and of the edges in $G$ are chosen uniformly at random from the set $[0,1]$.

## 4. Random graph properties

To bound the runtime of our algorithm in the next section, we first provide some properties of random graphs $G \in \mathcal{G}(n, p)$. The main result of this section will be Theorem 9. It bounds the quantity $\mu(p)$ which we define as the expected number of locally shortest paths and shortest paths passing a fixed edge of G. Let LSP denote the set of all locally shortest paths and SP the set of all shortest paths in $G$. We will use the following four lemmas.

Lemma 1 (Bollobás [2]). Let $G \in \mathcal{G}(n, p)$ with $p n<1 / 2$. Then, $\operatorname{Pr}[|C(G)| \leq 20 \log n]=1-\mathcal{O}\left(n^{-2}\right)$.
Lemma 2 (Bollobás [2]). For every $\alpha>0$ and $G \in \mathcal{G}(n, p)$ with $p n=\alpha \log n, \operatorname{Pr}[G$ is connected $]=1-\mathcal{O}\left(n^{1-2 \alpha}\right)$.
Lemma 3 (Chung and Lu[3]). For every $\varepsilon>0$ and $G \in \mathcal{g}(n, p)$ with $p n=1+\varepsilon, \operatorname{Pr}[\operatorname{diam}(G) \leq 2 \log n]=1-o\left(n^{-1}\right)$.
Lemma 4 (Nachmias and Peres [13]). Let $x \in G$ and $G \in \mathcal{G}(n, p)$ with $p n \leq 1+n^{-1 / 3}$. Then, $\boldsymbol{P r}\left[|C(x)|>2 n^{2 / 3}\right]=\mathcal{O}\left(n^{-1 / 3}\right)$.
The following lemma gives a general upper bound on the expected diameter of a random graph $G \in \mathscr{G}(n, p)$ for arbitrary $p$. Recall that we defined the diameter of a disconnected graph as the maximum diameter of its components.

Lemma 5. Let $G \in \mathscr{G}(n, p)$. Then, $\boldsymbol{E}[\operatorname{diam}(G)]=\mathcal{O}\left(n^{1 / 3}\right)$.
Proof. Let $G$ be a complete graph on $n$ vertices with edge weights uniformly distributed at random in [ 0,1 ]. Let $G_{\leq p}=$ ( $V, E_{\leq p}$ ) be the subgraph of $G$ containing all vertices but only those edges with weight less or equal $p$. Then $G_{\leq p}$ is a $g(n, p)$ graph. We apply Kruskal's algorithm [11] for the construction of a minimum spanning forest of $G$, i.e., we look at the edges in increasing weight order and integrate every edge that does not introduce a cycle in the current edge set. We stop this process if the current edge has weight greater than $p$ and denote the obtained subset of edges $E_{\text {Kruskal }, \leq p}$. Let us also denote the edge set of the spanning forest which is returned from the completed Kruskal algorithm by $E_{\text {Kruskal }}$. By Addario-Berry et al. [1], the expected diameter of $G_{\text {Kruskal }}:=\left(V, E_{\text {Kruskal }}\right)$ is of order $\Theta\left(n^{1 / 3}\right)$. Clearly, $E_{\text {Kruskal }, \leq p} \subseteq E_{\text {Kruskal }} \cap E_{\leq p}$. As $G_{\text {Kruskal }, \leq p}:=\left(V, E_{\text {Kruskal }, \leq p}\right)$ is a minimum spanning forest of $G_{\leq p}$, we get

$$
\boldsymbol{E}\left[\operatorname{diam}\left(G_{\leq p}\right)\right] \leq \boldsymbol{E}\left[\operatorname{diam}\left(G_{\text {Kruskal }, \leq p}\right)\right] \leq \boldsymbol{E}\left[\operatorname{diam}\left(G_{\text {Kruskal }}\right)\right]=\mathcal{O}\left(n^{1 / 3}\right)
$$

```
\(\operatorname{Update}(u, v, w)\)
\(\triangleright\) Phase 1: Delete edge \((u, v)\) and all paths containing edge \((u, v)\),
                store pairs of vertices affected by the update in list \(A\)
if \((u, v) \in P_{u v}\) then
    \(Q \leftarrow\{(u, v)\}\)
while \(Q \neq \emptyset\) do
    extract any \(\pi_{x y}\) from \(Q\)
    remove \(\pi_{x y}\) from \(P_{x y}, L\left(r\left(\pi_{x y}\right)\right)\), and \(R\left(\ell\left(\pi_{x y}\right)\right)\)
    if \(\pi_{x y} \in P_{x y}^{*}\) then
        remove \(\pi_{x y}\) from \(P_{x y}^{*}, L^{*}\left(r\left(\pi_{x y}\right)\right)\), and \(R^{*}\left(\ell\left(\pi_{x y}\right)\right)\)
        add \((x, y)\) to \(A\)
    add all paths in \(L\left(\pi_{x y}\right)\) to \(Q\)
    add all paths in \(R\left(\pi_{x y}\right)\) to \(Q\)
\(\triangleright\) Phase 2: Insert edge \((u, v)\) with weight \(w\)
if \(w<\infty\) then
    \(w_{u v}=w\)
    add \((u, v)\) to \(A\)
    add \((u, v)\) to \(P_{u v}, L\left(\pi_{v v}\right)\), and \(R\left(\pi_{u u}\right)\)
\(\triangleright\) Phase 3: Scan pairs in \(A\)
while \(A \neq \emptyset\) do
    extract any pair \((x, y)\) from \(A\)
    add \(\pi_{x y}\) with minimum \(w\left(\pi_{x y}\right)\) in \(P_{x y}\) to \(H\) (if any)
\(\triangleright\) Phase 4: Propagation loop
while \(H \neq \emptyset\) do
    extract path \(\pi_{x y}\) with minimum \(w\left(\pi_{x y}\right)\) from \(H\)
    if \(w\left(\pi_{x y}\right)\) is larger than the smallest weight in \(P_{x y}\) then continue
    add \(P_{x y}^{*}\) to \(Q\)
    add \(\pi_{x y}\) to \(P_{x y}^{*}, L^{*}\left(r\left(\pi_{x y}\right)\right)\), and \(R^{*}\left(\ell\left(\pi_{x y}\right)\right)\)
    for each \(\pi_{x^{\prime} b} \in L^{*}\left(\ell\left(\pi_{x y}\right)\right)\) do
        if \(\left(x^{\prime}, x\right) \circ \pi_{x y} \in L\left(\pi_{x y}\right)\) then continue
        \(\pi_{x^{\prime} y} \leftarrow\left(x^{\prime}, x\right) \circ \pi_{x y}\)
        \(w\left(\pi_{x^{\prime} y}\right) \leftarrow w_{x^{\prime} x}+d_{x y}\)
        \(\ell\left(\pi_{x^{\prime} y}\right) \leftarrow \pi_{x^{\prime} b}, r\left(\pi_{x^{\prime} y}\right) \leftarrow \pi_{x y}\)
        add \(\pi_{x^{\prime} y}\) to \(P_{x^{\prime} y}, L\left(\pi_{x y}\right), R\left(\pi_{x^{\prime} b}\right)\), and \(H\)
    for each \(\pi_{a y^{\prime}} \in R^{*}\left(r\left(\pi_{x y}\right)\right)\) do
        if \(\pi_{x y} \circ\left(y, y^{\prime}\right) \in R\left(\pi_{x y}\right)\) then continue
        \(\pi_{x y^{\prime}} \leftarrow \pi_{x y} \circ\left(y, y^{\prime}\right)\)
        \(w\left(\pi_{x y^{\prime}}\right) \leftarrow d_{x y}+w_{y y^{\prime}}\)
        \(\ell\left(\pi_{x y^{\prime}}\right) \leftarrow \pi_{x y}, r\left(\pi_{x y^{\prime}}\right) \leftarrow \pi_{a y^{\prime}}\)
        add \(\pi_{x y^{\prime}}\) to \(P_{x y^{\prime}}, L\left(\pi_{a y^{\prime}}\right), R\left(\pi_{x y}\right)\), and \(H\)
    \(\triangleright\) Phase 5: Delete all LSPs \(\pi\) that stopped being LSP
            because \(\ell(\pi)\) or \(r(\pi)\) stopped being SP
while \(Q \neq \emptyset\) do
    extract any \(\pi_{x y}\) from \(Q\)
    for each \(\pi_{x^{\prime} y} \in L\left(\pi_{x y}\right)\) do
        remove \(\pi_{x^{\prime} y}\) from \(P_{x^{\prime} y}, R\left(\left(x^{\prime}, x\right) \circ \ell\left(\pi_{x, y}\right)\right)\) and \(L\left(\pi_{x, y}\right)\)
        for each \(\pi_{x y^{\prime}} \in R\left(\pi_{x y}\right)\) do
        remove \(\pi_{x y^{\prime}}\) from \(P_{x y^{\prime}}, L\left(r\left(\pi_{x, y}\right) \circ\left(y, y^{\prime}\right)\right)\) and \(R\left(\pi_{x, y}\right)\)
```

Fig. 1. The slightly modified APSP algorithm of Demetrescu and Italiano [5].
To prove the desired bound on $\mu(p)$ we also need the following three technical lemmas.
Lemma 6. Let $G \in \mathcal{G}(n, p)$ with $p n \geq 4 \log n$. Then every shortest path in $G$ has weight $\mathcal{O}\left(\frac{\log ^{2} n}{n}\right)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$.
Proof. Let us consider two random graphs $G_{1}, G_{2} \in \mathcal{G}\left(n, \frac{2 \log n}{n}\right)$ on the same set of vertices and let $G \cup$ be the union of $G_{1}$ and $G_{2}$ (union of the edge sets). Then we get $G_{\cup} \in \mathcal{G}\left(n, \frac{4 \log n}{n}-\frac{4 \log ^{2} n}{n^{2}}\right.$ ). By Lemmas 2 and $3, G_{i}$ is connected and diam $\left(G_{i}\right) \leq 2 \log n$ with probability $1-\mathcal{O}\left(n^{-1}\right)$ for $i=1,2$. As the two random graphs are chosen independently, at least one of them is
connected and has diameter $\mathcal{O}(\log n)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. By construction, this also holds for $G_{\cup}$. Therefore all $G \in \mathcal{G}(n, p)$ with $p n \geq 4 \log n$ are connected and have a diameter of order $\mathcal{O}(\log n)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$.

We now prove that every shortest path in $G$ has a total weight of $\mathcal{O}\left(\frac{\log ^{2} n}{p n}\right)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. For this, recall that the edge weights of $G$ are chosen uniformly at random from $[0,1]$. Let us consider the subgraph $G_{0}=\left(V, E_{\leq} \frac{4 \log n}{p n}\right)$ consisting of all vertices but only those edges of $G$ with weight at most $\frac{4 \log n}{p n}$. This is a random graph $g\left(n, \frac{4 \log n}{n}\right)$ with weights chosen uniformly at random from $\left[0, \frac{4 \log n}{p n}\right]$. As we have shown above, $G_{0}$ is connected and has diameter of order $\mathcal{O}(\log n)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. This implies that every shortest path in $G_{0}$ has a total weight of $\mathcal{O}\left(\frac{\log ^{2} n}{p n}\right)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. This upper bound also holds with the same probability for all shortest paths in $G$.

Lemma 7. Let $G \in \mathcal{g}(n, p)$. The subgraph $G_{\mathrm{SP}}=\left(V, E_{\mathrm{SP}}\right)$ of all edges that are shortest paths in $G$ fulfills for all $\varepsilon>0, \Delta\left(G_{\mathrm{SP}}\right) \leq n^{\varepsilon}$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. In particular, $|\mathrm{LSP}| \leq|\mathrm{SP}| n^{\varepsilon}$ with probability $1-\mathcal{O}\left(n^{-2}\right)$.

Proof. Let us first consider the case $p n \geq 4 \log n$. We know from Lemma 6 that all elements of $E_{\text {SP }}$ have weight $\mathcal{O}\left(\frac{\log ^{2} n}{p n}\right)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. Thus, $G_{\text {SP }}$ is a subgraph of a random graph in $g\left(n, \frac{\log ^{2} n}{n}\right)$. Therefore, we can prove the first claim of the lemma by bounding $\Delta\left(G^{\prime}\right)$ for $G^{\prime} \in \mathcal{G}\left(n, \frac{\log ^{2} n}{n}\right)$. Using Stirling's formula, the probability for a vertex in $G^{\prime}$ to have a degree greater or equal $n^{\varepsilon}$ is at most

$$
\binom{n}{n^{\varepsilon}}\left(\frac{\log ^{2} n}{n}\right)^{n^{\varepsilon}} \leq \frac{\left(\log ^{2} n\right)^{n^{\varepsilon}}}{\sqrt{2 \pi n^{\varepsilon}}\left(\frac{n^{\varepsilon}}{e}\right)^{n^{\varepsilon}}} \leq n^{-\varepsilon n^{\varepsilon / 2}}
$$

for $n$ large enough. Hence, the probability for $\Delta\left(G^{\prime}\right)$ to be greater or equal $n^{\varepsilon}$ is at most

$$
\begin{aligned}
1-\left(1-n^{-\varepsilon n^{\varepsilon / 2}}\right)^{n} & \leq 1-\left(\left(1-n^{-\varepsilon n^{\varepsilon / 2}}\right)^{n^{\varepsilon n^{\varepsilon / 2}}-1}\right)^{2 n^{1-\varepsilon n^{\varepsilon / 2}}} \\
& \leq 1-e^{-2 n^{1-\varepsilon n^{\varepsilon / 2}} \leq n^{-n^{\varepsilon / 3}}}
\end{aligned}
$$

where we used $2 n^{1-\varepsilon n^{\varepsilon / 2}} \leq n^{-n^{\varepsilon / 3}}$ for $n$ large enough and $1+x \leq e^{x}$ for all $x \in \mathbb{R}$. This proves $\Delta\left(G_{\mathrm{SP}}\right) \leq n^{\varepsilon}$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. The second claim is a consequence of the fact, that each locally shortest path from vertex $x$ to vertex $y$ is uniquely determined by its first and also by its last edge. Moreover, every locally shortest path with at least 2 edges starts and ends with edges that are shortest paths themselves. Thus, there are at most $\mathcal{O}\left(n^{\varepsilon}\right)$ locally shortest paths for each shortest path with probability at least $1-\mathcal{O}\left(n^{-2}\right)$, which proves the lemma for $p n \geq 4 \log n$.

Let $p n<4 \log n$. We consider $G^{\prime} \in \mathcal{g}\left(n, \frac{4 \log n}{n}\right)$ with edge weights chosen randomly in $\left[0, \frac{4 \log n}{p n}\right]$. Although the weights of this graph are scaled up by the factor $\frac{4 \log n}{p n}$, we get $\Delta\left(G_{\mathrm{SP}}^{\prime}\right) \leq n^{\varepsilon}$ with probability $1-\mathcal{O}\left(n^{-2}\right)$, since the scaling has no effect on the subgraph $G_{\text {sp }}^{\prime}$ of all shortest path edges in $G^{\prime}$. Now the subgraph $G$ of all edges of $G^{\prime}$ with weight less or equal 1 is a $\mathcal{g}(n, p)$-graph with edge weights chosen uniformly at random from $\left[0, \frac{4 \log n}{p n}\right]$. Hence, every edge that is a shortest path in $G$ is also a shortest path in $G^{\prime}$. With this we get $\Delta\left(G_{\mathrm{SP}}\right) \leq n^{\varepsilon}$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. The second claim follows with the same arguments as in the case $p n \geq 4 \log n$.

Lemma 8. Let $G \in \mathcal{G}(n, p)$ with $p n \geq 1 / 2$. For all $\varepsilon>0, \mu(p)=\mathcal{O}\left(\boldsymbol{E}[\operatorname{diam}(G)]|\operatorname{sp}| n^{\varepsilon-2} / p\right)$.
Proof. We consider the subgraph $G^{\prime}=\left(V, E_{<1 /(2 p n)}\right)$ containing all vertices of $G$ but only those edges with weight less than $1 /(2 p n)$. Then by Lemma $1 G \in \mathcal{G}(n, 1 /(2 n))$ and the largest component of $G^{\prime}$ is of order $\mathcal{O}(\log n)$ with probability $1-\mathcal{O}\left(n^{-2}\right)$. Thus, every path in $G^{\prime}$ contains $\mathcal{O}(\log n)$ edges with probability $1-\mathcal{O}\left(n^{-2}\right)$. The expected weight of the heaviest element of LSP is at most $\boldsymbol{E}[\operatorname{diam}(G)]$. Moreover, in the case $p \geq \frac{4 \log n}{n}$ the expected weight of the heaviest element of LSP is at most $\mathcal{O}\left(\frac{\log ^{2} n}{p n}\right)$ as shown in Lemma 6. Thus, in expectation the largest number of edges with a weight greater or equal $1 /(2 p n)$ in an element of LSP is of order $\mathcal{O}(\boldsymbol{E}[\operatorname{diam} G] p n)$ and $\mathcal{O}\left(\log ^{2} n\right)$ if $p \geq \frac{4 \log n}{n}$. This implies an upper bound of $\mathcal{O}\left(\boldsymbol{E}[\operatorname{diam} G] \log ^{2} n\right)$ for the maximal number of edges with weight greater or equal $1 /(2 p n)$ in locally shortest paths of $G$ in expectation for all $p \geq 1 /(2 p n)$.

By Lemma 7 we know that the bound $\mid$ LSP $\mid=\mathcal{O}\left(|\mathrm{SP}| n^{\varepsilon / 2}\right)$ is violated with probability $\mathcal{O}\left(n^{-2}\right)$. In this case we can estimate the number of locally shortest paths and shortest paths in $G$ by $\mathcal{O}\left(n^{3}\right)$ (the first edge and the other endpoint of a locally shortest path determines the path uniquely) and the length of this paths trivially by $n-1$. Since the probability for this event is $\mathcal{O}\left(n^{-2}\right)$, the contribution to the expected number of edges in the multiset of all edges of all (locally) shortest paths is $\mathcal{O}\left(n^{2}\right)$. If $|\operatorname{LSP}|=\mathcal{O}\left(|S P| n^{\varepsilon / 2}\right)$, the maximal number of edges with weight greater or equal $1 /(2 p n)$ in locally shortest paths of $G$ is $\mathcal{O}(\boldsymbol{E}$ [diam $\left.G] \log ^{2} n\right)$ in expectation. Now in every (locally) shortest path in $G$ there can only be consecutive parts of edges of $G^{\prime}$ of order $\mathcal{O}(\log n)$ and they must be followed by an edge with weight greater or equal $1 /(2 p n)$. Since
there can only be $\mathcal{O}\left(\boldsymbol{E}[\operatorname{diam} G] \log ^{2} n\right)$ of these edges in the path in expectation, the total number of edges in the longest of all (locally) shortest paths is $\mathcal{O}\left(\boldsymbol{E}[\operatorname{diam} G] \log ^{3} n\right)$ in expectation. Thus, the multiset of all edges of all (locally) shortest paths contains $\mathcal{O}\left(|\operatorname{Sp}| \boldsymbol{E}[\operatorname{diam} G] n^{\varepsilon / 2} \log ^{3} n\right)$ edges. By Chernoff bounds $G$ has $\Theta\left(p n^{2}\right)$ edges. Therefore, the average number of (locally) shortest paths through a fixed edge is $\mathcal{O}\left(\boldsymbol{E}[\operatorname{diam} G]|S P| n^{\varepsilon-2} / p\right)$.

We are now well-prepared to prove the main theorem of this section. It bounds $\mu(p)$ which is the expected number of locally shortest paths and shortest paths passing a fixed edge.

Theorem 9. Let $G \in \mathcal{g}(n, p)$. For all $\varepsilon, \varepsilon^{\prime}>0$,
(i) $\mu(p)=\mathcal{O}(1)$ for $p n<1 / 2$,
(ii) $\mu(p)=\mathcal{O}\left(n^{2 / 3}\right)$ for $1 / 2 \leq p n \leq 1-n^{-1 / 3}$,
(iii) $\mu(p)=\mathcal{O}\left(n^{1+\varepsilon}\right)$ for $1-\bar{n}^{-1 / 3} \leq p n \leq 1+n^{-1 / 3}$,
(iv) $\mu(p)=\mathcal{O}\left(n^{4 / 3+\varepsilon}\right)$ for $1+n^{-1 / 3} \leq p n \leq 1+\varepsilon^{\prime}$,
(v) $\mu(p)=\mathcal{O}\left(n^{\varepsilon} / p\right)$ for $1+\varepsilon^{\prime} \leq p n$.

Proof. (i) We bound $\mu(p)$ by the total number of paths passing a fixed edge. Let us first estimate the expected number of paths of a fixed length $k$ in $G$ going through a fixed edge. There are $k$ possible positions of the fixed edge in a path of length $k$. Furthermore, we can choose the remaining $k-1$ vertices of such a path in $\prod_{i=1}^{k-1}(n-i)$ different ways. As edges are present in $G$ with probability $p$, the expected number of paths in $G$ that go through a fixed edge is bounded above by

$$
\sum_{k=1}^{n-1} k p^{k-1} \prod_{i=1}^{k-1}(n-i) \leq \sum_{k=1}^{n-1} k(p n)^{k-1} \leq(p n)^{-1} \sum_{k=1}^{n} k(p n)^{k}
$$

Thus, the expected number of paths in $G$ going through a fixed edge is at most

$$
(p n)^{-1} \sum_{k=1}^{n} \sum_{i=k}^{n}(p n)^{k}=\sum_{k=1}^{n} \frac{(p n)^{k-1}-(p n)^{n}}{1-p n} \leq \frac{1-(p n)^{n}}{(1-(p n))^{2}} \leq \frac{1}{(1-p n)^{2}}
$$

Thus, $\mu(p)=\mathcal{O}(1)$ for $p n \leq 1 / 2$.
(ii) Using the bound in (i), we get $\mu(p)=\mathcal{O}\left(1 /(1-p n)^{2}\right)=\mathcal{O}\left(n^{2 / 3}\right)$ for $p n \leq 1-n^{-1 / 3}$.
(iii) Applying Lemma 4, we get $|\mathrm{SP}|=\mathcal{O}\left(n^{5 / 3}\right)$ with probability $1-O\left(n^{-1 / 3}\right)$ and $|S P|=\mathcal{O}\left(n^{2}\right)$ otherwise. Combining this with Lemma 8 and Lemma 5 gives $\mu(p)=O\left(n^{1+\varepsilon}\right)$.
(iv) By Lemma 5, the expected diameter of $G$ is $\mathcal{O}\left(n^{1 / 3}\right)$. Thus, Lemma 8 yields $\mu(p)=\mathcal{O}\left(n^{4 / 3+\varepsilon}\right)$.
(v) By Lemma 3, we get $\boldsymbol{E}$ [diam $G]=O(\log n)$. Now Lemma 8 yields $\mu(p)=\mathcal{O}\left(n^{\varepsilon} / p\right)$ as $n^{\varepsilon}$ is asymptotically larger than $\log n$ for all $\varepsilon>0$.

## 5. Runtime analysis

In this section we describe the runtime of our algorithm in terms of the parameter $\mu(p)$. With this, the main result Theorem 11 is an immediate corollary of the bounds on $\mu(p)$ from the previous section.

Theorem 10. Let $G \in \mathcal{G}(n, p)$. The expected runtime of our algorithm for a random edge update on $G$ is $\mathcal{O}\left(\mu(p) n^{\varepsilon}\right)$ for all $\varepsilon>0$.
Proof. To bound the runtime, we will use the quantity $\mu(p)$ which is the expected number of locally shortest paths and shortest paths through a fixed edge $e$ of $G$. If the algorithm performs the deletion of the edge $e$, this is exactly the number of paths that stop being shortest or locally shortest. In the case of the insertion, we get (almost) the same picture by making a backwards analysis. Instead of the insertion of $e$ to $G=(V, E)$ we can also investigate the deletion of $e$ from the graph $G^{\prime}=(V, E \cup\{e\})$. Therefore the quantity $\mu(p)$ is also the expected number of paths in $G^{\prime}$ that start being shortest or locally shortest. The slight modification that $G^{\prime}$ contains one edge more than $G$ has no consequence for the order of $\mu(p)$.

We bound the runtime of all five phases separately. The algorithm is running through the first phase, only if the considered edge $e$ is already in the graph and has to be deleted or updated. If this is the case, the algorithm goes through the while loop for every locally shortest path of $G$ that contains the edge $e$ at most twice, since a locally shortest path can be added to $Q$ as a left- and as a right-extension. The only part of the while loop with more than constant runtime is the removing of the path $\pi_{x y}$ from the lists $P_{x y}, L\left(r\left(\pi_{x y}\right)\right)$, and $R\left(\ell\left(\pi_{x y}\right)\right)$. Since every locally shortest path is uniquely determined by the first (respectively the last edge), we can bound the runtime of Phase 1 by $\mathcal{O}\left(\mu(p) n^{\varepsilon}\right)$ using Lemma 7.

The runtime of the second phase is constant. The while loop in the third phase has an expected length of $\mathcal{O}(\mu(p))$. Since adding the path $\pi_{x y}$ to the priority queue $H$ costs $\mathcal{O}(\log n)$, the runtime of Phase 3 is $\mathcal{O}(\mu(p) \log n)$.

For the analysis of the runtime of Phase 4 it is crucial to observe that every line in the for-loops as well as every other line is executed $\mathcal{O}(\mu(p))$ times in expectation. Moreover, the algorithm has to add the extended paths $\pi_{x^{\prime} y}$ and $\pi_{x y^{\prime}}$ to lists of locally shortest path and the priority queue $H$ which is done in time less than $\mathcal{O}\left(n^{\varepsilon}\right)$ in every execution. Thus, the runtime of the algorithm in Phase 4 is $\mathcal{O}\left(\mu(p) n^{\varepsilon}\right)$.

If the algorithm performs an insertion or an update, the set $Q$ in Phase 5 contains all shortest paths of $G$ that stop being the shortest. If the algorithm performs a deletion, $Q$ is empty. Thus, the algorithm is running through the while loop $\mathcal{O}(\mu(p))$


Fig. 2. Experimental results for the algorithm of Demetrescu and Italiano [5]. We start with an empty graph with $n=100$ vertices and insert 3000 random edges. (a) shows the measured runtimes depending on the number of inserted edges. Analogously, (b) shows the measured runtimes for the deletion of 3000 edges in a random order until the empty graph is obtained again. The horizontal axes describe the current number of edges $m$. The vertical axes show the measured runtimes averaged over three million runs.
times in expectation. The for-loops are both performed $\mathcal{O}\left(n^{\varepsilon / 2}\right)$ times using Lemma 7 in the same way as in the beginning of this proof but with $\varepsilon / 2$ instead of $\varepsilon$. In the same way, we can bound the expected runtime of the lines in the for-loops by $\mathcal{O}\left(n^{\varepsilon / 2}\right)$. Altogether this gives an expected runtime of $\mathcal{O}\left(\mu(p) n^{\varepsilon}\right)$ in Phase 5.

With this we can now conclude our main result.

Theorem 11. Let $\mathcal{R}(p)$ denote the expected runtime for an edge update in a graph $G \in \mathcal{G}(n, p)$. For all $\varepsilon, \varepsilon^{\prime}>0$ we have shown that
(i) $\mathscr{R}(p)=\mathcal{O}\left(n^{\varepsilon}\right)$ for $p n<1 / 2$,
(ii) $\mathcal{R}(p)=\mathcal{O}\left(n^{2 / 3+\varepsilon}\right)$ for $p n \leq 1-n^{-1 / 3}$,
(iii) $\mathcal{R}(p)=\mathcal{O}\left(n^{1+\varepsilon}\right)$ for $p n \leq 1+n^{-1 / 3}$,
(iv) $\mathcal{R}(p)=\mathcal{O}\left(n^{4 / 3+\varepsilon}\right)$ for $p n \leq 1+\varepsilon^{\prime}$,
(v) $\mathcal{R}(p)=\mathcal{O}\left(n^{\varepsilon} / p\right)$ for $p n \geq 1+\varepsilon^{\prime}$.

Let us give an intuition how the properties of $g(n, p)$ change when more and more edges are inserted and how this affects $\mathcal{R}(p)$. In the early stage (i) of the random graph process, the graph consists of many small components of size $O(\log n)$ which are trees or unicyclic. There, it is very fast to update edges. Soon after in stage (ii), the components become larger and it becomes likely for a new edge to connect two of them. Therefore, the expected number of new (locally) shortest paths increases significantly. In stage (iii) and (iv) a giant component grows and the algorithm has to update many (locally) shortest paths whenever the giant component catches other components of the graph. In (v) the last isolated vertex joins the giant component and the graph becomes connected. As the process evolves, the minimum degree and the connectivity grows and it becomes less and less likely that an inserted edge is a shortest path. Thus, the expected insertion costs are also going down.

## 6. Empirical observations

To show that the theoretically observed behavior indeed occurs in practice, we also performed some experiments. For this, we used the original algorithm of Demetrescu and Italiano [5] available from http://www.dis.uniroma1.it/~demetres/ experim/dsp/. Analogous to the experiments of Demetrescu and Italiano [5], we used the D-LHP code without smoothing. All experiments are conducted on 2.4 GHz Opteron machines running Debian GNU/Linux. As the number of locally shortest paths between any pair of nodes has been reported to be very small [5], we assume that the experimental performance of our algorithm described in Section 3 should be similar to that of Demetrescu and Italiano.

We start with an empty graph with $n=100$ vertices and add 3000 edges in a random order. Fig. 2(a) shows the measured runtimes per insertion averaged over three million runs. Afterwards, we examine the opposite direction and remove all edges in a random order. The measured average runtimes per deletion are shown in Fig. 2(b). Note that as predicted in Theorem 11, both charts identify the largest update complexity shortly after the critical window.

## Acknowledgements

Thanks are due to Daniel Johannsen, Frank Neumann, and Yuval Peres for various helpful discussions. The first author was partially supported by a postdoctoral fellowship from the German Academic Exchange Service (DAAD).

## References

[1] L. Addario-Berry, N. Broutin, B. Reed, Critical random graphs and the structure of a minimum spanning tree, Random Struct. Algorithms 35 (2009) 323-347.
[2] B. Bollobás, Random Graphs, Cambridge Univ. Press, 2001.
[3] F. Chung, L. Lu, The diameter of sparse random graphs, Adv. Appl. Math. 26 (2001) 257-279.
[4] C. Demetrescu, G.F. Italiano, A new approach to dynamic all pairs shortest paths, J. ACM 51 (2004) 968-992.
[5] C. Demetrescu, G.F. Italiano, Experimental analysis of dynamic all pairs shortest path algorithms, ACM Trans. Algorithms 2 (2006) $578-601$.
[6] D. Eppstein, Z. Galil, G.F. Italiano, Dynamic graph algorithms, in: Algorithms and Theory of Computation Handbook, 1999, (chapter 8).
[7] P. Erdős, A. Rényi, On random graphs, Publ. Math Debrecen 6 (1959) 290-297.
[8] T. Friedrich, N. Hebbinghaus, Average update times for fully-dynamic all-pairs shortest paths, in: Proc. 19th International Symposium on Algorithms and Computation, ISAAC'08, in: Lecture Notes in Computer Science, vol. 5369, 2008, pp. 692-703.
[9] D. Frigioni, A. Marchetti-Spaccamela, U. Nanni, Fully dynamic shortest paths and negative cycles detection on digraphs with arbitrary arc weights, in: In Proc. 6th Annual European Symposium on Algorithms, ESA'98, in: Lecture Notes in Computer Science, vol. 1461, 1998, pp. $320-331$.
[10] V. King, Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs, in: Proc. 40th Annual Symposium on Foundations of Computer Science, FOCS'99, 1999, pp. 81-91.
[11] J.B. Kruskal, On the shortest spanning subtree of a graph and the traveling salesman problem, Proc. AMS 7 (1956) 48-50.
[12] A. Moffat, T. Takaoka, An all pairs shortest path algorithm with expected time $\mathcal{O}\left(n^{2} \log n\right)$, SIAM J. Comput. 16 (1987) 1023-1031.
[13] A. Nachmias, Y. Peres, The critical random graph, with martingales, Israel J. Math. 176 (2010) 29-43.
[14] S. Pettie, A new approach to all-pairs shortest paths on real-weighted graphs, Theoret. Comput. Sci. 312 (2004) 47-74.
[15] G. Ramalingam, T.W. Reps, On the computational complexity of dynamic graph problems, Theoret. Comput. Sci. 158 (1996) $233-277$.
[16] L. Roditty, U. Zwick, A fully dynamic reachability algorithm for directed graphs with an almost linear update time, in: Proc. 36th Annual ACM Symposium on Theory of Computing, STOC'04, 2004, pp. 184-191.
[17] L. Roditty, U. Zwick, On dynamic shortest paths problems, in: Proc. 12th Annual European Symposium on Algorithms, ESA'04, in: Lecture Notes in Computer Science, vol. 3221, 2004, pp. 580-591.
[18] L. Roditty, U. Zwick, Dynamic approximate all-pairs shortest paths in undirected graphs, in: Proc. 45th Annual Symposium on Foundations of Computer Science, FOCS'04, 2004, pp. 499-508.
[19] T. Takaoka, A new upper bound on the complexity of the all pairs shortest path problem, Inf. Process. Lett. 43 (1992) 195-199.
[20] M. Thorup, Worst-case update times for fully-dynamic all-pairs shortest paths, in: Proc. 37th Annual ACM Symposium on Theory of Computing, STOC'05, 2005, pp. 112-119.
[21] M. Thorup, Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles, in: Proc. 9th Scandinavian Workshop on Algorithm Theory, SWAT'04, in: Lecture Notes in Computer Science, vol. 3111, 2004, pp. 384-396.
[22] U. Zwick, All pairs shortest paths using bridging sets and rectangular matrix multiplication, J. ACM 49 (2002) 289-317.


[^0]:    4. A conference version of the paper [8] was presented at the 19th International Symposium on Algorithms and Computation

    * Corresponding author. Fax: +496819325 199.

