# Replacement Paths via Fast Matrix Multiplication 

Oren Weimann<br>Department of Computer Science and Applied Mathematics<br>The Weizmann Institute of Science<br>Rehovot, Israel<br>oren.weimann@weizmann.ac.il

Raphael Yuster<br>Department of Mathematics<br>University of Haifa<br>Haifa, Israel<br>raphy@math.haifa.ac.il


#### Abstract

Let G be a directed edge-weighted graph and let $P$ be a shortest path from s to $t$ in $G$. The replacement paths problem asks to compute, for every edge $e$ on $P$, the shortest s-to-t path that avoids e.

Apart from approximation algorithms and algorithms for special graph classes, the naive solution to this problem removing each edge $e$ on $P$ one at a time and computing the shortest s-to-t path each time - is surprisingly the only known solution for directed weighted graphs, even when the weights are integrals. In particular, although the related shortest paths problem has benefited from fast matrix multiplication, the replacement paths problem has not, and still required cubic time.

For an n-vertex graph with integral edge-lengths between -M and $M$, we give a randomized algorithm that uses fast matrix multiplication and is sub-cubic for appropriate values of M . We also show how to construct a distance sensitivity oracle in the same time bounds. A query ( $\mathbf{u}, \mathbf{v}, \mathrm{e}$ ) to this oracle requires sub-quadratic time and returns the length of the shortest u-to$v$ path that avoids the edge $e$. In fact, for any constant number of edge failures, we construct a data structure in sub-cubic time, that answer queries in sub-quadratic time. Our results also apply for avoiding vertices rather than edges.


## I. Introduction

In a network $G$ where edges occasionally fail, a shortest path $P$ between two vertices $s$ and $t$ may have to change accordingly. The replacement paths problem asks to compute, for every edge $e$ in $P$, the shortest $s$-to- $t$ path that avoids $e$. This generalization of the fundamental shortest paths problem is strongly motivated by two applications. In auction theory, the replacement paths problem is used to compute the Vickrey pricing of edges owned by selfish agents [1], [2]. Another application is that of computing the $k$ shortest simple paths between a pair of vertices. This problem reduces to running $k$ executions of a replacement paths algorithm, and has many applications itself [3]. The all pairs generalization of replacement paths (also called distance sensitivity oracles in the literature) asks to construct a data structure that upon query $(u, v, e)$ outputs the length of a shortest $u$-to- $v$ path that avoids the edge $e$. Other generalizations include avoiding more than one edge, or avoiding vertices rather than edges.

Replacement Paths: The naive solution to the replacement paths problem is to remove each edge $e$ on
$P$, one at a time, and compute the shortest $s$-to- $t$ path each time. This can be done in $O\left(m n+n^{2} \log n\right)$ time. Except for the slight improvement of Gotthilf and Lewenstein [4] to $O\left(m n+n^{2} \log \log n\right)$, assuming non-negative weights, no faster algorithms are known for weighted, directed graphs. Similarly, the fastest algorithm for finding $k$ shortest simple paths, given by Yen [5] and Lawler [6], uses $k$ replacement path computations and therefore runs in $O\left(k\left(m n+n^{2} \log \log n\right)\right)$ time. Hershberger et al. [7] gave an $\Omega(m \sqrt{n})$ lower bound for both these problems in the path-comparison model of Karger et al. [8].

Overcoming the $o(m n)$ bound for the replacement paths problem (and the $o(k m n)$ bound for $k$ shortest simple paths) has received a lot of attention recently. There were two directions taken. The first was to consider special graph classes. For undirected graphs, the problem is significantly easier. Malik et al. [9] presented an $\tilde{O}(m)$ time algorithm for the replacement paths problem, that was later extended to avoid vertices rather than edges [10]. Nardelli et al. [11] gave an $O(m \alpha(m, n))$ time algorithm for the problem, in a stronger model of computation, using the linear time single source shortest paths algorithm of Thorup [12]. For directed unweighted graphs, Roditty and Zwick [13] gave an $\widetilde{O}(m \sqrt{n})$ time randomized algorithm. Finally, for the class of directed planar graphs, Emek et al. [14] gave an $O\left(n \log ^{3} n\right)$ time algorithm, improved in [15] and later in [16] to $O(n \log n)$.

The second way of overcoming the $o(m n)$ bound was to settle for approximate distances. Roditty [17] showed that approximation can in fact beat the $O(k m n)$ bound for $\underset{\sim}{k}$ shortest simple paths. In particular, Roditty presented an $\tilde{O}(k m \sqrt{n})$-time 3/2-approximation algorithm for finding the $k$ shortest simple $s$-to- $t$ paths in directed graphs with positive edge lengths. Bernstein [18] improved this to a $(1+\varepsilon)$-approximation algorithm requiring $\tilde{O}(\mathrm{~km} / \varepsilon)$ time. Bernstein also gave the first approximation algorithm for the replacement paths problem. His replacement path algorithm is a $(1+\varepsilon)$-approximation running in $\tilde{O}(m \log (n C / c) / \varepsilon)$ time, where $C$ is the largest edge length in the graph and $c$ is the smallest.

All Pairs Replacement Paths: The more general all pairs replacement paths problem asks to construct a distance
sensitivity oracle for answering queries $(u, v, e)$ seeking the length of the shortest $u$-to- $v$ path that avoids the edge $e$. For directed weighted graphs, Demetrescu et al. [19] presented a distance sensitivity oracle that can be constructed in $\tilde{O}\left(m n^{2}\right)$ time and $O\left(n^{2} \log n\right)$ space, has constant query time, and can also handle queries for avoiding vertices rather than edges. Bernstein and Karger improved the preprocessing to $O\left(n^{2} \sqrt{m}\right)$ [20] and then further to $\tilde{O}(m n)$ [21], with unchanged constant query time. For the case of two failures (vertices or edges) Duan and Pettie [22] construct an oracle in polynomial time and $\tilde{O}\left(n^{2}\right)$ space, whose query time is $O(\log n)$. For more than two failures, Chechick et al. [23] showed that if we are willing to settle for approximate distances then an oracle constructed in polynomial time can handle $f$ failures in $\tilde{O}(f)$ query time.

Finally, we mention the (harder) related problem of maintaining all pairs shortest path queries while updating the graph (deleting and inserting edges and vertex). Demetrescu and Italiano [24] presented a data structure to maintain all pairs shortest path queries in $O(1)$ time with $\tilde{O}\left(n^{2}\right)$ amortized update. Their algorithm was slightly improved by Thorup [25]. For unweighted undirected graphs, Roditty and Zwick [26] obtained a $(1+\varepsilon)$-approximation algorithm with an expected amortized update time of $\tilde{O}(m n / t)$ and worstcase query time of $O(t)$. Thorup [27] obtained $O\left(n^{2.75}\right)$ worst case updates.

Our Results: Although several algorithms exploiting fast matrix multiplication are known for shortest path problems it was unknown how to use fast matrix multiplication for the tightly-connected replacement paths problem. We show how to exploit fast matrix multiplication for the replacement paths problem in directed graphs with integral (positive and negative) edge lengths. Our first main result is a randomized algorithm whose bounds are given by the following theorem.

Theorem 1 The replacement paths problem can be solved in $\tilde{O}\left(M n^{1+\frac{2}{3} \omega}\right)=O\left(M n^{2.584}\right)$ time with very high probability on an $n$-vertex directed graph with integral edgelengths in $\{-M, \ldots, M\}$.

There are a few things to notice about the bounds of Theorem 1. Using the current matrix multiplication exponent $\omega<2.376$ [28] we get a bound of $O\left(n^{2.584}\right)$ for fixed $M$, and sub-cubic time for any $M<n^{0.416}$. A surprising outcome is that if $\omega<2.25$, as may turn out to be the case, we get, for general graphs, better bounds than the RodittyZwick $O\left(n^{2.5}\right)$ algorithm that can only handle unweighted graphs, or graphs with small positive weights. In particular, if $\omega=2$ we get $O\left(M n^{2.333}\right)$.

We further show how to extend our method to the all pairs replacement paths problem, by constructing a distance
sensitivity oracle ${ }^{1}$ with sub-cubic construction and subquadratic (notably, sub-Dijkstra) query time as stated by the following theorem.

Theorem 2 For any $0<\alpha<1$, a distance sensitivity oracle can be constructed in $\tilde{O}\left(M n^{1-\alpha+\omega}\right)$ time on an $n$-vertex directed graph with integral edge-lengths in $\{-M, \ldots, M\}$. A query $(u, v, e)$ to this oracle requires $\tilde{O}\left(n^{1+\alpha}\right)$ time and returns the length of the shortest u-to-v path that avoids the edge $e$.

Using the same ideas of Theorem 2 we can get a more general distance sensitivity oracle that can handle multiple failures with the following bounds.

Corollary 1 For any $0<\alpha<1$, a distance sensitivity oracle for avoiding $f$ edges, can be constructed in $\tilde{O}\left(M n^{1-\alpha+\omega}\right)$ time on an $n$-vertex directed graph with integral edge-lengths in $\{-M, \ldots, M\}$. A query $(u, v, S)$ to this oracle requires $\tilde{O}\left(n^{2-(1-\alpha) / f}\right)$ time and returns the length of the shortest u-to-v path that avoids all the edges in $S$.

Finally, we mention that both our replacement paths algorithm and our distance sensitivity oracle can be made to work, in the same time bounds, for the case of failed vertices rather than edges.

Technique: In contrast to the naive replacement paths algorithm that performs $O(n)$ shortest path computations, we compute shortest paths on $o(n)$ random subgraphs. We start with random subgraphs $G_{j}$ that with very high probability capture all short replacement paths. We then choose a random subset of vertices $B$ such that with very high probability any long replacement path in $G_{j}$ decomposes into short subpaths with endpoints in $B$. These subpaths are captured by the all-pairs distances between vertices of $B$.

To compute these $B \times B$ distances, we use fast matrix multiplication and tweak an algorithm of Yuster and Zwick [29] so that it computes all $|B|^{2}$ distances faster than computing each distance individually. Solving the replacement paths problem then boils down to computing an $s$-to- $t$ shortest path in a new graph with $B$ vertices and (unbounded) edge-lengths corresponding to the minimal $B \times B$ distances between a few $G_{j}$ s. To compute the $s$-to- $t$ shortest path in this new graph, we can use a (costly) shortest path algorithm that can handle unbounded (and negative) lengths. We show how to reweigh the graph so that Dijkstra's algorithm can be used instead.

## II. The Replacement Paths Algorithm

In the replacement paths problem we are given a directed graph $G=(V, E)$ with integral (and possibly negative)

[^0]edge-lengths in $\{-M, \ldots, M\}$ and two distinct vertices $s$ and $t$. Let $P=\left(s=v_{0}, v_{1}, \ldots, v_{k}=t\right)$ be a shortest path from $s$ to $t$ where $1 \leq k<n$ (we assume no negative cycles; such cycles will be detected if they exist). We wish to compute paths $P_{1}, \ldots, P_{k}$ such that $P_{i}$ is a shortest path from $s$ to $t$ in the graph $G \backslash e_{i}$ where $e_{i}$ is the edge $\left(v_{i-1}, v_{i}\right)$. In this section, we show how to solve the replacement paths problem in the bounds of Theorem 1. We begin with a high level outline of our replacement paths algorithm. We then give a detailed description and analysis.

## A. Outline of the Algorithm.

The naive solution to the replacement paths problem would be to consider the subgraphs $G \backslash e_{i}$ for $i=1, \ldots, k$, and compute an $s$-to- $t$ shortest path in each of them. As the edge lengths are bounded by $M$, every $s$-to- $t$ path can be found in $\tilde{O}\left(M n^{\omega}\right)$ time using the single-source shortest-paths (SSSP) algorithm of Yuster-Zwick [29], or in $\tilde{O}(\sqrt{n} m)=\tilde{O}\left(n^{2.5}\right)$ time with the SSSP algorithm of Goldberg [30]. However, $k$ executions of an SSSP algorithm are too costly as $k$ can be $O(n)$.

The Random subgraphs $\boldsymbol{G}_{\dot{j}}$ : Instead of $k=O(n)$ graphs, we wish to construct $r=\tilde{O}\left(n^{1-\alpha}\right)$ graphs for some $0<\alpha<1$ to be chosen later. These graphs, $G_{1}, \ldots, G_{r}$, are random subgraphs of $G$ that are generated independently as follows: Each random subgraph $G_{j}$ is obtained from $G$ by removing every edge with probability $n^{\alpha-1}$.

Let $F_{e}$ denote the set of graphs $G_{j}$ that do not contain the edge $e$. We would like to have the property that with very high probability, for every $i$, the edges of $P_{i}$ are all present in at least one $G_{j} \in F_{e_{i}}$. This way, we could run an SSSP algorithm on every $G_{j}$ and report $P_{i}$ as the minimal $s$-to- $t$ shortest path in all $G_{j} \in F_{e_{i}}$. It turns out, that by choosing the appropriate $r$, the property holds for every $P_{i}$ that is sufficiently short (shorter than $n^{1-\alpha}$ ). The problem is how to handle the long $P_{i} \mathrm{~s}$. For these $P_{i} \mathrm{~s}$, the desired property is not guaranteed to hold with high probability. It does however hold for short subpaths of $P_{i}$.

We define an interval as a subpath of $P_{i}$ consisting of $n^{1-\alpha}$ consecutive vertices, so every $P_{i}$ induces at most $n$ (overlapping) intervals for a total of $n^{2}$ intervals. We show that with very high probability, the edges of any interval induced by $P_{i}$ are all present in at least one $G_{j} \in F_{e_{i}}$. However, we cannot assure that all the intervals of $P_{i}$ are in the same $G_{j}$. To overcome this, we pick a random subset $B \subseteq V$ of $\tilde{O}\left(n^{\alpha}\right)$ vertices and make sure that $s, t \in B$. We show that with very high probability each of the $n^{2}$ possible intervals has at least one vertex in $B$. This way, every $P_{i}$ decomposes into disjoint intervals whose endpoints are both in $B$.

For every $i=1, \ldots, k$ we construct the dense distance graph $G_{B}^{i}$ as follows: Its set of vertices is $B$, and the weight of the edge $(u, v)$ is the length of the shortest $u$-to- $v$ path in all $G_{j} \in F_{e_{i}}$. The shortest $s$-to- $t$ path in $G_{B}^{i}$ will give us the
replacement path $P_{i}$. We are thus left with two challenges: How to construct the $G_{B}^{i} \mathrm{~s}$, and how to compute the shortest $s$-to- $t$ path in every $G_{B}^{i}$.

The Dense distance graphs $G_{B}^{i}$ : Constructing the $G_{B}^{i} \mathrm{~s}$ boils down to computing, for every $G_{j}$ the distance between any two vertices in $B$. Naturally, this can be achieved by computing all-pairs shortest-paths on $G_{j}$, but notice that we are only interested in the $B \times B$ subset shortest paths. The algorithm of Yuster-Zwick [29] performs this task faster. It constructs, in $\tilde{O}\left(M n^{\omega}\right)$ time, an $n \times n$ matrix $D$ which has the property that the distance between any pair of vertices $u, v$ is equal to $\min _{\ell \in V}\{D[u, \ell]+D[\ell, v]\}$ and can therefore be computed from $D$ in $O(n)$ time. The entire $B \times B$ distances can thus be naively computed from $D$ in $O\left(|B|^{2} n\right)$ time.

We present an even faster way of doing this. We first extract from $D$ the rows $\left(D_{1}\right)$ and columns $\left(D_{2}\right)$ that correspond to $B$. We then prove that it is safe to only consider entries in $D_{1}$ and $D_{2}$ whose absolute value is bounded by $M n^{1-\alpha}$. Finally, the distance product $D_{1} \star D_{2}$, that gives us exactly the $B \times B$ distances we need, is computed using the Alon et al. [31] algorithm for distance product of matrices with an associated bound on the entries, which uses a related idea of Yuval [32].
$\boldsymbol{s}$-to- $\boldsymbol{t}$ shortest paths in $\boldsymbol{G}_{\boldsymbol{B}}^{\boldsymbol{i}}$ : We are left with the final challenge of computing the $s$-to- $t$ shortest paths in every $G_{B}^{i}$. Notice that the edge-lengths of $G_{B}^{i}$ may be positive or negative, but they are no longer bounded. Therefore, to compute the shortest $s$-to- $t$ paths, we may run Goldberg's SSSP algorithm multiple times, once for each $G_{B}^{i}$. This algorithm can handle large negative lengths. However, the similarity between the various $G_{B}^{i}$ s suggests a more efficient computation using the well known method of reduced lengths.

This method is useful in transforming a shortest-path problem involving positive and negative lengths into one involving only nonnegative lengths, which can then be solved using Dijkstra's SSSP algorithm. Traditionally, the edges of a given graph $G$ are first reweighed so they become nonnegative. Then, multiple runs of Dijkstra are performed on the same graph $G$ (now with nonnegative edge-lengths) but with different sources $s$. We introduce a variant of this method, where we reduce the lengths once and then Dijkstra is performed multiple times but on different graphs. Namely, once on each $G_{B}^{i}$.

We next fill out the missing details and analysis of the above outline. We focus on computing the length of the replacement paths, the actual paths can be easily found in the same time bounds.

As a first step, we compute some shortest $s$-to- $t$ path in $G$ in order to identify $e_{1}, \ldots, e_{k}$. This can be done in $\tilde{O}\left(M n^{\omega}\right)$ time using the SSSP algorithm of Yuster-Zwick.

## B. The Random Subgraphs $G_{j}$.

We generate the subgraphs $G_{1}, \ldots, G_{r}$ where $G_{j}$ is obtained from $G$ by removing every edge with probability $n^{\alpha-1}$. We choose $r=42 n^{1-\alpha} \log n$, and since $G$ may have $n^{2}$ edges, generating these subgraphs requires $O\left(r n^{2}\right)=$ $\tilde{O}\left(n^{3-\alpha}\right)$ time. Recall that $F_{e}$ denotes the set of graphs $G_{j}$ that do not contain the edge $e$, and let $f_{e}=\left|F_{e}\right|$. We start with the following lemma, stating that with very high probability $f_{e}$ is roughly equal to $\log n$ for all $e \in E$.

Lemma 1 The probability that $21 \log n<f_{e}<70 \log n$ for all $e \in E$ is at least $1-2 / n$.

Proof: We first show that for a single edge $e$, the probability that $f_{e}>21 \log n$ is at least $1-1 / n^{3}$. To see this, notice that the expectation of $f_{e}$ is precisely $E\left[f_{e}\right]=r n^{\alpha-1}=42 \log n$. So by Chernoff's bound (cf. [33]) we know that

$$
\operatorname{Pr}\left[f_{e}<E\left[f_{e}\right]-21 \log n\right]<e^{\frac{-(21 \log n)^{2}}{2 E\left[f_{e}\right]}}<\frac{1}{n^{3}}
$$

We therefore have, by union bound, that with very high (1$1 / n)$ probability $f_{e}>21 \log n$ for all $e \in E$. Similarly, by Chernoff's bound stating that for $a \geq(2 / 3) E\left[f_{e}\right]$ we have $\operatorname{Pr}\left[f_{e}<E\left[f_{e}\right]+a\right]<e^{-2 E\left[f_{e}\right] / 27}$ we have

$$
\operatorname{Pr}\left[f_{e}<E\left[f_{e}\right]+28 \log n\right]<e^{-2 E\left[f_{e}\right] / 27}<\frac{1}{n^{3}}
$$

So again, by union bound, we have that with very high (1$1 / n)$ probability $f_{e}<70 \log n$ for all $e \in E$.

Recall that $P_{i}$ is a shortest path from $s$ to $t$ in the graph $G \backslash e_{i}$, and an interval is a subpath of some $P_{i}$ consisting of $n^{1-\alpha}$ consecutive vertices. We say that an interval of $P_{i}$ survives if all the edges of this interval are present in some $G_{j} \in F_{e_{i}}$. We now show that all the $n^{2}$ possible intervals survive with very high probability as stated by the following lemma.

Lemma 2 The probability that all the intervals survive is at least $1-2 / n$.

Proof: Consider some specific interval $I$ of $P_{i}$, and some $G_{j} \in F_{e_{i}}$. The probability that all the $|I|-1$ edges of $I$ survived in $G_{j}$ is precisely $\left(1-n^{\alpha-1}\right)^{|I|-1} \geq$ $\left(1-n^{\alpha-1}\right)^{n^{1-\alpha}-1}>1 / e$. So the probability that $I$ does not survive in any $G_{j} \in F_{e_{i}}$ is less than $(1-1 / e)^{f_{e_{i}}}$. By Lemma 1, we can assume that $f_{e}>21 \log n$ and so $(1-1 / e)^{f_{e_{i}}}<(1-1 / e)^{21 \log n}<(1 / e)^{4 \log n}<1 / n^{4}$. In particular, by union bound, we get that with probability $1-1 / n^{2}$ all the possible $n^{2}$ intervals survive.

After establishing that all intervals survive, we need to show that with very high probability every interval has at least one vertex in $B$. We choose $B$ to be a random subset of $3 n^{\alpha} \log n$ vertices of $G$ and also require that $s, t \in B$.

Lemma 3 With probability at least $1-1 / n$ every interval contains some vertex in $B$.

Proof: Since $B$ is chosen randomly, the probability that a specific vertex $v$ does not belong to $B$ is exactly $1-$ $|B| / n=1-3 n^{\alpha-1} \log n$. So an entire interval does not belong to $B$ with probability $\left(1-3 n^{\alpha-1} \log n\right)^{n^{1-\alpha}}<1 / n^{3}$. There are at most $n^{2}$ intervals overall so by union bound we get that with probability $1-1 / n$ every interval contains a vertex in $B$.

## C. Constructing the Graphs $\boldsymbol{G}_{\boldsymbol{B}}^{i}$.

Recall that $G_{B}^{i}$ is a graph whose vertex set is $B$ (that includes $s$ and $t$ ), and the weight of the edge $(u, v)$ is the length of the shortest $u$-to- $v$ path in all $G_{j} \in F_{e_{i}}$. We can therefore construct $G_{B}^{i}$ by computing for every $G_{j}$ the $B \times B$ matrix $B_{j}$ storing the distances between any two vertices in $B$. Once we have $B_{1}, \ldots, B_{r}$ every edge-weight in $G_{B}^{i}$ can be computed by examining a single entry in every $B_{j}$ such that $G_{j} \in F_{e_{i}}$. By Lemma 1 we know that $\left|F_{e_{i}}\right|<70 \log n$ and so $G_{B}^{i}$ can be constructed in $O\left(|B|^{2} \log n\right)$ time. So we can construct all $G_{B}^{i} \mathrm{~S}$ in $O\left(|B|^{2} n \log n\right)=\tilde{O}\left(n^{2 \alpha+1}\right)$ time.

We are left with the problem of computing $B_{1}, \ldots, B_{r}$. We focus on computing a single $B_{j}$ of the $n$-vertex graph $G_{j}$. For two vertices $u, v$ let $c(u, v)$ denote the smallest number of edges on a shortest path from $u$ to $v$ and let $\delta(u, v)$ denote the distance from $u$ to $v$.

Lemma 4 ([29]) Given an n-vertex graph, the Yuster-Zwick algorithm constructs in $\tilde{O}\left(M n^{\omega}\right)$ time, an $n \times n$ matrix $D$ with the following properties: For any pair of vertices $i, j$ there exists a vertex $k$ on a shortest path realizing $c(i, j)$ so that $D[i, k]=\delta(i, k), D[k, j]=\delta(k, j)$, and $D[i, k]+$ $D[k, j]=\delta(i, j)$.

We first run the Yuster-Zwick algorithm on $G_{j}$ constructing the matrix $D$ as in Lemma 4. Consider the sub-matrix $D_{1}$ of $D$ which consists of taking only the rows that correspond to $B$. Let $D_{2}$ be the sub-matrix of $D$ that consists only of the columns that correspond to $B$. It is easy to see that $B_{j}$ is exactly the distance product $D_{1} \star D_{2}$ (the distance product $C=A \star B$ of two matrices $A$ and $B$ is defined as $\left.C[i, j]=\min _{k}\{A[i, k]+B[k, j]\}\right)$. To compute the distance product $D_{1} \star D_{2}$ we use the following result, first stated by Alon et al. [31], following a related idea of Yuval [32].

Lemma 5 ([31]) Let $A$ be an $n^{r} \times n^{s}$ matrix and let $B$ be an $n^{s} \times n^{t}$ matrix, both with elements taken from $\{-L, \ldots, L\} \cup\{+\infty\}$. Then, the distance product $A \star B$ can be computed in $\tilde{O}\left(L n^{\omega(r, s, t)}\right)$ time, where $\omega(r, s, t)$ is the matrix multiplication exponent of multiplying an $n^{r} \times n^{s}$ matrix with an $n^{s} \times n^{t}$ matrix.

We would like to compute $D_{1} \star D_{2}$ using the bounds of Lemma 5. The problem is that the elements of $D_{1}$ and $D_{2}$
are not bounded, in fact $L$ can be as large as $M n$. However, we claim that we only need to consider elements of $D_{1}$ and $D_{2}$ whose absolute value is less than $M n^{1-\alpha}$, the rest of the elements can be replaced by $+\infty$. This idea, together with the above lemmas, give the following.

Lemma 6 Every matrix $B_{j}$ can be computed in time $\tilde{O}\left(M n^{\omega}+M n^{1-\alpha} n^{\omega(\alpha, 1, \alpha)}\right)$.

Proof: We need to show that any element of $D_{1}$ and $D_{2}$ whose absolute value is greater than $M n^{1-\alpha}$ can be thought of as $+\infty$. To see this, consider a shortest path between two vertices $i, j \in B$. Recall that we assume every interval of $n^{1-\alpha}$ vertices has at least one vertex in $B$. Therefore, we are only interested in the $i$-to- $j$ shortest path if the number of its edges $c(i, j) \leq n^{1-\alpha}$. In this case, by Lemma 4, there must be some $k \leq n$ such that $c(i, k) \leq n^{1-\alpha}$, and $c(k, j) \leq n^{1-\alpha}$, and $D[i, k]+D[j, k]$ is the length of the $i$-to- $j$ shortest path. Since the absolute value of every edgelength is bounded by $M$, we get that $D[i, k] \leq M n^{1-\alpha}$ and $D[j, k] \leq M n^{1-\alpha}$. So the corresponding entries in $D_{1}$ and $D_{2}$ are bounded by $M n^{1-\alpha}$.

By Lemma 6, we can compute all the $B_{j}$ matrices in total time $\tilde{O}\left(r M n^{\omega}+r M n^{1-\alpha} n^{\omega(\alpha, 1, \alpha)}\right)$. Recalling that $r=$ $\tilde{O}\left(n^{1-\alpha}\right)$, that $\alpha \leq 1$, that $\omega \geq 2$, and that $\omega(\alpha, 1, \alpha) \leq$ $\omega \alpha+1-\alpha$ (See, e.g., Huang and Pan [34]), we have that $\tilde{O}\left(r M n^{\omega}+r M n^{1-\alpha} n^{\omega(\alpha, 1, \alpha)}\right)=\tilde{O}\left(M n^{1-\alpha+\omega}\right)$.

Corollary 2 All the $G_{B}^{i}$ graphs can be constructed in total time $\tilde{O}\left(M n^{1-\alpha+\omega}+n^{2 \alpha+1}\right)$.

## D. Computing s-to-t Shortest Paths in $G_{B}^{i}$

Finally, we need to compute the shortest $s$-to- $t$ path in every $G_{B}^{i}$. Notice that the edge-lengths of $G_{B}^{i}$ may be positive or negative, but they are no longer bounded by $M$. We may run Goldberg's SSSP algorithm (which can handle large integral negative lengths) on every $G_{B}^{i}$. However, we would like to use the faster Dijkstra algorithm. This algorithm can handle unbounded lengths but can not handle negative lengths. So, in order to use Dijkstra we must reweigh the edges so they become nonnegative, without changing the shortest $s$-to- $t$ path. We achieve this by using feasible price functions and reduced lengths.

For a directed graph $G=(V, E)$ with (possibly negative) edge-lengths $w(\cdot)$, a price function is a function $\phi$ from the vertices of $G$ to the reals. For an edge $(u, v)$, its reduced length with respect to $\phi$ is $w_{\phi}(u, v)=w(u, v)+\phi(u)-$ $\phi(v)$. A price function $\phi$ is feasible if $w_{\phi}(u, v) \geq 0$ for all edges $(u, v) \in E$. The idea behind feasible price functions is that for any two vertices $s, t \in V$, for any $s$-to- $t$ path $P$, $w_{\phi}(P)=w(P)+\phi(s)-\phi(t)$. This shows that an $s$-to- $t$ path is shortest with respect to $w_{\phi}(\cdot)$ iff it is shortest with respect to $w(\cdot)$. Moreover, the $s$-to- $t$ distance with respect to the
original lengths $w(\cdot)$ can be recovered by adding $\phi(t)-\phi(s)$ to the $s$-to- $t$ distance with respect to $w_{\phi}(\cdot)$.

The most popular feasible price function comes from single-source distances. Let $x$ be new vertex added to $G$ with an edge from $x$ to every other vertex of $G$ having weight 0 . Let $d(v)$ denote the length of the shortest path from $x$ to $v$ in $G \cup\{x\}$. Then for every edge $(u, v) \in E$, we have that $d(v) \leq d(u)+w((u, v))$, so $w_{d}((u, v)) \geq 0$ and thus $d(\cdot)$ is feasible. This means that knowing $d(\cdot)$, we can now use Dijkstra's SSSP algorithm on $G \cup\{x\}$ (with reduced lengths) from any source we choose and obtain the SSSP with respect to the original $G$. However, we would like to run Dijkstra not on $G$ but on every $G_{B}^{i}$. The following lemma states that we can use the same price function $d(v)$ on all $G_{B}^{i}$ s.

Lemma 7 The function $d(\cdot)$ is a feasible price function for every $G_{B}^{i}$.

Proof: Consider an edge $(u, v)$ in some graph $G_{B}^{i}$. We need to prove that $w_{d}((u, v)) \geq 0$. We know that $w(u, v)$ is the length of the shortest $u$-to- $v$ path in all $G_{j} \in F_{e_{i}}$. In particular, $w(u, v)$ is the length of some (not necessarily shortest) $u$-to- $v$ path $P$ in $G$. Consider the $x$-to- $v$ path in $G$ that is composed of the shortest $x$-to- $u$ path in $G$ and the path $P$. The length of this $x$-to- $v$ path is $d(u)+w(u, v)$ and therefore $d(v) \leq d(u)+w(u, v)$ so $w_{d}((u, v))=w(u, v)+$ $d(u)-d(v) \geq 0$.

Notice that we can compute $d(v)$ for every $v \in G$ in $\tilde{O}\left(M n^{\omega}\right)$ time using the Yuster-Zwick algorithm (in fact, if every vertex is reachable from $s$ then we can just use $x=s$ and recall that we have already computed the SSSP from $s$ in the beginning of the algorithm to identify $e_{1}, \ldots, e_{k}$ ). Running Dijkstra on $G_{B}^{i}$ is done in $O\left(|B|^{2}\right)$ time so on all $G_{B}^{i} \mathrm{~S}$ in $O\left(n|B|^{2}\right)=\tilde{O}\left(n^{1+2 \alpha}\right)$ time.

Total time complexity: From the above description, the total time complexity of our replacement paths algorithms is

$$
\tilde{O}\left(M n^{\omega}+n^{3-\alpha}+M n^{1-\alpha+\omega}+n^{1+2 \alpha}\right)
$$

Taking $\alpha$ to be $\omega / 3$ gives total time complexity $\tilde{O}\left(M n^{1+\frac{2}{3} \omega}\right)$ thus proving Theorem 1.

## III. Extensions to All Pairs Replacement Paths

In this section we describe the required changes to turn our replacement paths solution into an all pairs replacement paths solution. We begin with the case of avoiding a single edge. We present a distance sensitivity oracle that upon query ( $u, v, e$ ) outputs the length of the shortest $u$-to- $v$ path that avoids the edge $e$. We then show how to extend our distance sensitivity oracle to handle multiple failures. A query $(u, v, S)$ to this oracle returns the length of the shortest $u$-to- $v$ path that avoids all the edges in $S$.

## A. Handling a Single Failure

We now describe a distance sensitivity oracle for single edge failure in the bounds of Theorem 2. We divide the algorithm from the previous section into a preprocessing stage and a query stage. In the preprocessing stage, without yet knowing what edge will fail, we construct the random subgraphs $G_{j}$. We also choose the random subset of vertices $B$ and compute the $B \times B$ shortest path matrices $B_{j}$. In the query stage, knowing $(u, v, e)$, we first identify the set $F_{e}$ of all $G_{j}$ s that exclude $e$. We then add $u$ and $v$ to $B$ and accordingly extend every matrix $B_{j}$ such that $G_{j} \in F_{e}$. Using these extended $B_{j}$ matrices we can construct the graph $G_{B}^{e}$. Finally, we use reduced lengths and run Dijkstra's algorithm on $G_{B}^{e}$ (a graph with $B$ vertices, where the weight of the edge $(u, v)$ is the length of the shortest $u$-to- $v$ path in all $G_{j} \in F_{e}$ ).

The preprocessing stage is done exactly as in the previous section so the graphs $G_{j}$ are constructed in $\tilde{O}\left(n^{3-\alpha}\right)$ time. In particular, Lemmas 1, 2, 3, and 6 still hold and so computing all $B_{j}$ matrices can be done (as in Section II-C) in $\tilde{O}\left(M n^{1-\alpha+\omega}\right)$ time and $O\left(r \cdot|B|^{2}\right)=\tilde{O}\left(n^{1+\alpha}\right)$ space.

In the query stage, computing the set $F_{e}$ of all $G_{j}$ s that exclude $e$ can be done in $O(r)=\tilde{O}\left(n^{1-\alpha}\right)$ time by checking each of the $G_{j} \mathrm{~s}$ in constant time. We then need to modify every matrix $B_{j}$ where $G_{j} \in F_{e}$. By Lemma 1, there are only $O(\log n)$ such $B_{j}$ matrices as $\left|F_{e}\right|=O(\log n)$. Every such $B_{j}$ currently holds the $B \times B$ distances in $G_{j}$ and we would like to add to it the $\{u\} \times B$ and $B \times\{v\}$ distances. To do so, we recall that during the construction of $B_{j}$ in the preprocessing stage, we got an $n \times n$ matrix $D$ such that $D \star D$ is the all pairs distance matrix of $G_{j}$. We compute the $\{u\} \times B$ and the $B \times\{v\}$ distances from $D$. Each distance is computed in $O(n)$ time for a total of $O(|B| n)=\tilde{O}\left(n^{1+\alpha}\right)$ time. The next step is to construct $G_{B}^{e}$. We compute each of its $|B|^{2}$ edge-lengths by looking at a single entry in $\left|F_{e}\right|$ matrices $B_{j}$ for a total of $O\left(|B|^{2}\left|F_{e}\right|\right)=\tilde{O}\left(n^{2 \alpha}\right)$ time. Finally, a single execution of Dijkstra is done on the graph $G_{B}^{e}$ using reduced lengths in $O\left(|B|^{2}\right)$ time. The overall running time of the query stage is, therefore, $\tilde{O}\left(n^{1+\alpha}\right)$.

## B. Handling Multiple Failures

We would like to extend our distance sensitivity oracle from avoiding one edge to avoiding any fixed number of $f \geq$ 2 edges. A query $(u, v, S)$ now asks for the shortest $u$-to- $v$ path that avoids all edges in the set $S$ (where $f=|S|$ ). The oracle is essentially the same as the one we just described for avoiding a single edge. We outline the differences leading to the bounds of Corollary 1.

We begin with the randomized analysis. We now generate the random subgraphs $G_{1}, \ldots, G_{r}$ by removing every edge with probability $n^{(\alpha-1) / f}$. This time we choose $r=42 f n^{1-\alpha} \log n$. Let $F_{S}$ denote the set of graphs $G_{j}$ that do not contain any edges in $S$, and let $f_{S}=\left|F_{S}\right|$ :

- Lemma 1 should now state $21 f \log n<f_{S}<$ $70 f \log n$. The expectation $E\left[f_{S}\right]=r\left(n^{(\alpha-1) / f}\right)^{f}=$ $r n^{\alpha-1}$ just as before and the proof of Lemma 1 holds. We also change the length of an interval to be $n^{(1-\alpha) / f}$. An interval is now said to survive if all its edges appear in some $G_{j} \in F_{S}$.
- In Lemma 2, we now get that an interval $I$ survives in $G_{j}$ with probability $\left(1-n^{(\alpha-1) / f}\right)^{|I|-1}>1 / e$. So the probability that $I$ does not survive in any $G_{j} \in$ $F_{S}$ is less than $(1-1 / e)^{f_{S}}<1 / n^{6 f}$. Notice that the total number of possible intervals is no longer $n^{2}$, it is now $n^{3+f}$ since there are $n^{2}$ pairs of vertices, each pair is responsible for $n^{f}$ replacement paths, and each replacement path has up to $n$ intervals. Lemma 2 holds because, by union bound, we get that with probability $1-n^{3+f} / n^{6 f}>1-1 / n^{2}$ all the possible $n^{3+f}$ intervals survive.
- Finally, for Lemma 3, we change $|B|$ to be $3 f n^{1+(\alpha-1) / f} \log n$. A specific vertex $v$ does not belong to $B$ with probability $1-|B| / n=1-$ $3 f n^{(\alpha-1) / f} \log n$. So an entire interval does not belong to $B$ with probability $\left(1-3 f n^{(\alpha-1) / f} \log n\right)^{n^{(1-\alpha) / f}}<$ $1 / n^{3 f}$. There are at most $n^{3+f}$ intervals overall so by union bound we get that with probability $1-$ $n^{3+f} / n^{3 f}>1-1 / n$ every interval contains a vertex in $B$. Lemma 3 follows.

In the preprocessing stage, when computing $B_{j}$ (Section II-C), since an interval length is now $n^{(1-\alpha) / f}$ we now only need to consider elements of $D$ (See Lemma 4) whose absolute value is bounded by $M n^{(1-\alpha) / f}$. This fact is used to bound $L=M n^{(1-\alpha) / f}$ in Lemma 5. In the same lemma, since $|B|$ is now $3 f n^{1+(\alpha-1) / f} \log n$, the distance product of a $B \times n$ matrix and an $n \times B$ matrix is computed in $\tilde{O}\left(L n^{\omega(r, s, t)}\right)$ time with $\omega(r, s, t)=$ $\omega(1+(\alpha-1) / f, 1,1+(\alpha-1) / f)$. Lemma 6 should now state that constructing a single $B_{j}$ requires time $\tilde{O}\left(M n^{\omega}+\right.$ $\left.M n^{(1-\alpha) / f} n^{\omega(1+(\alpha-1) / f, 1,1+(\alpha-1) / f)}\right)$. Multiplying this by $r$ (recalling that $r=\tilde{O}\left(n^{1-\alpha}\right)$, that $\alpha \leq 1$, that $\omega \geq 2$, and that $\omega(x, 1, x) \leq \omega x+1-x)$ we get that the total time to construct all $B_{j} \mathrm{~s}$ is $\tilde{O}\left(M n^{1-\alpha+\omega}\right)$. The total space required for all $B_{j} \mathrm{~s}$ is $O\left(r \cdot|B|^{2}\right)=\tilde{O}\left(f^{3} n^{3-\alpha+(2 \alpha-2) / f}\right)$.

In the query stage, computing the set $F_{S}$ can be done in $O(f r)=\tilde{O}\left(n^{1-\alpha}\right)$ time by checking each of the $G_{j} \mathrm{~s}$ in $O(f)$ time. We then need to modify every matrix $B_{j}$ where $G_{j} \in F_{S}$. There are only $O(f \log n)$ such $B_{j}$ matrices as $\left|F_{S}\right|=O(f \log n)$. Adding the $\{u\} \times B$ and $B \times\{v\}$ distances to all $B_{j} \mathrm{~s}$ is done as in Section III-A in $O\left(\left|F_{S}\right| \cdot|B| n\right)=\tilde{O}\left(f^{2} n^{2-(1-\alpha) / f}\right)$ time. Finally, we construct $G_{B}^{S}$ : Its set of vertices is $B$, and the weight of the edge $(u, v)$ is the length of the shortest $u$-to- $v$ path in all $G_{j} \in F_{S}$. The construction of $G_{B}^{S}$ and the execution of Dijkstra is done similarly to Section II in $O\left(|B|^{2}\left|F_{S}\right|\right)=\tilde{O}\left(f^{3} n^{2-2(1-\alpha) / f}\right)$ time. The total query
time is thus $\tilde{O}\left(f n^{1-\alpha}+f^{2} n^{2-(1-\alpha) / f}+f^{3} n^{2-2(1-\alpha) / f}\right)=$ $\tilde{O}\left(n^{2-(1-\alpha) / f}\right)$. Corollary 1 follows.

## IV. Concluding Remarks

We have presented an algorithm for the replacement path problem in weighted directed graphs. Our algorithm uses fast matrix multiplication and runs in $\tilde{O}\left(M n^{1+\frac{2}{3} \omega}\right)$ time on an $n$-node graph with weights in $\{-M, \ldots, M\}$. We also showed how to extend this solution to an $\tilde{O}\left(M n^{1-\alpha+\omega}\right)$ time constructible oracle that upon query $(u, v, S)$ returns the length of the shortest $u$-to- $v$ path that avoids all the edges in the set $S$ in $\tilde{O}\left(n^{2-(1-\alpha) / f}\right)$ time.

Both of these solutions can be made to work in the case of avoiding vertices rather than edges. To achieve that, each random subgraph $G_{j}$ is now obtained from $G$ by removing every vertex (and its adjacent edges) with probability $n^{\alpha-1}$. Notice that the vertices of $B$ might also be removed. However, by a similar argument to Lemma 2, all the vertices of an interval survive in some $G_{j}$. In particular, by Lemma 3, some vertex of $B$ is contained in the interval and is thus not removed.

The main open question is whether we can solve the replacement paths problem in subcubic time on directed graphs with unbounded (i.e, independent of $M$ ) weights.

## REFERENCES

[1] N. Nisan and A. Ronen, "Algorithmic mechanism design," Games and Economic Behavior, vol. 35, pp. 166-196, 2001.
[2] J. Hershberger and S. Suri, "Vickrey prices and shortest paths: what is an edge worth?" in Proc. of the 42nd annual symposium on Foundations Of Computer Science (FOCS), 2001, pp. 252-259.
[3] D. Eppstein, "Finding the k shortest paths," SIAM J. Comput., vol. 28, pp. 652-673, 1999.
[4] Z. Gotthilf and M. Lewenstein, "Improved algorithms for the k simple shortest paths and the replacement paths problems," Information Processing Letters, vol. 109, no. 7, pp. 352-355, 2009.
[5] J. Yen, "Finding the k shortest loopless paths in a network," Management Science, vol. 17, pp. 712-716, 1971.
[6] E. Lawler, "A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem," Management Science, vol. 18, pp. 401-405, 1972.
[7] J. Hershberger, S. Suri, and A. Bhosle, "On the difficulty of some shortest path problems," in Proc. of the 20th Symposium on Theoretical Aspects of Computer Science (STACS), 2003, pp. 343-354.
[8] D. Karger, D. Koller, and S. J. Phillips, "Finding the hidden path: Time bounds for all-pairs shortest paths," SIAM J. Comput., vol. 22, no. 6, pp. 1199-1217, 1993.
[9] K. Malik, A. K. Mittal, and S. K. Gupta, "The k most vital arcs in the shortest path problem," Operations Research Letters, pp. 223-227, 1989.
[10] E. Nardelli, G. Proietti, and P. Widmayer, "Finding the most vital node of a shortest path," Theoretical Computer Science, vol. 296, no. 1, pp. 167-177, 2003.
[11] ——, "A faster computation of the most vital edge of a shortest path," Information Processing Letters, vol. 79, no. 2, pp. 81-85, 2001.
[12] M. Thorup, "Undirected single-source shortest paths with positive integer weights in linear time," Journal of the ACM, vol. 46, pp. 362-394, 1999.
[13] L. Roditty and U. Zwick, "Replacement paths and k simple shortest paths in unweighted directed graphs," in Proc. of the 32 nd annual International Colloquium on Automata, Languages and Programming (ICALP), 2005, pp. 249-260.
[14] Y. Emek, D. Peleg, and L. Roditty, "A near-linear time algorithm for computing replacement paths in planar directed graphs," in Proc. of the 19th ACM-SIAM Symposium On Discrete Algorithms (SODA), 2008, pp. 428-435.
[15] P. Klein, S. Mozes, and O. Weimann, "Shortest paths in directed planar graphs with negative lengths: a linear-space $o\left(n \log ^{2} n\right)$-time algorithm," in Proc. of the 20th ACM-SIAM Symposium On Discrete Algorithms (SODA), 2009, pp. 236245.
[16] C. Wulff-Nilsen, "Solving the replacement paths problem for planar directed graphs in $o(n \log n)$ time," in Proc. of the 21 st ACM-SIAM Symposium On Discrete Algorithms (SODA), 2010, pp. 756-765.
[17] L. Roditty, "On the k-simple shortest paths problem in weighted directed graphs," in Proc. of the 18th ACM-SIAM Symposium On Discrete Algorithms (SODA), 2007, pp. 9204928.
[18] A. Bernstein, "A nearly optimal algorithm for approximating replacement paths and k shortest simple paths in general graphs," in Proc. of the 21st ACM-SIAM Symposium On Discrete Algorithms (SODA), 2010, pp. 742-755.
[19] C. Demetrescu, M. Thorup, R. Chowdhury, and V. Ramachandran, "Oracles for distances avoiding a failed node or link," SIAM J. Comput., vol. 37, no. 5, pp. 1299-1318, 2008.
[20] A. Bernstein and D. Karger, "Improved distance sensitivity oracles via random sampling," in Proc. of the 19th ACMSIAM Symposium On Discrete Algorithms (SODA), 2008, pp. 34-43.
[21] ——, "A nearly optimal oracle for avoiding failed vertices and edges," in Proc. of the 41st ACM Symposium on Theory of Computing (STOC), 2009, pp. 101-110.
[22] R. Duan and S. Pettie, "Dual-failure distance and connectivity oracles," in Proc. of the 20th ACM-SIAM Symposium On Discrete Algorithms (SODA), 2009, pp. 506-515.
[23] S. Chechick, M. Langberg, D. Peleg, and L. Roditty, " $f-$ sensitivity distance oracles," Unpublished manuscript, 2009.
[24] C. Demetrescu and G. Italiano, "Experimental analysis of dynamic all pairs shortest path algorithms," in Proc. of the 15th ACM-SIAM Symposium On Discrete Algorithms (SODA), 2004, pp. 362-371.
[25] M. Thorup, "Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles," in Proc. of the 9th Scandinavian Workshop on Algorithm Theory (SWAT), 2004, pp. 384-396.
[26] L. Roditty and U. Zwick, "A fully dynamic reachability algorithm for directed graphs with an almost linear update time," in Proc. of the 36th ACM Symposium on Theory of Computing (STOC), 2004, pp. 184-191.
[27] M. Thorup, "Worst-case update times for fully-dynamic allpairs shortest paths," in Proc. of the 37th ACM Symposium on Theory of Computing (STOC), 2005, pp. 112-119.
[28] D. Coppersmith and S. Winograd, "Matrix multiplication via arithmetic progressions," Journal of Symbolic Computation, vol. 9, pp. 251-280, 1990.
[29] R. Yuster and U. Zwick, "Answering distance queries in directed graphs using fast matrix multiplication," in Proc. of the 46th annual symposium on Foundations Of Computer Science (FOCS), 2005, pp. 389-396.
[30] A. Goldberg, "Scaling algorithms for the shortest paths problem," in Proc. of the 4th ACM-SIAM Symposium On Discrete Algorithms (SODA), 1993, pp. 222-231.
[31] N. Alon, Z. Galil, and O. Margalit, "On the exponent of the all pairs shortest path problem," Journal of Computer and System Sciences, vol. 54, pp. 255-262, 1997.
[32] G. Yuval, "An algorithm for finding all shortest paths using $N^{2.81}$ infinite-precision multiplications," Information Processing Letters, vol. 4, pp. 155-156, 1976.
[33] N. Alon and J. Spencer, The probabilistic method, 2nd ed. Wiley-Interscience, 2000.
[34] X. Huang and V. Pan, "Fast rectangular matrix multiplications and applications," Journal of Complexity, vol. 14, pp. 257299, 1998.


[^0]:    ${ }^{1}$ We slightly abuse the term oracle here as oracles are usually thought of as having constant or logarithmic query time.

