

An Optimal Dual Fault Tolerant Reachability Oracle^{*†}

Keerti Choudhary

Department of Computer Science and Engineering, IIT Kanpur, Kanpur, India
keerti@cse.iitk.ac.in

Abstract

Let $G = (V, E)$ be an n -vertices m -edges directed graph. Let $s \in V$ be any designated source vertex. We address the problem of reporting the reachability information from s under two vertex failures. We show that it is possible to compute in polynomial time an $O(n)$ size data structure that for any query vertex v , and any pair of failed vertices f_1, f_2 , answers in $O(1)$ time whether or not there exists a path from s to v in $G \setminus \{f_1, f_2\}$.

For the simpler case of single vertex failure such a data structure can be obtained using the dominator-tree from the celebrated work of Lengauer and Tarjan [TOPLAS 1979, Vol. 1]. However, no efficient data structure was known in the past for handling more than one failures. We, in addition, also present a labeling scheme with $O(\log^3 n)$ -bit size labels such that for any $f_1, f_2, v \in V$, it is possible to determine in poly-logarithmic time if v is reachable from s in $G \setminus \{f_1, f_2\}$ using only the labels of f_1, f_2 and v .

Our data structure can also be seen as an efficient mechanism for verifying double-dominators. For any given $x, y, v \in V$ we can determine in $O(1)$ time if the pair (x, y) is a double-dominator of v . Earlier the best known method for this problem was using dominator chain from which verification of double-dominators of only a single vertex was possible.

1998 ACM Subject Classification G.2.2 Graph Algorithms

Keywords and phrases Fault tolerant, Directed graph, Reachability oracle, Labeling scheme

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.130

1 Introduction

Networks in most real life applications are prone to failures. These failures, though unpredictable, are transient due to some simultaneous repair process that is undertaken in the application. This motivates the research on designing fault tolerant structures for various graph problems. In the past few years, a lot of work has been done in designing fault tolerant structures for various fundamental graph problems, see [7, 9, 14].

We address the problem of building a compact data structure for answering reachability queries from a designated source on vertex failures. The only previous known result for this problem was for single failure. The single fault tolerant reachability is closely related to the notion of *dominators* as follows. Given a directed graph G and a source vertex s , we say that a vertex x dominates a vertex v if every path from s to v contains x . Lengauer and Tarjan [12] introduced a data structure called *dominator tree* which is a tree rooted at s such that for any v in G , the set of ancestors of v in the tree is precisely the set of dominators of v . Thus, for any two vertices x and v in G , v becomes unreachable from s on failure of x if and

* This work was partially supported by Google India PhD Fellowship Award.

† Full version of this article is available at <http://www.cse.iitk.ac.in/users/keerti/papers/icalp2016.pdf>.



© Keerti Choudhary;

licensed under Creative Commons License CC-BY

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).

Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;

Article No. 130; pp. 130:1–130:13



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



only if x is ancestor of v in the dominator tree. A lot of work has been done on computing dominators in optimal and near-optimal time, see [3, 10, 11, 12].

In this work, we focus on building efficient data structures for answering reachability queries upon two vertex failures. One simple solution to this problem is an $O(n)$ space and $O(n)$ query time data structure using the work of Baswana et. al [1] on fault tolerant reachability subgraphs (FTRS). They show that for every $k \geq 1$, we can compute in polynomial time a subgraph H of G with at most $2^k n$ edges that preserves the reachability information from s even after k edge/vertex failures. Such a subgraph is referred as a k -FTRS for G . Thus in this case a 2-FTRS serves as our data structure for answering reachability queries. To answer a reachability query after two vertex failures, we can run any standard graph traversal algorithm (BFS/DFS) from s in the 2-FTRS avoiding the failed vertices.

Another extreme solution to this problem is an $O(n^2)$ space data structure with $O(1)$ query time. This is possible by building n dominator trees, one for each graph $G \setminus \{y\}$, where $y \in V$. Now on failure of vertices x, y , a vertex v will be reachable from s if and only if x is not ancestor of v in the dominator tree of $G \setminus \{y\}$. To the best of our knowledge, these are the only known solutions for the dual failure case.

This brings us to the following central question of our work - Is it possible to achieve the best of the above two results? In other words, can we have an oracle of $O(n)$ space that still answers reachability queries in $O(1)$ time? We give an affirmative answer to this question. We also present a labeling scheme for this problem. Our results are summarized as follows:

Oracle: There exists a data structure of $O(n)$ size that given any two failing vertices f_1, f_2 and a query vertex v , takes $O(1)$ time to determine if v is reachable from s in $G \setminus \{f_1, f_2\}$.

Labeling Scheme: There exists a compact labeling scheme for answering reachability queries under two failures. Each vertex stores a label of $O(\log^3 n)$ bits such that for any two failing vertices f_1, f_2 and any destination vertex v , it is possible to determine whether v is reachable from s in $G \setminus \{f_1, f_2\}$ by processing the labels associated with f_1, f_2 and v only.

Our result also implies a data structure for the closely related problem of double dominator verification. A pair of vertices (x, y) is said to be *double-dominator* of a vertex v if each path from s to v contains either x or y , but none of x and y are dominators of v . Using our data structure together with the dominator-tree of Lengauer and Tarjan [12], one can obtain an $O(n)$ space data structure that for any given triplet $x, y, v \in V$ verifies in $O(1)$ time if (x, y) is double-dominator of v . The best previously known result for this could verify double-dominators only for a fixed s, v pair in $O(1)$ time using an $O(n)$ space data structure called dominator chain [16].

1.1 Related work

Demetrescu et al. [7] showed that any oracle for reporting distances from a single source, rather than just the reachability information upon vertex failures in a directed weighted graph must require $\Omega(m)$ space. Nonetheless, for the problem of reporting distances between an arbitrary pair of vertices, they give a construction of $O(n^2 \log n)$ size data structure that for any $u, v, x \in V$ reports the length of the shortest path from u to v avoiding x in constant time. Duan and Pettie [9] extended this result to dual failures by designing a data structure of $O(n^2 \log^3 n)$ space which could answer distance queries after any two vertex failures in an $O(\log n)$ time.

Parter and Peleg [14] addressed the problem of computing a sparse subgraph that preserves the distances from source s on failure of a single vertex. In particular, they show that for any given unweighted graph G we can compute a subgraph H with $O(n^{3/2})$ edges such that

for any two vertices v, x , the distance of v from s in the graph $H \setminus \{x\}$ is the same as that in $G \setminus \{x\}$. They also show that this bound is tight. Parter [13] extended this result to dual failure in undirected graphs by showing an upper bound of $O(n^{5/3})$, and also showed that this bound is tight. Baswana and Khanna [2] showed that for the undirected and single failure case if one is willing to settle for an approximation then there is a subgraph with $O(n \log n)$ edges that preserves the distances up to a multiplicative error of 3. Parter and Peleg [15] improved this result and obtained a subgraph with at most $3n$ edges.

Another closely related problem is the replacement paths problem. In this problem we are given a source s and a target t and for each edge e on the shortest path from s to t the algorithm computes the shortest path from s to t avoiding e . Many variants of this problem were studied along the years. For a recent overview see [17] and references therein.

The questions of finding graph spanners, approximate distance oracles and compact routing schemes that are resilient to f vertex or edge failures in undirected graphs were studied in [8, 5, 4].

1.2 Our Techniques

Consider a reachability tree T rooted at source s . Let v be any vertex in T , and P be the path from s to v in T . Let us first consider the simple case when only a single vertex, say x , fails in G . If x lies on P and v is still reachable from s , then we can find an alternate path consisting of - a prefix of P , followed by a “detour”, say D , avoiding P (and x), followed by a suffix of P . This simple decomposition can be used to easily handle reachability queries for one vertex failure. However, in the case of the dual failure, non trivialities arise when the second failing vertex lies on detour D . A natural direction to proceed from here is to define secondary detours which are disjoint from both P as well as D , but handling secondary detours is quite complicated.

So we take an alternative approach in which instead of initially starting with a single tree we begin with two independent trees - T_1 and T_2 , defined by Georgiadis and Tarjan [11]. They satisfy the condition that for any v , the path from s to v in two trees intersect only at the dominators of v . This allows us to reduce the problem to the case when exactly one failure is an ancestor of v in T_1 , and the other failure is ancestor of v in T_2 . Now let P_1, P_2 be the paths from s to v in T_1, T_2 , and let f_1, f_2 be the failed vertices lying respectively on P_1, P_2 . Then, the structure of an alternate path to v gets simplified as follows: It consists of a prefix of either P_1 or P_2 up to a vertex a (lying respectively before f_1 or f_2) followed by a detour which is disjoint from P_1, P_2 (and f_1, f_2), and followed again by a suffix of P_1 or P_2 . Note that the starting and terminating vertices of detour need not lie on the same tree-path.

So, we get a very clean and simple representation of detours. The main challenge lies in how to efficiently search for an appropriate detour avoiding f_1, f_2 . The problems arising and how they are tackled is discussed in detail in Section 4.

1.3 Organization of the paper

We describe notations and terminologies in Section 2. In Section 3, we briefly sketch the solution for the single failure case using detours. The overview of the paper is presented in Section 4. In Section 5, we present the reachability oracle for a simpler class of graphs that are 2-vertex strongly connected. In Section 6, we present the oracle for general graphs. The construction of the labeling scheme can be found in the full version of the paper.

2 Preliminaries

Given a directed graph $G = (V, E)$ on $n = |V|$ vertices and $m = |E|$ edges, and a source vertex $s \in V$, the following notations will be used throughout the paper.

- f_1, f_2 : A given pair of failed vertices.
- $par_T(x)$: The parent of vertex x in tree T .
- $depth_T(x)$: The depth of vertex x in tree T .
- $PATH_T(x, y)$: The path from vertex x to vertex y in tree T .
- $PATH_T(\bar{x}, y)$: $PATH_T(x, y) \setminus \{x\}$
- $PATH_T(x, \bar{y})$: $PATH_T(x, y) \setminus \{y\}$
- $PATH_T(\bar{x}, \bar{y})$: $PATH_T(x, y) \setminus \{x, y\}$
- $idom(x)$: The immediate dominator of vertex x . (See Definition 1).
- T_1, T_2 : A pair of independent trees for G rooted at s . (See Definition 2).
- $P[x, y]$: The subpath of path P lying between vertices x, y , assuming x precedes y on P .
- $P::Q$: The path formed by concatenating paths P and Q in G . Here it is assumed that the last vertex of P is the same as the first vertex of Q .

We now define immediate dominators and independent spanning trees which are crucial to our fault tolerant reachability oracle.

► **Definition 1** ([12]). A vertex w is said to be the *immediate dominator* of v if (i) w is a dominator of v , and (ii) every dominator of v (other than v itself) is also a dominator of w .

► **Definition 2** (Georgiadis and Tarjan [11]). Given a directed graph G and a designated source s , a pair of trees T_1, T_2 rooted at s are said to be *independent spanning trees* if for each $v \neq s$ the paths from s to v in T_1 and T_2 intersect only at the dominators of v . (It was shown by Georgiadis and Tarjan [11] that such a pair of trees can be computed in an $O(m)$ time).

Below we state a few basic properties of dominators in a directed graph.

► **Property 1.** Let T be a reachability tree rooted at s , and y_0, y_1 be vertices such that $y_0 = idom(y_1)$. Then for any $z \in PATH_T(\bar{y}_0, y_1)$, $idom(z)$ belongs to $PATH_T(y_0, y_1)$.


► **Property 2.** Let T be a reachability tree rooted at s , and y_1, y_2 be vertices such that y_1 is ancestor of y_2 , and $idom(y_1) = idom(y_2)$. Then for any $z \in PATH_T(y_1, y_2)$ either y_1 is a dominator of z or $idom(z) = idom(y_1)$.

For efficient implementation of our oracle, we use the following optimal result by Demaine et. al [6] for answering the range minima queries on tree paths.

► **Theorem 3** (Demaine et al. [6]). A tree T on n vertices and edge weights in the range $[0, n^3]$ can be preprocessed in $O(n \log n)$ time to build a data structure of $O(n)$ size so that given any $u, v \in T$, the edge of smallest weight on the tree path from u to v can be reported in $O(1)$ time.

The following corollary is immediate from Theorem 3.

► **Corollary 4.** Given a tree, say T , on n vertices, with each vertex assigned an integral weight in range $[0, n^3]$, we can obtain in polynomial time an $O(n)$ size data structure that for any two vertices x, y , outputs in $O(1)$ time the vertex with minimum weight on $PATH_T(\bar{x}, y)$.



■ **Figure 1** (i) Representation of a path bypassing f in the single vertex failure case; (ii) Representation of sets $S_A(v)$ and $S_B(v)$ when condition \mathcal{C} is satisfied for vertex v ; (iii) A path from $a \in S_A(v)$ to $b \in S_B(v)$ when v is reachable from s in $G \setminus \{f_1, f_2\}$.

3 Review of reachability oracle for single failure

In order to understand the dual fault tolerant reachability oracle we first briefly discuss the case of single failure. We here describe an alternative reachability oracle using detours instead of dominator tree. Let T be any reachability tree of G rooted at s , and f, v be respectively the failed and the query vertex. Also assume f is an ancestor of v in T . Notice that if v is reachable from s in $G \setminus \{f\}$, then there must exist a path starting from $PATH_T(s, \bar{f})$ and terminating at $PATH_T(\bar{f}, v)$ which, except for its endpoints, does not pass through any ancestor of v in T . (See Figure 1(i)). So for each $w \in V$, we can define a detour $D(w)$ to be a path starting from the highest possible ancestor of w in T and terminating at w such that none of the internal vertices of the path pass through an ancestor of w . Now on failure of f it suffices to search whether there exists a vertex lying in $PATH_T(\bar{f}, v)$ whose detour starts from an ancestor of f . This can be achieved by assigning to each vertex w a weight equal to the depth of the first vertex on $D(w)$. By doing this the problem of reachability under one vertex failure reduces to the problem of solving range minima on weighted trees, for which already an optimal solution exists. (See Corollary 4).


4 Overview

Let us consider the failure of a pair of vertices f_1, f_2 in G , and let v be the query vertex. Note that if any of the tree paths - $PATH_{T_1}(s, v)$ or $PATH_{T_2}(s, v)$ is intact, then v will be reachable from s . Also, if both $PATH_{T_1}(s, v)$ or $PATH_{T_2}(s, v)$ contains a common failed vertex, say f_1 , then v will not be reachable from s . This is because then f_1 would be a dominator of v . Thus the non-trivial case is when $PATH_{T_1}(s, v)$ contains only f_1 and $PATH_{T_2}(s, v)$ contains only f_2 , or the vice-versa. So whenever a query vertex v is given to us, we may assume that the following condition is satisfied.

$$\boxed{\mathcal{C} : f_1 \text{ lies on } PATH_{T_1}(s, v) \setminus PATH_{T_2}(s, v), \text{ and } f_2 \text{ lies on } PATH_{T_2}(s, v) \setminus PATH_{T_1}(s, v).}$$

Now consider the sets $S_A(v)$ and $S_B(v)$ as defined below. (For a better understanding of these sets see Figure 1(ii)).

- $S_A(v)$: Set of vertices lying either above f_1 on $PATH_{T_1}(s, v)$ or above f_2 on $PATH_{T_2}(s, v)$.
- $S_B(v)$: Set of vertices lying either below f_1 on $PATH_{T_1}(s, v)$ or below f_2 on $PATH_{T_2}(s, v)$.



■ **Figure 2** (i) A graph G with in-degree of each vertex bounded by two; (ii) A pair of independent spanning trees T_1 and T_2 for G ; (iii) A path P from $a \in S_A(v)$ to $b \in S_B(v)$ in $G \setminus \{f_1, f_2\}$.

It turns out that if v is reachable from s in $G \setminus \{f_1, f_2\}$, then there must exist a path from set $S_A(v)$ to $S_B(v)$ avoiding the vertices of both $\text{PATH}_{T_1}(s, v)$ and $\text{PATH}_{T_2}(s, v)$. This fact is formally stated in the following lemma (refer to the full version of the paper for its proof).

► **Lemma 5.** *Given a pair of failed vertices f_1, f_2 , a vertex v is reachable from s if and only if G contains a path P satisfying the following conditions. (See Figure 1(iii)).*


- C1.** *The first and last vertices of P lie respectively in sets $S_A(v)$ and $S_B(v)$.*
- C2.** *None of the internal vertices of P lies on $\text{PATH}_{T_1}(s, v)$ or $\text{PATH}_{T_2}(s, v)$.*

For simplicity we refer to a path satisfying the conditions $C1$ and $C2$ stated in the above lemma as an $S_{A,B}(v)$ path. In order to efficiently compute such a path we define a pair of detours $D^1(w)$ and $D^2(w)$ for each vertex $w \in V$ as follows.

- $D^i(w)$: a path starting from the highest possible ancestor of w in T_i and terminating at w such that none of the internal vertices of the path are ancestor of w in T_1 or T_2 .

Note that the detours $D^1(w)$ and $D^2(w)$ can be seen as a simple generalization of the detour $D(w)$ which was defined for the single failure case in Section 3. However, we show that this simple generalization is not sufficient to answer the reachability queries in dual failure. To understand this subtle point consider an $S_{A,B}(v)$ path with a, b as its endpoints. If the endpoint b is equal to v , then P could be simply either $D^1(v)$ or $D^2(v)$. The problem arises when $b \neq v$. This is because if b is an ancestor of v in T_1 , then P might contain vertices from $\text{PATH}_{T_2}(s, b)$. (Recall that the internal vertices of P are disjoint from $\text{PATH}_{T_2}(s, v)$, but not necessarily disjoint from $\text{PATH}_{T_2}(s, b)$). So in this case P can neither be $D^1(b)$ nor be $D^2(b)$. For a more clear insight into this consider the graph and its two independent spanning trees in Figure 2. Since in-degree of each vertex in the graph is at most two, $D^i(w)$ for each $w \in V$ is simply the incoming edge from $\text{par}_i(w)$ to w . Thus the path P connecting $S_A(v)$ to $S_B(v)$ is a concatenation of as many detours as there are number of edges in P . Determining whether a concatenation of all these single-edge detours can give us an $S_{A,B}(v)$ path is difficult to achieve in $O(1)$ time.

This shows that a simple generalization of detours from a single tree to two trees is not sufficient. To tackle the problem we extend the notion of detours to ‘Parent Detours’ and ‘Ancestor Detours’. These detours unlike the normal detour terminate at an appropriate ancestor of w in T_1 or T_2 . We formally define the parent-detours and ancestor-detours in the following sections, and show how they can be used to solve the problem of dual fault tolerant reachability.



■ **Figure 3** (i) Possibilities for path P when $b \in \text{PATH}_{T_1}(\bar{f}_1, \bar{v})$; (ii) Representation of $\text{PATH}_{T_2}(s, c)$ and $z = \text{LCA}(c, v)$ in T_2 ; (iii) Violation of assumption 2 if $P \cap \text{PATH}_{T_2}(z, c)$ is non-empty.

5 Reachability oracle for 2-vertex strongly connected graphs

In this section we describe an $O(n)$ space and $O(1)$ time reachability oracle for 2-vertex strongly connected graphs. By 2-vertex strongly connectedness we have that on removal of any vertex f ($f \neq s$), all the vertices in $G \setminus \{f\}$ are still reachable from s . Thus each vertex is dominated only by source s and by itself. This implies that for any vertex w , $\text{PATH}_{T_1}(s, w)$ and $\text{PATH}_{T_2}(s, w)$ intersects only at the endpoints s, w .

Consider a query vertex v which is reachable from s in $G \setminus \{f_1, f_2\}$. Let us assume that condition \mathcal{C} is satisfied for v . Let P be any $S_{A,B}(v)$ path, and a, b be respectively the first and last vertices on P . Without loss of generality we can assume that b lies on $\text{PATH}_{T_1}(s, v)$. See Figure 3(i). We make the following additional assumptions.

1. None of the $S_{A,B}(v)$ paths terminates at v (i.e. b cannot be v).
2. b is the lowest vertex on $\text{PATH}_{T_1}(s, v)$ at which an $S_{A,B}(v)$ path terminates.

► **Remark.** The assumption 1 is justified since if $b = v$, then v will be reachable from s using the detours $D^1(v)$ or $D^2(v)$.

We now state a lemma which provides the motivation for defining the parent detours.

► **Lemma 6.** Let a, b, P be as described above, and c be the child of b on $\text{PATH}_{T_1}(s, v)$. Then,

- (i) Vertex f_2 is an ancestor of c in T_2 .
- (ii) None of the internal vertices of P lie on $\text{PATH}_{T_1}(s, c)$ or $\text{PATH}_{T_2}(s, c)$.

Proof. Let z denote the LCA of vertices c and v in tree T_2 . (See Figure 3(ii)). Consider the path $Q = \text{PATH}_{T_2}(z, c)$. It is easy to see that none of the internal vertices of Q lies on $\text{PATH}_{T_2}(s, v)$. Also, the internal vertices of Q appearing on $\text{PATH}_{T_1}(s, v)$ must lie below c on $\text{PATH}_{T_1}(s, v)$. This is because, by definition of independent spanning trees, $\text{PATH}_{T_1}(s, c)$ and Q can intersect only at the vertices s and c .

We now prove claim 1. Let d be the first point of intersection of Q with $\text{PATH}_{T_1}(c, v)$. (See Figure 3(ii)). Then the internal vertices of $Q[z, d]$ are disjoint from both $\text{PATH}_{T_1}(s, v)$ and $\text{PATH}_{T_2}(s, v)$. Now if z is an ancestor of f_2 in T_2 , then $Q[z, d]$ forms an $S_{A,B}(v)$ path, terminating at descendant of b in T_1 , thereby violating assumption 2. Hence f_2 must be either same as z or an ancestor of z . This shows that f_2 is an ancestor of c in T_2 .

In order to prove claim 2, we first show that P is disjoint from Q . Let us suppose on the contrary, that there exists a vertex, say z' , belonging to $P \cap Q$. Also let d' be the first vertex of $Q[z', c]$ lying on $\text{PATH}_{T_1}(c, v)$. (See Figure 3(iii)). Then $P[a, z'] \cup Q[z', d']$ forms an

$S_{A,B}(v)$ path terminating at a descendant of b in T_1 . This again violates assumption 2. Thus $P \cap Q = \emptyset$. Now since the internal vertices of P are disjoint from $\text{PATH}_{T_1}(s, v)$, $\text{PATH}_{T_2}(s, v)$, they must be disjoint from $\text{PATH}_{T_1}(s, c)$ and $\text{PATH}_{T_2}(s, z) :: Q = \text{PATH}_{T_2}(s, c)$, as well. ◀

The above lemma implies that f_1 is an ancestor of c in T_1 , and f_2 is an ancestor of c in T_2 . Thus $S_A(v) = S_A(c)$. Hence we have the following corollary.

► **Corollary 7.** *P is an $S_{A,B}(c)$ path terminating at $\text{par}_{T_1}(c)$.*

In order to capture the above fact we define parent-detours for each $w \in V$ which instead of terminating at w terminates at either $\text{par}_{T_1}(w)$ or $\text{par}_{T_2}(w)$.

■ $\text{PD}_j^i(w)$: a path starting from the highest possible ancestor of w in T_i and terminating at $\text{par}_{T_j}(w)$ s.t. none of the internal vertices of the path lie on $\text{PATH}_{T_1}(s, w)$ or $\text{PATH}_{T_2}(s, w)$.

By above definition of parent-detour it follows that P can be replaced by either $\text{PD}_1^1(c)$ or $\text{PD}_1^2(c)$ depending upon whether it starts from an ancestor of c in T_1 or T_2 . Now let x_1 denote the child of f_1 in T_1 lying on $\text{PATH}_{T_1}(s, v)$, and x_2 denote the child of f_2 in T_2 lying on $\text{PATH}_{T_2}(s, v)$. Then the parent-detours of vertices x_1, x_2 may not be of any help, since they would terminate at f_1 and f_2 . However, the parent-detours of vertices in $S_B(v) \setminus \{x_1, x_2\}$ will suffice to determine whether v is reachable from s or not.

Notice that in above discussion, we observed that P is an $S_{A,B}(c)$ path terminating at $b = \text{par}_{T_1}(c)$. This shows that for vertices lying on $\text{PATH}_{T_1}(\bar{x}_1, v)$, we only need to worry about parent-detours terminating at $\text{par}_{T_1}(\cdot)$, i.e. $\text{PD}_1^1(\cdot)$ and $\text{PD}_1^2(\cdot)$. Whereas, for vertices on $\text{PATH}_{T_2}(\bar{x}_2, v)$, we need to worry about parent-detours terminating at $\text{par}_{T_2}(\cdot)$, i.e. $\text{PD}_2^1(\cdot)$ and $\text{PD}_2^2(\cdot)$. We thus have the following lemma.

► **Lemma 8.** *Let x_1 and x_2 be as defined above. A vertex v is reachable from s in $G \setminus \{f_1, f_2\}$ if and only if at least one of the following vertices lie in $S_A(v)$.*

- (i) *The first vertex of $D^1(v)$ or $D^2(v)$.*
- (ii) *The first vertex of either $\text{PD}_1^1(w)$ or $\text{PD}_1^2(w)$ for some $w \in \text{PATH}_{T_1}(\bar{x}_1, v)$.*
- (iii) *The first vertex of either $\text{PD}_2^1(w)$ or $\text{PD}_2^2(w)$ for some $w \in \text{PATH}_{T_2}(\bar{x}_2, v)$.*

5.1 Implementation of the oracle

We first introduce the following notations for detours and parent detours.

- $\beta^i(v)$: $\text{depth}_{T_i}(\text{first vertex on } D^i(v))$.
- $\gamma_j^i(v)$: $\text{depth}_{T_i}(\text{first vertex on } \text{PD}_j^i(v))$.

Now let f_1, f_2 be a given pair of failed vertices and v be a given query vertex. Our first step is to check if condition \mathcal{C} is satisfied. Recall that this requires only verifying the ancestor-descendant relationship in trees T_1 and T_2 . One simple method to achieve this for any given tree T is to perform the pre-order and the post-order traversal of T , and store the vertices in the order they are visited. Now x will be ancestor of y in T if and only if x appears before y in the pre-order traversal, and after y in the post-order traversal.

Algorithm 1 presents the pseudo-code for answering reachability query for a vertex v assuming condition \mathcal{C} is satisfied. This can be explained in words as follows. For $i = 1, 2$, we first check if $D^i(v)$ starts from an ancestor of f_i in T_i or not. This is done by comparing the value of $\beta^i(v)$ with the depth of f_i in T_i . Next we compute the vertices x_1, x_2 . Finally for $i, j \in \{1, 2\}$, we compute a vertex $w \in \text{PATH}_{T_j}(\bar{x}_j, v)$ for which $\gamma_j^i(\cdot)$ is minimum. If $\gamma_j^i(w)$ is less than the depth of f_i in T_i , then it implies that $\text{PD}_j^i(w)$ starts from an ancestor of f_i in T_i , so we return True. If we reach to the end of code, that means we have not been able to find any path for v , so we return False.

Algorithm 1: Oracle for reachability to v in 2-vertex strongly connected graphs.

```

1 if  $\beta^1(v) < \text{depth}_{T_1}(f_1)$  or  $\beta^2(v) < \text{depth}_{T_2}(f_2)$  then Return True;
2 ;
3  $x_1 \leftarrow$  the vertex with minimum depth on  $\text{PATH}_{T_1}(\bar{f}_1, v)$ ;
4  $x_2 \leftarrow$  the vertex with minimum depth on  $\text{PATH}_{T_2}(\bar{f}_2, v)$ ;
5 foreach  $i, j \in \{1, 2\}$  do
6    $w \leftarrow$  a vertex on  $\text{PATH}_{T_j}(\bar{x}_j, v)$  for which  $\gamma_j^i(\cdot)$  is minimum;
7   if  $\gamma_j^i(w) < \text{depth}_{T_i}(f_i)$  then Return True;
8   ;
9 end
10 Return False;
```

The above oracle can be easily implemented in $O(1)$ time, by having a total of six weight functions – one each for storing the depth of a vertex in trees T_1, T_2 , and the other four for storing the values $\gamma_j^i(\cdot)$, for $i, j \in \{1, 2\}$. By doing this the vertices x_1, x_2 can be computed in constant time since they are respectively the vertices with minimum depth on the paths $\text{PATH}_{T_1}(\bar{f}_1, v)$ and $\text{PATH}_{T_2}(\bar{f}_2, v)$. Also Step 4 can be carried out in an $O(1)$ time. We can thus state the following theorem.

► **Theorem 9.** *A 2-vertex strongly connected graph on n vertices can be preprocessed in polynomial time for a given source vertex s to build a data structure of $O(n)$ size such that for any query vertex v , and pair of failures f_1, f_2 , it takes $O(1)$ time to determine if there exists any path from s to v in $G \setminus \{f_1, f_2\}$.*

6 Reachability oracle for general graphs

In this section we explain the reachability oracle for general graphs. Consider a query vertex u in G . Let u_0, u_1, \dots, u_k be the dominators of u with $u_0 = s$ and $u_k = u$. Thus $\text{PATH}_{T_1}(s, u)$ and $\text{PATH}_{T_2}(s, u)$ intersect only at u_i 's. (See Figure 4(i)). As in Section 5, we assume that condition \mathcal{C} holds for u , so none of the u_i 's can be equal to f_1 or f_2 . Now let $i, j \in [1, k]$ be such that $f_1 \in \text{PATH}_{T_1}(\bar{u}_{i-1}, \bar{u}_i)$ and $f_2 \in \text{PATH}_{T_2}(\bar{u}_{j-1}, \bar{u}_j)$. It is easy to see that if $i \neq j$, then u is reachable from s by the path $\text{PATH}_{T_1}(s, u_{i-1})::\text{PATH}_{T_2}(u_{i-1}, u_i)::\text{PATH}_{T_1}(u_i, u)$. (See Figure 4(ii)). Thus we consider the case when $i = j$ ¹. For simplicity, we use symbols, v and $\text{idom}(v)$ to respectively denote the vertices u_i and u_{i-1} . Notice that in order to check reachability of u from s , it suffices to check if v is reachable from s in $G \setminus \{f_1, f_2\}$.

We now divide our analysis into various different cases as follows:

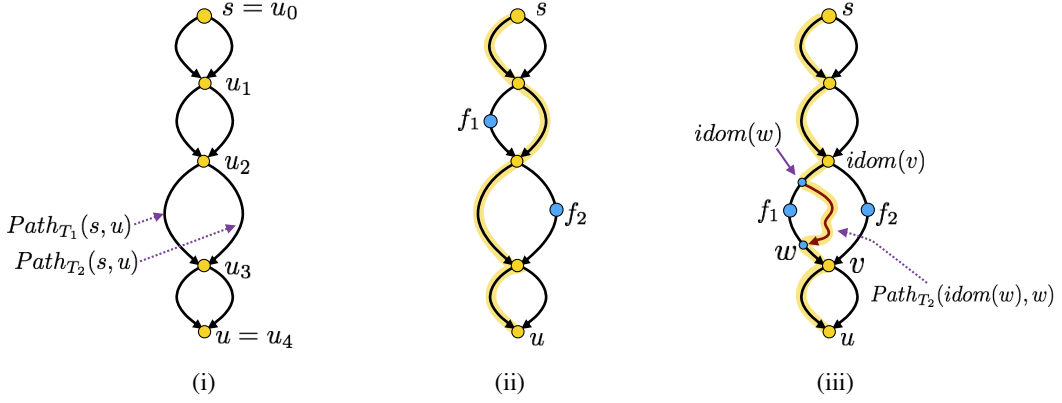
Case 1. *There exists an $S_{A,B}(v)$ path terminating at vertex v .*

In this case v will be reachable from s using either of the detours $D^1(v)$ or $D^2(v)$.

Case 2. *There exists a vertex $w \in S_B(v)$ for which $\text{idom}(w) \in S_A(v) \setminus \{\text{idom}(v)\}$.*

In this case also we can show that v is reachable from s by the following argument. Without loss of generality let us assume that w is an ancestor of v in T_1 . Since $\text{idom}(w)$ is an

¹ One can verify in $O(1)$ time whether $u_{i-1} = u_{j-1}$ (i.e. if $i = j$) since vertex $u_{i-1} = \text{LCA}(f_1, v)$ and vertex $u_{j-1} = \text{LCA}(f_2, v)$ in the dominator tree of G .



■ **Figure 4** (i) Representation of dominators of u ; (ii) A path from s to u (highlighted in yellow) when f_1 lies on $\text{PATH}_{T_1}(\bar{u}_1, \bar{u}_2)$ and f_2 lies on $\text{PATH}_{T_2}(\bar{u}_2, \bar{u}_3)$; (iii) A path from s to u (highlighted in yellow) when there exists a vertex $w \in S_B(v)$ for which $\text{idom}(w)$ lies in the set $S_A(v) \setminus \{\text{idom}(v)\}$.

ancestor of w in T_1 , it must lie on $\text{PATH}_{T_1}(\overline{\text{idom}(v)}, \bar{f}_1)$. (See Figure 4(iii)). Consider the path $Q = \text{PATH}_{T_2}(\text{idom}(w), w)$. Note that f_2 cannot lie on Q . This is because otherwise $\text{PATH}_{T_1}(\text{idom}(v), w)$ and $\text{PATH}_{T_2}(\text{idom}(v), f_2)::Q[f_2, w]$ will form two vertex disjoint paths from $\text{idom}(v)$ to w , which would violate the fact that $\text{idom}(w) \neq \text{idom}(v)$. Also f_1 cannot lie on Q , as Q is disjoint from $\text{PATH}_{T_1}(\overline{\text{idom}(w)}, \bar{w})$. Thus v is reachable from s by the path $\text{PATH}_{T_1}(s, \text{idom}(w))::Q::\text{PATH}_{T_1}(w, v)$.

Case 3. *None of the $S_{A,B}(v)$ path terminates at v , and there does not exist a vertex in $S_B(v)$ whose immediate dominator lies in $S_A(v) \setminus \{\text{idom}(v)\}$.*

This is the most non-trivial case of dual fault tolerant reachability oracle. We now provide analysis for this case.

Let us suppose v is reachable from s in $G \setminus \{f_1, f_2\}$. Then without loss of generality we can assume that there exists an $S_{A,B}(v)$ path (say P) terminating at an ancestor of v in T_1 . In case there are multiple $S_{A,B}(v)$ paths, then we take P to be that path which terminates at lowest vertex on $\text{PATH}_{T_1}(\bar{f}_1, \bar{v})$. Let a, b be respectively the first and last vertices on P . By Property 1, we know that $\text{idom}(b)$ cannot be an ancestor of $\text{idom}(v)$ in T_1 . Therefore, $\text{idom}(b)$ must be equal to $\text{idom}(v)$. This is because $\text{idom}(b)$ cannot lie in $S_A(v) \setminus \{\text{idom}(v)\}$, and if $\text{idom}(b)$ lies in $S_B(v)$ then $P \cap S_B(v)$ will contain both b and $\text{idom}(b)$, which would violate the definition of an $S_{A,B}(v)$ path.

Now consider the vertex c which is child of b on $\text{PATH}_{T_1}(s, v)$. It turns out that in a general graph the parent-detours of c may not be of any help. This is because the analysis for 2-vertex strongly connected graphs crucially exploited the fact that $\text{idom}(b) = \text{idom}(c) = s$. But in general graphs, if $\text{idom}(b)$ is not equal to $\text{idom}(c)$, then it can be shown that Lemma 6 no longer holds. To be more precise, we can show that the internal vertices of P might not be disjoint from $\text{PATH}_{T_2}(s, c)$.

However, the problem can be resolved if we take c to be the first descendant of b on $\text{PATH}_{T_1}(s, v)$ whose immediate dominator is the same as that of b . This motivates us to define the notion of pseudo-child (and pseudo-parent) as follows.

► **Definition 10.** Given a reachability tree T rooted at s , a vertex x is said to be pseudo-parent of y in T (and y is said to be pseudo-child of x) if x is the nearest ancestor of y in T whose immediate dominator is the same as that of y .

Note that in a 2-vertex strongly connected graph, the definition of pseudo-parent and pseudo-child degenerates to normal notion of parent and child. This is because, immediate dominator of all the vertices (other than s) in such graph is equal to s .

We now state a lemma which is an analogue of Lemma 6 for general graphs.

► **Lemma 11.** *Let a, b, P be as described above, and c be the pseudo-child of b on $\text{PATH}_{T_1}(s, v)$. Then,*

- (i) *Vertex f_2 is an ancestor of c in T_2 .*
- (ii) *None of the internal vertices of P lie on $\text{PATH}_{T_1}(s, c)$ or $\text{PATH}_{T_2}(s, c)$.*

As a corollary of the above lemma we get that P is an $S_{A,B}(c)$ path terminating at pseudo-parent of c in T_1 . We thus define ancestor-detours which are a generalization of parent-detours as follows.

- $\text{AD}_j^i(w)$: a path starting from the highest possible ancestor of w in T_i and terminating at pseudo-parent of w in tree T_j such that none of the internal vertices of the path lie on $\text{PATH}_{T_1}(s, w)$ or $\text{PATH}_{T_2}(s, w)$.

Now let x_1 be the first descendant of f_1 on $\text{PATH}_{T_1}(s, v)$ whose immediate dominator is equal to $\text{idom}(v)$. Similarly, let x_2 be the first descendant of f_2 on $\text{PATH}_{T_1}(s, v)$ whose immediate dominator is equal to $\text{idom}(v)$. Then the ancestor-detours of x_1, x_2 will not be of any help as they would terminate at either f_1, f_2 or their ancestors. Now as in Section 5, we can argue that ancestor-detours of vertices on $\text{PATH}_{T_1}(\bar{x}_1, v) \cup \text{PATH}_{T_2}(\bar{x}_2, v)$ suffice to answer the reachability query for vertex v . This completes the analysis of the third case.

We thus have the following lemma.

► **Lemma 12.** *Let v be a vertex satisfying condition \mathcal{C} such that $f_1 \in \text{PATH}_{T_1}(\overline{\text{idom}(v)}, v)$ and $f_2 \in \text{PATH}_{T_1}(\overline{\text{idom}(v)}, v)$. Also let x_1 and x_2 be as defined above. Then v is reachable from s in $G \setminus \{f_1, f_2\}$ if and only if either of the following statements holds true.*

- (i) [Case 1] *The first vertex of $D^1(v)$ or $D^2(v)$ lies in $S_A(v)$.*
- (ii) [Case 2] *There exists a vertex $w \in S_B(v)$ for which $\text{idom}(w) \in S_A(v) \setminus \{\text{idom}(v)\}$.*
- (iii) [Case 3] *There exists a vertex $w \in \text{PATH}_{T_1}(\bar{x}_1, v)$ such that $\text{idom}(w) = \text{idom}(v)$ and the first vertex of either $\text{AD}_1^1(w)$ or $\text{AD}_1^2(w)$ lies in $S_A(v)$.*
- (iv) [Case 3] *There exists a vertex $w \in \text{PATH}_{T_2}(\bar{x}_2, v)$ such that $\text{idom}(w) = \text{idom}(v)$ and the first vertex of either $\text{AD}_2^1(w)$ or $\text{AD}_2^2(w)$ lies in $S_A(v)$.*

6.1 Implementation of the oracle

We now explain the implementation of reachability oracle for general graphs. As in Section 5, we define the following notations.

- $\alpha^i(v)$: $\text{depth}_{T_i}(\text{idom}(v))$.
- $\beta^i(v)$: depth_{T_i} (first vertex on $D^i(v)$).
- $\gamma_j^i(v)$: depth_{T_i} (first vertex on $\text{AD}_j^i(v)$).

Let f_1, f_2 be a given pair of failed vertices and v be a given query step. We assume that condition \mathcal{C} is satisfied, and failures f_1, f_2 lies respectively on $\text{PATH}_{T_1}(\overline{\text{idom}(v)}, v)$ and $\text{PATH}_{T_2}(\overline{\text{idom}(v)}, v)$. We first check for $i = 1, 2$, if $D^i(v)$ starts from an ancestor of f_i in T_i or not. This is done by comparing the value of $\beta^i(v)$ with the depth of f_i in T_i .

Next we compute the vertices x_1, x_2 as follows. Recall that x_1 is the highest ancestor of v in $\text{PATH}_{T_1}(\bar{f}_1, v)$ whose immediate dominator is equal to $\text{idom}(v)$. So to obtain x_1 , we call the range minima query for vertices on $\text{PATH}_{T_1}(\bar{f}_1, v)$ with $\langle \alpha^1(\cdot), \text{depth}_{T_1}(\cdot) \rangle$ as the weight function. By comparing the value of $\alpha^1(\cdot)$, it is able to filter out those vertices in

130:12 An Optimal Oracle for Dual Fault Tolerant Reachability

$\text{PATH}_{T_1}(\bar{f}_1, v)$ whose immediate dominator is at minimum depth in T_1 , i.e. it is equal to $\text{idom}(v)$. After this it assigns x_1 to be that vertex which has minimum depth in T_1 . Vertex x_2 is computed in a similar manner.

Now notice that to find whether there exists a vertex in $S_B(v)$ whose immediate dominator lies in $S_A(v) \setminus \{\text{idom}(v)\}$, we only need to restrict ourself to paths $\text{PATH}_{T_1}(\bar{f}_1, \bar{x}_1)$ and $\text{PATH}_{T_1}(\bar{f}_1, \bar{x}_1)$. This is because Property 2 implies that immediate dominator of vertices in $\text{PATH}_{T_1}(x_1, v)$ is either equal to $\text{idom}(v)$ or lies in $\text{PATH}_{T_1}(x_1, v)$ itself. Similarly, for $\text{PATH}_{T_2}(x_2, v)$. So for $i = 1, 2$, we perform the range minima query to find a vertex, say w , on $\text{PATH}_{T_i}(\bar{f}_i, \bar{x}_i)$ for which $\alpha^i(w)$ is minimum. If $\alpha^i(w)$ is less than the depth of f_i in T_i , then we report that v is reachable from s .

Finally for $i, j \in \{1, 2\}$, we compute a vertex $w \in \text{PATH}_{T_j}(\bar{x}_j, v)$ for which $\langle \alpha^i(\cdot), \gamma_j^i(\cdot) \rangle$ is minimum. The term $\alpha^i(\cdot)$ is added in front so that we are able to filter out those vertices whose immediate dominator is equal to $\text{idom}(v)$. Now if $\gamma_j^i(w)$ is less than the depth of f_i in T_i , then it implies that $\text{AD}_j^i(w)$ starts from an ancestor of f_i in T_i , so we return True.

If we reach to the end of code, that means we have not been able to find any path for v , so we return False.

Algorithm 2: Oracle for reachability to v in general graphs.

```

1  if  $\beta^1(v) < \text{depth}_{T_1}(f_1)$  or  $\beta^2(v) < \text{depth}_{T_2}(f_2)$  then Return True;
2  ;
3   $x_1 \leftarrow$  a vertex on  $\text{PATH}_{T_1}(\bar{f}_1, v)$  for which  $\langle \alpha^1(\cdot), \text{depth}_{T_1}(\cdot) \rangle$  is minimum;
4   $x_2 \leftarrow$  a vertex on  $\text{PATH}_{T_2}(\bar{f}_2, v)$  for which  $\langle \alpha^2(\cdot), \text{depth}_{T_2}(\cdot) \rangle$  is minimum;
5  foreach  $i \in \{1, 2\}$  do
6  |    $w \leftarrow$  a vertex on  $\text{PATH}_{T_i}(\bar{f}_i, \bar{x}_i)$  for which  $\alpha^i(\cdot)$  is minimum;
7  |   if  $\alpha^i(w) < \text{depth}_{T_i}(f_i)$  then Return True;
8  |   ;
9  end
10 foreach  $i, j \in \{1, 2\}$  do
11 |    $w \leftarrow$  a vertex on  $\text{PATH}_{T_j}(\bar{x}_j, v)$  for which  $\langle \alpha^i(\cdot), \gamma_j^i(\cdot) \rangle$  is minimum;
12 |   if  $\gamma_j^i(w) < \text{depth}_{T_i}(f_i)$  then Return True;
13 |   ;
14 end
15 Return False;
```

As in Algorithm 1, we can argue that the Steps 2, 3, 5 and 9, can be implemented in $O(1)$ time. Thus Algorithm 2 takes constant time to answer reachability queries. We thus conclude with the following theorem.

► **Theorem 13.** *A directed graph $G = (V, E)$ on n vertices can be preprocessed in polynomial time for a given source vertex $s \in V$ to build a data structure of $O(n)$ size such that for any $f_1, f_2, v \in V$, it takes $O(1)$ time to determine if there exists a path from s to v in $G \setminus \{f_1, f_2\}$.*

Acknowledgements. I am very grateful to my advisor, Prof. Surender Baswana, for many helpful discussions and for reviewing this paper.

References

- 1 Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault tolerant subgraph for single source reachability: Generic and optimal. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 19-21, 2016 (to appear)*, 2016.
- 2 Surender Baswana and Neelesh Khanna. Approximate shortest paths avoiding a failed vertex: Near optimal data structures for undirected unweighted graphs. *Algorithmica*, 66(1):18–50, 2013.
- 3 Adam L. Buchsbaum, Loukas Georgiadis, Haim Kaplan, Anne Rogers, Robert Endre Tarjan, and Jeffery Westbrook. Linear-time algorithms for dominators and other path-evaluation problems. *SIAM J. Comput.*, 38(4):1533–1573, 2008.
- 4 Shiri Chechik. Fault-tolerant compact routing schemes for general graphs. *Inf. Comput.*, 222:36–44, 2013.
- 5 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. f -sensitivity distance oracles and routing schemes. *Algorithmica*, 63(4):861–882, 2012.
- 6 Erik D. Demaine, Gad M. Landau, and Oren Weimann. On cartesian trees and range minimum queries. *Algorithmica*, 68(3):610–625, 2014.
- 7 Camil Demetrescu, Mikkel Thorup, Rezaul Alam Chowdhury, and Vijaya Ramachandran. Oracles for distances avoiding a failed node or link. *SIAM J. Comput.*, 37(5):1299–1318, 2008.
- 8 Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 169–178, 2011.
- 9 Ran Duan and Seth Pettie. Dual-failure distance and connectivity oracles. In *SODA'09: Proceedings of 19th Annual ACM – SIAM Symposium on Discrete Algorithms*, pages 506–515, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- 10 Wojciech Fraczak, Loukas Georgiadis, Andrew Miller, and Robert Endre Tarjan. Finding dominators via disjoint set union. *J. Discrete Algorithms*, 23:2–20, 2013.
- 11 Loukas Georgiadis and Robert Endre Tarjan. Dominators, directed bipolar orders, and independent spanning trees. In *Automata, Languages, and Programming – 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*, pages 375–386, 2012.
- 12 Thomas Lengauer and Robert Endre Tarjan. A fast algorithm for finding dominators in a flowgraph. *ACM Trans. Program. Lang. Syst.*, 1(1):121–141, 1979.
- 13 Merav Parter. Dual failure resilient BFS structure. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21-23, 2015*, pages 481–490, 2015.
- 14 Merav Parter and David Peleg. Sparse fault-tolerant BFS trees. In *Algorithms – ESA 2013 – 21st Annual European Symposium, Sophia Antipolis, France, September 2-4, 2013. Proceedings*, pages 779–790, 2013.
- 15 Merav Parter and David Peleg. Fault tolerant approximate BFS structures. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1073–1092, 2014.
- 16 Maxim Teslenko and Elena Dubrova. An efficient algorithm for finding double-vertex dominators in circuit graphs. In *2005 Design, Automation and Test in Europe Conference and Exposition (DATE 2005), 7-11 March 2005, Munich, Germany*, pages 406–411, 2005.
- 17 Virginia Vassilevska Williams. Faster replacement paths. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1337–1346, 2011.