Optimal Parallel Shortest Paths in Small Treewidth Digraphs ⋆ (to appear in ESA'95)

Shiva Chaudhuri and Christos D. Zaroliagis

Max-Planck-Institut für Informatik, Im Stadtwald, D-66123 Saarbrücken, Germany
E-mail: {shiva, zaro}@mpi-sb.mpg.de

Abstract. We consider the problem of preprocessing an n-vertex digraph with real edge weights so that subsequent queries for the shortest path or distance between any two vertices can be efficiently answered. We give parallel algorithms for the EREW PRAM model of computation that depend on the *treewidth* of the input graph. When the treewidth is a constant, our algorithms can answer distance queries in $O(\alpha(n))$ time using a single processor, after a preprocessing of $O(\log^2 n)$ time and $O(n)$ work, where $\alpha(n)$ is the inverse of Ackermann's function. The class of constant treewidth graphs contains outerplanar graphs and series-parallel graphs, among others. To the best of our knowledge, these are the first parallel algorithms which achieve these bounds for any class of graphs except trees. We also give a dynamic algorithm which, after a change in an edge weight, updates our data structures in $O(\log n)$ time using $O(n^\beta)$ work, for any constant $0 < \beta < 1$. Moreover, we give an algorithm of independent interest: computing a shortest path tree, or finding a negative cycle in $O(\log^2 n)$ time using $O(n)$ work.

1 Introduction

Finding shortest paths in digraphs is a fundamental problem in network optimization [3]. Given an n-vertex digraph G with real edge weights, the shortest paths problem asks for paths of minimum weight between vertices in G. In the single-source problem we seek such paths from a specific vertex to all other vertices and in the all-pairs shortest paths (apsp) problem we seek such paths between every pair [18].

For general digraphs the best parallel algorithm for the apsp problem takes $O(\log^2 n)$ time using $O(n^3)$ work on an EREW PRAM [12]. In the case of planar digraphs there is an $O(\log^4 n)$-time, $O(n^2)$-work EREW PRAM algorithm [10]. An apsp algorithm must output paths between $\Omega(n^2)$ vertex pairs and thus requires this much work and space. For sparse digraphs a more efficient approach is to preprocess the digraph so that subsequently, *queries* can be efficiently answered. A query specifies two vertices and a *shortest path query* asks for a minimum weight path between them, while a *distance query* only asks for

⋆ This work was partially supported by the EU ESPRIT Basic Research Action No. 7141 (ALCOM II).

the weight of such a path. For example, for *outerplanar* digraphs, it was shown in [11] that after preprocessing requiring $O(\log n)$ time and $O(n \log n)$ work on a CREW PRAM, a distance query is answered in $O(\log n)$ time using a single processor and a shortest path query in $O(\log n)$ time using $O(L + \log n)$ work (where L is the number of edges of the reported path). In [11] it is also shown how distance queries in planar digraphs can be answered in $O(\log n + \log^2 q)$ time using $O(\log n + q)$ work, after polylog-time and $O(n \log n \log^* n + q^{1.5})$-work preprocessing on a CREW PRAM. These latter bounds are given in terms of a minimum number of faces q that collectively cover all vertices of the planar digraph. Note that q varies from 1 (outerplanar digraph) up to $\Theta(n)$.

The study of graphs using the *treewidth* as a parameter was pioneered by Robertson and Seymour [16, 17] and continued by many others (see e.g. [5, 7]). Informally, the treewidth is a measure of how close the structure of the graph is to a tree (see Section 2 for a formal definition). Graphs of treewidth t are also known as partial t-trees. These graphs have at most tn edges. Classifying graphs based on treewidth is useful because diverse properties of graphs can be captured by a single parameter. For instance, the class of graphs of bounded treewidth includes outerplanar graphs, series-parallel graphs, graphs with bounded bandwidth and cutwidth and many other classes [5, 7]. Thus, giving efficient algorithms parameterized by treewidth is an important step in the development of better algorithms for many natural classes of sparse graphs.

In this paper we consider the problem of preprocessing a digraph of small treewidth in parallel, so that afterwards, queries can be efficiently answered using a single processor. We also consider the dynamic version of the problem, where edge weights may change. In [9] sequential algorithms are given that, for digraphs of constant treewidth, after $O(n)$ time preprocessing answer a distance (resp. shortest path) query in $O(\alpha(n))$ (resp. $O(L\alpha(n))$) time[1]. After a change in an edge weight, the algorithm updates the data structure in $O(n^\beta)$ time, for any constant $0 < \beta < 1$. The main contribution of this paper is an algorithm that achieves optimal parallelization, on the EREW PRAM, of the above results. For digraphs of constant treewidth, after $O(\log^2 n)$ time and $O(n)$ work preprocessing our algorithm answers a distance (resp. shortest path) query in $O(\alpha(n))$ (resp. $O(L\alpha(n))$) time using a single processor. Updates can be performed in $O(\log n)$ time using $O(n^\beta)$ work for any constant $0 < \beta < 1$. This improves all previous parallel results for this class of graphs. Moreover, it improves the results in [11] for outerplanar digraphs in many ways: it improves the preprocessing and distance query bounds, it runs on the weakest PRAM model and it applies to a larger class of graphs. We note that the time bottleneck in preprocessing is the computation of the *tree-decomposition* (see Section 2) of the input graph. If an explicit tree-decomposition of the graph is also provided with the input, then the preprocessing time is $O(\log n)$.

As in [9], we give a tradeoff between the preprocessing work and the query time. For bounded treewidth digraphs, after $O(nI_k(n))$ preprocessing, we can

[1] $\alpha(n)$ is the inverse of Ackermann's function [1] and is a very slowly growing function.

answer queries in $O(k)$ time, for $1 \leq k \leq \alpha(n)$. $I_k(n)$ is a function that decreases rapidly with k (see Section 3). In particular $I_1(n) = \lceil \log n \rceil$ and $I_2(n) = \log^* n$.

A solution to the single-source problem consists of a *shortest path tree* rooted at a given vertex. A shortest path tree exists iff there is no negative weight cycle in the graph. In parallel computation, the best algorithm for constructing a shortest path tree (or finding a negative cycle) in a general digraph G takes as much time as computing apsp in G [12]. Some improvements have been made for outerplanar [11] and planar digraphs [10] with no negative cycles. In those papers, a shortest path tree can be computed in $O(\log^2 n)$ time, after a preprocessing of the input digraph. The preprocessing work of [10] is $O(n^{1.5})$ on an EREW PRAM, while the preprocessing work in [11] is $O(n \log n)$ on a CREW PRAM. Even with randomization allowed, and the weights restricted to being positive integers, for planar digraphs, the best polylog-time algorithm uses n processors (and hence $\Omega(n \log n)$ work) on an EREW PRAM. Although, on a CREW PRAM, a negative cycle in an outerplanar digraph can be found in $O(\log n \log^* n)$ time and $O(n)$ work, this algorithm does not construct the shortest path tree [14]. Hence, the work for finding a shortest path tree in polylog-time is $\Omega(n \log n)$, even for the case of outerplanar digraphs.

We give an algorithm to construct a shortest path tree (or find a negative cycle) in digraphs of constant treewidth that runs on an EREW PRAM in $O(\log^2 n)$ time using $O(n)$ work (Section 3). If a tree-decomposition is also provided with the input, then the algorithm runs in $O(\log n)$ time. To the best of our knowledge, this is the first deterministic parallel algorithm for the shortest path tree problem that achieves $O(n)$ work.

Our algorithms start by computing a tree-decomposition of the input digraph G. The tree decomposition of a graph with constant treewidth can be computed in $O(\log^2 n)$ time using $O(n)$ work on an EREW PRAM [8]. The main idea behind our algorithms is the following. We define a certain value for each node of the tree-decomposition of G, and an associative operator on these values. We then show that the shortest path problem reduces to computing the product of these values along paths in the tree-decomposition. (A similar idea was used in [2], to show that computing shortest paths reduces to computing the product of certain elements in a closed semiring.) Parallel algorithms to compute the product of node values along paths in a tree are given in [4]. Our preprocessing vs. query bound trade-off arises from a similar trade-off in [4]. The dynamization of our data structures is based on the above ideas and on a graph equipartitioning result which is of independent interest. We note that a similar approach is used in [9], however, the parallel algorithms presented in this paper require substantially different techniques.

2 Preliminaries

In this paper, we will be concerned with finding shortest paths or distances between vertices of a directed graph. Thus, we assume that we are given an *n-vertex weighted digraph* G, i.e. a digraph $G = (V(G), E(G))$ and a weight

function $wt : E(G) \longrightarrow \mathbb{R}$. We call $wt(u, v)$ the *weight* of the edge $\langle u, v\rangle$. The weight of a path in G is the sum of the weights of the edges on the path. For $u, v \in V(G)$, a *shortest path* in G from u to v is a path whose weight is minimum among all paths from u to v. The *distance* from u to v, written as $\delta(u, v)$ or $\delta_G(u, v)$, is the weight of a shortest path from u to v in G. A cycle in G is a (simple) path starting and ending at the same vertex. If the weight of a cycle in G is less than zero, then we will say that G contains a *negative cycle*. It is well-known [18] that shortest paths exist in G, iff G does not contain a negative cycle.

For a subgraph H of G, and vertices $x, y \in V(H)$, we shall denote by $\delta_H(x, y)$ the distance of a shortest path from x to y in H. A *shortest path tree* rooted at $v \in V(G)$, is a spanning tree such that $\forall w \in V(G)$, the tree path from v to w is a shortest path in G from v to w.

Let G be a (directed or undirected) graph and let $W \subseteq V(G)$. Then by $G[W]$ we shall denote the subgraph of G induced by W. Let V_1, V_2 and S be disjoint subsets of $V(G)$. We say that S is a *separator for* V_1 *and* V_2, or that S *separates* V_1 *from* V_2, iff every path from a vertex in V_1 (resp. V_2) to a vertex in V_2 (resp. V_1) passes through a vertex in S. Let H be a subgraph of G. A *cut-set for* H is a set of vertices $C(H) \subseteq V(H)$, whose removal separates H from the rest of the graph.

Often, we will want to focus on a subgraph induced by a subset of the vertices of a graph, however, we would like the distances between vertices in this subgraph to be the same as in the original graph. Let H be a digraph, with V_1, V_2 and U a partition of $V(H)$ such that U is a separator for V_1 and V_2. Let H_1 and H_2 be subgraphs of H such that $V(H_1) = V_1 \cup U$, $V(H_2) = V_2 \cup U$ and $E(H_1) \cup E(H_2) = E(H)$. We say that H_1' is a graph obtained by *absorbing* H_2 *into* H_1, if H_1' is obtained from H_1 by adding edges $\langle u, v\rangle$, with weight $\delta_{H_2}(u, v)$ or $\delta_H(u, v)$, for each pair $u, v \in U$. (In case of multiple edges, retain the one with minimum weight.) The following lemma, proved in [9], shows that absorbing a subgraph into another preserves distances.

Lemma 1. *Let H_1 and H_2 be subgraphs of H and let H_1' be obtained by absorbing H_2 into H_1. Then, for all $x, y \in V(H_1')$, $\delta_{H_1'}(x, y) = \delta_H(x, y)$.*

A *tree-decomposition* of a (directed or undirected) graph G is a pair (X, T) where $T = (V(T), E(T))$ is a tree and X is a family $\{X_i | i \in V(T)\}$ of subsets of $V(G)$, such that $\cup_{i \in V(T)} X_i = V(G)$ and also the following conditions hold:

- (*edge mapping*) $\forall (v, w) \in E(G)$, there exists an $i \in V(T)$ with $v \in X_i$ and $w \in X_i$.
- (*continuity*) $\forall i, j, k \in V(T)$, if j lies on the path from i to k in T, then $X_i \cap X_k \subseteq X_j$, or equivalently: $\forall v \in V(G)$, the nodes $\{i \in V(T) | v \in X_i\}$ induce a connected subtree of T.

The *treewidth* of a tree-decomposition is $\max_{i \in V(T)} |X_i| - 1$. The treewidth of G is the minimum treewidth over all possible tree-decompositions of G.

Fact 1 *[8] Given a constant $t \in \mathbb{N}$ and an n-vertex graph G, there exists an EREW PRAM algorithm, running in $O(\log^2 n)$ time using $O(n)$ work, which tests whether G has treewidth at most t and if so, outputs a tree-decomposition (X, T) of G with treewidth at most t.*

Fact 2 *[6, 8] Given a constant $t \in \mathbb{N}$ and a tree-decomposition of treewidth t of an n-vertex graph G, we can compute a rooted, binary tree-decomposition of G with depth $O(\log n)$ and treewidth at most $3t + 2$, in $O(\log n)$ time using $O(n)$ work on an EREW PRAM.*

We shall call the tree-decomposition found in Fact 2 *balanced*. Given a tree-decomposition of G, we can easily find separators in G, as the following proposition shows.

Proposition 1. *[17] Let G be a graph, (X, T), its tree-decomposition, $e = (i, j) \in E(T)$ and T_1 and T_2 the two subtrees obtained by removing e from T. Then $X_i \cap X_j$ separates $\cup_{m \in V(T_1)} X_m$ from $\cup_{m \in V(T_2)} X_m$.*

3 The Static Data Structures

For a function f let $f^{(1)}(n) = f(n)$; $f^{(i)}(n) = f(f^{(i-1)}(n))$, $i > 1$. Define $I_0(n) = \lceil \frac{n}{2} \rceil$ and $I_k(n) = \min\{j \mid I_{k-1}^{(j)}(n) \leq 1\}$, $k \geq 1$. The functions $I_k(n)$ decrease rapidly as k increases, in particular, I_1 behaves like $\log n$ and I_2 like $\log^* n$. Define $\alpha(n) = \min\{j \mid I_j(n) \leq 1\}$. The following was proved in [4].

Fact 3 *[4] Let $\bullet$ be an associative operator defined on a set S, such that for $x, y \in S$, $x \bullet y$ can be computed in $O(m)$ time and $O(w)$ work. Let T be a tree with n nodes such that each node is labelled with an element from S. Then: (i) for each $k \geq 1$, after $O(m \log n)$-time and $O(wnI_k(n))$-work preprocessing on an EREW PRAM, the composition of labels along any path in the tree can be computed in $O(wk)$ time by a single processor; and (ii) after $O(m \log n)$-time and $O(wn)$-work preprocessing on an EREW PRAM, the composition of labels along any path in the tree can be computed in $O(w\alpha(n))$ time by a single processor.*

The main idea of our algorithm is, as in [9], to reduce shortest path computations to the above problem. We define a certain value for each node of the tree-decomposition of G, as well as an associative operator on these values. We then show that shortest path computation reduces to computing products of those values along paths in the tree-decomposition. Then the rest follows by the above Fact.

Call a tuple (a, b, c) a *distance tuple* if a, b are arbitrary symbols and $c \in \mathbb{R}$. For two distance tuples, $(a_1, b_1, c_1), (a_2, b_2, c_2)$, define their product $(a_1, b_1, c_1) \otimes (a_2, b_2, c_2) = (a_1, b_2, c_1 + c_2)$ if $b_1 = a_2$ and as nonexistent otherwise.

For a set of distance tuples, M, define minmap(M) to be the set $\{(a, b, c) : (a, b, c) \in M$ and $\forall (a', b', c') \in M$ if $a' = a, b' = b$, then $c \leq c'\}$, i.e. among all tuples with the same first and second components, minmap retains only the tuples with the smallest third component.

Let M_1 and M_2 be sets of distance tuples. Define the operator $\circ$ by $M_1 \circ M_2 = \text{minmap}(M)$, where $M = \{x \otimes y : \ x \in M_1, \ y \in M_2\}$. It is not hard to show that $\circ$ is an associative operator.

Let G be a digraph with real edge weights. Note that in the above definition, if M_1 and M_2 have tuples of the form (a, b, x) where $a, b \in V(G)$ and x is the weight of a path from a to b, then $M_1 \circ M_2$ computes tuples (a, b, y) where y is the (shortest) distance from a to b using only the paths represented in M_1 and M_2.

For $X, Y \subseteq V(G)$, not necessarily distinct, define $P(X, Y) = \{(a, b, \delta_G(a, b)) : a \in X, b \in Y\}$. We will write $S(X)$ for $P(X, X)$ (by definition, $S(X)$ contains tuples $(x, x, 0), \ \forall x \in X$).

The following lemma, proved in [9], establishes the desired connection between computing shortest paths and products along tree paths of the operator $\circ$ defined above.

Lemma 2. *Let G be a weighted digraph and (X, T) its tree decomposition. For $i \in V(T)$, define $\gamma(i) = S(X_i)$. Let $v_1, \ldots, v_p$ be a path in T. Then $\gamma(v_1) \circ \ldots \circ \gamma(v_p) = P(X_{v_1}, X_{v_p})$.*

Therefore, it only remains to show how the γ values can be efficiently computed in parallel for each node of a tree-decomposition. This is shown in the next lemma. The following algorithm first converts the given tree-decomposition into a balanced one, and then repeatedly shrinks the tree, by absorbing the subgraphs corresponding to leaves. When the tree is reduced to a single node, the algorithm computes γ using brute force, for this node. The distances computed are the distances in the original graph, since distances are preserved during absorption. Then, it reverses the shrinking process and expands the tree, using the γ values already computed to compute γ values for the newly expanded nodes.

Lemma 3. *Let G be an n-vertex weighted digraph and let (X, T) be the tree-decomposition of G, of treewidth t. For each pair u, v such that $u, v \in X_i$ for some $i \in V(T)$, let $Dist(u, v) = \delta(u, v)$ and $Int(u, v) = x$ where x is some intermediate node (neither u nor v) on a shortest path from u to v. (If $wt(u, v) = \delta(u, v)$, then $Int(u, v) = null$.) Then in $O(\log n \log^2 t)$ time using $O(t^3 n)$ work on an EREW PRAM, we can either find a negative cycle in G, or compute the values $Dist(u, v)$ and $Int(u, v)$ for each such pair u, v.*

Proof. Initially all the values $Dist(u, v)$ are set to ∞ and $Int(u, v)$ to null. We give an inductive algorithm. First, convert (X, T) into a balanced tree-decomposition of G using Fact 2. Then, for each vertex of T, we compute its *level number*, which is one more than the level of its parent, with the root having level number 1. This computation can be done in $O(\log n)$ time and $O(n)$ work ([13], Theorem 3.4).

We use induction on the number of levels of T. Let d be the depth of T and N_d be the set of tree nodes at level d. For all nodes $z \in N_d$, we run the algorithm of [12] to solve the apsp problem in $G[X_z]$. This will take $O(\log^2 t)$ time and $O(|N_d| t^3)$ work. If there is a negative cycle in some $G[X_z]$, it will be

found by the algorithm of [12]; so henceforth assume that there is no $G[X_z]$ containing a negative cycle. Update the values for pairs $u, v \in X_z$ as follows: if the weight of the shortest path found is less than the current value of $Dist(u, v)$, then set $Dist(u, v)$ to the new value and $Int(u, v)$ to any intermediate vertex on the shortest path found. If $wt(u, v)$ is equal to the weight of the shortest path found, then set $Int(u, v) =$ null.

If $d = 1$ (which implies that $|V(T)| = 1$), we are done. Otherwise remove all nodes $z \in N_d$ from T and call the resulting tree T'. Let $V' = \cup_{i \in V(T')} X_i$ and construct G' by absorbing every $G[X_z]$ into $G[V']$, where the weight of each added edge $\langle u, v \rangle$ is $\delta_{G[X_z]}(u, v)$. Then, for any vertices $u, v \in V'$, $\delta_{G'}(u, v) = \delta_G(u, v)$, by Lemma 1. In particular, if G contains a negative cycle, so does G'. Note that $(X - Y, T')$ is a tree-decomposition for G', where $Y = \cup_{z \in N_d} X_z$.

Inductively run the algorithm on G'. If a negative cycle is found in G', then a negative cycle in G can be found by replacing any edges added during the absorption by their corresponding paths in the subgraphs $G[X_z]$. Hence, we may assume that G' does not contain a negative cycle.

For $a, b \in V'$, $Dist(a, b) = \delta_{G'}(a, b) = \delta_G(a, b)$, as desired. If $Int(a, b) = x \neq$ null, then x is an intermediate vertex on a shortest a to b path in G' and hence also in G, as desired. If $Int(a, b) =$ null, then $\langle a, b \rangle$ is a shortest path in G'. If $wt(a, b) > Dist(a, b)$, then this edge must have been added during the absorption. Correct the value $Int(a, b)$ by setting it to some intermediate vertex on the corresponding a to b shortest path found in $G[X_z]$. After this, all Int values are correct for $a, b \in V'$.

Construct a digraph G'' by absorbing $G[V']$ into every $G[X_z]$, with each added edge $\langle u, v \rangle$ having weight $\delta_G(u, v)$. By Lemma 1, $\delta_{G''}(x, y) = \delta_G(x, y)$, $\forall x, y \in X_z$. Run the algorithm of [12] on G'' to recompute all pairs shortest paths. Update the values $Dist(a, b)$ and $Int(a, b)$ for $a, b \in X_z$ as before. Now for each $z \in N_d$ we have: For $a, b \in X_z$, $Dist(a, b) = \delta_{G''}(a, b) = \delta_G(a, b)$ as desired. For $a, b \in V' \cap X_z$, $Int(a, b)$ is not changed since $Dist(a, b)$ is already $\delta_G(a, b)$. If either a or b does not belong to $V' \cap X_z$, $Int(a, b) =$ an intermediate vertex on a shortest path in G'' and hence in G, or $Int(a, b) =$ null in which case $wt(a, b) = \delta_G(a, b)$. Thus, the values computed are correct for all pairs a, b which completes the induction.

Concerning the resource bounds, it suffices to notice that the algorithm performs a bottom-up and a top-down traversal of T by processing the tree level-by-level and visiting every tree node at most twice. At each level, the algorithm takes $O(\log^2 t)$ time using $O(t^3)$ work per node. Hence, in total it takes $O(\log n \log^2 t)$ time and $O(t^3 n)$ work on an EREW PRAM. □

We are now ready for our main theorem.

Theorem 1. *For any integer t and any $k \geq 1$, let G be an n-vertex weighted digraph of treewidth at most t, whose tree-decomposition can be found in $T(n, t)$ parallel time using $W(n, t)$ work on an EREW PRAM. Then, the following hold on an EREW PRAM: (i) After $O(\log n \log^2 t + T(n, t))$ time and $O(t^3 n I_k(n) + W(n, t))$ work and space preprocessing, distance queries in G can be answered*

in $O(t^3k)$ time using a single processor. (ii) After $O(\log n \log^2 t + T(n,t))$ time and $O(t^3 n + W(n,t))$ work and space preprocessing, distance queries in G can be answered in $O(t^3\alpha(n))$ time using a single processor.

Proof. Using Fact 1, we first compute a tree-decomposition (X,T) of G. By Lemma 3, we either find a negative cycle (and in such a case we stop), or compute values $Dist(u,v)$ for u,v such that $u,v \in X_i$ for some $i \in V(T)$. From these values, we can easily compute $\gamma(i)$, $\forall i \in V(T)$. By Fact 3 we preprocess T so that product queries on γ can be answered. Given a query, $u,v \in V(G)$, let i,j be vertices of T such that $u \in X_i$ and $v \in X_j$. We ask for the product of the γ values on the path between i and j. By Lemma 2, the answer to this query contains the information about $\delta(u,v)$. The bounds follow easily by the ones given in Fact 3 and by the fact that the composition of any two γ values can be computed in $O(\log^2 t)$ time using $O(t^3)$ work. □

As in [9], a distance query time of Q yields a path query time of $O(t^3LQ)$, where L is the length of the reported path. (We omit the details for lack of space.) Hence, we can summarize all the preceding results as follows.

Theorem 2. *Let G be an n-vertex weighted digraph of constant treewidth and let $k \geq 1$ be any constant integer. Then, the following hold on an EREW PRAM: (i) After $O(\log^2 n)$ time and $O(nI_k(n))$ work and space preprocessing, distance (resp. shortest path) queries in G can be answered in $O(k)$ (resp. $O(kL)$) time using a single processor (where L is the length of the reported path). (ii) After $O(\log^2 n)$ time and $O(n)$ work and space preprocessing, distance (resp. shortest path) queries in G can be answered in $O(\alpha(n))$ (resp. $O(L\alpha(n))$) time using a single processor (where L is the length of the reported path).*

For the single source problem, we have the following theorem.

Theorem 3. *Let G be an n-vertex weighted digraph of constant treewidth and let $s \in V(G)$. Then, in $O(\log^2 n)$ time using $O(n)$ work on an EREW PRAM, we can either compute a shortest path tree rooted at s, or find a negative cycle in G (if exists). If the tree-decomposition of G is also provided with the input, then the computation takes $O(\log n)$ time.*

Proof. We first compute a tree-decomposition (X,T) of G of treewidth t, using Fact 1. Using Lemma 3, we either compute $Dist(u,v)$, for u,v such that $u,v \in X_i$, for some $i \in V(T)$, or find a negative cycle in G. If there is no negative cycle, we can easily compute $\gamma(i)$, $\forall i \in V(T)$. Let $i \in V(T)$ such that $s \in X_i$. Root T at node i and make it balanced, using Fact 2. Starting at i, perform a top-down traversal of T by visiting all nodes of T level-by-level. (This can be done in parallel as follows: each processor associated with some node of T forks two other processors and associates them with the children of its node.) At each node $j \in V(T)$ store the product of the γ values on the path from i to j. Since the composition of two γ values can be computed in $O(\log^2 t)$ time using $O(t^3)$ work on an EREW PRAM and each node is visited exactly once, the whole process takes $O(\log n \log^2 t)$ time using $O(t^3 n)$ work.

Let $y \in V(G)$ and let $j \in V(T)$ such that $y \in X_j$. By Lemma 2, the value stored at vertex j during the above mentioned top-down traversal, is $P(X_i, X_j)$ which contains the tuple $(s, y, \delta(s, y))$. Thus, we may assume that for each $y \in V(G)$, we have the value $\delta(s, y)$. For an edge $\langle v, u \rangle$ of G, define $h(v, u)$ to be the node z of T such that $v, u \in X_z$ and z is the closest such node to the root. (By the continuity condition, $h(v, u)$ is unique.) It is not hard to see that during the previous top-down traversal of T, we can found such nodes $h(v, u)$ for each edge $\langle v, u \rangle$ in G.

To construct the shortest path tree ST, we do the following. Starting at the root node i, we perform a second, level-by-level, top-down traversal of T. For a node $j \in V(T)$ at level $\ell \geq 1$, we check (sequentially) edges $\langle v, u \rangle$, where $v, u \in X_{h(v,u)}$ and v belongs to the shortest path tree constructed so far. (Initially, $j = i$ and $v = s$.) If $\delta(s, u) = \delta(s, v) + wt(v, u)$, then make v the parent of u in ST. If v, u belong also to any child of $X_{h(v,u)}$, then mark the edge $\langle v, u \rangle$ as being "examined" in the local memory of the processor associated with this child. Note that this last operation is needed in order to avoid concurrent access conflicts in the shared memory, in the case where there is another node $k \in V(T)$ at the same level with j for which $v, u \in X_k$.

A simple induction argument shows that the above procedure creates a shortest path tree rooted at s. It is again easy to see that each tree node is visited exactly once and that we need $O(t)$ time (using a single processor) in such a node. Hence, in total, ST can be constructed in $O(t \log n)$ time using $O(t^3 n)$ work. Since t is constant, the bounds claimed in the theorem follow. □

4 The Dynamic Algorithm

In this section we shall give our dynamic data structures and algorithms. The approach follows the one in [9], but the parallel implementation is rather different. We divide the digraph into subgraphs with disjoint edge sets and small cut-sets, and construct another (smaller) digraph – the reduced digraph – by absorbing each subgraph. The sizes of the subgraphs are chosen so that the subgraphs and the reduced digraph have size $O(\sqrt{n})$. We then construct a query data structure for each subgraph and for the reduced digraph. Queries can be efficiently answered by querying these data structures. Since the edge sets are disjoint, a change in the weight of an edge affects the data structure for only one subgraph. Then we update the data structure of this subgraph. This may result in new distances between vertices in its cut-set, which appear in the reduced digraph as changes in the weights of edges between these cut-set vertices. Since the cut-set is small, the weights of only a few edges in the reduced digraph change. The data structure for the reduced digraph is updated to reflect these changes. Thus an update in the original digraph is accomplished by a constant number of updates in subgraphs of size $O(\sqrt{n})$, which yields $O(\sqrt{n})$ update work. By recursively applying this idea, we get an update work of $O(n^\beta)$, for any constant $0 < \beta < 1$. In the following, we first give the graph partitioning results and then give the details of our algorithms.

4.1 Graph Equipartitions

Lemma 4. *Let T be a rooted binary tree on n nodes and let $1 \leq m \leq n$. Then, in $O(\log n)$ time and $O(n)$ work on an EREW PRAM, we can partition the nodes of T into at least n/m and at most $8n/m$ groups such that each group: (i) is a connected subtree; (ii) is connected to the rest of the tree through at most 3 edges; and (iii) has at most m nodes.*

Proof. The algorithm is a variant of the usual parallel tree contraction algorithm [13]. By adding a leaf as a child to each node that has one child, we obtain a tree in which each node is a leaf or has two children. Assign a weight of 1 to each node in the tree. Number the leaves of the tree from left to right using the Euler tour technique [13]. From now on assume that we have a tree with weights on the nodes adding up to n, in which each internal node has two children, and in which some of the leaves are numbered from left to right.

In parallel, for each odd numbered leaf that is a left child, if the sum of the weights of the leaf, its parent and its sibling is at most m, then shrink the edges connecting the leaf and its sibling to their parent. Assign the parent a weight equal to the sum of the weights of the three nodes. If the sibling is a leaf, it is even numbered. Assign this number to the parent (which is now a leaf in the modified tree). If the sum of the weights exceeds m, then delete the numbers (if they exist) from the leaf and the sibling.

Repeat the above for each odd numbered leaf that is a right child. After these two steps, all the numbered leaves in the tree have an even number. Divide each of these numbers by 2.

We repeat these two steps $\log n$ times. It is not hard to see that after the ith iteration, at most $l/2^i$ leaves have numbers, where l is the initial number of leaves. Thus, at the end, there are no numbered leaves. Throughout, the following invariant is maintained: if a leaf does not have a number, then the weights of the leaf, its parent and sibling add up to more than m. Call such a triple of leaf, parent and sibling an *overweight* group.

Each non-numbered leaf is contained in some overweight group, and no node can belong to more than two overweight groups. Thus, the sum of the weights of all the overweight groups is at most $2n$, hence the number of overweight groups is at most $2n/m$. Since each overweight group contains at most two non-numbered nodes, the total number of non-numbered leaves at the end is $4n/m$. Since each internal node has two children, the total number of nodes remaining in the tree is at most $8n/m$.

Each node v in the remaining tree is associated with the connected subtree induced by the nodes that were shrunk into v in the above process. These are the required groups. It is easy to see that v has a weight equal to the number of nodes in the associated subtree. Since this weight is at most m, there are at least n/m such connected subtrees. Also, as shown above, there are no more than $8n/m$ connected subtrees. It follows from the construction that each subtree is connected to the rest of the tree through at most 3 edges. □

Input-Output Conventions: We assume that the above algorithm has its input tree specified as a linked structure in n contiguous memory cells. The output it produces is in $O(n)$ contiguous memory cells, divided into contiguous blocks, each block containing one of the connected components in the same linked format, and one final block containing the compressed tree (i.e. the tree at the end of the shrinking process) in a linked format. This can be accomplished using standard EREW PRAM methods using $O(\log n)$ time and $O(n)$ work, which we now describe briefly.

By assigning the preorder number to each node in the compressed tree, we can assign a unique number between 1 and q (where q is the number of nodes in the compressed tree) to each connected subtree. Then, by solving a prefix summation problem on q elements, where the ith element is the number of nodes in subtree i, we can allocate contiguous memory blocks for the various subtrees. It remains to copy the subtrees into the appropriate blocks.

Since each node in the compressed tree knows the memory addresses allocated for its subtree, reversing the shrinking process, we can assign a unique memory address in the appropriate block to each node in a subtree. Now it is a simple matter for each node to copy itself into this address, and duplicate its link structure.

Definition 1. *Let λ, δ be positive integer constants and let $1 \leq m \leq n$. Then, given an n-vertex digraph G as well as its balanced tree-decomposition of treewidth t, we define an* (λ, δ, m)-equipartition of G *to be a partition of G into q subgraphs $H_1, \ldots, H_q$, where $n/m \leq q \leq \delta n/m$, along with the construction of another subgraph H' such that: (i) H_i has at most tm vertices and a cut-set $C(H_i)$ of size at most λt; (ii) H' is the induced subgraph on vertices $\cup_{i=1}^{q} C(H_i)$, augmented with edges $\langle x, y \rangle$, $x, y \in C(H_i)$ for each $1 \leq i \leq q$; and (iii) we have a tree-decomposition of treewidth t for each H_i and a tree-decomposition for H' of treewidth $3t$.*

The following lemma shows that an $(3, 8, m)$-equipartition can be efficiently computed.

Lemma 5. *Given an n-vertex digraph G along with its balanced tree-decomposition of treewidth t, we can compute an $(3, 8, m)$-equipartition of G in $O(\log n)$ time using $O(t^2 n)$ work on an EREW PRAM, where $1 \leq m \leq n$.*

Proof. Let (X, T) be the balanced tree decomposition of G. Then, by Fact 2, T has at most $2n$ nodes. Partition the nodes of T into $n/m \leq q \leq 8n/m$ connected components using Lemma 4. For each component T_i, $1 \leq i \leq q$, create a subgraph H_i, that is the induced subgraph of G on the vertices in $\cup_{v \in V(T_i)} X_v$. Note that T_i is a tree decomposition of H_i. The number of vertices in H_i is at most $t|V(T_i)| = tm$. Let v_1, v_2 and v_3 be the nodes through which T_i is connected to the other components. Then, $C(H_i) = X_{v_1} \cup X_{v_2} \cup X_{v_3}$, and $C(H_i)$ has at most $3t$ vertices. H' is constructed by constructing a clique on $C(H_i)$ for each $1 \leq i \leq q$. The tree decomposition for H' is constructed by shrinking each component T_i into a single node u and assigning $X_u = C(H_i)$. It

is easily verified that this is a tree decomposition of H' of width $3t$. Also, it is not hard to see that the work required for the above constructions is bounded by $O(t^2 n)$ and the time by $O(\log(tn))$. The EREW PRAM implementation can be easily done using the data structures described after the proof of Lemma 4. □

4.2 Data Structures and Algorithms

Let $PD(G, \{P_W, P_T\}, \{U_W, U_T\}, Q)$ be a parallel dynamic data stucture for a digraph G, where $O(P_W)$ (resp. $O(P_T)$) is the preprocessing work and space (resp. time) to be set up, $O(Q)$ is the time to answer a distance query using a single processor and $O(U_W)$ (resp. $O(U_T)$) is the work (resp. time) to update it after the modification of an edge-weight.

Theorem 4. *Assume that we are given an n-vertex weighted digraph G and its balanced tree decomposition of treewidth t. Then, for $r > 0$, we can construct, on an EREW PRAM, the following (with $A = 5t^3 11^{3r}$):*
(i) $PD(G, \{A^r n, 2^r \log n\}, \{A^r n^{(1/2)^{r-1}}, A^r \log n\}, A^r \alpha(n))$; and
(ii) $PD(G, \{A^r n I_k(n), 2^r \log n\}, \{A^r n^{(1/2)^{r-1}}, A^r \log n\}, A^r k)$, for $k \geq 1$.

Proof. We shall prove part (i). Part (ii) can be proved similarly. We use induction on r. If $r = 1$, then, the work and time allowed for updates exceeds the preprocessing, and the static data structure of Theorem 1 suffices, with updates implemented by simply recomputing the whole data structure. We use the notation $D(G, n, r, t)$ for $PD(G, \{A^r n, 2^r \log n\}, \{A^r n^{(1/2)^{r-1}}, A^r \log n\}, A^r \alpha(n))$.

Assume the theorem holds for $r' < r$. We show how to construct $D(G, n, r, t)$.

We first construct an $(3, 8, \sqrt{n})$-equipartition of G using Lemma 5, yielding H' and $H_1, \ldots, H_q$, $\sqrt{n} \leq q \leq 8\sqrt{n}$.

Define G_i to be H_i with all edges joining pairs of vertices in its cut-set deleted. Define G' to be H' with edges $\langle x, y \rangle$ weighted $\delta_{G_i}(x, y)$ for each pair $x, y \in C(G_i)$, $1 \leq i \leq q$. Replace multiple edges by the edge of minimum weight. Note that G' is exactly the graph obtained by absorbing $G_1, G_2, \ldots, G_q$ into the rest of the graph. By Lemma 1, it follows that $\delta_{G'}(x, y) = \delta_G(x, y)$, $\forall x, y \in V(G')$.

Let $u \in V(G_i), v \in V(G_j) - V(G_i)$. Then, any path from u to v must pass through a vertex in each of the cut-sets of G_i and G_j. Then we have $\delta_G(u, v) = \min\{\delta_{G_i}(u, x) + \delta_{G'}(x, y) + \delta_{G_j}(y, v) : \ x \in C(G_i), \ y \in C(G_j)\}$. Similarly, for $u, v \in V(G_i)$, we have $\delta_G(u, v) = \min\{\delta_{G_i}(u, v), \min\{\delta_{G_i}(u, x) + \delta_{G'}(x, y) + \delta_{G_i}(y, v) : \ x, y \in C(G_i)\}\}$. If we are able to make queries of the form $\delta_{G_i}(x, y)$ and $\delta_{G'}(x, y)$, the above directly yields a query algorithm for any pair of vertices x, y.

In the following, let $n_i = |V(G_i)|$ and $n' = |V(G')|$. The $(3, 8, \sqrt{n})$-equipartition of G gives us a tree-decomposition of treewidth t for each subgraph G_i, and a tree-decomposition of treewidth $3t$ for G'. We balance these tree-decompositions, yielding tree-decompositions for each G_i with treewidth at most $3t + 2 \leq 5t$ and for G' with treewidth at most $9t + 2 \leq 11t$. Inductively, we construct in parallel, $D(G_i, n_i, r - 1, 5t)$ for each $1 \leq i \leq q$, which enables us to answer queries of the

form $\delta_{G_i}(x,y)$, and $D(G', n', r-1, 11t)$ which enables us to answer queries of the form $\delta_{G'}(x,y)$.

The update procedure is the following: note that $E(G_i) \cap E(G_j) = \emptyset,\ i \neq j$ and $E(G_i) \cap E(G') = \emptyset$, i.e. each edge of G belongs to exactly one of the G_i's or to G'. Suppose the cost of an edge belonging to G_i is changed. Then, we update the data structure for G_i. This may result in new values for $\delta_{G_i}(x,y),\ x,y \in C(G_i)$. We query the updated data structure for $\delta_{G_i}(x,y),\ x,y \in C(G_i)$ and change the weights of the corresponding edges of G', updating the data structure for G' after each change. That the procedure is correct follows from the fact that changing the cost of an edge in G_i does not change $\delta_{G_j}(x,y),\ x,y \in C(G_j)$ when $j \neq i$. Thus, after we change, in G', the cost of edges $\langle x,y \rangle,\ x,y \in C(G_i)$, we have $\delta_{G'}(u,v) = \delta_G(u,v),\ u,v \in V(G')$, again, by repeated applications of Lemma 1. After the last update, the data structure for G' yields correct distances in G, between vertices in $V(G')$.

Now suppose we change the cost of an edge belonging to G'. Then the distances $\delta_{G_i}(x,y)$ do not change. Thus, in this case, we simply update the data structure for G'. This completes the description of the preprocessing and update algorithms.

The time and work required to set up this data structure is the time and work required to construct (1) the equipartitions of G_i's and G', and (2) the data structures of G_i's and G' inductively. By Lemma 5, (1) requires $O(\log n)$ time and $O(t^2 n)$ work. Then, writing $PW(r,t)n$ and $PT(r,t)\log n$ for the preprocessing work and time respectively, we have

$$PW(r,t)n \leq t^2 n + \sum_{i=1}^{q} PW(r-1,5t)n_i + PW(r-1,11t)n'$$

$$PT(r,t)\log n \leq \log n + \max\{PT(r-1,11t)\log n', PT(r-1,5t)\log N\}$$

where $N = \max\{n_1, \ldots, n_q\}$.

Querying involves taking the minimum of the results of the sub-queries specified in the query algorithm previously. Writing $Q(r,t)\alpha(n)$ for the query time, we have

$$Q(r,t)\alpha(n) \leq (5t)^2[2Q(r-1,5t)\alpha(N) + Q(r-1,11t)\alpha(n')]$$

During updates, in the worst case, there is one update in a graph G_i and then, at most $(5t)^2$ queries in G_i and updates in graph G'. Thus, with $UW(r,t)n^{(1/2)^{r-1}}$ and $UT(r,t)\log n$ representing the work and time respectively, we have

$$\begin{aligned} UW(r,t)n^{(1/2)^{r-1}} &\leq UW(r-1,5t)N^{(1/2)^{r-2}} \\ &\quad + (5t)^2[Q(r-1,5t)\alpha(N) + UW(r-1,11t)(n')^{(1/2)^{r-2}}] \end{aligned}$$

$$UT(r,t)\log n \leq UT(r-1,5t)\log N + (5t)^2[Q(r-1,5t)\alpha(N) + UT(r-1,11t)\log n']$$

It is easy to show that $n' \leq 88tn^{1/2}$, $\sum_{i=1}^{q} n_i \leq 5tn$ and $N = 5tn^{1/2}$. Using these facts and easy estimates, we obtain the following recurrences.

$$\begin{aligned} PW(r,t)n &\leq 2t^2 PW(r-1,11t) \\ PT(r,t) &\leq 2PT(r-1,11t) \\ Q(r,t) &\leq (5t)^3 Q(r-1,11t) \\ UW(r,t) &\leq (5t)^3 UW(r-1,11t) \\ UT(r,t) &\leq (5t)^3 UW(r-1,11t) \end{aligned}$$

from which the claimed bounds follow.

Thus we can construct $D(G,n,r,t)$, completing the induction. □

The following theorem shows how to obtain an update work of $O(n^\beta)$, for any constant $0 < \beta < 1$, in a digraph of constant treewidth.

Theorem 5. *Let $k \geq 1$ be any constant integer and let $0 < \beta < 1$ be any constant. Given an n-vertex weighted digraph G of constant treewidth, we can construct on an EREW PRAM: (i) $PD(G, \{n, \log^2 n\}, \{n^\beta, \log n\}, \alpha(n))$; and (ii) $PD(G, \{nI_k(n), \log^2 n\}, \{n^\beta, \log n\}, k)$.*

Proof. Using Facts 1 and 2, we can compute a balanced tree-decomposition of G in $O(\log^2 n)$ time and $O(n)$ work on an EREW PRAM. The rest of the proof follows now by Theorem 4, if we set $r = 1 - \log \beta$. □

The algorithms described above give answers to distance queries only. They can be modified to answer path queries as well. Also, before running our update procedure after a change in the weight of an edge, we have to assure that this change does not create a negative cycle in the digraph G. This can be easily tested as follows. Let $\langle u, v \rangle$ be an edge with weight $wt(u,v)$ and let $wt'(u,v)$ be its new weight. Clearly, the new weight $wt'(u,v)$ creates a negative cycle in G iff $\delta_G(v,u) + wt'(u,v) < 0$. This test takes time proportional to that of finding $\delta_G(v,u)$ and hence does not affect our update bound.

References

1. W. Ackermann, "Zum Hilbertschen Aufbau der reellen Zahlen", *Math. Ann.*, 99(1928), pp.118-133.
2. A. Aho, J. Hopcroft and J. Ullman, "The Design and Analysis of Computer Algorithms", Addison-Wesley, 1974.
3. R. Ahuja, T. Magnanti and J. Orlin, "Network Flows", Prentice-Hall, 1993.
4. N. Alon and B. Schieber, "Optimal Preprocessing for Answering On-line Product Queries", Tech. Rep. No. 71/87, Tel-Aviv University, 1987.
5. S. Arnborg, "Efficient Algorithms for Combinatorial Problems on Graphs with Bounded Decomposability - A Survey", *BIT*, 25 (1985), pp.2-23.
6. H. Bodlaender, "NC-algorithms for Graphs with Small Treewidth", *Proc. 14th WG'88*, LNCS 344, Springer-Verlag, pp.1-10, 1989.

7. H. Bodlaender, "A Tourist Guide through Treewidth", *Acta Cybernetica*, Vol.11, No.1-2, pp.1-21, 1993.
8. H. Bodlaender and T. Hagerup, "Parallel Algorithms with Optimal Speedup for Bounded Treewidth", *Proc. 22nd ICALP'95*, LNCS, Springer-Verlag, to appear.
9. S. Chaudhuri and C. Zaroliagis, "Shortest Path Queries in Digraphs of Small Treewidth", *Proc. 22nd ICALP'95*, LNCS, Springer-Verlag, to appear.
10. E. Cohen, "Efficient Parallel Shortest-paths in Digraphs with a Separator Decomposition", *Proc. 5th ACM SPAA*, 1993, pp.57-67.
11. H. Djidjev, G. Pantziou and C. Zaroliagis, "On-line and Dynamic Algorithms for Shortest Path Problems", *Proc. 12th STACS'95*, LNCS 900, pp. 193-204, Springer-Verlag, 1995.
12. Y. Han, V. Pan and J. Reif, "Efficient Parallel Algorithms for Computing All Pair Shortest Paths in Directed Graphs", *Proc. 4th ACM SPAA*, 1992, pp.353-362.
13. J. JáJá, "An Introduction to Parallel Algorithms", Addison-Wesley, 1992.
14. D. Kavvadias, G. Pantziou, P. Spirakis and C. Zaroliagis, "Efficient Sequential and Parallel Algorithms for the Negative Cycle Problem", *Proc. 5th ISAAC*, 1994, LNCS 834, pp.270-278, Springer-Verlag.
15. P. Klein and S. Subramanian, "A linear-processor polylog-time algorithm for shortest paths in planar graphs", *Proc. 34th IEEE Symp. on FOCS*, 1993, pp.259-270.
16. N. Robertson and P. Seymour, "Graph Minors I: Excluding a Forest", *J. Comb. Theory Series B*, 35 (1983), pp.39-61.
17. N. Robertson and P. Seymour, "Graph Minors II: Algorithmic Aspects of Treewidth", *J. Algorithms*, 7 (1986), pp.309-322.
18. R.E. Tarjan, "Data Structures and Network Algorithms", SIAM, 1983.