

On the Fine-Grained Complexity of Least Weight Subsequence in Multitrees and Bounded Treewidth DAGs

Jiawei Gao¹

University of California, San Diego, CA, USA
jiawei@cs.ucsd.edu

Abstract

This paper introduces a new technique that generalizes previously known fine-grained reductions from linear structures to graphs. Least Weight Subsequence (LWS) [30] is a class of highly sequential optimization problems with form $F(j) = \min_{i < j} [F(i) + c_{i,j}]$. They can be solved in quadratic time using dynamic programming, but it is not known whether these problems can be solved faster than $n^{2-o(1)}$ time. Surprisingly, each such problem is subquadratic time reducible to a highly parallel, non-dynamic programming problem [36]. In other words, if a “static” problem is faster than quadratic time, so is an LWS problem. For many instances of LWS, the sequential versions are equivalent to their static versions by subquadratic time reductions. The previous result applies to LWS on linear structures, and this paper extends this result to LWS on paths in sparse graphs, the Least Weight Subpath (LWSP) problems. When the graph is a multitree (i.e. a DAG where any pair of vertices can have at most one path) or when the graph is a DAG whose underlying undirected graph has constant treewidth, we show that LWSP on this graph is still subquadratically reducible to their corresponding static problems. For many instances, the graph versions are still equivalent to their static versions.

Moreover, this paper shows that if we can decide a property of form $\exists x \exists y P(x, y)$ in subquadratic time, where P is a quickly checkable property on a pair of elements, then on these classes of graphs, we can also in subquadratic time decide whether there exists a pair x, y in the transitive closure of the graph that also satisfy $P(x, y)$.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases fine-grained complexity, dynamic programming, graph reachability

Digital Object Identifier 10.4230/LIPIcs.IPEC.2019.16

Related Version A full version of the paper is available on ECCC : <https://ecc.ecc.weizmann.ac.il/report/2019/045/>.

Funding Work supported by a Simons Investigator Award from the Simons Foundation.

Acknowledgements I sincerely thank Russell Impagliazzo for his guidance and advice on this paper. I would like to thank Marco Carmosino and Jessica Sorrell for helpful comments. Also I would like to thank the anonymous reviewers for comments on an earlier version of this paper.

1 Introduction

1.1 Extending one-dimensional dynamic programming to graphs

Least Weight Subsequence (LWS) [30] is a type of dynamic programming problems: select a set of elements from a linearly ordered set so that the total cost incurred by the adjacent pairs of selected elements is optimized. It is defined as follows: Given elements x_0, \dots, x_n ,

¹ Now at Google.



© Jiawei Gao;

licensed under Creative Commons License CC-BY

14th International Symposium on Parameterized and Exact Computation (IPEC 2019).

Editors: Bart M. P. Jansen and Jan Arne Telle; Article No. 16; pp. 16:1–16:17

Leibniz International Proceedings in Informatics



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

and an $n \times n$ matrix C of costs $c_{i,j}$ for all pairs of indices $i < j$, compute F on all elements, defined by

$$F(j) = \begin{cases} 0, & \text{for } j = 1 \\ \min_{0 \leq i < j} [F(i) + c_{i,j}], & \text{for } j = 2, \dots, n \end{cases}$$

$F(j)$ is the optimal cost value from the first element up to the j -th element. We use the notation LWS_C to define the LWS problem with cost matrix C . The Airplane Refueling problem [30] is a well known example of LWS: Given the locations of airports on a line, find a subset of the airports for an airplane to add fuel, that minimizes the total cost. The cost of flying from the i -th to the j -th airport without stopping is defined by $c_{i,j}$. Other LWS examples include finding a longest chain satisfying a certain property, such as Longest Increasing Subsequence [25] and Longest Subset Chain [36]; breaking a linear structure into blocks, such as Pretty Printing [34]; variations of Subset Sum such as special versions of the Coin Change problem and the Knapsack problem [36]. These problems have $O(n^2)$ time algorithms using dynamic programming, and in many special cases it can be improved: when the cost satisfies the quadrangle inequality or some other properties, there are near linear time algorithms [50, 46, 26]. But for the general LWS, it is not known whether these problems can be solved faster than $n^{2-o(1)}$ time.

A general approach to understanding the fine-grained complexity of these problems was initiated in [36]. Many LWS problems have succinct representations of $c_{i,j}$. Usually C is defined implicitly by the data associated to each element, and the size of the data on each element is relatively small compared to n . Taking problems defined in [36] as examples, in **LowRankLWS**, $c_{i,j} = \langle \mu_i, \sigma_j \rangle$, where μ_i and σ_j are boolean vectors of length $d \ll n$ associated to each element that are given by the input. The **ChainLWS** problem has costs c_1, \dots, c_n defined by a boolean relation P so that $c_{i,j}$ equals c_j if $P(i, j)$ is true, and ∞ otherwise. P is computable by data associated to element i and element j . (For example, in **LongestSubsetChain**, $P(i, j)$ is true iff set S_i is contained in set S_j , where S_i and S_j are sets associated to elements i and j respectively.) So the goal of the problem becomes finding a longest chain of elements so that adjacent elements that are to be selected satisfy property P . When C can be represented succinctly, we can ask whether there exist subquadratic time algorithms for these problems, or try to find subquadratic time reductions between problems. [36] showed that in many LWS_C problems where C can be succinctly described in the input, the problem is subquadratic time reducible to a corresponding problem, which is called a **StaticLWS $_C$** problem. The problem **StaticLWS $_C$** is: given elements x_1, \dots, x_n , a cost matrix C , and values $F(i)$ on all $i \in \{1, \dots, n/2\}$, compute $F(j) = \min_{i \in \{n/2+1, \dots, n\}} [F(i) + c_{i,j}]$ for all $j \in \{n+1, \dots, 2n\}$. It is a parallel, batch version (with many values of j rather than a single one) of the LWS update rule applied sequentially one index at a time in the standard DP algorithm. The reduction from LWS_C to **StaticLWS $_C$** implies that a highly sequential problem can be reducible to a highly parallel one. If a **StaticLWS $_C$** problem can be solved faster than quadratic time, so can the corresponding LWS_C problem. Apart from one-directional reductions from general LWS_C to **StaticLWS $_C$** , [36] also proved subquadratic time equivalence between some concrete problems (**LowRankLWS** is equivalent to **MinInnerProduct**, **NestedBoxes** is equivalent to **VectorDomination**, **LongestSubsetChain** is equivalent to **OrthogonalVectors**, and **ChainLWS**, which is a generalization of **NestedBoxes** and **LongestSubsetChain**, is equivalent to **Selection**, a generalization of **VectorDomination** and **OrthogonalVectors**).

Some of the LWS problems can be naturally extended from lines to graphs. For example, on a road map, we wish to find a path for a vehicle, along which we wish to find a sequence of cities where the vehicle can rest and add fuel so that the total cost is minimized. The cost

of traveling between cities x and y without stopping is defined by cost $c_{x,y}$. Connections between cities could be a general graph, not just a line. Works about algorithms for special LWS problems on special classes of graphs include [11, 43, 24, 38].

Using a similar approach as [36], this paper extends the Least Weight Subsequence problems to the Least Weight Subpath (LWSP_C) problem whose objective is to find a least weight subsequence on a path of a given DAG $G = (V, E)$. Let there be a set V_0 containing vertices that can be the starting point of a subsequence in a path. The optimum value on each vertex is defined by:

$$F(v) = \begin{cases} \min(0, \min_{u \rightsquigarrow v} [F(u) + c_{u,v}]), & \text{for } v \in V_0 \\ \min_{u \rightsquigarrow v} [F(u) + c_{u,v}], & \text{for } v \notin v_0 \end{cases}$$

where $u \rightsquigarrow v$ means v is reachable from u . The goal of LWSP_C is to compute $F(v)$ for all vertices $v \in V$. Examples of LWSP_C problems will be given in Appendix B. LWSP_C can be solved in time $O(|V| \cdot |E|)$ by doing reversed depth/breadth first search from each vertex, and update the F value on the vertex accordingly. It is not known whether it has faster algorithms, even for Longest Increasing Subsequence, which is an LWS_C instance solvable in $O(n \log n)$ time on linear structures. If C is succinctly describable in similar ways as LowRankLWS, NestedBoxes, SubsetChain or ChainLWS, we wish to study if there are subquadratic time algorithms or subquadratic time reductions between problems.

For the cost matrix C , we consider that every vertex has some additional data so that $c_{x,y}$ can be computed by the data contained in x and y . Let the size of additional data associated to each vertex v be its weighted size $w(v)$. The weight of a vertex can be defined in different ways according to the problems. For example, in LowRankLWS, the weighted size of an element can be defined as the dimension of its associated vector; and in SubsetChain, the weighted size of an element is the size of its corresponding subset. We use $m = |E|$ as the number of graph edges. Let n be the number of vertices. We study the case where the graph is sparse, i.e. $m = n^{1+o(1)}$. Let the total weighted size of all vertices be N . For LWS_C and other problems without graphs, we use N as the input size. For LWSP_C and other problems on graphs, we use $M = \max(m, N)$ as the size of the input.

In this paper we will see that if we can improve the algorithm for StaticLWS_C to $N^{2-o(1)}$, then on some classes of graphs we can solve LWSP_C faster than $M^{2-o(1)}$ time.

1.2 Fine-grained complexity preliminaries

Fine-grained complexity studies the exact-time reductions between problems, and the completeness of problems in classes under exact-time reductions. These reductions have established conditional lower bounds for many interesting problems. The Orthogonal Vectors problem (OV) is a well-studied problem solvable in quadratic time. If the *Strong Exponential Time Hypothesis (SETH)* [31, 32] is true, then OV does not have truly subquadratic time algorithms [47]. The problem OV is defined as follows: Given n boolean vectors of dimension $d = \omega(\log n)$, and decide whether there is a pair of vectors whose inner product is zero. The best algorithm is in time $n^{2-\Omega(1/\log(d/\log n))}$ [7, 23]. The *Moderate-dimension OV conjecture (MDOVC)* states that for all $\epsilon > 0$, there are no $O(n^{2-\epsilon} \text{poly}(d))$ time algorithms that solve OV with vector dimension d . If this conjecture is true, then many interesting problems would get lower bounds, including dynamic programming problems such as Longest Common Subsequence [2, 20], Edit Distance [14, 5], Fréchet distance [18, 21, 22], Local Alignment [9], CFG Parsing and RNA Folding [1], Regular Expression Matching [15, 19], and also many graph problems [42, 8, 16]. There are also conditional hardness results about graph problems based on the hardness of All Pair Shortest Path [49, 4, 10, 39] and 3SUM [6, 35].

The *fine-grained reduction* was introduced in [49], which can preserve polynomial saving factors in the running time between problems. The statements for fine-grained complexity are usually like this: if there is some $\epsilon_2 > 0$ such that problem Π_2 of input size n is in $\text{TIME}((T_2(n))^{1-\epsilon_2})$, then problem Π_1 of input size n is in $\text{TIME}((T_1(n))^{1-\epsilon_1})$ for some ϵ_1 . If T_1 and T_2 are both $O(n^2)$ then this reduction is called a subquadratic reduction. Furthermore, the *exact-complexity reduction* is a more strict version that can preserve sub-polynomial savings factors between problems. We use $(\Pi_1, T_1(n)) \leq_{\text{EC}} (\Pi_2, T_2(n))$ to denote that there is a reduction from problem Π_1 to problem Π_2 so that if problem Π_2 is in $\text{TIME}(T_2(n))$, then problem Π_1 is in $\text{TIME}(T_1(n))$.

1.3 Introducing reachability to first-order model checking

Similar to extending LWS_C to paths in graphs, introducing transitive closure to first-order logic also which makes parallel problems become sequential. The first-order property (or first-order model checking) problem is to decide whether an input structure satisfies a fixed first-order logic formula φ . Although model checking for input formulas is PSPACE-complete [44, 45], when φ is fixed by the problem, it is solvable in polynomial time. We consider the class of problems where each problem is the model checking for a fixed formula φ . The sparse version of OV [27] is one of these problems, defined by the formula $\exists u \exists v \forall i \in [d](\neg \text{One}(u, i) \vee (\neg \text{One}(v, i)))$, where relation $\text{One}(u, i)$ is true iff the i -th coordinate of vector u is one.

If φ has k quantifiers ($k \geq 2$), then on input structures of n elements and m tuples of relations, it can be solved in time $O(n^{k-2}m)$ [28]. On dense graphs where $k \geq 9$, it can be solved in time $O(n^{k-3+\omega})$, where ω is the matrix multiplication exponent [48]. Here we study the case where the input structure is sparse, i.e. $m = n^{1+o(1)}$, and ask whether a three-quantifier first-order formula can be model checked in time faster than $m^{2-o(1)}$. The *first-order property conjecture (FOPC)* states that there exists integer $k \geq 2$, so that first-order model checking for $(k+1)$ -quantifier formulas cannot be solved in time $O(m^{k-\epsilon})$ for any $\epsilon > 0$. This conjecture is equivalent to MDOVC, since OV is proven to be a complete problem in the class of first-order model checking problems; in other words, any model checking problem of 3 quantifier formulas on sparse graphs is subquadratic time reducible to OV [28]. This means from improved algorithms for OV we can get improved algorithms for first-order model checking.

The first-order property problems are highly parallelizable. If we introduce the transitive closure (TC) operation on the relations, then these problems will become sequential. The transitive closure of a binary relation E can be considered as the reachability relation by edges of E in a graph. In a sparse structure, the TC of a relation may be dense. So it can be considered as a dense relation succinctly described in the input. In finite model theory, adding transitive closure significantly adds to the expressive power of first-order logic (First discovered by Fagin in 1974 according to [37], and then re-discovered by [12].) In fine-grained complexity, adding arbitrary transitive closure operations on the formulas strictly increases the hardness of the model checking problem. More precisely, [27] shows that SETH on constant depth circuits, which is a weaker conjecture than the SETH (which concerns k -CNF-SAT), implies the model checking for two-quantifier first-order formulas with transitive closure operations cannot be solved in time $O(m^{2-\epsilon})$ for any $\epsilon > 0$. This means this problem may stay hard even if the SETH on k -CNF-SAT is refuted.

However, we will see that for a class of three-quantifier formulas with transitive closure, model checking is no harder than OV under subquadratic time reductions.

We define problem Selection_P to be the decision problem for whether an input structure satisfies $(\exists x \in X)(\exists y \in Y)P(x, y)$. $P(x, y)$ is a fixed property specified by the problem that

can be decided in time $O(w(x) + w(y))$, where weighted size $w(x)$ is the size of additional data on element x . For example, OV is Selection_P where $P(x, y)$ iff x and y are a pair of orthogonal vectors. In this case $w(x)$ is defined as the length of vector x . (If we work on the sparse version of OV , the weighted size $w(x)$ is defined by the Hamming weight of x .)

On a directed graph $G = (V, E)$, we define Path_P to be the problem of deciding whether $(\exists x \in V)(\exists y \in V)[\text{TC}_E(x, y) \wedge P(x, y)]$, where TC_E is the transitive closure of relation E and $P(x, y)$ is a property on x, y fixed by the problem. That is, whether there exist two vertices x, y not only satisfying property P but also y is reachable from x by edges in E . We will give an example of Path_P in Appendix B. Also, we define ListPath_P to be the problem of listing all $x \in V$ such that $(\exists y \in V)[\text{TC}_E(x, y) \wedge P(x, y)]$.

Considering the model checking problems, we let PathFO_3 and ListPathFO_3 denote the class of Path_P and ListPath_P such that P is of form $\exists z\psi(x, y, z)$ or $\forall z\psi(x, y, z)$, where ψ is a quantifier-free formula in first-order logic. Later we will see that problems in PathFO_3 and ListPathFO_3 are no harder than OV . In these model checking problems, the weighted size of an element is the number of tuples in the input structure that the element is contained in.

Trivially, Selection_P on input size (N_1, N_2) can be decided in time $O(N_1N_2)$, where N_1 is the total weighted size of elements in X , and N_2 is the total weighted size of elements in Y . Path_P and ListPath_P on input size M and total vertex weighted size N are solvable time $O(MN)$ by depth/breadth first search from each vertex, where M is defined to be the maximum of N and the number of edges m . This paper will show that on some graphs, if Selection_P is in truly subquadratic time, so is Path_P and ListPath_P . Interestingly, by applying the same reduction techniques from Path_P to Selection_P , we can get a similar reduction from a dynamic programming problem on a graph to a static problem.

1.4 Main results

This paper works on two classes of graphs, both having some similarities to trees. The first class is where the graph G is a multitree. A *multitree* is a directed acyclic graph where the set of vertices reachable from any vertex form a tree. Or equivalently a DAG is a multitree if and only if on all pairs of vertices u, v , there is at most one path from u to v . In different contexts, multitrees are also called *strongly unambiguous graphs*, *mangroves* or *diamond-free posets* [29]. These graphs can be used to model computational paths in nondeterministic algorithms where there is at most one path connecting any two states [13]. The butterfly network, which is a widely-used model of the network topology in parallel computing, is an example of multitrees. We also work on multitrees of strongly connected component, which is a graph that when each strongly connected components are replaced by a single vertex, the graph becomes a multitree.

The second class of graphs is when we treat G as undirected by replacing all directed edges by undirected edges, the underlying graph has constant treewidth. *Treewidth* [40, 41] is an important parameter of graphs that describes how similar they are to trees.² On these classes of graphs, we have the following theorems.

► **Theorem 1** (Reductions between decision problems.). *Let $t(M) \geq 2^{\Omega(\sqrt{\log M})}$, and let the graph $G = (V, E)$ satisfy one of the following conditions:*

- *G is a multitree, or*
- *G is a multitree of strongly connected components, or*
- *The underlying undirected graph of G has constant treewidth,*

² Here we consider the undirected treewidth, where both the graph and the decomposition tree are undirected. It is different from *directed treewidth* defined for directed graphs by [33].

then, the following statements are true:

- If Selection_P is in time $N_1N_2/t(\min(N_1, N_2))$, then Path_P is in time $M^2/t(\text{poly}M)$.³
- If Path_P is in time $M^2/t(M)$, then ListPath_P is in time $M^2/t(\text{poly}M)$.
- When $P(x, y)$ is of form $\exists z\psi(x, y, z)$ or $\forall z\psi(x, y, z)$ where ψ is a quantifier-free first-order formula, Selection_P is in time $N_1N_2/t(\min(N_1, N_2))$ iff Path_P is in time $M^2/t(\text{poly}M)$ iff ListPath_P is in time $M^2/t(\text{poly}M)$.

This theorem implies that OV is hard for classes PathFO_3 and ListPathFO_3 . By the improved algorithm for OV [7, 23], we get improved algorithms for PathFO_3 and ListPathFO_3 :

► **Corollary 2** (Improved algorithms.). *Let the graph G be a multitree, or multitree of strongly connected components, or a DAG whose underlying undirected graph has constant treewidth. Then PathFO_3 and ListPathFO_3 are in time $M^2/2^{\Omega(\sqrt{\log M})}$.*

Next, we consider the dynamic programming problems. If the cost matrix C in LWSP_C is succinctly describable, we get the following reduction from LWSP_C to StaticLWS_C .

- **Theorem 3** (Reductions between optimization problems.). *On a multitree graph, or a DAG whose underlying undirected graph has constant treewidth, let $t(N) \geq 2^{\Omega(\sqrt{\log N})}$, then,*
1. *if StaticLWS_C of input size N is in time $N^2/t(N)$, then LWSP_C on input size M is in time $M^2/t(\text{poly}(M))$.*
 2. *if LWSP_C is in time $M^2/t(M)$, then LWS_C is in time $N^2/t(\text{poly}(N))$.*

If there is a reduction from a concrete StaticLWS_C problem to its corresponding LWS_C problem (e.g. there are reductions from MinInnerProduct to LowRankLWS , from VectorDomination to NestedBoxes and from OV to $\text{LongestSubsetChain}$ [36]), then the corresponding LWS_C , StaticLWS_C and LWSP_C problems are subquadratic-time equivalent. From the algorithm for OV [23] and SparseOV [28], we get improved algorithm for problem $\text{LongestSubsetChain}$:

► **Corollary 4** (Improved algorithm). *On a multitree or a DAG whose underlying undirected graph has constant treewidth, $\text{LongestSubsetChain}$ is in time $M^2/2^{\Omega(\sqrt{\log M})}$.*

The reduction uses a technique that decomposes multitrees into sub-structures where it is easy to decide whether vertices are reachable. So we also get reachability oracles using subquadratic space, that can answer reachability queries in sublinear time.

► **Theorem 5** (Reachability oracle). *On a multitree of strongly connected components, there exists a reachability oracle with subquadratic preprocessing time and space that has sublinear query time. On a multitree, the preprocessing time and space is $O(m^{5/3})$, and the query time is $O(m^{2/3})$.*

1.5 Organization

In Section 2 we prove the first part of Theorem 1, by reduction from Path_P to Selection_P on multitrees. The case for bounded treewidth DAGs will be presented in the full version. Section 3 proves Theorem 3 by presenting a reduction from LWSP_C to StaticLWS_C , and the proof of correctness will be in the full version. Section 4 discusses about open problems.

³ This reduction also applies to optimization versions of these two problems. Let Path_F be a problem to compute $\min_{x,y \in V, x \rightsquigarrow y} F(x, y)$ and Selection_F be a problem to compute $\min_{x \in X, y \in Y} F(x, y)$, where F is a function on x, y , instead of a boolean property. Then the same technique gives us a reduction from Path_F to Selection_F .

Appendix A lists the definitions of problems, and Appendix B shows some concrete problems as examples.

Due to space restrictions, several proofs had to be deferred to the full version, including the rest of Theorem 1, the subquadratic equivalence of Selection_P , Path_P and ListPath_P when P is a first-order property, and the reachability oracle for multitrees.

2 From sequential problems to parallel problems, on multitrees

We will prove the first part of Theorem 1 by showing that if $t(M) \geq 2^{\Omega(\sqrt{\log M})}$, then $(\text{Path}_P, M^2/t(\text{poly}M)) \leq_{\text{EC}} (\text{Selection}_P, N_1N_2/t(\min(N_1, N_2)))$. This section gives the reduction for multitrees and multitrees of strongly connected components. For constant treewidth graphs, the reduction will be shown in the full version.

2.1 The recursive algorithm

The algorithm uses a divide-and-conquer strategy. We will consider each strongly connected component as a single vertex, whose weighted size equals the total weighted size of the component. In the following algorithm, whenever querying Selection_P or exhaustively enumerating pairs of reachable vertices and testing P on them, we can extract all the vertices from a strongly connected component. Thus we will be working on a multitree, instead of a multitree of strongly connected components. Testing P on a pair of vertices (or strongly connected components) of total weighted sizes N_1, N_2 is in time $O(N_1N_2)$.

Let CutPath_P be a variation of Path_P . It is the property testing problem for $(\exists x \in S)(\exists y \in T)[TC_E(x, y) \wedge \varphi(x, y)]$, where (S, T) is a cut in the graph, such that all the edges between S and T are directed from S to T . CutPath_P on input size M and total vertex weighted size N can be solved in time $O(MN)$ if $P(x, y)$ is decidable in time $O(w(x) + w(y))$: start from each vertex and do depth/breadth first search, and on each pair of reachable vertices decide if P is satisfied.

► **Lemma 6.** *For $t(M) \geq 2^{\Omega(\sqrt{\log M})}$, if $\text{Selection}_P(N, N)$ is in time $N^2/t(N)$ and $\text{CutPath}_P(M)$ is in time $M^2/t(M)$, then $\text{Path}_P(M)$ is in time $M^2/t(\text{poly}(M))$.*

Proof. Let γ be a constant satisfying $0 < \gamma \leq 1/4$. Let $T_{\Pi}(M)$ be the running time of problem Π on a structure of total weighted size M . We show that there exists a constant c where $0 < c < 1$ so that if $T_{\text{Path}_P}(M')$ is at most $M'^2/t(M'^c)$ for all $M' < M$, then $T_{\text{Path}_P}(M) \leq M^2/t(M^c)$. We run the recursive algorithm as shown in Algorithm 1. The intuition is to divide the graph into a cut S, T , recursively compute Path_P on S and T , and deal with paths from S to T .

It would be good if the difference of total weighted sizes between S and T is at most M^γ . Otherwise, it means by the topological order, there is a vertex of weighted size at least M^γ in the middle, adding it to either S or T would make the size difference between S and T exceed M^γ . In this case, we use letter x to denote the vertex. We will deal with x separately. We temporarily set aside the time of recursively running Selection_P on x (when x is shrunk from a strongly connected component) in all the recursive calls, and consider the rest of the running time.

Algorithm 1 $\text{Path}_P(G)$ on a DAG

```

// Reducing  $\text{Path}_P$  to  $\text{Selection}_P$  and  $\text{CutPath}_P$ 
1 if  $G$  has only one vertex then return false.
2 Let  $M$  be the weighted size of the problem.
3 Topological sort all vertices.
4 Keep adding vertices to  $S$  by topological order, until the total weighted size of  $S$ 
  exceeds  $M/2$ . Let the rest of vertices be  $T$ .
5 if  $|S| - |T| > M^\gamma$  then
6    $\lfloor$  Let  $x$  be the last vertex added to  $S$ . Remove  $x$  from  $S$ .
7   Run  $\text{Path}_P$  on the subgraph induced by  $S$ .
8   Run  $\text{CutPath}_P(S, T)$ .
9   if  $x$  exists then
10    Run  $\text{CutPath}_P(S, x)$ .
11     $\lfloor$  If  $x$  is originally a strongly connected component, run  $\text{Selection}_P$  on it.
12    Run  $\text{CutPath}_P(x, T)$ 
13 Run  $\text{Path}_P$  on the subgraph induced by  $T$ .
14 if any one of the above three calls returns true then return true.

```

Let M_S and M_T be the sizes of sets S and T respectively. Without loss of generality, assume $M_S \geq M_T$, and let $\Delta = M_S - M_T$, which is at most M^γ . Then we have

$$\begin{aligned}
T_{\text{Path}_P}(M) &= T_{\text{Path}_P}(M_S) + T_{\text{Path}_P}(M_T) + 3T_{\text{CutPath}_P}(M) + O(M) \\
&= T_{\text{Path}_P}(M_T + \Delta) + T_{\text{Path}_P}(M_T) + 3T_{\text{CutPath}_P}(M) + O(M) \\
&\leq 2T_{\text{Path}_P}(M/2 + \Delta) + 3T_{\text{CutPath}_P}(M) + O(M) \\
&= 2(M/2 + \Delta)^2/t((M/2 + \Delta)^c) + 3M^2/t(M) + O(M).
\end{aligned}$$

Because $t(M) < M$ and is monotonically growing, The term $3M^2/t(M) + O(M)$ is bounded by $4M^2/t(M) \leq 16(M/2)^2/t(M) \leq 16(M/2 + \Delta)^2/t((M/2 + \Delta)^c)$. Thus the above formula is bounded $18(M/2 + \Delta)^2/t((M/2 + \Delta)^c)$. By picking small enough constant γ and c , this sum is less than $M^2/t(M^c)$.

For the time of running Selection_P on x where x is originally a strongly connected component, we consider all recursive calls of Path_P . Let the size of each such x be M_i . The total time would be $\sum_i M_i^2/t(M_i) < (\sum_i M_i^2)/t(M^\gamma)$. Because $\sum_i M_i \leq M$, the sum is at most $M^2/t(M^\gamma)$, a value subquadratic to M , with M being the input size of the outermost call of Path_P . \blacktriangleleft

2.2 A special case that can be exhaustively searched

The following lemma shows that if no vertex has both a lot of ancestors and a lot of descendants, then the total number of reachable pairs of vertices is subquadratic to m . This lemma holds for any DAG, not just for multitrees. We will use this lemma in the next subsection to show that in a subgraph where all vertices have few ancestors and descendants, we can test property P on all pairs of reachable vertices by brute force. Actually, we will use a weighted version of this lemma, which will be proved in the full version.

► **Lemma 7.** *If in a DAG $G = (V, E)$ of m edges, every vertex has either at most n_1 ancestors or at most n_2 descendants, then there are at most $(m \cdot n_1 \cdot n_2)$ pairs of vertices s, t such that s can reach t .*

In a DAG $G = (V, E)$ of m edges, let S, T be two disjoint sets of vertices where edges between S and T only direct from S to T . If every vertex has either at most n_1 ancestors in S or at most n_2 descendants in T , then there are at most $(m \cdot n_1 \cdot n_2)$ pairs of vertices $s \in S$ and $t \in T$ such that s can reach t .

Proof. We define the ancestors of an edge $e \in E$ to be the ancestors (or ancestors in S) of its incoming vertex, and its descendants to be the descendants (or descendants in T) of its outgoing vertex. Let the number of its ancestors and descendants be denoted by $anc(e)$ and $des(e)$ respectively.

For each edge e , it belongs to exactly one of the following three types:

Type A: If $anc(e) \leq n_1$ but $des(e) > n_2$, then let $count(e)$ be $anc(e)$.

Type B: If $des(e) \leq n_2$ but $anc(e) > n_1$, then let $count(e)$ be $des(e)$.

Type C: If $anc(e) \leq n_1$ and $des(e) \leq n_2$, then let $count(e)$ be $anc(e) \cdot des(e)$.

$\sum_{e \in E} count(e) \leq m \cdot n_1 \cdot n_2$ because the $count$ value on each edge is bounded by $n_1 \cdot n_2$. We will prove that this value upper bounds the number of reachable pairs of vertices.

For each pair of reachable vertices (u, v) (or (u, v) s.t. $u \in S$ and $v \in T$), let (e_1, \dots, e_p) be the path from u to v . Along the path, anc does not decrease, and des does not increase. A path belongs to exactly one of the following three types:

Type a: Along the path $anc(e_1) \leq anc(e_2) \leq \dots \leq anc(e_p) \leq n_1$, and $des(e_1) \geq des(e_2) \geq \dots \geq des(e_p) > n_2$. That is, all the edges are Type A.

Type b: Along the path $des(e_p) \leq des(e_{p-1}) \leq \dots \leq des(e_1) \leq n_2$, and $anc(e_p) \geq anc(e_{p-1}) \geq \dots \geq anc(e_1) > n_1$. That is, all the edges are Type B.

Type c: Along the path there is some edge e_i so that $anc(e_i) \leq n_1$ and $des(e_i) \leq n_2$. That is, it has at least one Type C edge.

There will not be other cases, for otherwise if a Type A edge directly connects to a Type B edge without a Type C edge in the middle, then the vertex joining these two edges would have more than n_1 ancestors and more than n_2 descendants.

If a path from u to v is Type a, then its last edge e_p is Type A. If it is Type b, then its first edge e_1 is Type B. If it is Type c, then there is some edge e_i in the path that is Type C. This means:

1. For each Type A edge e , $count(e)$ is at least the number of all Type a pairs (u, v) whose path has e as its last edge.
2. For each Type B edge e , $count(e)$ is at least the number of all Type b pairs (u, v) whose path has e as its first edge.
3. For each Type C edge e , $count(e)$ is at least the number of all Type c pairs (u, v) whose path contains e .

Therefore each path is counted at least once by the $count(e)$ of some edge e . ◀

2.3 Subroutine: reachability across a cut

Now we will show the reduction from $CutPath_P$ to $Selection_P$. The high level idea of $CutPath_P$ is that we think of the reachability relation on $S \times T$ as an $|S| \times |T|$ boolean matrix whose one-entries correspond to reachable pairs of vertices. If we could partition the matrix into all-one combinatorial rectangles, then we can decide all entries within these rectangles by a query to $Selection_P$, because in the same rectangle, all pairs are reachable.

▷ **Claim 8.** Consider the reachability matrix of on sets S and T . Let M_S and M_T be the sizes of S and T . If there is a way to partition the matrix into non-overlapping combinatorial rectangles $(S_1, T_1), \dots, (S_k, T_k)$ of sizes $(r_1, c_1), \dots, (r_k, c_k)$, and if there is some t so that computing each subproblem of size (r_i, c_i) takes time $r_i \cdot c_i / t(\min(r_i, c_i))$, and all $r_i \geq \ell$, and all $c_i \geq \ell$ for a threshold value ℓ , then all the computation takes total time $O(M_S \cdot M_T / t(\ell))$.

Algorithm 2 $\text{CutPath}_P(S, T)$ on a multitree

```

1 Compute the total weighted size of ancestors  $anc(v)$  and descendants  $des(v)$  for all
  vertices.
2 Insert all vertices with at least  $M^\alpha$  ancestors and  $M^\alpha$  descendants into linked list  $L$ .
3 while there exists a vertex  $v \in L$  do
  | // we call  $v$  a pivot vertex
4   Let  $A$  be the set of ancestors of  $v$  in  $S$ .
5   Let  $B$  be the set of descendants of  $v$  in  $T$ .
6   Add  $v$  to  $A$  if  $v \in S$ , otherwise add  $v$  to  $B$ .
7   Run  $\text{Selection}_P$  on  $(A, B)$ . If it returns true then return true.
8   for each  $a \in A$  do
9     | let  $des(a) = des(a) - |B|$ .
10    | if  $des(a) < M^\alpha$  and  $a \in L$  then remove  $a$  from  $L$ .
11  for each  $b \in B$  do
12    | let  $anc(b) = anc(b) - |A|$ .
13    | if  $anc(b) < M^\alpha$  and  $b \in L$  then remove  $b$  from  $L$ .
14  Remove  $v$  from the graph.
15 for each edge  $(s, t)$  crossing the cut  $(S, T)$  do
16  | Let  $A$  be the set of ancestors of  $s$  (including  $s$ ) in  $S$ .
17  | Let  $B$  be the set of descendants of  $t$  (including  $t$ ) in  $T$ .
18  | On all pairs of vertices  $(a, b)$  where  $a \in A, b \in B$ , check property  $P$ . If  $P$  is true
    | on any pair of  $(a, b)$  then return true.

```

Proof. Let the minimum of all r_i be r_{min} and the minimum of all c_i be c_{min} . Then the factor of time saved for computing each combinatorial rectangle is at least $t(\min(r_{min}, c_{min}))$, greater than $t(\ell)$. So the time spent on all rectangles is at most $O((\sum_{i=1}^t c_i)(\sum_{i=1}^t r_i)/t(\ell))$, also we have $(\sum_{i=1}^t c_i)(\sum_{i=1}^t r_i) \leq M_S \cdot M_T$ because the rectangles are contained inside the matrix of size $M_S \cdot M_T$ and they do not overlap. So the total time is $O(M_S \cdot M_T/t(\ell))$. \triangleleft

The algorithm $\text{CutPath}_P(S, T)$ is shown in Algorithm 2. It tries to cover the one-entries of the reachability matrix by combinatorial rectangles as many as possible. Finally, for the one-entries not covered, we go through them by exhaustive search, which takes less than quadratic time.

In the beginning, we can compute the total weighted size of ancestors (or descendants) of all vertices in the DAG in $O(M)$ time by going through all vertices by topological order (or reversed topological order).

In each query to $\text{Selection}_P(A, B)$, all vertices in A can reach all vertices in B , because they all go through v . For any pair of reachable vertices $s \in S, t \in T$, if they go through any pivot vertex, then the pair is queried to Selection_P . Otherwise it is left to the end, and checked by exhaustive search on all pairs of reachable vertices.

The calls to Selection_P correspond to non-overlapping all-one combinatorial rectangles in the reachability matrix. This is because the graph G is a multitree. For each call to Selection_P , the rectangle size is at least $M^\alpha \times M^\alpha$. Thus the total time for all the Selection_P calls is $O(M^2/t(M^\alpha))$ by Claim 8.

Each time we remove a pivot vertex v , there will be no more paths from set A to set B , for otherwise there would be two distinct paths connecting the same pair of vertices. Thus,

removing a v decreases the total number of weighted-pairs⁴ of reachable vertices by at least $M^\alpha \times M^\alpha$. There are $M \times M$ weighted-pairs of vertices, so the total weight (and thus the total number) of pivot vertices like v is at most $(M \times M)/(M^\alpha \times M^\alpha) = M^{2-2\alpha}$.

Each time we find a pivot vertex v , we update the total weighted size of descendants for all its ancestors, and update the total weighted size of ancestors for all its descendants. Because it has at least M^α ancestors and M^α descendants, the value decrease on each affected vertex is at least M^α . So each vertex has decreased its ancestors/descendants values for at most $M/M^\alpha = M^{1-\alpha}$ times. In other words, each vertex can be an ancestor/descendant of at most $M^{1-\alpha}$ pivot vertices. The total time to deal with all ancestors/descendants of all pivot vertices in the while loop is in $O(M \cdot M^{1-\alpha}) = O(M^{2-\alpha})$.

Finally, after the while loop, there are no vertices with both more than M^α ancestors and M^α descendants. In this case, by a weighted version of Lemma 7 (See the full version), the number of weighted-pairs of reachable vertices is bounded by $M \cdot M^\alpha \cdot M^\alpha = M^{1+2\alpha}$. So the total time to deal with these paths is $O(M^{1+2\alpha})$.

Thus the total running time is $O(M^2/t(M^\alpha) + M^{2-\alpha} + M^{1+2\alpha})$. By choosing α and γ to be appropriate constants, we get subquadratic running time.

If $t(M) = M^\epsilon$, then by choosing $\alpha = 1/(2 + \epsilon)$, we get running time $M^{2-\epsilon/(2+\epsilon)}$.

3 Application to Least Weight Subpath

In this section we will prove Theorem 3. The reduction from LWSP_C to StaticLWS_C uses the same structure as the reduction from Path_P to Selection_P in the proof of Theorem 1 shown in Section 2. Because in LWSP we only consider DAGs, there are no strongly connected components in the graph.

Process $\text{LWSP}_C(G, F_0)$ computes values of F on initial values F_0 defined on all vertices of G . On a given LWSP_C problem, we will reduce it to an asymmetric variation of StaticLWS_C . Process $\text{StaticLWS}_C(A, B, F_A)$ computes all the values of function F_B defined on domain B , given all the values of F_A defined on domain A , such that $F_B(b) = \min_{a \in A} [F_A(a) + c_{a,b}]$. Let N_A and N_B be the total weighted size of A and B respectively. It is easy to see that if StaticLWS_C on $|N_A| = |N_B|$ is in time $N_A^2/t(N_A)$, then StaticLWS_C on general A, B is in time $O(N_A \cdot N_B/t(\min(N_A, N_B)))$.

We also define process $\text{CutLWSP}_C(S, T, F_S)$, which computes all the values of F_T defined on domain T , given all the values of F_S on domain S , where $F_T(t) = \min_{s \in S, s \rightsquigarrow t} [F_S(s) + c_{s,t}]$.

The reduction algorithm is adapted from the reduction from Path_P to Selection_P . LWSP_C is analogous to Path_P , StaticLWS_C is analogous to Selection_P , and CutLWSP_C is analogous to CutPath_P . In Path_P , we divide the graph into two halves, recursively call Path_P on the subgraphs, and use CutPath_P to deal with paths from one side of the graph to the other side. Similarly in LWSP_C , we divide the graph into two halves, recursively compute function F on the source side of the graph, then based on these values we call CutPath_P to compute the initial values of function F on the sink side of the graph, and finally we recursively call LWSP_C on the sink side of the graph. In CutPath_P , we first identify large all-one rectangles in the reachability matrix, and then use Selection_P to solve them, and finally we go through all reachable pairs of vertices that are not covered by these rectangles. Similarly, in LWSP_C , we will use the similar method to identify large all-one rectangles in the reachability matrix and use StaticLWS_C to solve them, and finally we go through all reachable pairs of vertices and update F on each of them.

⁴ The number of weighted-pairs is defined to be the sum of $w(u) \cdot w(v)$ for all pairs of reachable vertices $u \rightsquigarrow v$.

■ **Algorithm 3** $\text{LWSP}_C(G = (V, E, V_0), F_0)$ on a DAG

```

1 if  $G$  has only one vertex  $v$  then
2   if  $v \in V_0$  then
3      $\lfloor$  return  $\min(0, F_0(v))$ .
4    $\rfloor$  return  $F_0$  on  $v$ .
5 Let  $M$  be the weighted size of the problem.
6 Topological sort all vertices.
7 Keep adding vertices to  $S$  by topological order, until the total weighted size of  $S$ 
  exceeds  $M/2$ . Let the rest of vertices be  $T$ .
8 if  $|S| - |T| > M^\gamma$  then
9    $\lfloor$  Let  $x$  be the last vertex added to  $S$ . Remove  $x$  from  $S$ .
10 Compute  $F$  on domain  $S$ , by  $F \leftarrow \text{LWSP}_C(G_S, F_0)$ , where  $G_S$  is the subgraph of  $G$ 
    induced by  $S$ .
11 Let  $F_T \leftarrow \text{CutLWSP}_C(S, T, F)$ .
12 For each vertex  $t \in T$ , let  $F_0(t) \leftarrow \min(F_0(t), F_T(t))$ .
13 if  $x$  exists then
14    $\lfloor$  Compute  $F_x \leftarrow \text{CutLWSP}_C(S, x, F)$  for vertex  $x$ .
15    $\lfloor$  Compute  $F$  on vertex  $x$  by  $F(x) \leftarrow \min(F_0(x), F_x(x))$ .
16    $\lfloor$  Let  $F'_T \leftarrow \text{CutLWSP}_C(x, T, F)$ .
17    $\lfloor$  For each vertex  $t \in T$ , let  $F_0(t) \leftarrow \min(F_0(t), F'_T(t))$ .
18 Compute  $F$  on domain  $T$ , by  $F \leftarrow \text{LWSP}_C(G_T, F_0)$ , where  $G_T$  is the subgraph of  $G$ 
    induced by  $T$ .
19 return  $F$  on domain  $V$ .
```

The algorithm LWSP_C is similar as Path_P (Algorithm 1), and is defined in Algorithm 3. Initially, we let $F(v) \leftarrow 0$ for all $v \in V_0$, and let $F(v) \leftarrow +\infty$ for all $v \notin V_0$. We run $\text{LWSP}_C(G, F_0)$ on the whole graph.

The algorithm $\text{CutLWSP}_C(S, T, F_S)$ is adapted from CutPath_P (Algorithm 2), with the following changes:

1. In the beginning, $F_T(t)$ is initialized to ∞ for all $t \in T$.
2. Each query to $\text{Selection}_P(A, B)$ in CutPath_P is replaced by
 - a. Compute F_B on domain B by $\text{StaticLWS}_C(A, B, F_S)$.
 - b. For each vertex b in B , let $F_T(b)$ be the minimum of the original $F_T(b)$ and $F_B(b)$.
3. Whenever processing a pair of vertices s, t such that s is can reach t in either the preprocessing phase or the final exhaustive search phase, we let $F_T(t) \leftarrow F_S(s) + c_{s,t}$ if $F_S(s) + c_{s,t} < F_T(t)$.
4. In the end, the process returns F_T , the target function on domain T .

The proof of correctness will be shown in the full version. The time complexity of this reduction algorithm follows from the argument of Section 2.

4 Open problems

One open problem is to study Path_P and LWSP_C on general DAGs. Also, we would like to consider the case where the graph is not sparse, where we can use $O(MN)$ as the baseline time complexity instead of $O(M^2)$.

It would also be desirable to study the fine-grained complexity of the DAG versions of other quadratic time solvable dynamic programming problems, e.g. the Longest Common Subsequence problem.

References

- 1 Amir Abboud, Arturs Backurs, Karl Bringmann, and Marvin Künnemann. Fine-grained complexity of analyzing compressed data: Quantifying improvements over decompress-and-solve. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 192–203. IEEE, 2017.
- 2 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 59–78. IEEE, 2015.
- 3 Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. SETH-based lower bounds for subset sum and bicriteria path. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 41–57. SIAM, 2019.
- 4 Amir Abboud, Fabrizio Grandoni, and Virginia Vassilevska Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1681–1697. SIAM, 2014.
- 5 Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. Simulating branching programs with edit distance and friends: or: a polylog shaved is a lower bound made. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 375–388. ACM, 2016.
- 6 Amir Abboud and Kevin Lewi. Exact weight subgraphs and the k-sum conjecture. In *International Colloquium on Automata, Languages, and Programming*, pages 1–12. Springer, 2013.
- 7 Amir Abboud, Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 218–230. SIAM, 2015.
- 8 Amir Abboud, Virginia Vassilevska Williams, and Joshua Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete Algorithms*, pages 377–391. SIAM, 2016.
- 9 Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In *International Colloquium on Automata, Languages, and Programming*, pages 39–51. Springer, 2014.
- 10 Udit Agarwal and Vijaya Ramachandran. Fine-grained Complexity for Sparse Graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 239–252, New York, NY, USA, 2018. ACM. doi:10.1145/3188745.3188888.
- 11 Alok Aggarwal, Baruch Schieber, and Takeshi Tokuyama. Finding a minimum-weight k-link path in graphs with the concave Monge property and applications. *Discrete & Computational Geometry*, 12(3):263–280, 1994.
- 12 Alfred V Aho and Jeffrey D Ullman. Universality of data retrieval languages. In *Proceedings of the 6th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 110–119. ACM, 1979.
- 13 Eric Allender and Klaus-Jörn Lange. $StUSPACE(\log n) \subseteq DSPACE(\log^2 n / \log \log n)$. In *International Symposium on Algorithms and Computation*, pages 193–202. Springer, 1996.
- 14 Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 51–58. ACM, 2015.

- 15 Arturs Backurs and Piotr Indyk. Which regular expression patterns are hard to match? In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 457–466. IEEE, 2016.
- 16 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Towards tight approximation bounds for graph diameter and eccentricities. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 267–280. ACM, 2018.
- 17 Hans L Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on computing*, 25(6):1305–1317, 1996.
- 18 Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 661–670. IEEE, 2014.
- 19 Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. A dichotomy for regular expression membership testing. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 307–318. IEEE, 2017.
- 20 Karl Bringmann and Marvin Künnemann. Quadratic conditional lower bounds for string problems and dynamic time warping. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 79–97. IEEE, 2015.
- 21 Karl Bringmann and Marvin Künnemann. Improved approximation for Fréchet distance on c-packed curves matching conditional lower bounds. *International Journal of Computational Geometry & Applications*, 27(01n02):85–119, 2017.
- 22 Karl Bringmann and Wolfgang Mulzer. Approximability of the discrete Fréchet distance. *Journal of Computational Geometry*, 7(2):46–76, 2015.
- 23 Timothy M Chan and Ryan Williams. Deterministic APSP, Orthogonal Vectors, and More: Quickly derandomizing Razborov-Smolensky. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1246–1255. SIAM, 2016.
- 24 S.C. Chen, J.Y. Wu, G.S. Huang, and R.C.T. Lee. Finding a Longest Increasing Subsequence on a Galled Tree. In *the 28th Workshop on Combinatorial Mathematics and Computation Theory, Penghu, Taiwan*, 2011.
- 25 Michael L Fredman. On computing the length of longest increasing subsequences. *Discrete Mathematics*, 11(1):29–35, 1975.
- 26 Zvi Galil and Kunsoo Park. A linear-time algorithm for concave one-dimensional dynamic programming. *Information Processing Letters*, 33(6):309–311, 1990. doi:10.1016/0020-0190(90)90215-J.
- 27 Jiawei Gao and Russell Impagliazzo. The Fine-Grained Complexity of Strengthenings of First-Order Logic. *Electronic Colloquium on Computational Complexity (ECCC)*, 26:9, 2019. URL: <https://eccc.weizmann.ac.il/report/2019/009>.
- 28 Jiawei Gao, Russell Impagliazzo, Antonina Kolokolova, and Ryan Williams. Completeness for First-order Properties on Sparse Structures with Algorithmic Applications. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 2162–2181, 2017.
- 29 Jerrold R Griggs, Wei-Tian Li, and Linyuan Lu. Diamond-free families. *Journal of Combinatorial Theory, Series A*, 119(2):310–322, 2012.
- 30 Daniel S Hirschberg and Lawrence L Larmore. The least weight subsequence problem. *SIAM Journal on Computing*, 16(4):628–638, 1987.
- 31 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- 32 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.
- 33 Thor Johnson, Neil Robertson, Paul D Seymour, and Robin Thomas. Directed tree-width. *Journal of Combinatorial Theory, Series B*, 82(1):138–154, 2001.

- 34 Donald E Knuth and Michael F Plass. Breaking paragraphs into lines. *Software: Practice and Experience*, 11(11):1119–1184, 1981.
- 35 Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3SUM conjecture. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1272–1287. SIAM, 2016.
- 36 Marvin Künnemann, Ramamohan Paturi, and Stefan Schneider. On the Fine-Grained Complexity of One-Dimensional Dynamic Programming. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:15, 2017.
- 37 Leonid Libkin. *Elements of finite model theory*. Springer Science & Business Media, 2013.
- 38 Guan-Yu Lin, Jia jie Liu, and Yue-Li Wang. Finding a Longest Increasing Subsequence from the Paths in a Complete Bipartite Graph. In *Proceedings of the 29th Workshop on Combinatorial Mathematics and Computation Theory*, 2012.
- 39 Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1236–1252. Society for Industrial and Applied Mathematics, 2018.
- 40 Neil Robertson and Paul D Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984.
- 41 Neil Robertson and P.D Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986. doi:10.1016/0196-6774(86)90023-4.
- 42 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 515–524. ACM, 2013.
- 43 Baruch Schieber. Computing a minimum weightk-link path in graphs with the concave monge property. *Journal of Algorithms*, 29(2):204–222, 1998.
- 44 Larry Joseph Stockmeyer. *The complexity of decision problems in automata theory and logic*. PhD thesis, Massachusetts Institute of Technology, 1974.
- 45 Moshe Y Vardi. The complexity of relational query languages. In *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, pages 137–146. ACM, 1982.
- 46 Robert Wilber. The concave least-weight subsequence problem revisited. *Journal of Algorithms*, 9(3):418–425, 1988.
- 47 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2):357–365, 2005.
- 48 Ryan Williams. Faster decision of first-order graph properties. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, page 80. ACM, 2014.
- 49 Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 645–654. IEEE, 2010.
- 50 F Frances Yao. Efficient dynamic programming using quadrangle inequalities. In *Proceedings of the twelfth annual ACM symposium on Theory of computing*, pages 429–435. ACM, 1980.

A List of problem definitions and class definitions

Here we list the main problems studied in this paper.

LWS_C: Given elements x_1, \dots, x_n and value $F(0) = 0$, compute $F(j) = \min_{0 \leq i < j} [F(i) + c_{i,j}]$ for all $j \in \{1, \dots, n\}$.

StaticLWS_C: Given elements x_1, \dots, x_{2n} and values of $F(i)$ on all $i \in \{1, \dots, n\}$, compute $F(j) = \min_{i \in \{1, \dots, n\}} [F(i) + c_{i,j}]$ for all $j \in \{n+1, \dots, 2n\}$.

16:16 On the Fine-Grained Complexity of LWS in Multitrees and Bounded Treewidth DAGs

LWSP_C: Given graph $G = (V, E)$ and starting vertex set $V_0 \subseteq V$, compute on each $v \in V$, the value of $F(v)$, where

$$F(v) = \begin{cases} \min(0, \min_{u \rightsquigarrow v} [F(u) + c_{u,v}]), & \text{for } v \in V_0 \\ \min_{u \rightsquigarrow v} [F(u) + c_{u,v}], & \text{for } v \notin V_0 \end{cases}$$

CutLWSP_C: On DAG G with a cut (S, T) where edges are only directed from S to T , given the values of function F_S on S , for all $t \in T$ compute $F_T(t) = \min_{s \in S, s \rightsquigarrow t} [F_S(s) + c_{s,t}]$.

Selection_P: On two sets X, Y , decide whether $(\exists x \in X)(\exists y \in Y)P(x, y)$.

Path_P: On graph $G = (V, E)$, decide whether $(\exists x \in V)(\exists y \in V)[\text{TC}_E(x, y) \wedge P(x, y)]$.

ListPath_P: On graph $G = (V, E)$, for all $x \in V$, decide whether $(\exists y \in V)[\text{TC}_E(x, y) \wedge P(x, y)]$.

CutPath_P: On graph $G = (V, E)$ with cut (S, T) where edges only direct from S to T , decide whether $(\exists x \in S)(\exists y \in T)[\text{TC}_E(x, y) \wedge P(x, y)]$.

PathFO₃: class of Path_P problems such that P is of form $\exists z\psi(x, y, z)$ or $\forall z\psi(x, y, z)$, where ψ is a quantifier-free logical formula.

ListPathFO₃: class of ListPath_P problems such that P is of form $\exists z\psi(x, y, z)$ or $\forall z\psi(x, y, z)$, where ψ is a quantifier-free logical formula.

B Problem examples

We give a list of problems that can be considered as instances of LWSP_C or Path_P.

Trip Planning (LWSP version of Airplane Refueling) On a DAG where vertices represent cities and edges are roads, we wish to find a path for a vehicle, along which we wish to find a sequence of cities where the vehicle can rest and add fuel so that the cost is minimized. The cost of traveling between cities x and y is defined by cost $c_{x,y}$. $c_{x,y}$ can be defined in multiple ways, e.g. $c_{x,y}$ is $\text{cost}(y)$ if $\text{dist}(x, y) \leq M$ and ∞ otherwise. $\text{dist}(x, y)$ is the distance between x, y that can be computed by the positions of x, y . M is the maximal distance the vehicle can travel without resting. $\text{cost}(y)$ is the cost for resting at position y .

Longest Subset Chain on graphs (LWSP version of Longest Subset Chain) On a DAG where each vertex corresponds to a set, we want to find a longest chain in a path of the graph such that each set is a subset of its successor. Here $c_{x,y}$ is -1 if S_x is a subset of S_y , and ∞ otherwise.

Multi-currency Coin Change (LWSP version of Coin Change) Consider there are two different currencies, so there are two sets of coins. We need to find a way to get value V_1 for currency #1 and value V_2 for currency #2, so that the total weight of coins is minimized. Each pair of values $v_1 \in \{0, \dots, V_1\}$ and $v_2 \in \{0, \dots, V_2\}$ can be considered as a vertex. We connect vertex (v_1, v_2) to (v'_1, v'_2) iff $v'_1 = v_1 + 1$ or $v'_2 = v_2 + 1$. The whole graph is a grid, and we wish to find a subsequence of a path from $(0, 0)$ to (V_1, V_2) so that the cost is minimized. The cost is defined by $C_{(v_1, v_2), (v'_1, v'_2)} = w_{1, v'_1 - v_1}$ and $C_{(v_1, v_2), (v_1, v'_2)} = w_{2, v'_2 - v_2}$, where $w_{i,j}$ is the weight of a coin of value j from currency # i .

Pretty Printing with alternative expressions (LWSP version of Pretty Printing) The

Pretty Printing problem is to break a paragraph into lines, so that each line have roughly the same length. If a line is too long or too short, then there is some cost depending on the line length. The goal of the problem is to minimize the cost.

For some text, it is hard to print prettily. For example, if there are long formulas in the text, then sometimes its line gets too wide, but if we move the formula into the next line, the original line has too few words. One solution for this issue is to use alternate

wording for the sentence, to rephrase a part of a sentence to its synonym. These sentences have different lengths, and formulas in some of them will be displayed better than others. These different ways can be considered as different paths in a graph, and we wish to find one sentence that has the minimal Pretty Printing cost.

A Path_P instance Say we have a set of words, and we want to find a word chain (a chain of words so that the last letter of the previous word is the same as the first letter of the next word) so that the first word and the last word satisfy some properties, e.g. they do not have similar meanings, they have the same length, they don't have the same letters on the same positions, etc. Each word corresponds to a vertex in the graph. For words that can be consecutive in a word chain, we add an edge to the words.