

1 On the Fine-grained Complexity of Least Weight 2 Subsequence in Multitrees and Bounded 3 Treewidth DAGs

4 **Jiawei Gao**

5 University of California, San Diego

6 jiawei@cs.ucsd.edu

7 — Abstract —

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

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

26 **2012 ACM Subject Classification** Theory of computation \rightarrow Problems, reductions and completeness

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

28 **Digital Object Identifier** 10.4230/LIPIcs.IPEC.2019.15

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

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

33 **1** Introduction

34 **1.1** Extending one-dimensional dynamic programming to graphs

35 Least Weight Subsequence (LWS) [30] is a type of dynamic programming problems: select a
36 set of elements from a linearly ordered set so that the total cost incurred by the adjacent
37 pairs of selected elements is optimized. It is defined as follows: Given elements x_0, \dots, x_n ,
38 and an $n \times n$ matrix C of costs $c_{i,j}$ for all pairs of indices $i < j$, compute F on all elements,
39 defined by

$$40 \quad F(j) = \begin{cases} 0, & \text{for } j = 1 \\ \min_{0 \leq i < j} [F(i) + c_{i,j}], & \text{for } j = 2, \dots, n \end{cases}$$



© J. 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. 15; pp. 15:1–15:23

Leibniz International Proceedings in Informatics



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

41 $F(j)$ is the optimal cost value from the first element up to the j -th element. We use the
 42 notation LWS_C to define the LWS problem with cost matrix C . The Airplane Refueling
 43 problem [30] is a well known example of LWS: Given the locations of airports on a line,
 44 find a subset of the airports for an airplane to add fuel, that minimizes the total cost. The
 45 cost of flying from the i -th to the j -th airport without stopping is defined by $c_{i,j}$. Other
 46 LWS examples include finding a longest chain satisfying a certain property, such as Longest
 47 Increasing Subsequence [25] and Longest Subset Chain [36]; breaking a linear structure
 48 into blocks, such as Pretty Printing [34]; variations of Subset Sum such as special versions
 49 of the Coin Change problem and the Knapsack problem[36]. These problems have $O(n^2)$
 50 time algorithms using dynamic programming, and in many special cases it can be improved:
 51 when the cost satisfies the quadrangle inequality or some other properties, there are near
 52 linear time algorithms [50, 46, 26]. But for the general LWS, it is not known whether these
 53 problems can be solved faster than $n^{2-o(1)}$ time.

54 A general approach to understanding the fine-grained complexity of these problems was
 55 initiated in [36]. Many LWS problems have succinct representations of $c_{i,j}$. Usually C is
 56 defined implicitly by the data associated to each element, and the size of the data on each
 57 element is relatively small compared to n . Taking problems defined in [36] as examples,
 58 in **LowRankLWS**, $c_{i,j} = \langle \mu_i, \sigma_j \rangle$, where μ_i and σ_j are boolean vectors of length $d \ll n$
 59 associated to each element that are given by the input. The **ChainLWS** problem has costs
 60 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 ∞
 61 otherwise. P is computable by data associated to element i and element j . (For example, in
 62 **LongestSubsetChain**, $P(i, j)$ is true iff set S_i is contained in set S_j , where S_i and S_j are sets
 63 associated to elements i and j respectively.) So the goal of the problem becomes finding a
 64 longest chain of elements so that adjacent elements that are to be selected satisfy property
 65 P . When C can be represented succinctly, we can ask whether there exist subquadratic time
 66 algorithms for these problems, or try to find subquadratic time reductions between problems.
 67 [36] showed that in many LWS_C problems where C can be succinctly described in the input,
 68 the problem is subquadratic time reducible to a corresponding problem, which is called a
 69 **StaticLWS $_C$** problem. The problem **StaticLWS $_C$** is: given elements x_1, \dots, x_n , a cost matrix C ,
 70 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
 71 $j \in \{n+1, \dots, 2n\}$. It is a parallel, batch version (with many values of j rather than a single
 72 one) of the LWS update rule applied sequentially one index at a time in the standard DP
 73 algorithm. The reduction from LWS_C to **StaticLWS $_C$** implies that a highly sequential problem
 74 can be reducible to a highly parallel one. If a **StaticLWS $_C$** problem can be solved faster
 75 than quadratic time, so can the corresponding LWS_C problem. Apart from one-directional
 76 reductions from general LWS_C to **StaticLWS $_C$** , [36] also proved subquadratic time equivalence
 77 between some concrete problems (**LowRankLWS** is equivalent to **MinInnerProduct**, **NestedBoxes**
 78 is equivalent to **VectorDomination**, **LongestSubsetChain** is equivalent to **OrthogonalVectors**, and
 79 **ChainLWS**, which is a generalization of **NestedBoxes** and **LongestSubsetChain**, is equivalent to
 80 **Selection**, a generalization of **VectorDomination** and **OrthogonalVectors**).

81 Some of the LWS problems can be naturally extended from lines to graphs. For example,
 82 on a road map, we wish to find a path for a vehicle, along which we wish to find a sequence
 83 of cities where the vehicle can rest and add fuel so that the total cost is minimized. The cost
 84 of traveling between cities x and y without stopping is defined by cost $c_{x,y}$. Connections
 85 between cities could be a general graph, not just a line. Works about algorithms for special
 86 LWS problems on special classes of graphs include [11, 43, 24, 38].

87 Using a similar approach as [36], this paper extends the Least Weight Subsequence
 88 problems to the Least Weight Subpath (**LWSP $_C$**) problem whose objective is to find a least

89 weight subsequence on a path of a given DAG $G = (V, E)$. Let there be a set V_0 containing
 90 vertices that can be the starting point of a subsequence in a path. The optimum value on
 91 each vertex is defined by:

$$92 \quad F(v) = \begin{cases} \min(0, \min_{u \rightsquigarrow v} [F(u) + c_{u,v}]), & \text{for } v \in V_0 \\ \min_{u \rightsquigarrow v} [F(u) + c_{u,v}], & \text{for } v \notin v_0 \end{cases}$$

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

101 For the cost matrix C , we consider that every vertex has some additional data so that
 102 $c_{x,y}$ can be computed by the data contained in x and y . Let the size of additional data
 103 associated to each vertex v be its weighted size $w(v)$. The weight of a vertex can be defined
 104 in different ways according to the problems. For example, in LowRankLWS , the weighted size
 105 of an element can be defined as the dimension of its associated vector; and in SubsetChain ,
 106 the weighted size of an element is the size of its corresponding subset. We use $m = |E|$ as the
 107 number of graph edges. Let n be the number of vertices. We study the case where the graph
 108 is sparse, i.e. $m = n^{1+o(1)}$. Let the total weighted size of all vertices be N . For LWS_C and
 109 other problems without graphs, we use N as the input size. For LWSP_C and other problems
 110 on graphs, we use $M = \max(m, N)$ as the size of the input.

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

113 1.2 Fine-grained complexity preliminaries

114 Fine-grained complexity studies the exact-time reductions between problems, and the com-
 115 pleteness of problems in classes under exact-time reductions. These reductions have estab-
 116 lished conditional lower bounds for many interesting problems. The Orthogonal Vectors
 117 problem (OV) is a well-studied problem solvable in quadratic time. If the *Strong Exponential*
 118 *Time Hypothesis (SETH)* [31, 32] is true, then OV does not have truly subquadratic time
 119 algorithms [47]. The problem OV is defined as follows: Given n boolean vectors of dimension
 120 $d = \omega(\log n)$, and decide whether there is a pair of vectors whose inner product is zero. The
 121 best algorithm is in time $n^{2-\Omega(1/\log(d/\log n))}$ [7, 23]. The *Moderate-dimension OV conjecture*
 122 (*MDOVC*) states that for all $\epsilon > 0$, there are no $O(n^{2-\epsilon} \text{poly}(d))$ time algorithms that solve
 123 OV with vector dimension d . If this conjecture is true, then many interesting problems
 124 would get lower bounds, including dynamic programming problems such as Longest Common
 125 Subsequence [2, 20], Edit Distance [14, 5], Fréchet distance [18, 21, 22], Local Alignment [9],
 126 CFG Parsing and RNA Folding [1], Regular Expression Matching [15, 19], and also many
 127 graph problems [42, 8, 16]. There are also conditional hardness results about graph problems
 128 based on the hardness of All Pair Shortest Path [49, 4, 10, 39] and 3SUM [6, 35].

129 The *fine-grained reduction* was introduced in [49], which can preserve polynomial saving
 130 factors in the running time between problems. The statements for fine-grained complexity
 131 are usually like this: if there is some $\epsilon_2 > 0$ such that problem Π_2 of input size n is in
 132 $\text{TIME}((T_2(n))^{1-\epsilon_2})$, then problem Π_1 of input size n is in $\text{TIME}((T_1(n))^{1-\epsilon_1})$ for some ϵ_1 . If
 133 T_1 and T_2 are both $O(n^2)$ then this reduction is called a subquadratic reduction. Furthermore,

134 the *exact-complexity reduction* is a more strict version that can preserve sub-polynomial
 135 savings factors between problems. We use $(\Pi_1, T_1(n)) \leq_{\text{EC}} (\Pi_2, T_2(n))$ to denote that there
 136 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
 137 problem Π_1 is in $\text{TIME}(T_1(n))$.

138 1.3 Introducing reachability to first-order model checking

139 Similar to extending LWS_C to paths in graphs, introducing transitive closure to first-order
 140 logic also which makes parallel problems become sequential. The first-order property (or
 141 first-order model checking) problem is to decide whether an input structure satisfies a fixed
 142 first-order logic formula φ . Although model checking for input formulas is PSPACE-complete
 143 [44, 45], when φ is fixed by the problem, it is solvable in polynomial time. We consider
 144 the class of problems where each problem is the model checking for a fixed formula φ .
 145 The sparse version of OV [27] is one of these problems, defined by the formula $\exists u \exists v \forall i \in$
 146 $[d](\neg \text{One}(u, i) \vee (\neg \text{One}(v, i)))$, where relation $\text{One}(u, i)$ is true iff the i -th coordinate of vector
 147 u is one.

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

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

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

175 We define problem Selection_P to be the decision problem for whether an input structure
 176 satisfies $(\exists x \in X)(\exists y \in Y)P(x, y)$. $P(x, y)$ is a fixed property specified by the problem that
 177 can be decided in time $O(w(x) + w(y))$, where weighted size $w(x)$ is the size of additional
 178 data on element x . For example, OV is Selection_P where $P(x, y)$ iff x and y are a pair of
 179 orthogonal vectors. In this case $w(x)$ is defined as the length of vector x . (If we work on the
 180 sparse version of OV , the weighted size $w(x)$ is defined by the Hamming weight of x .)

181 On a directed graph $G = (V, E)$, we define Path_P to be the problem of deciding whether
 182 $(\exists x \in V)(\exists y \in V)[\text{TC}_E(x, y) \wedge P(x, y)]$, where TC_E is the transitive closure of relation E
 183 and $P(x, y)$ is a property on x, y fixed by the problem. That is, whether there exist two
 184 vertices x, y not only satisfying property P but also y is reachable from x by edges in E . We
 185 will give an example of Path_P in Appendix B. Also, we define ListPath_P to be the problem
 186 of listing all $x \in V$ such that $(\exists y \in V)[\text{TC}_E(x, y) \wedge P(x, y)]$.

187 Considering the model checking problems, we let PathFO_3 and ListPathFO_3 denote the
 188 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
 189 a quantifier-free formula in first-order logic. Later we will see that problems in PathFO_3 and
 190 ListPathFO_3 are no harder than OV . In these model checking problems, the weighted size of
 191 an element is the number of tuples in the input structure that the element is contained in.

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

200 1.4 Main results

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

212 The second class of graphs is when we treat G as undirected by replacing all directed
 213 edges by undirected edges, the underlying graph has constant treewidth. *Treewidth* [40, 41]
 214 is an important parameter of graphs that describes how similar they are to trees.¹ On these
 215 classes of graphs, we have the following theorems.

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

- 218 ■ *G is a multitree, or*
- 219 ■ *G is a multitree of strongly connected components, or*
- 220 ■ *The underlying undirected graph of G has constant treewidth,*
 221 *then, the following statements are true:*

¹ 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].

- 222 ■ If Selection_P is in time $N_1N_2/t(\min(N_1, N_2))$, then Path_P is in time $M^2/t(\text{poly}M)$.²
 223 ■ If Path_P is in time $M^2/t(M)$, then ListPath_P is in time $M^2/t(\text{poly}M)$.
 224 ■ 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
 225 formula, Selection_P is in time $N_1N_2/t(\min(N_1, N_2))$ iff Path_P is in time $M^2/t(\text{poly}M)$
 226 iff ListPath_P is in time $M^2/t(\text{poly}M)$.

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

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

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

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

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

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

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

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

253 1.5 Organization

254 In Section 2 we prove the first part of Theorem 1, by reduction from Path_P to Selection_P
 255 on multitrees. The case for bounded treewidth DAGs will be presented in Appendix D.
 256 Section 3 proves Theorem 3 by presenting a reduction from LWSP_C to StaticLWS_C , and the
 257 proof of correctness will be left to Appendix E. Section 4 discusses about open problems.
 258 Appendix A lists the definitions of problems, and Appendix B shows some concrete problems

² 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 .

259 as examples. Appendix C gives a weighted version of Lemma 7. Appendix F proves the
 260 second part of Theorem 1 by reduction from ListPath_P to Path_P . Appendix G proves the last
 261 part of Theorem 1, the subquadratic equivalence of Selection_P , Path_P and ListPath_P when
 262 P is a first-order property. Appendix H talks about the reachability oracle for multitrees.

263 **2 From sequential problems to parallel problems, on multitrees**

264 We will prove the first part of Theorem 1 by showing that if $t(M) \geq 2^{\Omega(\sqrt{\log M})}$, then
 265 $(\text{Path}_P, M^2/t(\text{poly}M)) \leq_{\text{EC}} (\text{Selection}_P, N_1N_2/t(\min(N_1, N_2)))$. This section gives the re-
 266 duction for multitrees and multitrees of strongly connected components. For constant
 267 treewidth graphs, the reduction will be shown in Appendix D.

268 **2.1 The recursive algorithm**

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

276 Let CutPath_P be a variation of Path_P . It is the property testing problem for $(\exists x \in$
 277 $S)(\exists y \in T)[TC_E(x, y) \wedge \varphi(x, y)]$, where (S, T) is a cut in the graph, such that all the edges
 278 between S and T are directed from S to T . CutPath_P on input size M and total vertex
 279 weighted size N can be solved in time $O(MN)$ if $P(x, y)$ is decidable in time $O(w(x) + w(y))$:
 280 start from each vertex and do depth/breadth first search, and on each pair of reachable
 281 vertices decide if P is satisfied.

282 **► 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 $\text{CutPath}_P(M)$
 283 is in time $M^2/t(M)$, then $\text{Path}_P(M)$ is in time $M^2/t(\text{poly}(M))$.*

284 **Proof.** Let γ be a constant satisfying $0 < \gamma \leq 1/4$. Let $T_{\Pi}(M)$ be the running time of
 285 problem Π on a structure of total weighted size M . We show that there exists a constant
 286 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
 287 $T_{\text{Path}_P}(M) \leq M^2/t(M^c)$. We run the recursive algorithm as shown in Algorithm 1. The
 288 intuition is to divide the graph into a cut S, T , recursively compute Path_P on S and T , and
 289 deal with paths from S to T .

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

297 Let M_S and M_T be the sizes of sets S and T respectively. Without loss of generality,

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   └─ 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  └─ 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.

```

298 assume $M_S \geq M_T$, and let $\Delta = M_S - M_T$, which is at most M^γ . Then we have

$$\begin{aligned}
299 \quad T_{\text{Path}_P}(M) &= T_{\text{Path}_P}(M_S) + T_{\text{Path}_P}(M_T) + 3T_{\text{CutPath}_P}(M) + O(M) \\
300 \quad &= T_{\text{Path}_P}(M_T + \Delta) + T_{\text{Path}_P}(M_T) + 3T_{\text{CutPath}_P}(M) + O(M) \\
301 \quad &\leq 2T_{\text{Path}_P}(M/2 + \Delta) + 3T_{\text{CutPath}_P}(M) + O(M) \\
302 \quad &= 2(M/2 + \Delta)^2/t((M/2 + \Delta)^c) + 3M^2/t(M) + O(M). \\
303
\end{aligned}$$

304 Because $t(M) < M$ and is monotonically growing, The term $3M^2/t(M) + O(M)$ is bounded
305 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
306 is bounded $18(M/2 + \Delta)^2/t((M/2 + \Delta)^c)$. By picking small enough constant γ and c , this
307 sum is less than $M^2/t(M^c)$.

308 For the time of running Selection_P on x where x is originally a strongly connected
309 component, we consider all recursive calls of Path_P . Let the size of each such x be M_i . The
310 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
311 most $M^2/t(M^\gamma)$, a value subquadratic to M , with M being the input size of the outermost
312 call of Path_P . ◀

313 2.2 A special case that can be exhaustively searched

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

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

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

327 **Proof.** We define the ancestors of an edge $e \in E$ to be the ancestors (or ancestors in S) of
 328 its incoming vertex, and its descendants to be the descendants (or descendants in T) of its
 329 outgoing vertex. Let the number of its ancestors and descendants be denoted by $anc(e)$ and
 330 $des(e)$ respectively.

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

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

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

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

335 $\sum_{e \in E} count(e) \leq m \cdot n_1 \cdot n_2$ because the $count$ value on each edge is bounded by $n_1 \cdot n_2$. We
 336 will prove that this value upper bounds the number of reachable pairs of vertices.

337 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)
 338 be the path from u to v . Along the path, anc does not decrease, and des does not increase.
 339 A path belongs to exactly one of the following three types:

340 **Type a:** Along the path $anc(e_1) \leq anc(e_2) \leq \dots \leq anc(e_p) \leq n_1$, and $des(e_1) \geq des(e_2) \geq$
 341 $\dots \geq des(e_p) > n_2$. That is, all the edges are Type A.

342 **Type b:** Along the path $des(e_p) \leq des(e_{p-1}) \leq \dots \leq des(e_1) \leq n_2$, and $anc(e_p) \geq$
 343 $anc(e_{p-1}) \geq \dots \geq anc(e_1) > n_1$. That is, all the edges are Type B.

344 **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
 345 is, it has at least one Type C edge.

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

349 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
 350 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.
 351 This means:

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

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

359 2.3 Subroutine: reachability across a cut

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

365 ▷ **Claim 8.** Consider the reachability matrix of on sets S and T . Let M_S and M_T be the
 366 sizes of S and T . If there is a way to partition the matrix into non-overlapping combinatorial
 367 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

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.

```

368 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
369 all $c_i \geq \ell$ for a threshold value ℓ , then all the computation takes total time $O(M_S \cdot M_T / t(\ell))$.

370 **Proof.** Let the minimum of all r_i be r_{min} and the minimum of all c_i be c_{min} . Then the
371 factor of time saved for computing each combinatorial rectangle is at least $t(\min(r_{min}, c_{min}))$,
372 greater than $t(\ell)$. So the time spent on all rectangles is at most $O((\sum_{i=1}^t c_i)(\sum_{i=1}^t r_i) / t(\ell))$,
373 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
374 matrix of size $M_S \cdot M_T$ and they do not overlap. So the total time is $O(M_S \cdot M_T / t(\ell))$. ◀

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

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

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

386 The calls to Selection_P correspond to non-overlapping all-one combinatorial rectangles
387 in the reachability matrix. This is because the graph G is a multitree. For each call to
388 Selection_P , the rectangle size is at least $M^\alpha \times M^\alpha$. Thus the total time for all the Selection_P
389 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-pairs³ 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 Appendix C), 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 ,

³ 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   Compute  $F_x \leftarrow \text{CutLWSP}_C(S, x, F)$  for vertex  $x$ .
15   Compute  $F$  on vertex  $x$  by  $F(x) \leftarrow \min(F_0(x), F_x(x))$ .
16   Let  $F'_T \leftarrow \text{CutLWSP}_C(x, T, F)$ .
17   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  $V$ .

```

433 we will use the similar method to identify large all-one rectangles in the reachability matrix
434 and use StaticLWS_C to solve them, and finally we go through all reachable pairs of vertices
435 and update F on each of them.

436 The algorithm LWSP_C is similar as Path_P (Algorithm 1), and is defined in Algorithm 3.
437 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
438 $\text{LWSP}_C(G, F_0)$ on the whole graph.

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

- 441 1. In the beginning, $F_T(t)$ is initialized to ∞ for all $t \in T$.
- 442 2. Each query to $\text{Selection}_P(A, B)$ in CutPath_P is replaced by
 - 443 a. Compute F_B on domain B by $\text{StaticLWS}_C(A, B, F_S)$.
 - 444 b. For each vertex b in B , let $F_T(b)$ be the minimum of the original $F_T(b)$ and $F_B(b)$.
- 445 3. Whenever processing a pair of vertices s, t such that s is can reach t in either the
446 preprocessing phase or the final exhaustive search phase, we let $F_T(t) \leftarrow F_S(s) + c_{s,t}$ if
447 $F_S(s) + c_{s,t} < F_T(t)$.
- 448 4. In the end, the process returns F_T , the target function on domain T .

449 The proof of correctness will be shown in Appendix E. The time complexity of this
450 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. $\text{StUSPACE}(\log n) \subseteq \text{DSPACE}(\log^2 n / \log \log n)$. In *International Symposium on Algorithms and Computation*, pages 193–202. Springer, 1996.

- 500 **14** Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic
501 time (unless SETH is false). In *Proceedings of the forty-seventh annual ACM symposium on*
502 *Theory of computing*, pages 51–58. ACM, 2015.
- 503 **15** Arturs Backurs and Piotr Indyk. Which regular expression patterns are hard to match? In
504 *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages
505 457–466. IEEE, 2016.
- 506 **16** Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein.
507 Towards tight approximation bounds for graph diameter and eccentricities. In *Proceedings of*
508 *the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 267–280. ACM,
509 2018.
- 510 **17** Hans L Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth.
511 *SIAM Journal on computing*, 25(6):1305–1317, 1996.
- 512 **18** Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquad-
513 ratic algorithms unless SETH fails. In *Foundations of Computer Science (FOCS), 2014 IEEE*
514 *55th Annual Symposium on*, pages 661–670. IEEE, 2014.
- 515 **19** Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. A dichotomy for regular
516 expression membership testing. In *Foundations of Computer Science (FOCS), 2017 IEEE*
517 *58th Annual Symposium on*, pages 307–318. IEEE, 2017.
- 518 **20** Karl Bringmann and Marvin Künnemann. Quadratic conditional lower bounds for string
519 problems and dynamic time warping. In *Foundations of Computer Science (FOCS), 2015*
520 *IEEE 56th Annual Symposium on*, pages 79–97. IEEE, 2015.
- 521 **21** Karl Bringmann and Marvin Künnemann. Improved approximation for Fréchet distance on
522 c-packed curves matching conditional lower bounds. *International Journal of Computational*
523 *Geometry & Applications*, 27(01n02):85–119, 2017.
- 524 **22** Karl Bringmann and Wolfgang Mulzer. Approximability of the discrete Fréchet distance.
525 *Journal of Computational Geometry*, 7(2):46–76, 2015.
- 526 **23** Timothy M Chan and Ryan Williams. Deterministic APSP, Orthogonal Vectors, and More:
527 Quickly derandomizing Razborov-Smolensky. In *Proceedings of the Twenty-Seventh Annual*
528 *ACM-SIAM Symposium on Discrete Algorithms*, pages 1246–1255. SIAM, 2016.
- 529 **24** S.C. Chen, J.Y. Wu, G.S. Huang, and R.C.T. Lee. Finding a longest increasing subsequence on
530 a galled tree. In *the 28th Workshop on Combinatorial Mathematics and Computation Theory,*
531 *Penghu, Taiwan*, 2011.
- 532 **25** Michael L Fredman. On computing the length of longest increasing subsequences. *Discrete*
533 *Mathematics*, 11(1):29–35, 1975.
- 534 **26** Zvi Galil and Kunsoo Park. A linear-time algorithm for concave one-dimensional dynamic
535 programming. 1989.
- 536 **27** Jiawei Gao and Russell Impagliazzo. The fine-grained complexity of strengthenings of first-
537 order logic. *Electronic Colloquium on Computational Complexity (ECCC)*, 26:9, 2019. URL:
538 <https://ecc.ecc.weizmann.ac.il/report/2019/009>.
- 539 **28** Jiawei Gao, Russell Impagliazzo, Antonina Kolokolova, and Ryan Williams. Completeness for
540 first-order properties on sparse structures with algorithmic applications. In *Proceedings of*
541 *the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages
542 2162–2181, 2017.
- 543 **29** Jerrold R Griggs, Wei-Tian Li, and Linyuan Lu. Diamond-free families. *Journal of Combinat-*
544 *orial Theory, Series A*, 119(2):310–322, 2012.
- 545 **30** Daniel S Hirschberg and Lawrence L Larmore. The least weight subsequence problem. *SIAM*
546 *Journal on Computing*, 16(4):628–638, 1987.
- 547 **31** Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of*
548 *Computer and System Sciences*, 62(2):367–375, 2001.
- 549 **32** Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly
550 exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.

- 551 33 Thor Johnson, Neil Robertson, Paul D Seymour, and Robin Thomas. Directed tree-width.
552 *Journal of Combinatorial Theory, Series B*, 82(1):138–154, 2001.
- 553 34 Donald E Knuth and Michael F Plass. Breaking paragraphs into lines. *Software: Practice and*
554 *Experience*, 11(11):1119–1184, 1981.
- 555 35 Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3sum conjecture.
556 In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*,
557 pages 1272–1287. SIAM, 2016.
- 558 36 Marvin Künnemann, Ramamohan Paturi, and Stefan Schneider. On the Fine-Grained Com-
559 plexity of One-Dimensional Dynamic Programming. In *44th International Colloquium on*
560 *Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International*
561 *Proceedings in Informatics (LIPIcs)*, pages 21:1–21:15, 2017.
- 562 37 Leonid Libkin. *Elements of finite model theory*. Springer Science & Business Media, 2013.
- 563 38 Guan-Yu Lin, Jia jie Liu, and Yue-Li Wang. Finding a longest increasing subsequence from
564 the paths in a complete bipartite graph. 2012.
- 565 39 Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest
566 cycles and paths in sparse graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM*
567 *Symposium on Discrete Algorithms*, pages 1236–1252. Society for Industrial and Applied
568 Mathematics, 2018.
- 569 40 Neil Robertson and Paul D Seymour. Graph minors. iii. planar tree-width. *Journal of*
570 *Combinatorial Theory, Series B*, 36(1):49–64, 1984.
- 571 41 Neil Robertson and P.D Seymour. Graph minors. ii. algorithmic aspects of tree-width. *Journal*
572 *of Algorithms*, 7(3):309 – 322, 1986. doi:[https://doi.org/10.1016/0196-6774\(86\)90023-4](https://doi.org/10.1016/0196-6774(86)90023-4).
- 573 42 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the
574 diameter and radius of sparse graphs. In *Proceedings of the forty-fifth annual ACM symposium*
575 *on Theory of computing*, pages 515–524. ACM, 2013.
- 576 43 Baruch Schieber. Computing a minimum weightk-link path in graphs with the concave monge
577 property. *Journal of Algorithms*, 29(2):204–222, 1998.
- 578 44 Larry Joseph Stockmeyer. *The complexity of decision problems in automata theory and logic*.
579 PhD thesis, Massachusetts Institute of Technology, 1974.
- 580 45 Moshe Y Vardi. The complexity of relational query languages. In *Proceedings of the fourteenth*
581 *annual ACM symposium on Theory of computing*, pages 137–146. ACM, 1982.
- 582 46 Robert Wilber. The concave least-weight subsequence problem revisited. *Journal of Algorithms*,
583 9(3):418–425, 1988.
- 584 47 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications.
585 *Theoretical Computer Science*, 348(2):357–365, 2005.
- 586 48 Ryan Williams. Faster decision of first-order graph properties. In *Proceedings of the Joint*
587 *Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL)*
588 *and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*,
589 page 80. ACM, 2014.
- 590 49 Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix
591 and triangle problems. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE*
592 *Symposium on*, pages 645–654. IEEE, 2010.
- 593 50 F Frances Yao. Efficient dynamic programming using quadrangle inequalities. In *Proceedings*
594 *of the twelfth annual ACM symposium on Theory of computing*, pages 429–435. ACM, 1980.

595 **A** List of problem definitions and class definitions

596 Here we list the main problems studied in this paper.

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

599 **StaticLWS_C**: Given elements x_1, \dots, x_{2n} and values of $F(i)$ on all $i \in \{1, \dots, n\}$, compute
600 $F(j) = \min_{i \in \{1, \dots, n\}} [F(i) + c_{i,j}]$ for all $j \in \{n + 1, \dots, 2n\}$.

15:16 On the Fine-grained Complexity of LWS in Multitrees and Bounded Treewidth DAGs

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

$$603 \quad F(v) = \begin{cases} \min(0, \min_{u \rightsquigarrow v} [F(u) + c_{u,v}]), & \text{for } v \in V_0 \\ \min_{u \rightsquigarrow v} [F(u) + c_{u,v}], & \text{for } v \notin v_0 \end{cases}$$

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

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

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

608 **ListPath_P**: On graph $G = (V, E)$, for all $x \in V$, decide whether $(\exists y \in V)[\text{TC}_E(x, y) \wedge$
 609 $P(x, y)]$.

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

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

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

616 **B Problem examples**

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

618 **Trip Planning (LWSP version of Airplane Refueling)**

619 On a DAG where vertices represent cities and edges are roads, we wish to find a path for
 620 a vehicle, along which we wish to find a sequence of cities where the vehicle can rest and add
 621 fuel so that the cost is minimized. The cost of traveling between cities x and y is defined by
 622 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 ∞
 623 otherwise. $\text{dist}(x, y)$ is the distance between x, y that can be computed by the positions of
 624 x, y . M is the maximal distance the vehicle can travel without resting. $\text{cost}(y)$ is the cost
 625 for resting at position y .

626 **Longest Subset Chain on graphs (LWSP version of Longest Subset Chain)**

627 On a DAG where each vertex corresponds to a set, we want to find a longest chain in a
 628 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
 629 subset of S_y , and ∞ otherwise.

630 **Multi-currency Coin Change (LWSP version of Coin Change)**

631 Consider there are two different currencies, so there are two sets of coins. We need to
 632 find a way to get value V_1 for currency #1 and value V_2 for currency #2, so that the total
 633 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
 634 be considered as a vertex. We connect vertex (v_1, v_2) to (v'_1, v'_2) iff $v'_1 = v_1 + 1$ or $v'_2 = v_2 + 1$.
 635 The whole graph is a grid, and we wish to find a subsequence of a path from $(0, 0)$ to
 636 (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
 637 $C_{(v_1, v_2), (v_1, v'_2)} = w_{2, v'_2 - v_2}$, where $w_{i,j}$ is the weight of a coin of value j from currency # i .

638 **Pretty Printing with alternative expressions (LWSP version of Pretty Printing)**

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

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

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

649 A Path_P instance

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

656 C Weighted version of Lemma 7

657 ► **Lemma 9.** *If in a vertex-weighted DAG $G = (V, E)$ of m edges, every vertex has either*
 658 *ancestors of total weight at most n or descendants of total weight at most n , then there are*
 659 *at most $(m \cdot n^2)$ weighted-pairs of vertices (s, t) such that s can reach t .*

660 *In a vertex-weighted DAG $G = (V, E)$ of m edges, let S, T be two disjoint sets of vertices*
 661 *where edges between S and T only direct from S to T . If every vertex has either ancestors in*
 662 *S of total weight at most n or descendants in T of total weight at most n , then there are at*
 663 *most $(m \cdot n^2)$ weighted-pairs of vertices $s \in S$ and $t \in T$ such that s can reach t .*

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

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

670 For each edge $e = (v_1, v_2)$, it belongs to exactly one of the following three types:

671 **Type A:** If $anc(e) \leq n_1$ but $des(e) > n_2$, then let $count(e)$ be $anc(e) \cdot w(v_2)$.

672 **Type B:** If $des(e) \leq n_2$ but $anc(e) > n_1$, then let $count(e)$ be $w(v_1) \cdot des(e)$.

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

674 $\sum_{e \in E} count(e) \leq m \cdot n_1 \cdot n_2$ because the $count$ value on each edge is bounded by $n_1 \cdot n_2$. We
 675 will prove that this value upper bounds the number of weighted-pairs of reachable vertices.

676 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)
 677 be the path from u to v . Along the path, anc does not decrease, and dec does not increase.
 678 A path belongs to exactly one of the following three types:

679 **Type a:** Along the path $anc(e_1) \leq anc(e_2) \leq \dots \leq anc(e_p) \leq n_1$, and $des(e_1) \geq des(e_2) \geq$
 680 $\dots \geq des(e_p) > n_2$. That is, all the edges are Type A.

681 **Type b:** Along the path $des(e_p) \leq des(e_{p-1}) \leq \dots \leq des(e_1) \leq n_2$, and $anc(e_p) \geq$
 682 $anc(e_{p-1}) \geq \dots \geq anc(e_1) > n_1$. That is, all the edges are Type B.

683 **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
 684 is, it has at least one Type C edge.

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

15:18 On the Fine-grained Complexity of LWS in Multitrees and Bounded Treewidth DAGs

688 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
689 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.
690 This means:

- 691 1. For each Type A edge e , $count(e)$ is at least the weight product $w(u) \cdot w(v)$ of all Type a
692 pairs (u, v) whose path has e as its last edge.
- 693 2. For each Type B edge e , $count(e)$ is at least the weight product $w(u) \cdot w(v)$ of all Type b
694 pairs (u, v) whose path has e as its first edge.
- 695 3. For each Type C edge e , $count(e)$ is at least the weight product $w(u) \cdot w(v)$ of all Type c
696 pairs (u, v) whose path contains e .

697 Therefore the weight product of the endpoints of each path is counted at least once by the
698 $count(e)$ of some edge e . ◀

699 **D** CutPath $_P$ for bounded-treewidth DAGs

700 We prove the first part of Theorem 1 on DAGs whose underlying undirected graphs have
701 constant treewidth. The algorithm Path $_P$ for constant treewidth graphs is the same as the
702 one for multitrees. In this section we will show the reduction algorithm CutPath $_P$ for constant
703 treewidth graphs on a cut (S, T) .

704 Let \mathcal{T} be the decomposition tree of a graph G . Recall that by the definition of tree
705 decomposition, each node z of the tree corresponds to a set $\mathcal{B}(z)$ which is a subset of vertices
706 of G . Because the treewidth is constant, each set $\mathcal{B}(z)$ has a constant number of vertices.
707 Every vertex of G appears in at least one set of a tree node. Also, for every edge of G , there
708 is at least one tree node whose set contains both its endpoints. And if a vertex v appears
709 both in $\mathcal{B}(z_1)$ and $\mathcal{B}(z_2)$, then along the path from z_1 to z_2 , v must appear in all the sets
710 of the tree nodes. Here we consider the decomposition tree as rooted, where all edges are
711 directed from the root to leaves.

712 We use a similar reduction idea as Section 2.3. In the decomposition tree, each time we
713 find a node z to split the tree into two connected components. We first deal with all the
714 paths that go through the vertices in $\mathcal{B}(z)$. Any other path in the graph must be completely
715 contained in one of the connected components we have created. In the end, all connected
716 components are so small that we can go through all pairs of reachable vertices by exhaustive
717 search. The algorithm is defined in Algorithm 4.

718 The following claim uses a $1/3 - 2/3$ trick on trees:

719 \triangleright **Claim 10.** In a vertex-weighted rooted tree of total weight n , we can find a connected
720 subgraph of weighted size between $(1/3)n$ and $(2/3)n$ in $O(n)$ time.

721 **Proof.** For each node z in the tree, we will compute the weighted size of the subtree rooted
722 at z , denoted by $f(z)$. We compute $f(z)$ from the leaves up to the root, by a reversed
723 topological order. If z is a leaf then let $size(z) \leftarrow w(z)$ where $w(z)$ is the weight of z .

724 On each parent node p , we initially let $f(p) \leftarrow w(p)$, and then for each child c_i of p , add
725 the value $f(c_i)$ to $f(p)$. If before we add the $f(c_i)$ of certain child c_i to $f(p)$, $f(p) < (1/3)n$,
726 and after we add $f(c_i)$ to $f(p)$, $f(p) \geq (1/3)n$, then there are two cases:

727 If $f(p) \leq (2/3)n$, then the subgraph formed by p and its subtrees c_1, \dots, c_i is the connected
728 subgraph we want.

729 If $f(p) > (2/3)n$, then it must be $f(c_i) \geq (2/3)n - (1/3)n = (1/3)n$. That is, the subtree
730 rooted at c_i has weighted size between $(1/3)n$ and $(2/3)n$. But then we should have already
731 returned the subtree rooted c_i instead. So this case would not happen.

Algorithm 4: CutPath $_P(S, T)$ on constant treewidth DAG

```

1 Compute  $\mathcal{T}$ , the tree decomposition of the underlying undirected graph.
2 for each  $z$  in  $\mathcal{T}$  do
3    $\lfloor$  Let  $size(z)$  be the number of nodes of  $\mathcal{T}$ .
4 while there exists a tree node  $z$  in  $\mathcal{T}$  so that there is a connected subgraph of  $\mathcal{T}$ 
   rooted at  $z$  with weighted size between  $(1/3)size(z)$  and  $(2/3)size(z)$  do
   //  $z$  can be found in time  $O(size(z))$  by Claim 10.
5   for each  $v \in \mathcal{B}(z)$  do
6     // Deal with all paths going through  $v$ .
7     Let  $A$  be the set of ancestors of  $v$  in  $S$ .
8     Let  $B$  be the set of descendants of  $v$  in  $T$ .
9     Add  $v$  to  $A$  if  $v \in S$ , otherwise add  $v$  to  $B$ .
10    if both  $A$  and  $B$  have at least  $M^\alpha$  vertices then
11       $\lfloor$  Run Selection $_P$  on  $(A, B)$ . If it returns true then return true.
12    else
13       $\lfloor$  Exhaustively check  $P$  on all pairs of  $a \in A$  and  $b \in B$ . If  $P$  is true on any
14         $(a, b)$  then return true.
15      Remove  $v$  from the graph, and from the sets of all the tree nodes.
16    Remove  $z$  from  $\mathcal{T}$ .
17    for each tree node  $z'$  who was originally in the same connected component with  $z$ 
18      do
19         $\lfloor$  Update  $size(z')$  to be the new size of the connected component  $z'$  is in.
20 for each edge  $(s, t)$  crossing the cut $(S, T)$ , do
21   Let  $A$  be the set of ancestors of  $s$  (including  $s$ ) in  $S$ .
22   Let  $B$  be the set of descendants of  $t$  (including  $t$ ) in  $T$ .
23   On all pairs of vertices  $(a, b)$  where  $a \in A, b \in B$ , check property  $P$ . If  $P$  is true
24   on any pair of  $(a, b)$  then return true.

```

732 After we have added the sizes of all the children of p to $f(p)$, we have finished computing
733 $f(p)$. If $f(p)$ is still less than $1/3$, we will continue to let the next vertex by the reversed
734 topological order be the current parent. \blacktriangleleft

735 Next we will analyze the reduction algorithm. First, if a the treewidth of a graph is
736 constant, then the corresponding decomposition tree can be computed in linear time [17].

737 Unlike multitrees, here the calls to Selection $_P$ are not non-overlapping rectangles: different
738 v from the same $\mathcal{B}(z)$ may share the same ancestors or descendants. However, each time
739 after removing a z , the connected components of the decomposition tree correspond to non-
740 overlapping rectangles in the reachability matrix, and will not overlap with the rectangles
741 corresponding to the ancestors and descendants for any $v \in \mathcal{B}(z)$. Thus, the overlapping
742 only happens when dealing with the ancestors and descendants of different v from the same
743 $\mathcal{B}(z)$, and these Selection $_P$ rectangles will not overlap with other Selection $_P$ rectangles after
744 z is removed. Because in each non-overlapping rectangle corresponding to a connected
745 component, we only computed the Selection $_P$ for $|\mathcal{B}(z)|$ times, which is a constant. So by
746 Claim 8, the total time spent on all the calls to Selection $_P$ is still $O(M^2/t(M^\alpha))$.

747 When we remove all vertices $v \in \mathcal{B}(z)$, the graph vertices from sets of different connected

748 components of the decomposition tree are not reachable to each other. Because any path
749 from one connected component to another must go through some vertex in $\mathcal{B}(z)$.

750 Unlike multitree graphs, this time some vertex v in $\mathcal{B}(z)$ may have fewer than M^α
751 ancestors or descendants. If so, then we do exhaustive search on the sets of v 's ancestors and
752 descendants, since calling Selection_P will not save time. Each time we find a v , the connected
753 component of the decomposition tree that v belongs to loses at least $(1/3)\text{size}(v)$ of its vertices,
754 thus each vertex can be the ancestor/descendants of at most $O(\log_{3/2} M)$ such v 's. There
755 are at most M vertices in the graph, each of which can take part in at most M^α such paths
756 going through each such v . So the total time is $O(M \cdot \log_{3/2} M \cdot M^\alpha) = O(M^{1+\alpha} \cdot \log_{3/2} M)$.

757 Also, because each vertex can be the ancestor/descendants of at most $O(\log_{3/2} M)$ such
758 v 's, the total time for updating size for all of them is also bounded by $O(M \cdot \log_{3/2} M)$.

759 In the end, each remaining vertex has $O(M^\alpha)$ ancestors and $O(M^\alpha)$ descendants. The
760 total running time for the exhaustive search is $O(M \cdot M^\alpha \cdot M^\alpha) = O(M^{1+2\alpha})$ by Lemma 7.

761 The overall running time is $O(M^2/t(M^\alpha) + M^{1+\alpha} \cdot \log_{3/2} M + M^{1+2\alpha})$. By choosing α
762 and γ to be appropriate small constants, we get subquadratic running time.

763 **E Correctness of the LWSP_C algorithm**

764 For the correctness proof, we consider the case where there is no x between S and T . The
765 case where there is an x is similar.

766 **Correctness of CutLWSP_C .**

767 The correctness of CutLWSP_C follows from the correctness of CutPath_P . We claim that
768 after running $\text{CutLWSP}_C(S, T, F_S)$, for any vertex $t \in T$, there is $F_T(t) = \min_{s \in S, s \rightsquigarrow t} [F_S(s) +$
769 $c_{s,t}]$. Because for any pair $s \in S, t \in T$, such that s reachable to t , they are either
770 processed in a query to $\text{StaticLWS}_C(A, B)$ where $s \in A, t \in B$, or computed separately thus
771 $F_T(t) \leftarrow \min(F_T(t), F(s) + c_{s,t})$.

772 **Correctness of LWSP_C .**

773 The LWSP_C algorithm has the following facts:

- 774 1. Whenever a process LWSP_C on domain $V_1 \subseteq V$ returns, the values of F on V_1 are fixed
775 and will not be changed henceforth.
- 776 2. Whenever there is an edge from u to v , then the value of F on u is always fixed before the
777 value on v . So the final values of function F on all vertices are fixed by topological order.
- 778 3. Each time we call LWSP_C on a subset of vertices $V_1 \subseteq V$, the F values on all ancestors
779 of any vertex in V_1 that are not in V_1 have been fixed by some previous calls to LWSP_C .

780 Assume that when we call LWSP_C on subgraph with cut (S, T) , initially there is

$$781 \quad F_0(v) = \begin{cases} \min_{u \in R(v) \setminus (S \cup T), u \rightsquigarrow v} [F(u) + c_{u,v}], & \text{if } v \notin V_0 \\ \min(0, \min_{u \in R(v) \setminus (S \cup T), u \rightsquigarrow v} [F(u) + c_{u,v}]), & \text{if } v \in V_0 \end{cases} \quad (1)$$

782 where $R(v)$ is the set of vertices that can reach v . Then, if $\text{LWSP}_C(S, F_0)$ is correct, after
783 running $\text{LWSP}_C(S, F_0)$, for any $s \in S \setminus V_0$, there is $F(s) = \min_{u \in R(s) \setminus T, u \rightsquigarrow s} [F(u) + c_{u,s}]$. And
784 after running $\text{CutLWSP}_C(S, T, F)$, we have $F_T(t) = \min_{s \in S, s \rightsquigarrow t} [F(s) + c_{s,t}]$. Then after taking
785 $F_0(t) = \min(F_0(t), F_T(t))$ on all t , for any $t \in T \setminus V_0$, we get $F_0(t) = \min_{u \in R(t) \setminus T, u \rightsquigarrow t} [F(u) +$
786 $c_{u,t}]$. Similarly for any $t \in T \cap V_0$, $F_0(t)$ gets the the minimum of this value and 0. Therefore,
787 on each call of $\text{LWSP}_C(V_1, F_0)$ on a subset $V_1 \subset V$ with initial values F_0 , F_0 keeps the
788 invariant in formula (1).

F

 From listing problems to decision problems

In this section we prove the second part of Theorem 1, that ListPath_P is reducible to Path_P .

Consider a star graph, which is a graph with its vertex set partitioned in X, Y and another single vertex c . Every $x \in X$ is connected to c , and c is connected to every $y \in Y$. Let problem FindX_P be the following problem: on a star graph, find an $x \in X$ satisfying $(\exists y \in Y)P(x, y)$. We will prove that ListPath_P is reducible to FindX_P and FindX_P is reducible to Path_P .

► **Lemma 11.** *Let $t(M) \geq 2^{\Omega(\sqrt{\log M})}$. $(\text{ListPath}_P, M^2/(t(\text{poly}M))) \leq_{EC} (\text{FindX}_P, M^2/t(M))$*

Proof. We use a grouping reduction technique similar as the trick in [49] and [8].

We modify the algorithm for Path_P in Section 2 to get the algorithm for ListPath_P . That is, we divide the graph into two subgraphs and call ListPath_P recursively in a similar way as Path_P . Path_P needs to call Selection_P as queries, and in the counterpart of ListPath_P we will call FindX_P as queries.

Whenever we need to call $\text{Selection}_P(X, Y)$, we partition X and Y into groups of weighted size at most \sqrt{M} . Thus there are $O((|X|/\sqrt{M}) \times (|Y|/\sqrt{M}))$ groups. For each pair of group X_i, Y_j , we construct a star graph and call FindX_P on it. The star graph is constructed as follows: Connect every $x \in X_i$ to a dummy vertex c , and connect c to every $y \in Y_j$. Thus if there exist some satisfying x in X_i , FindX_P will find a satisfying x .

Every time a satisfying vertex x in X_i is found by FindX_P , we mark it and add it into the list of satisfying x , and then call the FindX_P on the same star again with x removed from the graph. We keep calling FindX_P on this graph, ignoring all marked vertices, until either all elements in X_i are marked and removed, or FindX_P cannot find a satisfying x .

Because there are at most M vertices that can be listed, there are at most M calls to FindX_P that returns a satisfying x . Each call has instance size \sqrt{M} . The running time is $O(M \cdot (\sqrt{M})^2/t(\sqrt{M}))$. The total time spent on the rest of the algorithm is the same as the running time of Path_P . ◀

► **Lemma 12.** *Let $t(M) \geq 2^{\Omega(\sqrt{\log M})}$. $(\text{FindX}_P, M^2/(t(\text{poly}M))) \leq_{EC} (\text{Path}_P, M^2/t(M))$*

Proof. First, we pick an arbitrary element $x_1 \in X$, and construct a graph by letting x_1 connect to all y in Y . Then we call Path_P on this graph. If it returns yes, then we return x_1 .

Otherwise, on the star graph we will replace the center vertex c by x_1 , remove the original x_1 , and call Path_P on this graph. After each call to Path_P , if it returns yes, we divide X in two halves and call Path_P again. Using binary search and shrinking the size of X by half each time, we will finally find a satisfying x . ◀

Lemmas 11 and 12 imply the reduction $(\text{ListPath}_P, M^2/(t(\text{poly}M))) \leq_{EC} (\text{Path}_P, M^2/t(M))$ for $t(M) \geq 2^{\Omega(\sqrt{\log M})}$.

G

 From parallel problems to sequential problems

We prove the third part of Theorem 1, the other direction of the reduction. The reduction from Path_P to ListPath_P is straightforward.

To reduce from Selection_P to Path_P , we can construct a graph with dummy vertex c in the middle, such that each x in set X is connected to c , and c is connected to every y in set Y . If P is expressible by first-order logic, then we will let c act like one of the y 's when computing $R(x, c)$, and act like of the x 's when computing predicates on $P(c, y)$. Let x_1 be an

831 arbitrary element in X , and y_1 be an arbitrary element in Y . We create c by merging x_1 and
 832 y_1 into a single element. c has all the relations x_1 and y_1 have. Thus, on any $x \in X, x \neq x_1$,
 833 the value of $P(x, c)$ is the same as $P(x, y_1)$. Symmetrically on any $y \in Y, y \neq y_1$, the value
 834 of $P(c, y)$ is the same as $P(x_1, y)$. Therefore, there exists x, y such that $P(x, y)$ is true iff
 835 Selection_P on this graph returns true.

836 In general, if we are allowed to define another property P' such that $P'(x, y) \leftarrow (P(x, y) \wedge$
 837 $(x \neq c) \wedge (y \neq c))$, we have a reduction from Selection_P to $\text{Path}_{P'}$.

838 **H** Reachability Oracle

839 This section presents a proof of Theorem 5. A reachability oracle on a graph takes in a
 840 pair of vertices u, v in the graph, and answers whether v is reachable from u . A naive
 841 approach is to use $O(n^2)$ space to store the reachability of all pairs of vertices. By adapting
 842 the Path_P algorithm on multitrees, we get sublinear time reachability oracles for multitrees
 843 using subquadratic space and subquadratic preprocessing time. If the graph is a multitree of
 844 strongly connected components, we can first treat each strongly connected component as a
 845 single vertex, whose weighted size is the total weighted size of all vertices in the component.

846 The reachability oracle for multitrees can be adapted directly from the Path_P algorithm.
 847 In the recursion tree of calling Path_P , we take down the final subproblem that each vertex
 848 belongs to, and when querying a pair of vertex, we find the Path_P instance corresponding to
 849 the least common ancestor of the two final subproblems corresponding to these vertices, and
 850 consider the CutPath_P process called by this Path_P instance.

851 Next we modify the CutPath_P algorithm. Among all pivot vertices, we call the ones
 852 who have no other pivot vertices as descendants “sink pivot vertices”. After computing the
 853 number of ancestors and descendants for all vertices, we can decide if a vertex is a sink in
 854 time linear to the degree of the vertex.

855 We create another graph G' . Similar to CutPath_P , we keep finding pivot vertices who
 856 have at least M^α ancestors and M^α descendants in the remaining graph, and then remove
 857 them. Whenever finding a pivot vertex v , we create edges from all its ancestors to v , and
 858 from v to all its descendants in G' .

859 Thus, when querying a pair of vertices a, b , if a can reach b , there are three cases:

- 860 ■ Case 1: b is a pivot vertex. Then there is an edge from a to b in G' .
- 861 ■ Case 2: The path from a to b goes through at least a pivot vertex. In this case, it must
 862 go through a sink pivot vertex. We decide if there is a sink pivot vertex v adjacent to
 863 a in G' who is also adjacent to b . Each vertex can be an ancestors of at most M/M^α
 864 sink pivot vertices, because each sink pivot vertex has more than M^α descendants, and
 865 different sink pivot vertices have disjoint set of descendants. So this checking can be done
 866 in time M/M^α .
- 867 ■ Case 3: The path from a to b does not go through any pivot vertex. Then we can do
 868 a DFS starting from a that traverses through at most M^α of a 's descendants in the
 869 remaining graph G to find b . The time taken is $O(M^\alpha)$.

870 Thus, the query time is $O(M/M^\alpha + M^\alpha)$, which is sublinear to M .

871 If the graph is a multitree, not a multitree of strongly connected components, then every
 872 vertex has unit weighted size. In this case, the modified CutPath_P process runs in time
 873 $O(m^2/m^\alpha + m^{2-\alpha} + m^{1+2\alpha})$, because now we do not use Selection_P thus the function $t(m)$ is
 874 $O(m)$, and there are no large-size vertices thus we can pick $\gamma = 0$. If we choose $\alpha = 1/3$, then
 875 the running time is $O(m^{5/3})$. The modified Path_P algorithm has running time satisfying the

⁸⁷⁶ recursion $T(m) = 2T(m/2) + O(m^{5/3})$, which is $O(m^{5/3})$. So the preprocessing time and
⁸⁷⁷ space is $O(m^{5/3})$, and the query time is $O(m^{2/3})$.