<https://helda.helsinki.fi>

Using Minimum Path Cover to Boost Dynamic Programming on DAGs : Co-linear Chaining Extended

Kuosmanen, Anna

Springer International Publishing AG
2018-04

Kuosmanen , A , Paavilainen , T , Gagie , T , Chikhi , R , Tomescu , A & Mäkinen , V 2018 ,
Using Minimum Path Cover to Boost Dynamic Programming on DAGs : Co-linear Chaining
Extended . in B J Raphael (ed.) , Research in Computational Molecular Biology : 22nd
Annual International Conference, RECOMB 2018, Paris, France, April 21-24, 2018,
Proceedings . Lecture Notes in Bioinformatics , vol. 10812 , Springer International Publishing
AG , Cham , pp. 105-121 , Annual International Conference on Research in Computational
Molecular Biology , Paris , France , 21/04/2018 . https://doi.org/10.1007/978-3-319-89929-9_7

<http://hdl.handle.net/10138/307569>

https://doi.org/10.1007/978-3-319-89929-9_7

unspecified

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Using Minimum Path Cover to Boost Dynamic Programming on DAGs: Co-Linear Chaining Extended

Anna Kuosmanen¹, Topi Paavilainen¹, Travis Gagie², Rayan Chikhi³, Alexandru Tomescu^{1,*}, and Veli Mäkinen^{1*,**}

¹ Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland

² Diego Portales University, Chile

³ CNRS, CRISAL, University of Lille 1, France

Abstract. Aligning sequencing reads on graph representations of genomes is an important ingredient of *pan-genomics* (Marshall et al. *Briefings in Bioinformatics*, 2016). Such approaches typically find a set of *local anchors* that indicate plausible matches between substrings of a read to subpaths of the graph. These anchor matches are then combined to form a (semi-local) alignment of the complete read on a subpath. *Co-linear chaining* is an algorithmically rigorous approach to combine the anchors. It is a well-known approach for the case of two sequences as inputs. Here we extend the approach so that one of the inputs can be a directed acyclic graph (DAGs), e.g. a *splicing graph* in transcriptomics or *variant graph* in pan-genomics.

The extension of co-linear chaining to DAGs turns out to have a tight connection to the *minimum path cover* problem that asks us to find a minimum-cardinality set of paths that cover all the nodes of a DAG. We study the case when the size k of a minimum path cover is small, which is often the case in practice. First, we propose an algorithm for finding a minimum path cover of a DAG (V, E) in $O(k|E| \log |V|)$ time, improving all known time-bounds when k is small and the DAG is not too dense. An immediate consequence is an improved space/time tradeoff for reachability queries in arbitrary directed graphs. Second, we introduce a general technique for extending dynamic programming (DP) algorithms from sequences to DAGs. This is enabled by our minimum path cover algorithm, and works by mimicking the DP algorithm for sequences on each path of the minimum path cover. This technique generally produces algorithms that are slower than their counterparts on sequences only by a factor k . We illustrate this on two classical problems extended to labeled DAGs: longest increasing subsequence and longest common subsequence. For the former we obtain an algorithm with running time $O(k|E| \log |V|)$. This matches the optimal solution to the classical problem variant when, e.g., the input sequence is modeled as a path. We obtain an analogous result for the longest common subsequence problem. Finally, we apply this technique to the co-linear chaining problem, that is a generalization of both of the above two problems. The algorithm for this problem turns out to be more involved, needing further ingredients, such as an FM-index tailored for large alphabets, and a two-dimensional range search tree modified to support range maximum queries. We implemented the new co-linear chaining approach. Experiments on splicing graphs show that the new method is efficient also in practice.

1 Introduction

A *path cover* of a DAG $G = (V, E)$ is a set of paths such that every node of G belongs to some path. A *minimum path cover* (MPC) is one having the minimum number of paths. The size of a MPC is also called the *width* of G . Many DAGs commonly used in genome research, such as graphs encoding human mutations [9] and graphs modeling gene transcripts [18], can consist, in the former case, of millions of nodes and, in the latter case, of thousands of nodes. However, they generally have a small width on average; for example, splicing graphs for most genes in human chromosome 2 have width at most 10 [40, Fig. 7]. To the best of our knowledge, among the many MPC algorithms [15, 20, 35, 31, 7, 8], there are only three whose complexities depends on the width of the DAG. Say the width of G is k . The first algorithm runs in time $O(|V||E| + k|V|^2)$ and can be obtained by slightly modifying an algorithm for finding a minimum chain cover

* Shared last author contribution

** Corresponding author: veli.makinen@helsinki.fi

in partial orders from [12]. The other two algorithms are due to Chen and Chen: the first one works in time $O(|V|^2 + k\sqrt{k}|V|)$ [7], and the second one works in time $O(\max(\sqrt{|V||E|}, k\sqrt{k}|V|))$ [8].

In this paper we present an MPC algorithm running in time $O(k|E|\log|V|)$. For example, for $k = o(\sqrt{|V|}/\log|V|)$ and $|E| = O(|V|^{3/2})$, this is better than all previous algorithms. Our algorithm is based on the following standard reduction of a minimum flow problem to a maximum flow problem (see e.g. [2]): (i) find a feasible flow/path cover satisfying all demands, and (ii) solve a maximum flow problem in a graph encoding how much flow can be removed from every edge. Our main insight is to solve step (i) by finding an approximate solution that is greater than the optimal one only by a $O(\log|V|)$ factor. Then, if we solve step (ii) with the Ford-Fulkerson algorithm, the number of iterations can be bounded by $O(k\log|V|)$.

We then proceed to show that some problems (like pattern matching) that admit efficient *sparse dynamic programming* solutions on sequences [11] can be extended to DAGs, so that their complexity increases only by the minimum path cover size k . Extending pattern matching to DAGs has been studied before [32,3,28]. For those edit distance -based formulations our approach does not yield an improvement, but on formulations involving sparse set of matching anchors [11] we can boost the naive solutions of their DAG extensions by exploiting a path cover. Namely, our improvement applies to many cases where a data structure over previously computed solutions is maintained and queried for computing the next value. Our new MPC algorithm enables this, as its complexity is generally of the same form as that of solving the extended problems. Given a path cover, our technique then computes so-called *forward propagation links* indicating how the partial solutions in each path in the cover must be synchronized.

To best illustrate the versatility of the technique itself, we show (in the Appendix) how to compute a longest increasing subsequence (LIS) in a labeled DAG, in time $O(k|E|\log|V|)$. This matches the optimal solution to the classical problem on a single sequence when, e.g., this is modeled as a path (where $k = 1$). We also illustrate our technique with the longest common subsequence (LCS) problem between a labeled DAG $G = (V, E)$ and a sequence S .

Finally, we consider the main problem of this paper—co-linear chaining (CLC)—first introduced in [27]. It has been proposed as a model of the sequence alignment problem that scales to massive inputs, and has been a subject of recent interest (see e.g. [33,41,43,44,45,26,36]). In the CLC problem, the input is directly assumed to be a set of N pairs of intervals in the two sequences that match (either exactly or approximately). The CLC alignment solution asks for a subset of these plausible pairs that maximizes the coverage in one of the sequences, and whose elements appear in increasing order in both sequences. The fastest algorithm for this problem runs in the optimal $O(N\log N)$ time [1].

We define a generalization of the CLC problem between a sequence and a labeled DAG. As motivation, we mention the problem of aligning a long sequence, or even an entire chromosome, inside a DAG storing all known mutations of a population with respect to a reference genome (such as the above-mentioned [9], or more specifically a linearized version of it [17]). Here, the N input pairs match intervals in the sequence with paths (also called *anchors*) in the DAG. This problem is not straightforward, as the topological order of the DAG might not follow the reachability order between the anchors. Existing tools for aligning DNA sequences to DAGs (BGREAT [24], vg [29]) rely on anchors but do not explicitly consider solving CLC optimally on the DAG.

The algorithm we propose uses the general framework mentioned above. Since it is more involved, we will develop it in stages. We first give a simple approach to solve a relaxed co-linear chaining problem using $O((|V| + |E|)N)$ time, then introduce the MPC approach that requires $O(k|E|\log|V| + kN\log N)$ time. As above, if the DAG is a labeled path representing a sequence, the running time of our algorithm is reduced to the best current solution for the co-linear chaining problem on sequences, $O(N\log N)$ [1]. We conclude (in the Appendix) with a Burrows-Wheeler technique to efficiently handle a special case that we omitted in this relaxed variant. We remark that one can reduce the LIS and LCS problems to the CLC problem to obtain the same running time bounds as mentioned earlier; these are given for the sake of comprehensiveness.

In the last section we discuss the anchor-finding preprocessing step. We implemented the new MPC-based co-linear chaining algorithm and conducted experiments on splicing graphs to show that the approach is practical, once anchors are given. Some future directions on how to incorporate practical anchors, and how to apply the techniques to transcript prediction, are discussed.

Notation. To simplify notation, for any DAG $G = (V, E)$ we will assume that V is always $\{1, \dots, |V|\}$ and that $1, \dots, |V|$ is a topological order on V (so that for every edge (u, v) we have $u < v$). We will also assume

that $|E| \geq |V| - 1$. A *labeled DAG* is a tuple (V, E, ℓ, Σ) where (V, E) is a DAG and $\ell : V \mapsto \Sigma$ assign to the nodes labels from Σ , Σ being an ordered alphabet.

For a node $v \in V$, we denote by $N^-(v)$ the set of in-neighbors of v and by $N^+(v)$ the set of out-neighbors of v . If there is a (possibly empty) path from node u to node v we say that u reaches v . We denote by $R^-(v)$ the set of nodes that reach v . We denote a set of consecutive integers with interval notation $[i..j]$, meaning $\{i, i+1, \dots, j\}$. For a pair of intervals $m = ([x..y], [c..d])$, we use $m.x$, $m.y$, $m.c$, and $m.d$ to denote the four respective endpoints. We also consider pairs of the form $m = (P, [c..d])$ where P is a path, and use $m.P$ to access P . The first node of P will be called its *startpoint*, and its last node will be called its *endpoint*. For a set M we may fix an order, to access an element as $M[i]$.

2 The MPC algorithm

In this section we assume basic familiarity with network flow concepts; see [2] for further details. In the *minimum flow problem*, we are given a directed graph $G = (V, E)$ with a single source and a single sink, with a *demand* $d : E \rightarrow \mathbb{Z}$ for every edge. The task is to find a flow of minimum value (the *value* is the sum of the flow on the edges exiting the source) that satisfies all demands (to be called *feasible*). The standard reduction from the minimum path cover problem to a minimum flow one (see, e.g. [30]) creates a new DAG G^* by replacing each node v with two nodes v^-, v^+ , adds the edge (v^-, v^+) and adds all in-neighbors of v as in-neighbors of v^- , and all out-neighbors of v as out-neighbors of v^+ . Finally, the reduction adds a global source with an out-going edge to every node, and a global sink with an in-coming edge from every node. Edges of type (v^-, v^+) get demand 1, and all other edges get demand 0. The value of the minimum flow equals k , the width of G , and any decomposition of it into source-to-sink paths induces a minimum path cover in G .

Our MPC algorithm is based on the following simple reduction of a minimum flow problem to a maximum flow one (see e.g. [2]): (i) find a feasible flow $f : E \rightarrow \mathbb{Z}$; (ii) transform this into a minimum feasible flow, by finding a maximum flow f' in G in which every $e \in E$ now has capacity $f(e) - d(e)$. The final minimum flow solution is obtained as $f(e) - f'(e)$, for every $e \in E$. Observe that this path cover induces a flow of value $O(k \log |V|)$. Thus, in step (ii) we need to shrink this flow into a flow of value k . If we run the Ford-Fulkerson algorithm, this means that there are $O(k \log |V|)$ successive augmenting paths, each of which can be found in time $O(E)$. This gives a time bound for step (ii) of $O(k|E| \log |V|)$.

We solve step (i) in time $O(k|E| \log |V|)$ by finding a path cover in G^* whose size is larger than k only by a relative factor $O(\log |V|)$. This is based on the classical greedy set cover algorithm, see e.g. [42, Chapter 2]: at each step, select a path covering most of the remaining uncovered nodes.

This approach is similar to the one from [12] for finding the minimum number k of chains to cover a partial order of size n . A *chain* is a set of pairwise comparable elements. The algorithm from [12] runs in time $O(kn^2)$, and it has the same feature as ours: it first finds a set of $O(k \log n)$ chains in the same way as us (longest chains covering most uncovered elements), and then in a second step reduces these to k . However, if we were to apply this algorithm to DAGs, it would run in time $O(|V||E| + k|V|^2)$, which is slower than our algorithm for small k . This is because it uses the classical reduction given by Fulkerson [15] to a bipartite graph, where each edge of the graph encodes a pair of elements in the relation. Since DAGs are not transitive in general, to use this reduction one needs first to compute the transitive closure of the DAG, in time $O(|V||E|)$. Such approximation-refinement approach has also been applied to other covering problems on graphs, such as a 2-hop cover [10].

We now show how to solve step (i) within the claimed running time, by dynamic programming.

Lemma 1. *Let $G = (V, E)$ be a DAG, and let k be the width of G . In time $O(k|E| \log |V|)$, we can compute a path cover $P_1, \dots, P_K$ of G , such that $K = O(k \log |V|)$.*

Proof. The algorithm works by choosing, at each step, a path that covers the most uncovered nodes. For every node $v \in V$, we store $\mathbf{m}[v] = 1$, if v is not covered by any path, and $\mathbf{m}[v] = 0$ otherwise. We also store $\mathbf{u}[v]$ as the largest number of uncovered nodes on a path starting at v . The values $\mathbf{u}[\cdot]$ are computed by dynamic programming, by traversing the nodes in inverse topological order and setting $\mathbf{u}[v] = \mathbf{m}[v] + \max_{w \in N^+(v)} \mathbf{u}[w]$. Initially we have $\mathbf{m}[v] = 1$ for all v . We then compute $\mathbf{u}[v]$ for all v , in time $O(|E|)$. By taking the node v with the maximum $\mathbf{u}[v]$, and tracing back along the optimal path starting at v , we obtain our first path in

Fig. 1. A path cover P_1, P_2, P_3 of a DAG. The forward links entering v from $\text{last2reach}[v, i]$ are shown with dotted black lines, for $i \in \{1, 2, 3\}$. We mark in gray the set $R^-(v)$ of nodes that reach v .

time $O(|E|)$. We then update $\mathfrak{m}[v] = 0$ for all nodes on this path, and iterate this process until all nodes are covered. This takes overall time $O(K|E|)$, where K is the number of paths found.

This algorithm analysis is identical to the one of the classical greedy set cover algorithm [42, Chapter 2], because the universe to be covered is V and each possible path in G is a possible covering set, which implies that $K = O(k \log |V|)$. $\square$

Combining Lemma 1 with the above-mentioned application of the Ford-Fulkerson algorithm, we obtain our first result:

Theorem 1. *Given a DAG $G = (V, E)$ of width k , the MPC problem on G can be solved in time $O(k|E| \log |V|)$.*

3 The dynamic programming framework

In this section we give an overview of the main ideas of our approach.

Suppose we have a problem involving DAGs that is solvable, for example by dynamic programming, by traversing the nodes in topological order. Thus, assume also that a partial solution at each node v is obtainable from all (and only) nodes of the DAG that can reach v , plus some other independent objects, such as another sequence. Furthermore, suppose that at each node v we need to query (and maintain) a data structure $\mathcal{T}$ that depends on $R^-(v)$ and such that the answer $\text{Query}(R^-(v))$ at v is decomposable as:

$$\text{Query}(R^-(v)) = \bigoplus_i \text{Query}(R_i^-(v)). \quad (1)$$

In the above, the sets $R_i^-(v)$ are such that $R^-(v) = \bigcup_i R_i^-(v)$, they are not necessarily disjoint, and $\bigoplus$ is some operation on the queries, such as min or max, that does not assume disjointness. It is understood that after the computation at v , we need to update $\mathcal{T}$. It is also understood that once we have updated $\mathcal{T}$ at v , we cannot query $\mathcal{T}$ for a node before v in topological order, because it would give an incorrect answer.

The first idea is to decompose the graph into a path cover $P_1, \dots, P_K$. As such, we decompose the computation only along these paths, in light of (1). We replace a single data structure $\mathcal{T}$ with K data structures $\mathcal{T}_1, \dots, \mathcal{T}_K$, and perform the operation from (1) on the results of the queries to these K data structures.

Our second idea concerns the order in which the nodes on these K paths are processed. Because the answer at v depends on $R^-(v)$, we cannot process the nodes on the K paths (and update the corresponding $\mathcal{T}_i$'s) in an arbitrary order. As such, for every path i and every node v , we distinguish the *last node* on path i that reaches v (if it exists). We will call this node $\text{last2reach}[v, i]$. See Figure 1 for an example. We note that this insight is the same as in [21], which symmetrically identified the *first* node on a chain i that can be reached from v (a *chain* is a subsequence of a path). The following observation is the first ingredient for using the decomposition (1).

Observation 1 *Let $P_1, \dots, P_K$ be a path cover of a DAG G , and let $v \in V(G)$. Let R_i denote the set of nodes of P_i from its beginning until $\text{last2reach}[v, i]$ inclusively (or the empty set, if $\text{last2reach}[v, i]$ does not exist). Then $R^-(v) = \bigcup_{i=1}^K R_i$.*

Proof. It is clear that $\bigcup_{i=1}^K R_i \subseteq R^-(v)$. To show the reverse inclusion, consider a node $u \in R^-(v)$. Since $P_1, \dots, P_K$ is a path cover, then u appears on some P_i . Since u reaches v , then u appears on P_i before $\text{last2reach}[v, i]$, or $u = \text{last2reach}[v, i]$. Therefore u appears on R_i , as desired. $\square$

This allows us to identify, for every node u , a set of *forward propagation links* $\text{forward}[u]$, where $(v, i) \in \text{forward}[u]$ holds for any node v and index i with $\text{last2reach}[v, i] = u$. These propagation links are the second ingredient in the correctness of the decomposition. Once we have computed the correct value at u , we update the corresponding data structures $\mathcal{T}_i$ for all paths i to which u belongs. We also propagate the query value of $\mathcal{T}_i$ in the decomposition (1) for all nodes v with $(v, i) \in \text{forward}[u]$. This means that when we come to process v , we have already correctly computed all terms in the decomposition (1) and it suffices to apply the operation $\oplus$ to these terms.

The next lemma shows how to compute the values last2reach (and, as a consequence, all forward propagation links), also by dynamic programming.

Lemma 2. *Let $G = (V, E)$ be a DAG, and let $P_1, \dots, P_K$ be a path cover of G . For every $v \in V$ and every $i \in [1..K]$, we can compute $\text{last2reach}[v, i]$ in overall time $O(K|E|)$.*

Proof. For each P_i and every node v on P_i , let $\text{index}[v, i]$ denote the position of v in P_i (say, starting from 1). Our algorithm actually computes $\text{last2reach}[v, i]$ as the index of this node in P_i . Initially, we set $\text{last2reach}[v, i] = -1$ for all v and i . At the end of the algorithm, $\text{last2reach}[v, i] = -1$ will hold precisely for those nodes v that cannot be reached by any node of P_i . We traverse the nodes in topological order. For every $i \in [1..K]$, we do as follows: if v is on P_i , then we set $\text{last2reach}[v, i] = \text{index}[v, i]$. Otherwise, we compute by dynamic programming $\text{last2reach}[v, i]$ as $\max_{u \in N^-(v)} \text{last2reach}[u, i]$. $\square$

An immediate application of Theorem 1 and of the values $\text{last2reach}[v, i]$ is for solving reachability queries (see Appendix A.1). Another simple application is an extension of the *longest increasing subsequence (LIS)* problem to labeled DAGs (Appendix A.2).

The LIS problem, the LCS problem of Section 4, as well as the CLC problem of Section 5 make use of the following standard data structure (see e.g. [25, p.20]).

Lemma 3. *The following two operations can be supported with a balanced binary search tree $\mathcal{T}$ in time $O(\log n)$, where n is the number of leaves in the tree.*

- $\text{update}(k, \text{val})$: For the leaf w with $\text{key}(w) = k$, update $\text{value}(w) = \text{val}$.
- $\text{RMaxQ}(l, r)$: Return $\max_{w : l \leq \text{key}(w) \leq r} \text{value}(w)$ (Range Maximum Query).

Moreover, the balanced binary search tree can be built in $O(n)$ time, given the n pairs $(\text{key}, \text{value})$ sorted by component key .

4 The LCS problem

Consider a labeled DAG $G = (V, E, \ell, \Sigma)$ and a sequence $S \in \Sigma^*$, where Σ is an ordered alphabet. We say that the *longest common subsequence (LCS)* between G and S is a longest subsequence C of any path label in G such that C is also a subsequence of S .

We will modify the LIS algorithm of Appendix A.2 minimally to find a LCS between a DAG G and a sequence S . The description is self-contained yet, for the interest of page limit, more dense than the LIS algorithm derivation. The purpose is to give an example of the general MPC-framework with fewer technical details than required in the main result of this paper concerning co-linear chaining.

For any $c \in \Sigma$, let $S(c)$ denote set $\{j \mid S[j] = c\}$. For each node v and each $j \in S(\ell(v))$, we aim to store in $\text{LLCS}[v, j]$ the length of the longest common subsequence between $S[1..j]$ and any label of path ending at v , among all subsequences having $\ell(v) = S[j]$ as the last symbol.

Assume we have a path cover of size K and $\text{forward}[u]$ computed for all $u \in V$. Assume also we have mapped Σ to $\{0, 1, 2, \dots, |S| + 1\}$ in $O((|V| + |S|) \log |S|)$ time (e.g. by sorting the symbols of S , binary searching labels of V , and then relabeling by ranks, with the exception that, if a node label does not appear in S , it is replaced by $|S| + 1$).

Fig. 2. In the co-linear chaining problem between two sequences T and R , we need to find a subset of pairs of intervals (i.e., anchors) so that (i) the selected intervals in each sequence appear in increasing order; and (ii) the selected intervals cover in R the maximum amount of positions. The figure shows an input for the problem, and highlights in gray an optimal subset of anchors. Figure taken from [25].

Let $\mathcal{T}_i$ be a search tree of Lemma 3 initialized with key-value pairs $(0, 0), (1, -\infty), (2, -\infty), \dots, (|S|, -\infty)$, for each $i \in [1..K]$. The algorithm proceeds in fixed topological ordering on G . At a node u , for every $(v, i) \in \text{forward}[u]$ we now update an array $\text{LLCS}[v, j]$ for all $j \in S(\ell(v))$ as follows: $\text{LLCS}[v, j] = \max(\text{LLCS}[v, j], \mathcal{T}_i.\text{RMaxQ}(0, j - 1) + 1)$. The update step of $\mathcal{T}_i$ when the algorithm reaches a node v , for each covering path i containing v , is done as $\mathcal{T}_i.\text{update}(j', \text{LLCS}[v, j'])$ for all j' with $j' < j$ and $j' \in S(\ell(v))$. Initialization is handled by the $(0, 0)$ key-value pair so that any (v, j) with $\ell(v) = S[j]$ can start a new common subsequence.

The final answer to the problem is $\max_{v \in V, j \in S(\ell(v))} \text{LLCS}[v, j]$, with the actual LCS to be found with a standard traceback. The algorithm runs in $O((|V| + |S|) \log |S| + K|M| \log |S|)$ time, where $M = \{(v, j) \mid v \in V, j \in [1..|S|], \ell(v) = S[j]\}$, and assuming a cover of K paths is given. Notice that $|M|$ can be $\Omega(|V||S|)$. With Theorem 1 plugged in, the total running time becomes $O(k|E| \log |V| + (|V| + |S|) \log |S| + k|M| \log |S|)$. Since the queries on the data structures are semi-open, one can use the more efficient data structure from [16] to improve the bound to $O(k|E| \log |V| + (|V| + |S|) \log |S| + k|M| \log \log |S|)$. The following theorem summarizes this result.

Theorem 2. Let $G = (V, E, \ell, \Sigma)$ be a labeled DAG of width k , and let $S \in \Sigma^*$, where Σ is an ordered alphabet. We can find a longest common subsequence between G and S in time $O(k|E| \log |V| + (|V| + |S|) \log |S| + k|M| \log \log |S|)$.

When G is a path, the bound improves to $O((|V| + |S|) \log |S| + |M| \log \log |S|)$, which nearly matches the fastest sparse dynamic programming algorithm for the LCS on two sequences [11] (with a difference in log log-factor due to a different data structure, which does not work for this order of computation).

5 Co-linear chaining

We start with a formal definition of the co-linear chaining problem (see Figure 2 for an illustration), following the notions introduced in [25, Section 15.4].

Problem 1 (Co-linear chaining (CLC)). Let T and R be two sequences over an alphabet Σ , and let M be a set of N pairs $([x..y], [c..d])$. Find an ordered subset $S = s_1 s_2 \dots s_p$ of pairs from M such that

- $s_{j-1}.y < s_j.y$ and $s_{j-1}.d < s_j.d$, for all $1 \leq j \leq p$, and
- S maximizes the *ordered coverage* of R , defined as

$$\text{coverage}(R, S) = |\{i \in [1..|R|] \mid i \in [s_j.c..s_j.d] \text{ for some } 1 \leq j \leq p\}|.$$

The definition of ordered coverage between two sequences is symmetric, as we can simply exchange the roles of T and R . But when solving the CLC problem between a DAG and a sequence, we must choose whether we want to maximize the ordered coverage on the sequence R or on the DAG G . We will consider the former variant.

First, we define the following *precedence relation*:

Definition 1. Given two paths P_1 and P_2 in a DAG G , we say that P_1 precedes P_2 , and write $P_1 \prec P_2$, if one of the following conditions holds:

- P_1 and P_2 do not share nodes and there is a path in G from the endpoint of P_1 to the startpoint of P_2 , or
- P_1 and P_2 have a suffix-prefix overlap and P_2 is not fully contained in P_1 ; that is, if $P_1 = (a_1, \dots, a_i)$ and $P_2 = (b_1, \dots, b_j)$ then there exists a $k \in \{\max(1, 2 + i - j), \dots, i\}$ such that $a_k = b_1, a_{k+1} = b_2, \dots, a_i = b_{1+i-k}$.

We then extend the formulation of Problem 1 to handle a sequence and a DAG.

Problem 2 (CLC between a sequence and a DAG). Let R be a sequence, let G be a labeled DAG, and let M be a set of N pairs $(P, [c..d])$, where P is a path in G and $c \leq d$ are non-negative integers. Find an ordered subset $S = s_1 s_2 \dots s_p$ of pairs from M such that

- for all $2 \leq j \leq p$, it holds that $s_{j-1}.P \prec s_j.P$ and $s_{j-1}.d < s_j.d$, and
- S maximizes the *ordered* coverage of R , analogously defined as $\text{coverage}(R, S) = |\{i \in [1..|R|] \mid i \in [s_j.c..s_j.d] \text{ for some } 1 \leq j \leq p\}|$.

To illustrate the main technique of this paper, let us for now only seek solutions where paths in consecutive pairs in a solution do not overlap in the DAG. Suffix-prefix overlaps between paths turn out to be challenging; we will postpone this case until Appendix B.

Problem 3 (Overlap-limited CLC between a sequence and a DAG). Let R be a sequence, let G be a labeled DAG, and let M be a set of N pairs $(P, [c..d])$, where P is a path in G and $c \leq d$ are non-negative integers (with the interpretation that $\ell(P)$ matches $R[c..d]$). Find an ordered subset $S = s_1 s_2 \dots s_p$ of pairs from M such that

- for all $2 \leq j \leq p$, it holds that there is a non-empty path from the last node of $s_{j-1}.P$ to the first node of $s_j.P$ and $s_{j-1}.d < s_j.d$, and
- S maximizes $\text{coverage}(R, S)$.

First, let us consider a trivial approach to solve Problem 3. Assume we have ordered in $O(|E| + N)$ time the N input pairs as $M[1], M[2], \dots, M[N]$, so that the endpoints of $M[1].P, M[2].P, \dots, M[N].P$ are in topological order, breaking ties arbitrarily. We denote by $C[j]$ the maximum ordered coverage of $R[1..M[j].d]$ using the pair $M[j]$ and any subset of pairs from $\{M[1], M[2], \dots, M[j-1]\}$.

Theorem 3. *Overlap-limited co-linear chaining between a sequence and a labeled DAG $G = (V, E, \ell, \Sigma)$ (Problem 3) on N input pairs can be solved in $O((|V| + |E|)N)$ time.*

Proof. First, we reverse the edges of G . Then we mark the nodes that correspond to the path endpoints for every pair. After this preprocessing we can start computing the maximum ordered coverage for the pairs as follows: for every pair $M[j]$ in topological order of their path endpoints for $j \in \{1, \dots, N\}$ we do a depth-first traversal starting at the startpoint of path $M[j].P$. Note that since the edges are reversed, the depth-first traversal checks only pairs whose paths are predecessors of $M[j].P$.

Whenever we encounter a node that corresponds to the path endpoint of a pair $M[j']$, we first examine whether it fulfills the criterion $M[j'].d < M[j].c$ (call this case (a)). The best ordered coverage using pair $M[j]$ after all such $M[j']$ is then

$$C^a[j] = \max_{j' : M[j'].d < M[j].c} \{C[j'] + (M[j].d - M[j].c + 1)\}, \quad (2)$$

where $C[j']$ is the best ordered coverage when using pairs $M[j']$ last.

If pair $M[j']$ does not fulfill the criterion for case (a), we then check whether $M[j].c \leq M[j'].d \leq M[j].d$ (call this case (b)). The best ordered coverage using pair $M[j]$ after all such $M[j']$ with $M[j'].c < M[j].c$ is then

$$C^b[j] = \max_{j' : M[j].c \leq M[j'].d \leq M[j].d} \{C[j'] + (M[j].d - M[j'].d)\}. \quad (3)$$

Inclusions, i.e. $M[j].c \leq M[j'].c$, can be left computed incorrectly in $C^b[j]$, since there is a better or equally good solution computed in $C^a[j]$ or $C^b[j]$ that does not use them [1].

Finally, we take $C[j] = \max(C^a[j], C^b[j])$. Depth-first traversal takes $O(|V| + |E|)$ time and is executed N times, for $O((|V| + |E|)N)$ total time. $\square$

However, we can do significantly better than $O((|V| + |E|)N)$ time. In the next sections we will describe how to apply the framework from Section 3 here.

5.1 Co-linear chaining on sequences revisited

We now describe the dynamic programming algorithm from [1] for the case of two sequences, as we will then reuse this same algorithm in our MPC approach.

First, sort input pairs in M by the coordinate y into the sequence $M[1], M[2], \dots, M[N]$, so that $M[i].y \leq M[j].y$ holds for all $i < j$. This will ensure that we consider the overlapping ranges in sequence T in the correct order. Then, we fill a table $C[1..N]$ analogous to that of Theorem 3 so that $C[j]$ gives the maximum ordered coverage of $R[1..M[j].d]$ using the pair $M[j]$ and any subset of pairs from $\{M[1], M[2], \dots, M[j-1]\}$. Hence, $\max_j C[j]$ gives the total maximum ordered coverage of R .

Consider Equations (2) and (3). Now we can use an *invariant technique* to convert these recurrence relations so that we can exploit the range maximum queries of Lemma 3:

$$\begin{aligned} C^a[j] &= (M[j].d - M[j].c + 1) + \max_{j' : M[j'].d < M[j].c} C[j'] \\ &= (M[j].d - M[j].c + 1) + \mathcal{T}.\text{RMaxQ}(0, M[j].c - 1), \\ C^b[j] &= M[j].d + \max_{j' : M[j].c \leq M[j'].d \leq M[j].d} \{C[j'] - M[j'].d\} \\ &= M[j].d + \mathcal{I}.\text{RMaxQ}(M[j].c, M[j].d), \\ C[j] &= \max(C^a[j], C^b[j]). \end{aligned}$$

For these to work correctly, we need to have properly updated the trees $\mathcal{T}$ and $\mathcal{I}$ for all $j' \in [1..j-1]$. That is, we need to call $\mathcal{T}.\text{update}(M[j'].d, C[j'])$ and $\mathcal{I}.\text{update}(M[j'].d, C[j'] - M[j'].d)$ after computing each $C[j']$. The running time is $O(N \log N)$.

Figure 2 illustrates the optimal chain on our schematic example. This chain can be extracted by modifying the algorithm to store traceback pointers.

Theorem 4 ([36,1]). *Problem 1 on N input pairs can be solved in the optimal $O(N \log N)$ time.*

5.2 Co-linear chaining on DAGs using a minimum path cover

Let us now modify the above algorithm to work with DAGs, using the main technique of this paper.

Theorem 5. *Problem 3 on a labeled DAG $G = (V, E, \ell, \Sigma)$ of width k and a set of N input pairs can be solved in time $O(k|E| \log |V| + kN \log N)$ time.*

Proof. Assume we have a path cover of size K and $\text{forward}[u]$ computed for all $u \in V$. For each path $i \in [1..K]$, we create two binary search trees $\mathcal{T}_i$ and $\mathcal{I}_i$. As a reminder, these trees correspond to coverages for pairs that do not, and do overlap, respectively, on the sequence. Moreover, recall that in Problem 3 we do not consider solutions where consecutive paths in the graph overlap.

As keys, we use $M[j].d$, for every pair $M[j]$, and additionally the key 0. The value of every key is initialized to $-\infty$.

After these preprocessing steps, we process the nodes in topological order, as detailed in Algorithm 1. If node v corresponds to the endpoint of some $M[j].P$, we update the trees $\mathcal{T}_i$ and $\mathcal{I}_i$ for all covering paths i containing node v . Then we follow all forward propagation links $(w, i) \in \text{forward}[v]$ and update $C[j]$ for each path $M[j].P$ starting at w , taking into account all pairs whose path endpoints are in covering path i . Before the main loop visits w , we have processed all forward propagation links to w , and the computation of $C[j]$ has taken all previous pairs into account, as in the naive algorithm, but now indirectly through the

Algorithm 1: Co-linear chaining between a sequence and a DAG using a path cover.

Input: DAG $G = (V, E)$, a path cover $P_1, P_2, \dots, P_K$ of G , and N pairs $M[1], M[2], \dots, M[N]$ of the form $(P, [c..d])$.

Output: The index j giving $\max_j C[j]$.

Use Lemma 2 to find all forward propagation links;

```

for  $i \leftarrow 1$  to  $K$  do
    Initialize search trees  $\mathcal{T}_i$  and  $\mathcal{I}_i$  with keys  $M[j].d$ ,  $1 \leq j \leq N$ , and with key 0, all keys associated with
    values  $-\infty$ ;
     $\mathcal{T}_i.\text{update}(0, 0)$ ;
     $\mathcal{I}_i.\text{update}(0, 0)$ ;
    /* Save to  $\text{start}[i]$  (respectively,  $\text{end}[i]$ ) the indexes of all pairs whose path starts
    (respectively, ends) at  $i$ . */
    for  $j \leftarrow 1$  to  $N$  do
         $\text{start}[M[j].P.\text{first}].\text{push}(j)$ ;
         $\text{end}[M[j].P.\text{last}].\text{push}(j)$ ;
    for  $v \in V$  in topological order do
        for  $j \in \text{end}[v]$  do
            /* Update the search trees for every path that covers  $v$ , stored in  $\text{paths}[v]$ . */
            for  $i \in \text{paths}[v]$  do
                 $\mathcal{T}_i.\text{update}(M[j].d, C[j])$ ;
                 $\mathcal{I}_i.\text{update}(M[j].d, C[j] - M[j].d)$ ;
            for  $(w, i) \in \text{forward}[v]$  do
                for  $j \in \text{start}[w]$  do
                     $C^a[j] \leftarrow (M[j].d - M[j].c + 1) + \mathcal{T}_i.\text{RMaxQ}(0, M[j].c - 1)$ ;
                     $C^b[j] \leftarrow M[j].d + \mathcal{I}_i.\text{RMaxQ}(M[j].c, M[j].d)$ ;
                     $C[j] \leftarrow \max(C[j], C^a[j], C^b[j])$ ;
    return  $\text{argmax}_j C[j]$ ;
```

K search trees. Exceptions are the pairs overlapping in the graph, which we omit in this problem statement. The forward propagation ensures that the search tree query results are indeed taking only reachable pairs into account. While $C[j]$ is already computed when visiting w , the startpoint of $M[j].P$, the added coverage with the pair is updated to the search trees only when visiting the endpoint.

There are NK forward propagation links, and both search trees are queried in $O(\log N)$ time. All the search trees containing a path endpoint of a pair are updated. Each endpoint can be contained in at most K paths, so this also gives the same bound $2NK$ on the number of updates. With Theorem 1 plugged in, we have $K = k$ and the total running time becomes $O(k|E| \log |V| + kN \log N)$. $\square$

Appendix B shows how to handle the case of path overlaps, giving the following result:

Theorem 6. *Let $G = (V, E, \ell, \Sigma)$ be a labeled DAG and let M be a set of N pairs of the form $(P, [c..d])$. The algorithms from Theorems 3 and 5 can be modified to solve Problem 2 with additional time $O(L \log^2 |V|)$ or $O(L + \#\text{overlaps})$, where L is at most the input length and $\#\text{overlaps}$ is the number of overlaps between the input paths.*

The bound $O(L + \#\text{overlaps})$ comes as a direct consequence of using a generalized suffix tree to compute the overlaps in advance [34, proof of Theorem 2]. With the overlaps given, one can process each in constant time to see if they give the maximum for $C[j]$. The other bound $O(L \log^2 |V|)$ comes from backward searching a subpath in a concatenation of all subpaths using FM-index. At each step a range can be identified that contains all subpaths with suffix-prefix overlap with the current one. This range limits the keys suitable in the binary search trees storing the already computed coverage values, and thus a two-dimensional range search is needed (see Appendix B).

k	#graphs	mean $ V $	MPC method	Naive method
1	75	7275	$18 \pm 27\text{ms}$	$5638 \pm 12378\text{ms}$
2	117	8109	$23 \pm 36\text{ms}$	$6355 \pm 17641\text{ms}$
3	93	8306	$27 \pm 41\text{ms}$	$6499 \pm 17940\text{ms}$
4	99	8933	$32 \pm 49\text{ms}$	$6864 \pm 17868\text{ms}$
5	48	9779	$40 \pm 59\text{ms}$	$8053 \pm 18742\text{ms}$
6	32	10265	$45 \pm 65\text{ms}$	$7934 \pm 16659\text{ms}$
7	16	9928	$41 \pm 59\text{ms}$	$6973 \pm 15345\text{ms}$
8	10	11052	$57 \pm 83\text{ms}$	$8731 \pm 17497\text{ms}$
9	4	9538	$52 \pm 77\text{ms}$	$6252 \pm 13906\text{ms}$
10	3	10833	$61 \pm 102\text{ms}$	$7055 \pm 16221\text{ms}$
11	2	11186	$50 \pm 70\text{ms}$	$5932 \pm 10548\text{ms}$
15	1	16848	$154 \pm 194\text{ms}$	$25253 \pm 43873\text{ms}$

N	mean $ V $	MPC method	Naive method
$(10^0 \dots 10^1]$	8681	$8 \pm 5\text{ms}$	$15 \pm 8\text{ms}$
$(10^1 \dots 10^2]$	8808	$8 \pm 5\text{ms}$	$79 \pm 68\text{ms}$
$(10^2 \dots 10^3]$	9732	$10 \pm 7\text{ms}$	$524 \pm 392\text{ms}$
$(10^3 \dots 10^4]$	6824	$70 \pm 22\text{ms}$	$15153 \pm 5875\text{ms}$
$(10^4 \dots 10^5]$	12235	$153 \pm 66\text{ms}$	$49482 \pm 31900\text{ms}$

Fig. 3. The average running times, and their standard deviation, (in milliseconds) of the two approaches for co-linear chaining between a sequence and a DAG (Problem 2), for all inputs of a certain width k (top), and with N belonging to a certain interval (below). Both approaches are given the same anchors; the time for finding them is not included.

6 Discussion and experiments

For applying our solutions to Problem 2 in practice, one first needs to find the alignment anchors. As explained in the problem formulation, alignment anchors are such pairs $(P, [c..d])$ where P is a path in G and $\ell(P)$ matches $R[c..d]$. With sequence inputs, such pairs are usually taken to be *maximal exact matches* (MEMs) and can be retrieved in small space in linear time [5,4]. It is largely an open problem how to retrieve MEMs between a sequence and a DAG efficiently: The case of length-limited MEMs is studied in [37], based on an extension of [38] with features such as suffix tree functionality. On the practical side, anchor finding has already been incorporated into tools for conducting alignment of a sequence to a DAG [24,29].

For the purpose of demonstrating the efficiency of our MPC-approach applied to co-linear chaining, we implemented a MEM-finding routine based on simple dynamic programming. We leave it for future work to incorporate a practical procedure (e.g. like those in [24,29]). We tested the time improvement of our MPC-approach (Theorem 5) over the trivial algorithm (Theorem 3) on the sequence graphs of annotated human genes. Out of all the 62219 genes in the HG38 annotation for all human chromosomes, we singled out 8628 genes such that their sequence graph had at least 5000 nodes. Out of these, we picked 500 genes at random.

The size of the graphs for these 500 genes varied between $|V| = 5023$ and $|V| = 30959$ vertices. Their width, i.e., the number of paths in the MPC, varied between $k = 1$ and $k = 15$. (The number of graphs for each value of k is listed in the column #graphs of the top table of Figure 3.) The number of anchors, N , for patterns of length 1000 varied between 10^1 and 10^5 . As shown in Figure 3, with small values of N , our MPC-based co-linear chaining algorithm was twice as fast as the trivial algorithm. When values of N were increased from 10^1 to 10^5 , the difference increased to two orders of magnitude.

The improved efficiency when compared to the naive approach gives reason to believe a practical sequence-to-DAG aligner can be engineered along the algorithmic foundations given here. Future work includes the incorporation of a practical anchor-finding method, and testing whether the complete scheme improves transcript prediction through improved finding of exon chains [23].

On the theoretical side, it remains open whether the MPC algorithm could benefit from a better initial approximation and/or one that is faster to compute. More generally, it remains open whether the overall bound $O(k|E| \log |V|)$ for the MPC problem can be improved.

Acknowledgements. We thank Djamel Belazzougui for pointers on backward step on large alphabet and Gonzalo Navarro for pointing out the connection to pattern matching on hypertexts. This work was funded in part by the Academy of Finland (grant 274977 to AIT and grants 284598 and 309048 to AK and to VM), and by Futurice Oy (to TP).

References

1. Mohamed Ibrahim Abouelhoda. A Chaining Algorithm for Mapping cDNA Sequences to Multiple Genomic Sequences. In *14th International Symposium on String Processing and Information Retrieval*, volume 4726 of *LNCS*, pages 1–13. Springer, 2007.
2. Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.
3. Amihod Amir, Moshe Lewenstein, and Noa Lewenstein. Pattern matching in hypertext. *J. Algorithms*, 35(1):82–99, 2000.
4. Djamel Belazzougui. Linear time construction of compressed text indices in compact space. In *Proc. Symposium on Theory of Computing STOC 2014*, pages 148–193. ACM, 2014.
5. Djamel Belazzougui, Fabio Cunial, Juha Kärkkäinen, and Veli Mäkinen. Versatile Succinct Representations of the Bidirectional Burrows-Wheeler Transform. In *Proc. 21st Annual European Symposium on Algorithms (ESA 2013)*, volume 8125 of *LNCS*, pages 133–144. Springer, 2013.
6. Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, 3rd ed. edition, 2008.
7. Y. Chen and Y. Chen. An Efficient Algorithm for Answering Graph Reachability Queries. In *2008 IEEE 24th International Conference on Data Engineering*, pages 893–902, April 2008.
8. Y. Chen and Y. Chen. On the graph decomposition. In *2014 IEEE Fourth International Conference on Big Data and Cloud Computing*, pages 777–784, Dec 2014.
9. Deanna M Church, Valerie A Schneider, Karyn Meltz Steinberg, Michael C Schatz, Aaron R Quinlan, Chen-Shan Chin, Paul A Kitts, Bronwen Aken, Gabor T Marth, Michael M Hoffman, et al. Extending reference assembly models. *Genome biology*, 16(1):13, 2015.
10. Edith Cohen, Eran Halperin, Haim Kaplan, and Uri Zwick. Reachability and Distance Queries via 2-Hop Labels. *SIAM J. Comput.*, 32(5):1338–1355, 2003.
11. David Eppstein, Zvi Galil, Raffaele Giancarlo, and Giuseppe F. Italiano. Sparse Dynamic Programming I: Linear Cost Functions. *J. ACM*, 39(3):519–545, July 1992.
12. Stefan Felsner, Vijay Raghavan, and Jeremy Spinrad. Recognition algorithms for orders of small width and graphs of small dilworth number. *Order*, 20(4):351–364, Nov 2003.
13. Paolo Ferragina and Giovanni Manzini. Indexing compressed text. *J. ACM*, 52(4):552–581, July 2005.
14. Michael L Fredman. On computing the length of longest increasing subsequences. *Discrete Mathematics*, 11(1):29–35, 1975.
15. D. R. Fulkerson. Note on Dilworth’s decomposition theorem for partially ordered sets. *Proceedings of the American Mathematical Society*, 7(4):701–702, 1956.
16. Harold N. Gabow, Jon Louis Bentley, and Robert E. Tarjan. Scaling and related techniques for geometry problems. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, STOC ’84, pages 135–143, New York, NY, USA, 1984. ACM.
17. David Haussler, Maciej Smuga-Otto, Benedict Paten, Adam M. Novak, Sergei Nikitin, Maria Zueva, and Dmitrii Miagkov. A Flow Procedure for the Linearization of Genome Sequence Graphs. In S. Cenk Sahinalp, editor, *Research in Computational Molecular Biology: 21st Annual International Conference, RECOMB 2017, Hong Kong, China, May 3-7, 2017, Proceedings*, pages 34–49, Cham, 2017. Springer International Publishing.
18. Steffen Heber, Max Alekseyev, Sing-Hoi Sze, Haixu Tang, and Pavel A. Pevzner. Splicing graphs and EST assembly problem. *Bioinformatics*, 18 Suppl 1:S181–S188, 2002.
19. Wing-Kai Hon, Kunihiko Sadakane, and Wing-Kin Sung. Breaking a time-and-space barrier in constructing full-text indices. *SIAM J. Comput.*, 38(6):2162–2178, 2009.
20. John E. Hopcroft and Richard M. Karp. An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs. *SIAM J. Comput.*, 2(4):225–231, 1973.
21. H. V. Jagadish. A Compression Technique to Materialize Transitive Closure. *ACM Trans. Database Syst.*, 15(4):558–598, December 1990.
22. Ruoming Jin, Ning Ruan, Yang Xiang, and Haixun Wang. Path-tree: An Efficient Reachability Indexing Scheme for Large Directed Graphs. *ACM Trans. Database Syst.*, 36(1):7:1–7:44, March 2011.
23. Anna Kuosmanen, Tuukka Norri, and Veli Mäkinen. Evaluating approaches to find exon chains based on long reads. *Briefing in Bioinformatics*, 2017.

24. Antoine Limasset, Bastien Cazaux, Eric Rivals, and Pierre Peterlongo. Read mapping on de Bruijn graphs. *BMC bioinformatics*, 17(1):237, 2016.
25. Veli Mäkinen, Djamal Belazzougui, Fabio Cunial, and Alexandru I. Tomescu. *Genome-Scale Algorithm Design*. Cambridge University Press, 2015.
26. Veli Mäkinen, Leena Salmela, and Johannes Ylinen. Normalized N50 assembly metric using gap-restricted co-linear chaining. *BMC Bioinformatics*, 13:255, 2012.
27. Gene Myers and Webb Miller. Chaining Multiple-Alignment Fragments in Sub-Quadratic Time. In Kenneth L. Clarkson, editor, *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, 22-24 January 1995, San Francisco, California.*, pages 38–47. ACM/SIAM, 1995.
28. Gonzalo Navarro. Improved approximate pattern matching on hypertext. *Theor. Comput. Sci.*, 237(1-2):455–463, 2000.
29. Adam M Novak, Erik Garrison, and Benedict Paten. A graph extension of the positional Burrows-Wheeler transform and its applications. In *International Workshop on Algorithms in Bioinformatics*, pages 246–256. Springer, 2016.
30. S. C. Ntafos and S. Louis Hakimi. On path cover problems in digraphs and applications to program testing. *IEEE Trans. Software Engrg.*, 5(5):520–529, 1979.
31. James B. Orlin. Max flows in $O(nm)$ time, or better. In *Proceedings of the 45th Annual ACM Symposium on the Theory of Computing, STOC '13*, pages 765–774, New York, NY, USA, 2013. ACM.
32. Kunsoo Park and Dong Kyue Kim. String matching in hypertext. In *CPM*, pages 318–329, 1995.
33. Rob Patro, Geet Duggal, Michael I Love, Rafael A Irizarry, and Carl Kingsford. Salmon provides fast and bias-aware quantification of transcript expression. *Nature Methods*, 14(4):417–419, 04 2017.
34. Romeo Rizzi, Alexandru I. Tomescu, and Veli Mäkinen. On the complexity of Minimum Path Cover with Subpath Constraints for multi-assembly. *BMC Bioinformatics*, 15(S-9):S5, 2014.
35. Claus-Peter Schnorr. An algorithm for transitive closure with linear expected time. *SIAM J. Comput.*, 7(2):127–133, 1978.
36. Tetsuo Shibuya and Igor Kurochkin. Match Chaining Algorithms for cDNA Mapping. In *Proc. Algorithms in Bioinformatics (WABI 2003)*, volume 2812 of *LNCs*, pages 462–475. Springer, 2003.
37. Jouni Sirén. Indexing Variation Graphs. In *2017 Proceedings of the Nineteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 13–27. SIAM, 2017.
38. Jouni Sirén, Niko Välimäki, and Veli Mäkinen. Indexing graphs for path queries with applications in genome research. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 11(2):375–388, 2014.
39. J. Su, Q. Zhu, H. Wei, and J. X. Yu. Reachability Querying: Can It Be Even Faster? *IEEE Transactions on Knowledge and Data Engineering*, 29(3):683–697, March 2017.
40. Alexandru I. Tomescu, Travis Gagie, Alexandru Popa, Romeo Rizzi, Anna Kuosmanen, and Veli Mäkinen. Explaining a Weighted DAG with Few Paths for Solving Genome-Guided Multi-Assembly. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 12(6):1345–1354, 2015.
41. Raluca Uricaru, Céline Michotey, Hélène Chiapello, and Eric Rivals. Yoc, a new strategy for pairwise alignment of collinear genomes. *BMC Bioinformatics*, 16(1):111, Apr 2015.
42. Vijay V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.
43. Michaël Vyverman, Bernard De Baets, Veerle Fack, and Peter Dawyndt. A long fragment aligner called alfalfa. *BMC Bioinformatics*, 16(1):159, May 2015.
44. Michaël Vyverman, Dieter De Smedt, Yao-Cheng Lin, Lieven Sterck, Bernard De Baets, Veerle Fack, and Peter Dawyndt. Fast and Accurate cDNA Mapping and Splice Site Identification. In *Proceedings of the International Conference on Bioinformatics Models, Methods and Algorithms (BIOSTEC 2014)*, pages 233–238, 2014.
45. Sebastian Wandelt and Ulf Leser. RRCA: ultra-fast multiple in-species genome alignments. In Adrian-Horia Dediu, Carlos Martín-Vide, and Bianca Truthe, editors, *Algorithms for Computational Biology - First International Conference, AlCoB 2014, Tarragona, Spain, July 1-3, 2014, Proceedings*, volume 8542 of *Lecture Notes in Computer Science*, pages 247–261. Springer, 2014.
46. Hilmi Yildirim, Vineet Chaoji, and Mohammed J. Zaki. GRAIL: Scalable Reachability Index for Large Graphs. *Proc. VLDB Endow.*, 3(1-2):276–284, September 2010.

Construction time	Index size	Query time	Reference
$O(k E \log V)$	$O(k V)$	$O(1)$	this paper
$O(V ^2 + k\sqrt{k} V)$	$O(k V)$	$O(1)$	[7]
$O(k E)$ or $O(V E)$	$O(k V)$	$O(\log^2 k)$	[22]
$O(k(V + E))$	$O(k V)$	$O(k)$ or $O(V + E)$	[46]

Table 1. Previous comparable space/time tradeoffs for solving reachability queries. Compiled from [39, Table 1].

A Two simple applications

A.1 Reachability queries

An immediate application of Theorem 1 and of the values `last2reach` $[v, i]$ is for solving reachability queries. If we have all these $O(k|V|)$ values, then we can answer in constant time whether a node y is reachable from a node x , as in [21]: we check $\text{index}[x, i] \leq \text{index}[\text{last2reach}[y, i], i]$, where `index` was defined in the proof of Lemma 2, i is a path containing x , and we take by convention $\text{index}[-1, i] = -1$. Recall also that reachability queries in an arbitrary graph can be reduced to solving reachability queries in its DAG of strongly connected components, because nodes in the same component are pairwise reachable. See Table 1 for existing tradeoffs for solving reachability queries.⁴

Corollary 1. *Let $G = (V, E)$ be an arbitrary directed graph and let the width of its DAG of strongly connected components be k . In time $O(k|E| \log |V|)$ we can construct from G an index of size $O(k|V|)$, so that for any $x, y \in V$ we can answer in $O(1)$ time whether y is reachable from x .*

A.2 The LIS problem

The LIS problem asks us to delete the minimum number of values from an input sequence $s_1 \dots s_n$ such that remaining values form a strictly increasing series of values. Here the input sequence is assumed to come from an ordered alphabet Σ . For example, on input sequence 1, 4, 2, 3, 7, 5, 6, from the alphabet $\Sigma = \{1, 2, 3, 4, 5, 6, 7\}$, the unique optimal solution is 1, 2, 3, 5, 6. Such a longest increasing subsequence can be found in the optimal $O(n \log n)$ time [14].

This optimal algorithm works as follows. We first map Σ to a subset of $\{1, 2, \dots, n\}$ with an order-preserving mapping, in $O(n \log n)$ time (by e.g., sorting the sequence elements, and relabeling by the ranks in the order of distinct values). We then store, at every index i of the input sequence, the value `LLIS` $[i]$ defined as the length of the longest strictly increasing subsequence ending at i and using the i -th symbol. The values `LLIS` $[i]$ can be computed by dynamic programming, by storing all previous key-value pairs $(s_j, \text{LLIS}[j])$ in a search tree $\mathcal{T}$ as in Lemma 3, and querying $\mathcal{T}.\text{RMaxQ}(0, s_i - 1)$.

Consider the following extension of the LIS problem to a labeled DAG $G = (V, E, \ell, \Sigma)$ of width k . For a path $P = (v_1, \dots, v_t)$ in G , let the *label* of P , denoted $\ell(P)$, be the concatenation of the labels of the nodes of P , namely $\ell(v_1) \dots \ell(v_t)$. Among all paths P in G , and among all subsequences of $\ell(P)$, we need to find a longest strictly increasing subsequence.

We now explain how to extend the previous dynamic programming algorithm for this problem. We analogously map Σ to a subset of $\{1, 2, \dots, |V|\}$ with an order-preserving mapping in $O(|V| \log |V|)$ time, as above. Recall that we assume $V = \{1, \dots, |V|\}$, where $1, \dots, |V|$ is a topological order. Assume also that we have K paths to cover V and `forward` $[u]$ is computed for all $u \in V$.

For each node v , we aim to analogously compute `LLIS` $[v]$ as the length of a longest strictly increasing subsequence of the labels of all paths ending at v , with the property that $\ell(v)$ is the last element of this subsequence.

⁴ Note that [7] incorrectly attributes to [21] query time $O(\log k)$, and as a consequence [39, 22] incorrectly mention query time $O(\log k)$ for [7].

For each $i \in [1..K]$, we let $\mathcal{T}_i$ be a search tree as in Lemma 3, initialized with key-value pairs $(0, 0), (1, -\infty), (2, -\infty), \dots, (|V|, -\infty)$. The algorithm proceeds in the fixed topological ordering. Assume now that we are at some position u , and have already updated all search trees associated with the covering paths going through u . For every $(v, i) \in \text{forward}[u]$, we update $\text{LLIS}[v] = \max(\text{LLIS}[v], \mathcal{T}_i.\text{RMaxQ}(0, \ell(v) - 1) + 1)$. Once the algorithm reaches v in the topological ordering, value $\text{LLIS}[v]$ has been updated from all u' such that $(v, i) \in \text{forward}[u']$. It remains to show how to update each $\mathcal{T}_i$ when reaching v , for all covering paths i on which v occurs. This is done as $\mathcal{T}_i.\text{update}(\ell(v), \text{LLIS}[v])$. Initialization is handled by the $(0, 0)$ key-value pair so that any position can start a new increasing subsequence. Figure 4 shows an example.

Fig. 4. Detailed steps for solving a LIS DAG instance, focusing on three nodes: u , u' , and v . In any topological order, node u will be processed before node u' , which in turn will be processed before v . After computing $\text{LLIS}[u] = 1$ we set $\text{LLIS}[v] = \max(1, \mathcal{T}_3.\text{RMaxQ}(0, 7) + 1) = 2$, since at that point $\mathcal{T}_3$ contains $(0, 0), (1, -\infty), \dots, (3, -\infty), (4, 1), (5, -\infty), \dots, (10, -\infty)$. After computing $\text{LLIS}[u'] = 3$, we set $\text{LLIS}[v] = \max(2, \mathcal{T}_2.\text{RMaxQ}(0, 7) + 1) = 4$, since at that point $\mathcal{T}_2$ contains $(0, 0), (1, -\infty), (2, -\infty), (3, 1), (4, -\infty), (5, -\infty), (6, 2), (7, 3), (8, -\infty), \dots, (10, -\infty)$. When we finally reach v we leave $\text{LLIS}[v] = \max(4, \mathcal{T}_1.\text{RMaxQ}(0, 7) + 1) = 4$, since at that point $\mathcal{T}_1$ contains $(0, 0), (1, -\infty), \dots, (4, -\infty), (5, 2), (6, -\infty), (7, 1), (8, -\infty), (9, 3), (10, -\infty)$.

The final answer to the problem is $\max_{v \in V} \text{LLIS}[v]$, with the actual LIS to be found with a standard traceback. The algorithm runs in $O(K|V| \log |V|)$ time. With Theorem 1 plugged in, we have $K = k$ and the total running time becomes $O(k|E| \log |V| + k|V| \log |V|) = O(k|E| \log |V|)$, under our assumption $|E| \geq |V| - 1$. The following theorem summarizes this result.

Theorem 7. *Let $G = (V, E, \ell, \Sigma)$ be a labeled DAG of width k , where Σ is an ordered alphabet. We can find a longest increasing subsequence in G in time $O(k|E| \log |V|)$.*

When the DAG is just a labeled path with $|E| = |V| - 1$ (modeling the standard LIS problem), then the algorithm from Lemma 1 returns one path ($K = 1$). The complexity is then $O(|V| \log |V|)$, matching the best possible bound for the standard LIS problem [14].

B Co-linear chaining with path overlaps

We now consider how to extend the algorithms we developed for Problem 3 to work for the more general case of Problem 2, where overlaps between paths are allowed in a solution. The detection and merging of such path overlaps has been studied in [34], and we tailor a similar approach for our purposes.

We use an *FM-index* [13] tailored for large alphabets [19], and a two-dimensional range search tree [6] modified to support range maximum queries. The former is used for obtaining all ranges $[i'..i]$ in the coverage array C such that all input pairs $M[i'], \dots, M[i]$ have a path $M[i''].P$, $i' \leq i'' \leq i$, overlapping with the path $M[j].P$ of j -th input pair $M[j]$. Here, the endpoint of the j -th input pair is at node v visited in topological order. This implies that the paths of the input pairs $[i'..i]$ have already been visited, and thus, by induction, that $C[i'..i]$ values have been correctly computed (subject to the modification we are about to study). The sequence ranges $[M[i''].c..M[i''].d]$ for all $i'' \in [i'..i]$ may be arbitrarily located with respect to interval $[M[j].c..M[j].d]$, so we need to maintain an analogous mechanism with search trees of type $\mathcal{T}$ and $\mathcal{I}$ as in our co-linear chaining algorithm based on a path cover. This time we cannot, in advance, separate

the input pairs to K paths with different search trees, but we have a dynamic setting, with interval $[i'..i]$ deciding which values should be taken into account. This is where a two-dimensional range search tree is used to support these queries in $O(\log^2 N)$ time: Figure 5 illustrates this.

Fig. 5. Answering two-dimensional range maximum queries. A search interval $[i'..i]$ is split to $O(\log(i - i' + 1))$ intervals / subtrees. In each subtree root we store a type $\mathcal{T}$ and a type $\mathcal{I}$ search tree of Lemma 3 containing values $C[i'']$ or $C[i''] - M[i''].d$, respectively, for indexes i'' in the corresponding subinterval.

In what follows we show that $O(L)$ queries are sufficient to take all overlaps into account throughout the algorithm execution (holding for both the trivial algorithm and for the one based on a path cover), where $L = \sum_i |M[i].P|$ —the sum of the path lengths—is at most the total input length. The construction will actually induce an order for the input pairs such that $O(L)$ queries are sufficient: Since the other parts of the algorithms do not use the order of input pairs directly, we can safely reorganize the input accordingly.

With this introduction, we are ready to consider how all the intervals $[i'..i]$ related to j -th pair are obtained. We build in $O(L \log \log |V|)$ time the FM-index version proposed in [19] of sequences $T = (\prod_i \#(M[i].P)^{-1})\#$, where $\#$ is a symbol not in alphabet $\{1, 2, \dots, |V|\}$ and considered smaller than other symbols, e.g. $\# = 0$, and X^{-1} denotes the reverse $X[|X|]X[|X| - 1] \dots X[1]$ of X .

For our purposes it is sufficient to know that the FM-index of T , when given an interval $I(X)$ corresponding to lexicographically-ordered suffixes that start with X , can determine the interval $I(cX)$ in $O(\log \log |V|)$ time [19]. This operation is called *backward step*.

We use the index to search $M[j].P$ in the forward direction by searching its reverse with backward steps. Consider we have found interval $I((M[j].P[1..k])^{-1})$, for some k , such that backward step $I(\#(M[j].P[1..k])^{-1})$ results in a non-empty interval $[i'..i]$. This interval $[i'..i]$ corresponds to all suffixes of T that have $\#(M[j].P[1..k])^{-1}$ as a prefix. That is, $[i'..i]$ corresponds to input pairs whose path suffix have a length k overlap with the path prefix of j -th input pair. For this interval to match with coverage array C , we just need to rearrange the input pairs according to their order in the first N rows of the array storing the lexicographic order of suffixes of T .

Since each backward step on the index may induce a range search on exactly one interval $[i'..i]$, the running time is dominated by the range queries.⁵ On the other hand, this also gives the bound L on the number of range queries, as claimed earlier.

Alternatively, one can omit the expensive range queries and process each overlapping pair separately, to compute in constant time its contribution to $C[j]$. This gives another bound $O(L \log \log |V| + \#\text{overlaps})$, where $\#\text{overlaps}$ is the number of overlaps between the input paths. This can be improved to $O(L + \#\text{overlaps})$ by using a generalized suffix tree to compute the overlaps in advance [34, proof of Theorem 2].

The result is summarized in Theorem 6 at page 9.

⁵ A simple wavelet tree based FM-index would provide the same bound, but in case the range search part is later improved, we used the best bound for the subroutine.