

Covering Pairs in Directed Acyclic Graphs

Niko Beerenwinkel¹, Stefano Beretta^{2,4}, Paola Bonizzoni², Riccardo Dondi³,
and Yuri Pirola²

¹ Dept. of Biosystems Science and Engineering, ETH Zurich, Basel, Switzerland,
`niko.beerenwinkel@bsse.ethz.ch`

² DISCO, Univ. degli Studi di Milano-Bicocca, Milan, Italy,
`{beretta,bonizzoni,pirola}@disco.unimib.it`

³ Dip. di Scienze Umane e Sociali, Univ. degli Studi di Bergamo, Bergamo, Italy,
`riccardo.dondi@unibg.it`

⁴ Inst. for Biomedical Technologies, National Research Council, Segrate, Italy

Abstract. The Minimum Path Cover problem on directed acyclic graphs (DAGs) is a classical problem that provides a clear and simple mathematical formulation for several applications in different areas and that has an efficient algorithmic solution. In this paper, we study the computational complexity of two constrained variants of Minimum Path Cover motivated by the recent introduction of next-generation sequencing technologies in bioinformatics. The first problem (MinPCR), given a DAG and a set of pairs of vertices, asks for a minimum cardinality set of paths “covering” all the vertices such that both vertices of each pair belong to the same path. For this problem, we show that, while it is NP-hard to compute if there exists a solution consisting of at most three paths, it is possible to decide in polynomial time whether a solution consisting of at most two paths exists. The second problem (MaxRPSP), given a DAG and a set of pairs of vertices, asks for a path containing the maximum number of the given pairs of vertices. We show its NP-hardness and also its W[1]-hardness when parametrized by the number of covered pairs. On the positive side, we give a fixed-parameter algorithm when the parameter is the maximum overlapping degree, a natural parameter in the bioinformatics applications of the problem.

1 Introduction

The *Minimum Path Cover* (MinPC) problem is a well-known problem in graph theory [7]. Given a *directed acyclic graph* (DAG), MinPC asks for a minimum-cardinality set Π of paths such that each vertex of G belongs to at least one path of Π . The problem can be solved in polynomial time with an algorithm based on the well-known Dilworth’s theorem for partially ordered sets [3], which allows to relate the size of a minimum path cover to that of a maximum matching in a bipartite graph obtained from the input DAG.

The Minimum Path Cover problem has important applications in several fields ranging from bioinformatics [1, 6, 11] to software testing [10]. In particular, in bioinformatics the Minimum Path Cover problem is applied to the reconstruction of a set of highly-similar sequences starting from a large set of their

short fragments (called *short reads*) [6, 11]. More precisely, each fragment is represented by a single vertex and two vertices are connected if the alignments of the corresponding reads on the genomic sequence overlap. In [11], the paths on such a graph represent putative transcripts and a minimum-cardinality set of paths “covering” all the vertices represents a set of protein isoforms which are likely to originate from the observed reads. On the other hand, in [6] the paths on such a graph represent the genomes of putative viral haplotypes and a minimum-cardinality set of paths covering the whole graph represents the likely structure of a viral population.

Recently, different constraints have motivated the definition of new variants of the minimum path cover problem. In [1], given a DAG D and a set P of required paths, the proposed problem asks for a minimum cardinality set of paths such that: (1) each vertex of the graph belongs to some path, and (2) each path in P is a subpath of a path of the solution. The authors have described a polynomial-time algorithm to solve this problem by collapsing each required path into a single vertex and then finding a minimum path cover on the resulting graph. Other constrained problems related to minimum path cover have been proposed in the context of social network analysis and, given an edge-colored graph, ask for the maximum number of vertex-disjoint uni-color paths that cover the vertices of the given graph [2, 12].

Some constrained variants of the minimum path cover problem have been introduced in the past by Ntafos and Hakimi in the context of software testing [10] and appear to be relevant for some sequence reconstruction problems of recent interest in bioinformatics. More precisely, in software testing each procedure to be tested is modeled by a graph where vertices correspond to single instructions and two vertices are connected if the corresponding instructions are executed sequentially. The test of the procedure should check each instruction at least once, hence a minimum path cover of the graph represents a minimum set of execution flows that allows to test all the instructions. Clearly, not all the execution flows are possible. For this reason, Ntafos and Hakimi proposed the concept of required pairs, which are pairs of vertices that a feasible solution must include in a path, and that of impossible pairs, which are pairs of vertices that a feasible solution must not include in the same path. In particular, one of the problems introduced by Ntafos and Hakimi is the *Minimum Required Pairs Cover* (MinRPC) problem where, given a DAG and a set of required pairs, the goal is to compute a minimum set of paths *covering* all the required pairs, *i.e.*, a minimum set of paths such that, for each required pair, at least one path contains both vertices of the pair.

The concept of required pairs is also relevant for sequence reconstruction problems in bioinformatics, as short reads are often sequenced in pairs (*paired-end reads*) and these pairs of reads must align to a single genetic sequence. As a consequence, each pair of vertices corresponding to paired-end reads must belong to the same path of the cover. Paired-end reads provide valuable information that, in principle, could greatly improve the accuracy of the reconstruction. However, they are often used only to filter out the reconstructed sequences that

do not meet such constraints, instead of directly exploiting them during the reconstruction process. Notice that MinRPC asks for a solution that covers only the required pairs, while in bioinformatics we are also interested in covering all the vertices. For this reason, we consider a variant of the Minimum Path Cover problem, called *Minimum Path Cover with Required Pairs* (MinPCR), that, given a DAG and a set of required pairs, asks for a minimum set of paths covering all the vertices and all the required pairs. Clearly, MinPCR is closely related to MinRPC. In fact, as we show in Section 2, the same reduction used in [10] to prove the NP-hardness of MinRPC can be applied to our problem, leading to its intractability.

In this paper, we continue the analysis of [10] by studying the complexity of path covering problems with required pairs. More precisely, we study how the complexity of these problems is influenced by two parameters relevant for the sequence reconstruction applications in bioinformatics: (1) the minimum number of paths covering all the vertices and all the required pairs and (2) the maximum *overlapping degree* (defined later). In the bioinformatics applications we discussed, the first parameter—the number of covering paths—is often small, thus an algorithm exponential in the size of the solution could be of interest. The second parameter we consider in this paper, the maximum overlapping degree, can be informally defined as follows. Two required pairs overlap when there exists a path that connects the vertices of the pairs, and the path cannot be split in two disjoint subpaths that separately connect the vertices of the two pairs. Then, the overlapping degree of a required pair is the number of required pairs that overlap with it. In the sequence reconstruction applications, as the distance between two paired-end reads is fixed, the maximum overlapping degree is small compared to the number of vertices, hence it is a natural parameter for investigating the computational complexity of the problem.

First, we investigate how the computational complexity of MinPCR is influenced by the first parameter. In this paper we prove that it is NP-complete to decide if there exists a solution of MinPCR consisting of at most three paths (via a reduction from the 3-Coloring problem). We complement this result by giving a polynomial-time algorithm for computing a solution with at most 2 paths, thus establishing a sharp tractability borderline for MinPCR when parameterized by the size of the solution. These results significantly improve the hardness result that Ntafos and Hakimi [10] presented for MinRPC (and that holds also for MinPCR), where the solution contains a number of paths which is polynomial in the size of the input.

Then, we investigate how the computational complexity of MinPCR is influenced by the second parameter, the overlapping degree. Unfortunately, MinPCR is NP-hard even if the maximum overlapping degree is 0. In fact, this can be easily obtained by modifying the reduction presented in [10] to hold also for restricted instances of MinPCR with no overlapping required pairs.

A natural heuristic approach for solving MinPCR is the one which computes a solution by iteratively adding a path that covers a maximum set of required pairs not yet covered by a path of the solution. This approach leads

to a natural combinatorial problem, the *Maximum Required Pairs with Single Path* (MaxRPSP) problem, that, given a DAG and a set of required pairs, asks for a path that covers the maximum number of required pairs. We investigate the complexity of MaxRPSP and we show that it is not only NP-hard, but also W[1]-hard when the parameter is the number of covered required pairs. Similarly as MinPCRP, we consider the MaxRPSP problem parameterized by the maximum overlapping degree but, differently from MinPCRP, we give a fixed-parameter algorithm for this case. This positive result shows a gap between the complexity of MaxRPSP and the complexity of MinPCRP when parameterized by the maximum overlapping degree.

The rest of the paper is organized as follows. First, in Section 2 we give some preliminary notions and we introduce the formal definitions of the two problems. In Section 3, we investigate the computational complexity of MinPCRP when the solution consists of a constant number of paths: we show that it is NP-complete to decide if there exists a solution of MinPCRP consisting of at most three paths, while the existence of a solution consisting of at most two paths can be computed in polynomial time. In Section 4, we investigate the computational complexity of MaxRPSP: we prove its W[1]-hardness when the parameter is the number of required pairs covered by the path (Section 4.1) and we give a fixed-parameter algorithm when the parameter is the maximum overlapping degree (Section 4.2).

2 Preliminaries

In this section, we introduce the basic notions used in the rest of the paper and we formally define the two combinatorial problems we are interested in.

While our problems deal with directed graphs, we consider both directed and undirected graphs. We denote an *undirected graph* as $G = (V, E)$ where V is the set of vertices and E is the set of (undirected) edges, and a *directed graph* as $D = (N, A)$ where N is the set of vertices and A is the set of (directed) arcs. We denote an edge of $G = (V, E)$ as $\{v, u\} \in E$ where $v, u \in V$. Moreover, we denote an arc of $D = (N, A)$ as $(v, u) \in A$ where $v, u \in N$.

Given a directed graph $D = (N, A)$, a *path* π from vertex v to vertex u , denoted as vu -path, is a sequence of vertices $\langle v_1, \dots, v_n \rangle$ such that $(v_i, v_{i+1}) \in A$, $v = v_1$ and $u = v_n$. We say that a vertex v *belongs to* a path $\pi = \langle v_1, \dots, v_n \rangle$, denoted as $v \in \pi$, if $v = v_i$, for some $1 \leq i \leq n$. Given a path $\pi = \langle v_1, \dots, v_n \rangle$, we say that a path $\pi' = \langle v_i, v_{i+1}, \dots, v_{j-1}, v_j \rangle$, with $1 \leq i \leq j \leq n$, is a subpath of π . Given a set $N' \subseteq N$ of vertices, a path π *covers* N' if every vertex of N' belongs to π .

In the paper, we consider a set R of pairs of vertices in N . We denote each pair as $[v_i, v_j]$, to avoid ambiguity with the notations of edges and arcs.

Now, we are able to give the definitions of the combinatorial problems we are interested in.

Problem 1. Minimum Path Cover with Required Pairs (MinPCRP)

Input: a directed acyclic graph $D = (N, A)$, a source $s \in N$, a sink $t \in N$, and




Fig. 1. Examples of overlapping required pairs. The two required pairs $[u', v']$ and $[u'', v'']$ are represented by dashed lines. In (a) the required pairs are *alternated*, while in (b) they are *nested*.

a set $R = \{[v_x, v_y] \mid v_x, v_y \in N, v_x \neq v_y\}$ of required pairs.

Output: a minimum cardinality set $\Pi = \{\pi_1, \dots, \pi_n\}$ of directed *st*-paths such that every vertex $v \in N$ belongs to at least one *st*-path $\pi_i \in \Pi$ and every required pair $[v_x, v_y] \in R$ belongs to at least one *st*-path $\pi_i \in \Pi$, i.e. v_x, v_y belongs to π_i .

Problem 2. Maximum Required Pairs with Single Path (MaxRPSP)

Input: a directed acyclic graph $D = (N, A)$, a source $s \in N$, a sink $t \in N$ and a set $R = \{[v_x, v_y] \mid v_x, v_y \in N, v_x \neq v_y\}$ of required pairs.

Output: an *st*-path π that covers a set $R' = \{[v_x, v_y] \mid v_x, v_y \in \pi\} \subseteq R$ of maximum cardinality.

Two required pairs $[u', v']$ and $[u'', v'']$ in R *overlap* if there exists a path π in D such that the four vertices appear in π in one of the following orders (assuming that the vertex u' appears before u'' in π), where v' and u'' are two distinct vertices of G (see Fig. 1):

- $\langle u', u'', v', v'' \rangle$ (the two required pairs are *alternated*);
- $\langle u', u'', v'', v' \rangle$ (the required pair $[u'', v'']$ is *nested* in $[u', v']$).

Notice that, from this definition, the required pairs $[x, y]$ and $[y, z]$ do not overlap.

Finally, consider a required pair $[u', v']$ of R . We define the *overlapping degree* of $[u', v']$ as the number of required pairs in R that overlap with $[u', v']$.

Hardness of MinPCRPs. As we mentioned in the introduction, MinPCRPs is related to a combinatorial problem which has been studied in the context of program testing [10], where it is shown to be NP-hard. More precisely, given a directed acyclic graph $D = (N, A)$, a source $s \in N$, a sink $t \in N$ and a set $R = \{[v_x, v_y] \mid v_x, v_y \in N, v_x \neq v_y\}$ of required pairs, the *Minimum Required Pairs Cover* (MinRPC) problem asks for a minimum cardinality set $\Pi = \{\pi_1, \dots, \pi_n\}$ of directed *st*-paths such that every required pair $[v_x, v_y] \in R$ belongs to at least one *st*-path $\pi_i \in \Pi$, i.e. $v_x, v_y \in \pi_i$.

MinRPC can be easily reduced to MinPCRPs due to the following property: each vertex of the graph D (input of MinRPC) must belong to at least one required pair. Indeed, if this condition does not hold for some vertex v , we can modify the graph D by contracting v (that is removing v and adding an edge (u, z) to A , for each $u, z \in N$ such that $(u, v), (v, z) \in A$). This implies that, since in an instance of MinRPC all the resulting vertices belong to some required

pair, a feasible solution of that problem must cover every vertex of the graph. Then, a solution of MinRPC is also a solution of MinPCRP, which implies that MinPCRP is NP-hard.

MinPCRP on *directed* graphs (not necessarily acyclic) is as hard as MinPCRP on DAGs. In fact, since each strongly connected component can be covered with a single path, we can replace them with single vertices, obtaining a DAG and without changing the size of the solution. Clearly, MinPCRP on general graphs and requiring that the covering paths are simple is as hard as the Hamiltonian path problem, which is NP-complete.

3 A Sharp Tractability Borderline for MinPCRP

In this section, we investigate the computational complexity of MinPCRP and we give a sharp tractability borderline for k -PCRP, the restriction of MinPCRP where we ask whether there exist k paths that cover all the vertices of the graph and all the set of required pairs. First, we show (Sect. 3.1) that 3-PCRP is NP-complete. This result implies that k -PCRP does not belong to the class XP⁵, so it is probably hopeless to look for an algorithm having complexity $O(n^k)$, and hence for a fixed-parameter algorithm in k . We complement this result by giving (Sect. 3.2) a polynomial time algorithm for 2-PCRP, thus defining a sharp borderline between tractable and intractable instances of MinPCRP.

3.1 Hardness of 3-PCRP

In this section we show that 3-PCRP is NP-complete. We prove this result via a reduction from the well-known 3-Coloring (3C) problem which, given an undirected (connected) graph $G = (V, E)$, asks for a coloring $c : V \rightarrow \{c_1, c_2, c_3\}$ of the vertices of G with exactly 3 colors, such that, for every $\{v_i, v_j\} \in E$, we have $c(v_i) \neq c(v_j)$.

Starting from an undirected graph $G = (V, E)$ (instance of 3C), we construct a corresponding instance $\langle D = (N, A), R \rangle$ of 3-PCRP as follows. For every subset $\{v_i, v_j\}$ of cardinality 2 of V , we define a graph $D_{i,j} = (N_{i,j}, A_{i,j})$ (in the following we assume that, for each $D_{i,j}$ associated with set $\{v_i, v_j\}$, $i < j$). The vertex set $N_{i,j}$ is $\{s^{i,j}, n_i^{i,j}, n_j^{i,j}, f^{i,j}, t^{i,j}\}$. The set $A_{i,j}$ of arcs connecting the vertices of $N_{i,j}$ can have two possible configurations, depending on the fact that $\{v_i, v_j\}$ belongs or does not belong to E . In the former case, that is $\{v_i, v_j\} \in E$, $D_{i,j}$ is in *configuration (1)* (see Fig. 2 (a)) and:

$$A_{i,j} = \{(s^{i,j}, n_i^{i,j}), (s^{i,j}, n_j^{i,j}), (s^{i,j}, f^{i,j}), (n_i^{i,j}, t^{i,j}), (n_j^{i,j}, t^{i,j}), (f^{i,j}, t^{i,j})\}$$

In the latter case, that is $\{v_i, v_j\} \notin E$, $D_{i,j}$ is in *configuration (2)* (see Fig. 2 (b)) and:

$$A_{i,j} = \{(s^{i,j}, n_i^{i,j}), (s^{i,j}, f^{i,j}), (n_i^{i,j}, n_j^{i,j}), (n_j^{i,j}, t^{i,j}), (f^{i,j}, t^{i,j})\}$$




Fig. 2. Example of the two configurations of subgraph $D_{i,j} = (N_{i,j}, A_{i,j})$ associated with a pair $\{v_i, v_j\}$ of vertices of a graph $G = (V, E)$.




Fig. 3. Example of graph $D = (N, A)$ associated with graph $G = (V, E)$. Grey boxes represent subgraphs $D_{i,j}$ in one of the two possible configurations of Fig. 2.

The whole graph $D = (N, A)$ is constructed by concatenating the graphs $D_{i,j}$ (for all $1 \leq i < j \leq n$) according to the lexicographic order of their indices i, j . The sink $t^{i,j}$ of each graph $D_{i,j}$ is connected to the source $s^{i',j'}$ of the graph $D_{i',j'}$ which immediately follows $D_{i,j}$. A distinguished vertex s is connected to the source of $D_{1,2}$ (i.e., the first subgraph), while the sink of $D_{n-1,n}$ (i.e., the last subgraph) is connected to a second distinguished vertex t . Fig. 3 depicts such a construction.

The set R of required pairs is defined as follows.

$$R = \{[s, f^{i,j}] \mid \{v_i, v_j\} \in E\} \cup \bigcup_{1 \leq i \leq n} R_i \quad \text{where } R_i = \{[n_i^{i,j}, n_i^{i,h}] \mid 1 \leq j \leq h \leq n\}$$

The following lemmas prove the correctness of the reduction.

Lemma 1. *Let $G = (V, E)$ be an undirected (connected) graph and let $\langle D = (N, A), R \rangle$ be the corresponding instance of 3-PCRP. Then, given a 3-coloring of G we can compute in polynomial time three paths of D that cover all its vertices and every required pair in R .*

Proof. Consider a 3-coloring of G and let $\{V_1, V_2, V_3\}$ be the tri-partition of V induced by the 3-coloring. We show how to compute in polynomial time three paths π_1, π_2, π_3 that cover all the vertices of D and every required pair in R . For each $v_i \in V_c$, path π_c passes through vertices $n_i^{i,j}$ of subgraphs $D_{i,j}$ for every $v_j \in V$, while for each subgraph $D_{p,q}$ such that $v_p, v_q \notin V_c$, π_c passes

⁵ We recall that the class XP contains those problems that, given a parameter k , can be solved in time $O(n^{f(k)})$

through vertices $f^{p,q}$. Notice that each π_c is well-defined, since when $n_i^{i,j}$, $n_j^{i,j}$ are associated with the same color c , $D_{i,j}$ is in configuration (2), hence the path can pass through both vertices $n_i^{i,j}$ and $n_j^{i,j}$.

We show that π_1, π_2, π_3 cover every required pair in R . Notice that for each $\{v_i, v_j\} \in E$, since v_i and v_j have different colors, by construction one of the paths π_1, π_2, π_3 passes through $n_i^{i,j}$, while another one passes through $n_j^{i,j}$. Now, we show that every required pair in R_i is covered. By construction, the vertices $n_i^{i,j}$ of D associated with the same vertex v_i of G belong to the same path π_c where c is the color of v_i . Therefore, all the required pairs in each R_i are covered by one of the three paths. Now, we show that π_1, π_2, π_3 cover the required pairs $\{\{s, f^{i,j}\} \mid 1 \leq i < j \leq n\}$. Indeed, consider a subgraph $D_{i,j}$, and assume w.l.o.g. that v_i has color c and that v_j has color c' . By construction, path $\pi_{c''}$ (with $c'' \notin \{c, c'\}$) passes through $f^{i,j}$. Then, π_1, π_2, π_3 cover all the required pairs in R .

Finally, in order to show that all the vertices of D are covered by at least one path, the only vertices that might be not covered are $s^{i,j}$ and $t^{i,j}$, for $1 \leq i < j \leq n$, since they do not belong to any required pair. However, these vertices are articulation points, hence all the three paths necessarily pass through them. \square

Lemma 2. *Let $G = (V, E)$ be an undirected graph and let $\langle D = (N, A), R \rangle$ be the corresponding instance of 3-PCR. Then, given three paths in D that cover all its vertices and every required pair in R we can compute in polynomial time a 3-coloring of G .*

Proof. Consider three paths π_1, π_2, π_3 of D that cover all the vertices of D and every required pair in R . We define the corresponding 3-coloring of the graph G , consisting of the colors c_1, c_2, c_3 .

First, we prove a property of the three paths π_1, π_2, π_3 . We show that, given a vertex $v_i \in V$, there exists at least one path among π_1, π_2, π_3 that covers all the required pairs in R_i . Consider a vertex $v_i \in V$. Since G is connected, it follows that there exists at least one vertex adjacent to v_i , w.l.o.g. v_j , such that $\{v_i, v_j\} \in E$. Now, consider the subgraph $D_{i,j}$. By construction, since $D_{i,j}$ has a configuration (1), a solution of MinPCR must contain three different paths, each one passing through one of the vertices $n_i^{i,j}$, $n_j^{i,j}$, $f^{i,j}$. Now, assume that path π_1 passes through $n_i^{i,j}$. Notice that π_2, π_3 cannot pass through $n_i^{i,j}$. But then, since π_1 is the only path that covers $n_i^{i,j}$ and since R_i contains a pair $[n_i^{i,j}, n_i^{i,h}]$, for each $h \neq j$, it follows that all the vertices $n_i^{i,h}$, $1 \leq h \leq n$, must belong to π_1 . It follows that, given a vertex $v_i \in V$, there exists one path in $\{\pi_1, \pi_2, \pi_3\}$ that covers all the required pairs in R_i . Moreover, since all the three paths pass through the vertices $s^{i,j}$ and $t^{i,j}$ for $1 \leq i < j \leq n$, then all the vertices of D are covered by $\{\pi_1, \pi_2, \pi_3\}$.

Now, we define a 3-coloring of G , where $C = \{c_1, c_2, c_3\}$ is the set of colors. If a required pair in R_i is covered by a path π_x , $1 \leq x \leq 3$, then we assign the color c_x to vertex v_i . Notice that the coloring is feasible, that is $c(v_i) \neq c(v_j)$ when $\{v_i, v_j\} \in E$. Indeed, consider two vertices v_i, v_j associated with the same color, and consider the two corresponding sets R_i, R_j of required pairs.

By construction, it follows that R_i, R_j are covered by the same path, say π_1 . Consider the subgraph $D_{i,j}$. Since R_i, R_j are both covered by π_1 , it follows that $D_{i,j}$ must have a configuration (2), hence $\{v_i, v_j\} \notin E$. Hence we have defined a 3-coloring of G . \square

As a consequence of the previous lemmas, we can easily prove the following result.

Theorem 1. *3-PCRPP is NP-complete.*

Proof. The NP-hardness of 3-PCRPP follows directly from Lemma 1 and Lemma 2 and from the NP-completeness of 3C [8]. 3-PCRPP is in NP, since, given three paths π_1, π_2, π_3 , we can verify in polynomial time that π_1, π_2, π_3 cover all the vertices of D and that every required pair in R is covered by some path in $\{\pi_1, \pi_2, \pi_3\}$. \square

3.2 A Polynomial Time Algorithm for 2-PCRPP

In this section we give a polynomial time algorithm for computing a solution of 2-PCRPP. Notice that 1-PCRPP can be easily solved in polynomial time, as there exists a solution of 1-PCRPP if and only if the reachability relation of the vertices of the input graph is a total order.

The algorithm for solving 2-PCRPP is based on a polynomial-time reduction to the 2-Clique Partition problem, which, given an undirected graph $G = (V, E)$, asks whether there exists a partition of V in two sets V_1, V_2 both inducing a clique in G . The 2-Clique Partition problem is polynomial-time solvable [8, probl. GT15]. To perform this reduction we assume that given $\langle D = (N, A), R \rangle$, instance of 2-PCRPP, every vertex of the graph D belongs to at least one required pair in R . Otherwise, we add to R the required pairs $[s, v_i]$ for all $v_i \in N$ that do not belong to any required pair. Therefore, a solution that covers all the required pairs in R covers also all the vertices, hence it is a feasible solution of 2-PCRPP. Moreover, notice that this transformation does not affect the solution of 2-PCRPP, since all the paths start from s and cover all the nodes of the graph, including the additional required pairs.

The algorithm, starting from an instance $\langle D = (N, A), R \rangle$ of 2-PCRPP, computes in polynomial time a corresponding undirected graph $G = (V, E)$ where:

- $V = \{v_c \mid c \in R\}$
- $E = \{\{v_{c_i}, v_{c_j}\} \mid \text{there exists a path in } D \text{ that covers both } c_i \text{ and } c_j\}$

Given a set of required pairs $R' \subseteq R$, we denote by $V(R')$ the corresponding set of vertices of G (i.e., $V(R') = \{v_c \mid c \in R'\}$).

The algorithm is based on the following fundamental property.

Lemma 3. *Given an instance $\langle D = (N, A), R \rangle$ of 2-PCRPP and the corresponding graph $G = (V, E)$, then there exists a path π that covers a set R' of required pairs if and only if $V(R')$ is a clique of G .*

Proof. We prove the lemma by induction on the number k of required pairs (vertices, resp.) of R' ($V(R')$, resp.).

When $k = 0$ the lemma trivially holds, in fact having no required pairs, *i.e.*, $R' = \emptyset$, induces an empty clique, *i.e.*, $V(\emptyset)$.

If $k = 1$, then we can assume that there exists at least one path in D that covers the only required pair c (otherwise no solution for 2-PCRPP exists), and $V(\{c\})$ induces a clique (of size 1) in G .

Now, assume that the lemma holds for every set of required pairs in R (or set of vertices of G) of size k , we show that it holds also for a set of required pairs in R (or set of vertices of G) of size $k + 1$.

Consider a path π that covers a set R' of $k + 1$ required pairs. We show that $V(R')$ induces a clique in G . Let c be a required pair in R' and let $R'' = R' \setminus \{c\}$. By induction hypothesis, $V(R'')$ is a clique of G . Since π passes through all the vertices belonging to required pairs of R' , it follows that there exists a path covering both the required pairs c_i and c , for every $c_i \in R''$. Hence, by construction, $\{v_{c_i}, v_c\} \in E$, for every $v_{c_i} \in V(R'')$, and so we can conclude that $V(R')$ is a clique of G .

Consider a clique $V(R')$ of size $k + 1$. We show that there exists a path covering the set R' of required pairs. Let $c = [n_x, n_y]$ be a required pair in R' and let $R'' = R' \setminus \{c\}$. Clearly, $V(R'')$ induces a clique of size k in G . By induction hypothesis, there exists a path π that covers all the required pairs in R'' . Starting from path π , we can compute (in polynomial time) a path π' that covers R' . Notice that either n_x or n_y does not belong to π , otherwise π would already cover the required pair c . Assume w.l.o.g. that n_x does not belong to π . Since for each vertex $v_{c_i} \in V(R'')$ there exists an edge $\{v_{c_i}, v_c\}$ in G , it follows that for each required pair $c_i \in R''$ there exists a path $\pi_{c_i, c}$ connecting the vertices of c_i and c (hence, in particular, $\pi_{c_i, c}$ passes through n_x). Consider the vertices of R'' in the total order induced by π . There exists a vertex n_z in π (possibly n_z is the source s) such that n_z belongs to a required pair $c_z \in R''$, and n_z is the last vertex in π of a required pair in R'' for which the path $\pi_{c_z, c}$ passes through n_z and then n_x . Let n_{z+1} be the successor of n_z in π and let c_{z+1} the required pair containing n_{z+1} . Notice that $\pi_{c, c_{z+1}}$ passes through n_x and then n_{z+1} . Now, we can compute a path π'' by concatenating the following subpaths: the subpath of π from s to n_z , the subpath of $\pi_{c_z, c}$ from n_z to n_x , the subpath of $\pi_{c_{z+1}, c}$ from n_x to n_{z+1} and the subpath of π from n_{z+1} to t . By construction it is easy to see that π is a subpath of π'' , hence π'' covers all the required pairs in R'' and passes through n_x . A similar construction can be applied to π'' to cover n_y (if n_y does not already belong to π), hence obtaining a path π' that covers every required pair in R' . \square

From Lemma 3, it follows that, in order to compute the existence of a solution of 2-PCRPP over the instance $\langle D = (N, A), R \rangle$ (in which every vertex of D belongs to at least one required pair in R), we have to compute if there exists a 2-Clique Partition of the corresponding graph G . Computing the existence of a 2-Clique Partition over a graph G is equivalent to compute if there exists a 2-Coloring of the complement graph G' (hence deciding if G' is bipartite), which is well-

known to be solvable in polynomial time [8, probl. GT15]. We can conclude that 2-PCRSP can be decided in polynomial time.

4 Parameterized Complexity of MaxRPSP

In this section, we consider the parameterized complexity of MaxRPSP. We show that although MaxRPSP is $W[1]$ -hard when parameterized by the optimum, *i.e.*, the number of required pairs covered by a single path (Section 4.1), the problem becomes fixed-parameter tractable if the maximum number of overlapping required pairs is a parameter (Section 4.2).

4.1 $W[1]$ -hardness of MaxRPSP Parameterized by the Optimum

In this section, we investigate the parameterized complexity of MaxRPSP when parameterized by the size of the solution, that is the maximum number of required pairs covered by a single path, and we prove that the problem is $W[1]$ -hard (notice that this result implies the NP-hardness of MaxRPSP). For details on parameterized reductions, we refer the reader to [4, 9].

We prove this result via a parameterized reduction from the Maximum Clique (MaxClique) problem. Given an undirected graph $G = (V, E)$, MaxClique asks for a clique $C \subseteq V$ of maximum size. Here, we consider the two decision versions of MaxClique and MaxRPSP, h -Clique and k -RPSP respectively, parameterized by the sizes of the respective solutions. For example, given an undirected graph $G = (V, E)$, the h -Clique problem consists of deciding if there exists a clique $C \subseteq V$ of size h . We recall that h -Clique is known to be $W[1]$ -hard [5].

First, we start by showing how to construct an instance of k -RPSP starting from an instance of h -Clique. Given an (undirected) graph $G = (V, E)$ with n vertices v_1, \dots, v_n , we construct the associated directed acyclic graph $D = (N, A)$ as follows. The set N of vertices is defined as:

$$N = \{v_i^z \mid v_i \in V, 1 \leq z \leq h\} \cup \{s, t\}$$

Informally, N consists of two distinguished vertices s, t and of h copies v_i^1, \dots, v_i^h of every vertex v_i of G .

The set of arcs A is defined as:

$$A = \{(v_i^z, v_j^{z+1}) \mid \{v_i, v_j\} \in E, 1 \leq z \leq h-1\} \cup \{(s, v_i^1), (v_i^h, t) \mid v_i \in V\}$$

Informally, we connect every two consecutive copies associated with vertices that are adjacent in G , the source vertex s to all the vertices v_i^1 , with $1 \leq i \leq n$, and all the vertices v_i^h , with $1 \leq i \leq n$, to the sink vertex t .

The set R of required pairs is defined as:

$$R = \{[v_i^x, v_j^y] \mid \{v_i, v_j\} \in E, 1 \leq x < y \leq h\}$$

Informally, for each edge $\{v_i, v_j\}$ of G there is a required pair $[v_i^x, v_j^y]$, $1 \leq x < y \leq h$, between every two different copies associated with v_i, v_j .




Fig. 4. Example of directed acyclic graph $D = (N, A)$ associated with an instance $G = (V, E)$ of the h -Clique problem. Each gray box highlight an independent set I_z composed of one copy of the vertices in V . Edges (v_1^z, v_j^{z+1}) , (v_i^z, v_j^{z+1}) , and (v_n^z, v_i^{z+1}) are some of the directed edges in A associated with edges $\{v_1, v_j\}$, $\{v_i, v_j\}$, $\{v_i, v_n\} \in E$.

By construction, the vertices in N (except for s and t) are partitioned into h independent sets $I_z = \{v_i^z \mid 1 \leq i \leq n\}$, with $1 \leq z \leq h$, each one containing a copy of every vertex of V . Moreover, the arcs of A only connect two vertices of consecutive subsets I_z and I_{z+1} , with $1 \leq z \leq h - 1$. Figure 4 presents an example of directed graph D associated with an undirected graph G .

Now, we are able to prove the main properties of the reduction.

Lemma 4. *Let $G = (V, E)$ be an undirected graph and $\langle D = (N, A), R \rangle$ be the associated instance of k -RPSP. Then: (1) starting from an h -clique in G we can compute in polynomial time an st -path π in D that covers $\binom{h}{2}$ required pairs of R ; (2) starting from an st -path π in D that covers $\binom{h}{2}$ required pairs we can compute in polynomial time an h -clique in G .*

Proof. (1) Starting from an h -clique C in G we show how to compute a path π in D that covers $\binom{h}{2}$ required pairs of R . Let $C = \{v_{i_1}, \dots, v_{i_h}\}$ be a clique of G and let $\langle v_{i_1}, \dots, v_{i_h} \rangle$ be an arbitrary ordering of C . Let $\pi_C = \langle s, v_{i_1}^1, \dots, v_{i_h}^h, t \rangle$ be a sequence of vertices obtained by selecting the vertex $v_{i_z}^z$ for each independent set I_z , with $1 \leq z \leq h$ (in addition to vertices s and t). Since C is a clique of G , by construction of D , every pair of vertices $(v_{i_z}^z, v_{i_{z+1}}^{z+1})$ is connected by an arc, hence π_C is an st -path of D . Moreover, the st -path π_C covers exactly $\binom{h}{2}$ required pairs of R because, by construction of R , there exists a pair between every two copies of vertices which are adjacent in G . More precisely, since the clique C has all the possible edges among its h vertices, the number of required pairs covered by the st -path π_C is $\binom{h}{2}$.

(2) Let π be an st -path in D that covers a set $R' \subseteq R$ of $\binom{h}{2}$ required pairs, then we show how to compute in polynomial time an h -clique C in G . Notice that, by construction of D , the st -path π must contain exactly one vertex v_i^z , $1 \leq i \leq n$ and $1 \leq z \leq h$, for each independent set I_z of D . By construction of set R , each vertex v_i^z of π appears in at most $h-1$ required pairs of R' . Hence, the total number of required pairs covered by the st -path π , which contains exactly h inner vertices v_i^z , is at most $\frac{h(h-1)}{2} = \binom{h}{2}$. Let C be the set $\{v_i \mid v_i^z \in \pi \setminus \{s, t\}\}$. We claim that C is an h -clique. First, we prove that C contains h vertices. Suppose to the contrary that C has less than h vertices. Then, there exist two vertices v_i^x and $v_{i'}^y$ of π that correspond to the same vertex v_i of C , that is $i' = i'' = i$. Since $[v_i^x, v_i^y] \notin R$, it follows that each v_i^x, v_i^y appears in at most $h-2$ required pairs of R' . As a consequence, the total number of required pairs covered by the st -path π is strictly less than $\binom{h}{2}$, violating the initial hypothesis that π covers $\binom{h}{2}$ required pairs of R . Hence C contains h vertices. As all the internal vertices of π (*i.e.*, all its vertices but s and t) represent distinct vertices of G , then all the required pairs covered by π represent distinct edges of G . The only undirected graph with h vertices and $\binom{h}{2}$ edges is the complete graph, hence C is an h -clique of G . \square

The W[1]-hardness of k -RPSP follows easily from Lemma 4.

Theorem 2. *k -RPSP is W[1]-hard when parameterized by the number of required pairs covered by an st -path.*

Proof. The result follows from Lemma 4 and from the W[1]-hardness of h -Clique when parameterized by h [5]. \square

4.2 An FPT Algorithm for MaxRPSP Parameterized by the Maximum Overlapping Degree

In this section we propose a parameterized algorithm for the MaxRPSP problem, where the parameter is the maximum overlapping degree of the required pairs in R . For the rest of the section, let $\langle D = (N, A), R \rangle$ be an instance of the MaxRPSP problem. A required pair $[u', v'] \in R$ is *maximal* if it is not nested in another required pair $[u'', v'']$.

For ease of exposition, we fix an order of the required pairs in R and we represent the i -th required pair of the ordering as $[v_i^1, v_i^2]$. Whenever no confusion arises, we will refer to that required pair as i -pair. Intuitively, we want that the order of the required pairs is “compatible” with the topological order of the vertices. More formally, given two distinct required pairs $[v_j^1, v_j^2]$ and $[v_i^1, v_i^2]$ with $j < i$, then pair $[v_j^1, v_j^2]$ is nested in $[v_i^1, v_i^2]$ or there does not exist a path π from s to v_j^2 that covers both the required pairs (that is, π passes through v_j^2 before v_i^2). Clearly, an order that satisfies this condition can be easily computed from the topological order of the vertices.

We present a parameterized algorithm based on dynamic programming for the MaxRPSP problem when the parameter p is the maximum number of overlapping required pairs. In fact, we can decompose a path π , starting in s , ending

in a vertex v , and covering k required pairs, into two subpaths: the first one— π_1 —starts in s , ends in a vertex v' , and covers k_1 required pairs, while the other one— π_2 —starts in v' , ends in v , and covers the remaining $k_2 = k - k_1$ required pairs (possibly using vertices of π_1). The key point to define the dynamic programming recurrence is that, for each required pair p , we keep track the set of required pairs overlapping p and covered by the path. To this aim, for each required pair $[v_i^1, v_i^2]$, we define the set $OP([v_i^1, v_i^2])$ as the set of vertices v such that v belongs to a required pair that overlaps $[v_i^1, v_i^2]$ and such that v_i^2 is reachable from v . By a slightly abuse of the notation, we consider that $OP([v_i^1, v_i^2])$ always contains vertex v_i^1 .

The recurrence relies on the following observation. Let π be a path covering a set P of required pairs and let $N(P)$ be the set of vertices belonging to the required pairs in P . Consider two required pairs $[v_i^1, v_i^2]$ and $[v_j^1, v_j^2]$ in P , with $j < i$. Then, either $[v_j^1, v_j^2]$ is nested in $[v_i^1, v_i^2]$ (hence the fact that π covers the pair $[v_j^1, v_j^2]$ can be checked by the recurrence looking only at the required pairs that overlap with $[v_i^1, v_i^2]$) or pairs $[v_i^1, v_i^2]$ and $[v_j^1, v_j^2]$ are alternated. In the latter case, since $[v_i^1, v_i^2]$ is in P , we only have to consider the vertices in the set $N(P) \cap OP([v_i^1, v_i^2]) \cap OP([v_j^1, v_j^2])$. Moreover, let p_i be the number of required pairs that overlap the required pair $[v_i^1, v_i^2]$, then $|OP([v_i^1, v_i^2])|$ is at most $2p_i$. Hence, the cardinality of set $N(P) \cap OP([v_i^1, v_i^2]) \cap OP([v_j^1, v_j^2])$ is bounded by $2 \max(p_i, p_j)$. Moreover, given two sets S and S' of vertices such that $S \subseteq OP([v_i^1, v_i^2])$ and $S' \subseteq OP([v_j^1, v_j^2])$, we say that S is in *agreement* with S' if $S \cap (OP([v_i^1, v_i^2]) \cap OP([v_j^1, v_j^2])) = S' \cap (OP([v_i^1, v_i^2]) \cap OP([v_j^1, v_j^2]))$. Informally, when S and S' are in agreement, they must contain the same subset of vertices of $OP([v_i^1, v_i^2]) \cap OP([v_j^1, v_j^2])$.

Let $P([v_i^1, v_i^2], S)$ denote the maximum number of required pairs covered by a path π ending in vertex v_i^2 and such that the set $S \subseteq OP([v_i^1, v_i^2])$ is covered by π . In the following we present the recurrence to compute $P([v_i^1, v_i^2], S)$. For ease of exposition we only focus on vertices that appear as second vertices of the required pairs. In fact, paths that do not end in such vertices are not able to cover new required pairs. Furthermore, for simplicity, we consider the source s as the second vertex of a fictitious required pair (with index 0) $[\perp, s]$ which does not overlap any other required pair. Such a fictitious required pair does not contribute to the total number of required pairs covered by the path.

The recurrence is:

$$P([v_i^1, v_i^2], S) = \max_{\substack{[v_j^1, v_j^2] \text{ not nested in } [v_i^1, v_i^2] \text{ and } j < i; \\ S' \text{ in agreement with } S; \\ \exists \text{ a path from } v_j^2 \text{ to } v_i^2 \\ \text{covering all vertices in } S \setminus S';}} \{P([v_j^1, v_j^2], S') + |Ov([v_i^1, v_i^2], S, S')|\} \quad (1)$$

where $Ov([v_i^1, v_i^2], S, S') = \{[v_h^1, v_h^2] \mid [v_h^1, v_h^2] \text{ is nested in } [v_i^1, v_i^2] \wedge v_h^1 \in S \wedge v_h^2 \in S \setminus S'\}$. Notice that each required pair is assumed to be nested in itself.

The base case of the recurrence is $P([\perp, s], \emptyset) = 0$.

The correctness of the recurrence derives from the following two lemmas.

Lemma 5. *If $P([v_i^1, v_i^2], S) = k$, then there exists a path π in D ending in v_i^2 , such that every vertex in S belongs to π and the number of required pairs covered by π is k .*

Proof. We prove the lemma by induction on the index i . It is easy to see that the base case holds. Assume that the lemma holds for index values less than i , we prove that the lemma holds for i . Let $P([v_i^1, v_i^2], S) = k$. By Eq. (1), there exists a vertex v_j^2 with $j < i$, such that $P([v_j^1, v_j^2], S') = k_1$ for some set S' in agreement with S . Assume that $|Ov([v_i^1, v_i^2], S, S')| = k_2$, with $k_1 + k_2 = k$. By induction hypothesis, since $P([v_j^1, v_j^2], S') = k_1$, there exists a path π' ending in v_j^2 , covering every vertex in S' , and such that π' covers k_1 required pairs. Furthermore, the k_2 covered required pairs have at least one vertex in $S \setminus S'$, hence the vertices of such required pairs belong to a path π'' which starts in v_j^2 and ends in v_i^2 (path π'' exists by hypothesis). But then, the path obtained by the concatenation of π' and π'' covers $k_1 + k_2$ required pairs. \square

Lemma 6. *Let π be a path in D ending in v_i^2 and covering k required pairs. Let S be the set of all the vertices belonging to required pairs covered by π and overlapping $[v_i^1, v_i^2]$. Then $P([v_i^1, v_i^2], S) \geq k$.*

Proof. We prove the lemma by induction on the index i . It is easy to see that the base case holds. Assume that the lemma holds for index values less than i , we prove that the lemma holds for i . Let π be a path, ending in v_i^2 , that covers k required pairs and let S be the set of vertices that belong to the required pairs covered by π and overlapping $[v_i^1, v_i^2]$. We claim that $P([v_i^1, v_i^2], S) \geq k$. Consider the rightmost vertex v_j^2 of π such that v_j^2 belongs to a required pair covered by π and not nested in the i -pair. Decompose path π into two parts: one— π' —from s to v_j^2 , and the other one— π'' —from v_j^2 to v_i^2 . Let S' be the set of vertices that belong to the required pairs covered by π and overlapping $[v_j^1, v_j^2]$. Let k' be the number of required pairs covered by π' and k'' be the number of the remaining required pairs covered by π (that is, $k = k' + k''$). First, notice that $k'' = |Ov([v_i^1, v_i^2], S, S')|$. By induction hypothesis $P([v_j^1, v_j^2], S') = k_1$ for some $k_1 \geq k'$. Moreover, by construction, S' is in agreement with S and the subpath of π from v_j^2 to v_i^2 covers all the vertices in $S \setminus S'$. As a consequence, by Eq. (1), $P([v_i^1, v_i^2], S)$ is at least $k_1 + k'' \geq k' + k'' = k$, which concludes the proof. \square

Let p be the maximum number of overlapping required pairs in D (that is, $p = \max_i \{p_i\}$). It follows that the number of possible subsets S is bounded by $O(2^p)$. Then, each entry $P([v_i^2], S)$ requires time $O(2^p n)$ to be computed, and, since there exist $O(2^p n)$ entries, the recurrence requires time $O(4^p n^2)$. From Lemma 5 and Lemma 6, it follows that an optimal solution for MaxRPSP can be obtained by looking for the maximum of the values $P([v_i^2], S)$. Hence, the overall time complexity of the algorithm is bounded by $O(4^p n^2)$.

References

1. Bao, E., Jiang, T., Girke, T.: BRANCH: boosting RNA-Seq assemblies with partial or related genomic sequences. *Bioinformatics* 29(10), 1250–1259 (2013)

2. Bonizzoni, P., Dondi, R., Pirola, Y.: Maximum disjoint paths on edge-colored graphs: Approximability and tractability. *Algorithms* 6(1), 1–11 (2013)
3. Dilworth, R.P.: A decomposition theorem for partially ordered sets. *Annals of Mathematics* 51(1), 161–166 (1950)
4. Downey, R., Fellows, M.: *Parameterized Complexity*. Springer Verlag (1999)
5. Downey, R.G., Fellows, M.R.: Fixed-parameter tractability and completeness II: On completeness for $W[1]$. *Theoretical Computer Science* 141(1&2), 109–131 (1995)
6. Eriksson, N., Pachter, L., Mitsuya, Y., Rhee, S.Y., Wang, C., Gharizadeh, B., Ronaghi, M., Shafer, R.W., Beerenwinkel, N.: Viral population estimation using pyrosequencing. *PLoS Comput Biol* 4(5), e1000074 (2008)
7. Ford, L.R., Fulkerson, D.R.: *Flows in Networks*. Princeton university press (1962)
8. Garey, M., Johnson, D.: *Computer and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman (1979)
9. Niedermeier, R.: *Invitation to Fixed-Parameter Algorithms*. Oxford University Press (2006)
10. Ntafos, S., Hakimi, S.: On path cover problems in digraphs and applications to program testing. *IEEE Transactions on Software Engineering* 5(5), 520–529 (1979)
11. Trapnell, C., Williams, B.A., Pertea, G., Mortazavi, A., Kwan, G., van Baren, M.J., Salzberg, S.L., Wold, B.J., Pachter, L.: Transcript assembly and quantification by RNA-Seq reveals unannotated transcripts and isoform switching during cell differentiation. *Nature Biotechnology* 28(5), 516–520 (2010)
12. Wu, B.Y.: On the maximum disjoint paths problem on edge-colored graphs. *Discrete Optimization* 9(1), 50–57 (2012)