Journal of Combinatorial Theory, Series B 106 (2014) 1–14

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series B

www.elsevier.com/locate/jctb

Vertex-disjoint directed and undirected cycles in general digraphs

Jørgen Bang-Jensen, Matthias Kriesell, Alessandro Maddaloni, Sven Simonsen

IMADA, University of Southern Denmark, Campusvej 55, DK-5230 Odense M, Denmark

ARTICLE INFO

Article history:
Received 29 June 2011
Available online 17 March 2014

Keywords:
Cycle
Dicycle
Disjoint cycle problem
Mixed problem
Cycle transversal number
Intercyclic digraphs

ABSTRACT

The *dicycle transversal number* $\tau(D)$ of a digraph D is the minimum size of a *dicycle transversal* of D, i.e. a set $T \subseteq V(D)$ such that $D - T$ is acyclic. We study the following problem: Given a digraph D, decide if there is a dicycle B in D and a cycle C in its underlying undirected graph $UG(D)$ such that $V(B) \cap V(C) = \emptyset$. It is known that there is a polynomial time algorithm for this problem when restricted to strongly connected graphs, which actually finds B, C if they exist. We generalize this to any class of digraphs D with either $\tau(D) \neq 1$ or $\tau(D) = 1$ and a bounded number of dicycle transversals, and show that the problem is $\mathcal{NP}$-complete for a special class of digraphs D with $\tau(D) = 1$ and, hence, in general.

1. Introduction

All graphs and digraphs are assumed to be finite, and they may contain loops or multiple arcs or edges. Notation follows [1], and we recall the most relevant concepts here. In order to distinguish between directed cycles in a digraph D and cycles in their *underlying graph* $UG(D)$ we use the term *dicycle* for a directed cycle in D and *cycle* for a cycle in $UG(D)$. Whenever we consider a (directed) path P containing vertices a, b such that a precedes b on P, we denote by $P[a,b]$ the subpath of P which starts in a and ends in b. Similarly, we denote by $P(a,b]$, $P[a,b)$, and $P(a,b)$, respectively, the subpath that starts in the successor of a on P and ends in b, starts in a and ends in the predecessor of b, and starts in the successor of a on P and ends in the predecessor of b, respectively. The same notation applies to dicycles.

A digraph D is *acyclic* if it does not contain a dicycle, and it is *intercyclic* if it does not contain two disjoint dicycles. A *dicycle transversal* of D is a set S of vertices of D such that $D - S$ is acyclic, and the *dicycle transversal number* $\tau(D)$ is defined to be the size of a smallest dicycle transversal. McCuaig characterized the intercyclic digraphs of minimal in- and out-degree at least 2 in terms of their dicycle

0095-8956/$ – see front matter
http://dx.doi.org/10.1016/j.jctb.2013.10.005

transversal number and designed a polynomial time algorithm that, for any digraph, either finds two disjoint cycles or a structural certificate for being intercyclic [10].

Theorem 1. *(See [10].) There exists a polynomial time algorithm which decides whether a given digraph is intercyclic and finds two disjoint cycles if it is not.*

The undirected graphs without two disjoint cycles have been characterized by Lovász [9], generalizing earlier statements of Dirac for the 3-connected case [7]. The characterization again implies a polynomial algorithm for finding such cycles if they exist, see e.g. [5].

Here we are concerned with the following problem.

Problem 1. Given a digraph D, decide if there is a dicycle B in D and a cycle C in $UG(D)$ with $V(B) \cap V(C) = \emptyset$.

The motivation for studying this problem comes from [2] where a mixed variant of the subdigraph homeomorphism problem has been studied. The problem of deciding if, for a given digraph D and $b, c \in V(D)$, there exist disjoint dicycles B, C in D with $b \in V(B)$ and $c \in V(C)$ is known to be $\mathcal{NP}$-complete by the classic dichotomy of Fortune, Hopcroft, and Wyllie on the *fixed directed subgraph homeomorphism problem* [8]: For some pattern digraph H, not part of the input, we want to decide for an input digraph D and an injection f from $V(H)$ to $V(D)$ if we can extend f on $V(H) \cup A(H)$ such that every loop at x maps to a cycle of D containing $f(x)$, every arc xy with $x \neq y$ maps to an $(f(x), f(y))$-path, and the resulting paths and cycles are *internally disjoint*, i.e. no internal vertex of either object is a vertex of another one.[1] The dichotomy then states that the problem is solvable in polynomial time if the arcs of H have the same initial vertex or if they have the same terminal vertex, and is $\mathcal{NP}$-complete in all other cases [8].

In [2], an extension of this has been studied, where H might be a *mixed graph*, having both arcs and edges, and the edges of H are asked to be mapped to cycles and paths[2] of $UG(D)$.

We find it surprising that, as a consequence of the resulting dichotomy, the problem is already $\mathcal{NP}$-complete as soon as there is both an arc and an edge in the pattern graph. In particular, the problem of deciding whether for a digraph D and $b, c \in V(D)$ there exists a cycle B in D and a cycle C in $UG(D)$ with $b \in V(B)$, $c \in V(C)$, and $V(B) \cap V(C) = \emptyset$, is $\mathcal{NP}$-complete. The proof shows that even the weaker problem to decide whether for a digraph D and $c \in V(D)$ there exists a cycle B in D and a cycle $C \in UG(D)$ with $c \in V(C)$ and $V(B) \cap V(C) = \emptyset$ is $\mathcal{NP}$-complete, even if we are assuming that, in addition, D is strongly connected.

So the question arose what happens if we do not prescribe vertices at all, leading to Problem 1.

The first two authors showed in [3] that Problem 1 is solvable in polynomial time when D is strongly connected. The solution turned out to be more complex than expected, and builds on McCuaig's results on intercyclic digraphs [10], Thomassen's results on 2-linkages in acyclic digraphs [11], and a new reduction algorithm for digraphs with dicycle transversal number 1.

Theorem 2. *(See [3].) There is a polynomial algorithm for Problem 1 restricted to strongly connected digraphs. Furthermore, one can find the desired cycles in polynomial time if they exist.*

In this paper, based on the complete characterization from [3] of those strongly connected digraphs with dicycle transversal number 2 which are no-instances for Problem 1 (see Theorem 5), we will show that there is a polynomial time algorithm for Problem 1 restricted to digraphs with $\tau(D) \geqslant 2$. After this we show that Problem 1 is $\mathcal{NP}$-complete for digraphs with $\tau(D) = 1$, and, hence, $\mathcal{NP}$-complete in general. Finally we show that the problem is still polynomial for the class

[1] Where, in the case of a cycle C of D assigned to a loop of H at x, we consider its internal vertices to be all but $f(x)$.

[2] We are always assuming that D and $UG(D)$ have the same set of vertices and arcs, respectively, i.e. they differ only by means of incidence relations.

of digraphs where the dicycle transversal number is 1 but the number of transversal vertices is constantly bounded.

The case $\tau(D) \geqslant 3$ is easily dealt with due to the following result from [3].

Theorem 3. *(See [3].) If D is a strongly connected digraph with $\tau(D) \geqslant 3$ then there is a dicycle B in D and a cycle C in $UG(D)$ with $V(B) \cap V(C) = \emptyset$, and we can find such cycles in polynomial time.*

Since a digraph with at least two non-trivial strong components of size greater than 1 has two disjoint dicycles and the dicycle transversal number of any digraph D is the sum of the dicycle transversal numbers of its strongly connected components, we get, as an immediate consequence:

Corollary 4. *There exists a polynomial time algorithm for Problem 1 restricted to digraphs with dicycle transversal number at least 3, which finds the desired cycles.*

Trivially, acyclic digraphs are no-instances to Problem 1, so let us assume that the digraphs D under consideration have at least one dicycle. McCuaig's algorithm from [10] finds two disjoint dicycles in D if they exist. If they do not exist, then we know that the digraphs D under consideration have exactly one non-trivial strong component D', where $\tau(D') = \tau(D) \in \{1, 2\}$. We then apply the algorithm from [3] to D'; if D' is a yes-instance to Problem 1 then so is D, and hence we can assume that D' is a no-instance to Problem 1. For $\tau(D) = 2$, we employ the complete characterization of no-instances in [3] and derive a polynomial time algorithm which takes the (undirected) cycles which are in D but not in D' into account to produce a correct answer. For digraphs with $\tau(D) = 1$ we give an algorithm with running time $\ell(D)^{k(D)} \cdot p(|V(D)|)$, where p is a polynomial, $k(D)$ is the number of dicycle transversal vertices of D, and $\ell(D)$ is the maximum number of disjoint paths between a pair of distinct transversal vertices. Since $\ell(D) \leqslant |V(D)|$, this is a polynomial time algorithm if we are in any class of digraphs with $\tau(D) = 1$ and a constantly bounded number of dicycle transversals. In Section 4 we give a proof that Problem 1 is $\mathcal{NP}$-complete for a certain class of digraphs with dicycle transversal number 1 (and hence in general) by providing a two-step reduction from 3-SAT to Problem 1.

2. Strongly connected digraphs with dicycle transversal number 2

In this section we describe the characterization in [3] of the strongly connected no-instances to Problem 1. They fall into three infinite classes called multiwheels, trivaults and vaults.

A *multiwheel* MW_p is obtained from a directed cycle $c_0c_1 \dots c_{p-1}c_0$, $p \geqslant 3$, by adding a new vertex v and adding, for each $i \in \{0, \dots, p-1\}$, ℓ_i arcs from v to c_i and k_i arcs from c_i to v where $\ell_i + k_i \geqslant 1$. A *split multiwheel* SMW_p is obtained from a multiwheel MW_p by replacing the central vertex v by two vertices v^+, v^-, adding the arc v^-v^+, and letting all arcs entering (leaving) v in MW_p enter v^- (leave v^+). See Fig. 1. The vertices v or v^+, v^- are called the *central vertices* of the multiwheel or split multiwheel, respectively.

A *trivault* is obtained from six disjoint digraphs R_i, L_i, $i \in \{0, 1, 2\}$, where each R_i is either a non-trivial out-star[3] with root b_i or a (b_i, x_i)-path and each L_i is either a non-trivial in-star with root c_i or a (y_i, c_i)-path, as follows:

(i) for each $i \in \{0, 1, 2\}$ either add a single arc from c_i to b_i or identify b_i, c_i,
(ii) for distinct $i, j \in \{0, 1, 2\}$, if R_i is a non-trivial out-star and L_j is a non-trivial in-star, add a single arc from each leaf of R_i to c_j and from b_i to every leaf of L_j and an arbitrary number of arcs (possibly 0) from b_i to c_j,

[3] An *out-star* is the digraph formed by a collection of arcs $uv_1, uv_2, \dots, uv_p$, none of which form a loop, for some $p \geqslant 1$. Similarly and in-star is a collection of arcs all of which have the same head u.

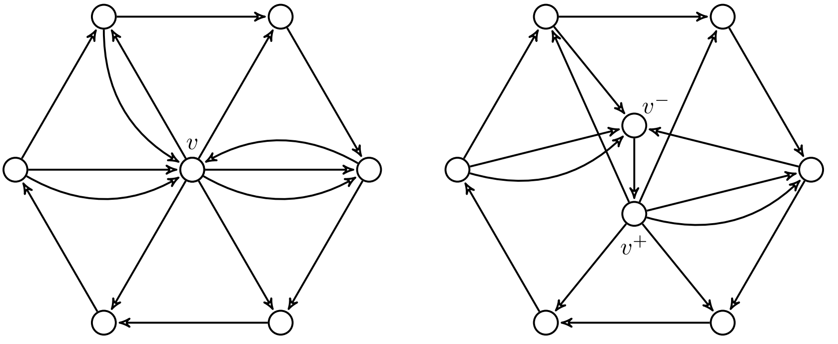

Fig. 1. The left part shows a multiwheel with centre v and the right one shows the corresponding split multiwheel obtained by splitting v into v^-, v^+.

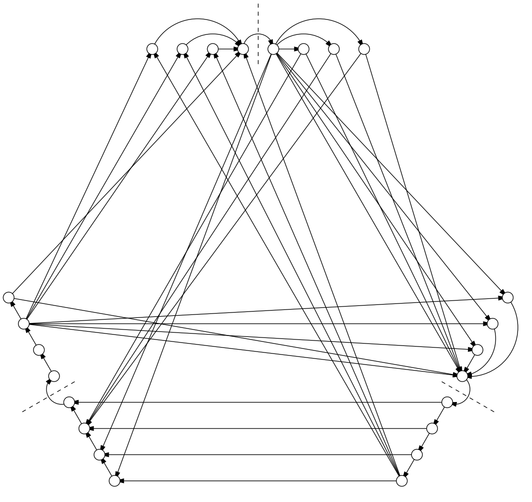

Fig. 2. A typical trivault. The dotted lines indicate either an arc $c_i b_i$ or that $b_i = c_i$.

(iii) for distinct $i, j \in \{0, 1, 2\}$, if R_i is a non-trivial out-star and L_j is a path, select $v \in L_j$ and add a single arc from each leaf of R_i to v, at least one arc from b_i to y_j, and an arbitrary number of arcs (possibly 0) from b_i to each $z \in L_j[y_j, v]$,

(iv) similarly, for distinct $i, j \in \{0, 1, 2\}$, if R_i is a path and L_j is a non-trivial in-star, select $v \in R_i$ and add a single arc from v to each leaf of L_j, at least one arc from x_i to c_j, and an arbitrary number of arcs (possibly 0) from each $z \in R_i[v, x_i]$ to c_j, and

(v) if, for distinct $i, j \in \{0, 1, 2\}$, R_i, L_j are paths, then add at least one arc from x_i to some vertex of L_j, and at least one arc from some vertex of R_i to y_j, and add an arbitrary number of arcs (possibly 0) from each $z \in R_i$ to each $w \in L_j$.

Fig. 2 shows a typical trivault.

We say that a trivault has a *niche* if there are distinct $i, j, k \in \{0, 1, 2\}$ such that either

(a) R_i, L_j are paths and there are arcs pq, rs such that p occurs before r on R_i and q occurs after s on L_j, or

(b) R_i is a path, containing an in-neighbour x of L_k such that there are at least two arcs from $R_i(x, x_i]$ to L_j, or

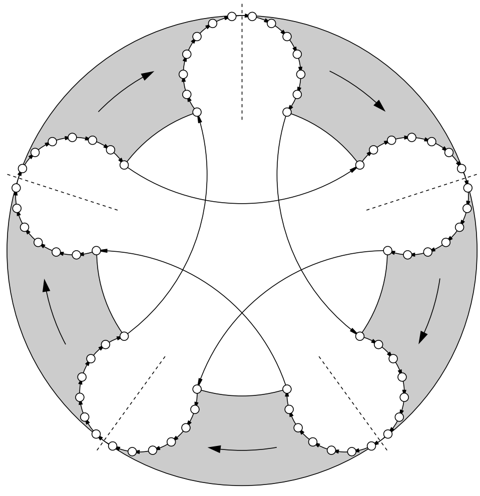

Fig. 3. A typical vault. The five central arcs must have multiplicity 1 and are the only arcs from P_i to P_{i+2}.

(c) L_i is a path containing an out-neighbour y of R_k such that there are at least two arcs from R_j to $L_i[y_i, y)$.

Observe that every trivault is strongly connected. It might contain a vertex v of in- and out-degree 1; however, this is either in some path $R_i - x_i$ or in some path $L_i - y_i$, and contracting any arc (on that path) incident with v produces, consequently, a trivault again; this smaller trivault will have a niche if and only if the original one had a niche. Hence we can consider every trivault as a subdivision of a (primal) trivault without vertices of in- and out-degree 1, which has a niche if and only if the primal trivault had one.

Finally, let us describe the vaults. Let $\ell \geqslant 5$ be odd, let $P_0, \dots, P_{\ell-1}$ be disjoint non-empty paths such that for each $i \in \{0, \dots, \ell-1\}$, P_i is an (a_i, d_i)-path and b_i, c_i are vertices of P_i such that either $b_i c_i$ is an arc on P_i or $b_i = c_i \in \{a_i, d_i\}$. Suppose that D is obtained from the disjoint union of the paths P_i, $i \in \{1, 2, \dots, \ell-1\}$, by

(i) adding at least one arc from some vertex in $P_i[c_i, d_i]$ to some vertex from $P_{i+1}[a_{i+1}, b_{i+1}]$ (multiarcs may occur), and
(ii) adding a single arc from d_i to a_{i+2}, for all $i \in \{0, \dots, \ell-1\}$,

where the indices are taken modulo ℓ. Any digraph of such a form is called a *vault*, and the P_i are called its *walls*. We say that the vault D has a *niche*, if there exist arcs pq, rs from some P_i to P_{i+1} such that p occurs before r on P_i and q occurs after s on P_{i+1}. In that case,

$$P_i[a_i, p]P_{i+1}[q, d_{i+1}]P_{i+3}[a_{i+3}, d_{i+3}] \dots P_{i-2}[a_{i-2}, d_{i-2}]a_i$$

is a dicycle of D, disjoint from the cycle of $UG(D)$ constituted by the path

$$P_i[r, d_i]P_{i+2}[a_{i+2}, d_{i+2}]P_{i+4}[a_{i+4}, d_{i+4}] \dots P_{i-1}[a_{i-1}, d_{i-1}]P_{i+1}[a_{i+1}, s]$$

and the arc rs. Fig. 3 shows a vault with $\ell = 5$, where all paths $P_i[a_i, b_i]$ or $P_i[c_i, d_i]$ have seven vertices; the grey areas indicate the set of arcs connecting $P_i[c_i, d_i]$ to $P_{i+1}[a_{i+1}, b_{i+1}]$, a niche would correspond to a pair of arcs which can be drawn without crossing in such an area. Vaults are strongly connected digraphs as they have a spanning dicycle. They may contain vertices of both in- and out-degree 1, but, as such vertices occur only as internal vertices of the P_i, we deduce that every vault D is a subdivision of a vault $\tilde{D}$ without vertices of in- and out-degree 1, where $\tilde{D}$ has a niche if and only if D has. Allowing $\ell = 3$ in the definition of vaults will produce some, but not all trivaults.

Now we are ready to state the characterization from [3] of the strongly connected no-instances.

Theorem 5. *(See [3].) Let $D = (V, A)$ be a strongly connected digraph with dicycle transversal number 2. In polynomial time we can either find a cycle B in D and a cycle C in $UG(D)$ with $V(B) \cap V(C) = \emptyset$ or show that D has no such cycles in which case D satisfies one of the following.*

(i) *D is a subdivision of either a multiwheel or a split multiwheel.*
(ii) *D is a subdivision of a trivault without a niche.*
(iii) *D is a subdivision of a vault without a niche.*

Furthermore, if D satisfies one of (i)–(iii)*, we can produce a certificate for this in polynomial time.*

3. Digraphs D with $\tau(D) = 2$

We now look at the case that our input digraph D to Problem 1 has dicycle transversal number 2. As we have mentioned in the introduction, we may assume that D is not strongly connected and has exactly one non-trivial component D', where, moreover, D' is a no-instance. By Theorem 5, D' is either the subdivision of a multiwheel (split or not), a niche-free trivault or a niche-free vault.

Clearly, if $UG(D - V(D'))$ contains a cycle then D is a yes-instance. Hence we may assume that $UG(D - V(D'))$ is a forest, that is, all connected components of $UG(D - V(D'))$ are trees. We observe that if D is a yes-instance, then there exists a cycle C in $UG(D)$ disjoint from some dicycle B in D' such that C traverses every connected component H of $UG(D - V(D'))$ at most once (for if C traverses H then we consider a component P of $C - V(H)$ and the — not necessarily distinct — neighbours h, h' of the end vertices of P in H on C, and replace the (h, h')-path $C - V(P)$ with the (h, h')-path in H as to obtain a cycle C' disjoint from B traversing H only once). Thus we loose no information by contracting every connected component of $D - V(D')$ to a single vertex, and reorienting all arcs between a vertex of $D - V(D')$ and D' so that they all start in D'. Hence, from now on, we may assume that $D - V(D')$ consists of independent vertices, which we call the *external vertices*. Since $\tau(D') = 2$ we may assume that D has no parallel arcs, due to the following.

Lemma 6. *If D has a pair of parallel arcs, then either D is a yes-instance or we may delete one of the parallel arcs without changing the problem.*

Proof. Suppose e and f are both arcs from u to v. If one of u, v belongs to $D - V(D')$ then D is a yes-instance since $\tau(D') = 2$, so we may assume that u, v are vertices of D'. Then (as D' itself is a no-instance) we can delete one of e, f without changing the problem. □

We now further simplify the problem by observing that each of the following operations can be applied to D without changing a no-instance into a yes-instance or vice versa. We repeat doing any one of these as long as possible, while always calling the resulting graph D and its non-trivial strong component D' and observing that the dicycle transversal number does not change either.[4]

(i) If there is more than one arc from u to v check if $\{u, v\}$ is a dicycle transversal. If not, then D is a yes-instance (take uvu as the undirected cycle). Otherwise we delete all but one copy of uv.
(ii) Delete all external vertices with degree at most 1 (they are on no cycle).
(iii) Contract the outgoing arc of a vertex v with $d^-_D(v) = 1 = d^+_D(v)$.

We now analyze connections between pairs of vertices in D' and external vertices. Since D' is a no-instance, any undirected cycle C (partly) certifying a yes-instance must use at least one external vertex. It is possible to show that if D is a yes-instance, then we can choose C such that it contains at

[4] Note that some of the operations we perform may create new parallel arcs, which is why we need (i).

most two external vertices. However, we will illustrate this only for the vault-case (in which case one external vertex always suffices), whereas for multiwheels and trivaults the total number of dicycles is polynomial in $|V(D)|$ hence allowing for a brute force algorithm.

By Theorem 5, D' is a subdivision of some graph D'_0, where D'_0 is either a multiwheel or a split multiwheel, a trivault without a niche or a vault without a niche. We proceed by distinguishing cases accordingly, starting with the ones that are easiest to analyze.

3.1. Multiwheels and split multiwheels

Assume first that D' is a subdivision of a multiwheel or split multiwheel, with cycle C and central vertex a or central vertices a^-, a^+, respectively. This case is particularly easy.

Theorem 7. *If the unique non-trivial strong component D' of D is a subdivision of a multiwheel or of a split multiwheel then Problem 1 is solvable in polynomial time.*

Proof. A dicycle in D' is either a subdivision of C, or a subdivision of a cycle W which is formed by two arcs ca, $a'c'$ where a, a' are central vertices (possibly the same) and $c, c' \in V(C)$ together with the unique (c', c)-path in C and the arc aa' if $a \neq a'$. Hence there are only $O(|V(D')|^2)$ dicycles, and for each such dicycle B we check if $UG(D) - V(B)$ contains a cycle. This leads straightforwardly to a cubic time algorithm as desired. □

3.2. Trivaults

Assume next that D' is a subdivision of a trivault D'_0, and let L_i, R_i, b_i, c_i for $i \in \{0, 1, 2\}$ be as in the definition of a trivault (with D'_0 instead of D). Again, we have good control on the dicycles:

Theorem 8. *If the unique non-trivial strong component D' of D is a subdivision of a trivault then Problem 1 is solvable in polynomial time.*

Proof. Set $X_i := L_i \cup R_i$ for $i \in \{0, 1, 2\}$. If a dicycle W in D'_0 contains a vertex of X_i then it enters X_i via an arc from a vertex of some R_j with $j \neq i$ to some $\ell \in L_i$, and it exits X_i via an arc from some $r \in R_i$ to some vertex from L_k with $k \neq i$. Moreover, W will contain the unique ℓ, r-path in X_i and, in particular, b_i and c_i — hence it cannot traverse X_i more than once. Therefore, every dicycle in D_0 is formed by either (i) a pair (a, b), (c, d) of arcs with $a \in R_i$, $b \in L_j$, $c \in R_j$, $d \in L_i$, where $i \neq j$ together with the unique (b, c)-path in X_j and the unique (d, a)-path in X_i, or (ii) a triple (a, b), (c, d), (e, f) with $a \in R_0$, $b \in L_1$, $c \in R_1$, $d \in L_2$, $e \in R_2$, $f \in L_0$ together with the unique (b, c)-path in X_1, the unique (d, e)-path in X_2, and the unique (f, a)-path in X_0, or (iii) a triple (a, b), (c, d), (e, f) with $a \in R_0$, $b \in L_2$, $c \in R_2$, $d \in L_1$, $e \in R_1$, $f \in L_0$ together with the unique (b, c)-path in X_2, the unique (d, e)-path in X_1, and the unique (f, a)-path in X_0. As the dicycles in D' are obtained by those in D'_0 by subdividing arcs, there are only $O(|E(D'_0)|^3)$ many dicycles in D, and we can construct them easily. For each such dicycle B we check if $UG(D) - V(B)$ contains a cycle. This leads straightforwardly to a polynomial time algorithm as desired. □

3.3. Vaults

Let us now treat the final and most complicated case which is when D' is a subdivision of a niche-free vault D'_0 with walls P_i and let a_i, b_i, c_i, d_i be vertices on P_i as in the definition of a vault, $i \in \{0, \dots, \ell - 1\}$ (all indices modulo ℓ). We start with a couple of new definitions.

A pair $(\{u, v\}, \alpha)$ is called a *k-clasp* if α is an external vertex, u, v are neighbours of α, and there exists a cycle C^* in $UG(D)$ containing u, v, α and at most k external vertices such that there exists a dicycle B^* in $D - V(C^*)$. By definition, there cannot be a 0-clasp, and by what we have seen before, u, v need to be distinct. Observe that, since D' is a no-instance, D is a yes-instance if and only if there exists a k-clasp for some k.

Given an arc $pq \in D'_0$, we denote by $\widehat{pq}$ the corresponding subdivision dipath in D' and call it, for brevity, a *link*. A link of length 1 is called trivial.

We may assume that consecutive arcs on P_i are not subdivided (so that the P_i are paths in D', too). If $\widehat{d_i a_{i+2}}$ is non-trivial then we enlarge the wall P_i by $\widehat{d_i a_{i+2}} - a_{i+2}$ and redefine d_i accordingly. Hence we may assume that non-trivial links always connect consecutive walls.

Lemma 9. *There is always a directed cycle avoiding any prescribed wall, but there is no directed cycle avoiding two consecutive walls.*

Proof. Note that the subdigraph consisting of the walls $P_{i-1}, P_{i+1}, P_{i+3}, \ldots, P_{i-4}, P_{i-2}$ and all links between them contains a directed cycle avoiding P_i. On the other hand a directed cycle avoiding walls P_i, P_{i+1}, if it existed, could not contain vertices of P_{i+2}, because each vertex of $V(P_{i+2})$ has in-degree 0 in $D' - V(P_i) \cup V(P_{i+1})$. Repeating this argument inductively one sees that no wall could be part of the cycle, hence such a cycle cannot exist. $\square$

Lemma 10. *If u, v are distinct neighbours of an external vertex α and u, v are either on the same wall or on distinct non-consecutive walls then $(\{u, v\}, \alpha)$ is a 1-clasp.*

Proof. If u and v are on the same wall P_i, then by Lemma 9, there is a directed cycle avoiding P_i, which is therefore disjoint from the undirected cycle containing α, u, v and using only vertices from $V(P_i) \cup \{\alpha\}$. If u and v are not on the same P_i it is possible to relabel the walls in such a way that u is on the wall P_0 and v is on wall P_{2k}, where $0 < 2k < \ell - 1$. The undirected cycle $P_0[u, d_0]P_2P_4 \ldots P_{2k-2}P_{2k}[a_{2k}, v]\alpha u$ is therefore disjoint from any directed cycle contained in the subdigraph induced by $P_1, P_3, P_5, \ldots, P_{\ell-4}, P_{\ell-2}, P_{\ell-1}$ and all the links between them. $\square$

Let b'_i (c'_i) be the last (first) vertex on P_i such that there exists a link from a vertex in P_{i-1} to b'_i (from c'_i to a vertex in P_{i+1}). A pair $(\{u, v\}, \alpha)$ is a *pin* if α is an external vertex, u, v are neighbours of α and, possibly after switching the names of u and v, there exists an $i \in \{0, \ldots, \ell - 1\}$ such that u is in $P_i[b'_i, d_i]$ and v is in $P_{i+1}[a_{i+1}, c'_{i+1}]$ and such that there is no link $\widehat{pq}$ with p in $P_i[a_i, u)$ and $q \in P_{i+1}(v, d_{i+1}]$.

The following theorem classifies all sets $\{u, v\}$ of two distinct vertices from D' with a common external neighbour α: Either $\{u, v\}$ is a dicycle transversal of D', or $(\{u, v\}, \alpha)$ is a 1-clasp. (Hence if there is no 1-clasp in D at all then all such $\{u, v\}$ are dicycle transversals of D', so that we cannot find a k-clasp for any k, and hence D is a no-instance for Problem 1.)

Theorem 11. *Let $u \neq v$ be vertices from D' with a common external neighbour α.*

(i) *$(\{u, v\}, \alpha)$ is a pin if and only if $\{u, v\}$ is a dicycle transversal of D'.*
(ii) *$(\{u, v\}, \alpha)$ is not a pin if and only if $(\{u, v\}, \alpha)$ is a 1-clasp.*

Proof. Since it is not possible that $\{u, v\}$ is a dicycle transversal of D' while $(\{u, v\}, \alpha)$ is a 1-clasp, it suffices to prove the only-if-parts of (i) and (ii).

For (i), suppose that $(\{u, v\}, \alpha)$ is a pin. By switching the names of u and v if necessary, we can obtain that u is on some wall P_i and v is on P_{i+1} (as in the definition of a pin). We show that the walls P_i and P_{i+1}, cannot be part of a dicycle which avoids u, v and then use Lemma 9 to conclude that $\{u, v\}$ is a dicycle transversal. If we remove u then, as u does not occur before b'_i, the vertex set of path starting from u's out-neighbour on P_i (if one exists) and ending at d_i has in-degree 0 in D' and hence cannot be contained in a directed cycle, so we can remove it during our search. Symmetrically, the vertex set of the path starting from a_{i+1} and ending at v's in-neighbour (if one exists) on P_{i+1} has out-degree 0 in D' and can be removed. At this stage the set consisting of the remaining vertices on P_i and the set consisting of the remaining vertices on P_{i+1} have 0 out-degree and in-degree respectively, hence they cannot be part of a dicycle. Now Lemma 9 implies that $\{u, v\}$ is a dicycle transversal.

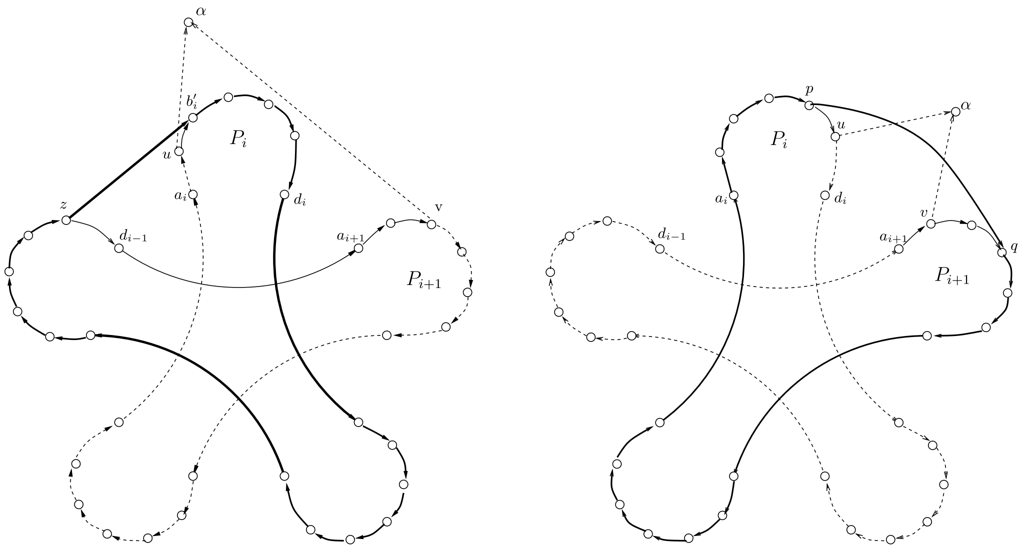

Fig. 4. Possible 1-clasps formed by an external vertex with two neighbours in consecutive walls of a vault. The directed cycle is indicated in bold and the other cycle by dashed arcs.

For (ii), suppose that $(\{u,v\},\alpha)$ is not a pin. We prove that $(\{u,v\},\alpha)$ is a 1-clasp. First consider the case of u, v both being on walls: If u, v are on the same wall or on distinct non-consecutive walls, then Lemma 10 guarantees that $(\{u,v\},\alpha)$ is a 1-clasp. So let us assume that there exists an $i \in \{0,\dots,\ell-1\}$ such that u is on P_i and $v \in P_{i+1}$. If u comes before b_i' on P_i then there is a link $\widehat{zb_i'}$, with $z \in P_{i-1}$, and the directed cycle $\widehat{zb_i'}P_i(b_i',d_i]P_{i+2}P_{i+4}\dots P_{i-1}[a_{i-1},z]$ is disjoint from the undirected cycle $P_{i+1}[v,d_{i+1}]P_{i+3}P_{i+5}\dots P_i[a_i,u]\alpha v$ (see the left part of Fig. 4). Symmetrically, if v comes after c_{i+1}' then there is a link $\widehat{c_{i+1}'z'}$, with $z' \in P_{i+2}$ and the directed cycle $\widehat{c_{i+1}'z'}P_{i+2}(z',d_{i+2}]P_{i+4}\dots P_{i+1}[a_{i+1},c_{i+1}']$ is disjoint from the undirected cycle $P_{i+1}[v,d_{i+1}]P_{i+3}\dots P_i[a_i,u]\alpha v$. Hence we may assume that u is in $P_i[b_i',d_i]$ and v is in $P_{i+1}[a_{i+1},c_{i+1}']$. Since $(\{u,v\},\alpha)$ is not a pin, there exists a link $\widehat{pq}$ with p coming before u on P_i and q coming after v on P_{i+1}. Now the directed cycle $\widehat{pq}P_{i+1}(q,d_{i+1}]P_{i+3}\dots P_i[a_i,p]$ is disjoint from the undirected cycle $P_i[u,d_i]P_{i+2}\dots P_{i+1}[a_{i+1},v]\alpha u$ (see the right part of Fig. 4), certifying that $(\{u,v\},\alpha)$ is a 1-clasp.

Now consider the case that one of u, v, say, u, is not on a wall and, hence, an internal vertex of a link $\widehat{u_1u_2}$ is between two consecutive walls. Define similarly v_1, v_2 if v is not on a wall, and $v_1 = v_2 = v$ otherwise (in which case $\widehat{v_1v_2}$ is the trivial path formed by $v_1 = v_2$). If there exist $g,h \in \{1,2\}$ such that u_g and v_h are on the same wall P_i then the subdigraph induced by $UG(D[V(P_i) \cup \{\alpha\} \cup V(\widehat{u_1u_2}) \cup V(\widehat{v_1v_2})])$ contains a cycle which avoids all walls except for P_i, and hence, by Lemma 9, $(\{u,v\},\alpha)$ is a 1-clasp. So assume there is no such choice for g, h.

Now it is easy to see that we may choose g, h such that u_g and v_h are on distinct non-consecutive walls. Thus we can relabel everything in the same way as in the proof of Lemma 10 having one of u_g, v_h on P_0 and the other on P_{2k}, where $0 < 2k < \ell-1$. If u_g in P_0 and v_h in P_{2k} then let U_g, respectively V_h denote the (u_g,u)-path contained in $\widehat{u_1u_2}$, respectively the (v,v_h)-path contained in $\widehat{v_1v_2}$ and let R be the path joining u_g and v_h through walls $P_0,P_2,\dots,P_{2k}$. Now $U_gRV_h\alpha u$ forms a cycle in $UG(D)$ disjoint from any directed cycle contained in the subdigraph induced by $P_1,P_3,\dots,P_{\ell-2},P_{\ell-1}$ and all the links between them. Likewise, if v_h in P_0 and u_g in P_{2k}, we have that $V_hRU_g\alpha v$ is the desired cycle in $UG(D)$ which is disjoint from any dicycle contained in the subdigraph induced by $P_1,P_3,\dots,P_{\ell-2},P_\ell$ and all links between these. □

Theorem 12. *If the unique non-trivial strong component D' of D is a subdivision of a vault then Problem 1 is solvable in polynomial time.*

Proof. We first reduce to the situation described immediately before Lemma 9. For every $\alpha \in D - V(D')$ consider the sets $\{u, v\}$ formed by two distinct neighbours u, v of α. For each such $(\{u, v\}, \alpha)$ it takes polynomial time to check if $(\{u, v\}, \alpha)$ is a pin (according to (i) of Theorem 11 this is equivalent to check whether $D' - \{u, v\}$ is acyclic). As soon as $(\{u, v\}, \alpha)$ is not a pin, it must be a 1-clasp and one gets the two cycles as in the proof of Theorem 11. If all $(\{u, v\}, \alpha)$ turn out to be pins, then no external vertex can be part of a solution and hence D is a no-instance. □

4. Digraphs D with $\tau(D) = 1$

The aim of this section is to prove that Problem 1 is $\mathcal{NP}$-complete for digraphs with transversal number 1 and an unbounded number of transversal vertices. We start with a quite different $\mathcal{NP}$-complete problem on bipartite graphs and then show how to reduce from this problem.

Problem 2. Let $G = (U, V, E)$ be a bipartite graph with colour classes U and V, and suppose that $V_1, \ldots, V_k$ form a partition of V. Decide if there exists a cycle C in G which avoids at least one vertex from each V_i.

Lemma 13. *Problem 2 is $\mathcal{NP}$-complete.*

Proof. We will show how to reduce 3-SAT to Problem 2 in polynomial time. Let $W[u, v, p, q]$ be the graph with vertices $\{u, v, y_1, y_2, \ldots, y_p, z_1, z_2, \ldots, z_q\}$ and the edges of the two (u, v)-paths $uy_1y_2\ldots y_pv$ and $uz_1z_2\ldots z_qv$. Graphs of this type will form the *variable gadgets*.

Let $\mathcal{F}$ be an instance of 3-SAT with variables $x_1, x_2, \ldots, x_n$ and clauses $C_1, C_2, \ldots, C_m$. We may assume without loss of generality that each variable x occurs at least once in either the negated or the non-negated form in $\mathcal{F}$. The ordering of the clauses $C_1, C_2, \ldots, C_m$ induces an ordering of the occurrences of a variable x and its negation $\bar{x}$ in these. With each variable x_i we associate a copy of $W[u_i, v_i, 2p_i + 1, 2q_i + 1]$ where x_i occurs p_i times and $\bar{x}_i$ occurs q_i times in the clauses of $\mathcal{F}$. Initially, these copies are assumed to be disjoint, but we chain them up by identifying v_i and u_{i+1} for each $i \in \{1, 2, \ldots, n-1\}$. Let $s = u_1$ and $t = v_n$. Let G' be the graph obtained in this way. Observe that G' is bipartite since each $W[u_i, v_i, 2p_i + 1, 2q_i + 1]$ is the union of two even length (u_i, v_i)-paths.

For each $i \in \{1, 2, \ldots, m\}$ we associate the clause C_i with three of the vertices $V_i = \{a_{i,1}, a_{i,2}, a_{i,3}\}$ (this is the *clause gadget*) from the graph G' above as follows: assume C_i contains variables x_j, x_k, x_ℓ (negated or not). If x_j is not negated in C_i and this is the rth occurrence of x_j (in the order of the clauses that use x_j), then we identify $a_{i,1}$ with $y_{j,2r-1}$ and if C_i contains $\bar{x}_j$ and this is the hth occurrence of $\bar{x}_j$, then we identify $a_{i,1}$ with $z_{j,2h-1}$. We proceed similarly with x_k, $a_{i,2}$ and x_ℓ, $a_{i,3}$, respectively. Thus G' contains all the vertices $a_{j,i}$, $j \in \{1, \ldots, m\}$, $i \in \{1, 2, 3\}$.

Claim. *G' contains an (s,t)-path P which avoids at least one vertex from $\{a_{j,1}, a_{j,2}, a_{j,3}\}$ for each $j \in \{1, \ldots, m\}$ if and only if $\mathcal{F}$ is satisfiable.*

For a proof, suppose P is an (s,t)-path which avoids at least one vertex from $\{a_{j,1}, a_{j,2}, a_{j,3}\}$ for each $j \in \{1, \ldots, m\}$. By construction of G', for each variable x_i, P traverses either the subpath $u_iy_{i,1}y_{i,2}\ldots y_{i,2p_i+1}v_i$ or the subpath $u_iz_{i,1}z_{i,2}\ldots z_{i,2q_i+1}v_i$. Now define a truth assignment by setting x_i false if and only if the first traversal occurs for i. This is a satisfying truth assignment for $\mathcal{F}$ since for any clause C_j at least one literal is avoided by P and hence becomes true by the assignment (the literals traversed become false and those not traversed become true). Conversely, given a truth assignment for $\mathcal{F}$ we can form P by routing it through all the false literals in the chain of variable gadgets. This proves the claim.

Now let $B = (U, V, E)$ be the bipartite graph with colour classes U, V which we obtain from G' by adding new vertices β_1, β_2 and the edges $s\beta_1$, $s\beta_2$, $\beta_1 t$, $\beta_2 t$. Here V is the vertex set $\{\beta_1, \beta_2\} \cup$

$\{y_{i,2j+1}\colon\ i \in \{1,\dots,n\},\ \ j \in \{0,1,\dots,p_i\}\} \cup \{z_{i,2j+1}\colon\ i \in \{1,\dots,n\},\ \ j \in \{0,1,\dots,q_i\}\}$, and U is the set of the remaining vertices. For each $i \in \{1,\dots,n\}$ let $V'_i = \{y_{i,2p_i+1}, z_{i,2q_i+1}\}$ and let $V'_{n+1} = \{\beta_1, \beta_2\}$. Then $V_1, V_2, \dots, V_m, V'_1, \dots, V'_n, V'_{n+1}$ form a partition of V.

By the construction of G that every cycle C distinct from the 4-cycle $s\beta_1 t\beta_2 s$ is either formed by one of the subgraphs $W[u_i, v_i, 2p_i+1, 2q_i+1]$ or consists of an (s,t)-path in G and one of the two (t,s)-paths $t\beta_1 s$, $t\beta_2 s$.

We show that G has a cycle C which avoids at least one vertex from each of the sets $V_1, V_2, \dots, V_m, V'_1, \dots, V'_n, V'_{n+1}$ if and only if $\mathcal{F}$ is satisfiable. This follows from our claim and the fact that the definition of V'_i, $i \in \{1,\dots,n\}$, and V'_{n+1} implies that the desired cycle exists if and only if G' has an (s,t)-path which avoids at least one vertex from $V_j = \{a_{j,1}, a_{j,2}, a_{j,3}\}$ for each $j \in \{1,\dots,m\}$. Note that the sets V'_i, $i \in \{1,\dots,n\}$, exclude cycles of the form $W[u_i, v_i, p_i, q_i]$ and V'_{n+1} excludes the cycle $s\beta_1 t\beta_2 s$. □

Theorem 14. *Problem 1 is $\mathcal{NP}$-complete.*

Proof. We describe a polynomial reduction from Problem 2 to Problem 1 restricted to the case of dicycle transversal number 1 and an unbounded number of transversal vertices.

Let $H = (U, V, E)$ be a bipartite graph with colour classes U, V where $U = \{b_1, \dots, b_r\}$, and $V = V_1 \cup V_2 \cup \dots \cup V_k$ with $V_i = \{p_{i,1}, \dots, p_{i,\ell_i}\}$, $\ell_i > 0$ for $i \in \{1,2,\dots,k\}$, and $V_i \cap V_j = \emptyset$ if $i \neq j$.

We construct a directed graph D from H in the following way: Direct all edges of H from V to U. Add vertices $v_0, \dots, v_{k+1}$. Add arcs $v_{i-1}p_{i,j}$ and $p_{i,j}v_i$ for all $i \in \{1,\dots,k\}$, $j \in \{1,\dots,l_i\}$. Finally add the arc $v_k v_0$.

We claim that D contains a dicycle B and a cycle C of $UG(D)$ which are disjoint if and only if there is a cycle in H avoiding at least one vertex of V_i for each $i \in \{1,2,\dots,k\}$. First suppose there is a cycle in H avoiding the vertex p_{i,a_i} of V_i for each $i \in \{1,2,\dots,k\}$. Then, by the construction of D, the same cycle will be a cycle in $UG(D)$. The cycle $v_0 p_{1,a_1} v_1 p_{2,a_2} \dots v_{k-1} p_{k,a_k} v_k v_0$ is vertex disjoint from this undirected cycle, and we are done.

Now suppose there is an undirected cycle C disjoint from some dicycle in D. Note that every dicycle in D is formed by the arc $v_k v_0$ and some (v_0, v_k)-path. The path is of the form $v_0 p_{1,a_1} v_1 \dots v_{k-1} p_{k,a_k} v_k$. Hence C does not contain any of the vertices $v_0, v_1, \dots, v_k$ and hence uses only $p_{i,j}$ or b_ℓ vertices and always alternates between them. Therefore C has a corresponding cycle in H, and this one avoids at least the vertex p_{i,a_i} from of the set V_i for each $i \in \{1,\dots,k\}$. □

5. Digraphs D with $\tau(D) = 1$ and a bounded number of dicycle transversals

Consider a digraph D with $\tau(D) = 1$. We show that if there is a bounded number of transversal vertices then our problem is polynomially decidable. We preprocess D as in the first paragraphs of Section 3 and then delete each arc connecting a transversal vertex with an external vertex. We also delete all but one copy of any collection of parallel arcs between two transversal vertices. None of the deleted arcs will ever be used to certify a yes-instance because every transversal vertex is contained in the directed cycle. After this process we delete external vertices with degree at most 1.

Let W be an arbitrary dicycle of D and let $a, a_1, \dots, a_{k-1}$ be the transversal vertices of D, in the order they show up on W. Since each a_i is a transversal, every dicycle must visit the transversal vertices in the same cyclic order as W. Build a new acyclic digraph $\tilde{D}$ by splitting a into an outgoing part a_0 and an ingoing part a_k. All arcs leaving (entering) a now leave a_0 (enter a_k). Given the preprocessed graph Problem 1 is equivalent to that of finding in $\tilde{D}$ a directed (a_0, a_k)-path disjoint from an undirected cycle. Note that all transversal vertices are (a_0, a_k)-separators in $\tilde{D}$, and every (a_0, a_k)-path contains $a_0, a_1, \dots, a_k$ in that order.

For $x \in \{1,\dots,k\}$, fix a largest system $\mathcal{P}^x$ of internally disjoint (a_{x-1}, a_x)-paths, say, $P^x_1, \dots, P^x_{\ell^x}$, and let $P^* := \bigcup_{x=1}^{k} \bigcup_{i=1}^{\ell_x} P^x_i$ be the digraph formed by the union of all these paths. Note that no vertex in $\tilde{D}$ except $a_1, \dots, a_{k-1}$ belongs to more than one system $\mathcal{P}^x$. Furthermore, by Menger's theorem, for every $x \in \{1,2,\dots,k\}$ we have

$$\text{either } \mathcal{P}^x \text{ is just the arc } a_{x-1}a_x \text{ or } \ell_x \geqslant 2. \tag{1}$$

Suppose that there exists an (a_0, a_k)-dipath C in $\tilde{D}$ and a cycle C' in $UG(\tilde{D})$ disjoint from C (that is, the original D is a yes-instance). We call such a pair C, C' a *solution*. In what follows we will assume that

$$\text{there is no solution } C,\ C' \text{ such that } A(C) \subseteq A\left(P^*\right). \tag{2}$$

Our algorithm below will find any solution of the excluded type in a preprocessing step. We now show that we can choose a solution C, C' such that C changes from one path to another at most once in any of the path systems $\mathcal{P}^x$ above. Let C, C' be a solution and such that the number of arcs of $C \cup C'$ which are not in the path system is as small as possible, that is,

$$\left|A\left(C \cup C'\right) \setminus A\left(P^*\right)\right| \text{ is minimum over all solutions.} \tag{3}$$

For each path P_i^x as defined above, let $Q_{i,1}^x, \ldots, Q_{i,h_i^x}^x$ be the connected components of $C \cap P_i^x$ ordered such that $Q_{i,j}^x$ is before $Q_{i,j'}^x$ on P_i^x if $j < j'$. Likewise, let $R_{i,1}^x, \ldots, R_{i,k_i^x}^x$ be the connected components of $P_i^x - V(C)$, if any, ordered in the same way as before. Let $b_{i,j}^x$ and $c_{i,j}^x$ be the first and the last vertex of $Q_{i,j}^x$, respectively. With this notation we have $a_{x-1} = b_{i,1}^x$ and $a_x = c_{i,h_i^x}^x$ for all i.

Claim 1. *For all x, i, the dipath C visits $Q_{i,1}^x, \ldots, Q_{i,h_i^x}^x$ in this order.*

Suppose C first visits $Q_{i,j'}^x$ and then some $Q_{i,j}^x$, with $j < j'$. Then $\tilde{D}$ contains the dicycle $C[b_{i,j'}^x, b_{i,j}^x]P_i^x[b_{i,j}^x, b_{i,j'}^x]$, contradicting that $\tilde{D}$ is acyclic. This proves Claim 1.

For a subdigraph H of a digraph D, let $d_D(H)$ denote the number of edges in D having exactly one end vertex in H.

Claim 2. *For all x, i either P_i^x is a subpath of C or for all $j \in \{1, \ldots, k_i^x\}$, we have $d_{C'}(R_{i,j}^x) = 2$.*

To see this, observe that if P_i^x is not a subpath of C, then $d_{C'}(R_{i,j}^x)$ is even, and positive, for otherwise, by replacing the $(b_{i,j}^x, b_{i,j+1}^x)$-subpath of C by $P_i^x[b_{i,j}^x, b_{i,j+1}^x]$ we get an (a_0, a_k)-path which is still disjoint from C' but gives a lower value for (3). If $d_{C'}(R_{i,j}^x) \geqslant 4$, then the digraph induced by $V(C') \cup V(R_{i,j}^x)$ contains a cycle C'' (which is disjoint from C by the definition of $R_{i,j}^x$) such that replacing C' by C'' yields a lower value for (3). This proves Claim 2.

Claim 3. *For all x, either P_i^x is a subpath of C or P_i^x does not contain the arc $c_{i,j}^x b_{i,j+1}^x$ for any $j \in \{1, \ldots, h_i^x - 1\}$.*

Otherwise we could replace the $(c_{i,j}^x, b_{i,j+1}^x)$-subpath of C by this arc and get, again, a smaller value for (3). This proves Claim 3.

We define a *bridge* as the subdigraph of $\tilde{D}$ formed by either a single arc of $A(\tilde{D}) - A(P^*)$ connecting two vertices of P^*, or the arcs incident with the vertices of a connected component of $UG(\tilde{D} - V(P^*))$. By (1) and (2), a bridge contains neither two interior vertices of any P_i^x, nor a cycle of $UG(\tilde{D})$.

A *switch* is a maximal subpath of C of length at least 1 such that all its edges and internal vertices belong to some bridge. It is then evident that a switch is a (v, w)-subpath of a single bridge where v is contained in some P_i^x and w is contained in some P_j^y. Since $\tilde{D}$ is acyclic, $y \geqslant x$, but if $y > x$ then C misses a_x, a contradiction. Hence $x = y$, and we call the switch, more specifically, an *x-switch*. We may achieve that $i \neq j$, for suppose that v, w are both from P_i^x. If P_i^x was the only path in $\mathcal{P}^x$ then, by (1) it has length 1 but then $v = a_{x-1}$ and $w = a_x$, so that our switch is internally disjoint from P_i^x, contradicting the maximality of $|\mathcal{P}^x|$. So $\mathcal{P}^x$ contains at least two paths P_i^x, P_j^x, $i \neq j$. If both v, w are internal vertices of P_i^x, then the switch and the subpath of P_i^x from v to w contains a cycle $\tilde{C}$ which is disjoint from some (a_0, a_k)-path C^* in P^*, contradicting (2). Hence we may assume that at least one of v, w is on P_j^x, too.

Claim 4. *For every $x \in \{1, 2, \ldots, k\}$, there is at most one x-switch.*

For suppose, to the contrary, that, for some x, there are at least two x-switches, and consider the first two along C. Suppose the first one is from P_i^x to P_j^x, where $i \neq j$. Then the second one is from P_j^x to some P_k^x where $j \neq k$. By Claim 3, $P_i^x \setminus C$ has at least one non-empty component, so consider $R_{i,1}^x$, and $P_j^x \setminus C$ has at least two, so consider $R_{j,1}^x$ and $R_{j,2}^x$. By Claim 2, exactly one of the two $(R_{j,1}^x, R_{j,2}^x)$-subpaths of C' misses $R_{i,1}^x$ (for otherwise $d_{C'}(R_{i,1}^x) \neq 2$). Let us denote this path by M. But then one could change C using $P_i^x[b_{i,1}^x, b_{i,2}^x]$ instead of $C[b_{i,1}^x, b_{i,2}^x]$, where the latter contains $Q_{j,2}^x$. Now $M \cup R_{j,1}^x \cup R_{j,2}^x \cup Q_{j,2}^x$ contains an undirected cycle C' disjoint from the new dicycle C, and together they achieve a lower value for (3). This contradiction proves Claim 4.

Now we have obtained enough structural information to construct a polynomial algorithm: D is a yes-instance if and only if it either contains a solution C, C' with C in P^* or it contains a solution C, C' which satisfies Claims 1 to 4 and these observations can be used to search for a solution efficiently.

Theorem 15. *For fixed k, there is a polynomial time algorithm that decides whether a given digraph D with $\tau(D) = 1$ and at most k dicycle transversal vertices has a dicycle B in D and a cycle C in $UG(D)$ with $V(B) \cap V(C) = \emptyset$, and finds these cycles if they exist.*

Proof. The problem is polynomially equivalent to finding C, C' in $\tilde{D}$ as in the first two paragraphs of this section (or decide that they do not exist). All further objects, in particular suitable maximal path systems $\mathcal{P}^x$, can be computed in polynomial time.

We first iterate through all k-tuples $\pi = (\pi_1, \ldots, \pi_k)$, where, for each x, π_x is a path from $\mathcal{P}^x$. There are less than $|V(D)|^k$ choices. For each π, set $C_\pi := \bigcup_{x=1}^k \pi_k$ and check if $UG(\tilde{D} \setminus C_\pi)$ has a cycle C'. All that can be done in polynomial time, and we stop (with a yes-instance) as soon as we find such a C'.

Now we are in a stage where (2) holds and hence a solution (if any exists) must use at least one switch and the considerations including Claim 4 guarantee that there exists a solution if and only if there is solution C, C' which uses at most one x-switch for each x.

For all pairs of arcs (e, f) such that e starts in some P_i^x and f ends in some P_j^x we check whether there is a dipath starting with e and ending with f without internal vertices from $V(P^*)$. This can be done in polynomial time, and such a path $H_{e,f}$ is uniquely determined because otherwise there would be a cycle C' in $UG(\tilde{D} \setminus P^*)$, contradicting that (2) holds. A path $H_{e,f}$ as above might serve as an x-switch for more than one pair of paths P_i^x, P_j^x if one and hence (by the maximality of $\mathcal{P}^x$, $x \in \{1, 2, \ldots, k\}$) only one of its end vertices is a transversal vertex; we can maintain a list of the options for each of them and this list has length at most $|V(D)|$. The number of hypothetical x-switches for each x is thus bounded by $|A(D)|^2$, hence we find all of them, plus their lists, in polynomial time.

Now we iterate through all k-tuples $\pi = (\pi_1, \ldots, \pi_k)$, where, for each x, π_x is either a path from $\mathcal{P}^x$ or a hypothetical x-switch $H_{e,f}$ connecting P_i^x, P_j^x with $i \neq j$. (Moreover, we may assume that not all of the π_x are paths from $\mathcal{P}^x$, as such a π has been considered earlier above.) There are far less than $(|A(D)|^2 + 1)^k$ choices for π here. For each π, construct a dipath C_π as follows: For each hypothetical x-switch π_x, say, starting at u and ending at v, take its union with the unique (a_{x-1}, u)- and the unique (v, a_x)-path in $\bigcup_{i=1}^{\ell_x} P_i^x$. Take the union of all these paths and of those π_x which have been selected as paths from $\mathcal{P}^x$ and call it C_π. It is clear that if C, C' as desired exist then $C = C_\pi$ for some π. Hence it suffices to check if $\tilde{D} \setminus C_\pi$ has a cycle C', for all C_π. All that can be done in polynomial time. □

6. Final remarks

The arc-disjoint analogue of Problem 1 seems at first glance simpler, but turns out to be $\mathcal{NP}$-complete when the digraphs in question may have arbitrarily many transversal arcs (a transversal arc of a digraph D is an arc belonging to every dicycle of D), as was recently shown by the authors of this paper. The algorithms for the polynomial cases use, among other things, the algorithms for the polynomial cases of Problem 1. The proof of the $\mathcal{NP}$-completeness result is more involved than that of this paper.

Theorem 16. *(See [4].) The problem of deciding whether a given input digraph D contains a dicycle B and a cycle C in $UG(D)$ such that $A(B) \cap A(C) = \emptyset$ is polynomially solvable for the class of strong digraphs and for every class of digraphs with a constantly bounded number of transversal arcs. In the remaining case, when the digraphs may have arbitrarily many transversal arcs, the problem is $\mathcal{NP}$-complete.*

It is well known that the following problems are both polynomially solvable (see e.g. [6] and [1]):

- Given a graph G; decide whether G contains a 2-factor, that is, a spanning collection of vertex disjoint cycles.
- Given a digraph D; decide whether D has a cycle factor, that is, a spanning collection of vertex disjoint dicycles.

Problem 3. What is the complexity of deciding, for a given digraph D, whether $UG(D)$ contains a 2-factor with cycles $C_1, C_2, \dots, C_k$ such that one of these cycles is a directed cycle of D?

Notice that, as far as we know, the complexity of deciding whether a digraph has a cycle factor with at least two cycles is open (see Problem 13.6.13 in [1]). This is why we do not formulate the problem as straightforward generalization (insisting that $k \geqslant 2$) of the problem studied in this paper.

References

[1] J. Bang-Jensen, G. Gutin, Digraphs. Theory, Algorithms and Applications, second edition, Springer Monogr. Math., Springer-Verlag, London, Ltd., London, 2009.

[2] J. Bang-Jensen, M. Kriesell, Disjoint directed and undirected paths and cycles in digraphs, Theoret. Comput. Sci. 46–49 (2009) 5138–5144.

[3] J. Bang-Jensen, M. Kriesell, On the problem of finding disjoint cycles and dicycles in a digraph, Combinatorica 31 (6) (2011) 639–668.

[4] J. Bang-Jensen, M. Kriesell, A. Maddaloni, S. Simonsen, Arc-disjoint directed and undirected cycles in a general digraph, 2013, submitted for publication.

[5] H. Bodlaender, On disjoint cycles, Internat. J. Found. Comput. Sci. 5 (1990) 59–68.

[6] J.A. Bondy, U.S.R. Murty, Graph Theory, Grad. Texts in Math., vol. 244, Springer, 2008.

[7] G.A. Dirac, Some results concerning the structure of graphs, Canad. Math. Bull. 6 (1963) 183–210.

[8] S. Fortune, J. Hopcroft, J. Wyllie, The directed subgraph homeomorphism problem, Theoret. Comput. Sci. 10 (1980) 111–121.

[9] L. Lovász, On graphs not containing independent circuits, Mat. Lapok 16 (1965) 289–299 (in Hungarian).

[10] W. McCuaig, Intercyclic digraphs, in: Graph Structure Theory, Seattle, WA, 1991, in: Contemp. Math., vol. 147, Amer. Math. Soc., Providence, RI, 1993, pp. 203–245.

[11] C. Thomassen, The 2-linkage problem for acyclic digraphs, Discrete Math. 55 (1985) 73–87.