Arc-Disjoint Directed and Undirected Cycles in Digraphs

Jørgen Bang-Jensen,[1,*] Matthias Kriesell,[2]
Alessandro Maddaloni,[3,†] and Sven Simonsen[4]

[1] *DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE*
UNIVERSITY OF SOUTHERN DENMARK
CAMPUSVEJ 55, DK–5230 ODENSE M, DENMARK
E-mail: jbj@imada.sdu.dk

[2] *INSTITUT FÜR MATHEMATIK*
TECHNISCHE UNIVERSITÄT ILMENAU
WEIMARER STRAßE 25, D–98693 ILMENAU, GERMANY
E-mail: matthias.kriesell@tu-ilmenau.de

[3] *TECIP INSTITUTE SCUOLA SUPERIORE SANT' ANNA*
VIA MORUZZI 1, PISA 56124, ITALY
E-mail: a.maddaloni@sssup.it

[4] *DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE*
UNIVERSITY OF SOUTHERN DENMARK
CAMPUSVEJ 55, DK–5230 ODENSE M, DENMARK
E-mail: contact@svensimonsen.com

Received September 25, 2013; Revised September 29, 2015

Published online 17 December 2015 in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/jgt.22006

*Contract grant sponsor: Danish council for independent research; contract grant number: 1323-00178B.
†Contract grant sponsor: Danish council for independent research; contract grant number: 1323-00178B.

Abstract: The **dicycle transversal number** $\tau(D)$ of a digraph D is the minimum size of a **dicycle transversal** of D, that is a set of vertices of D, whose removal from D makes it acyclic. An arc a of a digraph D with at least one cycle is a **transversal arc** if a is in every directed cycle of D (making $D - a$ acyclic). In [3] and [4], we completely characterized the complexity of 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 turns out that the problem is polynomially solvable for digraphs with a constantly bounded number of transversal vertices (including cases where $\tau(D) \geq 2$). In the remaining case (allowing arbitrarily many transversal vertices) the problem is NP-complete. In this article, we classify the complexity of the arc-analog of this problem, where we ask for a dicycle B and a cycle C that are arc-disjoint, but not necessarily vertex-disjoint. We prove that the problem is polynomially solvable for strong digraphs and for digraphs with a constantly bounded number of transversal arcs (but possibly an unbounded number of transversal vertices). In the remaining case (allowing arbitrarily many transversal arcs) the problem is NP-complete. J. Graph Theory 83: 406–420, 2016

Keywords: *cycle; dicycle; disjoint cycle problem; arc-disjoint cycle problem; mixed problem; cycle transversal number; transversal arc*

AMS classification: *05c38; 05c20; 05c85*

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 undirected cycles in its **underlying graph** $UG(D)$ we use the term **dicycle** for a directed cycle in D and **cycle** for an undirected 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 $\mathbf{P[a, b]}$ the subpath of P that starts in a and ends in b. The same notation applies to dicycles. An (s, t)-path is a directed path from s to t. A digraph is **strongly connected** (or just **strong**) if it contains an (s, t)-path for all choices of vertices s, t.

The **in-degree** $\mathbf{d^-(v)}$ (**out-degree** $\mathbf{d^+(v)}$) of a vertex v is the number of arcs of the form xv (vy) and the **degree** $\mathbf{d(v)}$ of v is the sum of its in- and out-degrees. We denote by $\delta(D)$ the minimum degree of a vertex in D.

A digraph D is **acyclic** if it does not contain a dicycle, and it is **(arc-) intercyclic** if it does not contain two (arc-)disjoint dicycles. A **dicycle transversal** of D is a set S of vertices of D whose removal makes D acyclic, and the **dicycle transversal number** $\tau(D)$ is defined to be the size of a smallest dicycle transversal. If D is not acyclic but $D - v$ ($D - a$) is acyclic for some vertex $v \in V(D)$ (arc $a \in A(D)$), then v (a) is a **transversal vertex (arc)** of D. McCuaig characterized the intercyclic digraphs of minimal in- and out-degree at least two in terms of their dicycle transversal number and designed a polynomial time algorithm that, for any digraph, either finds two disjoint cycles or a structural certificate for being intercyclic [12].

Theorem 1. *[12] There exists a polynomial time algorithm that 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 [11], generalizing earlier statements of DIRAC for the 3-connected case [9]. The characterization again implies a polynomial algorithm for finding such cycles if they exists, see for example [7].

In [3, 4] we completely characterized the complexity of the following problem that can be seen as lying in-between the two problems of two disjoint cycles in a graph and two disjoint dicycles in a digraph.

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 results are as follows.

Theorem 2. *[3] There is a polynomial algorithm for Problem 1 restricted to strong digraphs. Furthermore, one can find the desired cycles in polynomial time if they exist and in the case where $\tau(D) = 2$ and no solution exists, a structural certificate for this can be produced in polynomial time.*

Theorem 3. *[4] There is a polynomial time algorithm for Problem 1 restricted to digraphs with a constantly bounded number of transversal vertices. The problem is NP-complete for the class of all digraphs with $\tau(D) = 1$.*

In this article, we are concerned with the arc-disjoint version of Problem 1.

Problem 2. *Given a digraph D, decide if there is a dicycle B in D and a cycle C in $UG(D)$ with $A(B) \cap A(C) = \emptyset$.*

The motivation for studying this problem comes from our work on Problem 1 as well as from several other recent articles dealing with mixed structures in digraphs, see for example [2, 5]. Studying natural problems of this kind may provide more insight into structural properties of directed graphs that may not become visible from studying only the corresponding directed version of the problem. For instance, it is known that every 6-connected graph is 2-linked, that is, it has disjoint (s_1, t_1)-, (s_2, t_2)-paths for every choice of distinct vertices s_1, s_2, t_1, t_2 [14, 15], whereas there is no degree of strong connectivity that guaranties that a digraph is 2-linked [17]. Thus it is natural to ask whether there is a natural number k so that the underlying digraph of every k-strong digraph D satisfies that for every choice of distinct vertices s_1, s_2, t_1, t_2 D contains disjoint (s_1, t_1)-, (s_2, t_2)-paths P_1, P_2 with the additional property that P_1 is a directed (s_1, t_1)-path in D, whereas P_2 is not necessarily a directed path in D. Note that it was shown in [2] that deciding the existence of such a pair P_1, P_2 in a given input digraph is NP-complete.

As it is the case for Problem 1, the corresponding undirected version in which we seek two edge-disjoint cycles, and also the directed version where we seek arc-disjoint dicycles are both polynomially solvable. For the former a good characterization exists (see e.g. [7]). For the latter a polynomial algorithm follows from the fact that D has a pair of arc-disjoint dicycles if and only if its line-digraph $L(D)$ has a pair of vertex disjoint dicycles and thus the algorithm by MCCUAIG solves the problem. Note that this transformation does not preserve equivalent instances of Problem 2, in fact it does not even do so for undirected cycles as the line graph of a tree may contain arbitrarily many (edge-)disjoint cycles.

We will prove the following characterization.

Theorem 4. *Problem 2 is polynomially solvable for the class of strong digraphs and for every class of nonstrong 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 NP-complete.*

The proof is partitioned into four parts. First we cover the case of strong digraphs, then we show that for nonstrong digraphs, if we allow arbitrarily many transversal arcs, Problem 2 is NP-complete. Then we show that if there is at least one transversal vertex and we can bound the number of transversal arcs by some constant k, there is a polynomial algorithm (whose running time depends on k). Finally, we prove that if there is no transversal vertex, then the problem is again polynomially solvable.

2. THE CASE OF STRONG DIGRAPHS

Let D be a strongly connected instance of Problem 2. We may assume that D has no vertex with in-degree and out-degree one, since such a vertex can be suppressed without changing the problem (the new digraph is a yes-instance if and only if D is a yes-instance. Hence below we always assume that $\delta(D) \geq 3$.

Lemma 5. *Suppose D is a strongly connected no-instance of Problem 2 with minimum degree at least three. If C is a dicycle in D, then every vertex $v \in V(D) - V(C)$ has $d(v) = 3$. In particular, if $\tau(D) \geq 2$, then $UG(D)$ is a cubic graph.*

Proof. Suppose v is a vertex of $V(D) - V(C)$ with $d_D(v) \geq 4$. Because D is a strongly connected no-instance $D - A(C)$ must be a forest with all leaves in $V(C)$. So every arc going out of v is the initial arc of a $(v, V(C))$-path in $D - A(C)$, every arc entering v is the terminal arc of a unique $(V(C), v)$-path in $D - A(C)$ and, no matter how we choose these $d(v)$ paths, they will be pairwise internally disjoint. Call these paths between v and the cycle $P_1, \ldots, P_{d(v)}$, such that their endpoints $p_1, \ldots, p_{d(v)}$ appear in this order on C. Since D is strongly connected, at least one of these paths must be a $(v, V(C))$-path and at least one must be a $(V(C), v)$-path. Rotate the labeling of the paths until P_1 is a $(v, V(C))$-path and P_2 is a $(V(C), v)$-path. Since no p_i is an internal vertex of $C[p_1, p_2]$ the dicycle $vP_1C[p_1, p_2]P_2v$ is arc-disjoint from the cycle $vP_3C[p_3, p_4]P_4v$, contradicting that D is a no-instance for Problem 2. Thus we must have $d(v) = 3$ for all $v \in V(D) - V(C)$. The last assertion of the lemma follows from the fact that if $\tau(D) \geq 2$, then $D - v$ has a dicycle for every v and thus the argument above implies that $d(v) = 3$. □

Lemma 6. *Let D be a strong digraph with $\tau(D) = 1$ and k transversal vertices. In polynomial time we can either decide that D is a yes-instance of Problem 2 and produce the desired cycles or reduce D to a collection $\mathcal{H}$ of at most k smaller strong digraphs so that D is a yes-instance of Problem 2 if and only if at least one digraph in $\mathcal{H}$ is a yes-instance of Problem 1.*

Proof. By Lemma 5, in polynomial time, we can either determine that D is a yes-instance or conclude that every vertex v that is not a dicycle transversal has $d(v) = 3$ (as $D - v$ contains a dicycle). Let $k > 0$ be the number of dicycle transversals. By definition of a transversal vertex, every dicycle contains all transversal vertices. Furthermore, these transversal vertices occur in the same order on all dicycles. This is clear if $k \leq 2$ and if $k \geq 3$ and there are three transversal vertices x, y, z occurring in the order x, y, z, x on one dicycle and x, z, y, x on another dicycle, then x, y are in the same strong component of $D - z$, contradicting that z is on every dicycle. Let x be an arbitrary transversal vertex and let $x, x_1, \ldots, x_{k-1}$ be the order in which the transversal vertices occur on C (and every other dicycle). Split x into two vertices x_0 and x_k by letting all arcs entering (leaving) x enter x_k (leave x_0). The resulting digraph D^a is acyclic and we may assume that D^a does not contain a pair of arc-disjoint (x_0, x_k)-paths since such a pair easily yields a solution back in D. Thus we have that D is a yes-instance of Problem 2 if and only if D^a contains an (x_0, x_k)-path that is arc-disjoint from some cycle in $UG(D^a)$. As each x_i is a dicycle transversal of D and D^a is acyclic, there is no (x_i, x_j)-path with $j \neq i$ in $D^a - x_{i+1}$ for any $i \in \{0, \ldots, k-1\}$.

For $i \in \{1, \ldots, k\}$ let B_i denote the set of vertices z such that there is an (x_{i-1}, z)-path, but no (x_i, z)-path of length at least one in D^a. For simplicity we also use B_i to denote the subdigraph $D^a\langle B_i\rangle$ of D^a induced by B_i. A vertex z is an internal vertex of B_i if $z \in B_i - \{x_{i-1}, x_i\}$. By the remark above, every vertex of $V - \{x_0, x_1, \ldots, x_{k-1}\}$ is an internal vertex of a unique B_i. Furthermore, there is no arc from $B_i - \{x_i\}$ to B_j when $j \neq i$. Since D is strongly connected every vertex of B_i can reach x_i via a directed path inside B_i. By MENGER's Theorem, every B_i that is not just the set $\{x_{i-1}, x_i\}$ contains two internally disjoint (x_{i-1}, x_i)-paths. For each of those we fix two such paths $P_{i,1}, P_{i,2}$. We may assume that $P_{i,1}, P_{i,2}$ and are induced paths in D^a since otherwise D is clearly a yes-instance. Similarly, we may assume that there is no B_j containing three arc-disjoint (x_{j-1}, x_j)-paths.

If every B_i, $i \in \{1, \ldots, k\}$ does contain two internally disjoint (x_{i-1}, x_i)-paths we also have a yes, since then D^a contains two arc-disjoint (x_0, x_k)-paths, implying that D contains two arc-disjoint dicycles.

Note that if B_i contains any internal vertices (namely B_i is not just the arc $x_{i-1}x_i$), then it contains two internally disjoint (x_{i-1}, x_i)-paths, because no internal vertex of B_i is a transversal vertex. By Lemma 5, all internal vertices must have degree 3 (these are not transversal vertices). Furthermore they must all belong to either $P_{i,1}$ or $P_{i,2}$: Suppose $z \in B_i - V(P_{i,1}) \cup V(P_{i,2})$. As $d(z) = 3$ we see that $UG(B_i)$ contains two paths P, Q starting in z and both ending in vertices on $P_{i,j}$ while avoiding all vertices of $P_{i,3-j} - \{x_{i-1}, x_i\}$ for $j = 1$ or $j = 2$. Now any undirected cycle contained in the union of $P, Q, P_{i,j}$ is arc-disjoint from $P_{i,3-j}$ and hence D is a yes-instance.

By the information collected so far we have that $P_{i,1}, P_{i,2}$ are induced paths, every internal vertex z belongs to some $P_{i,r}$, has $d(z) = 3$ and its unique neighbour (out- or in-neighbour) outside $P_{i,r}$ is on $P_{i,3-r}$, $r \in [2]$. Furthermore, D^a does not contain a pair of arc-disjoint (x_0, x_k)-paths and each x_i is a cut-vertex of $UG(D^a)$ for $i \in \{1, \ldots, k-1\}$. Now it follows that D is a yes-instance if only if there is some $i \in \{1, \ldots, k\}$ such that the strong digraph D_i which we obtain from B_i by adding the arc x_ix_{i-1} is a yes-instance of Problem 1. Thus we can decide whether D is a yes-instance of Problem 2 by making at most $k \in O(n)$ calls to the polynomial algorithm for strongly connected instances of Problem 1. □

Theorem 7. *Problem 2 is polynomially solvable for strong digraphs. Furthermore, if D is a yes-instance, then a solution can be produced in polynomial time.*

Proof. Let D be a strongly connected instance of Problem 2. If $\tau(D) = 1$ we can use the polynomial algorithm that is implicit from the proof of Lemma 6 (possibly followed by at most $k \leq n$ calls of the polynomial algorithm for strong instances of Problem 1), so we may assume that $\tau(D) \geq 2$. Now it follows from the proof of Lemma 5 that, in polynomial time, we can either produce a solution or conclude that $UG(D)$ is a cubic graph (if $UG(D)$ is not cubic, then we take a cycle C, a vertex $v \in V - C$, with $d(v) \geq 4$ and build an undirected cycle disjoint from C with some of the four paths between v and C). In the case of $UG(D)$ being a cubic graph, cycles that are edge-disjoint are also vertex-disjoint, implying that D is a yes-instance of Problem 2 if and only if it is a yes-instance of Problem 1[1]. Thus we can apply the polynomial algorithm for Problem 1 to decide whether or not D is a yes-instance and produce a solution in polynomial time when one exists. □

3. THE CASE OF NONSTRONG DIGRAPHS

In this section, we treat the case when D is a nonstrong digraph. Clearly, if D is acyclic, it is a no-instance. Hence, given that we can check in linear time whether D is acyclic, we assume below D is not acyclic.

Since a digraph with at least two nontrivial strong components of size greater than one has two disjoint dicycles and the dicyle transversal number of any digraph D is the sum of the dicycle transversal numbers of its strong components, we may assume that D has precisely one nontrivial strong component D' that is a no-instance of Problem 2 and since every digraph D with $\tau(D) \geq 3$ is a yes-instance of Problem 1 [4] and hence also of Problem 2 we can assume that $\tau(D') \in \{1, 2\}$.

If $UG(D - V(D'))$ contains a cycle, D is clearly a yes-instance of Problem 1 and hence of Problem 2, so we may assume that $UG(D - V(D'))$ is a forest. Furthermore, since no dicycle can use a vertex outside of D', D is a yes-instance if and only if it has a solution B, C where B is a dicycle in D' and the intersection of C with any connected component of $UG(D - V(D'))$ is a path. Hence, by contracting each connected component of $UG(D - V(D'))$ to a single vertex and possibly reorienting some arcs, we may assume that $V(D) - V(D')$ is an independent set of vertices $y_1, y_2, \dots, y_p$ and that every arc with exactly one end in $V(D')$ enters $V(D')$.

3.1. Digraphs with vertex transversal number one

We start by showing that if we do not restrict the class of digraphs further than we just did above, then Problem 2 is NP-complete. As it was the case in our proof in [4] of the NP-completeness of Problem 1, in order to prove the $\mathcal{NP}$-completeness of Problem 2 we first consider a problem for undirected graphs and then show how to reduce a restricted (but still NP-complete) variant of this to Problem 2 in polynomial time.

[1]Note that it was shown in [3] that if $\tau(D) > 2$, then D is a yes-instance of Problem 1 and hence also of Problem 2.

Problem 3. *Given an undirected graph G and a collection $\mathcal{A}$ of disjoint subsets each consisting of two vertices of G, decide if there exists a cycle C in G such that $|X - V(C)| \geq 1$ for all $X \in \mathcal{A}$.*

We say that a cycle C in G is **legal** with respect to $\mathcal{A}$ if $X \not\subseteq V(C)$ holds for all $X \in \mathcal{A}$.

Theorem 8. *Problem 3 is NP-complete.*

Proof. Given an instance (G, k) of the independent set problem (i.e. the problem of deciding whether G contains an independent set of size k), we construct an equivalent instance $(G', \mathcal{A})$ of Problem 3 as follows:
Fix an ordering $v_1, ..., v_n$[2] of $V(G)$ and initialize $\mathcal{A} = \emptyset$. For every $i \in \{1, \dots, n\}$ we create a vertex x_i of G' and an (x_i, x_{i+1})-path P_i on $d_G(v_i) + 3$ vertices.

Let us say that we **block** a pair of paths P, Q, when we add to $\mathcal{A}$ a set $\{p, q\}$ such that p (q) is an internal vertex of P (Q) that does not belong to any of the sets in $\mathcal{A}$ (constructed up to that moment). Likewise we **block** a pair P, u of a path P and a vertex u, when we add to $\mathcal{A}$ a set $\{p, u\}$ such that p is an internal vertex of P so that neither p nor u belongs to any of the sets currently in $\mathcal{A}$. For every edge $v_i v_j \in E(G)$, we block the pair P_i, P_j. For every i we create a vertex u_i and the edge $x_i u_i$ and block the pair P_i, u_i. Note that the paths P_i, $i \in \{1, \dots, n\}$ are long enough to permit all these disjoint blocks.

Furthermore we create a set W of $n - k$ new vertices $w_1, ..., w_{n-k}$ and for every $1 \leq i \leq n, 1 \leq j \leq n - k$ we add a (u_i, w_j)-path U_{ij} and a (w_j, x_i)-path X_{ij}, both of length $2n(n - k)$. Now we block all possible pairs among the set $\{U_{ij}, X_{ij} | 1 \leq i \leq n, 1 \leq j \leq n - k\}$ except the pairs of the kind U_{ij}, X_{i+1j}. Again the paths are long enough to permit all these disjoint blocks. This concludes the description of the instance $(G', \mathcal{A})$ (see Figure 1). By construction it has size polynomial in $|V(G)|$ and we can construct it in polynomial time.

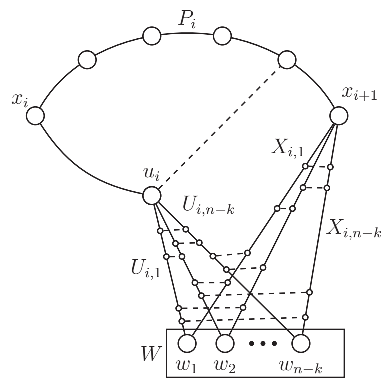

FIGURE 1. The choice gadget G' related to vertex v_i of G and its paths to W. Blocker pairs indicated by dashed edges. Note that blockers concerning other parts of the graph have been omitted.

We claim that every legal cycle in G' traverses $x_1, ..., x_n$ in this order and then returns to x_1. Indeed if $w_j \in W$ is a vertex of a legal cycle, this cycle contains U_{ij} and X_{i+1j}, for

[2]Throughout the proof, the indices i are meant cyclically "modulo" n, so that $n + 1 = 1$.

some i. So a legal cycle containing x_i must continue towards x_{i+1} either using P_i or the path $x_i u_i U_{ij} X_{i+1j}$, for some j (but it cannot use both of them, since P_i, $\{u_i\}$ is blocked).

Now assume $(G', \mathcal{A})$ has legal cycle C. By construction of G' and the fact that C is legal, the set $I := \{v_i | V(P_i) \subset V(C)\}$ is independent in G. Moreover $|I| \geq k$, since C can intersect W in at most $n - k$ vertices, implying that $V(P_i) \subset C$ for at least k different values of i.

Vice-versa if G has an independent set I of size k, we form a legal cycle by joining, for every $i \in \{1, \ldots, n\}$, x_i with x_{i+1} via the path P_i, if $v_i \in I$ and via the path $x_i U_{ij} w_j X_{i+1j}$, otherwise. Here $j = j(i)$ is determined by any injective mapping from $\{i \in \{1, \ldots, n\} | v_i \notin I\}$ to $\{1, \ldots, n - k\}$. Since $|V(G) - I| = n - k$ such an injective mapping exists (it is in fact bijective) and the cycle we form is legal. □

Theorem 9. *Problem 3 restricted to instances $(G, \mathcal{A})$ where $G = (U, W, E)$ is a bipartite graph with partite sets U, W and $W = \bigcup_{X \in \mathcal{A}} X$ is still NP-complete*

Proof. We prove this by turning a general instance $(G, \mathcal{A})$ of Problem 3 into an equivalent restricted bipartite instance H satisfying the above condition on the structure of $\mathcal{A}$.

Let $W = \bigcup_{X \in \mathcal{A}} X$ and notice that if there exists a cycle C such that $V(C) \cap W = \emptyset$, then G is a yes-instance and we can replace G with an arbitrary yes-instance H satisfying the condition in the statement of the theorem[3]. So we can assume that $G - W$ induces a forest and thus contracting every edge between the vertices in $G - W$ will only shorten a possible solution cycle. To get rid of the edges between vertices in W first observe that an edge e whose end vertices form a set in $\mathcal{A}$ can never be in any legal cycle, so we can safely delete it. We subdivide the remaining edges in W. Now our graph H is bipartite with the two partitions being $W = \bigcup_{X \in \mathcal{A}} X$ and $U = V(H) - W$. And the restricted instance H has a legal cycle if and only if the general instance had one. Clearly this transformation can be performed in polynomial time. □

Theorem 10. *Problem 2 is NP-complete.*

Proof. We prove this by reducing from an instance $(G = (U, W, E), \mathcal{A})$ of Problem 3 restricted to bipartite graphs where $W = \bigcup_{X \in \mathcal{A}} X$.

We get an equivalent instance of Problem 2 by constructing a digraph D as follows. Let $\mathcal{A} = \{\{p_1, q_1\}, \{p_2, q_2\}, \ldots, \{p_k, q_k\}\}$ and replace every p_i (q_i) in G by a path $P_i = u_{i,1} u_{i,2} \ldots u_{i,d_G(p_i)}$ $(Q_i = v_{i,1} v_{i,2} \ldots v_{i,d_G(q_i)})$ such that the path has one vertex for each neighbor of p_i (q_i) and the jth such vertex is adjacent to the jth neighbor of p_i (q_i) in G. This makes sure all internal vertices of P_i (Q_i) have underlying degree 3. Orient all edges from U to W, this makes sure that no dicycle can enter U, and orient all paths P_i and Q_i as directed paths from $u_{i,1}$ to $u_{i,d_G(p_i)}$ respectively, from $v_{i,1}$ to $v_{i,d_G(q_i)}$. Add vertices $s_1, \ldots, s_k$ and $t_1, \ldots, t_k$ to V and add the following arcs for each $i \in \{1, \ldots, k\}$: the arcs $s_i u_{i,1}$, $s_i v_{i,1}$ from s_i to the first vertices of P_i and Q_i, the arcs $u_{i,d_G(p_i)} t_i$, $v_{i,d_G(q_i)} t_i$ from the last vertices of P_i and Q_i to t_i, the arcs $t_i s_{i+1}$, $i \in [k-1]$ and finally the arc $t_k s_1$. Call the resulting digraph D (see Figure 2).

If there exists a legal cycle C in G we can construct the corresponding undirected cycle C_1 in D by keeping the edges of C and adding the relevant internal arcs of the paths P_i

[3]For instance let H be $K_{2,4}$ and let $\mathcal{A}$ consist of two disjoint sets that partition the larger side of the bipartition.

and Q_i. Since C was legal we know that $V(C_1) \cap V(P_i) = \emptyset$ or $V(C_1) \cap V(Q_i) = \emptyset$ for all $i \in \{1, \ldots, k\}$, so we can still create a dicycle C_2 that is disjoint from C_1 by using, for $i \in \{1, \ldots, k\}$, the path P_i if $V(C_1) \cap V(P_i) = \emptyset$ and the path Q_i otherwise and adding all arcs $t_i s_{i+1}$, $i \in [k-1]$ and the arc $t_k s_1$.

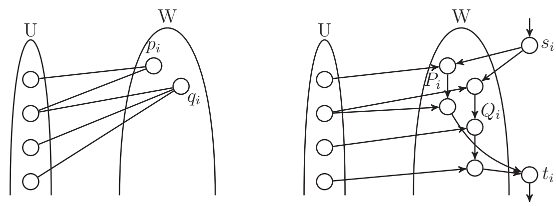

FIGURE 2. A piece of the construction of D. A set $\{p_i, q_i\} \in \mathcal{A}$ on the left gets blown up into paths P_i and Q_i on the right, so that every vertex has underlying degree 3. Then we orient each of P_i, OQi as directed paths from s_i to t_i.

Conversely, if there exists an undirected cycle C_1 arc-disjoint from a dicycle C_2 in D we know that $V(C_2) \cap U = \emptyset$ and since every vertex in $V(D) - U$ has underlying degree 3 we observe that C_1 and C_2 are vertex-disjoint. From the construction of D, t_i is the only out-neighbor of the set $(V(P_i) \cup V(Q_i)) - t_i$, which itself does not contain any dicycle. It follows that the only possible type of dicycle in D must be one that traverses $s_1, t_1, s_2, t_2, \ldots, s_k, t_k$ in that order and contains either P_i or Q_i for all $i \in \{1, \ldots, k\}$. This implies that for every $i \in [k]$ we have $s_i, t_i \notin C_1$ and that $V(P_i) \not\subset V(C_1)$ or $V(Q_i) \not\subset V(C_1)$. So the cycle C in G that uses the edges corresponding to arcs of C_1 going from U to W is legal since $V(P_i) \not\subset V(C_1)$ or $V(Q_i) \not\subset V(C_1)$ implies $p_i \notin V(C)$ or $q_i \notin V(C)$.□

3.2 Nonstrong digraphs with $\tau(D) = 1$ and a bounded number of transversal arcs

In this section, we will consider Problem 2 on nonstrong digraphs D with $\tau(D) = 1$ and a bounded number of transversal arcs. As Theorem 7 provides a polynomial algorithm that decides the problem for strong digraphs we assume here that there is exactly one nontrivial strong component D' and that D' itself is a no-instance of Problem 2.

Suppose that D has k transversal arcs (they are all arcs of D', in fact) and let them be called $u_1v_1, \ldots, u_kv_k$ according to the order in which they appear on every dicycle of D. Since these are the only transversal arcs, we know, by Menger's theorem, that between every pair v_i, u_{i+1} with $v_i \neq u_{i+1}$ there must exist at least two arc-disjoint (v_i, u_{i+1})-paths in D'. On the other hand if there are more than two we can combine two of these into an undirected cycle while still being able to use the remaining path to construct an arc-disjoint directed cycle, thus contradicting that D' was a no-instance. So if, for each i with $v_i \neq u_{i+1}$, we consider a maximum collection of arc-disjoint (v_i, u_{i+1})-paths there are exactly two paths in each of these collections. There may be many different such pairs so for each $i \in [k]$ with $v_i \neq u_{i+1}$ we fix two such paths P_i and Q_i[4].

[4]Note that P_i and Q_i may share internal vertices and these are all transversal vertices of D.

A (p, q)-path W, where $p \in P_i$ and $q \in Q_i$ for some $i \in \{1, \ldots, k\}$, is called a **(p, q)-switch** if it is internally disjoint from $V(P_i) \cup V(Q_i)$. Notice that all switches are paths in D' and that a switch may consist of just one vertex, if P_i and Q_i share an internal vertex a, this vertex forms an (a, a)-switch. If for any $p \in V(P_i)$ there are two different switches to Q_i, ending in q and q' respectively, then these paths can be combined with part of Q_i into an undirected cycle arc-disjoint from P_i, contradicting the assumption that D' was a no-instance. Similarly, a (p, q)-switch and a (p', q)-switch with p, p' on P_i and q on Q_i would lead to a contradiction. The same observation applies to switches that start on Q_i and terminate in P_i. So for every $p \in V(P_i)$ $(q \in V(Q_i))$ there exists at most one choice of switch between p and Q_i (P_i and q) with either its tail or head being p (q) or the switch is the vertex $p \in V(P_i) \cap V(Q_i)$ $(q \in V(P_i) \cap V(Q_i))$.

We say a directed cycle C **switches** paths between v_i and u_{i+1} if C contains an (x, y)-switch with $x, y \in V(P_i) \cup V(Q_i) - \{v_i, u_{i+1}\}$.

Lemma 11. *Suppose D is a yes-instance of Problem 2 and let C, C' be a solution such that C switches a minimum number of times in total between all pairs v_i and u_{i+1} for $i \in \{1, \ldots, k\}$. Then for any $h \in \{1, \ldots, k\}$ C switches at most once between v_h and u_{h+1}.*

Proof. Assume, by contradiction, that C switches at least twice between v_h and u_{h+1}. Now redefine C' to be an undirected cycle in $D \setminus A(C)$ that minimizes $|A(C') \setminus (A(P_h) \cup A(Q_h))|$. □

Claim 1. *For any choice of distinct vertices $a, b \in V(C) \cap V(P_h)$ such that $A(P_h[a, b]) \cap A(C) = \emptyset$, the set $A(P_h[a, b]) \cap A(C')$ is connected and non-empty.*

Proof of Claim. Suppose that $A(P_h[a, b]) \cap A(C') = \emptyset$. Then replacing $C[a, b]$ with $P_h[a, b]$ in C yields a solution with fewer switches contradicting minimality of C. At the same time if $A(P_h[a, b]) \cap A(C')$ has two connected components R_1, R_2, (see Figure 3 left side) then any (R_1, R_2)-path in C' together with $P_h[a, b]$ contains an undirected cycle C'' arc-disjoint from C with fewer arcs in $A(C'') \setminus (A(P_h) \cup A(Q_h))$ than C', contradiction. This proves the claim. □

Note that the claim also holds for Q_h instead of P_h.

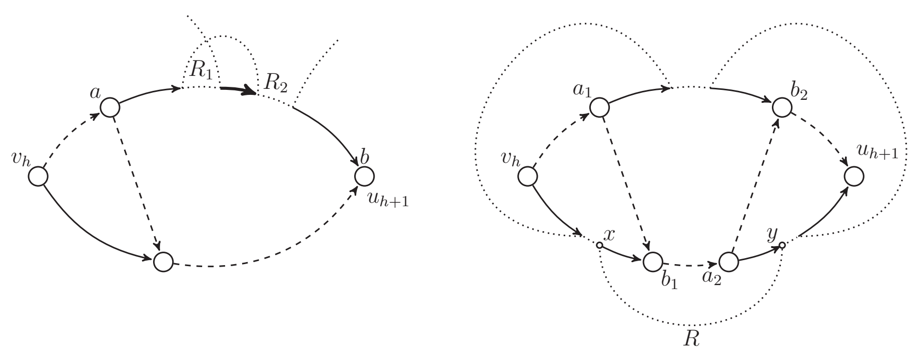

FIGURE 3. The dashed lines indicate the directed cycle C while dotted lines indicate the undirected cycle C'. On the left $A(P_h[a, b]) \cap A(C')$ has two components which the fat path connects, allowing for a "shorter" undirected cycle. On the right the two switches can be avoided by using $RQ[x, y]$ as the undirected cycle.

Without loss of generality we can assume that the first two switches of C after v_h (and by assumption before u_{h+1}) are an (a_1, b_1)-switch and an (a_2, b_2)-switch, where $a_1, b_2 \in V(P_h)$ and $b_1, a_2 \in V(Q_h)$. By the claim $P_h[a_1, b_2]$, $Q_h[v_h, b_1]$ and $Q_h[a_2, u_{h+1}]$ all have a nonempty intersection with C' and $A(P_h[a_1, b_2]) \cap A(C')$ must be connected (see Fig. 3 right side). Furthermore C' must contain two disjoint $(V(Q_h[v_h, b_1]), V(Q_h[a_2, u_{h+1}]))$-paths, one of which does not intersect $P_h[a_1, b_2]$. Let us call this path R and its endpoints x and y. Then R together with $Q_h[x, y]$ contains an undirected cycle arc-disjoint from the directed cycle obtained from C by replacing $C[a_1, b_2]$ with $P_h[a_1, b_2]$ and this dicycle has fewer switches than C, contradiction. □

Theorem 12. *Problem 2 can be solved in polynomial time on digraphs with $\tau(D) = 1$ and a bounded number of transversal arcs.*

Proof. Consider a solution C, C' as in Lemma 11. Recall that C must contain the arc u_iv_i plus a (v_i, u_{i+1})-path for $i = 1, ..., k$ in this order. There are at most $2 + \min(|V(P_i)|, |V(Q_i)|)$ such (v_i, u_{i+1})-paths that can be part of C, namely P_i, Q_i or a path that contains one switch between these two and the switches are completely determined by either of their endpoints. It follows that there are at most $O(n^k)$ minimum switching directed cycles.

We consider the following algorithm for Problem 2: for every choice of a minimum switching directed cycle C, remove the arcs of the cycle from the digraph and look for an undirected cycle in the remaining digraph. If a solution is found output YES, if no solution is found after all possible choices of minimum switching directed cycles have been checked, output NO. The correctness of the algorithm follows from Lemma 11 and the running time is $O(n^{k+2})$. □

3.3 Nonstrong digraphs with $\tau(D) = 2$

In this section we complete the proof of Theorem 4 by showing that Problem 2 is polynomially solvable for digraphs with transversal number 2 (and hence with no transversal arcs). We will use the polynomial algorithm for Problem 1 for digraphs with $\tau(D) = 2$ as a subroutine. To do so, we need to describe a class of digraphs that are called vaults in [3]. Let $\ell \geq 5$ be odd, let $P_0, \dots, P_{\ell-1}$ be disjoint nonempty 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_ic_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

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

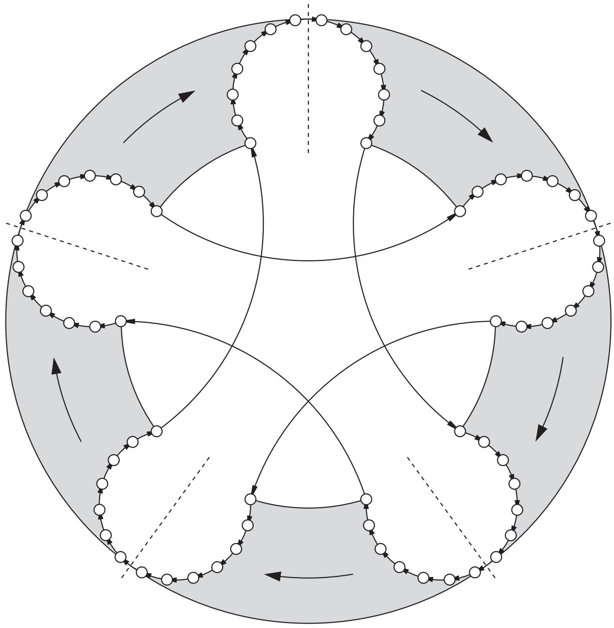

FIGURE 4. A typical vault. The five central arcs must have multiplicity 1 and are the only arcs from P_i to P_{i+2}.

and the arc *rs*. Figure 4 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 gray 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 that can be drawn without crossing in such an area. Vaults are strong 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.

Now, we are ready to state the result from [3] which we shall use.

Theorem 13. *[3] Let $D = (V, A)$ be a strong digraph with $\tau(D) = 2$. 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 has a at most $O(|V|^6)$ distinct dicycles.*
(ii) *D is a subdivision of a vault without a niche.*

Furthermore, if D satisfies one of (i)-(ii), we can produce a certificate for this in polynomial time including all the distinct dicycles in D if case (i) occurs. □

We also need the following lemma that is easy to establish.

Lemma 14. *[4] Let D be a vault with $\ell \geq 5$ walls. Then the following holds.*

(i) *D contains a dicycle that avoids all vertices of any given wall P_i.*
(ii) *For every choice of vertices x, x' such that x and x' belong to different and non-consecutive walls P_i and P_{i+2r}, where $r > 0$, D contains a dicycle that is disjoint from the (x, x')-path $P_i[x, d_i]P_{i+2} \ldots P_{i+2r-2}P_{i+2r}[a_{i+2r}, x']$.*

Concluding the proof of Theorem 4. We are now ready to prove the characterization given in Theorem 4: By Theorem 7, Problem 2 is polynomial for strong digraphs, by Theorem 10 it is NP-complete when the number of transversal arcs is arbitrary and by Theorem 12 it is polynomial when the number of transversal arcs is positive and bounded. When there is no transversal arc, if $\tau(D) \neq 2$, then either D has no dicycle or $\tau(D) \geq 3$ and D is a yes-instance[5], so we are left with the case of nonstrong digraphs with $\tau(D) = 2$.

Since we are only interested in a polynomial algorithm and are not concerned about obtaining the most efficient algorithm, it suffices to show how to handle the case (ii) of Theorem 13. By Lemma 5 we may assume that D' is a subdivision of a cubic vault S with an odd number $\ell > 3$ of walls without niches. Recall that, by the argument in the beginning of this section, $D - V(D')$ consists of one or more isolated vertices. If any such vertex v has three or more neighbors in $V(D')$, then either v has two neighbors on the same wall or two neighbors u, w on nonconsecutive walls. In both cases Lemma 14 would imply that D is a yes-instance, so, by the remark in the beginning of the section, we may assume that every vertex of $V(D) - V(D')$ has out-degree 2 and in-degree zero. By Lemma 14(i), each such vertex z is attached to distinct out-neighbors x, x' in $V(D')$. Fix such a triple z, x, x'. If x, x' belong to the same wall or to nonconsecutive walls then, by Lemma 14, we conclude that D is a yes-instance. So we may assume that x, x' are on consecutive walls P_i, P_{i+1}, respectively. These walls contain the subpaths $P_i[a_i, b_i]$, $P_i[c_i, d_i]$ and $P_{i+1}[a_{i+1}, b_{i+1}]$, $P_{i+1}[c_{i+1}, d_{i+1}]$, respectively. Since the vault has no niche, the connections between the two walls corresponding to the (possibly subdivided) arcs connecting them in the vault are as follows: If $r_{i,1}, \ldots, r_{i,k}$ denote the principal vertices (i.e. the nonsubdivision vertices) of $P_i[c_i, d_i]$ in this order, and $l_{i+1,1}, \ldots, l_{i+1,k'}$ denote the principal vertices of $P_{i+1}[a_{i+1}, b_{i+1}]$ in this order, then, as we have no niches and S is cubic, $k = k'$ and there are (possibly subdivided) arcs from $r_{i,j}$ to $l_{i+1,j}$. Suppose first that x is equal to one of the vertices $r_{i,j}$. Then, by Lemma 14, there is a directed cycle C using the entire wall P_i and avoiding P_{i+1} and the subdigraph induced by $\{x, z\} \cup V(P_{i+1}) \cup V(Q)$, where Q is the path in D' corresponding to the (possibly subdivided) arc $r_{i,j}l_{i+1,j}$ in S, contains an undirected cycle that is arc-disjoint from C, so D is a yes-instance.

Now suppose that x is a principal vertex of $P_i[a_i, b_i]$. Then we find a directed cycle C starting at x by following $P_i[x, d_i]$, then skipping every second wall (and fully using the others) by using the (possibly subdivided) arcs of the kind $d_p a_{p+2}$ of the vault, finally using the part of the wall P_{i-1} right before P_i in the cyclic order until the principal vertex w of P_{i-1} for which there is an arc wx of S connecting it to x and then traversing the subdivided path in D' corresponding to that arc to close up to x. The walls not used by C plus their respective (possibly subdivided) arcs of the form $d_j a_{j+2}$ plus z plus $P_i[a_i, x]$ constitute an undirected cycle arc-disjoint from C. So D is a yes-instance. It follows that x is not a principal vertex of P_i and, by symmetry, that x' is not a principal vertex of P_{i+1}. Hence x, x' are subdivision vertices. Moreover, if there was another vertex z'' outside D' distinct from z with outneighbors x, x'', then again x, x'' must be on consecutive walls, so either x'' is in P_{i+1} or x'' is in P_{i-1}. But then $x'zxz''x''$ is an undirected path edge-disjoint from D' connecting distinct vertices of the same wall P_{i+1} or nonconsecutive walls

[5]Recall that this is due to the fact [4] that if $\tau(D) \geq 3$, then D is a yes-instance of Problem 1.

P_{i-1}, P_{i+1}, and, by Lemma 14, D is a yes-instance. Hence the out-neighborhoods of the vertices outside D' are pairwise disjoint sets of subdivision vertices in D'. Consequently, the underlying undirected graph has no vertices of degree larger than 3 and D is a yes-instance of Problem 2 if and only if it is a yes-instance of Problem 1. Hence, we can apply the polynomial algorithm for Problem 1 for digraphs with $\tau(D) = 2$. □

4. FINAL REMARKS

We saw in this article that Problem 2 is NP-complete if we allow arbitrarily many transversal arcs in the input digraphs. One could ask whether this is due to the fact that the cycles may avoid using all vertices (in particular, we have no control on the lengths of the cycles in a solution).

On the other hand, it was shown in [5] that is it an NP-complete problem to decide whether the underlying graph of a given input digraph D contains a spanning tree T so that $D - A(T)$ still has some vertex that can reach all other vertices by directed paths. So there are also problems where we are looking for spanning structures that are NP-complete in the mixed version.

It could thus be interesting to study the complexity of the following problem.

Problem 4. *Given a digraph D, does $UG(D)$ contain arc-disjoint 2-regular spanning subgraphs $\mathcal{F}_1, \mathcal{F}_2$ (i.e. each $\mathcal{F}_i$ is a spanning collection of vertex disjoint cycles in $UG(D)$) with the property that $\mathcal{F}_1$ is a cycle factor of D (while $\mathcal{F}_2$ may contain not only dicycles)?.*

Note that by Petersen's theorem [13], the underlying graph $UG(D)$ of a digraph D has two arc-disjoint undirected cycle factors if and only if $UG(D)$ has a 4-regular spanning subgraph and this condition can be checked polynomially using matching techniques. Similarly, a digraph D has two arc-disjoint dicycle factors if and only if it contains a 2-regular spanning subdigraph and we can use flow techniques to check this in polynomial time. Therefore, also in this case both the homogeneous versions of Problem 4 are polynomially solvable.

REFERENCES

[1] J. Bang-Jensen and G. Gutin, Digraphs. Theory, Algorithms and Applications, 2nd edn., Springer Monographs in Mathematics, Springer-Verlag London, Ltd., London, 2009.

[2] J. Bang-Jensen and M. Kriesell, Disjoint directed and undirected paths and cycles in digraphs, Theoret Comput Sci 410 (2009), 5138–5144.

[3] J. Bang-Jensen and 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, and S. Simonsen, Disjoint directed and undirected cycles in a general digraph, J Combin Theory Ser B 106 (2014), 1–14.

[5] J. Bang-Jensen and A. Yeo, Arc-disjoint spanning sub(di)graphs in digraphs, Theoret Comput Sci 438 (2012), 48–54.

[6] H. L. Bodlaender, On disjoint cycles, Int J Found Comput Sci 5 (1994), 59–68.

[7] B. Bollobás. “Extremal Graph Theory”, Academic Press Inc., London, 1978, pp. xx+488.
[8] J. A. Bondy and U. S. R. Murty, Graph theory, Springer Graduate Texts in Mathematics 244 (2008), pp. xii+651.
[9] G. A. Dirac, Some results concerning the structure of graphs, Canad Math Bull 6 (1963), 183–210.
[10] S. Fortune, J. Hopcroft, and J. Wyllie, The directed subgraph homeomorphism problem, Theoret Comput Sci 10 (1980), 111–121.
[11] L. Lovász, On graphs not containing independent circuits (in Hungarian), Mat Lapok 16 (1965), 289–299.
[12] W. McCuaig, Intercyclic digraphs, Graph structure theory (Seattle, WA, 1991), Contemp. Math. 147, Amer. Math. Soc., Providence, RI, 1993, pp. 203–245.
[13] J. Petersen, Die Theorie der regulären Graphs, Acta Math 15 (1891), 193–220.
[14] P. Seymour, Disjoint paths in graphs, Discrete Math 29 (1980), 293–309.
[15] C. Thomassen, 2-linked graphs, Eur J Combin 1 (1980), 371–378.
[16] C. Thomassen, Disjoint cycles in digraphs, Combinatorica 3 (1983), 393–396.
[17] C. Thomassen, Highly connected non-2-linked digraphs, Combinatorica 11 (1991), 393–395.
[18] C. Thomassen, The 2-linkage problem for acyclic digraphs, Discrete Math 55 (1985), 73–87.