

INTRINSIC LINKING IN DIRECTED GRAPHS

JOEL STEPHEN FOISY, HUGH NELSON HOWARDS and NATALIE ROSE RICH*

(Received July 19, 2013, revised April 1, 2014)

Abstract

We extend the notion of intrinsic linking to directed graphs. We give methods of constructing intrinsically linked directed graphs, as well as complicated directed graphs that are not intrinsically linked. We prove that the double directed version of a graph G is intrinsically linked if and only if G is intrinsically linked. One Corollary is that \overline{J}_6 , the complete symmetric directed graph on 6 vertices (with 30 directed edges), is intrinsically linked. We further extend this to show that it is possible to find a subgraph of \overline{J}_6 by deleting 6 edges that is still intrinsically linked, but that no subgraph of \overline{J}_6 obtained by deleting 7 edges is intrinsically linked. We also show that \overline{J}_6 with an arbitrary edge deleted is intrinsically linked, but if the wrong two edges are chosen, \overline{J}_6 with two edges deleted can be embedded linklessly.

1. Introduction

Research in spatial graphs has been rapidly on the rise over the last fifteen years. It is interesting because of its elegance, depth, and accessibility. Since Conway and Gordon published their seminal paper in 1983 [4] showing that the complete graph on 6 vertices is intrinsically linked (also shown independently by Sachs [10]) and the complete graph on 7 vertices is intrinsically knotted, over 62 papers have been published referencing it. The topology of graphs is, of course, interesting because it touches on chemistry, networks, computers, etc. This paper takes the natural step of asking about the topology of directed graphs, sometimes called digraphs. Applications of graph theory are increasingly focused on directed graphs, from epidemiology to traffic patterns (see, for example [6], [1]). Many of these applications depend, not just on the abstract directed graph, but the topology and geometry of an actual realization of the graph. Examples abound, such as very large scale integration (known as VLSI) for circuits. As a result, understanding the topology of directed graphs is a worthy pursuit.

Recall that an undirected graph is *intrinsically linked* if every embedding of the graph in \mathbb{R}^3 contains at least two non-splittably linked cycles. A graph H is a *minor* of a graph G if H can be obtained from G by a sequence of vertex deletions, edge deletions and edge contractions. The set of all minor-minimal intrinsically linked graphs is given by the seven Petersen family graphs [4], [9], [10]. These graphs are

1991 Mathematics Subject Classification. 57M25, 57M15, 05C10.

*indicates student at time of research.

obtained from K_6 by $\Delta - Y$ and $Y - \Delta$ exchanges. Similarly, we define a directed graph \overline{G} to be *intrinsically linked* if it contains a nontrivial link consisting of a pair of consistently oriented directed cycles in every spatial embedding.

We prove that the double directed version of a graph G is intrinsically linked if and only if G is intrinsically linked. One Corollary is that $\overline{J_6}$, the complete symmetric directed graph on 6 vertices (with 30 directed edges), is intrinsically linked. We extend this to show that it is possible to find a subgraph of $\overline{J_6}$ by deleting 6 edges that is still intrinsically linked, but that no subgraph of $\overline{J_6}$ obtained by deleting 7 edges is intrinsically linked and that $\overline{J_6}$ with an arbitrary edge deleted is intrinsically linked, but if the wrong two edges are chosen, $\overline{J_6}$ with two edges deleted can be embedded linklessly. We further show that, unlike for undirected graphs, the edge contraction operation does not necessarily preserve the property of having a linkless embedding. Finally, we show that, again unlike for directed graphs, the $\Delta - Y$ operation does not necessarily preserve intrinsic linking.

One can imagine numerous interesting applications of intrinsic linking in directed graphs. One such application is to computer chips. Computer chips can be thought of as containing embedded directed graphs where signals are sent along wires that serve as edges. Diodes make the graph directed allowing the electricity to flow in only one direction on a given wire and gates serve as vertices. Intrinsic linking in the directed graph means that any way of building the chip will require the wires to cross at least twice. This is important because building chips where the wires cross adds to the expense and complexity of the chip. Understanding exactly how many links are intrinsic in the directed graph could help determine the difficulty and cost of building a specific chip.

Another potential application is to the idea of maximum upward planar embeddable subgraphs of directed graphs (see for example [2]). People are interested in efficiently finding the largest planar subgraph of a directed graph that can be embedded in \mathbb{R}^3 so that all the edges of the embedding point monotonically upward. To check if a given graph or subgraph has an upward planar embedding, it does not suffice to check if it has a cycle, however a directed cycle is an obstruction to such an embedding as is intrinsic linking in the underlying graph. Thus intrinsic linking in the directed graph would be a natural obstruction to study since it involves both the existence of directed cycles within the directed graph and intrinsic linking in the underlying graph.

The authors would like to thank Keir Lockridge and Danielle O'Donnol for helpful comments.

2. Preliminaries

We start by associating directed graphs and non-directed graphs with each other. Recall that in a directed graph each edge has one vertex that is its initial point or starting point and the other is said to be its terminal point or ending point. Instead of viewing edges abstractly as pairs of vertices, we view them as ordered pairs. For

graphs that are not directed, a *simple graph* is one which contains no loops (edges that run from a vertex to itself) or multiple edges (two edges associated to the same pair of vertices). A *simple directed graph* is one that contains no loops (edges that have the same initial and terminal vertex) or multiple edges (two edges associated to the same ordered pair of vertices). Note in this case there may be two edges associated to v_i and v_j , $i \neq j$, one with v_i as its initial point and v_j as its terminal point and one in the opposite direction. The *complete graph on n vertices*, K_n has exactly one edge for each distinct pair of vertices. The *complete symmetric directed graph on n vertices*, \overline{J}_n has exactly one edge for each ordered pair of distinct vertices. Note \overline{J}_n has twice as many edges as K_n .

To help clarify when we have put a direction on the edges of an undirected graph, say G , we will refer to the directed graph as \overline{G} . To differentiate when we are specifying a graph and when an embedding of a graph we will refer to the abstract graph as G , and the embedded version as Γ (or $\overline{\Gamma}$ for the embedded directed graph etc.).

Given a graph G , we define *the double directed version of G* , denoted $\overline{D(G)}$ by taking the vertices of G and for each edge of G associated to vertices v_i and v_j we include two edges, one with v_i as its initial point and v_j as its terminal point and one in the opposite direction. The embedded version of $\overline{D(G)}$ we will call $\overline{\Delta(\Gamma)}$.

We will follow the same convention on edges and cycles as we do with graphs. If e_i is an edge of G , and we put directions on the edges of G , then e_i will correspond to $\overline{e}_i \subset \overline{G}$. If we embed G or \overline{G} the edge e_i will then correspond to $\alpha_i \subset \Gamma$ or $\overline{\alpha}_i \subset \overline{\Gamma}$ respectively. If we double G to get directed graph $\overline{D(G)}$, then we assign the two directed edges \overline{e}_i and \overline{e}'_i to e_i (and, of course $\overline{\alpha}_i$ and $\overline{\alpha}'_i$ in $\overline{\Delta(\Gamma)}$). Finally for cycles in \overline{G} we will reserve the notation \overline{C} for a consistently oriented cycles where as C may be used to designate any cycle even if it is not directed or if it is directed, but perhaps not consistently oriented. If $C \subset G$, and we embed G to get Γ , we will refer the cycle that is the image of C as $\gamma \subset \Gamma$ and again use the oriented cycle notation as before, so $\overline{\gamma}$ would represent a consistently oriented, embedded cycle.

Given a directed graph \overline{G} , we can, of course, associate a non-directed graph G by keeping the same set of vertices and for each edge of \overline{G} associated to the ordered pair (v_i, v_j) we include an edge associated to the vertices v_i and v_j (essentially just ignoring the direction). We call G the *underlying graph* of \overline{G} . Note that \overline{G} , and G have the same number of edges here. We define the *underlying simple graph of \overline{G}* to be the largest simple subgraph of the underlying graph of \overline{G} . As examples, $\overline{J}_6 = \overline{D(K_6)}$ and K_6 is the underlying simple graph of \overline{J}_6 . The underlying graph of \overline{J}_6 , has two parallel edges for each edge of K_6 .

3. Results

Lemma 3.1. *For all n , given any complete (non-directed) graph K_n we may assign directions to the edges to yield a directed graph $\overline{K_n}$ and in which $\overline{K_n}$ contains no cycles.*

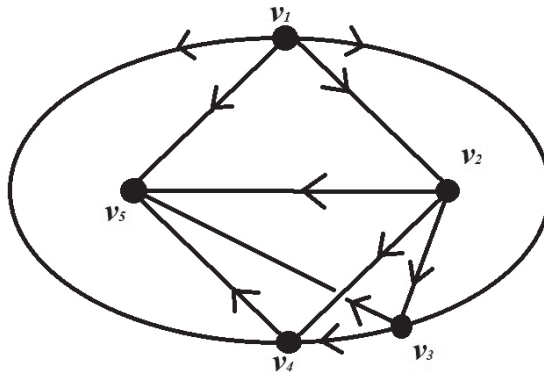


Fig. 1. Here K_5 is constructed with directed edges and with no cycles. The same construction will work for K_n in general.

It is important to note that $\overline{K_n} \neq \overline{J_6}$. It has half as many edges, only one edge for any distinct pair of vertices, instead of two. There are $2^{\binom{n}{2}}$ ways to orient the edges of $\overline{K_n}$, so up to graph isomorphism there are many distinct versions of $\overline{K_n}$, while there is a single K_6 and a single $\overline{J_6}$ up to isomorphism.

Proof of Lemma 3.1. To form a $\overline{K_n}$ with no cycles, we start by choosing directions for all edges containing vertex v_1 and direct all of them away from v_1 . Now we know no cycle can contain v_1 since every cycle must contain one edge entering a vertex and one leaving it. For v_2 direct all currently unlabelled edges containing v_2 away from v_2 . The only edge entering v_2 comes from v_1 , call it $\overline{e_1}$. The edge $\overline{e_1}$ cannot be part of a cycle since no cycle contains v_1 . No edge other than $\overline{e_1}$ enters v_2 , so v_2 cannot be part of a cycle. Now for v_3 , direct all currently unlabelled edges containing v_3 away from v_3 . The only edges entering v_3 come from v_1 and v_2 , call these $\overline{e_2}$ and $\overline{e_3}$. Now as before these two edges cannot be part of a cycle and no other edge enters v_3 , so v_3 cannot be part of a cycle. Continue in this manner until all edges are labelled (any edge from v_i to v_j with $i < j$ will have v_i as its initial point and v_j as its terminal point). We now see that the graph contains no cycles. We include a picture of $\overline{K_5}$ embedded in this manner, with no cycles, in Fig. 1. \square

Theorem 3.2. For all n , given any complete (non-directed) graph K_n we may assign directions to the edges to yield a directed graph $\overline{K_n}$ with $\overline{K_n}$ not intrinsically linked.

Proof. Label as in Lemma 3.1. Obviously a graph with no cycles cannot contain a pair of linked cycles. \square

We may now generalize to all (non-directed) graphs.

Corollary 3.3. *Every graph G is the underlying graph of a directed graph \overline{G} that is not intrinsically linked.*

Proof. Since every simple graph is a subgraph of a complete graph the result is immediate in that case. If G is not simple, we direct all edges of a maximal simple graph and then use the same direction for all parallel edges. Before considering loops, there can be no cycles. Embed the loops so they bound a disk in the complement of the graph (this may all be done in a small neighborhood of the edge's associated vertex). Now the directed graph will have cycles consisting of the loops, but the cycles cannot form a non-split link since they all bound disks that are disjoint on their interior. \square

Theorem 3.5 makes use of the following lemma introduced previously in [3] and [5]. Recall that a *theta curve* (or theta graph) is a graph with two vertices of valence three and three edges running between the two vertices. Note that a theta curve will always contain three cycles resulting from the three ways to pick two of its edges (the cycles will not, of course, be disjoint from each other).

Lemma 3.4. *Given disjoint embeddings of a cycle γ and a theta curve θ in S^3 , with θ containing three cycles γ_1 , γ_2 , and γ_3 , if γ has nonzero linking number with γ_i for some i , then γ also has nonzero linking number with γ_j for some $j \neq i$.*

Proof. First we show that the theta graph lies on an immersed sphere, with the cycles of the theta graph dividing it into three distinct disks. This can be proven in essentially the same way that one can prove that every knot bounds a disk (although the disk will, of course, not be embedded if the knot is not the unknot). In that situation one may take a homotopy of the knot through space to a circle, take the flat disk it bounds and then reverse the homotopy, extending it to the disk and knot as a pair. Here instead we take a homotopy of the theta graph onto an embedded sphere so that the graph is embedded on the sphere with the three cycles dividing the sphere into three disks which are disjoint on their interiors. Now reverse the homotopy extending it to the pair of the sphere and the graph. The sphere, of course, now may not be embedded since the cycles may be knotted, but this is not a problem.

Now γ must algebraically intersect the sphere zero times, showing that the sum of the linking numbers of γ with γ_1 , γ_2 and γ_3 with sign must add to zero. This shows that if one of the terms in the sum is nonzero, at least two of the terms must be nonzero. \square

Theorem 3.5. *If G is intrinsically linked then $\overline{D(G)}$ is intrinsically linked.*

Proof. For any embedding of $\overline{D(G)}$ (called $\overline{\Delta(\Gamma)}$ once embedded) the underlying graph must contain a link since the underlying graph for $\overline{D(G)}$ contains (numerous)

copies of G . Let the edges of one copy of G be labelled $\{\overline{\alpha}_1, \overline{\alpha}_2, \dots, \overline{\alpha}_n\}$. Let the complementary edges pointing in the opposite direction be labelled $\{\overline{\alpha}'_1, \overline{\alpha}'_2, \dots, \overline{\alpha}'_n\}$, forming another copy of G , sharing only the vertices, where we always choose the labels so that if $\overline{\alpha}_s$ runs from v_i to v_j in $\overline{\Delta}(\Gamma)$, then $\overline{\alpha}'_s$ runs from v_j to v_i in $\overline{\Delta}(\Gamma)$. The link from the first copy of G in the embedding of the underlying graph consists of a (non-directed) linked pair of cycles. If those cycles have consistent directions when viewed as part of $\overline{\Delta}(\Gamma)$ we are done. If not, without loss of generality let γ be one of the cycles and γ_1 be the other. Choose the labels so that γ_1 does not have all of its edges consistently directed (γ may fit this description, too, but may not).

Without loss of generality, let $\overline{\alpha}_1 \cup \overline{\alpha}_2 \cup \dots \cup \overline{\alpha}_j$ be all the edges of γ_1 . We are trying to find a consistently oriented directed cycle linked with γ . We arbitrarily choose a favorite direction for γ_1 (although for efficiency one would probably choose the direction that at least half the edges of γ_1 already pointed). We leave those edges already going in our selected direction fixed and focus on the edges going the other direction. We will swap out the edges, trying to get a linked cycle in the direction we have chosen. Let $\overline{\alpha}_i$ be an inconsistent edge relative to our desired direction. Now $\gamma_1 \cup \overline{\alpha}'_i$ is a theta curve and by Lemma 3.4 either $\overline{\alpha}_i \cup \overline{\alpha}'_i$ or $\overline{\alpha}_1 \cup \overline{\alpha}_2 \cup \dots \cup \overline{\alpha}_{i-1} \cup \overline{\alpha}'_i \cup \overline{\alpha}_{i+1} \cup \dots \cup \overline{\alpha}_j$ is linked with γ . We know $\overline{\alpha}_i \cup \overline{\alpha}'_i$ is consistently oriented. If it is linked with γ let $\overline{\gamma}'' = \overline{\alpha}_i \cup \overline{\alpha}'_i$. If not then we know $\overline{\alpha}_1 \cup \overline{\alpha}_2 \cup \dots \cup \overline{\alpha}_{i-1} \cup \overline{\alpha}'_i \cup \overline{\alpha}_{i+1} \cup \dots \cup \overline{\alpha}_j$, which we will call γ_2 , is the linked cycle. It follows that γ_2 has one more consistently oriented edge than γ_1 did. If all the edges are now consistent let $\gamma_2 = \overline{\gamma}''$, if not, pick one of the now smaller collection of edges pointing the wrong way, say $\overline{\alpha}_k$ and take the theta curve where we add $\overline{\alpha}'_k$ to $\overline{\alpha}_1 \cup \overline{\alpha}_2 \cup \dots \cup \overline{\alpha}_{i-1} \cup \overline{\alpha}'_i \cup \overline{\alpha}_{i+1} \cup \dots \cup \overline{\alpha}_j$. The process must terminate, since each time we either have a linked, consistently oriented, two cycle or we produce a cycle γ_{i+1} from γ_i , where γ_{i+1} is the same length as γ_i , has one more consistently oriented edge than γ_i did, and is still linked with γ . Repeat until we have a consistently oriented cycle that is linked with γ and call that new cycle $\overline{\gamma}''$.

Now if γ is consistently oriented, we are done. If not, we repeat now letting $\overline{\gamma}''$ play the role of the fixed cycle and arbitrarily picking a favorite direction for γ . We then follow the same type of process as before, taking an edge $\overline{\alpha}_s$ of γ that is inconsistently oriented and add in $\overline{\alpha}'_s$ to γ to get a theta curve. The same algorithm now applied to γ shows we have a consistently oriented cycle that is linked with $\overline{\gamma}''$. \square

Theorem 3.6. *If G is not intrinsically linked (knotted) then $\overline{D}(G)$ is not intrinsically linked (knotted).*

Proof. Let Γ be a linkless (knotless) embedding of G . Take edges parallel to the edges of Γ and orient one in each direction so that the edge from v_i to v_j and the edge from v_j to v_i bound a disk disjoint from the rest of the graph on its interior. The result is a linkless (knotless) embedding of $\overline{D}(G)$. \square

We now sharpen our result.

Theorem 3.7. *Let \overline{G} be \overline{J}_6 minus a consistently oriented cycle of length at least 3, then \overline{G} is intrinsically linked. The same is true if \overline{G} is \overline{J}_6 minus two disjoint consistently oriented triangles or a consistently oriented 4 cycle and a disjoint edge.*

Note that when we use the term cycle here, the first and last vertex are the same and otherwise no vertex is repeated more than once.

Proof. If we are deleting a single n cycle from \overline{J}_6 , then without loss of generality let the deleted cycle consist of edges labelled $\overline{e}_1, \overline{e}_2, \dots, \overline{e}_n$. Take a subgraph of \overline{G} , including edges $\overline{e}_1, \overline{e}_2, \dots, \overline{e}_n$ whose underlying graph is a K_6 (We can do this since we never deleted both of the edges between a given pair of vertices). Once embedded, this graph must have a (non-directed) linked pair of cycles. If none of $\overline{\alpha}_1, \overline{\alpha}_2, \dots, \overline{\alpha}_n$ are in the cycles then proceed as in the proof of Theorem 3.5. If some of them are, use the included edges to dictate the preferred direction of the cycles. Note that the preferred direction of a given cycle will be well defined. This is true because if one of the triangles includes only one of the edges, then we use the direction of that single edge. If one included two of them, they are adjacent edges in both the triangle and the n -cycle and thus must be consistently oriented in both, so all of the edges of $\overline{\alpha}_1, \overline{\alpha}_2, \dots, \overline{\alpha}_n$ that might be included in a given triangle will be consistently oriented in the triangle. We only consider swapping out inconsistent edges, so any edge $\overline{\alpha}_i$ that might be swapped out is not an element of $\{\overline{\alpha}_1, \overline{\alpha}_2, \dots, \overline{\alpha}_n\}$ and thus we know that both $\overline{\alpha}_i$ and $\overline{\alpha}'_i$ are in our remaining graph even after the n -cycle was deleted. The only way that one of the triangles could include 3 edges would be if $n = 3$ and the triangle was exactly $\{\overline{\alpha}_1, \overline{\alpha}_2, \overline{\alpha}_3\}$, but this triangle is already consistently oriented. If $n > 3$, we cannot use 3 of the edges in the same triangle since no subset of $\overline{\alpha}_1, \overline{\alpha}_2, \dots, \overline{\alpha}_n$ forms a triangle.

The argument in the cases of two disjoint consistently oriented triangles or a consistently oriented 4 cycle and a disjoint edge is similar. □

Corollary 3.8. *The graphs $\overline{J}_6 - \overline{e}_i$, $\overline{J}_6 - \overline{e}_i - \overline{e}_j$, and $\overline{J}_6 - \overline{e}_i - \overline{e}_j - \overline{e}_k$ are intrinsically linked if \overline{e}_i , \overline{e}_j , and \overline{e}_k do not share a vertex.*

Proof. Each of these have a subgraph that is \overline{J}_6 minus a consistently oriented Hamiltonian cycle. A graph with an intrinsically linked subgraph must be intrinsically linked. □

Theorem 3.9. *The graph $\overline{J}_6 - \overline{e}_i - \overline{e}_j$ is not intrinsically linked if \overline{e}_i and \overline{e}_j share exactly one vertex and are not consistently oriented with each other, but is intrinsically linked if they are consistently oriented. Further, $\overline{J}_6 - \overline{e}_i - \overline{e}_j$ is not intrinsically linked if \overline{e}_i and \overline{e}_j share both vertices (but with opposite orientations).*

Proof. Fig. 2 shows a linkless embeddings of $\overline{J}_6 - \overline{e}_i - \overline{e}_j$ where \overline{e}_i and \overline{e}_j share a vertex and are not consistently oriented. Fig. 3 shows a linkless embedding of $\overline{J}_6 - \overline{e}_i - \overline{e}_j$ where \overline{e}_i and \overline{e}_j share both vertices, but with opposite orientations ($\overline{e}_i = \overline{e}'_j$).

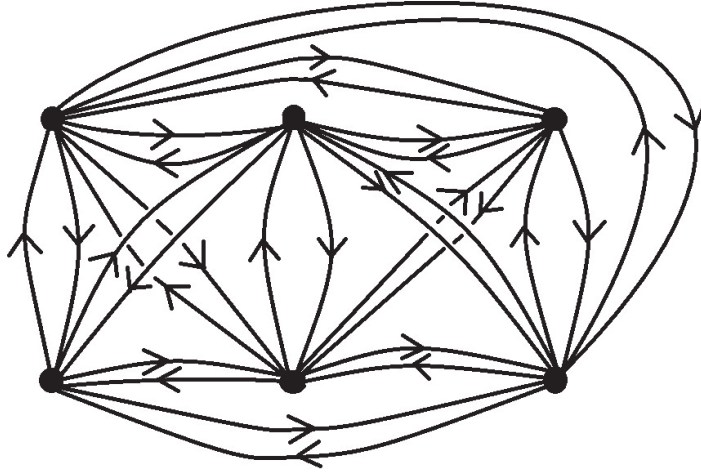


Fig. 2. We see a linkless embedding of $\overline{J_6}$ minus two edges between the same pair of vertices.

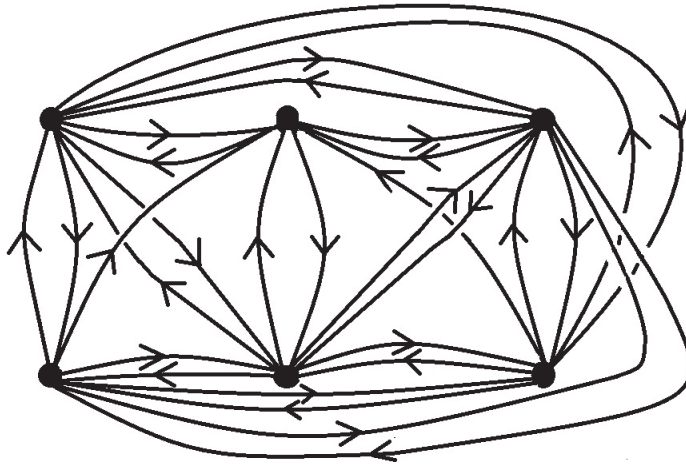


Fig. 3. We see a linkless embedding of $\overline{J_6}$ minus two inconsistently oriented edges that share exactly one vertex.

Corollary 3.8 implies that in all other cases $\overline{J_6 - e_i - e_j}$ is intrinsically linked. \square

Corollary 3.10. *The graph $\overline{J_6}$ with seven edges deleted is never intrinsically linked.*

Proof. Seven edges will have seven terminal points. Since we have only six vertices one vertex must be the terminal point of at least two of the edges removed. This means the two edges share a vertex and are not consistently oriented with each other. \square

4. Examining traditional moves that preserve intrinsic linking in non-directed graphs in the setting of directed graphs

We recall that for non-directed graphs the $\Delta - Y$ move preserves intrinsic linking [7]. It is natural to ask if the same is true for directed graphs. It is clear that if you replace a triangle with a Y and an added vertex it will not preserve intrinsic linking if you arbitrarily assign directions to the edges in the Y . This would fail, for example, if one picked a graph embedded with a unique link in it made up of two directed triangles. Now replace one of those triangles with a Y where all the edges of the Y have the added vertex as a terminal point. There is no link in the complement of the Y , but since no edge has its initial point at the added vertex there are no cycles and thus no links that run through the Y . One might wonder if one could choose the directions for the Y more wisely and preserve intrinsic linking. We answer that in the negative in the following theorem

Theorem 4.1. *The $\Delta - Y$ move, where we replace a directed triangle with by a directed Y , does not preserve intrinsic linking no matter how the three edges are oriented.*

Proof. We prove this by generating a specific counterexample, where the initial graph is intrinsically linked, but after a single $\Delta - Y$ move, no matter how the Y is oriented, the resulting graph is not. We start with an embedding of a graph $\overline{\Gamma}$ shown in Fig. 4. Note that this is a $\overline{J_6}$ with two disjoint, consistently oriented triangles removed so by Theorem 3.7 the graph is intrinsically linked. We adopt the convention in the figure that edges with no orientation drawn on them represent two edges embedded next to each other, but pointing in opposite directions. The unique link in this embedding is the cycle corresponding to vertices v_1, v_3 and v_5 linked with the cycle for vertices v_2, v_6 , and v_4 . We apply the $\Delta - Y$ move to the cycle v_2, v_6, v_4 to get the graph in Fig. 5. Note that no matter how we orient the edges (v_2, v_7) , (v_4, v_7) and (v_6, v_7) at least two of the edges will have to have either v_7 as their initial point or their terminal point and thus there will be at least one pair of edges in the Y that fail to be consistently oriented with each other. In Fig. 5 we have chosen to have (v_2, v_7) and (v_6, v_7) inconsistently oriented, but we can do this without loss of generality in the case where the Y contains exactly one pair of edges that are not consistently oriented

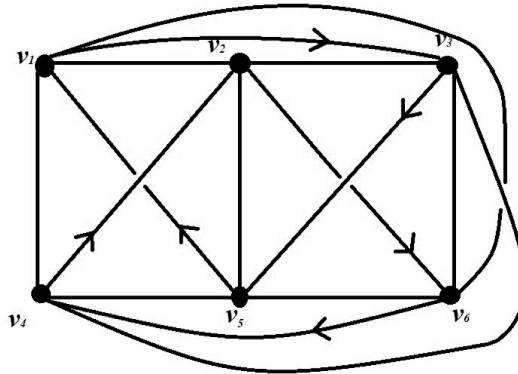


Fig. 4. An intrinsically linked directed graph $\bar{\Gamma}$ embedded so that it contains a unique link. Edges with no orientation drawn on them represent two edges embedded next to each other, but pointing in opposite directions.

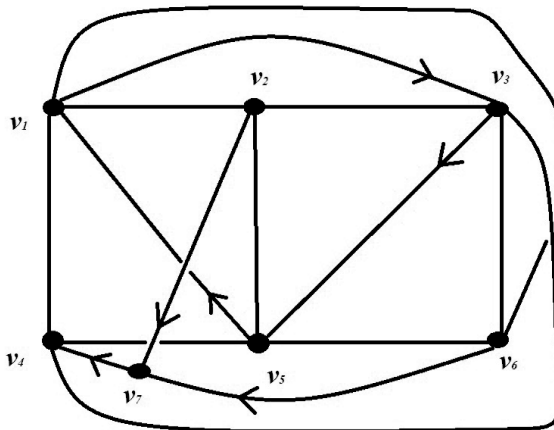


Fig. 5. A directed graph formed from Γ via a $\Delta - Y$ move in a linkless embedding. Edges with no orientation drawn on them represent two edges embedded next to each other, but pointing in opposite directions.

due to symmetries of the graph. We could, of course, reverse the directions of all of the edges in the Y , but the argument in that case is completely analogous. We also could orient it so that two or even all three pairs of edges are not consistently oriented instead of just one pair, but that only makes the argument easier, so we may focus on the single case shown in the figure.

Now note that any link in the new graph must use edge (v_2, v_7) because otherwise the only crossings left are from the edges between v_1 and v_6 and between v_3 and v_4 . Since (v_2, v_7) is consistently oriented with (v_7, v_4) , but not (v_6, v_7) , we know that any link must also involve (v_7, v_4) . A quick inspection shows there are no links left in the graph. (Once we delete the unused edge (v_6, v_7) , we can think of the graph as isotopic to $\bar{\Gamma}$ with edges (v_2, v_6) and (v_4, v_6) deleted and a single vertex of degree two inserted in the middle of the edge (v_2, v_4) .) As $\bar{\Gamma}$ had only one link and that link contained both edges (v_2, v_6) and (v_4, v_6) , clearly the new graph contains no links. \square

Recall that a *vertex expansion* of a vertex v in a graph G is achieved by replacing v with two vertices, v' and v'' , adding the edge (v', v'') , and connecting a subset of the edges that were incident to v to v' , and connecting the remaining edges that were incident to v to v'' . The reverse of this operation is *edge contraction*. We say that v is *expanded to a double edge* if both the edges (v', v'') and (v'', v') are added. The reverse operation is a *double edge contraction*.

We have the following result.

Theorem 4.2. *Vertex expansion does not preserve the property of being intrinsically linked. Equivalently, edge contraction does not preserve the property of having a linkless embedding.*

Proof. See Fig. 6. The directed graph on the left of the figure is intrinsically linked, by Theorem 3.7. The directed graph embedding shown on the right is linkless. \square

Note that if the added edge in Fig. 6 had the opposite direction, this embedding would have a link. We have shown that vertex expansion fails to preserve intrinsic linking if the direction of the edge is chosen arbitrarily. It is unclear if vertex expansion preserves intrinsic linking if we are allowed to pick the direction of the edge wisely.

Theorem 4.3. *Let \bar{G} have a linkless embedding, and let \bar{G}' be obtained from \bar{G} by a double edge contraction. Then \bar{G}' has a linkless embedding.*

Proof. Assume the hypotheses. Let $\bar{\Gamma}$ be a linkless embedding of \bar{G} . Let $\bar{\alpha}$ and $\bar{\alpha}'$ denote the double edge that is contracted to obtain embedding $\bar{\Gamma}'$ of \bar{G}' . By contracting $\bar{\alpha}$ down to a vertex, v , the result is an embedding of \bar{G}' , $\bar{\Gamma}'$, plus an extra loop from edge $\bar{\alpha}'$ based at vertex v that is not linked with any oriented cycle in $\bar{\Gamma}'$ (else

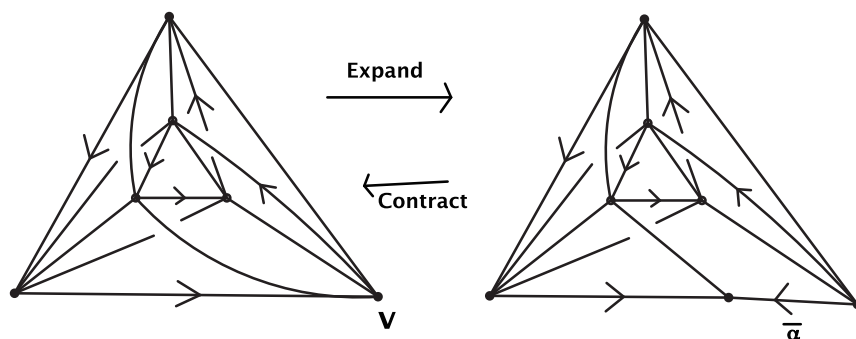


Fig. 6. The vertex v is expanded to the edge $\bar{\alpha}$. Edges with no orientation drawn on them represent two edges embedded next to each other, but pointing in opposite directions.

the cycle $(\bar{\alpha}, \bar{\alpha}')$ would have been linked in the embedding $\bar{\Gamma}$). Since $\bar{\Gamma}$ was linkless, the only possible linked cycle pairs in $\bar{\Gamma}'$ must contain a cycle that passes through the vertex v .

Suppose such a pair of oriented, linked cycles, $\bar{\gamma}$ and $\bar{\gamma}'$ exist. Without loss of generality, let $\bar{\gamma}$ be the cycle that contains v . Then $\bar{\gamma}$ pulls back to a (perhaps not consistently oriented) cycle containing $\bar{\alpha}$ in $\bar{\Gamma}$. (In theory, $\bar{\gamma}$ could pull back to a cycle that passes through just one vertex of $\bar{\alpha}$, but does not include $\bar{\alpha}$. In this case, the oriented non-split link pulls back to an oriented non-split link in $\bar{\Gamma}$, which is a contradiction.) By abuse of notation, we denote this cycle in $\bar{\Gamma}$ as γ . We know that $\bar{\gamma}' \subset \bar{\Gamma}'$ corresponds to an oriented cycle $\bar{\gamma}' \subset \bar{\Gamma}$ that is linked with γ . Since $\bar{\Gamma}$ is linkless, γ is indeed not consistently oriented. In particular the orientation at $\bar{\alpha}$ does not agree with the orientation of the other edges of γ . By Lemma 3.4, however, either the cycle of length two $(\bar{\alpha}, \bar{\alpha}')$ or the cycle formed by replacing $\bar{\alpha}$ in γ with $\bar{\alpha}'$ is linked with $\bar{\gamma}'$, and both of these linked cycles are consistently oriented in $\bar{\Gamma}$. This is a contradiction, thus $\bar{\Gamma}'$ is linkless showing \bar{G}' indeed has a linkless embedding. \square

We may state this equivalently as:

Theorem 4.4. *If \bar{G} is intrinsically linked and \bar{G}' is obtained from \bar{G} by a vertex expanded to a double edge, then \bar{G}' is also intrinsically linked.*

Corollary 4.5. *Suppose \bar{G} is intrinsically linked. If we take $\bar{\Gamma}$, an embedding of \bar{G} , and expand a vertex to a single new edge e , we may pick a direction on e such that the expanded, directed graph $\bar{\Gamma}^+$ contains a directed link.*

Proof. By Theorem 4.4, we know that expanding \bar{G} to get \bar{G}' with double edge expansion is intrinsically linked. Therefore if we pick any embedding of \bar{G}' with new

edges $\bar{\alpha}$ and $\bar{\alpha}'$ parallel to each other such that $\bar{\alpha} \cup \bar{\alpha}'$ bounds a disk in the complement of the graph, but with opposite directions to get $\bar{\Gamma}'$, we know $\bar{\Gamma}'$ contains a link. Since we picked $\bar{\alpha}$ and $\bar{\alpha}'$ so that $\bar{\alpha} \cup \bar{\alpha}'$ bounds a disk in the complement of the graph, they cannot form one of the components of the non-split link. Thus this link uses at most one of the two edges. Without loss of generality, let the edge used be $\bar{\alpha}$. Then let $\bar{\Gamma}^+$ be $\bar{\Gamma}'$ with $\bar{\alpha}'$ deleted. This graph is constructed to be isotopic to the graph we were interested in, and it contains a link. \square

Note that the direction we picked in this proof was dependent on the embedding, so this is slightly weaker than proving that vertex expansion with careful direction selection preserves intrinsic linking. Perhaps for a different embedding the preferred direction to get a link would be the opposite.

5. Open questions

We close with a few open questions. Since intrinsic linking in directed graphs is a new direction for this field, the options for open questions are rich and plentiful. Just as with the study on non-directed graphs, there should be extensions of this work that are accessible to undergraduate research and yet there also are deep difficult questions to answer.

QUESTION 5.1. Does vertex expansion (to a single directed edge) preserve intrinsic linking if we may carefully choose the direction of the expanded edge? In other words, if \bar{G} is intrinsically linked and \bar{G}' is obtained from \bar{G} by a vertex expanded to a single directed edge \bar{e}' , and \bar{G}'' is identical to \bar{G}' , except we replace \bar{e}' with \bar{e}'' , where \bar{e}'' has the same endpoints as \bar{e}' , but with opposite orientation, must \bar{G}' or \bar{G}'' be intrinsically linked?

If we take an embedding of \bar{J}_6 based on an embedding of K_6 which contains a unique non-split link and then for each α_i in the embedding of K_6 take $\bar{\alpha}_i$ and $\bar{\alpha}'_i$ in our embedding of \bar{J}_6 parallel to each other, we get an embedding of \bar{J}_6 with exactly four non-split links in it. Each of the linked triangles becomes two triangles with opposite orientations. The two pairs of triangles are linked, so there are four ways to pick from the two pairs to get a link. These links, of course, are not disjoint from each other, since they share vertices and even entire components, but they are distinct. There are currently no known embeddings of \bar{J}_6 with fewer links than this. This leads to more open questions.

QUESTION 5.2. How many distinct links must $D(G)$ have if G is intrinsically linked?

Parallel to work on non-directed graphs we ask the following important questions.

QUESTION 5.3. What is the complete list of minor minimal intrinsically linked directed graphs?

Note that by [8], such a list must be finite. Unfortunately, such a list will not completely characterize intrinsically linked directed graphs, as edge contraction does not necessarily preserve having a linkless embedding.

QUESTION 5.4. Is there a set of moves, like the $Y - \Delta$ and $\Delta - Y$ moves for non-directed graphs, that will generate the complete list of minor minimal intrinsically linked directed graphs?

Variations of many of the results that have been obtained for non-directed graphs could be attempted for directed graphs.

QUESTION 5.5. For what n does \overline{J}_n always contain a non-split 3 component link?

Note that in the case of non-directed graphs, K_9 was the smallest possible graph because each component needed to be at least a triangle, but in \overline{J}_n , link components may be bigons, so this bound does not hold. Even \overline{J}_6 can be embedded with a 3 component split link, but it is, of course, easy to find embeddings of \overline{J}_6 that do not contain a non-split 3-component link.

It is also natural to ask how intrinsic knotting fits into the realm of directed graphs

QUESTION 5.6. Is \overline{J}_7 intrinsically knotted?

QUESTION 5.7. What is the complete minor minimal set of intrinsically knotted directed graphs?

QUESTION 5.8. If G is intrinsically knotted, is $\overline{D(G)}$ intrinsically knotted?

Again many of the results that have been obtained for intrinsic knotting in non-directed graphs could be attempted for directed graphs.

Additionally, one might consider intrinsic knotting and linking in graphs where some edges are directed and some are not. In that case a cycle could use both directed edges and non-directed edges, in which case any directed edges in a given cycle would need to be consistently oriented, while non oriented edges could be used in any cycle.

References

- [1] J. Bang-Jensen and G. Gutin: Digraphs, second edition, Springer Monographs in Mathematics, Springer, London, 2009.

- [2] C. Binucci, W. Didimo and F. Giordano: *Maximum upward planar subgraphs of embedded planar digraphs*, *Comput. Geom.* **41** (2008), 230–246.
- [3] G. Bowlin and J. Foisy: *Some new intrinsically 3-linked graphs*, *J. Knot Theory Ramifications* **13** (2004), 1021–1027.
- [4] J.H. Conway and C.McA. Gordon: *Knots and links in spatial graphs*, *J. Graph Theory* **7** (1983), 445–453.
- [5] E. Flapan, R. Naimi and J. Pommersheim: *Intrinsically triple linked complete graphs*, *Topology Appl.* **115** (2001), 239–246.
- [6] P.P. Howards, E.F. Schisterman, C. Poole, J.S. Kaufman and C.R. Weinberg: *“Toward a clearer definition of confounding” revisited with directed acyclic graphs*, *American Journal of Epidemiology* **176** (2012), 506–511.
- [7] R. Motwani, A. Raghunathan and H. Saran: *Constructive results from graph minors: linkless embeddings*; in 29th Annual Symposium on Foundations of Computer Science (White Plains, NY, 1988), IEEE, 1988, 398–409.
- [8] N. Robertson and P.D. Seymour: *Graph minors. XX. Wagner’s conjecture*, *J. Combin. Theory Ser. B* **92** (2004), 325–357.
- [9] N. Robertson, P.D. Seymour and R. Thomas: *Sachs’ linkless embedding conjecture*, *J. Combin. Theory Ser. B* **64** (1995), 185–227.
- [10] H. Sachs: *On spatial representations of finite graphs*; in *Finite and Infinite Sets, I, II* (Eger, 1981), *Colloq. Math. Soc. János Bolyai* **37**, North-Holland, Amsterdam, 1984, 649–662.

Joel Stephen Foisy
Department of Mathematics
SUNY Potsdam
Potsdam, NY 13676
U.S.A.

Hugh Nelson Howards
Department of Mathematics
Wake Forest University
Winston-Salem, NC 27109
U.S.A.

Natalie Rose Rich
Department of Mathematics
Wake Forest University
Winston-Salem, NC 27109
U.S.A.