

On the Last New Vertex Visited by a Random Walk in a Directed Graph

Calum Buchanan Paul Horn Puck Rombach

February 14, 2023

Abstract

Consider a simple graph in which a random walk begins at a given vertex. It moves at each step with equal probability to any neighbor of its current vertex, and ends when it has visited every vertex. We call such a random walk a *random cover tour*. It is well known that cycles and complete graphs have the property that a random cover tour starting at any vertex is equally likely to end at any other vertex. Ronald Graham asked whether there are any other graphs with this property. In 1993, László Lovász and Peter Winkler showed that cycles and complete graphs are the only undirected graphs with this property. We strengthen this result by showing that cycles and complete graphs (with all edges considered bidirected) are the only directed graphs with this property.

Keywords: random walks, directed graphs, cover tours

2020 Mathematics Subject Classification: 05C81, 05C20

Let G be a connected digraph. A *cover tour* of G from a vertex u is a directed walk which begins at u and ends when it has visited every vertex of G . Lovász and Winkler showed that cycles and complete graphs are the only undirected graphs with the property that a random cover tour beginning at any given vertex is equally likely to end at any other vertex [1]. We address a question posed by Winkler of whether there are any further directed graphs with this property.

We will borrow the notation used in [1]; let $L(u, v)$ be the event that v is the last vertex to be visited by a random cover tour beginning at u . In the case of a bidirected cycle, one might be tempted to believe that a neighbor of the starting vertex should be less likely to be last visited than one further from the start. While this intuition does not hold for cycles, it does hold for many other connected graphs. It is also shown in [1] that, for any two nonadjacent vertices u and v of a bidirected graph G , there is a neighbor x of u such that $\mathbb{P}(L(x, v)) \leq \mathbb{P}(L(u, v))$, and this inequality is strict if the induced subgraph $G - \{u, v\}$ is connected. From the proof, it follows that non-neighbors of a given vertex u are at least as likely as neighbors to be last visited on a random walk from u . We generalize this notion in the following lemma (the proof is nearly identical to that of Theorem 2 in [1]).

Lemma 1. *Let G be a connected digraph, and let u be a vertex in G . If u does not have an edge to a vertex v , then there is an out-neighbor x of u such that $\mathbb{P}(L(x, v)) \leq \mathbb{P}(L(u, v))$. Furthermore, this inequality is strict if there is a directed walk from u to v which visits all vertices in G , but does not revisit u .*

Proof. Let u and v be vertices of a connected digraph G such that $uv \notin E(G)$. Let x_1, \dots, x_d denote the out-neighbors of u , and let $L(x_i; v, u)$ denote the event that u is the last vertex, and v the next-to-last vertex, in a random cover tour of G beginning at x_i . The event $L(u, v)$ is the disjoint union of the events $L(x_i, v)$ and $L(x_i; v, u)$; either the random cover tour from u to v visits u more than once, or it does not. Thus,

$$\mathbb{P}(L(u, v)) = \sum_{i=1}^d \frac{1}{d} (\mathbb{P}(L(x_i, v)) + \mathbb{P}(L(x_i; v, u))). \quad (1)$$

It follows that $\mathbb{P}(L(u, v))$ is at least the mean of the $\mathbb{P}(L(x_i, v))$ over all out-neighbors x_i of u .

The second statement also follows from equation (1). If it is possible for a cover tour to start at u and end at v without revisiting u , then $\mathbb{P}(L(x_i; v, u)) > 0$ for some $i \in \{1, \dots, d\}$. In this case, $\mathbb{P}(L(u, v))$ is strictly larger than the mean of the $\mathbb{P}(L(x_i, v))$. \square

In particular, if T is a cover tour from u to v in a digraph G with the property that a random cover tour from any given vertex is equally likely to end at any other vertex, then either $uv \in E(G)$ or u appears at least twice in T .

Theorem 2. *In any directed graph with the property that $\mathbb{P}(L(u, v)) = \mathbb{P}(L(u, w))$ for any three distinct vertices u, v , and w , every edge is bidirected.*

Proof. Suppose, for the sake of contradiction, that there exists a pair of vertices u and v in a digraph $G = (V, E)$, as described, such that $vu \in E$ and $uv \notin E$. By Lemma 1, any cover tour from u to v in G revisits u . Let T denote a shortest such cover tour, and let T' denote the closed walk $T + vu$.

We begin by considering the cover tour which starts at v , first takes the edge vu , and continues along T until the last seen vertex, v_1 . The vertex v appears only once in this cover tour, since v appears only once in T . By Lemma 1, we have $vv_1 \in E$. Notice also that v_1 appears only once in T . Otherwise, the walk that starts at u , follows T to the first copy of v_1 (visiting every vertex but v), and then follows T from the last copy of v_1 back to v is a cover tour that is shorter than T , which contradicts its minimality.

Now consider the cover tour contained in T' which starts at v_1 and ends at the last seen vertex, v_2 . As before, we have $v_1v_2 \in E$, and v_2 appears only once in T . We can continue in this way until we have chosen all of the vertices in G which appear only once in T , for T' is a closed walk containing all of the vertices of G , and each vertex which appears only once will be the last seen vertex on a walk from some v_i . Let $U = \{v, v_1, \dots, v_k\}$ denote this set of chosen vertices, and let W denote the closed walk that visits them in order. Note that W visits the vertices in U in opposite order to T . (See Figure 1.)




Figure 1: A minimum cover tour T from u to v is depicted by a dashed line, moving in clockwise direction. The edge vu is included to depict the closed walk $T' = T + vu$. The grey arrows indicate the closed walk W , which visits all vertices in U , in counterclockwise direction.

We note that any vertex x in $V \setminus U$ has a copy in T , and the minimality of T ensures that no two copies of x appear in the same interval (v_1, v) , (v, v_k) , or (v_{i+1}, v_i) for $i \in \{1, \dots, k-1\}$. We will now show that $v_1v \in E$ and $v_{i+1}v_i \in E$ for each $i \in \{1, \dots, k-1\}$. However, the minimality of T then implies that there are no vertices between v_1 and v in T ; that is, v_1v is an edge in T . Similarly, each edge $v_{i+1}v_i$ is in T . From this we derive our final contradiction, for u appears at least twice in T by assumption, and the two copies cannot both appear in the interval (v_k, v) .

We will find a cover tour in G from v_1 to v , which does not revisit v_1 . By Lemma 1, this implies that $v_1v \in E$. First, in T' we label the (v_2, v_1) interval A , the (v_1, v) interval B , and the (v, v_k) interval C . Since vertices in $V \setminus U$ appear multiple times on T , but only once in each interval, we use subscripts to differentiate copies. For example, u_C indicates the vertex u in the position on T' in the interval C , as drawn in Figure 1. (We will only need labels for these three intervals.) Further, for any $x, y \in V$, with interval subscripts if necessary, we use the notation $x \xrightarrow{T'} y$ to indicate the walk from x to y along T' , and for v_i, v_j in U we use $v_i \xrightarrow{W} v_j$ to indicate the walk from v_i to v_j along W .

We now present a walk in two parts; the desired cover tour from v_1 to v which does not revisit v_1 is a concatenation of these parts. As a general template, we take Part (I) to be

$$(v_1 \xrightarrow{W} v_k) \cup (v_k \xrightarrow{T'} v_2),$$

and Part (II) to be

$$(v_2 \xrightarrow{W} v).$$

If there no vertices in $V \setminus U$ that appear only in the intervals A , B or C then the concatenation of Parts (I) and (II) yields the desired cover tour. Otherwise, we modify Parts (I) and (II) depending on the intervals in which such vertices appear.

We label vertices in $V \setminus U$ which appear on T' only in A , B , or C as type AB , BC , or AC , depending on which two intervals they appear in. If a vertex appears in all three intervals, we may ignore one copy and label it arbitrarily. Part (I) of our walk starts at v_1 and ends at v_2 , and it covers all vertices of type BC . If there are no BC vertices, then Part (I) is as previously described. Otherwise, let x be the last BC vertex in the interval B . In this case, Part (I) is

$$(v_1 \xrightarrow{T'} x_B) \cup (x_C \xrightarrow{T'} v_2).$$

Since x was the last BC vertex in B , Part (I) of the walk has now covered all BC vertices.

Part (II) of our walk covers all AB and AC vertices. If there are none, then Part (II) is as previously described. Otherwise, let y be the last AB/AC vertex in A . If y is AB , then Part (II) is

$$(v_2 \xrightarrow{T'} y_A) \cup (y_B \xrightarrow{T'} v).$$

If y is AC , then Part (II) is

$$(v_2 \xrightarrow{T'} y_A) \cup (y_C \xrightarrow{T'} v_k) \cup (v_k \xrightarrow{W} v).$$

Since y was the last unseen vertex in A , Part (II) of the walk has now covered all AB and AC vertices. Concatenating Part (I) and Part (II) yields a cover tour from v_1 to v which does not revisit v_1 , and therefore $vv_1 \in E$ (and T). A similar argument shows that $v_{i+1}v_i \in E(G)$ (and T) for all $i \in \{1, \dots, k-1\}$, which completes the proof. \square

Acknowledgements

This work is the result of a collaboration at the 2021 Virtual Masamu Advanced Study Institute (MASI). Paul Horn was partially supported by Simons Collaboration grant 525039. We also thank Daniel Velleman for helpful comments.

References

- [1] L. Lovász, P. Winkler, A Note on the Last New Vertex Visited by a Random Walk, *J. Graph Theory*, Vol. 17, No. 5, (1993) 593-596.