The Prism of the Acyclic Orientation Graph is Hamiltonian

GARA PRUESSE *

Department of Computer Science and Electrical Engineering

University of Vermont, Burlington, VT, 05405-0156, USA

email: `gara@cs.uvm.edu`

FRANK RUSKEY †

Department of Computer Science

University of Victoria, Victoria, B.C., V8W 3P6, Canada

email: `fruskey@csr.uvic.ca`

Submitted: December 20, 1994; Accepted: March 13, 1995.

Abstract

Every connected simple graph G has an acyclic orientation. Define a graph $AO(G)$ whose vertices are the acyclic orientations of G and whose edges join orientations that differ by reversing the direction of a single edge. It was known previously that $AO(G)$ is connected but not necessarily Hamiltonian. However, Squire [3] proved that the square $AO(G)^2$ is Hamiltonian. We prove the slightly stronger result that the prism $AO(G) \times e$ is Hamiltonian.

If G is a mixed graph (some edges directed, but not necessarily all), then $AO(G)$ can be defined as before. The graph $AO(G)$ is again connected but we give examples showing that the prism is not necessarily Hamiltonian.

*This material is based upon work supported by the National Science Foundation under Grant No. NSF OSR-9350540.

†Research supported in part by the Natural Sciences and Engineering Research Council of Canada under Grant A3379.

1 Introduction

Every connected simple graph G has an acyclic orientation. Define a graph $AO(G)$ whose vertices are the acyclic orientations of G and whose edges join orientations that differ by reversing the direction of a single edge. Squire, Savage, and West [4] show that $AO(G)$ is connected and bipartite and discuss the Hamiltonicity of $AO(G)$ for well-known types of graphs such as trees, cycles, wheels, and ladders. Squire [3] proved that the square $AO(G)^2$ is Hamiltonian. Our main goal in this paper is to prove the slightly stronger result that the prism $AO(G) \times e$ is Hamiltonian. If the prism of a bipartite graph is Hamiltonian, then the square of the graph is Hamiltonian.

If G is a mixed graph (some edges directed, but not necessarily all), then $AO(G)$ can be defined as before. Mixed graphs are interesting because special cases of them reduce to known problems and results. For example, there is a one-to-one correspondence between the mixed complete graphs (every edge is present, either directed or not) and directed acyclic graphs. If G is a mixed complete graph, then a topological sort of the maximum directed subgraph of G corresponds to an acyclic orientation of G. Pruesse and Ruskey [1] showed that the prism of $AO(G)$ is Hamiltonian in the case of mixed complete graphs. It is therefore natural to ask whether the prism is Hamiltonian for all mixed graphs, since it is true for the cases of no fixed edges (this paper) and when all edges are present ([1]). The graph $AO(G)$ is again connected and bipartite; however, the final section of this paper describes an infinite class whose prisms are not Hamiltonian.

We now define a few terms and introduce some notation. A *poset* $\mathcal{P} = (S, R)$ is a reflexive, anti-symmetric, transitive relation $R(\mathcal{P})$ on a set $S(\mathcal{P})$. We let $\mathcal{P} - x$ denote the poset on the set $S(\mathcal{P}) \setminus x$, together with all the relations of $R(\mathcal{P})$ that do not involve x. An *ideal* of a poset $\mathcal{P}$ is a set $I \subseteq S(\mathcal{P})$ for which $y \in I$ and $x \prec y$ imply that $x \in I$. The *ideal graph* $J(\mathcal{P})$ has vertices which are the ideals of $\mathcal{P}$ and edges connecting those ideals that differ by one element; in other words, it is the Hasse diagram of the lattice of ideals of a poset, regarded as an undirected graph. In order to simplify notation we will usually write x for the singleton set $\{x\}$ and use juxtaposition to indicate union of sets. E.g., if E is a set and x and element, Ex denotes $E \cup \{x\}$.

For a graph G, the *square* of G, denoted G^2, is the graph on the same vertex set as G and where $(u, v) \in E(G^2)$ if and only if the distance between u and v in G is 1 or 2. The *prism* of a graph G, denoted $G \times e$, is obtained by taking two copies of G and adding edges between each corresponding pair of vertices in the two copies.

We think of the vertices in the prism as being signed; with each vertex v in G there is associated, in $G \times e$, a plus vertex $+v$ and a minus vertex $-v$. All of the plus vertices occur in one copy of G, the minus vertices in the other copy of G.

The following lemma, from [1], relates the hamitonicity of prisms and squares for bipartite graphs.

Lemma 1 *If G is bipartite and $G \times e$ is hamiltonian, then G^2 is hamiltonian.*

Pruesse and Ruskey [2] proved that the prism of the ideal graph is Hamiltonian.

Theorem 1 *For every poset $\mathcal{P}$ there is a Hamilton path in $J(\mathcal{P}) \times e$ from $+\emptyset$ to $-\emptyset$.*

It also proves useful to have a Hamilton path in the prism that ends at $S(\mathcal{P})$.

Theorem 2 *For every poset $\mathcal{P}$ there is a Hamilton path in $J(\mathcal{P}) \times e$ of the form $+\emptyset, -\emptyset, \ldots, \pm E$, where $E = S(\mathcal{P})$ and the sign of the final vertex depends on the parity of $|E|$.*

Proof: Let $n = |E|$. We argue by induction on n. If $n = 1$, then $+\emptyset, -\emptyset, -x, +x$ satisfies the conditions of the theorem. Note that this path also satisfies the following property: Between $\pm I$ and $\mp I$ only subsets of I occur. We argue that this property, called the *subset property*, is maintained inductively.

Let $n > 1$ and let x be a minimal element of $\mathcal{P}$. Inductively, there is a Hamilton path $H = +\emptyset, -\emptyset, \ldots, \pm E$ in $\mathcal{P} - x$ which has the subset property. The path $P = -x, +x, \ldots, \mp Ex$ in $\mathcal{P}$, obtained by adding x to each vertex of H and reversing all signs, includes all ideals that contain x. We now splice in those ideals that don't contain x. First, replace the edge $[-x, +x]$ with the path $+\emptyset, -\emptyset, -x, +x$. Now, starting at the vertex following $+x$, find the first vertex $\pm Ix$ for which I is an ideal of $\mathcal{P}$, and replace the subpath $\pm Ix = \alpha_1 x, \alpha_2 x, \ldots, \alpha_p x = \mp Ix$ in P by

$$\alpha_1 x, \alpha_1, \alpha_p, \ldots, \alpha_2, \alpha_2 x, \ldots, \alpha_p x. \tag{1}$$

Observe that, by the subset property, α_i for $i = 1, 2, \ldots, p$ are all ideals of $\mathcal{P}$. Repeat the above process, starting at the vertex following $\mp Ix$, and so on, and call the resulting Hamilton path H'. Note that if H has the subset property, then so does H', because each subpath (1) has the subset property. □

These two theorems will be used in the proof of the next section.

2 The Prism is Hamiltonian

Theorem 3 *If* G *is a simple graph, then* $AO(G) \times e$ *is Hamiltonian.*

Proof: Our proof is by induction on $n = |V(G)|$. If $n = 1$, then $AO(G) \times e$ is a single edge, which we regard as being Hamiltonian, as it has a Hamilton path whose end points are adjacent.

For $n > 1$, pick an arbitrary vertex v and remove it from G. Inductively, there is a Hamiltonian cycle H in $(AO(G-v)) \times e$. Consider a generic acyclic orientation of $G-v$, call it Γ, and define a poset $\mathcal{P}_\Gamma$ as follows. The elements, $E = S(\mathcal{P}_\Gamma)$, of $\mathcal{P}_\Gamma$ are $N(v)$, the vertices in the open neighborhood of v. Note that E depends only on v, and not otherwise on Γ. The relations of $\mathcal{P}_\Gamma$ are

$$R(\mathcal{P}_\Gamma) = \{x \prec y \mid \text{ there is a directed path from } y \text{ to } x \text{ in } \Gamma\}.$$

Note that Γ is a partial acyclic orientation of G; Γ orients all edges not incident with v. Orienting the remaining edges of G is called *extending* Γ.

The key observation is that there is a one-to-one correspondence between the ideals of $\mathcal{P}_\Gamma$ and the acyclic orientations of G that extend Γ. Let I be an ideal of $\mathcal{P}_\Gamma$, and orient all edges of the form $[u, v]$ either $u \leftarrow v$ or $u \rightarrow v$ depending on whether $u \in I$ or $u \notin I$, respectively; this yields an orientation of G that extends Γ, and which we will denote $\langle I \rangle_\Gamma$. We show that for each ideal I of $\mathcal{P}_\Gamma$, $\langle I \rangle_\Gamma$ is acyclic. Since Γ is acyclic, it is sufficient to show that v is not in any cycle. In the trivial cases $\langle \emptyset \rangle_\Gamma$ means that all edges are directed into v and $\langle E \rangle_\Gamma$ means that all edges are directed away from v, so acyclicity holds. In any case, if $u_1 u_2 \cdots u_k$ is any directed path in Γ with endpoints in $N(v)$, then $u_1 \in I$ implies $u_k \in I$, and hence v has a directed edge to u_1 only if it also has one to u_k. It follows that v is not in any directed cycle in $\langle I \rangle_\Gamma$.

The argument in the other direction (i.e., that each acyclic orientation of Γ is $\langle I \rangle_\Gamma$ for some $I \in J(\mathcal{P}_\Gamma)$), is also straightforward.

The addition or removal of a single element from an ideal to obtain another ideal corresponds to changing the orientation of an edge. Thus by traversing the ideals of $J(\mathcal{P}_\Gamma) \times e$ we are following a path in a certain subgraph of $AO(G) \times e$, namely that subgraph induced by all extensions of Γ.

We observe that H can be partitioned into

- [**positive pairs**] edges of the form $[+\Gamma, +\Gamma']$, and
- [**mates**] edges of the form $[+\Gamma, -\Gamma]$, and

- [**loners**] vertices of the form $-\Gamma$.

One way to obtain such a partitioning is to treat H as a sequence from which we "shell" (remove from the front) pairs and singletons of vertices. First rotate H as necessary so that the last vertex in H is negative. Now begin shelling: if the sequence begins with a positive vertex, shell a pair; if it begins with a negative vertex, shell a "loner". Since the sequence ends with a negative vertex, every element will be shelled, and hence H is partitioned into positive pairs, mates, and loners.

Our task now is to replace each orientation Γ of $(G-v)$ on H with its extensions in G. In the case of "mates", replace the edge $[+\Gamma, -\Gamma]$ by the path $+\langle\emptyset\rangle_\Gamma, \ldots, -\langle\emptyset\rangle_\Gamma$, where $+\emptyset, \ldots, -\emptyset$ is the Hamilton cycle in $J(\mathcal{P}_\Gamma) \times e$ guaranteed by Theorem 1.

Otherwise, $+\Gamma$ occurs in a "positive pair", and $-\Gamma$ occurs as a "loner". In the path $-\langle\emptyset\rangle_\Gamma, +\langle\emptyset\rangle_\Gamma, \ldots, \mp\langle E\rangle_\Gamma$ in $J(\mathcal{P}_\Gamma) \times e$ (whose existence is guaranteed by Theorem 2), we remove the first edge, obtaining 2 paths, $-\langle\emptyset\rangle_\Gamma$ and $+\langle\emptyset\rangle_\Gamma, \ldots, \mp\langle E\rangle_\Gamma$. For each positive pair, replace the edge $[+\Gamma, +\Gamma']$ by the path

$$+\langle\emptyset\rangle_\Gamma, \ldots, +\langle E\rangle_\Gamma, +\langle E\rangle_{\Gamma'}, \ldots, +\langle\emptyset\rangle_{\Gamma'},$$

if $|E|$ is even, or

$$+\langle\emptyset\rangle_\Gamma, \ldots, -\langle E\rangle_\Gamma, -\langle E\rangle_{\Gamma'}, \ldots, +\langle\emptyset\rangle_{\Gamma'},$$

if $|E|$ is odd.

In the case of "loners", replace the vertex $-\Gamma$ with the vertex $-\langle\emptyset\rangle_\Gamma$.

After doing the above substitutions H has been transformed into a Hamilton cycle in $AO(G) \times e$. □

The above proof does not extend to the case of mixed graphs. The proof uses the fact that directing all v's edges away from v extends Γ acyclicly. However, this is no longer the case when some edges incident with v have fixed orientation; indeed, the cardinality of the maximum set of edges that can be directed away from v can change greatly by changing the orientation of one edge of Γ.

3 The Example

Let $K_{n,n}$ denote the complete bipartite graph and M be a perfect matching in that graph. Define the mixed graph G_n to be $K_{n,n}$ with all edges directed from one partite set to the other, except for the edges in the matching M, which are left undirected.

Lemma 2 *For all* $n \geq 1$*, we have* $AO(G_n) = K_{1,n}$.

Proof: Call the two partite sets A and B and assume that the directed edges are all of the form $a \to b$, where $a \in A$ and $b \in B$. Note that any two edges of M directed from B to A will induce a 4-cycle. Thus at most one of the edges in M can be directed from B to A and still preserve the acyclic property. Each orientation with an edge from B to A is only adjacent to the orientation with no edges going from B to A. □

For $n \geq 3$, the graph $K_{1,n} \times e$ is not Hamiltonian since the removal of the two central vertices results in n connected components. Since the square of $K_{1,n}$ is Hamiltonian, we can still ask whether there is a mixed graph G such that $AO(G)^2$ is not Hamiltonian. Squire [3] has extended our Lemma 2 to give such an example. An open problem is to find an algorithm for generating acyclic orientations in time proportional to the number of orientations. If this problem is pursued based on the proof of this paper then a first step is finding an algorithm for generating the ideals of a poset in time proportional to the number of ideals, which is an open problem.

References

[1] G. Pruesse and F. Ruskey, *Generating Linear Extensions Fast*, SIAM J. Computing, 23 (1994) 373-386.

[2] G. Pruesse and F. Ruskey, *Gray Codes from Antimatroids*, Order, 10 (1993) 239-252.

[3] M. Squire, *Two New Gray Codes for Acyclic Orientations*, 94-14, Dept. Computer Science, N.C. State Univ, 1994.

[4] C. Savage, M. Squire, and D. West, *Gray Code results for Acyclic Orientations*, Congressus Numerantium, 96 (1993) 185-204.