Hamiltonian Cycles in Solid Grid Graphs
(Extended Abstract)

Christopher Umans*

Computer Science Division
U.C. Berkeley
umans@cs.berkeley.edu

William Lenhart

Computer Science Department
Williams College
lenhart@cs.williams.edu

Abstract

A grid graph *is a finite node-induced subgraph of the infinite two-dimensional integer grid. A* solid grid graph *is a grid graph without holes. For general grid graphs, the Hamiltonian cycle problem is known to be $\mathcal{NP}$-complete. We give a polynomial-time algorithm for the Hamiltonian cycle problem in solid grid graphs, resolving a longstanding open question posed in [IPS82]. In fact, our algorithm can identify Hamiltonian cycles in* quad-quad graphs*, a class of graphs that properly includes solid grid graphs.*

1 Introduction

A *grid graph* is a finite node-induced subgraph of the infinite two-dimensional integer grid. A *solid grid graph* is a grid graph all of whose bounded faces have area one. The study of Hamiltonian cycles in grid graphs was initiated by Itai, Papadimitriou and Szwarcfiter [IPS82], who proved that the problem for general grid graphs is $\mathcal{NP}$-complete, and gave a polynomial-time algorithm for rectangular solid grid graphs. The question of whether a polynomial-time algorithm exists for general solid grid graphs has remained unresolved until now.

Previous work has yielded efficient algorithms for certain smaller subclasses of grid graphs. Everett [Eve86] extended the results in [IPS82] to solve the Hamiltonian path problem in several specific restrictions of grid graphs; her algorithms generally decompose the graph into rectangular sub-problems. More recently, Afrati [Afr94] gave a linear time algorithm for *restricted grids*, a subclass of solid grid graphs which are essentially left-aligned stacks of decreasing-length grid strips. However, none of the previously used methods work on general solid grid graphs, for which the problem seems significantly more difficult. Cho and Zelikovsky [CZ95], who study spanning closed trails (a relaxation of Hamiltonian cycles) in a broad subclass of grid graphs called *polymino*, have even conjectured that for solid grid graphs, deciding Hamiltonicity is $\mathcal{NP}$-complete.

We present a polynomial-time algorithm that finds Hamiltonian cycles in solid grid graphs using the well-known technique of *cycle merging*. Given an input graph G, we first find a *2-factor*, which is a spanning subgraph for which all vertices have degree two. The 2-factor is a set of disjoint cycles that exactly cover the vertices of G; a Hamiltonian cycle is a 2-factor with a single component. We then repeatedly identify a transformation of the 2-factor that reduces the number of components. This process either identifies a Hamiltonian cycle or terminates with multiple components if one does not exist.

Our algorithm can be applied to a generalization of solid grid graphs which are "locally" solid grid graphs but may not be fully embeddable in the integer grid without overlap. We call these graphs *quad-quad graphs* because they can be characterized as follows: a quad-quad graph is a planar graph with maximum degree four whose bounded faces are quadrilaterals and whose dual (excluding the vertex corresponding to the unbounded face) also has only quadrilateral bounded faces. Figure 1 shows a general grid graph (a), a solid grid graph (b), and a quad-quad graph that is not a solid grid graph (c).

The remainder of the paper is organized as follows: §2 defines *alternating strip sequences* and shows that they can be used to reduce the number of components of a 2-factor; §3 proves their existence in Hamiltonian quad-quad graphs; §4 shows that the search for alternating strip sequences can be restricted to *static alternating strip sequences*, which can be located efficiently. Finally, §5 presents the algorithm and a brief analysis of its time complexity.

*Research supported in part by an NSF Graduate Research Fellowship.

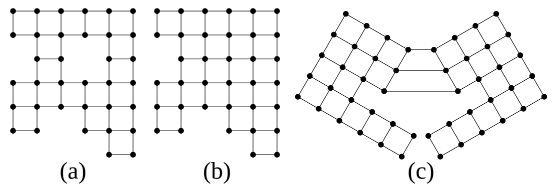

Figure 1: Examples of a general grid graph (a), a solid grid graph (b), and a quad-quad graph (c).

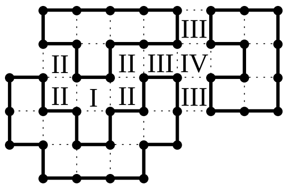

Figure 2: The four types of boundary cells.

2 Alternating Strip Sequences

The first step of the algorithm is to identify a 2-factor F of the input graph G. Clearly, if one does not exist, then the graph is not Hamiltonian. Otherwise, we attempt to transform F by merging cycles. Most of the remainder of the paper deals with structures involved in this transformation process. For ease of exposition, we assume, unless otherwise noted, that these structures reside in quad-quad graph G with respect to 2-factor F; in figures, we represent them as subgraphs of G in which the edges are colored dark if they belong to F and light otherwise. Also, we sometimes refer to edges that are in F as *dark* edges and those that are not in F as *light* edges. In this section we define the basic components used to transform F.

We begin with some definitions: A *border* is a cut $(S, V(G) - S)$, for which (1) no edge of F crosses the cut and (2) the node-induced subgraphs of S and $V(G) - S$ are connected. A *cell* is a bounded face, and a *boundary cell* is a cell that includes at least one edge that crosses a border. A *boundary* is the set of boundary cells associated with a given border. Boundary cells may be classified by the pattern of light and dark edges on their perimeter. There are four such configurations, and we number them in figure 2. Finally, an *alternating cycle* is a cycle whose edges are alternately in F and not in F.

In general, we can modify a 2-factor F by *flipping* an alternating cycle C – those edges that are in F are removed and those that are not in F are added, leaving a new 2-factor, which we denote $F \oplus C$. We now define a particular class of alternating cycles, called *alternating strips*, which, when flipped, have well-defined effects on the the number of components in the 2-factor.

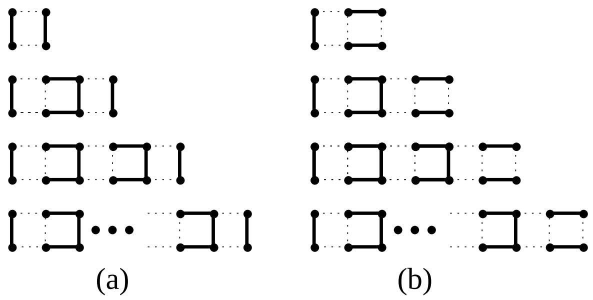

Figure 3: Odd (a) and even (b) alternating strips.

Definition 1 *An* **alternating strip** *is one of the subgraphs shown in figure 3. The subgraphs pictured in (a) are* **odd alternating strips***; those in (b) are* **even alternating strips**.

We say that an alternating strip *begins* at the leftmost cell of a row in figure 3 and *ends* at the rightmost cell. An alternating strip's *length* is the number of cells it includes. By a slight abuse of notation, we denote by $F \oplus s$ the 2-factor obtained by flipping the perimeter of alternating strip s. To establish the effect of this transformation on the number of components in the 2-factor, we need to consider flipping the strip one cell at a time, from the end to the beginning. If we proceed in this manner, each successive cell is an alternating cycle when it is flipped, so we maintain a 2-factor throughout the flipping process. Also, it should be clear that flipping the alternating strip in stages is equivalent to flipping it all at once, since any interior edges of the strip that are flipped are flipped twice.

Observation 1 *Flipping a Type III boundary cell (which is also an alternating strip of length one) reduces the number of components in F by one.*

Lemma 2 *Flipping the last two cells of an odd alternating strip of length at least three leaves the number of components of the 2-factor unchanged.*

Proof: By enumeration of cases, shown in figure 4(b) and (c). ■

Lemma 3 *Let s be an alternating strip that begins on a boundary cell. Then flipping s reduces the number of components of the 2-factor by one if s is an odd*

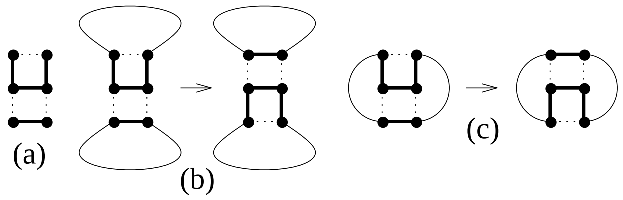

Figure 4: The last two cells of an odd alternating strip of length at least three.

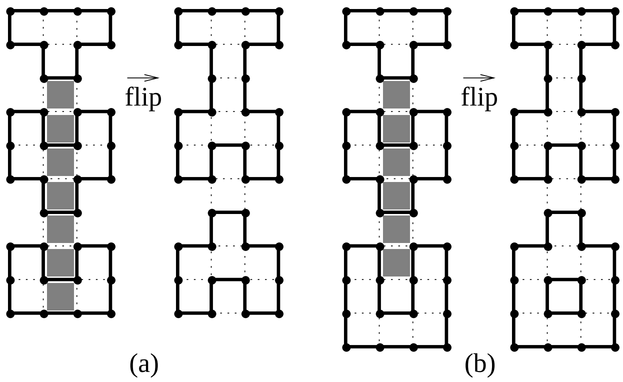

Figure 5: The effect of flipping an odd alternating strip (a) and an even alternating strip (b) as specified in Lemma 3. The boundary between the nested cycles in (b) is a boundary created by flipping the even strip.

alternating strip and leaves the number of components unchanged if s is an even alternating strip that does not end on a boundary cell.

Proof sketch: As noted, the lemma is immediate for the strip of length one; for odd length strips longer than one, we use the previous lemma and induction on the length of the strip. For even length strips, flipping the end cell increases the number of components by one; flipping the remaining odd length strip decreases the number of components by one, resulting in a net change of zero. ∎

Figure 5 illustrates the effect of flipping the two kinds of alternating strips dealt with in Lemma 3.

After flipping an even alternating strip s of the type specified in the lemma, the end cell becomes a Type I boundary cell (with respect to the new 2-factor). We omit the proof, which consists mainly of verifying that any new boundary cell created by flipping *only* the end cell of s remains a boundary cell after flipping the remainder of the strip. We say that a boundary including this Type I cell is a boundary *created* by s.

It is a consequence of later theorems that if G is Hamiltonian, it must contain an alternating strip that begins on a boundary. Unfortunately, this may be an even alternating strip, which will not improve (i.e. reduce the number of components in) F when flipped. However, by the same argument, an alternating strip must also exist beginning on a "new" boundary created by flipping the first strip. If it is odd, then the two strips flipped in succession will improve F. Otherwise, we can repeatedly find additional strips, each beginning on a boundary created by flipping the previous strip. We would like to be able to identify such a sequence that ends with an odd alternating strip, so that flipping the sequence of strips in succession improves F. This observation motivates the following definition:

Definition 2 *An* **alternating strip sequence** *is defined as follows:*

1. *an odd alternating strip is an alternating strip sequence, and*
2. *if G contains no Type III boundary cells, s is an even alternating strip in G, and $(a_1, a_2, \ldots, a_n)$ is an alternating strip sequence in G with respect to $F \oplus s$ whose first strip begins on a boundary created by s, then $(s, a_1, a_2, \ldots, a_n)$ is an alternating strip sequence in G.*

An example of an alternating strip sequence can be seen in figure 12. The following is immediate from the definition and Lemma 3:

Lemma 4 *Flipping the alternating strips of an alternating strip sequence in succession results in a new 2-factor with one fewer component than the original 2-factor.*

The critical relationship between alternating strip sequences and Hamiltonicity is seen in the following theorem (to be proved in the next section).

Theorem 5 (Existence of Alternating Strip Sequences) *Let F be a 2-factor of G with more than one component. If G is Hamiltonian then G with respect to F contains an alternating strip sequence.*

3 Existence Theorem

Our strategy in the proof of Theorem 5 is the following: first, we show that any 2-factor F can be transformed into any other by flipping a set of cycles in G that alternate with respect to F. Assuming the graph is Hamiltonian, we consider a set S of cycles that will

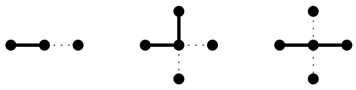

Figure 6: Legal configurations of edges in $F_1 \oplus F_2$. Dark edges belong to F_1, light edges to F_2.

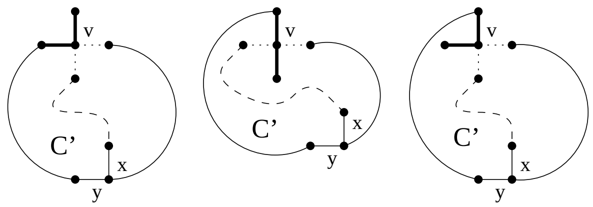

Figure 7: Possible configurations of alternating cycle C with a crossing cycle at vertex v.

transform F into a Hamiltonian cycle H. We show that for each boundary, at least one such cycle must cross it, and that we can perform a cell-to-cell walk from this crossing, along a portion of the boundary, and then along an alternating strip that begins on that boundary. We then show that flipping this alternating strip results in a new 2-factor that is "closer" to H than F. We repeatedly identify successive strips in an alternating strip sequence in this manner, eventually terminating with an odd alternating strip since we are getting progressively closer to H (we must merge cycles at some point, which can only be done by flipping an odd alternating strip).

The cell-to-cell walk and the distance measure that we use as the induction variable require an auxiliary directed "dependency graph;" this and other preliminaries are covered in the next three subsections.

3.1 Alternating Cycles

Two edge-disjoint cycles are *nested* if all the interior faces of one are interior faces of the other; two edge-disjoint cycles *cross* if they share at least one interior face and are not nested. The *area* of a cycle is the number of faces it encloses.

Lemma 6 *Let F_1 and F_2 be 2-factors of G and let $S = F_1 \oplus F_2$. Then S can be partitioned into edge-disjoint non-crossing cycles that alternate with respect to F_1.*

Proof: Figure 6 shows the possible edge configurations around each vertex in S. Other configurations imply that one or both of F_1 and F_2 are not 2-factors. It is therefore possible to walk along alternating edges of S until an alternating cycle is identified. This process can be repeated until all of S is partitioned into alternating cycles.

To show that the cycles can be chosen to be non-crossing, we prove that an alternating cycle C in S of minimal area crosses no other cycle in S. Suppose for the purpose of contradiction that some cycle in S does cross C. The point of crossing, v, must occur at a degree 4 vertex, giving rise to the three possible configurations in figure 7. It is then possible to perform an alternating walk along the dashed path until crossing C again at a degree 4 vertex. We know that edges x and y alternate; otherwise cycle C' as shown in the figure is an odd cycle in a bipartite graph. So cycle C' is an alternating cycle with area strictly smaller than C, contradicting our initial assumption. ■

3.2 The Dependency Graph

We first partition $E(G)$ into horizontal and vertical edge sets H and V and define two related sets $E_1(F)$ and $E_2(F)$ as follows:

$$E_1(F) \;=\; \{e \mid e \in H \wedge e \in F\} \cup \{e \mid e \in V \wedge e \notin F\}$$
$$E_2(F) \;=\; \{e \mid e \in V \wedge e \in F\} \cup \{e \mid e \in H \wedge e \notin F\}.$$

We define G^+, a variant of G^* (the standard dual of G,) as G^* minus the unbounded face and augmented with edges crossing each edge of G that borders the unbounded face. Each of these edges requires an additional endpoint in the vertex set. Let $V_1 \cup V_2 = V(G^+)$ be a bipartition of the vertices of G^+. The dependency graph for G with respect to F is defined as follows:

Definition 3 *The* **dependency graph** *for G with respect to F, denoted G_F^+, is obtained by directing the edges of G^+ so that $(u,v) \in E(G_F^+)$ if and only if $u \in V_k$ and (u,v) crosses some $e \in E_k(F)$ for $k \in \{1,2\}\}$.*

Because $E_1(F)$ and $E_2(F)$ partition $E(G)$ and $V_1 \cup V_2$ is a bipartition of the vertices of G^+, G_F^+ is well defined in that each edge of G^+ is assigned a unique direction. The next lemma describes an important relationship between the dependency graph and alternating cycles.

Lemma 7 *A cycle C in G alternates with respect to F if and only if every edge of G_F^+ that crosses C is directed uniformly with respect to the interior of C.*

Proof: An enumeration of cases shows edges of G_F^+ to be uniformly directed along pairs of successive edges of an alternating path in G with respect to F. Conversely, pairs of uniformly directed dependency graph arcs crossing adjacent edges of G imply primal edges that alternate with respect to F. ■

Figure 8 gives an example of a dependency graph G_F^+ overlaid on the associated graph G (with edges shaded with respect to F).

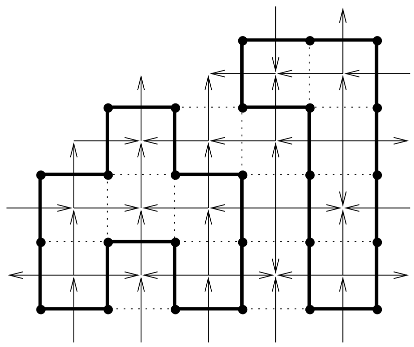

Figure 8: An example of a dependency graph.

3.3 Distance Between 2-Factors

Let F_1 and F_2 be 2-factors of G, and let S be a partition of $F_1 \oplus F_2$ into edge-disjoint non-crossing cycles that alternate with respect to F_1. We label each cycle with a 1 or -1 with respect to S by repeated application of the following rule: label with a 1 all outermost cycles; label with a -1 each cycle nested immediately within an outermost cycle that is crossed by arcs of $G_{F_1}^+$ oriented in the opposite direction as the arcs crossing the outer cycle; remove all labeled cycles.

For ease of notation, let $\ell(C, S)$ be the label of cycle C with respect to S as defined above. We now define the distance between F_1 and F_2. This measure depends on the particular partition S chosen from among the valid partitions of $F_1 \oplus F_2$, so we define the distance as the minimum over all such partitions:

Definition 4 *The* **distance** *from F_1 to F_2 is defined as:*

$$d(F_1, F_2) = \min_S \sum_{C \in S} \ell(C, S) \cdot area(C).$$

It is difficult to describe the intuition behind this measure of closeness; it is best seen simply as the formula that best facilitates the proof of Theorem 5.

3.4 Proof of Theorem 5

The theorem holds trivially if G contains a Type III boundary cell. We prove the non-trivial case by induction:

Induction Hypothesis: Let H be a Hamiltonian cycle in G, let F be a 2-factor of G with more than one component, and let B be a boundary in G. If G contains no Type III boundary cells then G contains an alternating strip sequence whose first strip begins on B.

Since the presence of a Type IV boundary cell implies the presence of a Type III boundary cell we can assume that G contains only Type I and Type II boundary cells.

Let S be a partition of $F \oplus H$ into edge-disjoint non-crossing alternating cycles that minimizes the distance formula. Some cycle $C \in S$ must cross B, for if none did, the boundary would remain in $F \oplus (F \oplus H) = H$, contradicting the choice of H to be Hamiltonian. Specifically, there must be an edge $e \in C$ between two components separated by B. We now describe a directed walk in G_F^+, beginning on the arc that crosses e, that identifies an alternating strip beginning on B. Phase I of the walk is along B; the walk proceeds outward with respect to C's interior if $l(C, S) = -1$ and inward if $l(C, S) = 1$. This phase lasts until a Type I boundary cell is encountered. Figure 9(a, b) depicts the two possible walks in this phase.

Claim 1 *Phase I must reach a Type I boundary cell.*

Proof: By assumption there are only Type I and Type II boundary cells in G with respect to F. Suppose for the purpose of contradiction that no Type I boundary cell is encountered, so the walk is entirely along Type II boundary cells. The dependency arcs along this walk form a directed path. If the walk proceeds toward the interior of C, then it eventually exits C, contradicting Lemma 7. If the walk proceeds outward, then $l(C, S) = -1$, and by the labeling procedure, C must be immediately nested within a cycle $C' \in S$ with $l(C', S) = 1$. The walk cannot cross C' since, by the labeling procedure, the arcs crossing C' and C are oriented oppositely with respect to their interiors. Therefore the walk must cycle back to edge e (recall the definitions of borders and boundaries), crossing C a second time, and contradicting Lemma 7. ■

Phase II of the walk proceeds in a straight line away from the dark edge of the Type I boundary cell that ended Phase I, as pictured in 9(c). This phase continues until we reach a dependency arc oriented against the directed path, as shown in the figure.

Claim 2 *Phase II proceeds along exactly the cells of an alternating strip.*

Proof: For convenience, we refer to the graph as it is oriented in figure 9(c) when we use terms such as

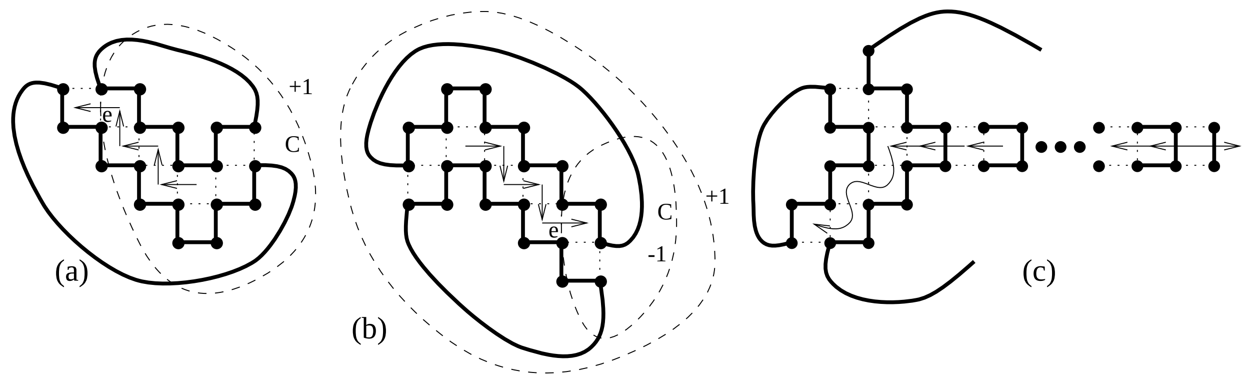

Figure 9: The walk along boundary B from edge e. The two cases of Phase I are depicted in (a) and (b); Phase II is shown in (c).

"left," "right," and "top." Also, we number the sequence of cells encountered along Phase II of the walk from left to right, starting with the Type I boundary cell as number one.

Since cell one is a Type I boundary cell, we know that its top, bottom, and right edges are light. Therefore, the top and bottom edges of cell two must be dark, for otherwise one of the rightmost vertices of cell one would have degree less than 2 in F. If the right edge of cell two were light, then the walk would end, since the dependency arc crossing it would point right. Otherwise (as pictured), the right edge of cell two is dark, and we can now argue that that the top and bottom edges of cell three must be light, for otherwise one of the rightmost vertices of cell two would have degree greater than 2 in F. Again, if the walk does not end, the rightmost edge of cell three must be light, which forces the top and bottom edges of cell four to be dark. As long as the directed walk continues, the sets of top, bottom and right edges of each successive cell alternate between being light and dark.

Recall from the proof of claim 1 that the entire directed walk occurs within an alternating cycle with label "1". Phase II must therefore terminate at an alternating cell as shown, since otherwise the walk would cross this enclosing cycle with an arc directed in contradiction of Lemma 7. The cells covered by Phase II of the walk are exactly the cells of an alternating strip. ■

Call this alternating strip s; we claim that flipping it results in a 2-factor that is closer to H than F.

Claim 3 $d(F \oplus s, H) < d(F, H)$.

Proof: We consider flipping s one cell at a time, starting with the ending cell c. We first note that c is on the *immediate interior* of some cycle $C \in S$ with $l(C, S) = 1$, by which we mean c is within C but is enclosed by no cycles of S nested within C. This is a consequence of the initial choice of the direction of the walk, the fact that the walk is directed, the labeling procedure, and Lemma 7. Since cell c is itself an alternating cycle, the four arcs of G_F^+ that cross it are all oriented uniformly by Lemma 7, and by the choice of the walk, they are oriented in the same direction as the arcs crossing C with respect to the interiors of c and C. We now examine the effect of flipping c on the alternating cycles of S. If c shares no edges with any cycles in S, then in $F \oplus c$, c itself is a new alternating cycle with $l(c, S) = -1$ and area one. If c shares edges only with C, then in $F \oplus c$ the area of C decreases by one. If c shares edges only with nested cycles, then they must have label "-1", and in $F \oplus c$ their collective area increases by one. Finally if c shares edges with both C and a nested cycle C', then in $F \oplus c$, C and C' are merged into one cycle with label "1" and area one less than the difference in areas of the original cycles. This general characterization can be verified for all specific cases; figure 10 shows examples of each of the four cases described. The important fact is that the distance formula is reduced by one in $F \oplus c$.

After flipping c, we note that the remainder of s is itself an alternating strip s', which can be identified by the same walk used to identify s, ending one cell sooner. As s is flipped one cell at a time, we decrease the value of the distance formula by one each time, so we have that $d(F \oplus s, H) < d(F, H)$. A final technical detail requires mention: in each step of this proof we exhibit a partition S' of $F \oplus c$ based on local changes to the previous partition S of F. It may be that S' does not minimize the distance formula, but then the

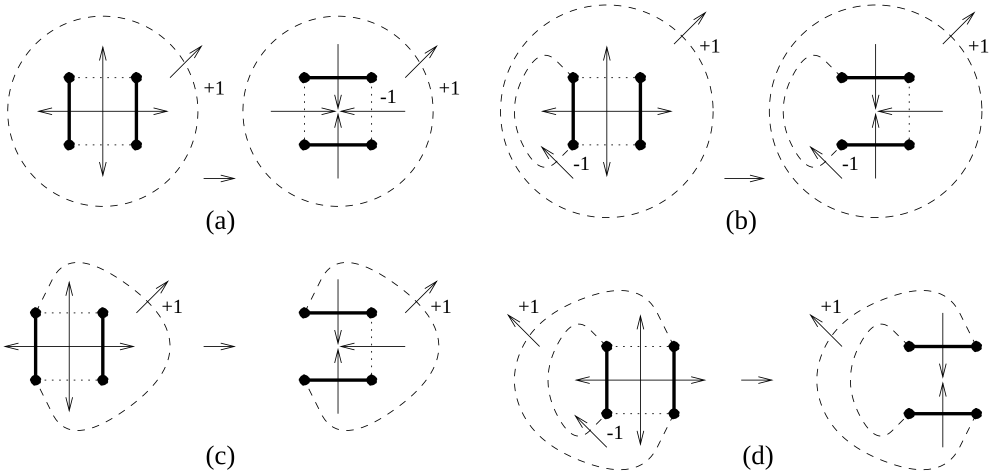

Figure 10: Representative examples of the four ways flipping the cell at the end of alternating strip s affects the partition into alternating cycles and the cycle labeling. Arrows represent arcs in the dependency graph.

true minimizing partition can only make the distance smaller still. ■

If s is an odd alternating strip, then we are done; if $F \oplus s$ contains any newly created Type III boundary cells, then we can show that there must be a Type III boundary cell on a boundary created by s (proof omitted). In this case s followed by a new Type III boundary cell is an alternating strip sequence and we are done. If neither of these two terminal cases hold, then there are no Type III boundary cells in $F' = F \oplus s$, so by induction, we can find an alternating strip sequence beginning on a boundary created by s. Appending this sequence to s forms an alternating strip sequence in G that begins on boundary B as required.

This induction must terminate, because the distance to H is reduced at each step. Since $d(F, H) = 0$ implies that $F = H$, we must have merged components at some point, so we cannot have identified a sequence of *only* even alternating strips.

4 Static Alternating Strip Sequences

Alternating strip sequences can be described as *dynamic* or *static* depending on whether the constituent strips overlap or share edges. If overlaps or shared edges are allowed the individual alternating strips might not be recognizable as such until prior strips in the sequence have been flipped. Since the algorithm has no *a priori* knowledge of the sequence, this situation is problematic. In this section we define static alternating strip sequences and show that their existence follows from the existence of general alternating strip sequences.

Definition 5 *An alternating strip sequence $A = (a_1, a_2, \ldots, a_n)$ is* **static** *if the following hold: 1) the a_i's are area-disjoint, and 2) no two strips share an edge, with the possible exception of any two consecutive strips a_i and a_{i+1}, which may share the single edge e in the manner shown in figure 11.*

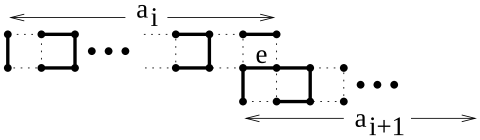

Figure 11: The only allowed shared edge in a static alternating strip sequence, pictured immediately before a_i is to be flipped.

Figure 12 gives examples of non-static and static alternating strip sequences.

We prove the existence of static alternating strip sequences by showing that we can "cut out" extraneous portions of alternating strip sequences until a static sequence remains. In discarding these portions, we need to ensure that boundaries whose existence is required later in the sequence are preserved. Fortunately, a boundary that arises in the course of flipping an alternating strip sequence is, at worst, only "shifted" by flipping an even alternating strip that might pass through it. This observation is formalized in the following technical lemma (proof omitted), which is used

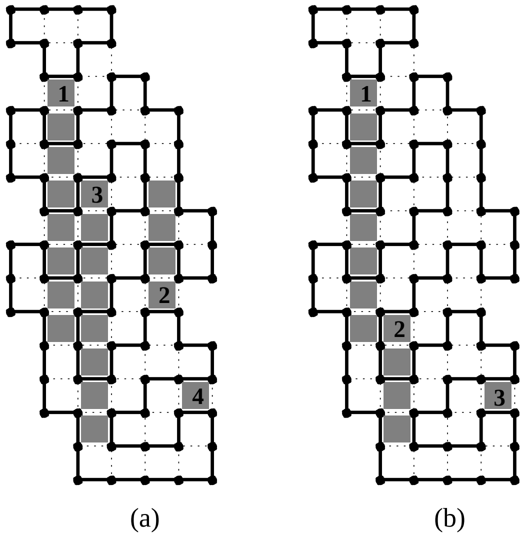

Figure 12: An alternating strip sequence that is not static (a), and the static alternating strip sequence derived from it (b). The numbers label the beginning cells of the strips in the order that they appear in the sequence.

heavily in the full proof of the theorem. To state it, we need the following notation: by G_F^- we denote the graph obtained by starting with G^* (the standard dual of G) and deleting all edges that cross an edge in F.

Lemma 8 *Let G contain no Type III boundary cells, let s be an even alternating strip in G, and let x and y be two vertices in G_F^- that correspond to cells not in s. If there exists a path p between x and y in G_F^- ($G_{F\oplus s}^-$) that does not include the vertex corresponding to the ending (beginning) cell of s, then there exists a path p' between x and y in $G_{F\oplus s}^-$ (G_F^-). Furthermore, if v is a vertex on p that corresponds to a cell in G that is not in s, then v is a vertex on p' in $G_{F\oplus s}^-$ (G_F^-).*

We need one more definition and lemma (proof omitted). The *total area* of an alternating strip sequence is the sum of the areas of its constituent strips.

Lemma 9 *If G contains no Type III boundary cells then a boundary B is uniquely determined by any boundary cell in B.*

Theorem 10 (Existence of Static Alternating Strip Sequences) *Let G contain no Type III boundary cells and let $A = (a_1, a_2, \ldots, a_k)$ be an alternating strip sequence that begins on boundary B. Then there exists a static alternating strip sequence A' that begins on B with total area no greater than that of A.*

Proof sketch: The proof is by induction on the total area of A. If A consists of a single (odd) alternating strip, the theorem follows trivially. Also, if there is a Type III boundary cell c on a boundary created by a_1, then the alternating strip sequence (a_1, c) is static, so we may assume to the contrary. By induction, then, there exists a static alternating strip sequence $S = (s_1, s_2, \ldots, s_n)$ in G with respect to $F \oplus a_1$, that begins on a boundary created by a_1 and whose total area is no greater than that of $(a_2, a_3, \ldots, a_k)$.

Consider the largest i for which s_i shares area or edges with a_1. If there is no such i, then $A' = (a_1, s_1, s_2, \ldots, s_n)$ is a static alternating strip sequence as required by the theorem. Otherwise, there are four cases to consider, depending on how s_i and a_1 overlap. (Note that it is a non-trivial fact, addressed at the end of this proof sketch, that these are the only four cases.) In the figures depicting the cases, we show only dark and light edges whose presence can be deduced from the degree constraint of the 2-factor and the fact that there are no Type III boundary cells on B or on a boundary created by a_1. For convenience, we let F' be the 2-factor after flipping a_1 and the strips of S up to s_i, so $F' = F \oplus a_1 \oplus s_1 \oplus s_2 \oplus \cdots \oplus s_{i-1}$. We now consider each of the four cases:

1. Strips a_1 and s_i share area. Figure 13(a) shows a_1 before it is flipped and (b) shows a_1 after it is flipped and just before s_i is to be flipped. It is not possible for s_i to cross a_1 perpendicularly; therefore it must be aligned with a_1 and in particular it must end at cell c in the figure. We now show that the sequence $A' = (s_{i+1}, s_{i+2}, \ldots, s_n)$ satisfies the theorem. The path between x and y in $G_{F'\oplus s_i}^-$ corresponds to the boundary created by s_i, on which s_{i+1} begins. By repeated application of Lemma 8, there must be an $x - y$ path in G_F^- that includes the vertex corresponding to the beginning cell of s_{i+1}. By Lemma 9, this is exactly boundary B in G with respect to F, so s_{i+1} begins on boundary B as required. By a similar application of Lemma 8, all strips s_j for $j > i+1$ begin on the boundary created after flipping the strips $s_{i+1}, s_{i+2} \ldots s_{j-1}$, so A' is a valid static alternating strip sequence.

2. Strips a_1 and s_i share a "side" edge. Figure 13(c) shows s_i just before it is to be flipped. Strip s_i must include cell c; call the portion that begins at c and proceeds to s_i's ending cell s_i'. By an application of Lemma 8 and Lemma 9 similar to that in case 1, $A' = (a_1, s_i', s_{i+1}, s_{i+2}, \ldots, s_n)$ is a static alternating strip sequence.

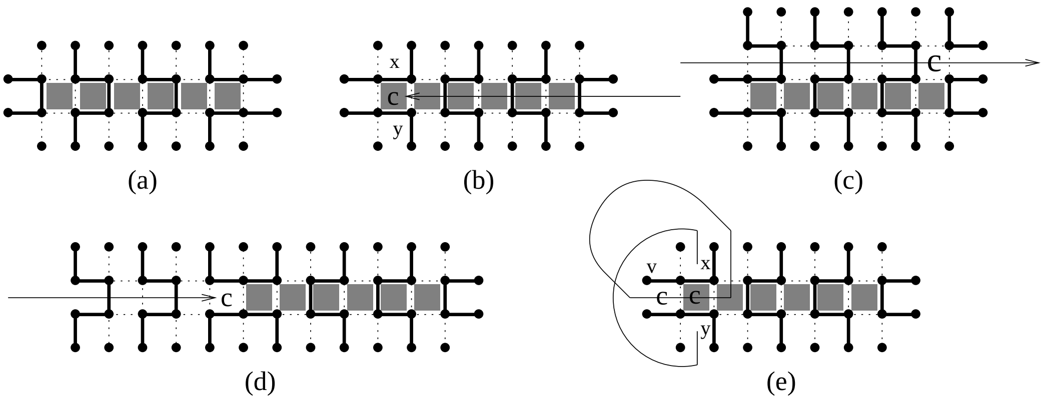

Figure 13: Strip a_1 before flipping (a); cases 1-3: a_1 and s_i just before s_i is to be flipped (b-d); case 4: a_1 just before s_i is to be flipped; s_i is one of the two cells marked with c (e).

3. Strips a_1 and s_i share the "top" edge of the beginning cell of a_1. Figure 13(d) shows this configuration. As in previous cases, we uses Lemma 8 and Lemma 9 to show that $A' = (s_{i+1}, s_{i+2}, \ldots, s_n)$ satisfies the theorem.

4. Strip s_i is a Type III boundary cell (so $i = n$) that shares the "top" edge of a_1 as in case 3, or overlaps with the beginning cell of a_1. In figure 13(e) s_i is one of the two cells marked with c. We argue that this configuration is not possible. The solid lines in the figure are a path and a cycle in $G^-_{F'}$. The path between x and y is obtained by repeated application of Lemma 8 to the path between x and y in G^-_F that corresponds to boundary B. The cycle passing through the "c"s corresponds to the boundary in $G_{F'}$ on which the strip s_i begins. But now vertex v cannot be on a cycle in the 2-factor, which contradicts F' being a 2-factor.

In other possible configurations for which s_i and a_1 share area or edges, the pattern of light and dark edges required by s_i and a_1 is inconsistent. We also use the fact that the sequence S is a *static* alternating strip sequence to argue that no strips s_j for $j < i$ can alter the configurations shown in figure 13 by more than possibly the single shared edge allowed in a static alternating strip sequence. This limits the possible configurations to the ones discussed. ■

5 The Algorithm

Given a procedure for finding static alternating strip sequences, the algorithm follows easily. First it finds an initial 2-factor (rejecting the non-Hamiltonian graph if one does not exist) and then it repeatedly locates and flips static alternating strip sequences until either the 2-factor has a single component, or no such sequence can be found.

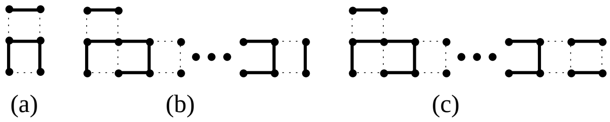

Figure 14: Odd chain alternating strip of length one (a) and general odd (b) and even (c) chain alternating strips.

We now describe in detail how to identify static alternating strip sequences. If there is a Type III boundary cell in G, then that is the static alternating strip sequence returned. Otherwise, we build an auxiliary graph H from the input graph G and a 2-factor F of G, in which vertices correspond to certain cells in G, and paths represent potential portions of alternating strip sequences. Graph H has the property that a particular shortest path must correspond to a static alternating strip sequence. It is built as follows:

1. Identify all alternating strips in G, and alternating strips with a dark edge in the single location that might be shared with the previous strip in a static sequence. We call these latter strips *chain alternating strips*; they are pictured in figure 14. An alternating cell must be adjacent to the beginning cell of a chain alternating strip as this is the only configuration in which they can appear in a static alternating strip sequence. For each strip identified in this step, we add a directed edge

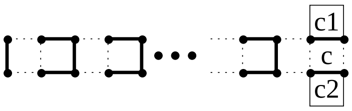

Figure 15: The cells adjacent to the ending cell of an even alternating strip.

(u, v) to H with vertex u labeled either **BEGIN** or **CHAIN** and vertex v labeled either **ODD** or **EVEN** by the obvious rules. For future reference, u is associated with the beginning cell of the strip and v is associated with the ending cell. The weight of edge (u, v) is the length of the strip.

2. Compute all-pairs-shortest-paths in G_F^-. For each **EVEN** vertex c added to H in the previous step, let c_1 and c_2 be the two vertices in G_F^- that correspond to the cells on either side of c as pictured in figure 15. Note that if c shares edges with the unbounded face, one or both of the adjacent cells may not be in the graph. In this case c_1 or/and c_2 should be the vertex in G_F^- corresponding to the unbounded face. Let p be the shortest path between c_1 and c_2. For every **BEGIN** vertex v in H associated with a cell whose vertex in G_F^- lies on p, we add the directed edge (c, v) to H with weight zero. If there exists a path between c_1 and c_2 in G_F^-, then for every **CHAIN** vertex v in H associated with a cell whose vertex in G_F^- is either c_1 or c_2, we add the directed edge (c, v) to H with weight zero.

3. Compute shortest paths in H. The shortest path between a **BEGIN** vertex in H that is associated with a boundary cell in G and an **ODD** vertex in H corresponds to a static alternating strip sequence. If no such path exists, then there is no alternating strip sequence in G.

We now sketch the proof of correctness. Step 1 clearly identifies all subgraphs that might be alternating strips in a static alternating strip sequence. The path p in step 2 corresponds exactly to the boundary that is created by flipping the even alternating strip s whose ending cell is c. To show that the sequence identified is static, we note that every path in H that corresponds to an alternating strip sequence has length equal to the total area of the sequence. If we suppose that the sequence identified by the algorithm is not static then by the techniques used in the proof of Theorem 10, we can show that a shorter path exists, which is a contradiction. Also, using Lemma 8, and techniques from the proof of Theorem 10, we can show that the boundaries that we associate with paths in G_F^- are not corrupted by flipping strips early in the sequence.

Finally, by Lemma 4 and Theorems 5 and 10, the main loop of the algorithm terminates with a single-component 2-factor if and only if G is Hamiltonian.

5.1 Analysis

The initial 2-factor can be found efficiently, and the procedure to identify static alternating strip sequences takes $O(n^3)$ time. The latter step is repeated at most n times. The purpose of this paper is to prove polynomiality; accordingly, this time bound is quite loose, and in the interest of clarity, the algorithm is not presented in its most efficient implementation.

We now have the main result:

Theorem 11 *The Hamiltonian cycle problem for quad-quad graphs (and solid grid graphs) is in $\mathcal{P}$.*

6 Conclusion

We have given a constructive polynomial-time algorithm for the Hamiltonian cycle problem in quad-quad graphs, a non-trivial superclass of solid grid graphs. It remains open whether a polynomial-time algorithm exists for solid grid graphs with *some* holes of restricted form. In studying such a problem, we believe the dependency graph of section 3.2 and its relation to alternating cycles may provide a starting point for determining which holes might be admissible. Another open problem stated in [Afr94] is the question of whether the Hamiltonian circuit problem for general solid grid graphs is $\mathcal{P}$-hard.

We would like to thank Christos Papadimitriou for several useful comments on an earlier draft of this paper.

References

[Afr94] Foto Afrati. The Hamilton circuit problem on grids. *Informatique Théorique et Applications*, 28(6):567–582, 1994.

[CZ95] Hwan-Gue Cho and Alexander Zelikovsky. Spanning closed trail and Hamiltonian cycle in grid graphs. In *Algorithms and Computations. 6th International Symposium, ISAAC 95*, pages 342–351, 1995.

[Eve86] Hazel Everett. Hamiltonian paths in nonrectangular grid graphs. Master's thesis, University of Saskatchewan, 1986.

[IPS82] Alon Itai, Christos H. Papadimitriou, and Jayme Luiz Szwarcfiter. Hamilton paths in grid graphs. *SIAM Journal of Computing*, 11(4):676–686, 1982.