

Floodlight Illumination of Infinite Wedges

Matthew Cary Atri Rudra Ashish Sabharwal Erik Vee

Computer Science and Engineering
University of Washington, Box 352350
Seattle, WA 98195-2350, U.S.A.

{cary, atri, ashish, env}@cs.washington.edu

TECHNICAL REPORT UW-CSE-2004-10-04

December 29, 2004

Abstract

The floodlight illumination problem asks whether there exists a one-to-one placement of n floodlights illuminating infinite wedges of angles $\alpha_1, \dots, \alpha_n$ at n locations p_1, \dots, p_n in a plane such that a given infinite wedge W of angle θ located at point q is completely illuminated by the floodlights. We prove that this problem is NP-hard, closing an open problem from CCCG 2001 [5]. In fact, we show that the problem is NP-complete even when $\alpha_i = \alpha$ for all $1 \leq i \leq n$ (the *uniform* case) and $\theta = \sum_{i=1}^n \alpha_i$ (the *tight* case). On the positive side, we define sufficient conditions on the locations of floodlights for which there are efficient algorithms to find an illumination. We discuss various approximate solutions and show that computing any *finite* approximation is NP-hard while ε -*angle* approximations can be obtained efficiently.

1 Introduction

Illumination problems generalize the well-known *art gallery* problem (see, e.g., [11, 12]). The task is to mount lights at various points so that a given region, typically a non-convex polygon, is completely illuminated. The mount points can be fixed in advance or not. The region may need to be illuminated from outside (like a soccer field) or from inside (like an indoor gallery). The lights may behave like ideal light bulbs, illuminating all directions equally, or like floodlights, illuminating a certain angle in a certain direction. The model of floodlights is the one used in this paper. It is quite natural and captures scenarios involving guards or security cameras with restricted angle of vision. Illumination algorithms using floodlights have focused in the past on illuminating the interior of orthogonal polygons [7, 1] and general polygons with restrictions on the floodlights used [2, 8, 6, 15]. There has also been work on *stage illumination* where one tries to illuminate lines rather than polygons [4].

The problem of illumination of *infinite wedges* by floodlights was introduced by Bose et al. [3]. As expected, the problem asks one to place and orient a given set of floodlights so that an infinite wedge W of some angle θ located at a point q is completely illuminated. Because of hardness of verification issues surrounding non-algebraic numbers, it is not clear whether the general problem is in the class NP. In fact, it is not obvious that it even has an exponential time solution. However, Steiger and Streinu [13] proved that it can indeed be solved in exponential time by formulating it as a bounded quantifier expression in Tarski's algebra [14] and using the result of Grigor'ev [10] on the complexity of deciding the truth value of such expressions. They also proved that the tight case, where the angles of floodlights sum exactly to the angle of the wedge to be illuminated, is the dual of a monotone matching problem with lines and slabs, up to a permutation of the floodlights. The exact complexity of this problem has been open [5].

In this paper we resolve this open problem by showing that tight floodlight illumination along with monotone matching is NP-complete, and the general problem NP-hard. After defining various floodlight illumination problems and the monotone matching problem and recapitulating the duality argument of Steiger and Streinu [13] in §2, we show

the NP-hardness of monotone matching in §3, implying the NP-hardness of the tight floodlight illumination problem. While we know of an NP-hardness reduction from the propositional satisfiability problem 3SAT to monotone matching, the one we give here is from an interesting graph problem, that of finding a *directed disjoint cycle cover*, which we also prove to be NP-complete. While this may be of independent interest, using this latter problem leads to a simplified and more natural proof for the hardness of monotone matching.

Although the general problem is NP-hard, several special cases can be solved efficiently. In §4 we outline conditions for an efficient solution in the tight case. We characterize and list many natural configuration classes of mount points for which a simple greedy algorithm works correctly. Finally, we discuss in §4.4 what types of approximations are possible. We show one natural type to be NP-hard and another to be easy.

2 Preliminaries

2.1 Floodlight Illumination

Definition 1. FLOODLIGHT ILLUMINATION Problem [3]:

Instance: n sites p_1, \dots, p_n , n angles $\alpha_1, \dots, \alpha_n$, and a generalized wedge W .

Question: Viewing the angles as spans of floodlights, is there an assignment of angles to sites along with angle orientations, that illuminates W ?

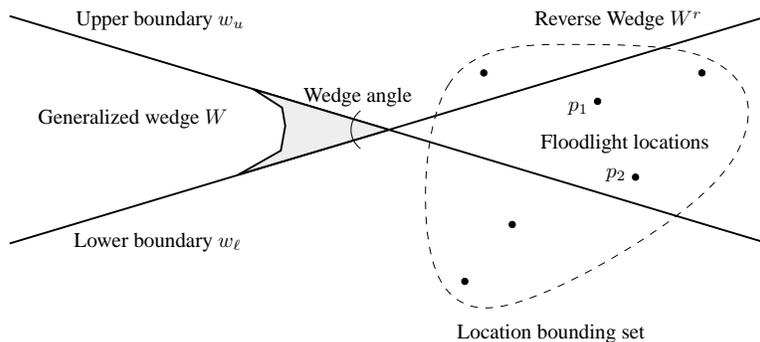


Figure 1: Basic definitions

Refer to Figure 1 for notation. We will assume the axis of the wedge lies on the x -axis, the wedge points to the left (*i.e.* negative x -axis), and unless otherwise mentioned, the floodlight angles are tight with respect to the wedge angle. We will consider the problem of illuminating a *generalized wedge*, which is simply a wedge with a continuous finite region adjacent to its apex removed. We will assume without loss of generality that all floodlights have non-zero angles of illumination.

There are several important special cases of the problem. In the *tight* illumination problem, the sum of floodlight spans $\sum \alpha_i$ equals the wedge angle. In this case, as described by Steiger and Streinu [13], a solution can be described by two permutations σ and τ . The first permutation σ orders the sites, so that we can talk about the first floodlight as the one placed at $p_{\sigma(1)}$, the second placed at $p_{\sigma(2)}$, etc. The second permutation τ assigns angles to the floodlights, so that the first floodlight has angle $\alpha_{\tau(1)}$, etc. The floodlight orientations are then inferred from the floodlight ordering. The first floodlight is oriented to illuminate the upper boundary w_u so that its upper ray is parallel to w_u and its lower ray is r_1 . The second floodlight is then oriented so its upper ray is parallel to r_1 , and so on. Because the sum of floodlight angles equals the wedge angle, any solution can be described in this fashion; see Steiger and Streinu [13] for further details.

The tight problem can be further specialized to the *restricted* problem, where τ is restricted to be the identity permutation. A different specialization is the *uniform* problem, where in addition to being tight, $\alpha_i = \alpha_j$ for all i, j . It is seen that the uniform problem is also restricted, as in this case all choices of permutation τ are equivalent. All three of these problems, the tight, restricted and uniform cases, are in NP as shown by Steiger and Streinu [13].

We show that the restricted and uniform versions have the same complexity as the monotone matching problem. We then show that monotone matching is NP-complete, implying that all versions of the floodlight illumination problem mentioned are NP-hard.

2.2 Monotone Matching

Suppose we are given n lines in the plane, $n + 1$ vertical lines defining n finite width vertical slabs, and two points, one on the leftmost vertical line and one on the rightmost. Call this an *arrangement of lines and slabs*, and denote it by (L, S, λ, ρ) , where $L \equiv \{(m_1, c_1), \dots, (m_n, c_n)\}$ is the set of lines $y = m_i x + c_i$, $S \equiv \{s_1, \dots, s_{n+1}\}$ is the set of vertical lines $x = s_i$ forming slabs, and λ and ρ are the two special points on the lines $x = s_1$ and $x = s_{n+1}$, respectively. A *monotone matching* in (L, S, λ, ρ) is a set of n line segments, each a portion of a unique line and spanning a unique slab, such that the following holds: (1) the left end point of the first segment is above λ , (2) the left endpoint of each subsequent segment is above the right endpoint of the segment in the previous slab, and (3) ρ is above the right endpoint of the last segment.

Definition 2. MONOTONE MATCHING Problem [13]:

Instance: An arrangement (L, S, λ, ρ) of n lines, $n + 1$ slabs, and two special points in the plane.

Question: Does this arrangement contain a monotone matching?

Analogous to the floodlight illumination case, we define the specialized *uniform* version UNIFORM MONOTONE MATCHING to be the problem where all slabs have the same width. We prove this variant to be NP-complete in §3.2.

2.3 Duality Between Floodlight Illumination and Monotone Matching

The restricted floodlight problem can be related to the monotone matching problem through duality as described by Steiger and Streinu [13]. The dual of a point p with coordinates (a, b) is the line ℓ_p with equation $y = ax + b$; the dual of a line ℓ with equation $y = mx + b$ is the point p_ℓ with coordinates $(-m, b)$. It is well known that this dual transformation preserves incidence and height ordering; i.e. if p intersects ℓ then their duals also intersect, and if p is above ℓ , then ℓ_p is above p_ℓ .

We now describe the dual of the floodlight illumination problem using the notation of Figure 1. w_l and w_u are dual to points λ and ρ ; as w_l has larger slope in the orientation of the figure, its dual λ has smaller x coordinate. The points in the vertical line containing λ are dual to lines that are parallel to w_l . The vertical strip between λ and ρ corresponds to the wedge angle. The line q between λ and ρ is dual to the intersection of w_l and w_u . The segment of q between λ and ρ corresponds to the lines with slope less than w_l and greater than w_u that have common intersection with w_l and w_u ; this is exactly the set of lines that form the wedge and reverse wedge of W .

Each site p_i corresponds to a line h_i which together make up the set of lines L . As we are in the restricted version of the problem, the angle of the first floodlight is α_1 . As described above, the tightness of the problem implies that the first floodlight must be oriented so that its upper ray is parallel to w_u . This corresponds in the dual to a vertical slab S_1 beginning at ρ and extending to the left a width proportional to α_1 (if S_1 extends from s_1 to s_2 , then $\alpha = \tan^{-1} x_2 - \tan^{-1} x_1$). The next floodlight then corresponds to a slab S_2 extending to the left of S_1 , continuing to the final floodlight which is a vertical slab S_n ending at λ .

A solution to the restricted problem is an assignment of sites to floodlights. In the dual this is a 1-1 assignment of lines h_i to slabs S_j . The illumination wedge of the first floodlight must overlap w_u , which corresponds to the right endpoint of the segment of the line h_i^1 assigned to s_1 being above ρ . Continuing, the right endpoint of the segment associated with S_2 must start above the left endpoint of the segment of S_1 , and so on, until the left endpoint of the segment at S_n is below λ .

If we flip the dual problem from left to right, in deference to those of us who read from left to right, we see we have reduced the restricted illumination problem to the monotone matching problem with lines.

Note that the unrestricted tight illumination problem corresponds to an extended matching problem where the widths of slabs are given and must be arranged in a partition of the slab between λ and ρ and then a matching found. The uniform illumination problem corresponds to the uniform matching problem, where the slabs are all of the same width, making, in particular, their arrangement immaterial.

2.4 Directed Disjoint Cycle Cover

As a tool for our main result, we prove NP-completeness for an interesting problem on directed graphs which naturally reduces to monotone matching. We also prove that this problem remains NP-complete even when the vertices of the graph are restricted to small degrees. This latter version is not critical but makes our NP-completeness proof for MONOTONE MATCHING cleaner.

Definition 3. DIRECTED DISJOINT CYCLE COVER Problem:

Instance: A directed graph $G = (V, E)$.

Question: Is there a directed disjoint cycle cover for G , i.e., a set $\mathcal{C} = \{C_1, \dots, C_k\}$ of vertex disjoint directed cycles in G such that every vertex $v \in V$ is in some cycle $C_i \in \mathcal{C}$?

Theorem 4. DIRECTED DISJOINT CYCLE COVER is NP-complete, even for graphs with indegree and outdegree each bounded above by 3, as well as for graphs with outdegree exactly 3 and indegree at most 4.

The theorem is proved by a reduction from VERTEX COVER in a manner very similar to the reduction to HAMILTONIAN CYCLE given by Garey and Johnson [9]. Our argument for the hardness of the bounded degree cases, however, is much simpler than those known for other similar graph problems. Proof details may be found in the Appendix. We use the second of the two degree restrictions in our reduction for monotone matching. For simplicity, we allow directed graphs to have self loops, though our proofs can be modified to work even for graphs with no self loops. A self loop is assumed to contribute one to both the indegree and the outdegree of the corresponding vertex.

3 NP-hardness of Monotone Matching and Floodlight Illumination

3.1 Monotone Matching with Pseudolines

We begin by defining a variant of the monotone matching problem that uses *pseudolines* instead of lines. A pseudoline is a curve in \mathbb{R}^2 that intersects any vertical line in exactly one point. A collection of pseudolines is a set of pseudolines no two of which intersect more than once. For computational purposes, we shall assume that the point of intersection of two pseudolines can be computed efficiently from their input representation. Arrangements of pseudolines and points, and monotone matchings with pseudolines are defined analogous to the case of straight lines. The pseudolines used in our constructions will all be piecewise linear functions.

Definition 5. PSEUDOLINE MONOTONE MATCHING Problem:

Instance: An arrangement (PL, S, λ, ρ) of n pseudolines, $n + 1$ slabs, and two special points in the plane.

Question: Does this arrangement contain a monotone matching?

Theorem 6. PSEUDOLINE MONOTONE MATCHING is NP-complete.

We begin by defining some configurations, or gadgets, of pseudolines that will be useful for our proof. Our construction for the straight line matching problem will follow by reconstructing these gadgets from straight lines.

The most important gadget is the *forcing gadget*, shown in Figure 2. This is a sequence of slabs associated with pseudolines that forces the line used previous to the gadget to end below a chosen point, and the line used after the gadget to start above another chosen point. Let R denote the *region of interest* for the monotone matching problem, i.e., a rectangular region covering horizontally all the slabs and covering vertically all intersections of pseudolines in PL with these slabs. The *sky* region begins above R and the *ground* region begins below R .

Lemma 7 (Forcing Gadget Property). *Refer to Figure 2 for the labeling of forcing gadgets. Given a monotone matching instance, if for all pairs F and F' of forcing gadgets composed of h, ℓ, s and h', ℓ', s' , respectively, if h and h' do not intersect, and point s of F is below h' if F' follows F , then in any valid matching, h is used in slab A and ℓ is used in slab B . This holds for any number of additional lines, as long as they appear above the ground and below the sky.*

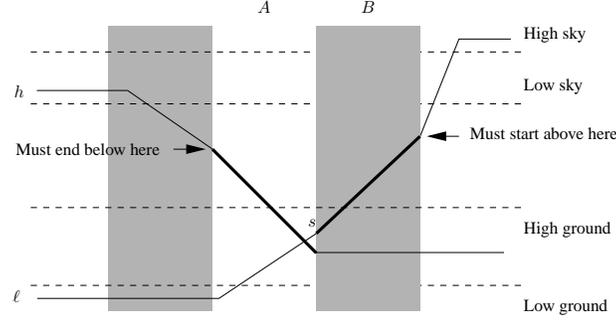


Figure 2: The forcing gadget. The arrows show how any lines used before or after the gadget are constrained.

Proof. ℓ cannot be used before slab B , as it begins in the ground, no other lines intersect ℓ before B , and the ground is below the starting point. If ℓ is used after slab B , then the line used subsequently must begin in the high sky. As the only lines in the high sky are from previous forcing gadgets, and there are no intersections in the high sky, then all future lines chosen will be in the high sky. As the high sky is above ρ , any such matching must end above ρ , and so will not be a valid matching.

Now, given that ℓ must be used in slab B , as its starting point s is in high ground, the only precursors it could have is h , some h' from a forcing gadget before F , or some line from the low ground. As argued in the beginning of this proof, no lines in the low ground can be used. We prove that the precursor could not come from an h' preceding F by induction on the number of forcing gadgets. If F is the first forcing gadget, then h must be used; if there are forcing gadgets before F , then we can inductively assume that all the corresponding h' were used in their corresponding gadgets, so that h must be used in slab A .

Finally, as our arguments were based on the presence or absence of pseudolines in the ground and sky, any number of additional lines appearing between the ground and sky do not affect the behavior of the forcing gadget. \square

Proof of Theorem 6. As a potential matching can easily be verified in polynomial-time, this problem is in NP.

The proof of NP-hardness is by a reduction from DIRECTED DISJOINT CYCLE COVER, which, by Theorem 4, is NP-complete. Suppose we are given a directed graph G . We will construct (PL, S, λ, ρ) such that G has a disjoint cycle cover iff $(PL, S, \lambda, \rho) \in$ PSEUDOLINE MONOTONE MATCHING. By Theorem 4, we may assume the outdegree of all vertices in G is exactly 3, and the indegree is at most 4. W.l.o.g. we will assume that every vertex in G has indegree at least one, for if not then G does not have a disjoint cycle cover and can be mapped to a trivial instance of monotone matching with no solution.

To do this we need two graph-related gadgets. Given a directed graph $G = (V, E)$, we will have gadgets $\mathbf{In}(v)$ and $\mathbf{Out}(u)$ for $u, v \in V$ as shown in Figure 3. Let $\mathcal{I}(v) \subset E$ be in the in-edges of v , and let $\mathcal{O}(u) \subset E$ be the out-edges of u . By our choice of G , $|\mathcal{I}(v)| \leq 4$ and $|\mathcal{O}(u)| = 3$. The gadget $\mathbf{In}(v)$ will allow us to select exactly one predecessor of v and exactly one successor of u , leading to a directed disjoint cycle cover as long as the predecessors and successors are consistent. Selecting a successor will actually be done by selecting and ruling out exactly $|\mathcal{O}(u)| - 1$ out-edges of u .



Figure 3: Graph gadgets $\mathbf{In}(v)$ and $\mathbf{Out}(u)$. The arrows denote locations set by forcing gadgets placed between the graph gadgets. $|\mathcal{I}(v)| \leq 4$ and $|\mathcal{O}(u)| = 3$.

Let $n = |V|$ and $m = |E|$. We will use m primary pseudolines, each corresponding to an edge in E . We will

abuse notation and talk of an edge $e \in E$ as being used in a particular slab; this will mean that the primary pseudoline corresponding to e is used at that slab. There will be a number of auxiliary pseudolines used in forcing gadgets. The $\mathbf{Out}(\cdot)$ gadgets and $\mathbf{In}(\cdot)$ gadgets will be arranged in sequence as shown in Figure 4. The primary pseudoline corresponding to edge (u, v) will first pass through $\mathbf{Out}(u)$, and then pass through $\mathbf{In}(v)$. Each pseudoline will intersect another pseudoline at most once, depending on the relationship between their $\mathbf{In}(\cdot)$ and $\mathbf{Out}(\cdot)$ gadgets.

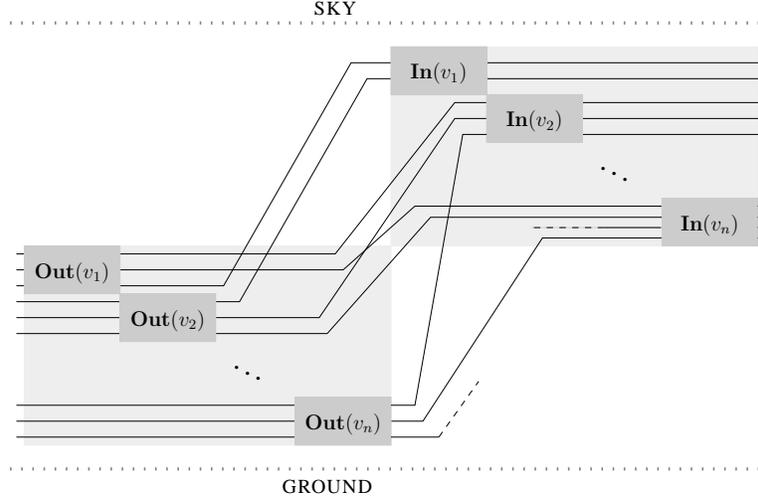


Figure 4: Overall view of the reduction from DIRECTED DISJOINT CYCLE COVER. Forcing gadgets appear between $\mathbf{In}(\cdot)$ and $\mathbf{Out}(\cdot)$ gadgets, to match the arrows in Figure 3, except for $\mathbf{Out}(v_1)$ and $\mathbf{In}(v_n)$, where the first and last arrows, respectively, determine the starting and ending points λ and ρ for the monotone matching instance.

We claim that when arranged as in Figure 4 along with appropriate forcing gadgets for each $\mathbf{In}(\cdot)$ and $\mathbf{Out}(\cdot)$ gadget, exactly one $e \in \mathcal{I}(v_i)$ is used in $\mathbf{In}(v_i)$ and exactly one $e \in \mathcal{O}(v_i)$ is not used in $\mathbf{Out}(v_i)$, for $1 \leq i \leq n$. To see this, refer to Figure 3. Because of the forcing gadgets, exactly one pseudoline must be used in $\mathbf{In}(v_i)$ and must correspond to an edge $e \in \mathcal{I}(v_i)$. Similarly, exactly two pseudolines must be used in $\mathbf{Out}(v_i)$, leaving out exactly one pseudoline corresponding to an edge $e \in \mathcal{O}(v_i)$. Furthermore, the edges selected in this way are consistent, i.e., if edge (u, v) is selected in $\mathbf{In}(v)$, then it is also the only edge left unselected in $\mathbf{Out}(u)$.

A directed disjoint cycle cover of G is equivalent to a permutation π on the vertices, where $\pi(v)$ is the predecessor of v in the cycle containing v . If such a permutation exists, then a monotone matching exists, by not selecting the edge at $\mathbf{Out}(u)$ corresponding to $\pi^{-1}(u)$, and selecting the edge corresponding to $\pi(v)$ at $\mathbf{In}(v)$. Conversely, if a monotone matching exists, then the permutation π can be recovered by setting $\pi(v)$ equal to the edge that is used in $\mathbf{In}(v)$. This completes the reduction. \square

3.2 Uniform Monotone Matching

We now prove NP-completeness of the uniform case where all slabs are of the same width. As stated earlier, NP-hardness of this problem subsumes NP-hardness of PSEUDOLINE MONOTONE MATCHING proved in the previous section. The argument, however, is more involved and reuses the key concepts developed in the earlier proof.

Theorem 8. UNIFORM MONOTONE MATCHING is NP-complete.

Proof. As with Theorem 6, the problem is clearly in NP.

We prove NP-hardness by a reduction from DIRECTED DISJOINT CYCLE COVER with graphs of outdegree exactly 3 and indegree at most 4. By Theorem 4, this latter problem is NP-complete. Suppose we are given a graph G with degrees bounded as stated above. We will construct (L, S, λ, ρ) such that G has a directed disjoint cycle cover iff $(L, S, \lambda, \rho) \in \text{UNIFORM MONOTONE MATCHING}$.

The reduction from DIRECTED DISJOINT CYCLE COVER is similar to the one given in the proof of Theorem 6 except for the construction of the new forcing gadgets and addition of a “buffer” zone between the **In** and **Out** gadgets to convert pseudolines into straight lines. The basic idea of the buffer zone is to make the lines behave as parallel (*i.e.* non-intersecting) lines in the region of interest as required in the graph gadgets of Figure 3. We first spell out the details of construction of the lines corresponding to the edges of G and then specify how to construct the forcing gadgets with lines. For the rest of the proof, w.l.o.g. we will assume that the slab width $w = 1$.

The graph gadgets are arranged as shown in Figure 5 with a buffer area consisting of $b - 4n$ slabs (or $\frac{b-4n}{2}$ forcing gadgets). We will use a coordinate system as shown in Figure 5 for the ease of exposition. The graph gadget **In**(v_j) has its forced start point at $(b + 3j - 3, n - j)$ and forced end point at $(b + 3j, n - j + 1)$ while the graph gadget **Out**(v_i) has its forced start and end points at $(4i - 4, -i)$ and $(4i, 1 - i)$, respectively. The line corresponding to edge (v_i, v_j) connects the points $(-1, -i)$ and $(b + 3j - 3, n - j)$.

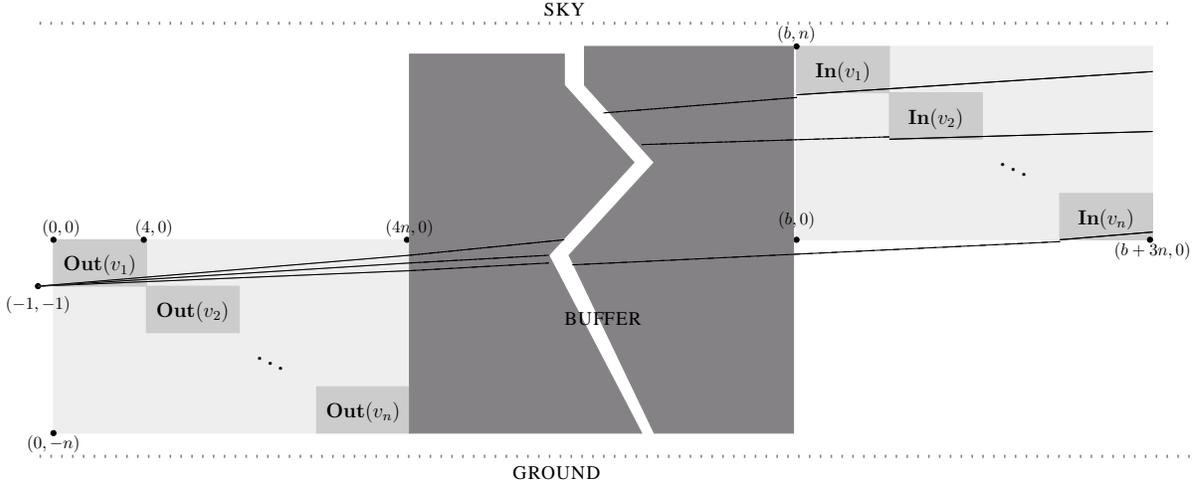


Figure 5: Overall view of the reduction from DIRECTED DISJOINT CYCLE COVER for the uniform case. Forcing gadgets appear in the buffer area and between **In**(\cdot) and **Out**(\cdot) gadgets, to match the arrows in Figure 3, except for **Out**(v_1) and **In**(v_n), where the first and last arrows, respectively, determine the starting and ending points λ and ρ for the uniform monotone matching instance. b is $10n^2$. For clarity, only the lines corresponding to outgoing edges from v_1 are shown: the edges are (v_1, v_1) , (v_1, v_2) and (v_1, v_n) .

Fix $b = 10n^2$. We now show that the lines corresponding to the outgoing edges from any graph gadget do not intersect each other in the horizontal intervals $[0, 4n]$ and $[b, b + 3n]$, that is, they are effectively parallel as needed in Figure 3. For any edge (v_i, v_j) the slope of the corresponding line is bounded from above by $m = \frac{n - (-n)}{b} = \frac{2n}{b}$. The maximum y coordinate for the line corresponding to (v_i, v_j) in the horizontal interval $[0, 4n]$ is given by $y_m = -i + m \cdot 4n$. Substituting the value of m and b , we have $y_m < -i + 1$. Thus, no two lines intersect in the interval $[0, 4n]$. One can similarly show that no two lines intersect in the interval $[b, b + 3n]$.

To complete the proof we specify the forcing gadgets. We need $\frac{b}{2}$ forcing gadgets (n each from the **In** and **Out** gadgets and $\frac{b-4n}{2}$ from the buffer area). The scheme is presented in Figure 6. For the k^{th} forcing gadget, the h line connects p_0 , which is at height $n^4 - 2kn$ and p_1 , the forced end point at the boundary of slab A . The ℓ line of the gadget connects q_0 (which is one higher than p_2 , the intersection of h with the boundary of slab B) and q_1 , the forced start point at the boundary of slab B .

We claim that the height of q_2 is (strictly) bounded from above and below by $n^4 - 2kn$ and $n^4 - 2(k + 1)n$, respectively. By our overall construction in Figure 5, p_1 is always higher than q_1 and their difference is bounded from above by n . The bounds on the height of q_2 follow from recalling that each slab has width one.

We now argue that in a valid uniform monotone matching, h is chosen in slab A and line ℓ is chosen in slab B as was proved in Lemma 7. We first show that ℓ is selected in slab B . For the sake of contradiction assume that ℓ was

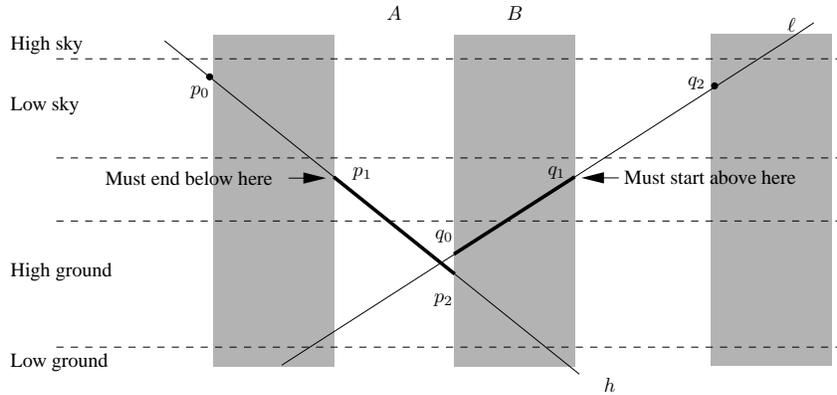


Figure 6: The forcing gadget with straight lines. The arrows show how any lines used before or after the gadget are constrained. q_0 is higher than p_2 by 1. For the k^{th} gadget, p_0 is at height $n^4 - 2kn$.

chosen after slab B . By the lower bound on the height of q_2 , we cannot come down from the sky as all the h' lines of the subsequent gadgets are below q_2 which is a contradiction as the ending point is below the sky. If ℓ were chosen before slab B , its precursor cannot include any ℓ' line from a previous forcing gadget as then the starting point λ would have to be below ground, which is not the case. It also cannot have any h' line from a previous forcing gadget as a precursor as all those lines are higher than ℓ by construction. Thus, ℓ is selected in slab B . From the above argument the only precursors of ℓ could be h or some h' from some previous forcing gadget which, as before by induction, is impossible. Thus, h is selected in slab A and l in slab B , fulfilling the requirements of the forcing gadget. \square

Corollary 9. MONOTONE MATCHING is NP-complete.

Proof. Membership in NP follows from an argument identical to the one for UNIFORM MONOTONE MATCHING. NP-hardness follows from Theorem 8 because UNIFORM MONOTONE MATCHING trivially reduces to this problem. \square

3.3 Floodlight Illumination

Combining the discussion in §2.3 with Theorem 8, we immediately have our main result:

Theorem 10 (Main Hardness Result). FLOODLIGHT ILLUMINATION is NP-hard. The tight, restricted, and uniform versions of the problem are NP-complete.

4 Illumination Algorithms for the Tight Case

In this section we look at algorithms to solve the floodlight illumination problem in the tight case. We characterize several special cases that can be solved in polynomial time. We also give approximation algorithms for the problem. For this section, we will use the notations and definitions from §2.1. We begin with some properties of floodlight illuminations which will be used later in this section but may also be of independent interest.

4.1 Properties of Floodlight Illumination

The first lemma shows how the position of certain floodlights are fixed by the problem instance. Next we prove a necessary condition for the existence of a solution. We then state a lemma which will be used in proving hardness of certain kind of approximations to the floodlight illumination problem. We end with a brief mention of a variant of the problem that is easy to solve.

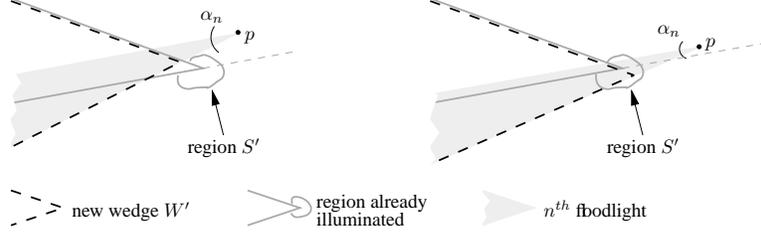


Figure 7: Two cases for the shape of overall illuminated region in the tight case

Lemma 11. *In the tight case, any floodlight illuminating at infinity the upper boundary w_u of a generalized wedge W must be parallel to and located above w_u .*

Proof. Let the floodlight f illuminating w_u be mounted at location p_u . As we are considering the tight case, the upper boundary f_u of the region illuminated by f must be parallel to w_u (see discussion in §2.1. If p_u is below w_u , f will not illuminate an infinite slice S of W including the boundary w_u . \square

Lemma 12. *A tight floodlight illumination problem on generalized wedges has a solution only if there is at least one location in the reverse wedge.*

Proof. Suppose there are no floodlight locations in the reverse wedge. Referring to §2.3, consider the dual of the problem where we work from right to left assigning locations. The location lines in the dual fall into three categories: those above ρ and λ (the duals of w_u and w_l), corresponding to locations above the wedges; those below the duals ρ and λ , corresponding to locations below the wedge; and those that run between ρ and λ , corresponding to locations in the forward wedge. As we are in the tight case, if there are points in the forward wedge, we cannot possibly have a solution. Hence all points are either above or below both of the dual starting and end points. In particular, no line above the dual points intersects a line below the dual points in the problem slab defined by the dual points. This means that we cannot use any of the lines below the starting point and hence, we cannot have a matching as there are not enough lines which can be used. \square

Lemma 13. *Suppose the sum of the angles of n floodlights is α . If they illuminate a wedge W of angle α , then the overall illuminated region is of the form $W' \cup S$, where $W' \supseteq W$ is a wedge of angle α aligned with W and S is a finite region.*

Proof. We prove by induction on n the weaker statement that ignores the requirement $W' \supseteq W$ above. This will, however, suffice for the proof because if W' did not include all of W , $W' \cup S$ would also not include all of W since S is finite. This will contradict the precondition.

For the base case of $n = 1$, $W' = W$ and $S = \emptyset$. For $n > 1$, note that the floodlight angles are tight relative to W . Hence, the first $n - 1$ floodlights must together cover some wedge of angle $\alpha - \alpha_n$, where α_n is the angle of the n^{th} floodlight (see Figure 7). By induction, the region illuminated by the first $n - 1$ floodlights is of the form $W'' \cup S'$, where W'' is a wedge of angle $\alpha - \alpha_n$ and S' is finite. Again, since the angles are tight, the only way to extend this region to cover W of angle α is to mount the n^{th} floodlight f at a location p that is above the lower boundary w_l of the already illuminated region and have its upper boundary f_u aligned with w_l . Let W' be the wedge of angle α defined by the upper boundary w_u of the already illuminated region and the lower boundary f_l of f . As seen from the two cases in Figure 7, the overall illuminated region is $W' \cup S$, where S is finite. \square

Finally, we mention a relaxation of the problem which makes it easy. Two **movable locations** can always solve any tight problem instance: assign two arbitrarily chosen first and last floodlights (the ones parallel to the wedge boundaries) to the movable locations, and move these locations back and inside the reverse wedge far enough so that every other location is within the reverse of the residual wedge. Now use Fact 14.

4.2 A Greedy Algorithm

We briefly describe a greedy algorithm given by Steiger and Streinu [13] for the floodlight illumination problem which takes an additional input: the order in which the floodlight angles are chosen, that is, permutation σ from §2.1. Note that for the uniform case, where each floodlight angle is the same, the permutation σ does not come into play. At each step the greedy algorithm assigns the current floodlight angle in σ to the position p which would leave the maximum number of position in the reverse wedge of the residual wedge obtained by placing the current floodlight angle on p . There is also a natural interpretation of the algorithm in the dual monotone matching problem, where one chooses a line for a slab that maximizes the number of choices for the next slab. All references to a greedy algorithm in the subsequent sections are to this algorithm. We complete the description of the greedy algorithm by stating a simple property of the algorithm:

Fact 14. *If all floodlight locations are contained inside the reverse wedge W^r , then the greedy illumination algorithm successfully illuminates W after any assignment of floodlights to positions. Equivalently, if the first special point in a monotone matching problem is below all lines and the second is above all, the greedy segment matching algorithm successfully finds a matching.*

4.3 Special Floodlight Location Configurations

In this section we consider special illumination problems where the floodlight locations are restricted to obey certain properties. This allows us to characterize cases where the greedy segment matching algorithm produces the right answer for the corresponding tight floodlight illumination problem.

Definition 15. Locations p_1, p_2, \dots, p_n are *angle-separated with respect to wedge W* of angle α if $p_i \notin W_j$ for every $1 \leq i \neq j \leq n$, where W_j is the wedge of angle α located at p_j and aligned with W (see Figure 8). Locations in a floodlight illumination problem are *angle-separated* if they are angle-separated with respect to the wedge to be illuminated.

We note that an equivalent way of defining angle-separation with respect to a wedge W with boundary slopes m_u and m_l is to require that for all location pairs (p, p') , the line joining p and p' has slope not in $[m_l, m_u]$. While the former definition is natural for the proof of the following lemma, this latter definition might be more convenient for algorithmic implementations.

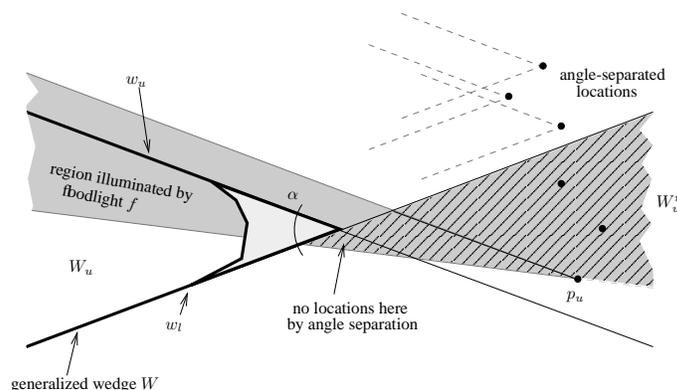


Figure 8: Angle separation implies all location must be in W^r for a solution to exist.

Lemma 16. *If the locations in a tight floodlight illumination problem on generalized wedges are angle-separated, there is no solution unless all locations are contained in the reverse wedge.*

Proof. Consider the dual (refer to §2.3). The wedge W_j for point p_j corresponds in the dual to the segment of the line ℓ_{p_j} between the starting and ending points, that is, the problem slab. If there is some point which is not in the reverse wedge then there exists a line above the end point or a line below the starting point. This implies that one cannot use at least one line in the matching and thus, a matching cannot exist. \square

Proposition 17 (Sufficient Condition for Greedy). *If the locations in a tight floodlight illumination problem on generalized wedges are angle-separated, the greedy segment matching algorithm always produces the right answer.*

Proof. If all locations are contained inside the reverse wedge, by Fact 14, a solution is always found by the greedy segment matching algorithm for any assignment of floodlights to locations. On the other hand, if at least one location is outside the reverse wedge, by Fact 16, there is no solution and the greedy algorithm, of course, doesn't find one. \square

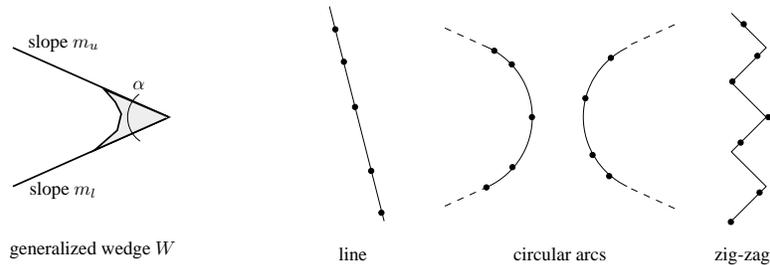


Figure 9: A few natural floodlight location configurations for which the problem is easy.

It follows that the floodlight illumination problem for generalized wedge W of angle α and with boundaries of slopes m_u and m_l is easy to solve when, for instance, all locations are on a straight line whose slope is not in $[m_l, m_u]$, or on a circular arc whose endpoint tangents have slope not in $[m_l, m_u]$, or in a vertical zig-zag pattern with angle greater than α , etc. (see Figure 9).

4.4 Approximate Illuminations

Theorem 10 motivates the study of approximation algorithms for the floodlight illumination problem. There are several natural ways to approximate wedge illumination. After giving precise definitions of some of these, we prove a negative result that illuminating all but a finite portion of a wedge is, in the tight case, not any easier than illuminating the whole wedge. This motivates other reasonable notions of approximation that leave unlit relatively small but infinite regions of the wedge.

Definition 18. Let \mathcal{F} be an illumination of a wedge W .

1. \mathcal{F} is a *finite-approximation* if it illuminates $W \setminus S$, where S is a finite region.
2. \mathcal{F} is an ε *angle-approximation* if it illuminates $W \setminus S_\varepsilon$, where S_ε is a union of wedges whose total angle is at most ε .

Lemma 19. *There is a solution to a tight floodlight illumination problem on a wedge W if and only if there is no finite-approximation to it.*

Proof. We prove the sufficient condition. Suppose there is a finite-approximate illumination \mathcal{F} for W . Let W be of angle α . By definition, \mathcal{F} must illuminate a wedge W^* of angle α that is aligned with W but is possibly contained strictly within W . Since the floodlight angles are tight relative to W^* , by Lemma 13, the overall region illuminated by \mathcal{F} is of the form $W' \cup S$, where W' is a wedge of angle α aligned with W^* (and hence with W) and S is finite. If $W \not\subseteq W'$, then $W' \setminus W$ is an infinite region R . As S is finite, $W' \cup S$ will not cover an infinite portion of this infinite region R of W , contradicting the fact that \mathcal{F} illuminates all but a finite region of W . It follows that $W \subseteq W'$, implying that W is completely illuminated by \mathcal{F} and providing an exact solution. The other direction of the proof is trivial. \square

This Lemma implies that computing a finite-approximation is NP-hard because computing the exact solution is. It also implies that there is a solution to the tight floodlight problem on a generalized wedge W iff there is a solution to the tight floodlight problem on the underlying normal wedge W' . In this sense, generalized wedges don't make the problem any harder. However, they provide a convenient tool for analysis, allowing, for instance, stronger inductive claims.

Note that ε angle-approximate illumination means all but ε of the wedge is illuminated "at infinity". It would be interesting to find an algorithm for the stronger approximation where the resulting illuminated area is a smaller wedge but located at the same apex as W .

Lemma 20. *For any $\varepsilon > 0$, an ε angle-approximation to the tight floodlight problem can be found efficiently.*

Proof. An ε angle-approximation can be achieved by adding two movable locations p_a and p_b to the floodlight locations, adding two floodlights f_a and f_b of angle $\varepsilon/2$ each, reducing any one original floodlight angle by ε , and proceeding as follows. Mount floodlight f_a at location p_a , orient it so that its upper boundary is parallel to and illuminates the upper boundary w_u of W , and move it far and low enough in W^r so that all other locations are above its lower boundary. Perform a similar operation on f_b and p_b starting with the lower boundary w_l of W . The region W' of W not illuminated by these two floodlights is a generalized wedge of angle $\alpha - \varepsilon$, where α is the wedge angle of W . Further, all remaining locations live within W^r . By Lemma 14, we can illuminate W^r exactly using the remaining floodlights. Now remove f_a and f_b , and add angle ε back to the floodlight whose angle was reduced. This completes the ε angle-approximation. \square

Theorem 21 (Approximations). *For the tight floodlight illumination problem, computing a finite-approximation is NP-hard, whereas for any ε an ε angle-approximation can be found in polynomial time.*

Acknowledgments

The authors would like to thank Neva Cherniavsky for introducing the problem to us, and Gidon Shavit and William Pentney for helpful discussions.

References

- [1] J. Abello, V. Estivill-Castro, T. C. Shermer, and J. Urrutia. Illumination of orthogonal polygons with orthogonal floodlights. *International Journal of Computational Geometry and Applications*, 8(1):25–38, 1998.
- [2] J. Bagga, L. Gewali, and D. Glasser. The complexity of illumination polygons by α -flood-lights. In *Proceedings of the 8th Canadian Conference on Computational Geometry*, pages 337–342, Ottawa, Canada, Aug. 1996.
- [3] P. Bose, L. Guibas, A. Lubiw, M. Overmars, D. Souvaine, and J. Urrutia. The floodlight problem. *International Journal of Computational Geometry and Applications*, 7(1/2):153–163, 1997.
- [4] J. Czyzowicz, E. Rivera-Campo, and J. Urrutia. Optimal floodlight illumination of stages. In *Proceedings of the 5th Canadian Conference on Computational Geometry*, pages 393–398, Waterloo, Canada, Aug. 1993.
- [5] E. Demaine and J. O'Rourke. Open problems from CCCG 2001. In *Proceedings of the 13th Canadian Conference on Computational Geometry*, Waterloo, Canada, Aug. 2001. <http://www.cs.uleth.ca/~wismath/cccg/papers/open.pdf>.
- [6] V. Estivill-Castro, J. O'Rourke, J. Urrutia, and D. Xu. Illumination of polygons with vertex lights. *Information Processing Letters*, 56(1):9–13, 1995.
- [7] V. Estivill-Castro and J. Urrutia. Optimal floodlight illumination of orthogonal art galleries. In *Proceedings of the 6th Canadian Conference on Computational Geometry*, pages 81–86, Saskatchewan, Canada, Aug. 1994.
- [8] V. Estivill-Castro and J. Urrutia. Two-floodlight illumination of convex polygons. In *Proc. 4th Intl. Workshop on Algorithms and Data Structures*, volume 955 of LNCS, pages 62–73, Kingston, Canada, Aug. 1995.

- [9] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
- [10] D. Y. Grigor'ev. Complexity of deciding Tarski algebra. *Journal of Symbolic Computation*, 5:65–108, 1998.
- [11] J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.
- [12] T. C. Shermer. Recent results in art galleries. *Proceedings of the IEEE*, 80(9):1384–1399, 1992.
- [13] W. L. Steiger and I. Streinu. Illumination by floodlights. *Computational Geometry*, 10:57–70, 1998.
- [14] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, 1948.
- [15] C. D. Tóth. Illuminating polygons with vertex π -floodlights. In *Proceedings, Part I, of the International Conference on Computational Science (ICCS)*, volume 2073 of LNCS, pages 772–784, San Francisco, May 2001.

Appendix

Proof of Theorem 4. We begin by giving a proof of the following Lemma that proves NP-completeness of the problem without degree restrictions, and then use it to prove the Theorem:

Lemma 22. DIRECTED DISJOINT CYCLE COVER is NP-complete.

Proof. The problem is clearly in NP because given a description of a set $\{C_1, \dots, C_k\}$ as a certificate of membership, one can easily verify in polynomial time that each $C_i, 1 \leq i \leq k$, is a directed cycle in the given graph G , the C_i 's are vertex disjoint, and they together cover all vertices of G .

Let $\text{VERTEXCOVER} = \{\langle G, k \rangle \mid G \text{ is a directed graph with a vertex cover of size } k\}$. VERTEXCOVER is NP-complete [9]. Our proof of NP-hardness is by a reduction from VERTEXCOVER . The reduction is very similar to the proof of NP-hardness of the $\text{HAMILTONIAN CIRCUIT}$ problem given by Garey and Johnson [9]. Our description of the reduction will closely follow the one they present. Let $\langle G, k \rangle$ be an instance of VERTEXCOVER . We will construct a directed graph G' such that G has a vertex cover of size k iff G' has a directed disjoint cycle cover. In fact, G' will be such that it has a directed disjoint cycle cover iff it has a directed Hamiltonian circuit. Hence, our proof gives an alternate argument for the NP-hardness of $\text{DIRECTED HAMILTONIAN CIRCUIT}$ [9].

The construction can be viewed in terms of component gadgets connected together by communication links. First, the graph G' has k “selector” vertices a_1, a_2, \dots, a_k , which will be used to select k vertices from the vertex set V for G . Second, for each edge in $e \in E$, G' contains a “cover-testing” component that will be used to ensure that at least one endpoint of e is among the selected k vertices. The component for $e = \{u, v\}$ is illustrated in Figure 10. It has 4 vertices, $V'_e = \{(u, e, \text{top}), (u, e, \text{bot}), (v, e, \text{top}), (v, e, \text{bot})\}$, and 6 directed edges, E'_e , shown in the figure.

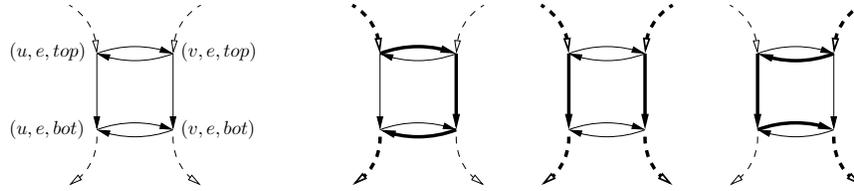


Figure 10: Cover-testing component for edge $e = \{u, v\}$, along with the three possible configurations of a directed disjoint cycle cover on this component

It is easily seen that any disjoint cycle cover for G' must involve edges of E'_e in one of the three configurations shown in Figure 10. This is because any other choice of edges will leave at least one of the gadget vertices impossible to cover. Thus, for example, if the circuit “enters” this component at (u, e, top) , it must “exit” at (u, e, bot) and visit either all four vertices in V'_e or just the two.

Additional edges in our overall construction will serve to join pairs of cover-testing components or to join a cover-testing component to a selector vertex. For each vertex $v \in V$, let the edges incident on v be ordered (arbitrarily) as $e_{v[1]}, e_{v[2]}, \dots, e_{v[\text{deg}(v)]}$, where $\text{deg}(v)$ denotes the degree of v in G . All the cover-testing components corresponding to these edges (having v as endpoint) are joined together by the following connecting edges: $E'_v = \{((v, e_{v[i]}, \text{bot}), (v, e_{v[i+1]}, \text{top})) : 1 \leq i < \text{deg}(v)\}$. As shown in Figure 11, this creates a single path in G' that includes exactly the vertices (x, y, z) having $x = v$.

The final connecting edges in G' join the first and last vertices from each of these paths to every one of the selector vertices, a_1, a_2, \dots, a_k . These edges are: $E'' = \{(a_i, (v, e_{v[1]}, \text{top})), ((v, e_{v[\text{deg}(v)]}, \text{bot}), a_i) : 1 \leq i \leq k, v \in V\}$. The final graph G' has $V = \{a_i : 1 \leq i \leq k\} \cup (\bigcup_{e \in E} V'_e)$ and $E' = (\bigcup_{e \in E} E'_e) \cup (\bigcup_{v \in V} E'_v) \cup E''$. G' can clearly be constructed from G and k in polynomial time.

We now argue the correctness of the construction. Suppose first that G' has a directed disjoint cycle cover \mathcal{C} . Pick any cycle $C \in \mathcal{C}$. C must traverse vertices of G' in a particular way, namely, it must start w.l.o.g. at a selector vertex a_i , go through all cover-testing components corresponding to the edges incident on a particular vertex $v \in V$, and either return to a_i to finish the cycle or return to $a_j, j \neq i$, to continue covering in this manner before finally returning to a_i . This is because (A) each cover-testing component must be covered in one of the three configurations shown

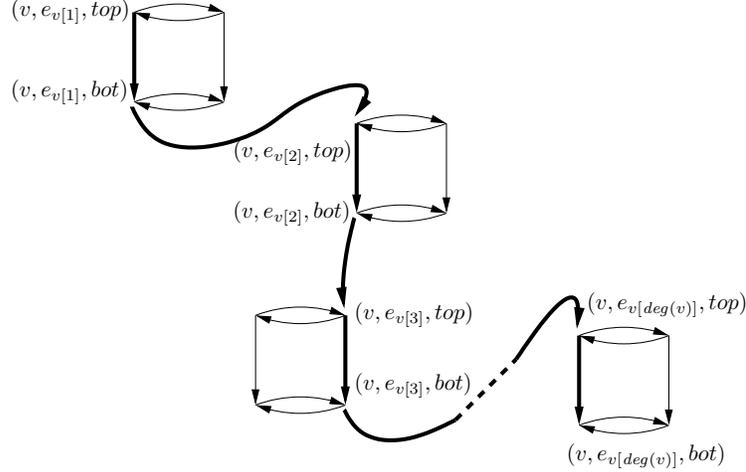


Figure 11: Path joining all cover-testing components for edges from E having vertex v as an endpoint

in Figure 10, and (B) edges in E'_v force C to traverse all cover-testing components of edges incident on v in order. Hence, the k vertices from $\{a_1, a_2, \dots, a_k\}$ divide C into k paths, each corresponding to a distinct vertex $v \in V$. Since C must include all vertices from every one of the cover-testing components, and since vertices from the cover-testing component for edge $e \in E$ can be traversed only by a path corresponding to an endpoint of e , every edge in E must have at least one endpoint among those k selected vertices. Therefore, this set of k vertices forms the desired vertex cover for G of size k .

Conversely, suppose $\{v_1, v_2, \dots, v_k\} = VC \subseteq V$ is a vertex cover for G with $|VC| = k$. It is readily seen that the following edges form a directed disjoint cycle cover of G' . From the cover-testing component representing each edge $e = \{u, v\} \in E$, choose the edges specified in Figure 10 depending on whether $\{u, v\} \cap VC$ equals $\{u\}$, $\{u, v\}$ or $\{v\}$. One of these three possibilities must hold because VC is a vertex cover for G . Next, choose all the edges in E'_{v_i} for $1 \leq i \leq k$. Finally, choose the edges $(a_i, (v_i, e_{v_i[1]}, top))$ for $1 \leq i \leq k$, edges $((v_i, e_{v_i[deg(v_i)]}, bot), a_{i+1})$ for $1 \leq i < k$, and the edge $((v_k, e_{v_k[deg(v_k)]}, bot), a_1)$. This forms a directed Hamiltonian circuit for G' and, in particular, a directed disjoint cycle cover. \square

We return to the proof of Theorem 4. The restricted degree problem is in NP because, by Lemma 22, the unbounded degree version is in NP. For NP-hardness, suppose we are given an instance $\langle G, k \rangle$ of VERTEX COVER. Use the construction in the proof of Lemma 22 to obtain a graph G' which has a directed disjoint cycle cover iff G has a vertex cover of size k . Notice that each of the k selector vertices a_i , $1 \leq i \leq k$, of G' have indegree as well as outdegree $n = |V(G)|$ corresponding to the edges in E'' , and the rest of the vertices have indegree as well as outdegree 2. We will describe a polynomial time process to convert G' into a directed graph G'' with indegree and outdegree each bounded above by 3. G'' will have a directed disjoint cycle cover iff G' does. Further, we will give a very simple way to convert G'' into a graph G^* that has outdegree exactly 3 and indegree bounded above by 4. G^* will have a directed disjoint cycle cover iff G'' does. This will finish the proof of the Theorem.

The construction of G'' is based on the “degree gadget” shown in Figure 12. Start by making G'' the subgraph of G' obtained by deleting all edges involving the k selector vertices a_1, \dots, a_k . Let v_1, v_2, \dots, v_n be the vertices of G . The i^{th} selector vertex a_i in G' has incoming edges from $(v_j, e_{v_j[deg(v_j)]}, bot)$ and outgoing edges to $(v_j, e_{v_j[1]}, top)$ for $1 \leq j \leq k$. Create vertices $V_{a_i}^{in} = \{b_{i,j}^{in} \mid 1 \leq j \leq n\}$ and $V_{a_i}^{out} = \{b_{i,j}^{out} \mid 1 \leq j \leq n\}$ in G'' and connect them as shown in Figure 12. Notice that vertex a_i in G'' has indegree and outdegree exactly 1, while the other degree gadget vertices have indegree and outdegree either 2 or 3. Overall, the vertices of G'' are $V(G') \cup \bigcup_{j=1}^k (V_{a_j}^{in} \cup V_{a_j}^{out})$, all of which have indegree as well as outdegree bounded above by 3.

We now argue that G'' has a directed disjoint cycle cover iff G' does. By construction, cycle covers in the two graphs behave identically on vertices other than the selector vertices and those of the degree gadgets. The only way to

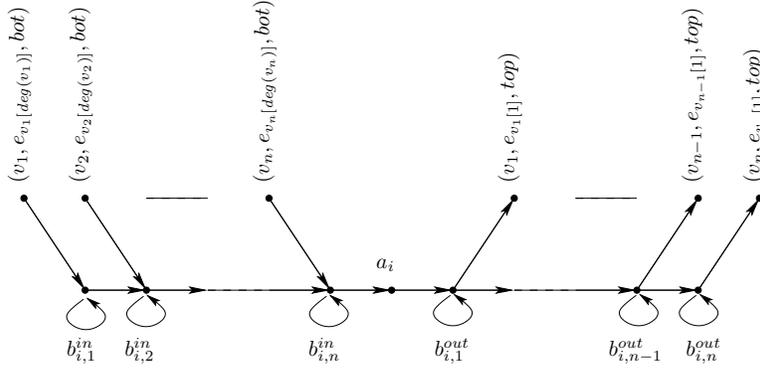


Figure 12: Degree gadget for the i^{th} selector vertex

cover a_i in G' is to enter from $(v_x, e_{v_x[\deg(v_x)]}, \text{bot})$ and leave to $(v_y, e_{v_y[1]}, \text{top})$ for some $v_x, v_y \in V(G)$. Similarly, the only way to cover a_i in G'' is to enter from $(v_x, e_{v_x[\deg(v_x)]}, \text{bot})$, traverse through $b_x^{in}, \dots, b_n^{in}$, cover a_i , traverse through $b_1^{out}, \dots, b_y^{out}$, and leave to $(v_y, e_{v_y[1]}, \text{top})$ for some $v_x, v_y \in V(G)$. The remaining degree gadget vertices in $V_{a_i}^{in} \cup V_{a_i}^{out}$ must be covered by self loops. This shows that cycle covers in G' and G'' behave essentially identically even on the selector vertices and the degree gadgets. Hence, G' has a directed disjoint cycle cover iff G'' has one.

G^* is constructed from G'' as follows. Start with $G^* = G''$. Let K_4 denote the directed complete graph on 4 vertices. Note that each vertex of K_4 has indegree as well as outdegree exactly 3. Let $p = |V(G'')|$ and q be the sum of the outdegrees of the vertices of G'' . Create $3p - q$ copies of K_4 and add them to G^* . For each outdegree 2 vertex v in G , add an outgoing edge to an arbitrarily selected vertex of an unmarked copy K_4 and mark that copy (see Figure 13). Similarly, do this for each outdegree 1 vertex, connecting it to two copies of K_4 . The number of edges added thus is $3p - q$, which exactly matches the number of copies of K_4 created in G^* . Every vertex of G^* now has outdegree exactly 3. The indegree of the original vertices from G'' hasn't changed and is therefore at most 3. The indegree of exactly one vertex in each copy of K_4 has changed, and it has increased to 4. Hence, G^* has the desired degree properties.

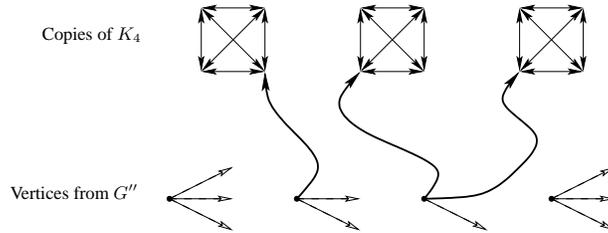


Figure 13: Using copies of K_4 to make outdegrees exactly 3

Since the copies of K_4 do not have any outgoing edge, the only cycles they can be involved in a disjoint cycle cover stay within that copy. This ensures that if G^* has a disjoint cycle cover, then G'' does too. On the other hand, if G'' has a disjoint cycle cover, then one for G^* consists of cycles corresponding to the cover in G'' along with one 4-cycle for each copy of the K_4 . Hence, G^* has a directed disjoint cycle cover iff G'' does. \square