

Games on Triangulations [March 6]

Oswin Aichholzer¹, David Bremner², Erik D. Demaine³,
Ferran Hurtado⁴, Evangelos Kranakis⁵, Hannes Krasser⁶,
Suneeta Ramaswami⁷, Saurabh Sethia⁸, and Jorge Urrutia⁹

¹ IICM - Software Technology, Graz University of Technology

² Faculty of Computer Science, University of New Brunswick

³ Laboratory for Computer Science, Massachusetts Institute of Technology

⁴ Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya

⁵ School of Computer Science, Carleton University

⁶ Institute for Theoretical Computer Science, Graz University of Technology

⁷ Computer Science Department, Rutgers University,

⁸ Department of Computer Science, Oregon State University

⁹ Instituto de Matemáticas, Universidad Nacional Autónoma de México

Abstract. We analyze several perfect-information combinatorial games played on planar triangulations. We introduce three broad categories of such games—constructing, transforming, and marking triangulations—and several specific games within each category. In various situations of each game, we develop polynomial-time algorithms to determine who wins a given game position under optimal play, and to find a winning strategy. Along the way, we show connections to existing combinatorial games such as Kayles and Nimstring (a variation on Dots-and-Boxes).

1 Introduction

A *triangulation* of a planar point set S is a simplicial decomposition of its convex hull whose vertices are precisely the points in S . We assume the point set S is in general position, i.e., no three points in S are collinear. In this work, we consider several perfect-information combinatorial games involving the vertices, edges (straight-line segments), and faces (triangles) of some triangulation. We present games in which two players \mathcal{R} (ed) and \mathcal{B} (lue) play in turns, as well as *solitaire* games for one player. In some *bichromatic* versions, player \mathcal{R} will use red and player \mathcal{B} will use blue, respectively, to color some element of the triangulation. In *monochromatic* variations, all players (or the single player) use the same color, green.

Games on triangulations come in three main flavors:

- *Constructing (a triangulation)*. The players construct a triangulation $T(S)$ on a given point set S . Starting from no edges, players \mathcal{R} and \mathcal{B} play in turn by drawing one or more edges in each move. In some variations, the game stops as soon as some structure is achieved. In other cases, the game stops when the triangulation is complete, and the last move or possibly some scoring decides then who wins.

- *Transforming (a triangulation)*. A triangulation $T(S)$ on top of S is initially given, all edges originally colored black. In each turn, a player applies some local transformation to the current triangulation, such as a flip, resulting in a new triangulation with some edges possibly recolored. The game stops when a specific configuration is achieved or no more moves are possible.
- *Marking (a triangulation)*. A triangulation $T(S)$ on top of S is initially given, all edges and nodes originally colored black. In each turn, some of its elements are marked (e.g., colored) in a game-specific way. The game stops when some configuration of marked elements is achieved (possibly the whole triangulation) or no more moves are possible.

We describe in Section 2 several games in each of these three categories. Our goal for each game is to characterize who wins the game, and to design efficient algorithms to determine the winner and compute a winning strategy. We present several such results in Sections 5–7. Our results contrast much of the work in algorithmic combinatorial game theory (see [6]), where games turn out to be computationally challenging in the sense of NP-hardness, PSPACE-hardness, EXPTIME-hardness, or EXPSPACE-hardness. Our results are all positive: despite the challenging nature of our games, we are able to obtain polynomial-time algorithms to solve many cases of them.

Besides beauty and entertainment, games keep attracting the interest of mathematicians and computer scientists because they also have applications to modeling several areas and because they often reveal deep mathematical properties of the underlying structures, in our case the combinatorics of planar triangulations.

Before we can summarize our results in Section 3, we define several particular games on triangulations in Section 2. Then, after some brief background and terminology in combinatorial games (Section 4), we study games xxx

2 Particular Games

We describe in this section the rules of several specific games that we have studied. See Table 1 for a high-level summary. Results on most of these games are given in the rest of the paper, while a few of them are considered for future research. Many other games in this general family also remain to be considered.

2.1 Constructing

2.1.1 Monochromatic Complete Triangulation. The players construct a triangulation $T(S)$ on a given point set S . Starting from no edges, players \mathcal{R} and \mathcal{B} play in turn by drawing one edge in each move. Each time a player completes one or more empty triangles, the player wins the corresponding number of points, and it is again her turn (an “extra move”). Once the triangulation is complete, the game stops and the player who owns more triangles is the winner.

Table 1. High-level summary of the games we consider.

Game	Category	Players	Chromatic
2.1.1 Monochromatic Complete Triangulation	Constructing	2	Mono
2.1.2 Monochromatic Triangle	Constructing	2	Mono
2.1.3 Bichromatic Complete Triangulation	Constructing	2	Bi
2.1.4 Bichromatic Triangle	Constructing	2	Bi
2.2.1 Monochromatic Flipping	Transforming	2	Mono
2.2.2 Monochromatic Flipping to Triangle	Transforming	2	Mono
2.2.3 Bichromatic Flipping	Transforming	2	Bi
2.2.4 All-Green Solitaire	Transforming	1	Mono
2.2.5 Green-Wins Solitaire	Transforming	1	Mono
2.3.1 Triangulation Coloring Game	Marking	2	Mono
2.3.2 Bichromatic Coloring Game	Marking	2	Bi
2.3.3 Four-Cycle Game	Marking	2	Mono
2.3.4 Nimstring Game	Marking	2	Mono

2.1.2 Monochromatic Triangle. Starts as in 2.1.1, but has a different stopping condition: the first player who completes one empty triangle is the winner.

2.1.3 Bichromatic Complete Triangulation. As in 2.1.1, but the two players use red and blue edges, and only monochromatic triangles count.

2.1.4 Bichromatic Triangle. As in 2.1.2, but with red and blue edges. The first empty triangle must be monochromatic.

2.2 Transforming

2.2.1 Monochromatic Flipping. Two players start with a triangulation whose edges are initially black. Each move consists of choosing a black edge, flipping it, and coloring the new edge green. The winner is determined by normal play.

2.2.2 Monochromatic Flipping to Triangle. As in 2.2.1, except now the winner is the player who completes the first empty green triangle.

2.2.3 Bichromatic Flipping. Two players play in turn, selecting a flippable black edge e of $T(S)$ and flipping it. Then e as well as any still-black boundary edges of the enclosing quadrilateral become red if it was player \mathcal{R} 's turn, and blue if it was player \mathcal{B} 's move. The game stops if no more flips are possible. The player who owns more edges of her color wins.

2.2.4 All-Green Solitaire. In each move, the player flips a flippable black edge e of $T(S)$; then e becomes green, as do the four boundary edges of the enclosing quadrilateral. The goal of the game is to color all edges green, or upon failure, to color as many edges as possible green.

2.2.5 Green-Wins Solitaire. As in 2.2.4, but the goal of the game is to obtain more green edges than black edges, or more generally, at least a specified ratio of green edges to black edges.

2.3 Marking

2.3.1 Triangulation Coloring Game. Two players move in turn by coloring a black edge of $T(S)$ green. The first player who completes an empty green triangle wins.

2.3.2 Bichromatic Coloring Game. Two players \mathcal{R} and \mathcal{B} move in turn by coloring red respectively blue a black edge of $T(S)$. The first player who completes an empty monochromatic triangle wins.

2.3.3 Four-Cycle Game. Same as 2.3.1, but the goal is to create an empty quadrilateral.

2.3.4 Nimstring Game. *Nimstring* is a game defined in Winning Ways [2] as a special case of the classic children’s (but nonetheless deep) combinatorial game Dots and Boxes [1, 2]. In the context of triangulations, players in *Nimstring* alternate coloring edges green in an initially black triangulation, and whenever one or more green triangles are completed, the completing player must move again (an “extra move”). The winner is determined by normal play, meaning that the goal is to make the last complete move. Thus, the player marking the last edge of the triangulation actually loses, because that last edge completes one or two triangles, and the player is forced to move again, which is impossible.

3 Summary of Results

All of the games defined in the previous section are in PSPACE, because the number of moves is always polynomial in the size of the triangulation: a polynomial-space algorithm can thus perform a depth-first search through the game tree. We conjecture that many of the games are also NP-hard or PSPACE-complete for general triangulations. Therefore, we focus our attention to special classes of triangulations where positive results are possible:

- *Convex position.* Often we consider the case in which the points are in convex position. In such situations, the *dual* of the triangulation (with a vertex for each triangle and an edge connecting each adjacent pair of triangles) is a tree with maximum degree 3.
- *Few inner points.* More general than points in convex position is if the number of *inner points* (points interior to the convex hull) is small, e.g., at most a constant.
- *Fans.* A fan triangulation has all triangles incident to a common point on the convex hull; all points are in convex position.

- *Wheels*. A wheel triangulation has exactly one inner point (the center), and all triangles are incident to this center point.
- *Serpentine triangulations*. A triangulation is *serpentine* if its dual is a simple path. Equivalently, a serpentine triangulation has no *inner triangles*, triangles with all three edges interior to the convex hull. The points are necessarily in convex position.
- *Simple-branching triangulations*. A triangulation is *simple-branching* if no two inner triangles share an edge. In the dual, which is necessarily a tree, this definition means that degree-3 vertices are separated by paths of at least one degree-2 vertex.

Table 2 summarizes which games we solve in which cases, in terms of developing algorithms to either decide the winner or to find a winning strategy. Each cell denotes the running time of the algorithm. $O(1)$ in the Winner column means that the winner can be determined in constant time by simply knowing the number of vertices, edges, and triangles. The Strategy column distinguishes whether the strategy can be computed quickly per move or the time bound is amortized over the entire game.

Table 2. Summary of our results.

Game	Triangulations	Winner	Strategy	Sec.
2.1.1 Monochromatic Complete Triang.	Convex position	$O(1)$	$O(n)/\text{move}$	5.3
2.1.2 Monochromatic Triangle	Convex position	$O(1)$	$O(n)/\text{move}$	5.2
2.1.3 Bichromatic Complete Triang.	—	—	—	
2.1.4 Bichromatic Triangle	—	—	—	
2.2.1 Monochromatic Flipping	Convex position	$O(1)$	$O(1)/\text{move}$	6.1
2.2.2 Monochromatic Flipping to Triangle	Serpentine	$O(1)$	Poly.	6.2
2.2.3 Bichromatic Flipping	—	—	—	
2.2.4 All-Green Solitaire	Convex position	$O(n)$	$O(n)$	7.1
2.2.5 Green-Wins Solitaire	General	bounds on ratio		7.2
2.3.1 Triangulation Coloring Game	Simple-branching	$O(n)$	$O(n)$	5.1
	Other	Poly.	Poly.	5.1
2.3.2 Bichromatic Coloring Game	≤ 2 inner points	$O(1)$	$O(n)/\text{move}$	5.5
2.3.3 Four-Cycle Game	Serpentine	$O(1)$	Poly.	5.6
	Wheel	$O(1)$	Poly.	5.6
2.3.4 Nimstring Game	Even fan	$O(1)$	—	5.4
	Odd wheel	$O(1)$	—	5.4

4 Background in Combinatorial Games

Games on triangulations belong to the more general area of *combinatorial games* which typically involve two players, \mathcal{R} and \mathcal{B} . Before we start proving our results, we need to define a few more terms from combinatorial game theory. For

more information, refer to the books [2, 4], the survey [6], and the introduction/bibliography [7] which contains over 900 references.

We consider games with *perfect information* (no hidden information as in many card games) and without chance moves (like rolling dice). In such a game, a *game position* consists of a set of options for \mathcal{R} 's moves and a set of options for \mathcal{B} 's moves. Each option is itself a game position, representing the result of the move.

Many of the games we consider in this paper (the monochromatic games) are also *impartial* in the sense that the options for \mathcal{R} are the same as the options for \mathcal{B} . In this case, a game position is simply a set of game positions, and can thus be viewed as a tree. The leaves of this tree correspond to the empty set, meaning that no options can be played; this game is called the *zero game*, denoted 0.

In general, each leaf might be assigned a label of whether the current player reaching that node is a winner or loser, or the players tied. However, a common and natural assumption is that the zero game is a losing position, because the next player to move has no move to make. We usually make this assumption, called *normal play*, so that the goal is to make the last move. In contrast, *misère play* is just the opposite: the last player able to move loses. In more complicated games, the winner is determined by comparing scores.

Any impartial perfect-information combinatorial game without ties has one of two *outcomes* under optimal play (when the players do their best to win): a *first-player win* or a *second-player win*. In other words, whoever moves first can force herself to reach a winning leaf, or else whoever moves second can force herself to reach a winning leaf, no matter how the other player moves throughout the game. Such forcing procedures are called *winning strategies*. For example, under normal play, the game 0 is a second-player win, and the game having a single move to 0 is a first-player win, in both cases no matter how the players move. More generally, impartial games may have a third outcome: that one player can force a *tie*.

The Sprague-Grundy theory of impartial games (see e.g. [2, Chapter 3]) says that, under normal play, every impartial perfect-information combinatorial game is equivalent to the classic game of *Nim*. In (single-pile) Nim, there is a pile of $i \geq 0$ beans, denoted $*i$, and players alternate removing any positive number of beans from the pile. Only the empty pile $*0$ results in a second-player win (because the first player has no move); for any other pile, the first player can force a win by removing all the beans. If a game is equivalent to $*i$, then i is called the *Nim value* of the game.

Given two or more games, their *sum* is the game in which, at each move, a player chooses one subgame to move in, and makes a single move in that subgame. In this sense, sums are *disjunctive*: a player makes exactly one move at each turn. Games often split into sums of independent games in this way, and combinatorial game theory explains how the sum relates to its parts. In particular, if we sum two games with Nim values i and j , then the resulting game has Nim value equal to the bitwise XOR of i and j .

A particularly useful game in the context of sums and Nimstring games in Section 5 is the *loony game*, denoted \mathfrak{D} . This game has the property that its sum with any other game is again \mathfrak{D} , and that the outcome of \mathfrak{D} is a first-player win. We will see examples of the loony game in Section 5.4.

5 Drawing Triangles and Empty Cycles

In this section we consider games dealing with the drawing of empty cycles (typically triangles) of a triangulation. We show that some of these games are equivalent to famous combinatorial games and describe their solutions.

5.1 Triangulation Coloring Game and Kayles

Obviously the Triangulation Coloring Game (2.3.1) terminates after a linear number of moves and there are no ties. For point sets S in convex position and several classes of triangulations $T(S)$ of S we will show a one-to-one relation to seemingly unrelated games on piles of beans, which will provide us with an optimal winning strategy for these settings. To this end, consider the dual of the triangulation $T(S)$, i.e., the graph G_T with a vertex per triangle of $T(S)$ and an edge between each pair of vertices corresponding to triangles of $T(S)$ that share a diagonal. An *inner* triangle of $T(S)$ consists entirely of diagonals of $T(S)$, and therefore it does not use an edge of the convex hull of S . Thus, exactly those vertices of G_T corresponding to inner triangles have degree three, whereas all other vertices have degree one (ears of the triangulation) or two.

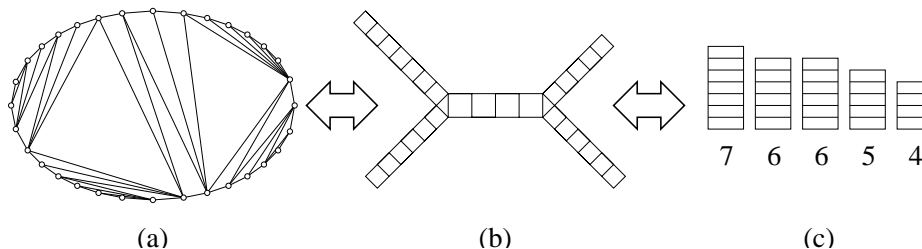


Fig. 1. Different incarnations of the triangulation coloring game

Motivation for considering the dual graph of $T(S)$ stems from the following observation. Coloring an edge of a triangle Δ for which one edge has already been colored leads immediately to a winning move for the opponent: she just has to color the third edge of Δ . Thus we call any triangle Δ with one colored edge ‘taken’, because we can never color another edge of Δ unless we are ready to lose. Thus, coloring an edge e of $T(S)$ means, in the dual setting, marking either a single vertex of G (if e was on the convex hull) or two adjacent vertices (if e was a diagonal) as taken. Vertices already marked cannot be marked again and

whoever marks the last vertex will win. Figure 1(b) shows G as a set of connected arrays of boxes, where marking a vertex of G might be seen as drawing a cross inside the corresponding box.

If the triangulation is serpentine (equivalently, a single array of boxes without branches), we can show that the first player has a winning strategy by applying a symmetry principle (called the *Tweedledum-Tweedledee Argument* in [2]). For an odd number of triangles, she first takes the central triangle by coloring the edge of this triangle that belongs to the convex hull. For an even number of triangles, she first takes both triangles adjacent to the central diagonal by coloring this diagonal. In both cases, the remainders are two combinatorially identical triangulations (two equal-sized box arrays), in which all possible moves can be played independently. Thus the winning strategy of player one is just to mimic any of her opponent's moves by simply coloring the corresponding edge in the other triangulation. This strategy ensures that she always can make a valid move, forcing the second player finally to color a second edge of an already taken triangle, leading to a winning position for the first player.

If $T(S)$ contains inner triangles, the dual is a tree and the problem of finding optimal strategies is more involved. We consider this situation with the simple-branching restriction that no two inner triangles share a common diagonal; see for example Figure 1(a). The main observation for our game is that all inner triangles can be ignored: Consider an inner triangle Δ and observe first that it cannot be taken on its own, because Δ does not have an edge from the convex hull of S . Thus the situation after the three neighbors of Δ have been taken is the same, regardless of whether Δ was taken together with one of them: Δ is blocked in any case. In the dual setting, this observation means that we can remove the vertex of G corresponding to Δ (plus adjacent edges) without changing the game. Drawing G with blocks as in Figure 1(b), we can thus remove the 'triangular' blocks and consider the remaining block arrays independently. Instead of playing with these arrays, we might as well deal with integers reflecting the length of each array; see Figure 1(c).

Surprisingly, this setting turns out to be an incarnation of a well-known taking-and-breaking game played on heaps of beans or sets of coins, called *Kayles* [2]. This game was introduced by Dudeney and independently also by Sam Loyd, who originally called it 'Rip Van Winkle's Game'. The following description is taken from [2, Chapter 4]: *Each player, when it is his turn to move, may take 1 or 2 beans from a heap, and, if he likes, split what is left of that heap into two smaller heaps.*

Any triangulation of a convex set S without inner triangles sharing a common diagonal can be represented by Kayles, while the reverse transformation is less general. The number of heaps which can be represented by a single legal triangulation has to be odd because any inner triangle has degree three. During a game of course any number of heaps may occur, as the triangulation may split into several independent parts. A generalization is thus to play the game on more than one point set from the very beginning.

Because Kayles is impartial, the Sprague-Grundy theory described in Section 4 applies, so the game is completely described by its sequence of Nim values for a single pile of size n . It has been shown that this Nim sequence has a periodicity of length 12, with 14 irregularities occurring, the last for $n = 70$; see Table 5.1. To compute the Nim value for a game with several heaps, we can xor-add up the Nim values (given by Table 5.1) for the individual heaps. Moreover, in this case, we can xor-add up just a four-bit vector, corresponding to the four ‘magic’ heap sizes 1, 2, 5, and 27, respectively, where powers of two in the Nim sequence appear for the first time. For example, a single heap of size 42 (Nim value 7) is equivalent to the situation of three heaps with sizes 1, 2, and 5, respectively, reflecting the ‘ones’ in the 4-bit representation of 7.

Table 3. Nim values for Kayles: the Nim sequence has periodicity 12 and there are 14 exceptional numbers

$\mathcal{G}(n) = K[n \bmod 12], K[0, \dots, 11] = (4, 1, 2, 8, 1, 4, 7, 2, 1, 8, 2, 7)$									
exceptional values:									
$\mathcal{G}(0) = 0$	$\mathcal{G}(3) = 3$	$\mathcal{G}(6) = 3$	$\mathcal{G}(9) = 4$	$\mathcal{G}(11) = 6$	$\mathcal{G}(15) = 7$	$\mathcal{G}(18) = 3$			
$\mathcal{G}(21) = 4$	$\mathcal{G}(22) = 6$	$\mathcal{G}(28) = 5$	$\mathcal{G}(34) = 6$	$\mathcal{G}(39) = 3$	$\mathcal{G}(57) = 4$	$\mathcal{G}(70) = 6$			

It follows that, in time linear in the number of heaps, a position can be determined to be either a first-player win (nonzero Nim value) or a second-player win (zero Nim value). Any move from a second-player-winning position leads to a first-player-winning position; and for any first-player-winning position there is always at least one move that leads to a second-player-winning position. A winning strategy just needs to follow such moves, because after one move, the players effectively reverse roles. Because any position has at most a linear number of possible moves, we conclude that for the triangulation-coloring game a winning move (if it exists) can be found in time linear in the size of the triangulation. It is interesting to note that there are no zeros in the Nim sequence of Kayles. This reflects the fact that when starting with a single integer number the first player can always win, as has been pointed out above for triangulations without inner triangles.

From the previous discussion we obtain the following result:

Theorem 1. *Deciding whether the Triangulation Coloring Game on a simple-branching triangulation on n points in convex position is a first-player win or a second-player win, as well as finding moves leading to an optimal strategy, can be solved in time linear in the size of the triangulation.*

At this point it is worth mentioning that there is a version of Kayles played on graphs: two players play in turn by selecting a vertex of a given graph G that must be nonadjacent to (and different from) any previously chosen vertex.

The last player that can select a vertex, completing a maximal independent set, is the winner. Deciding which player has a winning strategy is known to be PSPACE-complete [9].

Now, given a triangulation T on a point set S , let us define a graph $EG(T)$ having a vertex per each edge in T and an adjacency between any two nodes whose corresponding edges in T belong to the same triangle; an example is shown in Figure 2. From the preceding paragraphs, it is clear that playing the Triangulation Coloring Game on T is equivalent to playing Kayles on $EG(T)$.

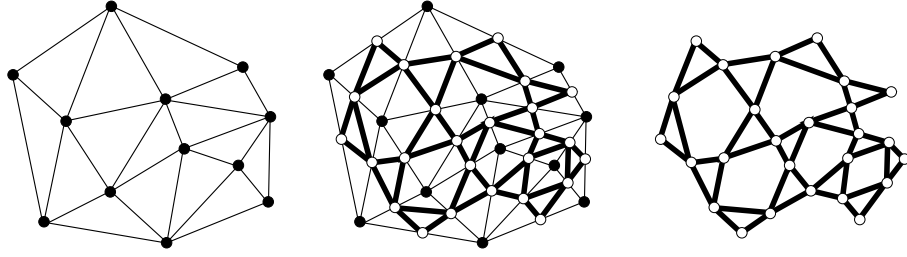


Fig. 2. A triangulation T of a point set (left) and the graph $EG(T)$ associated to adjacent edges (right)

While such a reduction does not prove hardness of the Triangulation Coloring Game, it does transfer any solutions to special cases of Kayles. In [3] it is shown that there are polynomial-time algorithms to determine the winner in Kayles on graphs with bounded asteroidal number, on cocomparability graphs, and on circular-arc graphs. Theorem 1 can be rephrased as a similar (and computationally efficient) result for outerplanar graphs in which every block is a triangle and blocks that contain three articulation vertices do not share any of them.

5.2 Monochromatic Triangle Game and Dawson's Kayles

We present an optimal strategy for Monochromatic Triangle (2.1.2) provided S is a convex set. First observe that an edge should be drawn only if it connects two vertices that have not been used before. Otherwise, a vertex p of degree at least two occurs, leading to a winning move for the opponent: she just has to close the triangle formed by two neighboring edges of p . (Note that it is important here that we consider point sets in convex position.) In other words, when drawing a diagonal pq in S , the two vertices p and q are taken for the rest of the game. Moreover, pq splits S into two independent subsets with cardinality n_1 and n_2 , respectively, such that $n_1 + n_2 = n - 2$. The player who draws the last edge according to these observations will win the game with her succeeding move.

Because no edge can be drawn in sets of cardinality of at most one, we have just shown that our game is an incarnation of a known game called *Dawson's Kayles*, a cousin of Kayles [2]. In terms of bowling, the game reads as follows:

A row of n pins is given and the only legal move is to knock down two adjacent pins. Afterwards, one or two shorter rows of pins remain, and single pins are removed immediately. Whoever makes the last strike wins.

In more mathematical terms, the game is defined by a set of k integers n_1, \dots, n_k . A move consists of choosing one n_i , $1 \leq i \leq k$, reducing it by two to \hat{n}_i and eventually replacing it afterwards by two numbers n'_i and n''_i , $n'_i + n''_i = \hat{n}_i$. Any $n_i \leq 1$ can be removed from the set, because it cannot be used for further moves. Whoever can make the last legal move wins. Note the case that n_i is not split after reduction corresponds to drawing an edge on the convex hull of S (or the respective subset).

Table 4. Nim values for Dawson's Kayles: the Nim sequence has periodicity 34 and there are 8 exceptional numbers

$\mathcal{G}(n) = K[n \bmod 34], K[0, \dots, 33] =$				
(4, 8, 1, 1, 2, 0, 3, 1, 1, 0, 3, 3, 2, 2, 4, 4, 5, 5, 9, 3, 3, 0, 1, 1, 3, 0, 2, 1, 1, 0, 4, 5, 3, 7)				
exceptional values:				
$\mathcal{G}(0) = 0$	$\mathcal{G}(1) = 0$	$\mathcal{G}(15) = 0$	$\mathcal{G}(17) = 2$	$\mathcal{G}(18) = 2$
$\mathcal{G}(32) = 2$	$\mathcal{G}(35) = 0$	$\mathcal{G}(52) = 2$		

Sprague-Grundy theory also applies to Dawson's Kayles. It has been shown that its Nim sequence has a periodicity of length 34, with 8 irregularities occurring, the last for $n = 52$; see Table 5.2. As with Kayles, to compute the Nim value for a position consisting of k heaps, i.e., to xor-add up the Nim values given by Table 5.2 for the k heap sizes, a vector with four bits is sufficient, corresponding to the heap sizes 2, 4, 14, and 69.

Theorem 2. *The Monochromatic Triangle game on n points in convex position is a second-player win when $n \equiv 5, 9, 21, 25, 29 \pmod{34}$ and for the special cases $n = 15$ and $n = 35$; otherwise it is a first-player win. Each move in a winning strategy can be computed in time linear in the size of the triangulation.*

[I think you should be able to improve this to $O(n)$ time total... xxx
—ERIK]

For n even, this result was clear from the very beginning, as in this case the first player, say \mathcal{R} , may start by drawing a diagonal d leaving $(n-2)/2$ points on each side and apply the symmetry principle: for every move of \mathcal{B} , player \mathcal{R} either makes a winning move, if available, or mimics her opponent's last move on the opposite side of d .

5.3 Monochromatic Complete Triangulation Game

In this subsection we consider the triangulation-construction game 2.1.1. In this context, we show by direct arguments that for a set S in convex position a greedy

strategy is optimal for this game where, depending on the parity of n , the first player can always win (odd n) or either player can force a tie (even n).

Theorem 3. *The outcome of the Monochromatic Complete Triangulation Game on n points in convex position is a first-player win for n odd, and a tie for n even.*

Let us call two edges sharing a common point p an *open triangle* if we can build a valid triangle (no intersections with other edges occur) by connecting the two endpoints not adjacent to p by inserting a third edge, called *closing* edge. Obviously closing edges are drawn between vertices of the same connected component. When drawing an edge connecting two formerly different components we call it an i -edge, $i \in \{0, 1, 2\}$, if i of its endpoints already have at least one other incident edge. Thus, 0-edges connect isolated points while, by convexity, 1-edges produce one additional open triangle, and 2-edges give rise to two additional open triangles. Because we have n points overall, the total number of i -edges throughout a game is $n - 1$.

Note that in addition to these two types of edges there exist so-called *redundant* edges: connecting two points from the same connected component, but not closing a triangle. This happens if a cycle of length at least 4 occurs, containing several open triangles. We first argue that any optimal strategy uses no redundant edges, i.e., open triangles will be closed immediately. Otherwise the opponent might close the triangles, getting the points, and continue afterwards with the same number of possible i -edges. Here it is crucial to observe that when an edge connects two different connected components, it is not important for the strategy which points of these components are used, since when closing all open triangles of the new connected component everything within its convex hull is triangulated. Thus for analyzing strategies not the exact shape but only the number of connected components counts.

The greedy strategy works as follows. As long as there are closing edges, draw them. Recall that after closing a triangle, it remains the same player's turn. Then draw an i -edge for the smallest possible i .

To analyze our strategy, let e_i denote the number of i -edges drawn during an entire game. The first time a point of S is used, it is either part of a 0-edge or a 1-edge. Also, a 0-edge uses two previously unused points, whereas a 1-edge uses one previously unused point. Thus, $2e_0 + e_1 = n$. Moreover, $e_2 = e_0 - 1$ because $e_0 + e_1 + e_2 = n - 1$. Further observe that if there are no open triangles left and a player plays an i -edge then her opponent can, and will by the observations above, close exactly i open triangles in her next move. (Note that only i triangles can be closed only because S is in convex position.) Thus, the goal of a player is to globally minimize the sum of i over all i -moves she makes.

We split the remaining proof of the theorem into three parts:

Lemma 1. *For n odd, player \mathcal{R} can win by playing greedily.*

Proof. We have $e_0 + e_1 + e_2 = n - 1$ which is even. Thus, there will be $(n - 1)/2$ rounds of both players picking i -edges. In each round, player \mathcal{R} picks first and

greedily, and hence in each round \mathcal{R} wins or ties with \mathcal{B} . For a tie to occur, \mathcal{B} must tie with \mathcal{R} in all rounds, but that requires that e_0, e_1, e_2 all be even, which is not possible because $e_2 = e_0 - 1$. \square

Lemma 2. *For n even, player \mathcal{B} can force a tie by playing greedily.*

Proof. We have $e_0 + e_1 + e_2 = n - 1$ which is odd. If the first move for player \mathcal{R} is an i -edge, then player \mathcal{B} wins i points. After this first move, there are $n - 2$ i -edges remaining. Players \mathcal{B} and \mathcal{R} will pick these i -edges alternately in $(n - 2)/2$ rounds. In each round, player \mathcal{B} picks first and greedily, and hence in each round \mathcal{B} either wins or ties with \mathcal{R} . Thus, \mathcal{B} either wins or ties overall by playing greedily. \square

Lemma 3. *For n even, player \mathcal{R} can force a tie.*

Proof. Here we diverge from the greedy strategy, because if e_0 ended up even, then player \mathcal{B} would win by two triangles by playing greedily (only e_2 is odd). Instead, \mathcal{R} employs a symmetry strategy to ensure that e_0 ends up odd, so that both e_1 and e_2 are even, leading to a tie. Player \mathcal{R} begins by playing a diagonal splitting S into two equal sets (recall that n is even). Then as long as \mathcal{B} does not unnecessarily leave triangles open, she plays symmetrically: close open triangles and mimic whatever \mathcal{B} has done in the opposite part of S . In this way, it is guaranteed that e_0 will be odd: the first diagonal plus two times the number of 0-edges \mathcal{B} has drawn. If at some point \mathcal{B} does not close an open triangle, then \mathcal{R} closes it and starts playing according to the ordinary greedy strategy. Because \mathcal{R} now won a triangle from \mathcal{B} , the scoring difference changed by two and thus with the greedy strategy \mathcal{R} will win for e_0 odd and get at least a tie for e_0 even. \square

[What's the running time? Definitely $O(n)$ per move. Can it be done in $O(1)$ /move? —ERIK] xxx

5.4 Nimstring Game and Dots and Boxes

Nimstring is a game defined in *Winning Ways* [2, pp. 518–520] as closely related to the classic children's (but nonetheless deep) combinatorial game *Dots and Boxes* [1, 2]. In the context of triangulations, players in *Nimstring* alternate *marking* edges one-by-one, and whenever a triangle has all three of its edges marked, the completing player is awarded an extra move and must move again. Completing two triangles with one stroke awards only a single extra move. The winner is determined by normal play, meaning that the goal is to make the last entire move. Thus, the player marking the last edge of the triangulation actually loses, because that last edge completes one or two triangles, and the player is forced to move again, which is impossible.

An equivalent view of the game involves “coins” and “strings.” Place a coin on each triangle. Tie a string between coins of adjacent triangles. In addition, for each edge that is adjacent to only one triangle (a *boundary edge*), tie a string

between that triangle's coin and the “floor.” A triangle with multiple boundary edges has multiple strings tying its coin to the floor, so that every coin has exactly three strings attaching it to other coins and/or the floor. Now the *String-and-Coins* game [2, pp. 516–517] is played as follows. At each turn, the current player breaks a string. If a coin (or two coins) become free, the player picks up the coin and is forced to break another string. Following normal play, the player who cannot complete her turn loses. In other words, the player who picks up the last coin loses, because she is required to break another string but none remain.

Returning to the triangulation setting, we call a triangle *completable* if it has exactly two of its edges marked. Such triangles can be completed to award an extra move. (In the Strings-and-Coins setting, a coin is *completable* if it has exactly one string remaining.) Winning Ways [2, pp. 521–522] characterizes when completable triangles should be completed: the current player should always complete a completable triangle unless the position is loony, which is detectable in linear time.

In this context, a configuration is loony if it contains one of the loony subconfigurations shown in Figure 3. These subconfigurations are loony because the current player can either play e_2 and complete her turn, or play e_1 and then e_2 and then anywhere else in the triangulation. Thus, the current player effectively gets to choose whether to play first or second in the rest of the triangulation, one of which necessarily leads to a win. Thus, in principle, loony positions are first-player wins, although the winning move might not be transparent.

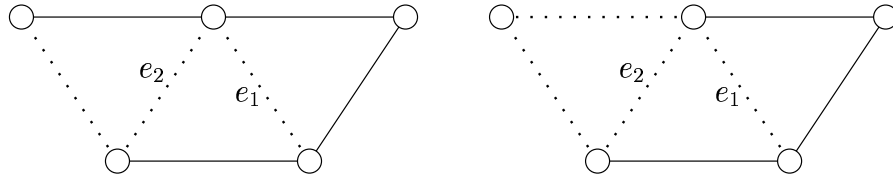


Fig. 3. Loony Nimstring subconfigurations: one completable triangle, incident via e_1 to a triangle with exactly one marked edge, incident via e_2 to an incompletable triangle.

We might hope that Nimstring is a first-player win precisely if the number of edges (submoves) minus the number of triangles (free moves) is odd. By Euler’s formula, this condition is equivalent to the number of vertices being even. However, this hope turns out to be false in general. The idea is that the player losing in terms of parity can switch the parity by forcing the other player to create a loony position.

We give partial characterizations of the winning player for two special classes of triangulations: fans and wheels.

Fans. A *fan* is a triangulation whose dual is a path and whose triangles share a common vertex called the *center*. We call the edges incident to the center the *spokes*, and call the edges not incident to the center the *rim edges*. See Figure 4.

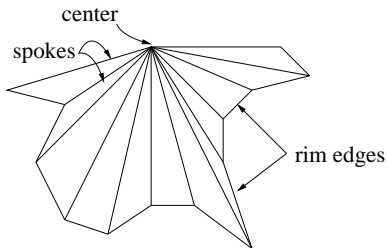


Fig. 4. A fan triangulation.

Theorem 4. *Nimstring in a fan with an even number of vertices is a first-player win.*

Proof. A winning strategy for the first player is as follows. Her first move is to mark the middle spoke. This move divides the fan into two symmetric parts. The first player then plays roughly symmetrically to the second player: at each move, the first player completes any completable triangles, and then marks the reflection of the second player's last move through the middle spoke. This symmetric play continues until either the end of the game, in which case the first player wins by parity, or the second player creates a lune.

If the second player ever creates a lune, we know that there is a strategy for a first-player win. (The first player will not continue to play symmetrically, because that will create a lune for the second player.) Because the first player is playing symmetric to the second player (except for triangle completion), the first player cannot create a lune first. If the second player never creates a lune, then a four cycle is never marked before its diagonal, so marking an edge never led to simultaneous completion of two triangles. Hence the total number of moves before the game ends is the number of edges minus the number of triangles in the fan, which is $(2n - 3) - (n - 2) = n - 1$ where n is the number of vertices. Because n is even, the second player will not be able to complete her turn at the end of game and hence will lose. \square

It remains open who wins with fans of an odd number of vertices. Counter-intuitively, it is not always the second player. See Table 5 for the solutions for small fans. Surprisingly, the misère version of the game seems simpler, depending only on parity; see Table 6.

Wheels. A *wheel* is a triangulation whose dual is a cycle. All the triangles of a wheel share a common vertex called the *center*. As with fans, we call the edges incident to the center the *spokes*, and the edges not incident to the center the *rim edges*. See Figure 5.

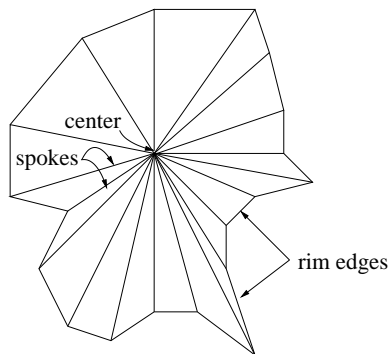
Theorem 5. *Nimstring in a wheel with an odd number of vertices is a second-player win.*

Table 5. Experimental results on Nimstring with normal play.

Triangles in fan	Outcome	Nimber	Winning moves
1	2nd	0	None
2	1st	1	All
3	2nd	0	None
4	1st	1	All
5	2nd	0	None
6	1st	1	Middle 3 spokes
7	1st	3	2nd rim or symmetric
8	1st	2	Middle spoke
9	1st	1	4th spoke or symmetric

Table 6. Experimental results on Nimstring with misère play.

Triangles in fan	Outcome	Nimber	Winning moves
1	1st	1	All
2	2nd	0	None
3	1st	1	All
4	2nd	0	None
5	1st	1	All
6	2nd	0	None
7	1st	1	All
8	2nd	0	None
9	1st	1	All
10	2nd	0	None

**Fig. 5.** A wheel triangulation.

Proof. A wheel with even number of triangles is already symmetric in the sense that every spoke has a well-defined opposite spoke and every rim edge has a well defined opposite rim edge, corresponding to rotation by 180° . The second player thus plays roughly symmetrically to the first player: at each move, the second player completes any completable triangles, and then marks the opposite of the second player's last move. This symmetric play continues until either the end of the game, in which case the second player wins by parity, or the first player creates a lune.

If the first player ever creates a lune, we know that there is a strategy for a second-player win. Because the second player is playing symmetric to the first player, the second player cannot create a lune first. If the first player never creates a lune, then a four cycle is never marked before its diagonal, so marking an edge never led to simultaneous completion of two triangles. Hence the total number of moves before the game ends is the number of edges minus the number of triangles in the wheel, which is $(2n - 2) - (n - 1) = n - 1$ where n is the number of vertices. Because n is odd, the first player will not be able to complete her turn at the end of game and hence will lose. \square

5.5 Bichromatic Coloring Game

We recall that *Bichromatic Coloring Game* 2.3.2 is defined as follows: Two players \mathcal{R} and \mathcal{B} move in turn by coloring red respectively blue a black edge of a triangulation $T(S)$ on top of a point set S . The first player who completes an empty monochromatic triangle wins.

A tie may occur for any given pair S and $T(S)$. This follows from the observation below that every triangulation of a simple polygon with (possibly) inner points can be edge-colored with two colors without attaining a monochromatic empty triangle. The result holds if both colors are equally often used (up to one for an odd number of edges of $T(S)$) as in a tied game.

Lemma 4. *Every triangulation of a simple polygon with (possibly) inner points can be edge-colored with two colors without attaining a monochromatic empty triangle. The result holds if both colors are equally often used (up to one edge for an odd number of edges).*

Proof. We prove the lemma by induction on the number of vertices. For the induction base we have $n \leq 3$ and therefore nothing to prove. For $n \geq 4$ remove the rightmost (and in case there are more such points the topmost) point p and its adjacent edges. What remains is either a single simple polygon with (possible) inner points or a set of such polygons. By induction the set(s) can be colored according to the lemma. Now re-insert p and all adjacent edges and alternate the coloring of these edges alternatively. Note that p was chosen to be a vertex of the boundary and thus even for an odd number of adjacent edges no two re-inserted edges of the same color are neighbored, i.e. can belong to a triangle of the triangulation. In other words the new edges cannot introduce a monochromatic triangle. If the number of adjacent edges of p is even the balance

between the number of red and blue edges is unchanged. Otherwise we chose the color previously less used (by -1) to 'start' the alternating coloring for the edges of p . Use a random color if both colors have been used similar often. In all cases we get a balance of 0 (even number of edges) or ± 1 (odd number of edges). \square

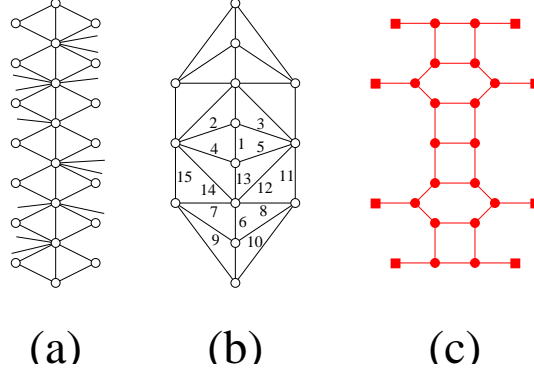


Fig. 6. A diamond-path without chords (a), as a triangulation (b), and its dual (c).

Despite the result of Lemma 4 we can show that for special sets there is still a winning strategy for the first player, say \mathcal{R} . As an example we consider the so-called *diamond-path without chords*, consisting of a chain of seven central edges and several adjacent protruding edges forming seven triangles on each side and possibly additional edges in between them, cf. Fig. 6(a). No chords within the chain are allowed, i.e., any protruding edge has exactly one (its inner) endpoint in common with inner edges, and two protruding edges share their outer endpoints if and only if they build one of the triangles of the chain. A diamond-path without chords might occur as a subset in a triangulation or by its own, see Figure 6 (b) and (c), where the latter drawing shows the dual structure of.

The following is a winning strategy for player \mathcal{R} , using edge numbers according to Figure 6 (b) as shortcuts. When giving a sequence of moves for player \mathcal{R} we display forced moves for player \mathcal{B} in brackets. Player \mathcal{R} starts with edge 1 and forces \mathcal{B} to answer with 2,3,4, or 5. Otherwise \mathcal{R} could start to build a star around one of the end-points of 1 (in direction away from \mathcal{B} 's first move), always forcing edges of \mathcal{B} . When finally the cycle is closed \mathcal{R} will win. Thus w.l.o.g. let \mathcal{B} 's move be 2. Now \mathcal{R} answers with 6, forcing \mathcal{B} to play 7,8,9, or 10. If \mathcal{B} plays 7 or 9 then \mathcal{R} wins by the sequence 8, (10), 12, (11), 5, (3,13). If \mathcal{B} 's move is 8 or 10 then the sequence 7, (9), 14, (15), 13, (4), 5, (3,12) gives a winning situation for \mathcal{R} .

A complementary result is that if S contains at most 2 inner points then player \mathcal{B} can force a tie.

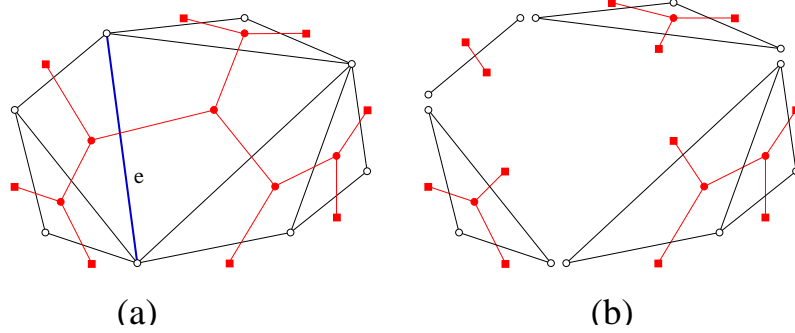


Fig. 7. Coloring an edge e blue splits the dual graph.

We first show this result to be true for S in convex position. Consider the dual of S , which is a tree in this case, where edges of the convex hull of S correspond to leaves, see Figure 7(a). When coloring an edge of the triangulation red we color the related edge of the dual graph similar. Vertices of the dual graph corresponding to triangles have degree three and to get a red triangle for player \mathcal{R} means to color these three edges red. When player \mathcal{B} colors an edge blue this “takes away” the two adjacent (one, if it was a convex hull edge) triangles for player \mathcal{R} . In the dual graph we remove the corresponding edge and its two end-points, where eventually other edges adjacent to these end-points get leaves instead. See Figure 7 (a) and (b) for an example. Thus the dual graph shows only black and red edges, and player \mathcal{B} s moves split the graph into independent sub graphs. Observe that as long as no sub graph contains more than two red edges \mathcal{R} cannot win.

We are now ready to give \mathcal{B} s defense strategy, described in the dual setting: Suppose it is \mathcal{B} s turn and there is at most one sub graph with two red edges, all other sub graphs containing at most one red edge. Then \mathcal{B} obviously can split the sub graph containing two red edges into at least two graphs, each containing at most one red edge. Since \mathcal{R} can add only one red edge at a time we get a similar setting after \mathcal{R} s move, and therefore \mathcal{B} can force a tie by repeating the same strategy until all edges are colored.

When starting with a set S in convex position there is exactly one edge colored red after \mathcal{R} s first move and by applying the above defense strategy \mathcal{B} can force a tie.

If S contains one inner point, then the dual graph contains one cycle. The first move of \mathcal{B} has to color an edge adjacent to the inner vertex blue and thus splits the corresponding cycle of the dual graph. Thus after \mathcal{R} s second move we can again apply the above defense strategy.

Suppose now S contains two inner points and thus two cycles in the dual graph. If the triangulation contains an edge e connecting the two inner points then in her first move \mathcal{B} colors e or one of the other four edges of the two triangles

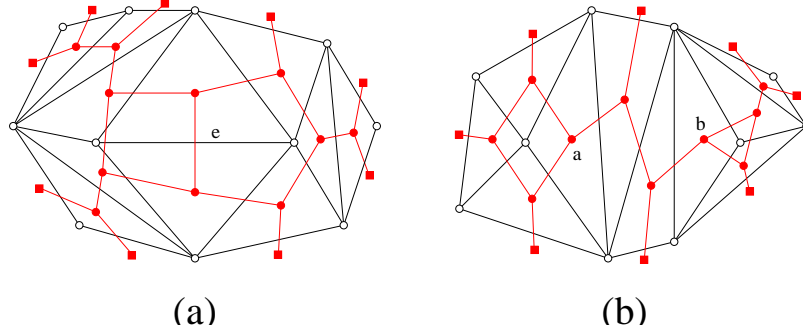


Fig. 8. Splitting two cycles of the dual graph.

adjacent to e blue. Thus both cycles are split, cf. Figure 8(a). After \mathcal{R} 's second move we can therefore again apply \mathcal{B} 's defense strategy.

So suppose now there is no such edge e , see Fig. 8(b) for the remainder. The two cycles of the dual graph, let us call them left and right cycle, respectively, are connected by a path via the vertices a and b , $a \neq b$. W.l.o.g. let \mathcal{R} 's first move be to the left of b . Then \mathcal{B} colors an edge of the triangle corresponding to a blue, splitting the left cycle. As long as \mathcal{R} continues to play to the left of b player \mathcal{B} replies (locally) according to her defense strategy. On \mathcal{R} 's first move to the right of b player \mathcal{B} colors an edge of the triangle corresponding to b . Note that at least one of the three possible edges must still be black and that afterwards not only the right cycle is split, but also \mathcal{R} 's latest edge is separated from all other red edges. Thus after \mathcal{R} 's next move we (globally) continue with the usual defense strategy.

[We are sure that the same technique can be used to obtain similar result for 3 (or more? how close to 6 can we come?) inner vertices (Can this be done without considering too many cases?). —OSWIN & FERRAN]

5.6 Four-Cycle Game

Recall game 2.3.3: players alternate coloring edges green in an initially black triangulation, and the first player to complete a green empty four-cycle wins. For any triangulation with at least two triangles, this game cannot end in a tie. In the rest of this section, we take “four-cycle” to mean an empty four-cycle, i.e., a four-cycle consisting of edges from two adjacent triangles.

We make two observations based on the idea of symmetry strategies. The general idea is to consider an edge-automorphism (i.e., a permutation of the edges) of the triangulation where the losing player is the first to play outside of the set of edges held fixed by the automorphism, since this permits the other player to mimic their play.

Theorem 6. *On a serpentine triangulation with at least two triangles, the four-cycle game is a first-player win.*

Proof. For any path, there is a natural *reversing* edge-automorphism that maps the first edge to the last, the second edge to the penultimate, and so on. We consider the following edge-automorphism induced by reversing the dual path. Triangles are mapped according the dual path reversal; within a triangle internal (resp. external) edges are mapped to internal (resp. external) edges. This definition leaves some ambiguity for the external edges of the terminal triangles on the path; either choice will work for the mimicking strategy. If the dual path has an odd number of edges, the fixed set is the triangulation edge dual to the center dual path edge. If the dual path has an even number of edges, the fixed set is the boundary edge of the center triangle. The first player's strategy is to play initially in the fixed set, and subsequently to mimic the second player's strategy by choosing the corresponding edge (image or preimage) under the automorphism. Naturally if three edges of a four-cycle are coloured at the beginning of the first player's turn, she chooses the fourth edge to win. Suppose the second player completes a four cycle at turn t . This means at turn $t - 1$ the first player played the third edge of a four-cycle. Since there is only one edge in the fixed set, this play must have been a mimicking play, i.e. the second player must have already played three edges of a four cycle, in which case the first player would have already won. \square

Theorem 7. *On a wheel triangulation with at least four vertices, the four-cycle game is second-player win.*

Proof. The edge-automorphism is not fixed in this case, but rather chosen by the second player in response to the first player's first move. We imagine taking a regular drawing of the wheel where the rim forms a regular convex n -gon and the central vertex is drawn at the centroid c of the polygon. Given such a drawing, the *antipodal* edge to an edge e is the unique edge whose interior is intersected by the line through c and the midpoint of e . Depending on the parity of the wheel, the antipodal edge to a rim (external) edge is either another rim edge or a spoke (internal) edge, and vice-versa. If the first player marks edge e in the first move, the second player replies by marking the antipodal edge. The edge-automorphism is defined by reflection about a line through the centroid such that the two marked edges reflect onto themselves (and hence form the fixed set). The rest of the argument of the proof of the previous theorem then carries forward, since the fixed set in this case has only two edges. \square

[How efficiently can a winning strategy be computed? It's certainly polynomial, but detecting whether there is a winning move (a completable 4-cycle) might take more than linear time? —ERIK] xxx

6 Flipping Games for Two Players

In the previous sections we have seen solutions for games that involve constructing and marking triangulations. We consider now two games belonging to the remaining family, transforming triangulations.

6.1 Monochromatic Flipping Game

In the Monochromatic Flipping game (2.2.1) two players start on top of a triangulation with initially black edges. Each move consists of choosing a black edge, flipping it, and coloring the new edge green. The winner is determined by normal play, and in convex position it is straightforward to see the outcome:

Theorem 8. *Monochromatic Flipping on n points in convex position is a first-player win if n is even and a second-player win if n is odd.*

Proof. All $n - 3$ diagonals in a triangulation of a convex point set are flippable. Thus the game ends after exactly $n - 3$ turns. \square

6.2 Monochromatic Flipping to Triangle Game

The same rules apply in the *Monochromatic Flipping to Triangle* game (2.2.2), MFT for short, except now the winner is who completes the first empty green triangle. Though when no green triangle is created we might end up with a tie (note that edges of the convex hull are never colored), MFT seems to be more involved and thus of much more attraction from a players point of view. We analyze the case of a serpentine triangulation of points in convex position; we remind that this means that the dual graph of the initial triangulation is a path. Notice that all diagonals are flippable in any triangulation of points in convex position.

We have the following result, which will be refined later:

Proposition 1. *MFT in a serpentine triangulation of n points in convex position is a tie or first-player win for n even, and a tie or second-player win for n odd.*

Proof. The dual of a serpentine triangulation T consists of a path π . The underlying point set is partitioned by π into two sides in an obvious way, where the two points incident to the two ears of T are assigned arbitrarily. An inner edge of T is said to cross π if and only if its end points get assigned to different sides in this way. A *playable* edge is a black edge of T crossing π . Observe that playable edges are exactly the edges which can be chosen during the game to be flipped and colored. Further note that although the triangulation T changes during the game the assignment of points remains fixed.

When performing a move with a playable edge e of T we denote the newly introduced edge by e' , see Figure 6.2. If e' crosses π then the two triangles adjacent to e form a quadrilateral such that flipping e to e' does not affect the

serpentine character of T . If e' does not cross π then e' closes exactly one triangle Δ that lies completely on one side of π . If Δ constitutes a green triangle we have a winner and the game is over. Otherwise the part of T split off by e' (including $\Delta \setminus e'$) is irrelevant for the remainder of the game since it entirely consists of non-playable edges. So we remove this part from T , and therefore T stays serpentine and includes all remaining playable edges. Note that in this case we get a green edge on the convex hull of T , namely e' .

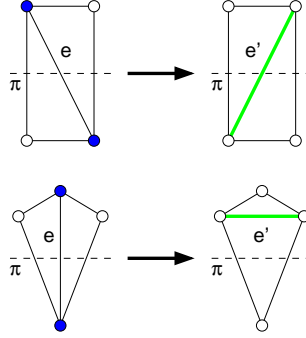


Fig. 9. Flipping does not retain the dual path.

We define a mapping from T to a string $\{b, g\}^*$ as follows: Without loss of generality fix an arbitrary direction of π and consider all edges of T crossing π in the induced ordered sequence. Assign to each of these edges a character, namely b for black (playable) and g for green edges, respectively. For a green edge on the convex hull of T we insert another g according to its position between two edges crossing π . Black edges of the convex hull are not encoded into the string.

Next consider the influences of a valid move upon the corresponding string. Let $i_1, i_2 \in \{0, 1\}$ and $s_1, s_2 \in \{b, g\}^*$. In the following strings b denotes the edge being chosen for the move and g^i is solely used to encode (possible) green edges of the convex hull. Similar as above we have to distinguish between two different types of flips. In the first case a sequence $s_1 g^{i_1} b g^{i_2} s_2$ changes to $s_1 g^{i_2} g g^{i_1} s_2 = s_1 g^{i_1+i_2+1} s_2$. The second case, when flipping to a non-crossing edge, might reduce the number of green edges on the convex hull and changes $s_1 g^{i_1} b g^{i_2} s_2$ to $s_1 g s_2$. In both cases only local changes are made, namely a single letter changes from b to g , and in some cases redundant gs are omitted.

We are now ready to analyze the game. We call a move *winning move* if it completes a green triangle. For a winning move only one edge e can be flipped and colored at a time, thus two of the edges of the resulting green triangle have been colored before and are incident to e . The crucial fact is that a winning move thus has to deal with a sub-sequence gbg , cf. Figure 6.2 for different cases. In other words there always is exactly one playable edge, corresponding to b , between two green edges. Assuming that no player makes voluntarily a bad move in a way such that the next player can make a winning move, we use this

observation to predict the winner of the game. Consider a situation where a player is forced to make a move leading to a winning move for her opponent. Since for every winning move there has to be a single b between two gs , and every move influences only one letter b , changing it to g , such a situation only arises if *each* b has exactly one twin b , i.e. is part of a sub-string $gbbg$. This implies that at this time the number of bs in the string has to be even.

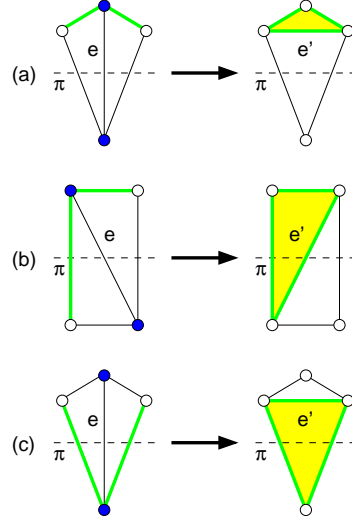


Fig. 10. Different cases of strings before playing a winning move are (a) s_1gbgs_2 , (b) $s_1gg^{i_1}bgs_2$, and (c) $s_1gg^{i_1}bg^{i_2}gs_2$.

In other words only for an odd number, say $2i - 1$, of bs (corresponding to playable edges) the previous player might have been forced to offer a winning move. Since any triangulation of a convex set of n points has $n - 3$ diagonals, all of which are initially playable, $(n - 3) - (2i - 1) = n - 2i - 2$ edges have already been played in this situation. By parity it follows that for n even the game is a tie or the first player wins and for n odd the game is a tie or a second-player win. \square

In order to improve on the above proposition, we need some technical lemmas. For a serpentine triangulation of a convex point set of size $n \geq 5$ we split its diagonals in two classes. We call a diagonal with exactly one endpoint of degree 3 a *fan edge* and all other diagonals *zig-zag edges*. Note that both endpoints of a zig-zag edge have degree at least 4. A special class of serpentine triangulations are *fan triangulations*, i.e., all of its diagonals are fan edges. In a fan triangulation all fan edges have one vertex in common.

Lemma 5. *MFT in a fan triangulation of a set S of $n \geq 8$ points is a first-player win for n even, and a second-player win for n odd.*

Proof. In a fan triangulation there exists a point $p \in S$ which is incident to all fan edges. S can be split by the dual path π in a way that all points $S \setminus p$ are opposite to p with respect to π . Flipping a fan edge e gives a new edge e' that lies completely on one side of π and closes a triangle Δ not including p . Similar to the proof of Proposition 1 we remove $\Delta \setminus e'$ s.t. the remaining triangulation stays serpentine and contains all playable edges. Note that all remaining diagonals are still fan edges. Thus all flips are of the same type throughout the whole game.

As argued before, a winning move always plays on a substring gbg in the corresponding encoding. In our situation each gbg indeed leads to a winning move because of the special type of flips. Thus a player winning according to parity just has to create a substring $gb^i g$ for $i \geq 2$, which eventually leads to gbg , to force a winning situation. For even n at least 5 flippable edges are necessary for the first player (always start e.g. with the left-most playable edge). Thus n even and $n \geq 8$ gives a first-player win. If n is odd the second player needs at least 6 flippable edges, i.e. n odd and $n \geq 9$ is a second-player win. \square

The collection of all zig-zag edges forms a path (e_1, \dots, e_m) . Each pair e_k, e_{k+1} , $1 \leq k \leq m-1$, shares exactly one common point which we call the *second point* of e_k and the *first point* of e_{k+1} . Any connected subgraph of this path is called a *zig-zag path*. A diagonal is *between two zig-zag edges* e_r and e_s if it is between e_r and e_s in the sequence of intersections with the dual path.

Lemma 6. *Consider MFT in a serpentine triangulation T of a convex point set. Let $\zeta = (e_1, \dots, e_m)$ be a zig-zag path in T s.t.*

1. m is even and $m \geq 2$,
2. e_1 and e_m are colored, and
3. all edges between e_1 and e_m are playable, and there exists at least one such edge.

Then MFT on T cannot be played as a tie.

Proof. By induction on the number of playable edges between e_1 and e_m . To start the induction we assume there is only one playable edge e between e_1 and e_m . Then $m = 2$, and e has to be a fan edge. Thus flipping e to e' gives the green triangle e_1, e_2, e' .

If there exist more than one playable edge between e_1 and e_m their number is reduced by playing any one of them. We show that afterwards we can apply induction to a zig-zag path fulfilling properties 1 to 3. There are different cases according to m and the type of flip.

Case $m = 2$. Let e be a playable edge by property 3 of ζ which can be flipped to e' . Observe that e is a fan edge. Flipping e either ends the game by closing a green triangle, or reduces the number of playable edges incident to e_1 and e_2 . We proceed by induction.

Case $m \geq 4$. There are three different cases of flips (see Figure 6.2):

- Flipping a fan edge does not change properties 1 to 3 for the remaining serpentine part of T , see the proof of Lemma 5. Apply induction.

- Flipping a zig-zag edge e_k which is not incident to e_1 or e_m , i.e. $3 \leq k \leq m-2$, alters the corresponding zig-zag path ζ to ζ' . Observe that any zig-zag path is a bipartite graph with partitions induced by its dual path π . Remove the three edges e_{k-1}, e_k, e_{k+1} from ζ . This leaves two disconnected zig-zag paths, $\zeta_1 = (e_1, \dots, e_{k-2})$ and $\zeta_2 = (e_{k+2}, \dots, e_m)$, both are part of ζ' . The second point of e_{k-2} in ζ_1 and the first point of e_{k+2} in ζ_2 are on different sides of π . Thus we need an odd number of zig-zag edges to combine ζ_1 and ζ_2 to ζ' . Therefore the parities of $|\zeta|$ and $|\zeta'|$ concur, and thus are even. Nevertheless, ζ' does not fulfill property 3 since it includes an additional non-playable edge. Removing this edge from ζ' splits ζ' in two zig-zag paths ζ'_1 and ζ'_2 . For $|\zeta'|$ is even, $|\zeta'_1|$ and $|\zeta'_2|$ are of different parity. Either $|\zeta'_1|$ or $|\zeta'_2|$ is even, and thus meets properties 1 to 3. Therefore induction can be applied.
- Without loss of generality consider flipping zig-zag edge e_2 to e'_2 , changing ζ to ζ' . Similarly as above, remove e_1, e_2, e_3 from ζ . This time only one zig-zag path $\zeta_1 = (e_4, \dots, e_m)$ of odd size remains, which is part of ζ' . Observe that e'_2 is also a zig-zag edge in ζ' . In addition, the second point of e'_2 in ζ' and the first point of e_4 in ζ' are on the same side of π . Using the same arguments as above, an even number of zig-zag edges has to lie between e_2 and e_4 in ζ' . Consider the part $\zeta'_1 = (e'_2, \dots, e_4, \dots, e_m)$ of ζ' . Observe that $|\zeta'_1|$ has to be even, and apply induction to it.

□

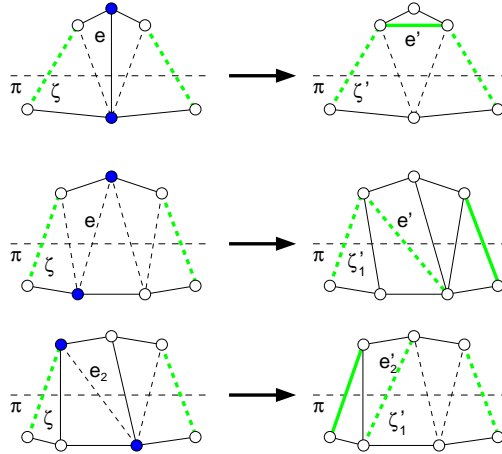


Fig. 11. Different cases of flips. Dashed edges show the zig-zag paths.

We are now ready for proving the theorem which improves on Proposition 1.

Theorem 9. *There is a constant N such that MFT in a serpentine triangulation of $n \geq N$ points in convex position is a first-player win for n even, and a second-player win for n odd.*

Proof. First consider the case for even n , i.e. MFT is a tie or a first-player win according to Proposition 1. If the serpentine triangulation T includes a fan of at least 5 flippable edges we know by Lemma 5 that the first player can force a win. Note that playing incident zig-zag edges to the fan does not interfere with the fan. Otherwise, each fan in T consists of at most 4 flippable edges. Thus there are $O(n)$ zig-zag edges. For n large enough, there exists a gadget such that the first player can win the game as described in Lemma 6. The argument for n odd is similar. \square

[How efficiently can the winning strategy be computed? I assume xxx polynomial time... —ERIK]

7 Solitaire Games

In this section we consider flipping games with only one player.

7.1 All-Green Solitaire

In each move, the player flips a flippable black edge e of $T(S)$; then e becomes green, as do the four boundary edges of the enclosing quadrilateral. The goal of the game is to color all edges green. This is not always possible, as can be seen from the example of a triangulated regular pentagon. More generally, we can ask the harder question of how to maximize the number of edges colored green. Our next result settles this question in the convex case:

Theorem 10. *The maximum number of edges that can be colored green by flips in the All-Green Solitaire Game can be computed in $O(n)$ time, and the sequence of flips can be computed in $O(n)$ time. In particular, in $O(n)$ time we can compute whether the player can win the All-Green Solitaire Game and if so find a winning strategy.*

Proof. Let S be the triangulated subpolygon to the right of a given oriented diagonal d . There are two diagonals d_1 and d_2 in S , that form a triangle together with d , which we orient leaving d to their left, as shown in Figure 12. Let us denote by S^1 and S^2 the subpolygons these diagonals define (we follow the counterclockwise order). When the notation is iterated we write simply $S^{i,j}$ instead of $(S^i)^j$.

Define $n(S)$ to be the number of edges that can be colored green, by flipping edges strictly interior to S , assuming all of these edges are initially black. We can compute $n(S)$ by a simple recursion. In the base case, S is a single triangle and thus $n(S) = 0$. In the general case, among the edges d_1 and d_2 , we can either flip d_1 , or flip d_2 , or flip neither. Whether we flip these edges now or later does

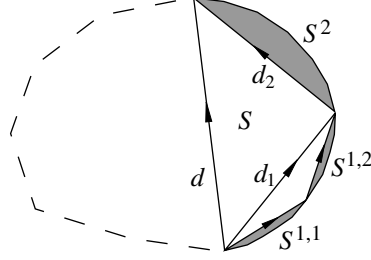


Fig. 12. Illustrating the technique for the All-Green Solitaire Game in convex position.

not matter, and we cannot flip both. If we flip one, we color 5 edges green, but prevent further flipping of these edges. Thus, $n(S)$ is given by the recursion

$$n(S) = \max \left\{ \underbrace{n(S^{1,1}) + n(S^{1,2}) + n(S^2) + 5}_{\text{flip } d_1}, \underbrace{n(S^{2,1}) + n(S^{2,2}) + n(S^1) + 5}_{\text{flip } d_2}, \underbrace{n(S^1) + n(S^2)}_{\text{flip neither}} \right\}.$$

By tracing back through this recursion, we can obtain an optimal strategy in $O(n)$ time. \square

7.2 Green-Wins Solitaire

Suppose the rules are the same as All-Green Solitaire, but the goal of the game is to obtain at the end more green edges than black edges. It is an open question whether this can always be done. In our next result we give worst-case bounds on how many green edges we can always guarantee:

Theorem 11. *The player of the Green-Wins Solitaire Game can obtain from any given triangulation on n points at least $1/6$ of the edges to be green at the end of the game. There are triangulated point sets such that no sequence of flips of black edges colors green more than $5/9$ of the edges. (In the above fractions we ignore additive constants.)*

Proof. For the lower bound, it is known that any triangulation of n points contains at least $\frac{n-4}{6}$ independently flippable edges, in the sense that no two of them are sides of the same triangle [8]. Each one of these edges will color 5 edges by its flip. A green non-flipped edge might get counted twice this way, thus we get at least $\frac{n-4}{6} + \frac{4}{2} \times \frac{n-4}{6} = \frac{n-4}{2}$ colored edges. As there are at most $3n$ edges, and we have colored at least $n/2$ edges (we disregard additive constants for both numbers), we have got at least $1/6$ of the edges to be green, as claimed.

As for the upper bound, we define a triangulated convex polygon $C(t)$ as follows (see Figure 13). The vertices of $C(t)$ are placed on an arc of circle with

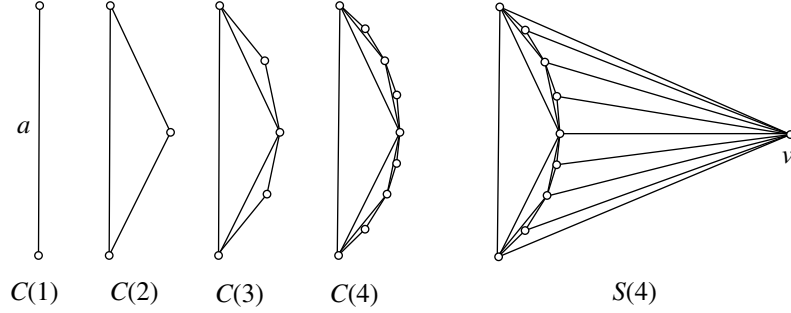


Fig. 13. Recursive construction of triangulated point sets $S(t)$

central angle below π . Take $C(1)$ equal to the chord a associated with the arc, add a triangle with the third vertex in the arc in order to get $C(2)$. Attach externally triangles to the two outer chords of $C(2)$ for constructing $C(3)$, and iterate this process in order to obtain $C(t)$. The number of vertices of $C(t)$ is $2^{t-1} + 1$. Now let v be a point that sees completely the circular arc from outside the circle; a triangulated point set $S(t)$ is defined by connecting v to all the vertices of $C(t)$. The edges in $S(t)$ are those in the boundary of $C(t)$, plus its diagonals, plus the edges incident to v , therefore their total number is

$$e(t) = (2^{t-1} + 1) + (2^{t-1} + 1 - 3) + (2^{t-1} + 1) = 3 \cdot 2^{t-1}.$$

Notice that no boundary edge of $C(t)$ and no edge incident to v can ever be flipped, therefore the edges in $S(t)$ incident to v are never colored green. On the other hand observe that if we suppress a from $S(t)$ we obtain two instances of $S(t-1)$; despite the fact that the two copies share an edge, the coloring process behaves independently.

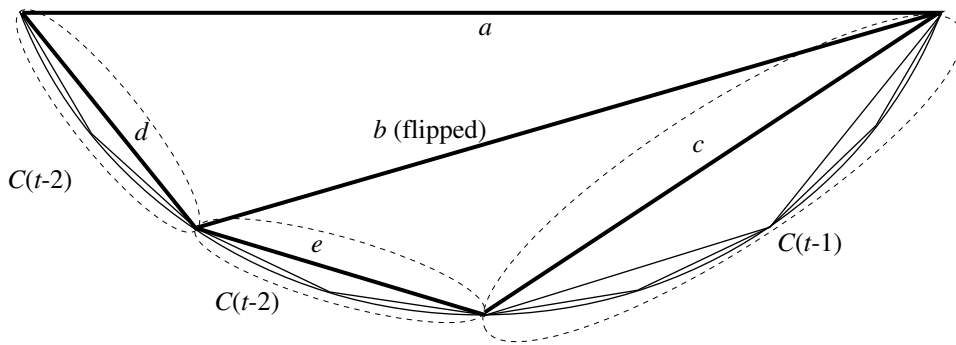


Fig. 14. Bounding the value $g_2(t)$. The five solid lines are green once b has been flipped

Let $g(t)$ be the maximum number of green edges that can be obtained from $S(t)$ after any flip sequence of black edges; we show next that $g(t)/e(t)$ approaches $5/9$ for large values of t . The numbers $g(1) = g(2) = 0$ and $g(3) = 5$ are directly computable. Let b and c be the edges that together with a form a triangle in $S(t)$. As $S(t)$ contains two copies of $S(t-1)$ which we can color independently, we have

$$2 \cdot g(t-1) \leq g(t). \quad (1)$$

Let $g_1(t)$ be the maximum number of green edges achievable from $S(t)$ when neither b nor c are flipped. We have

$$g_1(t) = 2 \cdot g(t-1). \quad (2)$$

Let $g_2(t)$ be the maximum number of green edges achievable from $S(t)$ when either b or c are flipped (refer to Figure 14). Assume it is b , and notice that in this case none of the edges d , e and c has been flipped before b is, as otherwise b would be green and disallowed to be flipped. Therefore the maximum number of green edges we can achieve this way is

$$g_2(t) \leq g(t-1) + 2 \cdot g(t-2) + 5. \quad (3)$$

Combining the above equation with the fact that $2 \cdot g(t-2) \leq g(t-1)$ that we know from (1), we obtain $g_2(t) \leq 2 \cdot g(t-1) + 5$. On the other hand the equality (2) directly gives that $g_1(t) \leq 2 \cdot g(t-1) + 5$. Hence we have

$$g(t) = \max(g_1(t), g_2(t)) \leq 2 \cdot g(t-1) + 5, \quad (4)$$

and from this we get

$$g_1(t) \leq g(t-1) + g(t-1) \leq g(t-1) + 2 \cdot g(t-2) + 5. \quad (5)$$

Using equations (3) and (5) we arrive to

$$g(t) = \max(g_1(t), g_2(t)) \leq g(t-1) + 2 \cdot g(t-2) + 5,$$

a recursive inequality which solves to

$$g(t) \leq \frac{5}{6} \cdot (2^t - (-1)^t - 3).$$

Therefore

$$\frac{g(t)}{e(t)} \leq \frac{5}{6} \cdot \frac{2^t - (-1)^t - 3}{3 \cdot 2^{t-1}} \xrightarrow{t \rightarrow \infty} \frac{5}{9}$$

as claimed. \square

Finally, it is quite easy to prove that the game is not very exciting in convex position:

Theorem 12. *The player of the Green-Wins Solitaire Game can always win for any given triangulation on $n \geq 4$ points in convex position.*

Proof. The number of edges in any triangulation is $2n - 3$, therefore we have to prove that we can always achieve at least $n - 1$ green edges after a suitable sequence of flips.

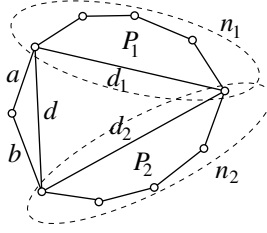


Fig. 15. Illustrating a winning strategy for the Green-Wins Solitaire in convex position

We proceed by induction. The cases $n = 4, 5, 6$ are easily checked directly, hence we can assume $n \geq 7$. Let a and b be consecutive boundary edges of an ear of the triangulation, and let d be the diagonal which completes a triangle with a and b (refer to Figure 15). Let d_1 and d_2 be the edges of the other triangle which shares the diagonal d . Consider the polygons P_1 and P_2 respectively separated by these diagonals from the whole polygon, and let n_1 and n_2 be their respective number of vertices, where $n_1 + n_2 = n$. We can assume that $n_1 \leq n_2$.

If $n_1 = 2$, we flip d and apply induction to P_2 . In this way we obtain at least $4 + (n_2 - 1)$ green edges, and

$$4 + (n_2 - 1) = 4 + (n - 2) - 1 = n + 1 > n - 1$$

as desired. If $n_1 = 3$, we proceed in the same way and obtain at least $4 + (n_2 - 1)$ green edges. Now

$$4 + (n_2 - 1) = 4 + (n - 3) - 1 = n > n - 1.$$

Finally, if $n_2 \geq n_1 \geq 4$, we flip d and apply induction both to P_1 and P_2 . In this way we obtain at least $3 + (n_1 - 1) + (n_2 - 1)$ green edges, where

$$3 + (n_1 - 1) + (n_2 - 1) = n_1 + n_2 + 1 = n + 1 > n - 1.$$

□

8 Conclusion

We have analyzed a variety of games on various classes of triangulations. Of course, many open problems remain, involving more general classes of triangulations, and other games on triangulations. One open problem in the opposite extreme is to obtain negative results: NP-hardness or PSPACE-completeness of some of our games on general triangulations.

Acknowledgments

Ferran Hurtado is partially supported by Projects MEC-DGES-SEUID PB98-0933, MCYT-FEDER BFM2002-0557 and Gen. Cat 2001SGR00224. Research of Evangelos Kranakis is supported in part by NSERC and MITACS grants. Research of Hannes Krasser is supported by the FWF (Austrian Fonds zur Förderung der Wissenschaftlichen Forschung). Suneeta Ramaswami is partially supported by a Rutgers University ISATC pilot project grant. Research of Jorge Urrutia is supported in part by CONACYT grant 37540-A and grant PAPIIT.

References

1. E. R. Berlekamp: The Dots and Boxes Game: Sophisticated Child's Play. Academic Press (1982)
2. E. R. Berlekamp, J. H. Conway, R. K. Guy: Winning Ways for your Mathematical Plays. Academic Press (1982). Second edition in print, A K Peters Ltd., (2001)
3. H. L. Boedlander, D. Kratsch: Kayles and Nimbers. *J. of Algorithms* **43** (2002) 106–119
4. J. H. Conway: On Numbers and Games. Academic Press (1976). Second edition, A K Peters Ltd. (2002)
5. D. G. Corneil, S. Olariu, L. Stewart: Asteroidal triple-free graphs. *SIAM J. Discrete Math.* **10** (1997) 399–430
6. E. D. Demaine: Playing games with algorithms: Algorithmic combinatorial game theory. In: Proc. 26th Symp. on Math Found. in Comp. Sci., Lect. Notes in Comp. Sci., Springer-Verlag (2001) 18–32
7. A. S. Fraenkel: Combinatorial Games: Selected Bibliography with a Succinct Gourmet Introduction. *Electronic Journal of Combinatorics*, <http://www.wisdom.weizmann.ac.il/~fraenkel>
8. J. Galtier, F. Hurtado, M. Noy, S. Pérennes, J. Urrutia: Simultaneous edge flipping in triangulations. To appear in *International Journal of Computational Geometry and Applications*.
9. T. J. Schaefer: On the complexity of some two-person perfect-information games. *J. Comput. Syst. Sci.* **16** (1978) 185–225