Playing with Triangulations

Oswin Aichholzer[1], David Bremner[2], Erik D. Demaine[3],
Ferran Hurtado[4], Evangelos Kranakis[5], Hannes Krasser[6],
Suneeta Ramaswami[7], Saurabh Sethia[8], and Jorge Urrutia[9]

[1] IICM - Softwaretechnology, Graz University of Technology,
Inffeldgasse 16b/I A-8010 Graz, Austria
oaich@igi.tu-graz.ac.at
[2] Faculty of Computer Science, University of New Brunswick,
P.O. Box 4400 Fredericton, NB E3B 5A3, Canada
bremner@unb.ca
[3] Laboratory for Computer Science, Massachusetts Institute of Technology,
200 Technology Square, Cambridge, MA 02139, USA
edemaine@mit.edu
[4] Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya,
Pau Gargallo,5, 08028 Barcelona, Spain
hurtado@ma2.upc.es
[5] School of Computer Science, Carleton University,
1125 Colonel By Drive, Ottawa, Ontario K1S 5B6, Canada
kranakis@scs.carleton.ca
[6] Institute for Theoretical Computer Science, Graz University of Technology,
Inffeldgasse 16b/I, A-8010 Graz, Austria
hkrasser@igi.tu-graz.ac.at
[7] Computer Science Department, Rutgers University,
321 Business & Science Building, Camden, NJ 08102, USA
rsuneeta@crab.rutgers.edu
[8] Department of Computer Science, Oregon State University,
Corvallis, OR 97331-3202, USA
saurabh@cs.orst.edu
[9] Instituto de Matemáticas, Universidad Nacional Autónoma de México,
Área inv. científ., Circ. Ext., Ciudad Univ., Coyoacán 04510, México D.F., México
urrutia@math.unam.mx

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. In various situations, we develop polynomial-time algorithms to determine who wins a given game under optimal play, and to find a winning strategy. Along the way, we show connections to existing combinatorial games such as Kayles.

1 Introduction

Let S be a set of n points in the plane, which we assume to be in general position, i.e., no three points of S lie on the same line. A *triangulation* of S is a simplicial decomposition of its convex hull having S as vertex set.

In this work we consider several games involving the vertices, edges (straight-line segments) and faces (triangles) of some triangulation. We present games where 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 (maybe the single one) 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, the last move or possibly some counting decides then who is the winner.
- *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, 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.

For each of the variety of games described in Section 3, we are interested in characterizing who wins the game, and designing efficient algorithms to determine the winner and compute a winning strategy. We present several such results in Sections 4–6 to give a taste of the area, and leave further details to the companion paper [1].

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.

2 Combinatorial Games

Games on triangulations belong to the more general area of *combinatorial games* which typically involve two players, $\mathcal{R}$ and $\mathcal{B}$. We define next a few more terms from combinatorial game theory that we will use in this paper. For more information, refer to the books [3, 5] and the survey [7]. The paper [8] contains a list of more than 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.

Most 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. [3, 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.

3 Examples of Games

We describe next the rules of several specific games that we have studied.

3.1 Constructing

3.1.1 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 triangle(s), it is (they are) given to this player 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.
3.1.2 Monochromatic Triangle. Starts as in 3.1.1, but has a different stopping condition: the first player who completes one empty triangle is the winner.
3.1.3 Bichromatic Complete Triangulation. As in 3.1.1, but the two players use red and blue edges. Only monochromatic triangles count.
3.1.4 Bichromatic Triangle. As in 3.1.2, but with red and blue edges. The first empty triangle must be monochromatic.

3.2 Transforming

3.2.1 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.
3.2.2 Monochromatic Flipping to Triangle. Same rules as for 3.2.1, except now the winner is who completes the first empty green triangle.
3.2.3 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.
3.2.4 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.
3.2.5 Green-Wins Solitaire. As in 3.2.4, but the goal of the game is to obtain more green edges than black edges.

3.3 Marking

3.3.1 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.
3.3.2 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.
3.3.3 Four-Cycle Game. Same as 3.3.1 but the goal is to get an empty quadrilateral.
3.3.4 Nimstring Game. *Nimstring* is a game defined in Winning Ways [3] as a special case of the classic children's (but nonetheless deep) combinatorial game Dots and Boxes [2, 3]. In the context of triangulations, players in Nimstring alternate *marking* one-by-one the edges of a given triangulation (i.e., coloring

green an edge, initially black), and whenever a triangle has all three of its edges marked, the completing player is awarded an extra move and must move again. 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.

4 Metamorphosis of Some Games

In this section we show that some of the above games are equivalent to famous combinatorial games and describe their solutions.

4.1 Triangulation Coloring Game

Obviously the Triangulation Coloring Game (3.3.1) 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.

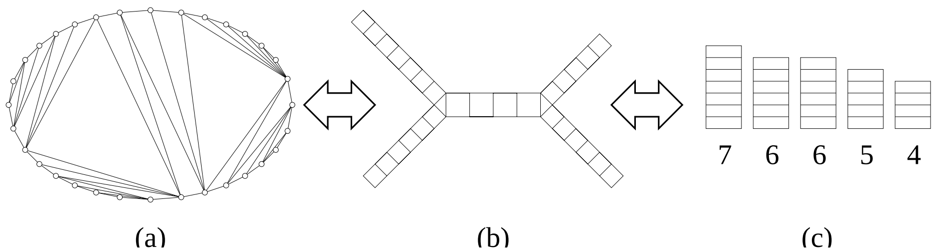

Fig. 1. Different incarnations of the triangulation coloring game

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*, i.e., its dual is a simple path (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 [3]). 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 the 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 only restriction that no two inner triangles share a common diagonal; see for example Figure 1(a). We say that in this case the corresponding triangulation is a *simple-branching* triangulation. 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* [3]. 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 [3, 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 2 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 4.1. To compute the Nim value for a game with several heaps, we can xor-add up the Nim values (given by Table 4.1) 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 1. Nim values for Kayles: the Nim sequence has periodicity 12 and there are 14 exceptional numbers

$\mathcal{G}(n) = K[n \text{ modulo } 12]$, $K[0, \dots, 11] = (4, 1, 2, 8, 1, 4, 7, 2, 1, 8, 2, 7)$
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 [10].

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 $EG(T)$.

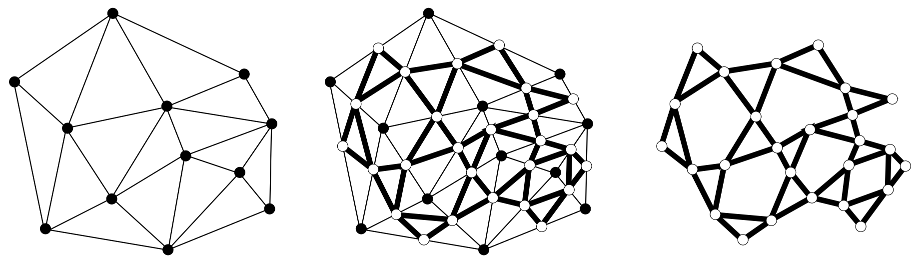

Fig. 2. A triangulation T of a point set (left) and the graph $EG(T)$ associated to adjacent edges (right)

While such a reduction does not prove hardness of the Triangulation Coloring Game, it does transfer any solutions to special cases of Kayles. In [4] 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.

4.2 Monochromatic Triangle

We present an optimal strategy for Monochromatic Triangle (3.1.2) 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 [3]. 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 $n_1, \ldots, n_k$. 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 2. Nim values for Dawson's Kayles: the Nim sequence has periodicity 34 and there are 8 exceptional numbers

$\mathcal{G}(n) = K[n \text{ modulo } 34]$, $K[0, \ldots, 33] =$
$(4, 8, 1, 1, 2, 0, 3, 1, 1, 0, 3, 3, 2, 2, 4, 4, 5, 5, 9, 3, 3, 0, 1, 1, 3, 0, 2, 1, 1, 0, 4, 5, 3, 7)$
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 4.2. 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 4.2 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.*

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 Monochromatic Complete Triangulation

In this section we consider the triangulation-construction game 3.1.1. 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 $e_2 = e_0 - 1$. □

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. □

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. □

6 Solitaire Games

In this section we consider games in which there is only one player.

6.1 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.

Notice that this is not always possible, as can be seen from the example of a triangulated regular pentagon. Two questions are of interest: (1) Characterize classes of triangulations for which it is possible; (2) give an efficient algorithm to find an appropriate flip sequence for such triangulations. Our next result settles the second question in the convex case:

Theorem 4. *Whether the player can win the All-Green Solitaire Game for a given triangulation of n points in convex position can be decided in time $O(n)$. When the player can win, a winning sequence of moves can be found within the same time bound.*

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 3. 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 $(S^i)^j$.

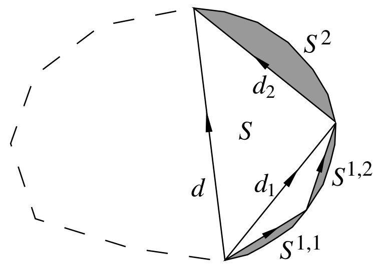

Fig. 3. Illustrating the technique for the All-Green Solitaire Game in convex position

Let us define the logic value $b(S)$ to be true if, and only if, it is possible to color green all the edges of the polygon S, by flipping black diagonals strictly interior to S, assuming that all edges in S, including d, are initially black. Let the logic value $g(S)$ to be true if, and only if, it is possible to color green all the edges of the polygon S, by flipping black diagonals strictly interior to S, assuming that all edges in S are initially black with the only exception of d, which is initially green. Notice that $b(S)$ implies $g(S)$. In the simplest case S consists of merely a single boundary edge (d), then $b(S)$ is false and $g(S)$ is true.

We use the symbols "$\wedge$" and "$\vee$" for the logic operators "and" and "or", respectively. If d is initially black, at some point we have to flip either d_1 or d_2 to give the color green to d. Say it is d_1, then none of the boundary edges of the quadrilateral that has d_1 as diagonal may have been flipped before, therefore in order to complete a green coloration we must have $g(S^{1,1}) \wedge g(S^{1,2}) \wedge g(S^2)$. Hence

$$b(S) = \left[g(S^{1,1}) \wedge g(S^{1,2}) \wedge g(S^2)\right] \vee \left[g(S^{2,1}) \wedge g(S^{2,2}) \wedge g(S^1)\right] .$$

For $g(S)$ to be true it is obviously sufficient that $b(S)$ is true; if $b(S)$ is false we must have both $b(S_1)$ and $b(S_2)$ to be true, and use the additional fact that d is initially green. Therefore we have

$$g(S) = [b(S_1) \wedge b(S_2)] \vee b(S) =$$
$$= [b(S_1) \wedge b(S_2)] \vee \left[g(S^{1,1}) \wedge g(S^{1,2}) \wedge g(S^2)\right] \vee \left[g(S^{2,1}) \wedge g(S^{2,2}) \wedge g(S^1)\right] .$$

Let us call *jump* of an oriented diagonal d the number of convex hull edges that d has on its right side. For example, if d is an edge at the convex hull the jump is 0 or $n-1$ when it is oriented counterclockwise and clockwise, respectively.

Notice that in the above formulas for $b(S)$ and $g(S)$ these values depend on similar values associated to diagonals with strictly smaller jump. Therefore if we process the polygons to the right of the oriented diagonals ordered by increasing jump, the computation of $g(S)$ and $b(S)$ is carried out in constant time for a fixed S, and in time $O(n)$ for the whole triangulation.

Once the computation is complete, we can consider in turn each diagonal e of the polygon as a tentative candidate for starting the All-Green Solitaire Game. This would color green e and the four boundary edges e_1, e_2, e_3 and e_4, of the enclosing quadrilateral, oriented counterclockwise. If we call P_1, P_2, P_3 and P_4 the polygons to the right of these diagonals, we would be able to complete a green coloring if, and only if, the four values $g(S_i)$ are true. As this is checked in constant time after our precomputation, the overall exploration can be done in time $O(n)$.

If a first move in a winning sequence is found, the same technique can easily be adapted to explore down the recurrence of $b(S)$ and $g(S)$ and find the entire sequence of winning moves in $O(n)$ additional time. Basically, depending on the case that caused either $b(S)$ or $g(S)$ to be true, you flip the appropriate edge to reduce to the appropriate subcases and recurse down the dual tree. In this way, $O(1)$ time is spent per triangle / dual node. For example, if $g(S)$ was set true because $g(S^{1,1}) \wedge g(S^{1,2}) \wedge g(S^2)$, then one should flip d_1 and then recurse in the subproblems $g(S^{1,1})$, $g(S^{1,2})$, and $g(S^2)$. □

6.2 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 bounds on how many green edges we can always guarantee.

Theorem 5. *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 provides more than 5/9 of the edges to be green at the end. (In the above fractions we don't pay attention to 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 [9]. 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.

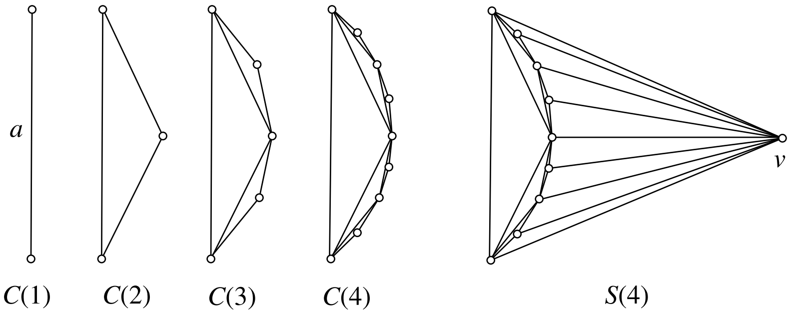

Fig. 4. Recursive construction of triangulated point sets $S(t)$

As for the upper bound, we define a triangulated convex polygon $C(t)$ as follows (see Figure 4). The vertices of $C(t)$ are placed on an arc of circle with 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.

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

$$2 \cdot g(t-1) \leq g(t). \tag{1}$$

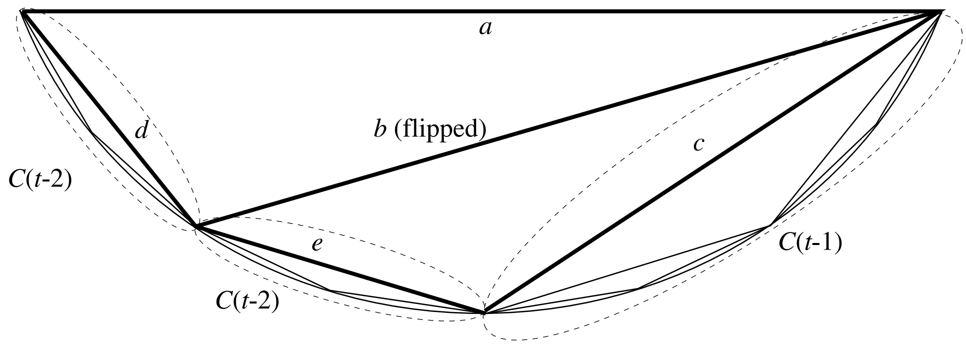

Fig. 5. Bounding the value $g_2(t)$. The five solid lines are green once b has been flipped

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

$$g_1(t) = 2 \cdot g(t-1). \tag{2}$$

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 5). 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

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

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

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

and from this we get

$$g_1(t) \leq g(t-1) + g(t-1) \leq g(t-1) + 2 \cdot g(t-2) + 5. \tag{5}$$

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

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

as claimed. □

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

Theorem 6. *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.

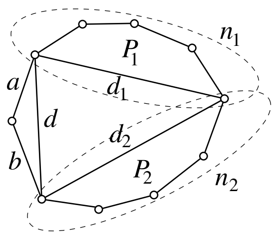

Fig. 6. Illustrating a winning strategy for the Green-Wins Solitaire in convex position

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 6). 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

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

□

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. O. Aichholzer, D. Bremner, E. D. Demaine, F. Hurtado, E. Kranakis, H. Krasser, S. Ramaswami, S. Sethia, J. Urrutia: Games on Triangulations: Several Variations. Manuscript in preparation
2. E. R. Berlekamp: The Dots and Boxes Game: Sophisticated Child's Play. Academic Press (1982)
3. 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)
4. H. L. Boedlander, D. Kratsch: Kayles and Nimbers. J. of Algorithms **43** (2002) 106–119
5. J. H. Conway: On Numbers and Games. Academic Press (1976). Second edition, A K Peters Ltd. (2002)
6. D. G. Corneil, S. Olariu, L. Stewart: Asteroidal triple-free graphs. SIAM J. Discrete Math. **10** (1997) 399–430
7. 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
8. A. S. Fraenkel: Combinatorial Games: Selected Bibliography with a Succinct Gourmet Introduction. Electronic Journal of Combinatorics, `http://www.wisdom.weizmann.ac.il/~fraenkel`
9. J. Galtier, F. Hurtado, M. Noy, S. Pérennes, J. Urrutia: Simultaneous edge flipping in triangulations. Submitted
10. T. J. Schaefer: On the complexity of some two-person perfect-information games. J. Comput. Syst. Sci. **16** (1978) 185–225