Games of No Chance 3
MSRI Publications
Volume **56**, 2009

New results in loopy games

AARON N. SIEGEL

ABSTRACT. We strengthen the usual notion of simplest form for stoppers and show that under the stronger definition, equivalence coincides with graph-isomorphism. We then show that the game graph of a canonical stopper contains no 2- or 3-cycles, but may contain n-cycles for all $n \geq 4$.

We also introduce several new methods for simplifying games γ whose graphs contain alternating cycles. These include a generalization of dominated and reversible moves.

1. Introduction

A *loopy game* is a combinatorial game in which repetition is permitted. The history and basic theory of loopy games are discussed in [Siegel 2009]. In this article we focus on two fundamental problems left unresolved by *Winning Ways*.

Long irreducible cycles. The first problem concerns the cycles that appear in the game graph of a stopper. Conway showed that every stopper s admits a simplest form [Conway 1978], so one would expect that certain cycles are intrinsic to the play of s. All canonical stoppers discussed in *Winning Ways* are *plumtrees*: their graphs contain only 1-cycles. It is therefore natural to ask whether longer canonical cycles are possible, and to attempt to characterize the structure of such cycles.

Conway defined the *simplest form* of s to be a representation with no dominated or reversible moves. This is not quite strong enough for our purposes, as illustrated by the example t shown in Figure 1. Certainly t has no dominated or reversible options, but it is easy to check that $t = \textbf{on}$. Thus while t technically has a 2-cycle, it is *reducible* in the sense that t has an alternate representation with just a 1-cycle.

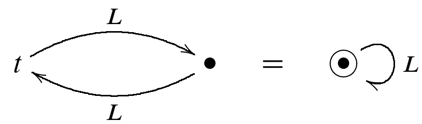

Figure 1. A 2-cycle that reduces to **on**.

In Section 3 of this paper, we introduce a stronger notion of simplicity, the *graph-canonical form* of a stopper. We show that if s and t are stoppers in graph-canonical form and $s = t$, then s and t have isomorphic game graphs. Then in Section 4, we investigate the types of cycles that can appear in a graph-canonical stopper s. We show that every such cycle of length $n > 1$ must contain at least two edges of each color. This rules out 2-cycles and 3-cycles; however, we give examples of graph-canonical stoppers with n-cycles for all $n \geq 4$.

Simplification of alternating cycles. The second problem concerns the simplification of games with alternating cycles. If γ is an arbitrary loopy game, it is desirable to know whether γ is stopper-sided, and if so to compute its sides. Previously, this problem was addressed by the technique known as *sidling* [Berlekamp et al. 2001; Conway 1978; Moews 1996], which produces a sequence of approximations to the sides of γ. If the sidling sequences converge, then they necessarily converge to the sides of γ, but there are many important cases in which they fail to converge.

In Section 5, we introduce generalizations of dominated and reversible options that apply to arbitrary loopy games. These can be used to obtain useful simplifications of γ^+ and γ^-. Often, the simplified forms are already stoppers, even in cases where sidling fails. In addition, the new methods are computationally more efficient than sidling procedures.

Finally, the Appendix (page 228) describes algorithms for comparing arbitrary games. A simplification engine can be built on these algorithms by using the techniques of Section 5. All of these algorithms and techniques have been implemented in *CGSuite* (see http://www.cgsuite.org/), with important applications to the analysis of actual games (see [Siegel 2009] for further discussion).

2. Preliminaries

We assume the reader is familiar with the theory of loopy games, as presented in *Winning Ways*. Sufficient background can be obtained from [Siegel 2009] in this volume. We briefly summarize some of the most relevant facts.

We denote loopy games by Greek letters $\gamma, \delta, \alpha, \beta, \ldots$. If every infinite play of γ is drawn, then γ is said to be *free*, and this will be the assumption when nothing is said to the contrary. When γ is free, we denote by γ^+ and γ^- the matching games with draws redefined as wins for Left and Right, respectively.

Infinite play in a sum $\alpha + \beta + \cdots + \gamma$ is assumed to be drawn unless the same player wins on every component in which play is infinite. In particular, if γ and δ are free, then the following are equivalent:

(i) $\gamma^+ \le \delta^+$;
(ii) Left can survive $\delta^+ - \gamma^+$ playing second;
(iii) Left, playing second in $\delta - \gamma$, can guarantee that *either* he gets the last move, *or* infinitely many moves occur in the δ component.

(i) $\iff$ (ii) by the definition of $\le$, and (ii) $\iff$ (iii) by the definition of sum (and the fact that $-\gamma^+ = (-\gamma)^-$). (iii) is a key characterization, and it will be used repeatedly in the proofs and algorithms that follow.

If s and t are *stoppers*, then the following conditions are all equivalent:

$$s \le t; \quad s^+ \le t^+; \quad s^- \le t^-; \quad s^- \le t^+; \quad \text{Left can survive } t - s \text{ playing second.}$$

Finally, throughout this paper we will assume that all games have a *finite* number of positions. Some results generalize to games with infinitely many positions; but it is usually clear when this is the case, and since the generalization will not be needed it is simpler to keep things finite.

Strategies. Often we will know that Left can survive some game γ and wish to show that he can survive a closely related game γ'. (For example, γ' might be obtained by eliminating a dominated option of γ.) In the loopfree case, this is typically handled by examining relationships between the followers of γ. However, when γ is loopy, altering the options of γ might also affect the structure of its followers. Because of this interdependence, we will usually need to take a global view of the structure of γ, and here it is useful to reason in terms of strategies.

DEFINITION 1. Let γ be a loopy game and let $\mathscr{A}$ denote the set of followers of γ. A *Left strategy for* γ is a partial mapping $S : \mathscr{A} \to \mathscr{A}$ such that, whenever $\delta \in \mathscr{A}$ has a Left option, then $S(\delta)$ is defined, and $S(\delta) =$ some δ^L.

We refer to $S(\delta)$ as the move *recommended by* S.

DEFINITION 2. Let S be a Left strategy for γ. S is a *first-player survival (winning) strategy* if Left, playing first from γ, survives (wins) every line of play in which he plays according to S.

Right strategies and second-player strategies are defined analogously. We say that Left (Right) survives (wins) γ playing first (second) if there exists an appropriate strategy.

DEFINITION 3. Let S be a Left strategy for γ. S is a *complete survival strategy* if, for each $\delta \in \mathscr{A}$ that Left can survive as first player, he can survive δ by playing according to S.

Note that a complete survival strategy recommends good moves from *every* follower of γ, even those that would never be encountered if γ itself were played according to S. Complete survival strategies always exist; this can be established by "pasting together" survival strategies.

LEMMA 4. *Let γ be any loopy game. Then there exists a complete Left survival strategy for γ.*

PROOF. First we inductively construct a sequence of strategies S_n, as follows. Let S_0 be a first-player Left survival strategy for any subposition γ_0 from which Left has a survival move. Given S_n and γ_n, let A_n be the set of positions that can be reached, with Left to move, by some line of play proceeding from γ_n, throughout which Left plays according to S_n. If $\bigcup_{i \le n} A_i$ contains every follower of γ from which Left has a survival move, then stop. Otherwise, choose any $\gamma_{n+1} \notin \bigcup_{i \le n} A_i$ from which Left has a survival move, and let S_{n+1} be the corresponding first-player Left survival strategy. Now define a strategy S by

$$S(\delta) = S_n(\delta) \text{ where } n \text{ is least such that } \delta \in A_n.$$

($S(\delta)$ may be chosen arbitrarily if $\delta \notin A_n$ for any n.) We claim that S is a complete Left survival strategy for γ.

To see this, let δ be some follower of γ from which Left has a survival move, and suppose Left plays δ according to S. Let $\delta = \delta_0, \delta_1, \delta_2, \ldots$ be the consecutive positions reached with Left to move (so $\delta_{i+1} = (S(\delta_i))^R$ for each i). We first show that Left has a survival move from each δ_i. This is obviously true for δ_0. For the inductive step, let n be least such that $\delta_i \in A_n$. Then $S(\delta_i) = S_n(\delta_i)$; since S_n is a survival strategy for δ_i, and $\delta_{i+1} = (S_n(\delta_i))^R$, Left has a survival move from δ_{i+1}.

If play is finite, we are done: Left must have made the last move. Otherwise, consider any δ_i, and let n be least such that $\delta_i \in A_n$. Since δ_{i+1} is reached from δ_i by play according to S_n, we also have $\delta_{i+1} \in A_n$. It follows that, for some n_0 and i_0, we have

$$S(\delta_i) = S_{n_0}(\delta_i) \text{ for all } i \ge i_0.$$

Since S_{n_0} is a survival strategy for δ_{i_0}, and the outcome does not depend on any finite initial segment of moves, Left has survived. □

Graphs. Throughout this paper, a *graph* will be a directed graph with separate Left and Right edge sets. We will use calligraphic letters $\mathcal{G}$, $\mathcal{H}$, ... to denote graphs.

Just as every game has an associated graph, we can define games by specifying a graph and a start vertex. Given a graph $\mathcal{G}$ and a vertex v of $\mathcal{G}$, let $\mathcal{G}|v$ be the graph obtained by removing from $\mathcal{G}$ all vertices not reachable from v. Denote by $\mathcal{G}_v$ the free game whose graph is $\mathcal{G}|v$ and whose start vertex is v. Note that a game is not the same as its graph; this distinction will often be essential. Thus when we write $\mathcal{G}_u = \mathcal{G}_v$, we mean that $\mathcal{G}_u$ and $\mathcal{G}_v$ are game-theoretically equal in the sense of the usual order-relation, whereas $u = v$ means that u and v represent the exact same vertex. Clearly $u = v$ implies $\mathcal{G}_u = \mathcal{G}_v$, but the converse certainly need not be true.

DEFINITION 5. A path directed from u to v is an *alternating path* if its edges alternate colors. The path is *Left-alternating* or *Right-alternating* if the first edge out of u is blue or red, respectively. An *alternating cycle* is an alternating path of even length that starts and ends at the same vertex. We say that an edge is *cyclic* if it belongs to an alternating cycle, and a graph is *alternating cycle-free* if it contains no alternating cycles (equivalently, no cyclic edges).

Note that s is a stopper if and only if its graph is alternating cycle-free.

If u and v are vertices of a graph $\mathcal{G}$, we write $u \xrightarrow{L} v$ to indicate that $\mathcal{G}$ has a Left edge directed from u to v; likewise $u \xrightarrow{R} v$ indicates a Right edge. We sometimes write $e : u \xrightarrow{L} v$ to mean that e is the (unique) Left edge directed from u to v.

3. Fusion

Recall the *simplest form theorem* for stoppers [Berlekamp et al. 2001; Conway 1978; Siegel 2009]:

THEOREM 6 (SIMPLEST FORM THEOREM). *Let s and t be stoppers. Assume that $s = t$, and that neither s nor t has any dominated or reversible options. Then for every s^L there is a t^L with $s^L = t^L$, and vice versa; and likewise for Right options.*

If s and t satisfy this criterion along with all their followers, then they are equivalent in play. However, their graphs might still differ fundamentally. Consider the two examples s and t shown in Figure 2. $s = t =$ **over**, and neither game has any dominated or reversible options, but their representations are clearly different.

A further simplification solves this problem. Suppose s is a stopper whose game graph contains two equivalent vertices, u and v, and assume that no followers of s have any dominated or reversible options. Then we can replace u

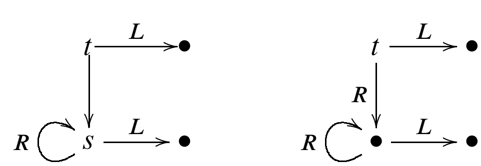

Figure 2. Two forms of **over**.

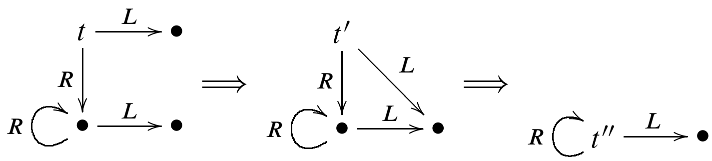

Figure 3. Fusion further simplifies stoppers.

and v with a single vertex, redirecting edges as appropriate, without changing the value of s or any of its followers. Repeated application of this "fusion" process ultimately produces a game with no two equivalent vertices, and this representation is unique up to graph isomorphism. In the example above, t can be reduced to s with two applications of fusion, as illustrated in Figure 3.

LEMMA 7 (FUSION LEMMA). *Let $\mathcal{G}$ be alternating cycle-free, with no dominated or reversible edges. Suppose u, v are two distinct vertices of $\mathcal{G}$ and $\mathcal{G}_u = \mathcal{G}_v$. Let $\mathcal{H}$ be the graph obtained by deleting v and replacing every edge $a \to v$ with an edge $a \to u$ of the same color. Then $\mathcal{H}$ is alternating cycle-free and $\mathcal{G}_w = \mathcal{H}_w$ for every vertex $w \neq v$.*

A cautionary note: fusion might fail when s is not a stopper, or when s is a stopper but is not in simplest form. Figure 4 gives an example: $\gamma = \delta = 2\ \&\ 0$, but if we fuse δ to γ, then the resulting vertex has value $3\ \&\ 0$.

PROOF OF LEMMA 7. First we show that $\mathcal{H}$ is alternating cycle-free. Assume instead (for contradiction) that $\mathcal{H}$ contains an alternating cycle. We can assume the cycle involves a redirected edge, since otherwise it would already be present in $\mathcal{G}$. So the cycle involves u, and we can assume without loss of generality that it is Left-alternating out of u. We will construct a sequence $(v_n)_{n=0}^{\infty}$ of vertices of $\mathcal{G}$ such that for all n, $\mathcal{G}_{v_n} = \mathcal{G}_{v_{n+1}}$ and there is an even-length Left-alternating path from v_n to v_{n+1}.

Let $v_0 = u$, $v_1 = v$. Since $\mathcal{H}$ contains a Left-alternating cycle out of u that involves a redirected edge, $\mathcal{G}$ must contain an even-length Left-alternating path from u to v. This establishes the base case.

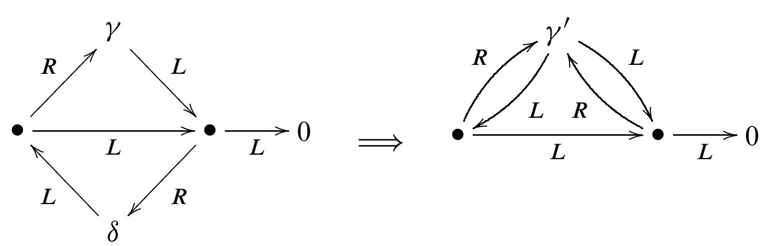

Figure 4. An example where fusion fails.

Now given v_n and v_{n+1}, we construct v_{n+2} as follows. We know that v_{n+1} is a Left-alternating follower of v_n. But $\mathcal{G}_{v_{n+1}} = \mathcal{G}_{v_n}$, so by repeated application of the Simplest Form Theorem, there is a Left-alternating follower v_{n+2} of v_{n+1} satisfying $\mathcal{G}_{v_{n+2}} = \mathcal{G}_{v_{n+1}}$. Since the path from v_n to v_{n+1} has even length, so does the path from v_{n+1} to v_{n+2}. This defines $(v_n)_{n=0}^{\infty}$.

But $\mathcal{G}$ is finite, so there must be some $m < n$ with $v_m = v_n$. It follows that there is an alternating cycle in $\mathcal{G}$ involving v_n, contradicting the assumption that $\mathcal{G}$ is alternating cycle-free. This shows that $\mathcal{H}$ is alternating cycle-free.

Next fix w, and let $s = \mathcal{G}_w$, $t = \mathcal{H}_w$. We wish to show that $s = t$. Since both are stoppers, it suffices to show that Left, playing second, never runs out of moves in $s - t$ or $t - s$. We will prove the $s - t$ case; the proof for $t - s$ is similar.

Let S be a complete Left survival strategy for $s - s$. Define the strategy S' for $s - t$ as follows: S' is equivalent to S except when S recommends a move from $\mathcal{G}_a - \mathcal{G}_b$ to $\mathcal{G}_a - \mathcal{G}_v$. In that case, S' recommends a move from $\mathcal{G}_a - \mathcal{H}_b$ to $\mathcal{G}_a - \mathcal{H}_u$. We claim that S' is a second-player Left survival strategy for $s - t$. To see this, note that whenever $\mathcal{G}_a \geq \mathcal{G}_v$, then also $\mathcal{G}_a \geq \mathcal{G}_u$. Since S is a complete survival strategy, this implies that if Left plays second from $s - t$, then any position $\mathcal{G}_a - \mathcal{H}_b$ reached according to S' will satisfy $\mathcal{G}_a \geq \mathcal{G}_b$. Therefore Left, playing according to S', will never run out of moves. This completes the proof. □

DEFINITION 8. A stopper s is said to be in *graph-canonical form* if s is in simplest form and $\mathcal{G}_u \neq \mathcal{G}_v$ for any two vertices $u \neq v$ of s.

THEOREM 9. *Suppose s, t are stoppers in graph-canonical form with $s = t$. Then the game graphs of s and t are isomorphic.*

PROOF. Let $s = \mathcal{G}_u$, $t = \mathcal{H}_v$. For every vertex a of $\mathcal{G}$, we know that there is a vertex b of $\mathcal{H}$ with $\mathcal{G}_a = \mathcal{H}_b$, and vice versa. ($b$ can be obtained by repeated application of the Simplest Form Theorem.) Since $\mathcal{G}$ and $\mathcal{H}$ contain no equivalent vertices, it follows that there is a bijection $f : V(\mathcal{G}) \to V(\mathcal{H})$ with $f(u) = v$ such that $\mathcal{G}_a = \mathcal{H}_{f(a)}$ for all vertices a of $\mathcal{G}$.

To see that f is a graph-homomorphism, suppose $\mathcal{G}$ contains a Left edge $a \xrightarrow{L} a'$. Write $b = f(a)$, so that $\mathcal{G}_a = \mathcal{H}_b$. Since $\mathcal{G}_a \leq \mathcal{H}_b$, Right has a survival response from $\mathcal{G}_{a'} - \mathcal{H}_b$. It cannot be to any $\mathcal{G}_{a'}^R$, since this would imply that

$$\mathcal{G}_a^{LR} = \mathcal{G}_{a'}^R \leq \mathcal{H}_b = \mathcal{G}_a,$$

contradicting the assumption that $\mathcal{G}$ contains no reversible moves. So $\mathcal{G}_{a'} \leq \mathcal{H}_{b'}$ for some vertex b' of $\mathcal{H}$ with $b \xrightarrow{L} b'$.

Now since $\mathcal{G}_a \geq \mathcal{H}_b$, Left has a survival response from $\mathcal{G}_a - \mathcal{H}_{b'}$. It cannot be to any $\mathcal{H}_{b'}^R$, since (as above) this would imply that $\mathcal{H}_b \geq \mathcal{H}_b^{LR}$, contradicting the assumption that $\mathcal{H}$ has no reversible moves. So $\mathcal{G}_a^L \geq \mathcal{H}_{b'}$ for some $\mathcal{G}_a^L$. Thus $\mathcal{G}_a^L \geq \mathcal{H}_{b'} \geq \mathcal{G}_{a'}$, and since $\mathcal{G}$ contains no dominated options, $\mathcal{G}_a^L = \mathcal{H}_{b'} = \mathcal{G}_{a'}$. Therefore $f(a') = b'$, so $\mathcal{H}$ contains a Left edge $f(a) \xrightarrow{L} f(a')$. The proof for Right edges is identical. □

4. Long irreducible cycles

In this section, we show that if s is a stopper in graph-canonical form, then every cycle in s of length greater than one must contain at least two edges of each color. In particular, s contains no 2- or 3-cycles. Longer cycles are possible, however: the game τ shown in Figure 5 is in graph-canonical form and has a 4-cycle. Soon we will see that there exist graph-canonical stoppers t with n-cycles for all $n \geq 4$. Such cycles are *irreducible* in the sense that any representation of t must contain at least an n-cycle.

DEFINITION 10. Let $\mathcal{G}$ be a graph. A cycle in $\mathcal{G}$ is *long* if it contains at least two edges. A cycle in $\mathcal{G}$ is *monochromatic* if all edges in the cycle are the same color; *bichromatic* otherwise.

LEMMA 11. *Let s be a stopper in graph-canonical form. Then s contains no long monochromatic cycles.*

PROOF. By symmetry, it suffices to prove the lemma for cycles consisting entirely of blue edges. So let $s_0, s_1, \ldots, s_n$ be a sequence of subpositions of s, with $s_{i+1} = s_i^L$ for $0 \leq i < n$ and $s_0 = s_n$. We will show that

$$s_0 \leq s_1 \leq s_2 \leq \cdots \leq s_n = s_0,$$

so in fact all subpositions in the sequence must be equivalent.

Left's survival strategy for $s_{i+1} - s_i$ is simple. As long as Right moves around the cycle in the $-s_i$ component, Left does the same in s_{i+1}, staying one move ahead of her. This continues until Right chooses to break the cycle. At that point the position must be either $s_{j+1}^R - s_j$ or $s_{j+1} - s_j^{L'}$ ($s_j^{L'} \neq s_{j+1}$), for some j. In the first case, we have

$$s_{j+1}^R = s_j^{LR},$$

and since s_j has no reversible options, this implies that $s_{j+1}^R \not\leq s_j$. So Left must have a winning move from $s_{j+1}^R - s_j$. Likewise, in the second case, we have

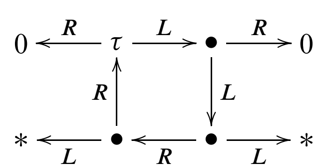

Figure 5. A stopper that is not equivalent to any plumtree.

$s_{j+1} = s_j^L$, and since s_j has no dominated options, this implies that $s_{j+1} \not\leq s_j^{L'}$. So again Left has a winning move; and we have shown that he can survive any line of play.

This shows that each $s_i \leq s_{i+1}$, and hence

$$s_0 = s_1 = s_2 = \cdots = s_n. \qquad \square$$

LEMMA 12. *Let s be a stopper in graph-canonical form. Then s contains no long cycles with just a single red edge.*

PROOF. Toward a contradiction, let s_0, s_1, …, s_n $(n \geq 2)$ be a sequence of subpositions of s, with $s_{i+1} = s_i^L$ for $0 \leq i < n$ and $s_0 = s_n^R$. We first show that

$$s_0 \leq s_1 \leq s_2 \leq \cdots \leq s_{n-1}. \qquad (\dagger)$$

To show that $s_i \leq s_{i+1}$, we proceed just as in the previous lemma; the only difference occurs when Right has moved to the position $s_n - s_n$. Then Left responds by playing to $s_n - s_0$. If Right continues to $s_0 - s_0$, then Left plays to $s_1 - s_0$ and resumes moving around the cycle as before; while if Right makes any other move, then the absence of any dominated or reversible options hands the win to Left, as in Lemma 11.

This proves (†), so in particular $s_0 \leq s_{n-1}$. But $s_0 = s_n^R = s_{n-1}^{LR}$, contradicting the assumption that s_{n-1} has no reversible moves. This completes the proof. $\square$

By symmetry, if s is a stopper in graph-canonical form, then s contains no long cycles with just a single blue edge. Therefore every long cycle in s must include at least two edges of each color.

Unicycles

DEFINITION 13. A stopper s is said to be a *unicycle* provided that:

(i) The graph of s has just one cycle; and
(ii) Each position on the cycle has just two options: a move to the next position on the cycle, and a move for the *other* player to a loopfree game.

We say that s is an *n-unicycle* if its cycle is an n-cycle.

For example, τ (Figure 5) is a 4-unicycle. In fact, there exist n-unicycles for all $n \geq 4$. Figure 6 gives an elegant example for all $n \geq 6$, in which 0 is the only loopfree subposition. Figure 7 is an interesting 13-unicycle: 0 and $*$ are the only loopfree subpositions; furthermore, the cycle is alternating except for the single pair of consecutive Left edges. The 13-unicycle generalizes to a $(4n+1)$-unicycle for all $n \geq 1$ (in particular, this gives an example of a 5-unicycle).

We can classify unicycles more precisely by considering the specific sequence of blue and red edges associated to each cycle. For example, τ has the pattern LLRR. Then a *P-unicycle* is a unicycle whose cycle matches the pattern P.

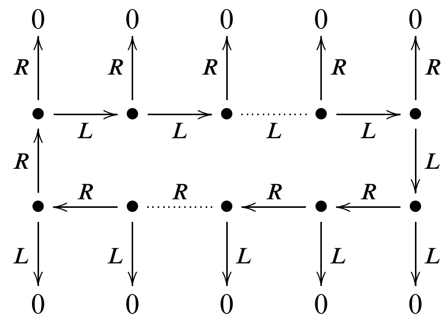

Figure 6. A particularly elegant n-unicycle ($n \geq 6$). It is assumed that there are *at least three* blue edges and *at least three* red edges in the cycle, though there need not be equally many of each color.

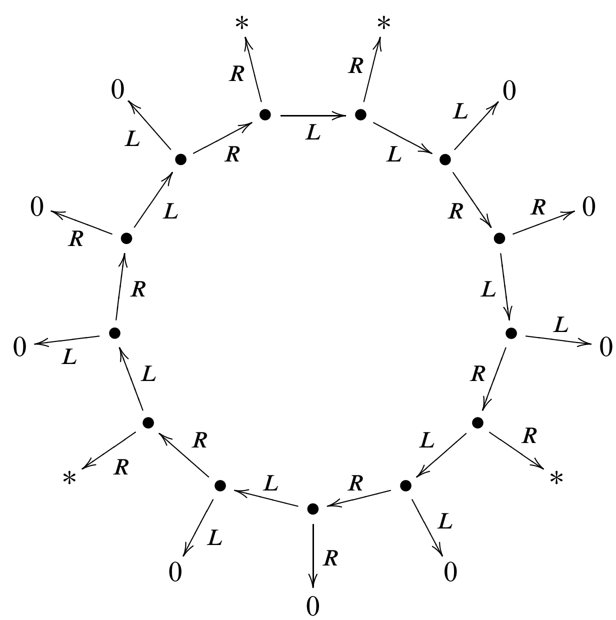

Figure 7. An "almost alternating" 13-unicycle. This generalizes to a $(4n+1)$-unicycle by continuing the pattern: three exits to 0 followed by one exit to $*$.

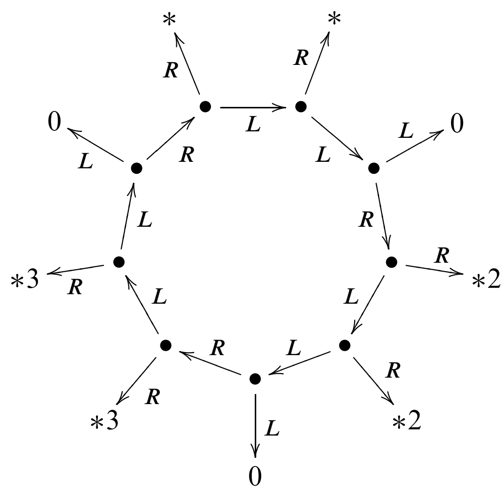

Figure 8. A 9-unicycle whose pattern cannot be realized if the exits are restricted to 0, $*$, and $*2$.

By Lemmas 11 and 12, we know that if there exists a P-unicycle, then P must have at least two edges of each color. Furthermore, P cannot be strictly alternating, since every unicycle is a stopper. As it turns out, these are the only restrictions up to length 9: if P has at most nine edges and meets both restrictions, then there exists a P-unicycle whose loopfree subpositions are all nimbers. $P =$ LLRLLRLLR is an interesting example: Figure 8 gives a P-unicycle with exits to 0, $*$, $*2$ and $*3$, but there are no P-unicycles with exits restricted to 0, $*$ and $*2$ (or any other combination of just three nimbers). All of these facts can be verified using *CGSuite*.

The same is true for patterns of length 10, with one possible exception: $Q =$ LLLRLRRRLR. It appears that there are no Q-unicycles whose exits are restricted to nimbers. However, if exits to arbitrary loopfree games are allowed, then the question remains open.

OPEN PROBLEM. Determine the patterns P for which there exists a P-unicycle. In particular, is there an LLLRLRRRLR-unicycle?

Note that the *number* of patterns of length n is equal to the number of directed binary necklaces of length n. This is sequence A000031 in Sloane's encyclopedia (http://www.research.att.com/~njas/sequences/) and is given by

$$\frac{1}{n}\sum_{d|n} 2^{n/d}\varphi(d),$$

where φ is the Euler phi-function.

5. Simplification of alternating cycles

This section introduces a suitable generalization of dominated and reversible moves to games with alternating cycles. All of the results are stated in terms of γ^+, but of course they dualize to γ^-.

DEFINITION 14. Let γ be a free loopy game. Then:

(a) A Left option γ^L is said to be *onside-dominated* if $(\gamma^{L'})^+ \geq (\gamma^L)^+$ for some other $\gamma^{L'}$.
(b) A Right option γ^R is said to be *onside-dominated* if $(\gamma^{R'})^+ \leq (\gamma^R)^+$ for some other $\gamma^{R'}$ such that no alternating cycle contains the edge $\gamma \xrightarrow{R} \gamma^{R'}$.
(c) A Right option γ^R is said to be *onside-reversible* if $(\gamma^{RL})^+ \geq \gamma^+$ for some γ^{RL}.
(d) A Left option γ^L is said to be *onside-reversible* if $(\gamma^{LR})^+ \leq \gamma^+$ for some γ^{LR} such that no alternating cycle contains the edges $\gamma \xrightarrow{L} \gamma^L \xrightarrow{R} \gamma^{LR}$.

The additional constraints in Definitions 14(b) and (d) are necessary, as demonstrated by examples such as Bach's Carousel [Berlekamp et al. 2001]. Of course, the point of these definitions is the following Lemma.

LEMMA 15. *Let γ be a free loopy game and let δ be any follower of γ. Suppose γ' is obtained from γ by either*:

(a) *Eliminating some onside-dominated option of δ*; *or*
(b) *Bypassing some onside-reversible option of δ.*

Then $\gamma^+ = (\gamma')^+$.

PROOF. We prove the lemma for onside-dominated Right options and onside-reversible Left options; the remaining cases are easier.

(a) Suppose that $(\delta^{R'})^+ \leq (\delta^R)^+$ and γ' is obtained by eliminating $\delta \xrightarrow{R} \delta^R$. Clearly $\gamma^+ \leq (\gamma')^+$, so we must show that $\gamma^+ \geq (\gamma')^+$. Let S be a complete Left survival strategy for $\gamma^+ - \gamma^+$, and define S' as follows: S' is identical to S, except that any recommendation from $-\delta^+$ to $-(\delta^R)^+$ is replaced by a recommendation to $-(\delta^{R'})^+$.

If Left plays according to S', then since S is a complete survival strategy and $(\delta^{R'})^+ \leq (\delta^R)^+$, the position $\alpha^+ - \beta^+$ reached after Left's move will always satisfy $\alpha^+ \geq \beta^+$. Therefore Left never runs out of moves. To complete the proof, we need to show that the play, if infinite, was not ultimately confined to the negative component. So assume that play was infinite. First suppose that Left was forced to deviate only finitely many times from S. Then after a finite initial sequence of moves, Left followed the survival strategy S. Therefore there must have been infinitely many plays in the positive component.

But by the assumptions of Definition 14, δ is *not* a Left-alternating follower of the dominating option $\delta^{R'}$. Thus between any two deviations from S, there must occur at least one play in the positive component. So if Left deviated infinitely many times from S, then again, infinitely many plays must have occurred in the positive component. This shows that S' is also a Left survival strategy for $\gamma^+ - \gamma^+$. Since S' never makes use of the edge $\delta \to \delta^R$, it also suffices for $\gamma^+ - (\gamma')^+$. This completes the proof.

(b) Suppose that $(\delta^{LR})^+ \leq \delta^+$ and γ' is obtained by bypassing δ^L through δ^{LR}. Let S be a complete Left survival strategy for $\gamma^+ - \gamma^+$, and consider the game $\gamma^+ - (\gamma')^+$. Note that whenever Left can survive some $\alpha^+ - \delta^+$, then $\alpha^+ \geq \delta^+ \geq (\delta^{LR})^+$, so he can also survive $\alpha^+ - (\delta^{LR})^+$. Thus he has a survival response to each $\alpha^+ - (\delta^{LRL})^+$. It follows that Left never runs out of moves if he simply plays $\gamma^+ - (\gamma')^+$ according to S. But each time Right plays from $-\delta^+$ to some $-(\delta^{LRL})^+$, the assumptions of Definition 14 guarantee a move in the positive component before the next time $-\delta^+$ is reached. By an argument similar to (a), S suffices as a Left survival strategy for $\gamma^+ - (\gamma')^+$.

To complete the proof, we must define a second-player Left survival strategy S' for $(\gamma')^+ - \gamma^+$. Let S' be identical to S, except at positions of the form $(\delta')^+ - \beta^+$, where δ' is the subposition of γ' corresponding to δ. Then there are two cases.

Case 1: If Left has a survival move from $(\delta^{LR})^+ - \beta^+$, then let

$$S'\big((\delta')^+ - \beta^+\big) = S\big((\delta^{LR})^+ - \beta^+\big).$$

That is, S' makes the same recommendation from $(\delta')^+ - \beta^+$ that S makes from $(\delta^{LR})^+ - \beta^+$. This is always valid, by definition of bypassing a reversible move.

Case 2: Otherwise, we have $(\delta^L)^+ \not\geq \beta^+$, so Left's move from $\delta^+ - \beta^+$ to $(\delta^L)^+ - \beta^+$ is losing, and therefore S does *not* recommend it (except possibly when *every* Left move from $\delta^+ - \beta^+$ is losing). In this case, S' simply follows the recommendation given by S.

If Left plays $(\gamma')^+ - \gamma^+$ according to S', then he never runs out of moves. As before, to complete the proof we must show that the play, if infinite, was not ultimately confined to the negative component. The proof is much the same as in (a): we show that each deviation from S must have been followed by a play in the positive component.

But Left only deviates from S at *Case 1* positions of the form $(\delta')^+ - \beta^+$. Until some move is made in the positive component, Left's plays in $-\beta^+$ are identical to those recommended by S from $(\delta^{LR})^+ - \beta^+$. Since *Case 1* states that Left can survive from $(\delta^{LR})^+ - \beta^+$, and since S is a complete survival strategy, this implies that some move must eventually occur in the positive component. □

We can also generalize the Fusion Lemma.

LEMMA 16 (GENERALIZED FUSION LEMMA). *Let $\mathcal{G}$ be an arbitrary graph. Suppose u, v are two distinct vertices of $\mathcal{G}$ with $\mathcal{G}_u^+ = \mathcal{G}_v^+$, and assume that there is no alternating path from u to v of even length. Let $\mathcal{H}$ be the graph obtained by deleting v and replacing every edge $a \to v$ with an edge $a \to u$ of the same color. Then $\mathcal{G}_w^+ = \mathcal{H}_w^+$ for every vertex $w \neq v$.*

SKETCH OF PROOF. The proof is similar to that of Lemma 15, so we just sketch it. In playing $\mathcal{G}_a^+ - \mathcal{H}_b^+$ (or $\mathcal{H}_a^+ - \mathcal{G}_b^+$), Left follows a fixed strategy for $\mathcal{G}_a^+ - \mathcal{G}_b^+$, moving to $-\mathcal{H}_u^+$ ($\mathcal{H}_u^+$) whenever a move to $-\mathcal{G}_v^+$ ($\mathcal{G}_v^+$) is recommended. The assumptions on u and v ensure that fusion introduces no "new" alternating cycles, so two deviations in the negative component imply an intervening move in the positive one. □

Appendix:
Algorithms for comparing games

The most basic computational task is the comparison of games, since comparisons form the basis for all simplifications. When G and H are loopfree, a straightforward recursion can determine whether $G \le H$. Where loopy games are concerned, the situation is more complicated. Recall that if s and t are stoppers, then $s \le t$ just if Left, playing second, can survive $t - s$. There is a relatively simple algorithm for testing this condition. If $s = \mathcal{G}_u$ and $t = \mathcal{H}_v$, then the basic idea is to determine those vertices of the direct sum $\mathcal{G} \oplus \mathcal{H}$ from which Right can force a win. Since this might depend on who has the move, we consider separately the pairs (A, L) and (A, R), where A is a vertex of $\mathcal{G} \oplus \mathcal{H}$; we will refer to such pairs as *states*. It is convenient to define an associated *state graph*:

DEFINITION 17. Let $\mathcal{G}$ be a game graph. Then the *state graph* $\mathcal{S}$ of $\mathcal{G}$ is the (monochromatic) directed graph defined as follows. The vertices of $\mathcal{S}$ are pairs (A, L) and (A, R), where A is a vertex of $\mathcal{G}$. Its edges are constituted as follows:

- $\mathcal{S}$ contains an edge $(A, L) \to (B, R)$ if and only if $\mathcal{G}$ contains a Left edge $A \to B$.
- $\mathcal{S}$ contains an edge $(A, R) \to (B, L)$ if and only if $\mathcal{G}$ contains a Right edge $A \to B$.
- $\mathcal{S}$ contains no edges $(A, L) \to (B, L)$ (or $(A, R) \to (B, R)$), for any A, B.

When we speak of predecessors, successors or outedges of a state (A, X), we mean predecessors, successors or outedges of (A, X) in the state graph.

Begin by marking as LOSING all states (A, L) with no successors. Then iteratively:

- Mark as LOSING all states (A, R) with a LOSING successor.
- Mark as LOSING all states (A, L) from which *all* successors are marked LOSING.

Stop when no further vertices can be marked.

Algorithm 1. Comparing stoppers.

The algorithm for comparing stoppers is summarized as Algorithm 1. Starting from those states (A, L) with no successors, the states of $\mathcal{G} \oplus \mathcal{H}$ from which Right can force a win are iteratively identified. Then $s \le t$ just if $(u \oplus v, R)$ is unmarked: if Right can win from $u \oplus v$, then he can do so in n moves, for some n; but then $(u \oplus v, R)$ will be marked on the n-th stage of the iteration.

This idea is not new. Three decades ago, Fraenkel and Perl [1975] gave a similar procedure for determining the $\mathscr{P}$- and $\mathscr{N}$-positions of an impartial loopy game. The partisan version of the algorithm was introduced several years later by Shaki [1979]. It was rediscovered independently and brought to my attention by Michael Albert (personal communication, 2004).

The algorithm can be refined to guarantee that each state is examined at most once per outedge. The improved version is summarized as Algorithm 2. A huge advantage of this refinement is that it allows substantial prunings. Traversing the states "top-down," and stopping as soon as a winner is determined, yields significant time savings when prunings are desirable. Note that wins for *both* players are determined, and not just for Right; occasionally this will quickly identify Left as the winner and permit an early pruning.

Comparing general games. If γ, δ are arbitrary loopy games, then the comparison process is substantially more difficult. Recall that $\gamma \le \delta$ if and only if Left, playing second, can survive both $\delta^+ - \gamma^+$ and $\delta^- - \gamma^-$; see [Siegel 2009]. For clarity, and since the two cases are exactly symmetric, we consider just $\delta^+ - \gamma^+$.

Now Left survives $\delta^+ - \gamma^+$ if and only if either

(a) he gets the last move, or
(b) infinitely many plays occur in δ.

Thus if play is infinite, but is entirely confined to the $-\gamma^+$ component, then Left has lost. We can eliminate condition (a) from consideration by first applying the stopper-comparison algorithm (Algorithm 2); the remaining task is to identify those states from which Right can keep the play indefinitely in $-\gamma^+$.

The solution is to make several passes through the graph. At the start of each pass, some states will already be marked as a WIN FOR R, and the goal

Visit each state (A, X) *at most once* (in any order) and perform the following steps:

(1) Mark (A, X) VISITED.

(2)

- If any successor of (A, X) is already marked as a WIN FOR X, then mark (A, X) as a WIN FOR X.
- If every successor of (A, X) is already marked as a WIN FOR Y $(Y \neq X)$, then mark (A, X) as a WIN FOR Y.

(3) If we just marked (A, X) as a win for either player, then examine each predecessor (B, Y) of (A, X) such that

- (B, Y) is marked VISITED; and
- the winner of (B, Y) has not been determined.

If we marked (A, X) as a WIN FOR Y, then immediately mark (B, Y) as a WIN FOR Y. If we marked (A, X) as a WIN FOR X, then rescan the successors of (B, Y), and if they are all marked as a WIN FOR X, then mark (B, Y) as a WIN FOR X.

If this determines the winner of (B, Y), repeat step 3 with (B, Y) in place of (A, X).

Algorithm 2. Comparing stoppers, refined.

is to identify new ones. Now suppose that, from some state (A, X), Right can guarantee that *either* a state marked WIN FOR R will be reached, *or* no further plays will ever occur in δ. Clearly (A, X) must be a WIN FOR R as well. Call a state BAD if it meets this test; GOOD otherwise. During each pass through the graph, we first identify all GOOD states, and then mark each BAD state as a WIN FOR R. The algorithm terminates when a pass completes with no new states identified as a WIN FOR R.

The procedure for identifying GOOD states is straightforward. For example, suppose that for some state (A, L), there exists an outedge *in* δ to a state that is not known to be a WIN FOR R. Then (A, L) can be marked GOOD immediately. The GOOD markers can then be back-propagated just as WIN markers were in the stoppers case.

In the worst case, each pass would identify just one GOOD state, so the algorithm is ostensibly $O(|V| \cdot |E|)$, where $|V|$ is the number of vertices and $|E|$ the number of edges in the state graph. In practice, however, more than a few passes are rarely necessary, and the algorithm is effectively $O(|E|)$.

The algorithm is summarized in detail as Algorithm 3.

First execute Algorithm 2 to identify states from which one of the players can force a win in finite time. Then:

(1) Visit each vertex A *at most once* and perform the following steps.

(a) If the winner of (A, L) is not yet determined, and either:

- (A, L) has an outedge *in* δ to a state whose winner is not yet determined; or
- (A, L) has an outedge to a state marked GOOD,

then mark (A, L) GOOD.

(b) If the winner of (A, R) is not yet determined, and every successor of (A, R) *in* γ is marked either WIN FOR L or GOOD, then mark (A, R) GOOD.

(c) If either of the previous steps caused a state (A, X) to be marked GOOD, then examine all γ-predecessors (B, Y) of (A, X) such that:

- (B, Y) is marked VISITED; and
- The winner of (B, Y) is not yet determined; and
- (B, Y) is not marked GOOD.

If $Y = L$, then immediately mark (B, Y) GOOD. If $Y = R$, then rescan the γ-successors of (B, Y), and if they are all marked either WIN FOR L or GOOD, then mark (B, Y) GOOD.
If this causes (B, Y) to be marked GOOD, then repeat step 1(c) with (B, Y) in place of (A, X).

(2) Visit each state (A, X) a second time and perform the following steps:

(a) If the winner of (A, X) is not yet determined, and (A, X) is *not* marked GOOD, then mark (A, X) as a WIN FOR R.

(b) If the previous step caused a state (A, X) to be marked as a WIN FOR R, then examine all VISITED predecessors (B, Y) of (A, X) whose winner is not yet determined.
If $Y = R$, then immediately mark (B, Y) as a WIN FOR R. If $Y = L$, then rescan the successors of (B, Y), and if they are all marked as a WIN FOR R, then mark (B, Y) as a WIN FOR R.
If this determines the winner of (B, Y), then repeat step 2(b) with (B, Y) in place of (A, X).

(3) Clear all VISITED and GOOD markers. If the previous step caused any new states to be marked as a WIN FOR R, then repeat starting with step 1. Otherwise, stop.

Algorithm 3. Testing whether Left can survive $\delta^{+} - \gamma^{+}$.

Acknowledgement

I thank Michael Albert for his considerable assistance in developing the algorithms presented in this appendix.

References

[Berlekamp et al. 2001] E. R. Berlekamp, J. H. Conway, and R. K. Guy, *Winning Ways for Your Mathematical Plays*, Second ed., A. K. Peters, Ltd., Natick, MA, 2001.

[Conway 1978] J. H. Conway, "Loopy games.", pp. 55–74 in *Advances in Graph Theory*, edited by B. Bollobás, Ann. Discrete Math. **3**, 1978.

[Fraenkel and Perl 1975] A. S. Fraenkel and Y. Perl, "Constructions in combinatorial games with cycles.", pp. 667–699 in *Infinite and Finite Sets, Vol. 2*, edited by A. Hajnal et al., Colloq. Math. Soc. János Bolyai **10**, North-Holland, 1975.

[Moews 1996] D. J. Moews, "Loopy games and Go.", pp. 259–272 in *Games of No Chance*, edited by R. J. Nowakowski, MSRI Publications **29**, Cambridge University Press, New York, 1996.

[Shaki 1979] A. Shaki, "Algebraic solutions of partizan games with cycles.", *Math. Proc. Cambridge Philos. Soc.* **85**:2 (1979), 227–246.

[Siegel 2009] A. N. Siegel, "Coping with cycles", 2009. In this volume.

AARON N. SIEGEL
aaron.n.siegel@gmail.com