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

Coping with cycles

AARON N. SIEGEL

ABSTRACT. *Loopy games* are combinatorial games in which repetition is permitted. The possibility of nonterminating play inevitably raises difficulties, and several theories have addressed these by imposing a variety of assumptions on the games under consideration. In this article we survey some significant results on partizan loopy games, focusing on the theory developed in the 1970s by Conway, Bach and Norton.

1. Introduction

A substantial portion of combinatorial games research focuses on games without repetition — those that are guaranteed to terminate after some finite number of moves. Such games are highly tractable, both theoretically and computationally, and the full force of the classical partizan theory can be brought to bear upon them. The great success of this theory has produced a vast body of splendid results, but it has also resulted in an unjust neglect of games *with* repetition.

In the late 1970s, John Conway and his students, Clive Bach and Simon Norton, introduced a disjunctive theory of partizan games with repetition — called *loopy games* because their game graphs may contain cycles. They showed that in many interesting cases, such games admit canonical forms. The past few years have witnessed some significant applications of this theory, to games as diverse as Fox and Geese, Hare and Hounds, Entrepreneurial Chess, and one-dimensional Phutball. In light of these advances, it is time for a reappraisal of the theory with an eye to the future.

A short history. The first disjunctive theory of loopy games is due to Cedric A. B. Smith and Aviezri Fraenkel. They showed (independently) that the usual Sprague–Grundy theory generalizes well to loopy games. In particular, many impartial loopy games are equivalent to nimbers, and the remainder are characterized by their nimber-valued options. Over a period of several decades, Fraenkel and his students explored this theory in depth. They constructed numerous examples and studied both their solutions and their computational complexity.

The partizan theory was introduced by Robert Li, who studied *Zugzwang games*, those in which it is a disadvantage to move. Li showed that Zugzwang games are completely characterized by a certain pair of ordinary numbers. Soon thereafter, Conway, Bach and Norton extended Li's theory to a much broader class of games. They showed that many loopy games γ — including most positions encountered in actual play — decompose into a pair of much simpler games, called the *sides* of γ. Their theory was published in the first edition of *Winning Ways*, together with a handful of examples, most notably the children's game *Fox and Geese*.

Intermittent progress was made over the next twenty years, but it was not until 2003 that loopy games saw a full-fledged revival. John Tromp and Jonathan Welton had recently detected an error in the *Winning Ways* analysis of Fox and Geese, and Berlekamp set out to repair it. His corrected analysis appears in the second edition of *Winning Ways*. Berlekamp's effort led to the development of new algorithms, which in turn paved the way for a re-examination of several other loopy games mentioned in *Winning Ways*.

In this survey, the *Winning Ways* theory is introduced first, so that earlier developments — notably those of Smith, Fraenkel and Li — can be presented in the modern context. Section 2 is an expository overview of some interesting properties of loopy games, with a focus on Fox and Geese. Much of that material is formalized in Section 3, and in Section 4 we tackle the theory of sides as it appears in *Winning Ways*. Each of these sections also addresses some related topics. Section 5 discusses several specific partizan games that have been successfully analyzed with this theory. In Section 6, we discuss the generalized Sprague–Grundy theory and its relationship to partizan games. Section 7 gives an overview of the Smith–Flanigan results on conjunctive and selective sums. Finally, in Section 8 we survey the development of algorithms for loopy games.

Two topics are notably absent from this survey. The first is the immense body of work on loopy impartial games, assembled over several decades by Aviezri Fraenkel and his students. Their work includes an extensive theoretical and algorithmic analysis of the generalized Sprague–Grundy theory; many beautiful examples; and connections to other fields, including combinatorial number theory and error-correcting codes. The present article is focused mainly on the partizan theory, and so does not do justice to their achievement; a forthcoming book by Fraenkel surveys this material in far more detail and accuracy than we could hope to achieve here.

The combinatorial theory of Go is another major omission. This might seem surprising, since Go is without question the most significant loopy game that has been subjected to a combinatorial analysis. However, there is good reason for its omission. Although Go is fundamentally disjunctive in nature, its unique *koban* rule implies an interrelationship between all components on the board. This gives rise to a rich and fascinating temperature theory that has been explored by many researchers, including Berlekamp, Fraser, Kao, Kim, Müller, Nakamura, Snatzke, Spight, and Takizawa, to list just a few. However, this temperature theory appears to be incompatible with the canonical theory that is the focus of our discussion. Because Go is so prominent, its body of results is vast; yet because it is so singular, these appear disconnected from other theories of loopy games. Thus while Go desperately deserves its own survey, this article is not the appropriate place for it.

This apparent dichotomy also raises the first — and arguably the most important — open problem of this survey.

OPEN PROBLEM. *Formulate a temperature theory that applies to all loopy games.*

Notation and preliminaries. Following *Winning Ways*, we denote loopy games by loopy letters $\gamma, \delta, \alpha, \beta, \ldots$. If γ is loopy, we define the associated *game graph* $\mathcal{G}$ as follows. $\mathcal{G}$ has one vertex, V_α, for each subposition α of γ (including γ itself), and there is an edge directed from V_α to V_β just if there is a legal move from α to β. When γ is partizan, we color the edge b*L*ue, *R*ed, or gr*E*en, depending on whether *L*eft, *R*ight, or *E*ither player may move from α to β.

An abbreviated notation is often useful. In many loopy games, repetition is limited to simple pass moves. In such cases we can borrow the usual brace-and-slash notation used to describe loopfree games, enhanced with the additional symbol **pass**. For example, if we write $\gamma = \{0 \mid \textbf{pass}\}$, we mean that Left has a move from γ to 0, and that Right has a move from γ back to γ. Likewise, if $\delta = \{0 \mid \textbf{pass} \parallel -1\}$, this means that Right has a move from δ to -1, and that Left has a move from δ to $\{0 \mid \textbf{pass}\} = \gamma$. For comparison, the game graph of δ is shown in Figure 1.

The main complication introduced by loopy games is the possibility of nonterminating play. The simplest way to resolve this issue is to declare all infinite

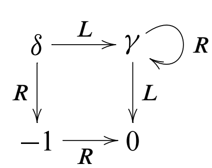

Figure 1. The game graph of $\delta = \{0 \mid \textbf{pass} \parallel -1\}$.

plays *drawn*, and this will be our assumption throughout Sections 2 and 3. We will often say that a player *survives* the play of a game if he achieves at least a draw.

2. Loops large and small

Fox and Geese is an old children's game played on an ordinary checkerboard. Four geese are arranged against a single fox as in Figure 2. The geese (controlled by Left) move as ordinary checkers, one space diagonally in the forward direction, while the fox (controlled by Right) moves as a checker king — one space in any diagonal direction. Neither animal may move onto an occupied square, and there are no jumps or captures. The geese try to trap the fox, while the fox tries to escape.

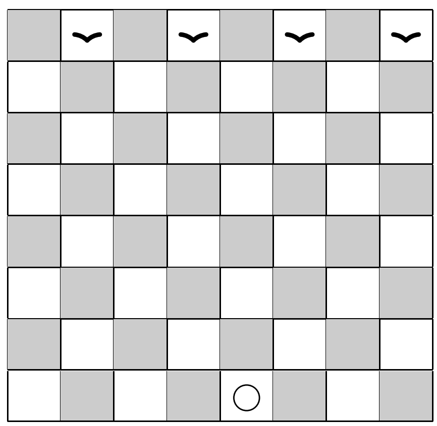

Figure 2. The usual starting position in Fox and Geese.

Fox and Geese has a curious feature: the game must end if played in isolation, because the geese will eventually run out of moves, whether or not they trap the fox. However, from a combinatorial perspective the game is certainly loopy. The *fox* may return to a previous location, and this results in local repetition if Left's intervening moves occur in a different component.

Before turning to a more formal treatment of loopy games and canonical forms, let us briefly investigate the behavior of Fox and Geese. Consider first the happy affair of an escaped fox (Figure 3). The geese have exhausted their supply of moves, and though Left has a tall Hackenbush stalk at his disposal, his situation is hopeless. Inevitably, he will run out of moves, and the fox will still be dancing about the checkerboard, none the worse for wear.

It is clear that an escaped fox α is more favorable to Right than any (finite) Hackenbush stalk we might assemble. In an informal sense, we have established

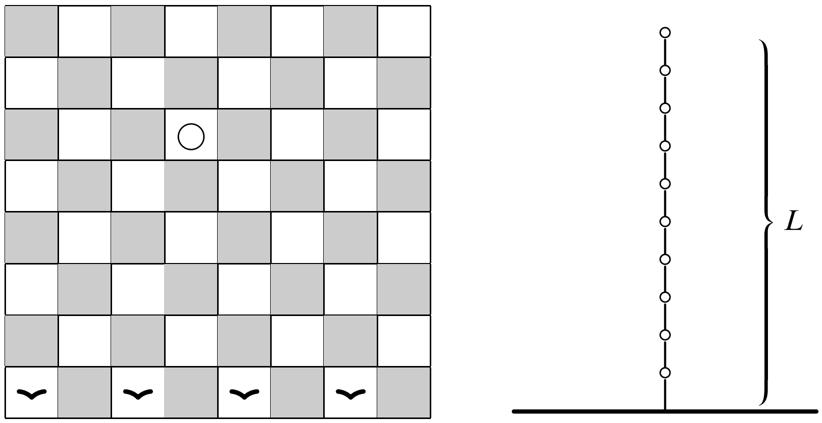

Figure 3. Hackenbush is hopeless facing an escaped fox.

that

$$\alpha < -n,$$

for every integer n.

It is equally clear that the fox's precise location on the checkerboard is irrelevant; all that matters is that she has an indefinite supply of moves at her disposal. The many distinct positions that arise as she moves about the board are all equivalent, and α can be written as a single pass move for Right: $\alpha = \{\,|\,\mathbf{pass}\}$, with game graph shown in Figure 4.

Figure 4. The games $\alpha = \mathbf{off}$, $\beta = \mathbf{on}$, and $\delta = \mathbf{dud}$.

The game α is normally known as **off**, and its inverse — from which *Left* can pass — is naturally enough called **on**.[1] One might expect that $\mathbf{on} + \mathbf{off} = 0$, but this is not the case: in their sum either player may pass, so that $\mathbf{on} + \mathbf{off}$ is a draw, while 0 is a second-player win.

In fact it is easy to see that $\mathbf{on} + \mathbf{off} + \gamma$ is drawn, no matter what game γ we include in the sum: both players have an inexhaustible supply of moves; so neither has anything to fear. Therefore $\mathbf{on} + \mathbf{off}$ is a **d**eathless **u**niversal **d**raw, which we abbreviate by **dud**, and we have the identity

$$\mathbf{dud} + \gamma = \mathbf{dud}$$

for all γ.

[1] The name is set-theoretic in origin: **ON** is standard notation for the class of all ordinal numbers, and the game **on** behaves much like an ill-founded relation, an entity that exceeds all the ordinals.

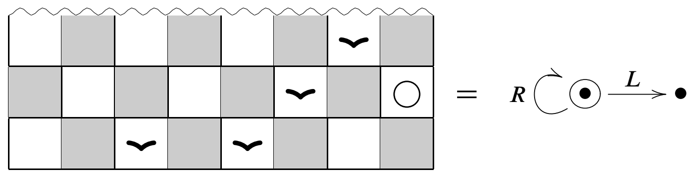

Figure 5. A trapped fox has value **over**.

Soon we shall put all of this on formal footing, but first consider one more example to illustrate the remarkable properties of loopy games. In Figure 5 the Fox is trapped. She is forced to shuttle indefinitely between the two lower-right-hand spaces, and at any moment the geese may choose to end the game. It is clear this game is positive, for Left may win at any time. Its abbreviated graph, known as **over**, is also pictured in Figure 5.

Just how large is **over**? The reader might wish to confirm that, for any n,

$$n \cdot \uparrow < \textbf{over} < \frac{1}{2^n},$$

by showing that Left can win the appropriate differences. **over** is larger than every loopfree infinitesimal, but smaller than every positive number.

3. Stoppers

When γ is loopy, there are typically three possible outcomes: Left wins (if he gets the last move); Right wins (if she gets the last move); or a Draw (if play never terminates). This divides loopy games into nine outcome classes, since the outcome might depend on who moves first:

		Left moves first		
		Left wins	Draw	Right wins
Right moves first	Left wins	$\mathscr{L}$	$\hat{\mathscr{P}}$	$\mathscr{P}$
	Draw	$\hat{\mathscr{N}}$	$\mathscr{D}$	$\check{\mathscr{P}}$
	Right wins	$\mathscr{N}$	$\check{\mathscr{N}}$	$\mathscr{R}$

We denote by $o(\gamma)$ the outcome class of γ. The outcome classes are naturally partially-ordered as shown in Figure 6.

As always in combinatorial game theory, we define equality by indistinguishability in sums:

$$\gamma = \delta \quad \text{if} \quad o(\gamma + \alpha) = o(\delta + \alpha) \text{ for all loopy games } \alpha.$$

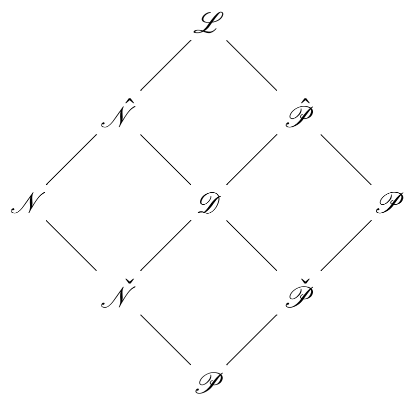

Figure 6. The partial order of loopy outcome classes.

As remarked in Section 2, it is *not* always true that $\gamma - \gamma = 0$. Second player can always assure a draw by playing the mirror image strategy, but in general this does not guarantee a win. For this reason, loopy games do not form a group, and we are forced to consider instead the monoid of loopy games, equipped with the natural partial order:

$$\gamma \geq \delta \quad \text{if} \quad o(\gamma + \alpha) \geq o(\delta + \alpha) \text{ for all loopy games } \alpha.$$

The theory of loopy games is motivated by two fundamental questions.

- Does every loopy game admit a unique simplest form, analogous to the canonical form of a loopfree game?
- Can one specify an *effective* equivalent definition of $\gamma \geq \delta$?

It turns out that both of these questions are easiest to answer for an important special class of loopy games called *stoppers*. They can also be resolved quite nicely for a larger class, the *stopper-sided games*, that encompasses most positions arising in studies of actual (playable) games.

A loopy game γ is a *stopper* if there is no infinite *alternating* sequence of play proceeding from any subposition of γ. The games **on**, **off**, and **over**, which we met in Section 2, are all stoppers, but **dud** is not. Further, *every* position that arises in Fox and Geese is a stopper, since the geese are constrained to make finitely many moves throughout the game.

If γ is a stopper, then γ is guaranteed to terminate when played in isolation. This property is central to the following characterization.

THEOREM 1 (CONWAY). *Let γ, δ be stoppers. Then*

$$\gamma \geq \delta \textit{ iff Left, playing second, can survive } \gamma - \delta.$$

PROOF. For the forward direction, suppose $\gamma \geq \delta$, and let $\alpha = -\delta$. Certainly Left can survive $\delta + \alpha$, by playing the mirror image strategy; then it follows directly from the definition of $\geq$ that he can survive $\gamma + \alpha$.

For the reverse direction, fix any loopy game α. We must show that:

(i) If Left can survive $\delta + \alpha$ playing first (second), then he can survive $\gamma + \alpha$ playing first (second).
(ii) If Left can win $\delta + \alpha$ playing first (second), then he can win $\gamma + \alpha$ playing first (second).

First suppose that Left is second player in case (i). We describe a strategy for playing $\gamma + \alpha$ that guarantees at least a draw.

Before play begins, Left constructs two dummy games: one copy of $\delta + \alpha$, and one copy of $\gamma - \delta$. Whenever Right makes a move in $\gamma + \alpha$, Left copies the move to the appropriate dummy game: if Right moves in the γ component, Left copies the move to $\gamma - \delta$; while if Right moves in the α component, Left copies the move to $\delta + \alpha$.

Now Left responds with his survival move in the dummy game. If this move is in the δ or $-\delta$ component, Left immediately makes the mirror image move in the *other* dummy game, and responds accordingly. Successive responses in the δ and $-\delta$ components produce an alternating sequence of moves proceeding from a subposition of δ. Since δ is a stopper, this cannot go on forever, and eventually Left's response must occur in the γ or α component. At that point Left copies it back to $\gamma + \alpha$ and awaits Right's next move.

If Left keeps to this strategy, he will never run out of moves in $\gamma + \alpha$. This proves case (i). In case (ii), Left uses the same technique, but follows his *winning* strategy in $\delta + \alpha$. This guarantees that eventually, $\delta + \alpha$ will reach a terminal position. At that point the α component of $\gamma + \alpha$ is terminal; therefore, since γ is a stopper, it must eventually terminate as well. So the game will necessarily end, and since Left has survived, Right cannot have made the last move.

If Left is first player, the argument is exactly the same. He makes his initial move in the $\delta + \alpha$ dummy component, according to his first-player survival (or winning) strategy for $\delta + \alpha$, and continues accordingly. □

Stoppers also admit a clean canonical theory: if γ is a stopper, then we can eliminate dominated options and bypass reversible ones, just as for loopfree games. The proofs are straightforward applications of Theorem 1.

A stopper is in *simplest form* if it has no dominated or reversible options.

THEOREM 2 (CONWAY). *If γ and δ are stoppers in simplest form with $\gamma = \delta$, then for every γ^L there is a δ^L with $\gamma^L = \delta^L$, and vice versa*; *and likewise for Right options*.

PROOF. See [5, Section 10]. □

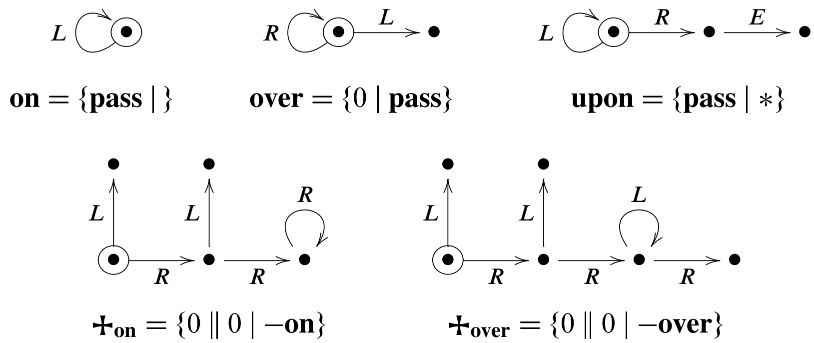

Figure 7. Simple stoppers.

Several simple stoppers. Figure 7 shows some of the simplest stoppers in canonical form. The reader might wish to verify some of their remarkable properties, which clarify their behavior in the partial-order of games:

- $\mathbf{on} \geq \gamma$ for all γ.
- $n \cdot \uparrow < \mathbf{over} < 2^{-n}$ for all n.
- $\uparrow^{\rightarrow n} < \mathbf{upon} < \uparrow^{\rightarrow n} + \uparrow^{n}$ for all n.
- $+_{\mathbf{over}} \leq \gamma$ for every all-small game $\gamma > 0$, but $+_{\mathbf{over}} > +_x$ for every number $x > 0$.
- $+_{\mathbf{on}}$ is the smallest positive game: if $\gamma > 0$, then $+_{\mathbf{on}} \leq \gamma$.

With the exception of $+_{\mathbf{over}}$, all of these values arise frequently in playable games. Also common is $\mathbf{upon} + *$, which has the canonical form $\{0, \mathbf{pass} \mid 0\}$.

In all of these examples, the only loops are simple pass moves (1-cycles). Stoppers with longer cycles exist, but are much less common in nature. A typical example is the game τ shown in Figure 8, which has a 4-cycle in canonical form.

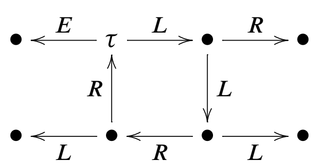

Figure 8. A stopper with a canonical 4-cycle.

Stoppers in canonical form can never have 2-cycles or 3-cycles; see [30] (this volume) for a proof, together with examples of stoppers with canonical n-cycles for all $n \geq 4$.

Idempotent	**Loopfree Games Absorbed**
$\mathbf{on} = \{\mathbf{pass} \mid \}$	All games
$\mathbf{over} = \{0 \mid \mathbf{pass}\}$	All infinitesimals
$\mathbf{star}_n = \{0 \parallel 0, *n \mid 0, \mathbf{pass}\}$ $(n \geq 2)$	$*n$ and $\uparrow^2$, but not $*m$ for any $m \neq n$
$\mathscr{I}_n = \{0 \parallel 0, \mathbf{pass} \mid 0, \downarrow_{[n-2]}*\}$ $(n \geq 2)$ $\mathscr{J}_n = \{0 \parallel 0, \downarrow_{[n-1]}* \mid 0, \mathbf{pass}\}$ $(n \geq 2)$	$\uparrow^n$ but not $\uparrow^{n-1}$
$\uparrow^{\mathbf{on}} = \{0 \parallel 0 \mid 0, \mathbf{pass}\}$	"Almost tiny" all-smalls (such as $\{0 \| 0 \mid \Downarrow\}$), but not $\uparrow^n$ for any n
$+_{\mathbf{over}} = \{0 \parallel 0 \mid \mathbf{under}\}$	All tinies, but no all-smalls
$+_{x\mathbf{under}} = \{0 \parallel 0 \mid -x\mathbf{over}\}$ $(x > 0)$	$+_{x\downarrow n}$ for all n, but not $+_{x-2^{-n}}$
$\mathscr{T}_x = \{0 \parallel 0 \mid -x, \mathbf{pass}\}$ $(x > 0)$	$+_y$ for all $y > x$, but not $+_x$
$+_{x\mathbf{over}} = \{0 \parallel 0 \mid -x\mathbf{under}\}$ $(x > 0)$	$+_{x+2^{-n}}$ for all n, but not $+_{x\uparrow n}$
$+_{\mathbf{on}} = \{0 \parallel 0 \mid \mathbf{off}\}$	None (except 0)

Figure 9. A variety of idempotents.

Idempotents. It is easy to see that $\mathbf{on} + \mathbf{on} = \mathbf{on}$: certainly $\mathbf{on} + \mathbf{on} \geq \mathbf{on}$, but we also know that $\mathbf{on} \geq \gamma$ for all γ. Slightly less obvious is the fact that $\mathbf{over} + \mathbf{over} = \mathbf{over}$, and here Theorem 1 is useful. To show that $\mathbf{over} + \mathbf{over} \leq \mathbf{over}$, we need simply exhibit a second-player survival strategy for Left in

$$\mathbf{over} + \mathbf{under} + \mathbf{under},$$

where $\mathbf{under} = -\mathbf{over} = \{\mathbf{pass} \mid 0\}$.

This is not difficult: so long as any **under** components remain, Left makes pass moves. This guarantees that, if Right ever destroys both **under** components (by moving from **under** to 0), the **over** component will still be present. Therefore, if Right destroys both **under** components, Left can win the game by moving from **over** to 0.

This example illustrates a striking feature of the monoid of loopy games: the presence of explicit idempotents. Figure 9, reproduced from [27], lists many more. Each idempotent ι is listed together with some of the loopfree games that it absorbs (where ι *absorbs* γ if $\iota + \gamma = \iota$). It's also worth noting that each idempotent ι in Figure 9 has a "negative variant" $-\iota$ and a "neutral variant" $\iota - \iota$, both of which are also idempotents (though of course, $\iota - \iota$ is not a stopper).

Berlekamp [2] describes several other idempotents that do not appear to have explicit representations as loopy games. These include ☆ and $\mathscr{E}_t$, which play central roles in the atomic weight and orthodox theories, respectively. It would be interesting to describe a formal system that encompasses these in addition to the idempotents of Figure 9.

Pseudonumbers. The *psuedonumbers* form an interesting subclass of infinite stoppers.

DEFINITION 3. A stopper x is said to be a *pseudonumber* if, for every follower y of x (including x itself), we have $y^L < y^R$ for all y^L, y^R.

So a surreal number is just a well-founded pseudonumber. It is not hard to show that x is a pseudonumber if and only if, for every follower y of x, each $y^L \le y$ and $y \le$ each y^R. Then as a consequence of Li's Theorem (Theorem 9 in Section 4, below), the only *finite* pseudonumbers are **on**, **off**, and the dyadic rationals and their sums with **over** and **under**. However, there are many infinite pseudonumbers. A typical example is the game

$$\widehat{\mathbb{Z}} = \{0, 1, 2, \ldots \mid \textbf{pass}\} = \omega : \textbf{off}.$$

It is not hard to check that $\widehat{\mathbb{Z}} \ge n$ for any integer n. Furthermore, it is the *least* pseudonumber with this property: if $y \ge n$ for all n, then $y \ge \widehat{\mathbb{Z}}$. Therefore $\widehat{\mathbb{Z}}$ is a least upper bound for the integers. This generalizes:

THEOREM 4. *The pseudonumbers are totally ordered by* $\ge$. *Furthermore*, *every set* $X = \{x, y, z, \ldots\}$ *of pseudonumbers has a least upper bound*, *given by*

$$\widehat{X} = \{x, y, z, \ldots \mid \textbf{pass}\} = \{x, y, z, \ldots \mid \} : \textbf{off}.$$

PROOF. See [27, Section 1.8]. □

Contrast this with surreal numbers, which certainly do *not* admit tight bounds. However, while they acquire some analytic structure, pseudonumbers lose the rich algebraic structure of the surreal numbers: they are not even closed under addition, since (say) **on** + **off** is not a stopper.

Pseudonumbers might seem fanciful, but astonishingly, Berlekamp and Pearson recently discovered positions in Entrepreneurial Chess with offside $\widehat{\mathbb{Z}}$ (see Section 5 for a description). Like all good numbers, $\widehat{\mathbb{Z}}$ also makes an appearance in Blue-Red Hackenbush (Figure 10).

4. Sides

As we have seen, stoppers generalize the canonical theory of loopfree games in a straightforward way. Most loopy games, however, are not stoppers.

A typical example is the game *Hare and Hounds*, which has experienced occasional bouts of popularity dating back to the late nineteenth century. The

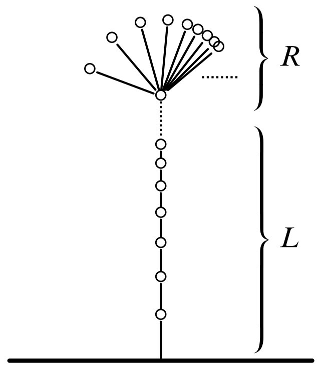

Figure 10. A Blue-Red Hackenbush position of value $\widehat{\mathbb{Z}} = \omega : \mathbf{off}$.

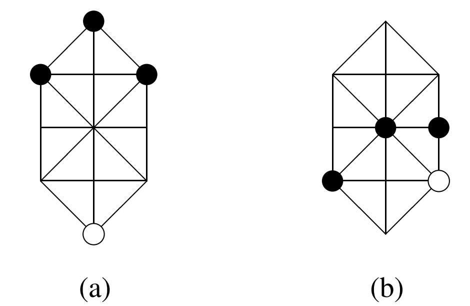

Figure 11. Hare and Hounds: (a) the starting position; (b) an endgame position.

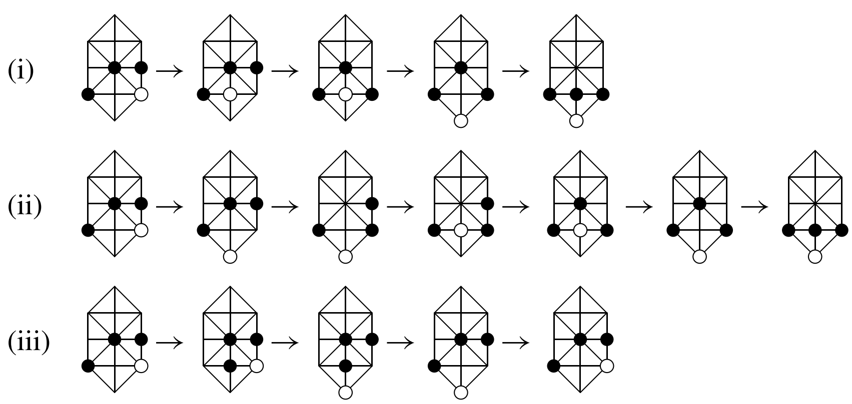

Figure 12. The position of Figure 11(b) is a second-player win.

game can played on an $n \times 3$ board for any odd $n \geq 5$; the starting position on the 5×3 board is shown in Figure 11(a).

The play resembles Fox and Geese. Left controls three hounds (black circles) and Right the lone hare (white circle). Each player, on his turn, may move any one of his units to an adjacent unoccupied intersection. The only restriction is that the hounds may never retreat — they can only advance or move sideways. There are no jumps or captures. The hounds win if they trap the hare (that is, if it is Right's turn and she has no moves available); the hare wins if this never happens.

Since the hounds are allowed to move sideways, Hare and Hounds is not always a stopper. It has another notable feature: if play never terminates, then the game is declared a win for Right. This differs from the other games we have studied, in which infinite plays are drawn. However, we will see that it actually makes the game simpler, since it causes many positions to reduce to stoppers that otherwise would not.

For example, consider the endgame position γ of Figure 11(b). If Right makes either of her available moves, then the hounds can certainly trap her; see Figure 12(i) and (ii). Conversely, if Left moves first, then the hare can evade capture indefinitely by following the pattern shown in Figure 12(iii). (The reader might wish to verify that if the *hounds* ever deviate from this pattern, then the hare can escape outright.) Therefore γ is a second-player win, and we conclude that $\gamma = 0$.

In the late 1970s, Conway, Bach and Norton made a breakthrough in the study of loopy games [5]. They observed, first of all, that games such as Hare and Hounds — where infinite plays are wins for one of the players — can often be brought into the theory of stoppers in a coherent way. Furthermore, their presence actually simplifies the analysis of games where infinite plays are drawn.

To understand this relationship, let γ be an arbitrary loopy game with infinite plays drawn, and suppose we wish to know whether Left can win γ. Then we might as well assume that infinite plays are wins for Right. Likewise, if we wish to know whether Left can survive γ, then we might as well assume that infinite plays are wins for Left. Therefore, we can determine the outcome class of γ by considering each of these two variants in turn. As it turns out, the variants often reduce to stoppers, even when γ itself does not; and in such cases, this reduction yields a substantial simplification.

Therefore, we now drop the assumption that all infinite plays are drawn. We assume that each game γ comes equipped with one of three winning conditions: Left wins infinite plays; Right wins infinite plays; or infinite plays drawn. We say that γ is *free* if infinite plays are draws and *fixed* otherwise.

When γ is free, we denote by γ^+ and γ^- the matching fixed games with infinite plays redefined as wins for Left and Right, respectively. When γ is fixed, we simply put $\gamma^+ = \gamma^- = \gamma$.

If infinite play occurs in a sum

$$\alpha + \beta + \cdots + \gamma,$$

we assume that Left (Right) wins the sum if he wins on *every* component in which play is infinite. If there are any draws, or if several components with infinite play are split between the players, the outcome of the sum is a draw.

When we consider the definition of $\geq$, we suppose now that α ranges over all fixed games in addition to free ones:

$$\gamma \geq \delta \quad \text{if} \quad o(\gamma+\alpha) \geq o(\delta+\alpha) \text{ for all } \textit{fixed or free} \text{ loopy games } \alpha.$$

The main result is the following, called the *Swivel Chair Theorem* in Winning Ways. It is a direct generalization of Theorem 1.

THEOREM 5 (SWIVEL CHAIR THEOREM). *The following are equivalent, for any loopy games* γ, δ:

(i) $\gamma \geq \delta$;
(ii) *Left, playing second, can survive both* $\gamma^+ - \delta^+$ *and* $\gamma^- - \delta^-$.

PROOF. See [3, Chapter 11] or [5, Section 2]; it's very similar to the proof of Theorem 1. □

Note the key implication of Theorem 5: how γ compares with other games depends only on γ^+ and γ^-. Thus when γ^+ and γ^- are equivalent to stoppers s^+ and t^-, the behavior of γ is completely characterized by s and t. In such cases we call s and t the *sides* of γ (the *onside* and *offside* respectively), and we say that γ is *stopper-sided*. It is customary to write

$$\gamma = s \,\&\, t,$$

and with s and t in simplest form, this should be regarded as a genuine canonical representation for γ.

For example, consider the game **dud** = {**pass** | **pass**}. We know that $\mathbf{on}^+ \geq \mathbf{dud}^+$ (since $\mathbf{on}^+$ is the largest game of all). But also, Left can survive the game

$$\mathbf{dud}^+ - \mathbf{on}^+$$

by passing indefinitely in the **dud** component, where he wins infinite plays. We conclude that $\mathbf{dud}^+ = \mathbf{on}^+$, and by a symmetric argument $\mathbf{dud}^- = \mathbf{off}^-$. This gives the identity

$$\mathbf{dud} = \mathbf{on} \,\&\, \mathbf{off}.$$

If $\gamma = s \,\&\, t$, then the outcome class of γ is determined by those of s and t. Since s and t are stoppers, their outcomes fall into the usual classes: positive, negative, fuzzy or zero. This yields a total of sixteen possibilities for γ. However, since $\gamma^+ \geq \gamma^-$, we know that $s^+ \geq t^-$; and since s and t are stoppers, this implies that $s \geq t$. This restriction rules out seven possibilities, leaving the remaining nine in one-to-one correspondence with the nine outcome classes discussed in Section 3. This correspondence is summarized in Figure 13.

		s			
		> 0	$\not\gtrless 0$	$= 0$	< 0
t	< 0	$\mathscr{D}$	$\check{\mathscr{N}}$	$\check{\mathscr{P}}$	$\mathscr{R}$
	$\not\gtrless 0$	$\hat{\mathscr{N}}$	$\mathscr{N}$	–	–
	$= 0$	$\hat{\mathscr{P}}$	–	$\mathscr{P}$	–
	> 0	$\mathscr{L}$	–	–	–

Figure 13. The outcome class of $\gamma = s \ \& \ t$ is determined by those of s and t.

The sides of γ therefore carry a great amount of information. Given their applicability, it is natural to ask how they might be computed in general. *Winning Ways* introduced a method called *sidling* that yields a sequence of increasingly good approximations to the sides of γ. Sometimes this sequence converges to the true onside and offside; but more often than not, it fails to converge. Nonetheless, sidling has been applied to obtain some interesting results, notably by David Moews in his 1993 thesis [21] and a subsequent article on Go [22].[2]

More recent discoveries include effective methods for computing sides (when they exist); see [30] in this volume for discussion.

Carousels. Stopper-sided decompositions are both useful and extremely common. However, there do exist loopy games that are not stopper-sided. In the 1970s, Clive Bach produced the first example of such a game, known as *Bach's Carousel*, by specifying its game graph explicitly. Much more recently, similar "carousels" have been discovered on 11×3 boards in Hare and Hounds. See Figure 14 for an example and [27] for further discussion.

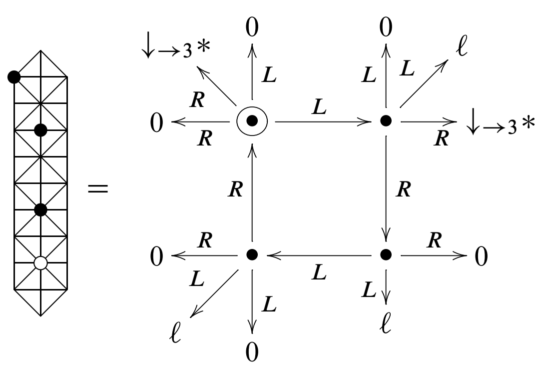

Figure 14. A carousel in Hare and Hounds. Here $\ell = \{0, \downarrow_{\rightarrow 2}{*} \mid 0, \downarrow_{\rightarrow 2}{*}\}$.

[2]In order to bring Go into the canonical theory, Moews considered Go positions together with explicitly kobanned moves. Means and temperatures as defined by Berlekamp cannot be recovered from the resulting canonical forms. Nonetheless, Moews' analysis yields interesting information about Go positions that is not captured by thermography.

What about sums of stoppers? *Winning Ways* gives an example, due to Bach, of certain *infinite* stoppers whose sum is not stopper-sided. But the following question remains open:

QUESTION. *Is the sum of* finite *stoppers necessarily stopper-sided*?

Finally, the following question was posed in *Winning Ways* and remains open.

QUESTION (BERLEKAMP–CONWAY–GUY). *Is there an alternative notion of simplest form that works for all finite loopy games*?

Degrees, classes, and varieties. When γ is loopy, it is often the case that $\gamma - \gamma \neq 0$. Provided $\gamma - \gamma$ is stopper-sided, we define the *degree of loopiness* (or *degree*) γ° by

$$\gamma^\circ = \text{Onside}(\gamma - \gamma).$$

If γ is equivalent to a loopfree game, then $\gamma^\circ = 0$; otherwise $\gamma^\circ > 0$. For example, it is not hard to check that $\mathbf{on}^\circ = \mathbf{on}$, $\mathbf{over}^\circ = \mathbf{over}$, and $\mathbf{upon}^\circ = {\uparrow}^{\mathbf{on}} = \{0 \mid -\mathbf{upon}{*}\}$.

For a fixed idempotent ι, the games of degree ι tend to group naturally into *classes* and *varieties* that interact in predictable ways. These were investigated in *Winning Ways* for the idempotent

$$\spadesuit = \left\{0^2 \,\middle\|\, \{\mathbf{on} \mid 0^4\}\right\}.$$

However, since the publication of the first edition of *Winning Ways*, there has been little effort to move the theory forward. For this reason, we omit a full discussion and instead refer the reader to *Winning Ways*. It is perhaps time to study classes and varieties in more detail, in light of recent discoveries concerning other aspects of loopy games.

OPEN PROBLEM. *Investigate the class structure of each idempotent in Figure* 9.

Zugzwang games. Although the theory of sides is due to Conway and his students, its acknowledged inspiration is an earlier study by Robert Li, a student of Berlekamp's in the 1970s [20]. Li investigated so-called *Zugzwang games* — those in which it is a disadvantage to move — and found that they generalize ordinary numbers in a straightforward way.

DEFINITION 6. γ is a *Zugzwang game* if, for every follower δ of γ, each $\delta^L < \delta$ and $\delta <$ each δ^R.

Li's Theorem completely classifies all loopy Zugzwang games:

THEOREM 7 (LI). *Let γ be a loopy game. Then the following are equivalent*:

(a) γ *is equal to some Zugzwang game*;

(b) *There exist dyadic rationals x and y, $x \geq y$, such that*

$$\gamma = x \mathbin{\&} y.$$

PROOF. See [20, Section 4]. □

Li also studied a mild generalization of Zugzwang games, which he called *weak Zugzwang games*.

DEFINITION 8. γ is a *weak Zugzwang game* if, for every follower δ of γ, each $\delta^L \leq \delta$ and $\delta \leq$ each δ^R.

Note that for *loopfree* games G, the weak and strong Zugzwang notions coincide, since necessarily $G \neq G^L, G^R$. For loopy games, however, there are several further weak Zugzwang games.

THEOREM 9 (LI). *Let γ be a loopy game. Then the following are equivalent*:

(a) *γ is equal to some weak Zugzwang game*;
(b) *$\gamma = x \mathbin{\&} y$, where $x \geq y$ and each of x, y is one of the following*:

(i) **on**;

(ii) **off**;

(iii) *A dyadic rational*;

(iv) *$z +$ **over** for some dyadic rational z*; *or*

(v) *$z +$ **under** for some dyadic rational z*.

PROOF. See [20, Section 6]. □

Li's results are intrinsically interesting, and also quite remarkable, given that he had none of the modern machinery of loopy games at his disposal.

5. Some specific partizan games

Several partizan games have been successfully analyzed using the disjunctive theory. We briefly survey the most important examples.

Fox and Geese. This game has been largely solved by Berlekamp, who showed that the *critical position* of Figure 15 has the exact value $1 + 2^{-(n-8)}$, where $n \geq 8$. *CGSuite* has confirmed that the 8×8 starting position (Figure 2) has value $2 +$ **over**. Many other interesting values arise; these are summarized in *Winning Ways* and in slightly more detail in [27].

Berlekamp's analysis leaves little to be discovered about Fox and Geese proper. Nonetheless, we can ask interesting questions about certain variants of the game. Murray [23] describes a variant from Ceylon, *Koti keliya*, which is played with six geese ("dogs" or "cattle") on the 12×12 board, with the fox

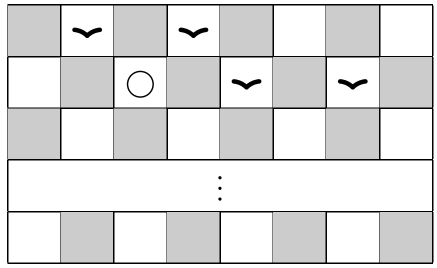

Figure 15. This critical position on an $n \times 8$ board has value $1 + 2^{-(n-8)}$ $(n \geq 8)$.

("leopard") permitted two moves per turn. It is unclear whether these moves must be in the same direction. Although a full solution to the 12×12 board appears to be out of reach computationally, it is interesting to observe how the fox's increased mobility affects play on smaller boards. As one might expect, it is far easier for the fox to escape, and positions whose values are large *negative* numbers become quite common. In fact, the following conjecture seems justified:

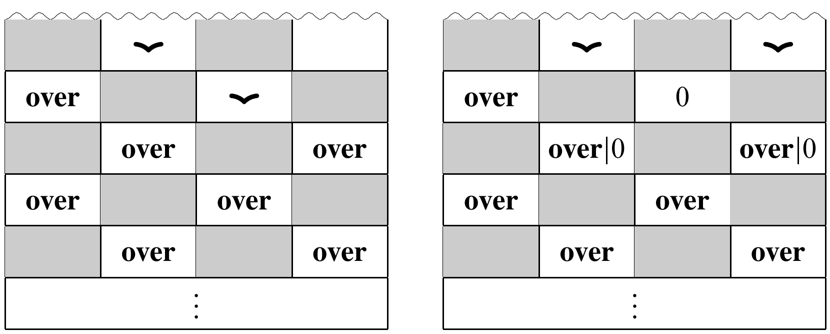

Figure 16. Conjectured values of $n \times 4$ Fox and Geese $(n \geq 5)$.

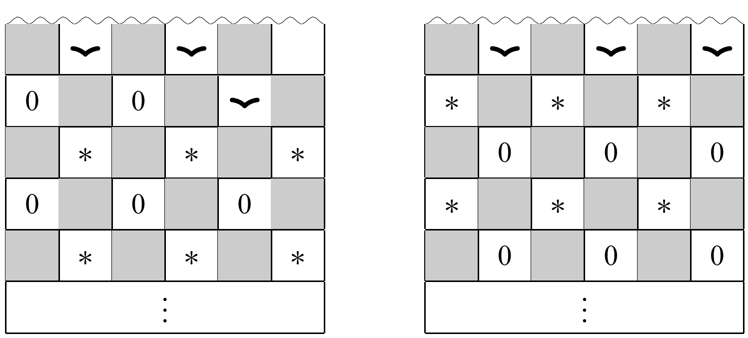

Figure 17. Conjectured values of $n \times 6$ Fox and Geese $(n \geq 8)$.

CONJECTURE. *The critical position of Figure* 15, *played with Ceylonese rules*, *has value* $-2n+11$ *for all* $n \geq 6$.

Finally, there is overwhelming experimental evidence for the following two conjectures.

CONJECTURE. *The diagrams of Figure* 16 *are valid on the* $n \times 4$ *Fox and Geese board*, *for all* $n \geq 5$. *Furthermore*, *the range of values that appear on the* $n \times 4$ *board can be classified completely*.

(In the diagrams of this and the next figure, the geese are fixed, and conjectured values are shown for each possible placement of the fox.)

CONJECTURE. *The diagrams of Figure* 17 *are valid on the* $n \times 6$ *Fox and Geese board*, *for all* $n \geq 8$.

Backsliding Toads and Frogs. *Backsliding Toads and Frogs* was introduced in *Winning Ways*. The game is played on a $1 \times n$ strip populated by several toads (controlled by Left, facing right) and frogs (controlled by Right, facing left). See Figure 18 for a typical starting position. There are two types of moves. Either player may *slide* one of his animals one space in either direction. Alternatively, he may choose to *jump* in the facing direction (toads to the right, frogs to the left). Players must jump over exactly one enemy (never a friendly animal) and must land on an unoccupied space. Jumps do not result in capture.

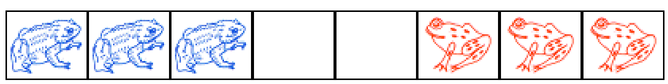

Figure 18. A typical starting position in Backsliding Toads and Frogs.

Readers familiar with ordinary Toads and Frogs will recognize the only difference between the two games: in the ordinary version, the animals are constrained to slide in the facing direction; in the loopy variant, they may slide backwards as well. This single difference has a monumental impact on the values that arise. The most obvious effect is that almost all positions in the Backsliding variant are loopy; for example, the position of Figure 18 has the remarkable value

$$\{\mathbf{on} \parallel 0 \mid -\tfrac{1}{2}\} \;\&\; \{\tfrac{1}{2} \mid 0 \parallel \mathbf{off}\}.$$

Positions in the Backsliding variant tend to have substantially *simpler* canonical forms than those in the loopfree version. For example, Erickson [8] noted that in ordinary Toads and Frogs, the “natural starting positions” of the form $\mathrm{T}^m\square^k\mathrm{F}^n$ are often quite complicated. In the Backsliding version, the *only* values (among all possibilities for k, m, n) are 0, $*$, **on**, **off**, **dud**, **on** & {**on** | **off**}, {**on** | **off**} & **off**, and the single anomalous value given above.

Nonetheless, Backsliding Toads and Frogs exhibits positions of value n and 2^{-n}, as well as positions of *temperature* n and 2^{-n}, for all $n \geq 0$. See [27, Chapter 3] or [29] for a complete discussion.

Hare and Hounds. Hare and Hounds exhibits asymptotic behavior much like Fox and Geese: the position shown in Figure 19, on a $(4n+5)\times 3$ or $(4n+7)\times 3$ board, has the exact value $-n$.

The mathematical analysis of Hare and Hounds began in the 1960s, when Berlekamp demonstrated a winning strategy for the hare on large boards. He was close to proving that Figure 19 has value $-n$, but the canonical theory had not yet been invented.

Hare and Hounds exhibits many interesting values, including $*2$ (rare among partizan games); $\uparrow^2$, $\uparrow^3$, and $\uparrow^4$ (but not, it seems, $\uparrow^5$); and a bewildering variety of stoppers. See [27, Chapter 4] or [28] (this volume) for examples of these, as well as a proof that Figure 19 has value $-n$.

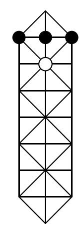

Figure 19. This critical position on the $(4n+5)\times 3$ or $(4n+7)\times 3$ board (shown here on the 9×3 board) has value $-n$.

Chess. Noam Elkies has observed several loopy values in Chess (in addition to many loopfree ones). See [6] for his constructions of **over** and **tis** $= 1 \ \& \ 0$. More recently, Elkies has produced positions of values **upon** and $+_{\mathbf{on}}$ [7]; see Figure 20. (The kings have been omitted from these diagrams in order to focus on the essential features of each position, but they can easily be restored without affecting the positions' values, using techniques outlined by Elkies [6].)

Entrepreneurial Chess. Entrepreneurial Chess is played on a quarter-infinite chessboard, with just the two kings and a White rook (Figure 21). In addition to his ordinary king moves, Left (Black) has the additional option of "cashing out." When he cashes out, the entire position is replaced by the integer n, where n is the sum of the *rank* and *column* values indicated in the diagram. Thus Left stands to gain by advancing his king as far to the upper-right as possible; and Right, with his rook, will eventually be able to stop him.

Entrepreneurial Chess was invented by Berlekamp, and has been studied extensively by Berlekamp and Pearson [4]. They have discovered many interesting

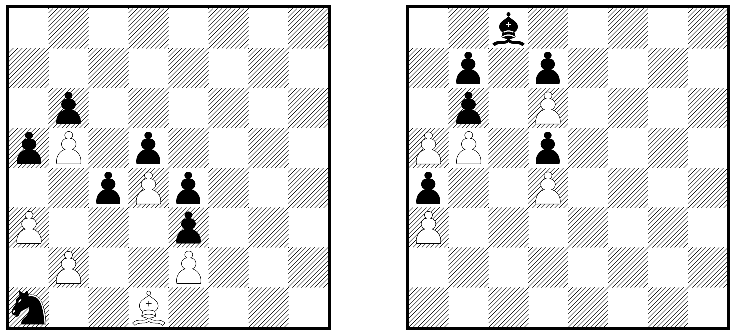

$\mathbf{upon} = \{\mathbf{pass}|*\}$ $\quad$ $+_{\mathbf{on}} = \{0\,\|\,0|\mathbf{off}\}$

Figure 20. Loopy values in chess.

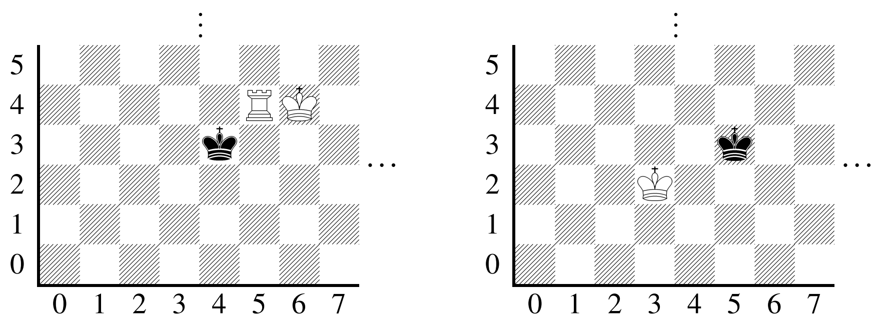

Figure 21. Entrepreneurial Chess.

values. For example, the position shown in Figure 21, left, has value $7 + \mathbf{over}$: Left can cash out for 7 points at any time, and in the meantime Right is constrained to shuttle his king between the squares adjacent to his rook. Berlekamp and Pearson's results also include a detailed temperature analysis of a wide class of positions.

A particularly interesting position γ arises in the pathological case when Left has captured Right's rook, as in Figure 21, right. The onside of γ is **on**, since Left need never cash out. Now consider the offside. Left must cash out *eventually*, since infinite plays are wins for Right, but he can defer this action for as long as necessary. Thus we have the remarkable identity

$$\gamma = \mathbf{on}\ \&\ \widehat{\mathbb{Z}},$$

where $\widehat{\mathbb{Z}} = \{0, 1, 2, \ldots \mid \mathbf{pass}\}$ is the pseudonumber defined in Section 3. This identity can be verified formally using the theory presented in Section 4.

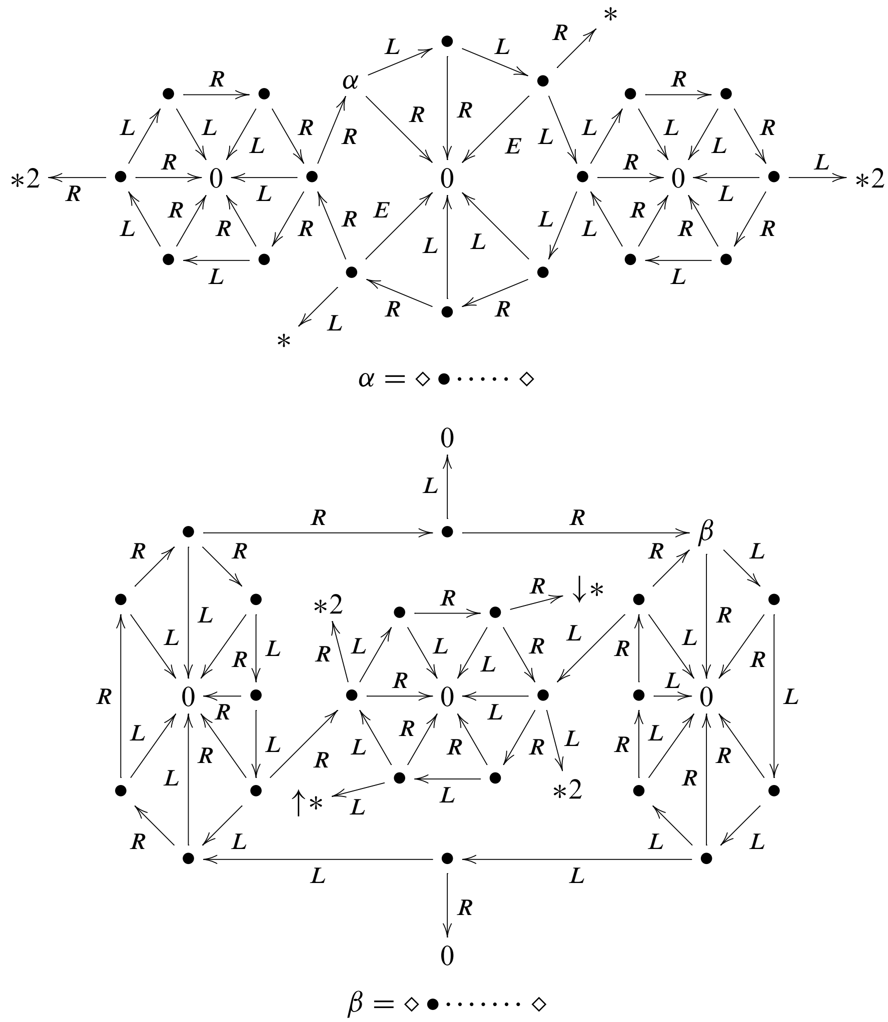

Figure 22. Some cycles that arise naturally in One-Dimensional Phutball.

One-Dimensional Phutball. Some extraordinary loopy positions in 1D Phutball were discovered jointly by Richard Nowakowski, Paul Ottaway, and myself. A few of these are shown in Figure 22. The game of Phutball, and the notation used here to describe the positions, are explained in [19]. It is interesting that although Phutball obviously allows for alternating cycles, all positions yet studied are equivalent to stoppers.

QUESTION. *Is every position in One-Dimensional Phutball equivalent to a stopper*?

These Phutball positions contain the most complicated loops yet detected. Moreover, the corresponding position on the 1×12 board ($\diamond \bullet \cdots\cdots\cdots \diamond$) is a stopper whose canonical game graph has 168 vertices and a 23-cycle. However, all of these examples are "tame" in the sense that every cycle alternates just once between Left and Right edges. It is possible to construct "wild" stoppers with more complicated cycles (see [30] in this volume), but nonetheless we have the following open problem.

OPEN PROBLEM. *Find a position in an actual combinatorial game* (*Phutball or otherwise*) *whose canonical form is a stopper containing a wild cycle*.

6. Impartial loopy games

Not surprisingly, *impartial* loopy games were studied long before partizan ones. In 1966, ten years before the publication of *On Numbers and Games*, Cedric A. B. Smith generalized the Sprague–Grundy theory to games with cycles.

For γ to be impartial, of course, infinite plays must be considered draws. We therefore have three outcome classes: the usual $\mathscr{N}$- and $\mathscr{P}$-positions, and also $\mathscr{D}$-positions (called $\mathscr{O}$-positions in *Winning Ways*).

Now consider an arbitrary impartial game γ. If all the options of γ are known to be nimbers $*a$, $*b$, $*c, \ldots$, then certainly $\gamma = *n$, where $n = \operatorname{mex}(a, b, c, \ldots)$: the usual Sprague–Grundy argument applies. But some games γ are equivalent to nimbers even though some of their options are not.

For instance, consider the example of Figure 23. It is not hard to see that $\gamma = *2$: in $\gamma + *2$, second player wins by mirroring moves to 0 or $*$; while if first player moves to $\delta + *2$, second player reverses to $*2 + *2 = 0$. However, the subposition δ is not equivalent to any nimber, since first player can always *draw* $\delta + *n$ by moving to the infinite loop.

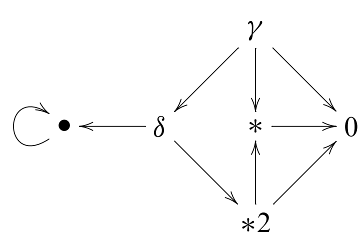

Figure 23. δ is not a nimber, but $\gamma = *2$.

Roughly speaking, $\gamma = *2$ because 2 is the mex of its nimber-valued options, and all *other* options reverse out, in the usual sense, to positions of value $*2$. Care is needed, however, to avoid circular definitions: the analysis of Figure 23

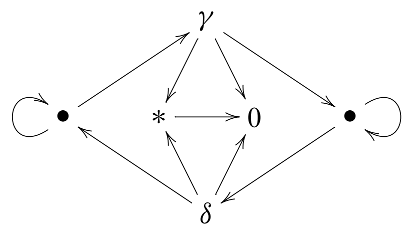

Figure 24. It's tempting to declare $\gamma = \delta = *2$ (cf. Figure 23); but $\gamma + *2$ and $\delta + *2$ are both draws.

works because the reversing move is "already known" to be $*2$. Indeed, Figure 24 shows that we cannot indiscriminately draw conclusions about the value of γ without a definite starting point.

These concerns led Smith to formulate the key notion of a rank function. The idea, motivated by Figures 23 and 24, is that we can safely assign Grundy values to subpositions of γ provided they are ranked in order of precedence. Formally,

DEFINITION 10. Let γ be a game, let $\mathscr{A}$ be the set of all followers of γ, and fix a partial function $G : \mathscr{A} \to \mathbb{N}$. Then G is a *Grundy function* if there exists a map $R : \mathscr{A} \to \mathbb{N}$ (a *rank function* for G) such that:

(i) If $G(\alpha) = n$ and $k < n$, then there is some option β of α with $R(\beta) < R(\alpha)$ and $G(\beta) = k$.
(ii) If $G(\alpha) = n$ and β is any option of α with $R(\beta) < R(\alpha)$, then $G(\beta) \neq n$.
(iii) Suppose $G(\alpha) = n$ and β is any option of α. If $G(\beta)$ is undefined, or if $R(\beta) \geq R(\alpha)$, then there exists an option δ of β with $R(\delta) < R(\alpha)$ and $G(\delta) = n$.

Conditions (i) and (ii) imply that $G(\alpha)$ obeys the mex rule, taken over all options of α with strictly lower rank. Condition (iii) implies that any remaining options reverse out to positions of lower rank than α. The main result is that there is a unique *maximal* Grundy function associated to γ (where G is maximal if its domain cannot be expanded).

THEOREM 11 (SMITH). *Let $G, H : \mathscr{A} \to \mathbb{N}$ be two Grundy functions for γ. If G and H are maximal, then $G = H$.*

PROOF. See [31, Section 9]. □

So we can safely refer to *the* Grundy function G of γ. It is a remarkable fact that G completely characterizes the behavior of γ.

LEMMA 12 (SMITH). *Let γ be a game with Grundy function G. If $G(\gamma) = n$, then $\gamma = *n$; if $G(\gamma)$ is undefined, then γ is not equal to any nimber.*

PROOF. See [31, Section 9]. □

When $G(\gamma)$ is undefined, we write

$$\gamma = \infty_{abc\cdots}$$

to mean that the nimber-valued followers of γ are exactly $*a$, $*b$, $*c$, ... We can now describe the outcome class of any sum of impartial games.
THEOREM 13 (SMITH).

(a) $\infty_{abc\cdots} + *n$ *is an* $\mathscr{N}$*-position if* n *is one of* $a, b, c, \ldots$; *otherwise it's a* $\mathscr{D}$*-position.*
(b) $\infty_{abc\cdots} + \infty_{def\cdots} + \cdots$ *is always a* $\mathscr{D}$*-position.*

PROOF. See [31, Section 9]. □

The parallel between Smith's theory and the classical Sprague–Grundy theory breaks down in one important respect. If $\gamma = *n$, then we can be quite certain that $\gamma + X$ and $*n + X$ have the same outcomes, *even when* X *is partizan*. However, there exist games α and β, both "equal to" ∞_0, whose outcomes are distinguished by a certain partizan game (see Figure 25). There is no contradiction: α and β indeed behave identically, provided they occur in sums *comprised entirely of impartial games*. One could say that the Sprague–Grundy theory embeds nicely in the partizan theory, while the Smith generalization does not.

$\alpha = \{0, \mathbf{dud} \mid 0, \mathbf{dud}\}$ $\qquad\qquad$ $\beta = \{0, \alpha \mid 0, \alpha\}$

Figure 25. Right can draw $\alpha + 1$ moving first, while $\beta + 1$ is a win for Left, no matter who moves first. Therefore $\alpha \neq \beta$; yet no *impartial* game distinguishes them.

This is an interesting fact, one that does not seem to appear elsewhere in the literature; and it raises an equally intriguing question:

OPEN PROBLEM. *Classify all impartial loopy games*, *relative to all partizan ones*.

Additional subtraction games. *Additional subtraction games* are just like ordinary subtraction games, except that their subtraction sets may contain *negative* numbers (so that players are permitted to *add* to a nonempty heap in certain fixed quantities). Such games are more interesting than one might expect. Several examples are mentioned in *Winning Ways*, and several related classes of games were studied by Fraenkel and Perl [12] and Fraenkel and Tassa [14] in the 1970s. The additional subtraction games cry out for further investigation.

OPEN PROBLEM. *Extend the analysis of additional subtraction games.*

The annihilation game. The *annihilation game* is an impartial game played on an arbitrary directed graph. At the start of the game, tokens are placed on the vertices of the graph, at most one per vertex. A move consists of sliding a token to an adjacent vertex, and whenever two tokens occupy the same vertex, they are both immediately removed from the game (the annihilation rule).

If the game is played on a loopfree graph, then the annihilation rule has no effect, since identical loopfree games ordinarily sum to zero. On loopy graphs, however, the effect is significant.

The annihilation game was proposed by Conway in the 1970s. Shortly thereafter, it was solved by Aviezri Fraenkel and his student Yaacov Yesha [16]. They specified a polynomial-time algorithm for determining the generalized Sprague–Grundy values of arbitrary positions. Interested readers should consult Fraenkel and Yesha's 1982 paper on the subject [17].

Infinite impartial games. The Smith–Fraenkel results completely resolve the disjunctive theory of finite impartial games. It is therefore natural to seek generalizations of the theory to infinite games. In the infinite case, one must allow ordinal-valued Grundy functions, even among loopfree games: for example, the game

$$*\omega = \{0, *, *2, *3, \dots\}$$

has Grundy value ω.

In the same paper that introduced the loopy Sprague–Grundy theory [31], Smith noted that his results generalize in a completely straightforward manner to infinite games with ordinal-valued Grundy functions. The definitions and theorems are essentially the same, with the functions G and R permitted to take on arbitrary ordinal values.

A more substantive result is due to Fraenkel and Rahat [13]. They identified a class of infinite loopy games whose Grundy values are nonetheless guaranteed to be finite. Their result can be summarized as follows:

DEFINITION 14. Let $\mathcal{G}$ be a graph. A *path* of $\mathcal{G}$ (of length n) is a sequence of *distinct* vertices

$$V_0, V_1, V_2, \dots, V_n$$

such that there is an edge directed from each V_i to V_{i+1}. We say that the path *starts at* V_0.

DEFINITION 15. Let $\mathcal{G}$ be a graph. A vertex V is said to be *path-bounded* if there is an integer N such that every path starting at V has length $\leq N$. $\mathcal{G}$ is said to be *locally path-bounded* if every vertex of $\mathcal{G}$ is path-bounded. (There need not exist a single bound that extends uniformly over all vertices.)

Note that all loopfree graphs are locally path-bounded.

THEOREM 16 (FRAENKEL–RAHAT). *Let γ be a (possibly infinite) impartial game. If the graph of γ is locally path-bounded, then the Grundy function for γ is finite wherever it is defined.*

PROOF. See [13, Section 3]. □

7. Conjunctive and selective sums

Although disjunctive sums have received the most attention, several authors have investigated the behavior of loopy games under other types of compound. The two most prominent are *conjunctive* and *selective* sums:

- In the *conjunctive sum* $\alpha \wedge \beta \wedge \cdots \wedge \gamma$, a player must move in every component. If any component is terminal, then there are no legal moves.
- In the *selective sum* $\alpha \vee \beta \vee \cdots \vee \gamma$, a player may move in any number of components (but at least one).

This line of research, like so many others, was pioneered by Cedric Smith [31], who focused on the impartial case. Smith's results are best described in terms of the *Steinhaus remoteness* of a position. If γ is a loopy game, we define the remoteness $R(\delta)$, for each follower δ of γ, as follows:

DEFINITION 17. Let γ be an impartial game, let $\mathscr{A}$ be the set of all followers of γ, and fix a function $R : \mathscr{A} \to \mathbb{N} \cup \{\infty\}$. Then R is a *remoteness function* provided that, for each $\delta \in \mathscr{A}$:

- If δ is terminal, then $R(\delta) = 0$.
- If $R(\alpha)$ is *even* for at least one option α of δ, then

$$R(\delta) = 1 + \min\{R(\alpha) : \alpha \text{ is an option of } \delta \text{ with } R(\alpha) \text{ even}\}.$$

- If $R(\alpha)$ is *odd* for every option α of δ, then

$$R(\delta) = 1 + \max\{R(\alpha) : \alpha \text{ is an option of } \delta\}.$$

It is not hard to check that every game admits a unique remoteness function R. The remoteness function tells us quite a bit about γ: it's a $\mathscr{P}$-position if $R(\gamma)$ is even, an $\mathscr{N}$-position if $R(\gamma)$ is odd, and a $\mathscr{D}$-position if $R(\gamma) = \infty$.

Furthermore, if the winning player strives to achive victory as quickly as possible, and the losing player tries to postpone defeat for as long as possible, then the magnitude of $R(\gamma)$ determines exactly how long the game will last.

Smith's main results are summarized by the following theorem.

THEOREM 18 (SMITH). *Let α, β, ..., γ be impartial loopy games. Then*:

(a) $$R(\alpha \wedge \beta \wedge \cdots \wedge \gamma) = \min\{R(\alpha), R(\beta), \ldots, R(\gamma)\}.$$

(b) *If $R(\alpha)$, $R(\beta)$, ..., $R(\gamma)$ are all even, then*

$$R(\alpha \vee \beta \vee \cdots \vee \gamma) = R(\alpha) + R(\beta) + \cdots + R(\gamma).$$

If $R(\alpha)$, $R(\beta)$, ..., $R(\gamma)$ are all finite, and k of them are odd ($k \geq 1$), *then*

$$R(\alpha \vee \beta \vee \cdots \vee \gamma) = R(\alpha) + R(\beta) + \cdots + R(\gamma) - k + 1.$$

Finally, if any of $R(\alpha)$, $R(\beta)$, ..., $R(\gamma)$ is infinite, then

$$R(\alpha \vee \beta \vee \cdots \vee \gamma) = \infty.$$

PROOF. See [31, Sections 6 and 7]. □

Theorem 18 enables us to find the outcome of any conjunctive or selective sum, provided we know the remoteness of each component. The remoteness function can therefore be regarded as an analogue of the Grundy function.

Partizan games. Smith's results were substantially extended by Alan Flanigan, who studied partizan loopy games under conjunctive and selective sums, as well as two additional types of compound, the *continued conjunctive* and *shortened selective* sums. We summarize Flanigan's results for conjunctive sums here. The remaining cases are beyond the scope of this paper; interested readers should consult Flanigan's 1979 thesis [9] and two subsequent papers [10; 11].

First note that we can define *partizan remoteness functions* R^L and R^R for γ. They are defined just as in the impartial case; but we only consider moves for the player in question, minimaxing over the *opponent's* remoteness function applied to each option.

DEFINITION 19. Let γ be a partizan game, let $\mathscr{A}$ be the set of all followers of γ, and fix functions $R^L, R^R : \mathscr{A} \to \mathbb{N} \cup \{\infty\}$. Then R^L, R^R are *partizan remoteness functions* provided that the following conditions (and their equivalents with Left and Right interchanged) are satisfied for each $\delta \in \mathscr{A}$:

- If δ has no Left options, then $R^L(\delta) = 0$.
- If $R^R(\delta^L)$ is *even* for at least one δ^L, then

$$R^L(\delta) = 1 + \min\{R^R(\delta^L) : R^R(\delta^L) \text{ is even}\}.$$

- If $R^R(\delta^L)$ is *odd* for every δ^L, then

$$R^L(\delta) = 1 + \max\{R^R(\delta^L)\}.$$

Smith's result for conjunctive sums is virtually unchanged in the partizan context.

THEOREM 20 (SMITH–FLANIGAN). *Let $\alpha, \beta, \ldots, \gamma$ be (partizan) loopy games. Then for $X = L, R$, we have*

$$R^X(\alpha \wedge \beta \wedge \cdots \wedge \gamma) = \min\{R^X(\alpha), R^X(\beta), \ldots, R^X(\gamma)\}.$$

PROOF. See [9, Chapter II.2]. □

Since the outcome class of γ is determined by the parities of $R^L(\gamma)$ and $R^R(\gamma)$, this is all we need to know.

Flanigan also noted that the analysis of conjunctive sums (but not selective sums) extends to infinite games: one can suitably define ordinal-valued remoteness functions, taking suprema instead of maxima when R is odd; then Theorem 20 generalizes verbatim.

8. Algorithms and computation

Computation is an essential part of combinatorial game theory. This is particularly true in the study of loopy games, since they are especially difficult to analyze by hand.

The basic algorithm for determining the outcome class of an impartial loopy game was introduced by Fraenkel and Perl [12] in 1975. The strategy is to iterate over all vertices V of the game graph of γ, assigning labels as summarized in Algorithm 1.

THEOREM 21 (FRAENKEL–PERL). *Algorithm* 1 *correctly labels the subspositions of γ according to their outcome classes, and concludes in time $O(n^2)$ in the number of vertices.*

PROOF. See [12, Section 3]. □

In fact, Fraenkel observes that we can improve slightly upon Algorithm 1: traverse the vertices of γ just once; and whenever a label is assigned to V, reexamine all unlabeled predecessors of V. With this modification, the algorithm runs in time $O(n)$ in the number of edges. Since game graphs tend to have relatively low edge density, this will usually be an improvement.

For each vertex V of the game graph of γ:

- If all options of V have been labeled $\mathscr{N}$, then label V by $\mathscr{P}$. (This includes the case where V is terminal.)
- If any option of V has been labeled $\mathscr{P}$, then label V by $\mathscr{N}$.

The algorithm continues until no more vertices can be labeled, whereupon all remaining vertices are labeled by $\mathscr{D}$.

Algorithm 1. Computing the outcome class of an impartial game γ.

Let $\mathcal{G}$ be the game graph of γ.

(i) Put $k = 0$.

(ii) For each vertex V of $\mathcal{G}$:

- If all options of V have been labeled $\mathscr{N}$, then label V by $\mathscr{P}$.
- If any option of V has been labeled $\mathscr{P}$, then label V by $\mathscr{N}$.

(iii) For each unlabeled vertex V, all of whose options are now labeled: if each option of V has an option labeled $\mathscr{P}$, then label V by $\mathscr{P}$ as well.

(iv) Label all remaining (unlabeled) vertices by $\mathscr{D}$.

(v) For each vertex V labeled $\mathscr{P}$, define $G(V) = k$ and remove V from $\mathcal{G}$.

(vi) If all remaining vertices of $\mathcal{G}$ are labeled $\mathscr{D}$, then stop: we are done.

(vii) Clear all $\mathscr{N}$ labels (but retain all $\mathscr{D}$ labels).

(viii) Put $k = k + 1$ and return to Step 2.

Algorithm 2. Computing the generalized Sprague–Grundy value of γ.

Fraenkel and Perl have also given an algorithm for computing the generalized Sprague–Grundy values of impartial loopy games (Algorithm 2); see Fraenkel and Yesha [18] for further discussion.

THEOREM 22 (FRAENKEL–PERL). *Algorithm* 2 *correctly defines the maximal Grundy function for* γ, *and concludes in time* $O(n^3)$ *in the number of vertices.*

PROOF. See [12, Section 4]. □

Algorithm 1 is virtually unchanged in the partizan case. Given a game γ with graph $\mathcal{G}$, one first constructs the corresponding *state graph* $\mathcal{S}$. The vertices of $\mathcal{S}$ consist of pairs (V, X), where V is a vertex of $\mathcal{G}$ and X is either L or R. There is an edge directed from (U, L) to (V, R) just if there is a Left edge directed from U to V, and so on. Algorithm 1 can then be applied directly to $\mathcal{S}$. This was noticed independently by Shaki [26], Fraenkel and Tassa [15], and Michael Albert [1].

Comparison. Algorithm 1 suffices to compare stoppers. Recall from Section 3 that if γ and δ are stoppers, then $\gamma \geq \delta$ if and only if Left, playing second, can survive $\gamma - \delta$. So to test whether $\gamma \geq \delta$, we simply compute the state graph of $\gamma - \delta$ and apply Algorithm 1. If V is the start vertex (corresponding to $\gamma - \delta$ itself), then $\gamma \geq \delta$ if and only if (V, R) is not marked $\mathscr{N}$.

One can extend these ideas in order to compare arbitrary games, but the algorithms are somewhat more involved. See [30] in this volume for a discussion.

Simplification and strong equivalence. Fraenkel and Tassa [15] studied various simplification techniques in detail. They identified certain situations in which one can safely simplify an arbitrary (free) loopy game γ. These techniques yield a good algorithm for determining whether γ is equivalent to a loopfree game. We summarize their results.

DEFINITION 23. Let γ be a free loopy game.

(a) A Left option γ^{L} is *strongly dominated* if Left, playing second, can *win* the game $\gamma^{L'} - \gamma^{L}$ for some other Left option $\gamma^{L'}$.
(b) A Left option γ^{L} is *strongly reversible* if Left, playing second, can *win* the game $\gamma - \gamma^{LR}$ for some Right option γ^{LR}.
(c) If δ is any free loopy game, then γ and δ are *strongly equivalent* if either player can *win* $\gamma - \delta$ playing second. In this case we write $\gamma \stackrel{*}{=} \delta$.

Strongly dominated and strongly reversible Right options are defined analogously.

Note that $\gamma \stackrel{*}{=} \gamma$ if and only if $\gamma - \gamma = 0$, i.e., if and only if γ is equivalent to a loopfree game.

THEOREM 24 (FRAENKEL–TASSA). *Let γ be a free loopy game and let δ be any follower of γ. Let γ' be obtained from γ by either*:

(a) *Replacing δ with a strongly equivalent game δ'*; *or*
(b) *Eliminating a strongly dominated option of δ*; *or*
(c) *Bypassing a strongly reversible option of δ*.

Then $\gamma = \gamma'$.

THEOREM 25 (FRAENKEL–TASSA). *Let γ be a free loopy game and assume that*:

(i) *γ is equivalent to a loopfree game* (*i.e.*, $\gamma - \gamma = 0$); *and*
(ii) *No follower of γ has any strongly dominated or strongly reversible options*.

Then γ is itself loopfree.

THEOREM 26 (FRAENKEL–TASSA). *Let γ be a free loopy game. If, for each subposition of γ, we repeatedly eliminate strongly dominated options and bypass strongly reversible ones, then the process is guaranteed to terminate. We will eventually arrive at a form for γ that contains no strongly dominated or strongly reversible options*.

Thus if γ is equivalent to a loopfree game, then Theorems 24 through 26 yield an algorithm for computing its canonical form: eliminate strongly dominated options and bypass strongly reversible ones until none remain.

Theorem 24 fails if the strong notions of domination and reversibility are replaced by their naive weakenings. This is a major obstacle to developing a general canonical theory of loopy games. These issues are discussed at length, and partially resolved, in [30] in this volume.

References

[1] M. Albert. Personal communication, 2004.

[2] E. R. Berlekamp. Idempotents among partisan games. In Nowakowski [25], pages 3–23.

[3] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning Ways for Your Mathematical Plays*. A. K. Peters, Ltd., Natick, MA, second edition, 2001.

[4] E. R. Berlekamp and M. Pearson. Entrepreneurial chess. In this volume.

[5] J. H. Conway. Loopy games. In B. Bollobás, editor, *Advances in Graph Theory*, number 3 in Ann. Discrete Math., pages 55–74, 1978.

[6] N. D. Elkies. On numbers and endgames: Combinatorial game theory in chess endgames. In Nowakowski [24], pages 135–150.

[7] N. D. Elkies. Personal communication, 2005.

[8] J. Erickson. New Toads and Frogs results. In Nowakowski [24], pages 299–310.

[9] J. A. Flanigan. *An Analysis of Some Take-Away and Loopy Partizan Graph Games*. PhD thesis, University of California, Los Angeles, 1979.

[10] J. A. Flanigan. Selective sums of loopy partizan graph games. *Internat. J. Game Theory*, 10:1–10, 1981.

[11] J. A. Flanigan. Slow joins of loopy games. *J. Combin. Theory, Ser. A*, 34(1):46–59, 1983.

[12] A. S. Fraenkel and Y. Perl. Constructions in combinatorial games with cycles. In A. Hajnal, R. Rado, and V. T. Sós, editors, *Infinite and Finite Sets, Vol. 2*, number 10 in Colloq. Math. Soc. János Bolyai, pages 667–699. North-Holland, 1975.

[13] A. S. Fraenkel and O. Rahat. Infinite cyclic impartial games. *Theoretical Comp. Sci.*, 252:13–23, 2001.

[14] A. S. Fraenkel and U. Tassa. Strategy for a class of games with dynamic ties. *Comput. Math. Appl.*, 1:237–254, 1975.

[15] A. S. Fraenkel and U. Tassa. Strategies for compounds of partizan games. *Math. Proc. Cambridge Philos. Soc.*, 92:193–204, 1982.

[16] A. S. Fraenkel and Y. Yesha. Theory of annihilation games. *Bull. Amer. Math. Soc.*, 82:775–777, 1976.

[17] A. S. Fraenkel and Y. Yesha. Theory of annihilation games, I. *J. Combin. Theory, Ser. B*, 33:60–82, 1982.

[18] A. S. Fraenkel and Y. Yesha. The generalized Sprague–Grundy function and its invariance under certain mappings. *J. Combin. Theory, Ser. A*, 43:165–177, 1986.

[19] J. P. Grossman and R. J. Nowakowski. One-dimensional Phutball. In Nowakowski [25], pages 361–367.

[20] S.-Y. R. Li. Sums of zuchswang games. *J. Combin. Theory, Ser. A*, 21:52–67, 1976.

[21] D. J. Moews. *On Some Combinatorial Games Connected with Go*. PhD thesis, University of California at Berkeley, 1993.

[22] D. J. Moews. Loopy games and Go. In Nowakowski [24], pages 259–272.

[23] H. J. R. Murray. *A History of Board-Games Other Than Chess*. Oxford University Press, 1952.

[24] R. J. Nowakowski, editor. *Games of No Chance*. Number 29 in MSRI Publications. Cambridge University Press, Cambridge, 1996.

[25] R. J. Nowakowski, editor. *More Games of No Chance*. Number 42 in MSRI Publications. Cambridge University Press, Cambridge, 2002.

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

[27] A. N. Siegel. *Loopy Games and Computation*. PhD thesis, University of California at Berkeley, 2005.

[28] A. N. Siegel. Loopy and loopfree canonical values in Hare and Hounds. To appear.

[29] A. N. Siegel. Backsliding Toads and Frogs. In this volume.

[30] A. N. Siegel. New results in loopy games. In this volume.

[31] C. A. B. Smith. Graphs and composite games. *J. Combin. Theory, Ser. A*, 1:51–81, 1966.

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