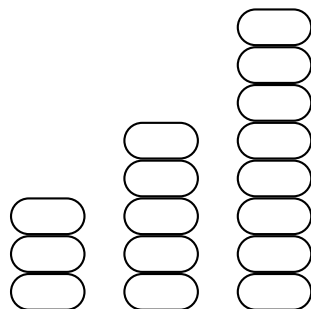# Combinatorial Games

1. Introduction

Combinatorial games are two-player games with no hidden information and no chance elements. They include child's play such as Tic-Tac-Toe and Dots and Boxes; mathematical abstractions "played" on arbitrary graphs or grids or posets; and some of the deepest and best-known board games in the world, such as Go and Chess.

The mathematical theory of combinatorial games pursues several interrelated objectives, including:

- exact solutions to particular games, usually in the form of an algebraic description of their outcomes;

- an understanding of the general combinatorial structure of games; and

- hardness results, suggesting that for certain games, or in certain situations, no concise solution exists.

In all of the games considered in this book, there will be just two players, **Left** and **Right**, who play alternately and whose moves affect the **position** in a manner defined by the rules of the game. Both players will have complete knowledge of the game state at all times ("no hidden information"), and the effect of each move will be entirely known before the move is made ("no chance elements"). We'll defer a formal definition until after we've introduced some examples.

Figure 1.1. A typical NIM position, with heaps of 3, 5, and 8 tokens.

NIM

NIM was among the earliest games to be subjected to a complete combinatorial analysis. It is played with several heaps of tokens; a typical position is shown in Figure 1.1. On her turn a player may remove any number of tokens (but necessarily at least one) from any *one* heap. The removed tokens are discarded and play continues. Whoever removes the last token wins.

Obviously any position containing just one (nonempty) heap is a first-player win, since whoever moves first will simply discard the entire heap. What about two-heap positions? If both heaps contain the *same number* of tokens, then the position is easily seen to be a second-player win: second player can simply respond to each of his opponent's moves with an identical move on the other heap, and (by induction) the game must eventually end. This is known as the **mirror-image strategy**. Conversely, a position containing exactly two *unequal* heaps is a first-player win, since first player can equalize the heaps and thereafter follow the mirror-image strategy.

A complete strategy for the game was discovered by the Harvard geometer Charles Bouton in 1902. Bouton's strategy is sufficiently simple to describe that we might as well do so now. Given two nonnegative integers a and b, the **nim-sum** $a \oplus b$ is given by the following computation: first write a and b in binary; then "add without carrying." For example, here's the calculation of $29 \oplus 21 \oplus 11 = 3$:

$$
\begin{array}{r}
11101 \\
\oplus\ 10101 \\
\oplus\ \ \ 1011 \\
\hline
=\ 00011
\end{array}
$$

If G is a NIM position with heaps $a_1, a_2, \ldots, a_k$, then the **nim value** of G is defined to be

$$a_1 \oplus a_2 \oplus \cdots \oplus a_k.$$

Call G a **zero position** if its nim value is zero. Then the strategy for Nim is neatly described:

Theorem 1.1 (Bouton's Theorem). *Let G be a* Nim *position.*

- *If G is a zero position, then every move from G leads to a nonzero position.*
- *If G is not a zero position, then there exists a move from G to a zero position.*

Therefore, if G is a zero position, second player can guarantee a win by reverting each of his opponent's moves to a new zero position. Since every move reduces the total number of tokens in play, the game will eventually end and second player will have the last move.

Likewise, if G is not a zero position, then first player can guarantee a win, simply by moving to any zero position.

For example, suppose G has three heaps, of sizes 29, 21, and 11. A moment ago we noted that $29 \oplus 21 \oplus 11 = 3$. This is nonzero, so the position must be a first-player win, and it is not hard to see that the unique winning move is to take three coins from the heap of size 11, leaving $29 \oplus 21 \oplus 8 = 0$.

Proof of Bouton's Theorem. First suppose $a_1 \oplus a_2 \oplus \cdots \oplus a_k = 0$ and consider a typical move from a_1 to a_1'. Necessarily $a_1 \neq a_1'$, so

$$a_1' \oplus a_2 \oplus \cdots \oplus a_k \neq a_1 \oplus a_2 \oplus \cdots \oplus a_k = 0.$$

Conversely, suppose $x = a_1 \oplus \cdots \oplus a_k \neq 0$. Consider the *most-significant bit*, say the j^{th} bit, of the binary representation of x. At least one of the a_i must have its j^{th} bit equal to 1; assume without loss of generality that it is a_1. Put $a_1' = x \oplus a_1$. Necessarily $a_1' < a_1$, since its j^{th} bit is equal to 0 and it agrees with a_1 on all higher-order bits. So there is a move available from a_1 to a_1', and

$$a_1' \oplus \cdots \oplus a_k = x \oplus a_1 \oplus \cdots \oplus a_k = x \oplus x = 0. \qquad \square$$

Outcomes and Solutions

Every Nim position G has a well-defined outcome: if G is a zero position, then second player can force a win; otherwise first player can force a win. This divides all Nim positions into two disjoint **outcome classes**.

The positions of most combinatorial games can likewise be partitioned into various outcome classes. While it's easy to prove that such outcomes *exist* (and we'll shortly do so for a wide class of games), actually computing them tends to be much harder. We can use Bouton's Theorem to determine

the outcome of any NIM position; extending this sort of analysis to more complicated games is a central goal of combinatorial game theory.

Note that Bouton's Theorem applies not to any particular starting position, but to all possible NIM positions, with any number of arbitrarily large heaps. When we say that NIM has been "solved," we mean that an efficient rule exists for calculating the outcome of *any* position. For most games (including NIM) it's straightforward to translate such a rule into a perfect winning strategy.

DAWSON'S KAYLES

DAWSON'S KAYLES is played with several rows of boxes, such as those shown in Figure 1.2(a). On his turn, a player must remove exactly two *adjacent* boxes from one of the strips. For example, he might play as shown in Figure 1.2(b), splitting the strip of size 10 into two smaller strips and leaving the position (c). Eventually all remaining boxes will be contained in strips of length one (or perhaps, no boxes will be left at all), after which no further moves are available.

There are two ways to play DAWSON'S KAYLES:

- In **normal play**, the player who makes the last move *wins*. This is the same convention used in NIM.

- In **misère play**, the player who makes the last move *loses*.
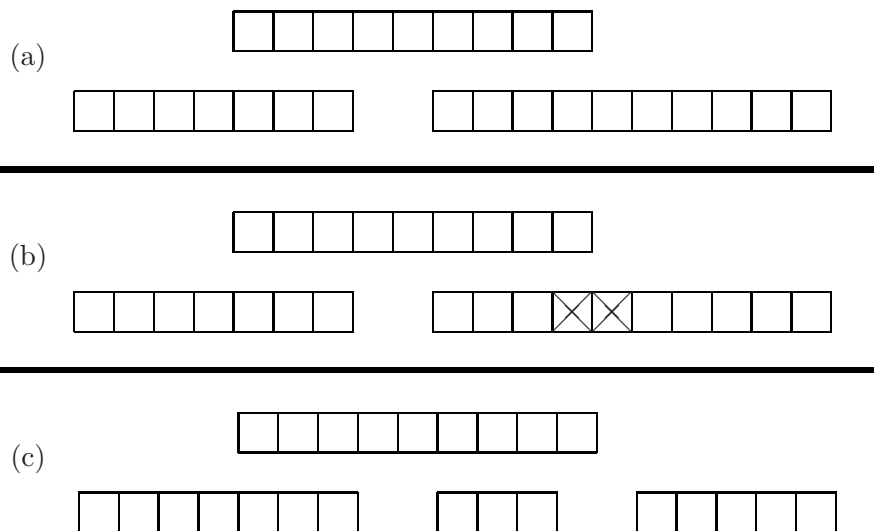
Figure 1.2. (a) A typical position in DAWSON'S KAYLES; (b) a typical move; (c) the resulting position.

A complete solution to the normal-play variant (on an arbitrary number of strips, of arbitrarily large size) was found more than fifty years ago by Richard Guy and Cedric Smith. They showed that a strip of $n + 34$ boxes behaves just like a strip of n for almost all n, the only exceptions being $n = 0, 15, 17, 32, 35, 52$. It therefore suffices to consider strips of a certain bounded size, and these can be analyzed exhaustively by generalizing the strategy for NIM. (Skip ahead to Figure IV.2.1 on page 186 for a full presentation of the Guy–Smith solution.)

DAWSON'S KAYLES, however, was first proposed as a misère game, and in misère play the situation is sharply different: no complete solution is known, despite significant effort by many authors. In fact misère DAWSON'S KAYLES is among the oldest unsolved problems in combinatorial game theory.[1]

At first glance it seems paradoxical that normal and misère play should so diverge in difficulty. After all, there is one sense in which the two forms of the game are equally hard: a given position G has exactly the same number of subpositions, regardless of the play convention. A brute-force analysis that computes the outcome of G by examining every subposition should therefore be no harder in misère play than in normal play.

The trouble is that the complexity of a brute-force analysis grows exponentially with the size of the position. Therefore, such an approach is not practical, nor does it grant much insight into the structure of the game. The key to the normal-play solution is the observation that any two strips in a DAWSON'S KAYLES position are **disjoint**—no single move can affect them both. In normal play, it is relatively easy to quantify the interactions between disjoint components. Each strip can therefore be analyzed independently, and this is what makes the Guy–Smith solution work.

Conversely, in misère play there is no simple characterization of the interaction between disjoint components. So although normal and misère play are similar from a brute-force perspective, the *combinatorics* of misère play are fundamentally more difficult. This combinatorial complexity is what distinguishes the two conventions.
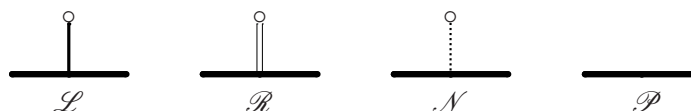
HACKENBUSH

HACKENBUSH is played on an undirected graph with edges colored blue, red, or green, such as shown in Figure 1.3 on the next page. Since this book is printed in black and white, we draw soLid lines for bLue edges, paRallel lines for Red edges, and dottEd lines for grEen edges.

---

[1]The game was introduced in 1935 by T. R. Dawson, in a slightly different form known as DAWSON'S CHESS. The two forms of the game are isomorphic—see Exercise 1.5.
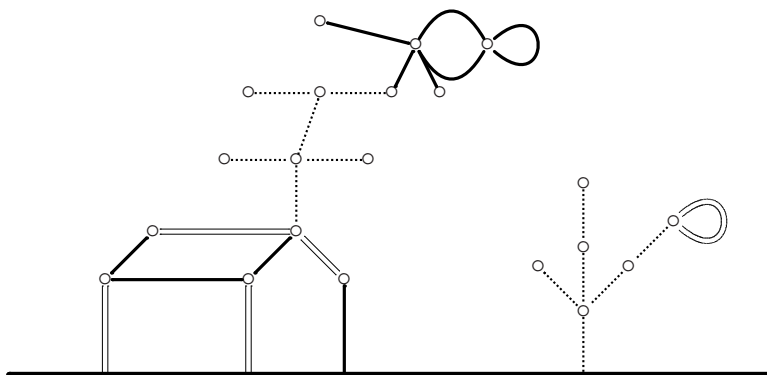
Figure 1.3. A typical HACKENBUSH position.

Figure 1.4. Four HACKENBUSH positions with distinct outcome classes.

The thick horizontal bar should be regarded as a single vertex of the graph, the **ground**. On her turn, Left may select any blue (solid) edge and remove it from the game; likewise Right may remove any red (parallel) edge. Green (dotted) edges may be removed by either player. As soon as an edge is removed, *all edges no longer connected to the ground are also removed.* Whoever chops the last edge wins.

There is one major difference between HACKENBUSH and the two previous examples. From a given HACKENBUSH position, the moves available to Left are distinct from those available to Right (except in the special case where every edge is green). Conversely, in both NIM and DAWSON'S KAYLES, both players have the same moves available at all times. Games such as HACKENBUSH, which distinguish between Left and Right, are called **partizan**; NIM-like games are **impartial**.

Impartial games like NIM and DAWSON'S KAYLES have just two outcome classes, but HACKENBUSH exhibits four:

> $\mathscr{N}$ first player (the *N*ext player) can force a win;
>
> $\mathscr{P}$ second player (the *P*revious player) can force a win;
>
> $\mathscr{L}$ *L*eft can force a win, no matter who moves first;
>
> $\mathscr{R}$ *R*ight can force a win, no matter who moves first.

$\mathscr{N}$ and $\mathscr{P}$ are familiar from NIM—they contain the nonzero and zero positions, respectively. But $\mathscr{L}$ and $\mathscr{R}$ are unique to partizan games, since they are biased in favor of one player or the other. Figure 1.4 illustrates simple HACKENBUSH positions representing each of the four possibilities.

Elements of $\mathscr{N}$ are called $\mathscr{N}$-**positions**, and likewise for $\mathscr{P}$, $\mathscr{L}$, and $\mathscr{R}$.
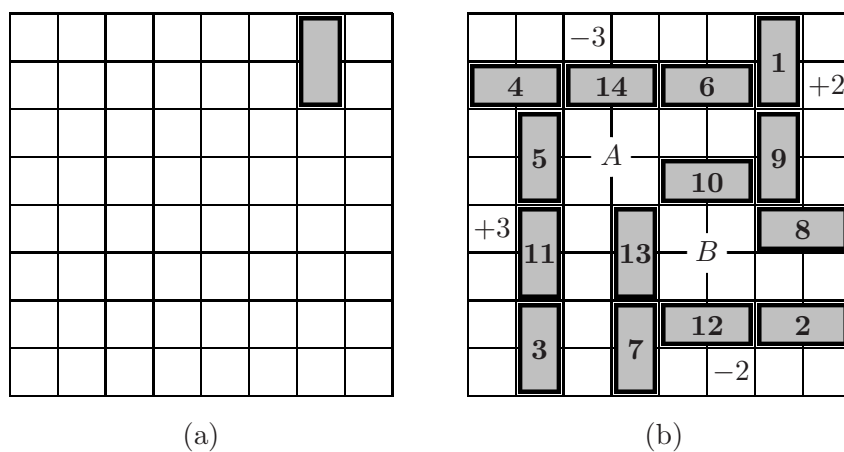
DOMINEERING

DOMINEERING (sometimes called CROSSCRAM) was invented by Göran Andersson around 1973. The game is usually played on an 8×8 checkerboard. Left and Right alternately place dominoes on the board, with the following restrictions:

- Each domino must cover exactly two adjacent squares of the grid. Dominoes may never overlap nor hang off the edge of the board.
- Left can only place verticaLly oriented dominoes; Right can only place hoRizontal ones.

The player who makes the last move wins. Since play is alternating, this is equivalent to placing the most dominoes, except that ties are broken in favor of second player.

DOMINEERING is a territorial game. Players will typically aim for moves like the one shown in Figure 1.5(a), which carves out space for another vertical domino and effectively locks Right out of the northeast corner.

Well-played games of DOMINEERING continue along this theme. Figure 1.5(b) shows Game 3 of the first (and so far only) World Domineering Championship match. After fourteen moves the board decomposes into six

(a)   (b)

Figure 1.5. (a) A typical opening move in DOMINEERING; (b) the first fourteen moves of Wolfe–Calistrate 1994, Round 3.

independent regions. Four of them are simple $1 \times n$ or $n \times 1$ strips, each of which is clearly worth an integer number of dominoes no matter how the game progresses. The remaining two regions can't be assigned numerical values in quite the same way: the final score on region A depends upon who makes the first move *within A*, and likewise for B.

Nevertheless, it's easy to check that Left's best move on region A yields a net score of $+2$, and Right's best move there yields a net score of 0, and we can indicate this by writing $A = \{2 \mid 0\}$. By similar reasoning we find that $B = \{0 \mid -2\}$. If Left moves on A, Right can respond on B, and vice versa; so Right can guarantee a net score of 0 on the combination $A + B$. Since the other four regions also cancel each other out, we conclude that Left and Right will ultimately place equally many dominoes. Therefore Right, as second player, will win.

Games, Options, Rulesets

Throughout this text, we'll use the term **game** to refer to an *individual position* in a combinatorial game, and we'll call a system of playable rules a **ruleset** instead. For example, DOMINEERING is a ruleset, and the specific DOMINEERING position pictured in Figure 1.5(b) is a game. This convention reflects the fact that individual positions are the primitive objects of our theory. We'll keep the distinction clear by using capital Roman letters

$$G, H, \ldots$$

to denote combinatorial games and capital Greek letters

$$\Gamma, \Lambda, \ldots$$

to denote rulesets.

The following definitions are straightforward and fundamental.

Definition 1.2. Let G and H be combinatorial games. H is a **Left option** (resp. **Right option**) of G if Left (resp. Right) can move directly from G to H.

We usually denote a typical Left option of G by G^L and a typical Right option by G^R.

Definition 1.3. Let G and H be combinatorial games. H is a **subposition** of G if there exists a sequence of consecutive moves (for either player, not necessarily alternating) leading from G to H.

For example, Figure 1.5(b) is a subposition of Figure 1.5(a).

The definition of subposition includes the case where the sequence of moves is empty, so that every game G is considered a subposition of itself. We'll use the term **proper subposition** in order to exclude this case.

Definition 1.4. Let G be a combinatorial game.

(a) A **run** of G of length $k < \infty$ is a sequence of positions

$$G_0, G_1, G_2, \ldots, G_k$$

such that $G_0 = G$, and each G_{i+1} is an option of G_i. Likewise, a run of G of length ∞ is a sequence

$$G_0, G_1, G_2, \ldots.$$

(b) An **alternating run** is a run in which successive moves alternate between Left and Right, i.e., if G_{i+1} is a Left option of G_i, then G_{i+2} is a Right option of G_{i+1}, etc.

(c) An alternating run of length k is said to be a **play** of G if either $k = \infty$ or else G_k has no options for the player to move.

The Fundamental Theorem

We'll now show that a wide class of games can be partitioned into the four outcome classes $\mathscr{L}$, $\mathscr{R}$, $\mathscr{N}$, and $\mathscr{P}$. A combinatorial game G is **short** provided that:

- G has just finitely many distinct subpositions; and
- G admits no infinite run.

It's easy to check that every position in NIM, DAWSON'S KAYLES, HACKENBUSH, and DOMINEERING is short. For example, suppose G is a DOMINEERING position with n empty squares. Then G has at most 2^n distinct subpositions, and infinite runs are impossible since every move strictly decreases the number of empty spaces.

Theorem 1.5 (Fundamental Theorem of Combinatorial Game Theory)**.** *Let G be a short combinatorial game, and assume normal play. Either Left can force a win playing first on G or else Right can force a win playing second, but not both.*

Proof. Consider a typical Left option of G, say G^L. Since G is short, G^L must have strictly fewer subpositions than G. Therefore we may assume by induction (and symmetry) that either Right can force a win playing first on G^L, or else Left can force a win playing second.

If Right can win *all* such G^L playing first, then certainly he can win G playing second regardless of Left's opening move. Conversely, if Left can win

any such G^L playing second, then he can win G by moving to it. Exactly one of these two possibilities must hold.   □

The proof of Theorem 1.5 is straightforward, but it contains several interesting features. First, the theorem has an obvious dual, in which "Left" and "Right" are interchanged. This is true of many theorems about partizan games, and we will usually state and prove just "one side" of each result.

Alert readers will also notice that while induction is central to the proof, there is no explicit mention of a base case. This is deliberate: there is an *implicit* base case in which Left has no options at all—in which case the inductive hypothesis is satisfied vacuously. Throughout this book we'll see many more inductive arguments whose base case is collapsed into the induction.

The Fundamental Theorem shows that every short game belongs to one of the four normal-play outcome classes $\mathscr{L}$, $\mathscr{R}$, $\mathscr{N}$, or $\mathscr{P}$. We denote by $o(G)$ the outcome class of G.

Disjunctive Sum

All four games—Nim, Dawson's Kayles, Hackenbush, and Domineer-ing—share a fundamental property: a typical position can be broken down into several independent components. For example, the three nim-heaps in Figure 1.1 are disjoint, in the sense that no single move may affect more than one of them. The outcome of the full position G therefore depends only on the structure of its individual heaps, and it makes sense to write

$$G = H_3 + H_5 + H_7,$$

where H_n denotes a heap of size n.

Similarly, the Dawson's Kayles position of Figure 1.2(a) consists of three disjoint strips. Moreover, after the opening move shown in Figure 1.2(b), the strip of length 10 is divided into two smaller strips. Here we may speak of a move from

$$S_7 + S_9 + S_{10}$$

to

$$S_7 + S_9 + S_3 + S_5,$$

where S_n is a strip of length n, and it is clear that any position can be written this way. In both Nim and Dawson's Kayles, the problem of finding the *outcome* of an arbitrary position reduces to understanding the *structure* of individual heaps (or strips).

Likewise, in the HACKENBUSH position of Figure 1.3 the house and the flower can never interact; and in the DOMINEERING position of Figure 1.5(b) there are six disjoint components.

This modularity is the driving force of combinatorial game theory. Given a combinatorial game G, it is often impractical to undertake a brute-force analysis of G itself. Instead, we analyze the components of G individually and attempt to extract information that will tell us something useful when they are pieced back together. The precise nature of this "information" is the main topic of this book.

To obtain a cohesive theory, it's necessary to formalize the notion of "piecing together" a game's components. If G and H are any two combinatorial games, we define

$$G + H,$$

the **disjunctive sum** of G and H, as follows. Copies of G and H are placed side-by-side. On his turn, a player must move in *either G or H* (but not both). Likewise, in the sum

$$G_1 + G_2 + \cdots + G_k,$$

a player must move in *exactly one* of the components G_i.

It should be clear that disjunctive sum accurately models the decompositions that we've observed in all four examples. Here, rather than "decomposing" complex positions, we regard the components themselves as primitive building blocks. This approach is far more flexible. We can certainly assemble the components of a position G back up into G itself; but we can also form new sums $G + H$, where H is unrelated to G and perhaps even drawn from a different ruleset. For example,
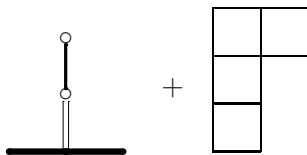
is a perfectly reasonable combinatorial game.

This flexibility has far-reaching consequences. It helps to expose shared structure across many rulesets, and it permits direct quantitative comparisons between them.

The Fundamental Equivalence

Now for the point! Let G and H be short games, and write

(†) $\qquad G = H \quad \text{if} \quad o(G + X) = o(H + X) \text{ for all } X.$

Here X ranges over *all conceivable* short games—including, for example, all Domineering positions, even if G or H happens to be a Hackenbush position. In such cases the definition of $+$ is the same as if we were summing components drawn from the same ruleset. We'll make all of this precise in Chapter II, when we give a formal definition of short game.

The relation given by (†) is the **fundamental equivalence**, and in various forms it will drive virtually every idea that we explore in this book. It's easily seen to be an equivalence relation, and the equivalence class of G modulo $=$ is known as the **game value** of G. More precisely, it is the **normal-play short partizan game value** of G, but normally we'll drop the cumbersome adjectives when it's clear which type of value we're talking about. The set of game values is denoted by

$$\mathbb{G}$$

and we'll spend most of Chapters II and III investigating its structure.

Notice that equality is a *defined relation* given by the fundamental equivalence. This is a standard convention in combinatorial game theory. When we wish to indicate that G and H have isomorphic structure, we'll write the notation

$$G \cong H$$

instead. Isomorphism is a stronger condition than equality: it is often the case that $G = H$ even when $G \ncong H$. (Conversely, $G \cong H$ always implies $G = H$.)

There is a strong analogy between $\mathbb{G}$ and the set of rational numbers $\mathbb{Q}$. The statement

$$G = H$$

for distinct games G and H should be no more troubling than the assertion

$$\frac{2}{4} = \frac{3}{6}.$$

Here G and H are distinct *games* that represent the same *game value*, just as $2/4$ and $3/6$ are distinct *ratios* that represent the same *rational number*.

Example. We have the following Hackenbush identity:

This is slighly less trivial than it seems, since from the left-hand side G Left's only move (up to symmetry) is to

while from the right-hand side H she can move to either

$$H^{L_1} = \qquad \text{or} \quad H^{L_2} = \qquad .$$

So in fact $G \ncong H$. But it's easy to see that in any sum

$$H + X,$$

Left's extra option H^{L_2} is no help. Any time she could win with a move to $H^{L_2} + X$, she might just as well play to $H^{L_1} + X$ instead and then ignore the extra stalk for the remainder of play.

Example. The fundamental equivalence provides a means to compare positions from different rulesets. For example, it's easy to see that

$$\qquad = \qquad .$$

In fact the two positions are isomorphic, since they both have one Left option to the empty game and no Right option. There's nothing profound in this observation, of course, but you might be surprised to learn that

Figure 1.6. A nontrivial identity between HACKENBUSH and DOMINEERING.

Let's take a moment to appreciate what this means. The DOMINEERING grid H in Figure 1.6 is a complicated game—there are fifteen distinct opening moves (nine for Left and six for Right), even after accounting for symmetries; and we can imagine that in a sum

$$H + X$$

it might become quite difficult to work out all the possibilities needed to determine the outcome. But knowing that H is equal to the much simpler HACKENBUSH stalk G, we can instead compute the outcome of

$$G + X,$$

and the fundamental equivalence gives $o(H + X) = o(G + X)$. In Chapter II we'll develop several tools for establishing identities like the one in Figure 1.6.

Other Kinds of Values

The definition of equality

$$G = H \quad \text{if} \quad o(G + X) = o(H + X) \text{ for all } X$$

is sensitive to three primitive parameters:

- the definition of the *outcome* $o(G)$;
- the definition of the *sum* $G + H$; and
- the domain of the *for all* quantifier over X.

In arriving at $\mathbb{G}$ we took $o(G)$ to be the *normal-play outcome* of G; $G + H$ to be the *disjunctive sum* of G and H; and *for all* to range over the set of short partizan games. But there are many other natural choices for all three primitives. For example, $o(G)$ might just as well be defined as the *misère-play outcome* of G, and we'll discuss variant meanings of "sum" and "for all" in Section 4.

Astonishingly, in virtually all cases—no matter how we choose to fix meanings for these three parameters—the fundamental equivalence remains the same:

$$G = H \quad \text{if} \quad o(G + X) = o(H + X) \text{ for all } X.$$

Each time we vary the meanings of "outcome," "sum," and "for all," the fundamental equivalence yields a different theory—but it is always a coherent theory. The fundamental equivalence is the glue that holds combinatorial game theory together, and we'll turn to it for guidance again and again throughout this book.

Exercises

1.1   Determine all winning moves from the Nim position with heaps 18, 22, and 29.

1.2   Let G be a Nim position. Let n be the number of *distinct* winning moves on G. Prove that either $n = 0$ or n is odd.

1.3   Determine the strategy for misère Nim.

1.4   Let G be an impartial game. Suppose that there exists an option H of G such that every option J of H is also an option of G. Prove that G must be an $\mathcal{N}$-position.

1.5   Dawson's Chess is played with two rows of pawns on a $3 \times n$ board:
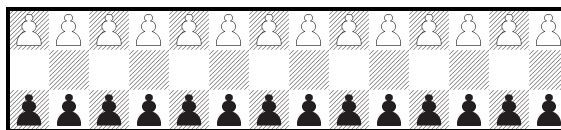
The pawns move and capture like ordinary CHESS pawns, except that capture is *mandatory*. (If a choice of captures is available, then the player may select either one.) Show that DAWSON'S CHESS on a $3 \times n$ board is isomorphic to DAWSON'S KAYLES with $n + 1$ boxes, regardless of which play convention (normal or misère) is observed. (This is the original form of the game proposed by Dawson [**Daw73**].)

1.6   No position in BLUE–RED HACKENBUSH is an $\mathcal{N}$-position.

1.7   Determine the outcomes of 4×4, 5×4, and 5×5 DOMINEERING.

1.8   Let S_L and S_R be sets of positive integers. The **partizan subtraction game** on S_L and S_R is played with a single heap of n tokens. On her move, Left must remove k tokens for some $k \in S_L$; likewise, on his move Right removes k tokens for some $k \in S_R$. Denote by $o(H_n)$ the (normal-play) outcome of a heap of n tokens. Prove that the sequence $n \mapsto o(H_n)$ is periodic, i.e., there is some $p > 0$ such that $o(H_{n+p}) = o(H_n)$ for all sufficiently large n.

2. HACKENBUSH: A Detailed Example

Many of the ideas and principles of combinatorial game theory are neatly illustrated by various situations that arise in HACKENBUSH. The simplest position is the game with no options at all,

which is called 0, the **empty game**. We'll write

$$0 = \{\,|\,\}$$

to indicate that there are no options for either player.

Introducing a single blue edge gives the game

$$\underline{\quad} = \left\{ \underline{\quad\quad} \;\middle|\; \right\} = \{0 \mid \}$$

which is an $\mathcal{L}$-position, since Left can win immediately by moving to 0, while Right has no move at all. We'll call it 1, as it behaves like one spare move for Left.

What can one say about positions like the one in Figure 2.1, in which every disjoint component is monochromatically red or blue?

Clearly Left will always play to remove just one blue edge; there is nothing to be gained by deleting more than one at a time. Likewise Right will always play to delete a single red segment. Each player is powerless to disrupt the other's plans, so the win will ultimately go to whoever started with more edges.

We can therefore replace every blue component with a positive integer number of edges, and every red component with a negative integer, and total
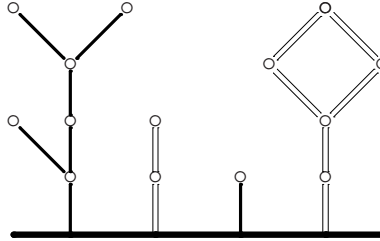
Figure 2.1. A HACKENBUSH position in which every disjoint component is monochromatic.

up the result. It's easy to see that for a position G of total value n,

$$o(G) = \begin{cases} \mathscr{L} & \text{if } n > 0; \\ \mathscr{R} & \text{if } n < 0; \\ \mathscr{P} & \text{if } n = 0. \end{cases}$$

In the third case there are an equal number of red and blue edges, and the win goes to whoever plays second.

For example, in Figure 2.1 the relevant number is

$$6 - 2 + 1 - 6 = -1,$$

so Right wins no matter who moves first.

These integer values are a sort of partizan analogue of the nim values that we saw in Section 1. Unlike nim values, they're *signed* and they add like ordinary integers. Also unlike nim values, they needn't be integers! Consider the game

$$G = \quad \text{} \quad = \left\{ \text{} \;\middle|\; \text{} \right\} = \{0 \mid 1\}.$$

What is the value of G? It's weaker for Left than 1, since Right has more flexibility. In fact Right can win quite handily on $G + (-1)$:

no matter who moves first, just by moving preferentially on G. We'll soon prove the remarkable identity

so that two copies of G behave like one copy of 1, and we may say that G has value $\frac{1}{2}$.

Zero Positions

A **zero position** is a game G with $G = 0$, that is, such that

$$o(G + X) = o(X) \quad \text{for all } X.$$

If G is a zero position, then it can be ignored in any sum where it appears. Clearly every zero position is a $\mathscr{P}$-position, since $o(G) = o(0) = \mathscr{P}$. Remarkably, the converse is also true.

Theorem 2.1. *If $o(G) = \mathscr{P}$, then G is a zero position.*

Proof. We must show that

$$o(G + X) = o(X)$$

for every short game X. By symmetry, it suffices to show that whenever Left can win X playing first (resp. second), she can also win $G + X$ playing first (resp. second).

But suppose Left can win X playing second. Then on $G + X$, she can respond to each of Right's moves with a winning move on the same component. This guarantees that Left will never run out of moves, and since $G + X$ is short, the game must eventually end. Thus Left can win $G + X$ playing second.

Likewise, suppose Left can win X playing first, say to X^L. Then on $G + X$, she can make an opening move to $G + X^L$ and thereafter follow the strategy just described. $\qquad\square$

Now let $-G$ denote the **negative** of G, in which the roles of Left and Right are reversed. In Hackenbush this means interchanging blue and red edges, leaving green ones unaffected. Symbolically, we may write

$$-G = \left\{ -G^R \mid -G^L \right\}$$

and we'll write $G - H$ as shorthand for $G + (-H)$. Clearly $o(G - G) = \mathscr{P}$, following the mirror image strategy, so in fact $G - G = 0$.

Theorem 2.2. *If $o(G - H) = \mathscr{P}$, then $G = H$.*

Proof. We have $G - H = 0$, so $G - H + H = H$. But $H - H = 0$, so also $G - H + H = G$. Therefore $G = H$. $\qquad\square$

Therefore in order to prove that G and H are equal, it suffices to show that their difference is a second-player win. Let's apply this technique to Hackenbush.

Half a Point

To prove the identity

$$\text{(figure)} \quad + \quad \text{(figure)} \quad = \quad \text{(figure)}$$

it suffices to show that the difference $\frac{1}{2} + \frac{1}{2} - 1$ is a second-player win:

$$\text{(figure)} \quad + \quad \text{(figure)} \quad + \quad \text{(figure)} \quad .$$

If either player moves on either copy of $\frac{1}{2}$, then her opponent can respond on the other, leaving a position of value $1 + 0 - 1 = 0$. In other words, the two highlighted moves below "cancel each other out":

$$\text{(figure)} \quad + \quad \text{(figure)} \quad + \quad \text{(figure)} \quad .$$

The only remaining option is Right's opening move on -1, but this clearly loses, since it leaves an $\mathscr{L}$-position.
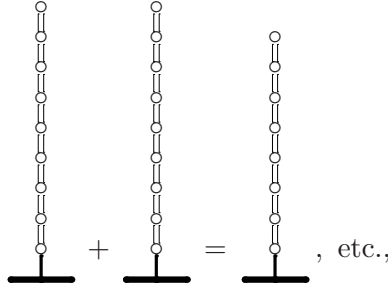
What about the game

$$\text{(figure)} \quad ?$$

A similar argument shows that

$$\text{(figure)} \quad + \quad \text{(figure)} \quad = \quad \text{(figure)}$$

so that we might call this game $\frac{1}{4}$. Proceeding like this, we obtain games of all values $1/2^n$:

$$+ \qquad = \qquad , \text{ etc.,}$$

and we can construct games of value $m/2^n$ by adding together an appropriate number of copies of $1/2^n$.

Now writing out the options of $\frac{1}{4}$, we see that

$$\frac{1}{4} = \qquad = \left\{ 0 \mid \frac{1}{2}, 1 \right\}.$$
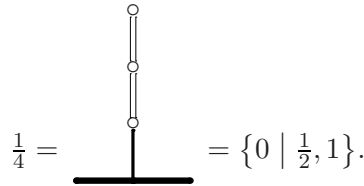
However it's clear that Right should always prefer to chop the higher of his two edges first. The reasoning is the same as in the monochromatic case: if he can't eliminate any of Left's edges, then Right should strive to remove as few of his own as possible. In fact it's not hard to check directly that any time Right can win a sum of games $1 + X$ (playing first or second), then he can also win $\frac{1}{2} + X$, so that he will always be content to replace the component 1 by $\frac{1}{2}$.

So Right's move from $\frac{1}{4}$ to 1 is **dominated** by $\frac{1}{2}$, and we can write

$$\tfrac{1}{4} = \left\{ 0 \mid \tfrac{1}{2} \right\},$$

which is easily verified by showing that $\frac{1}{4} - \left\{ 0 \mid \frac{1}{2} \right\}$ is a second-player win.

A similar line of reasoning shows that

$$2^{-(n+1)} = \left\{ 0 \mid 2^{-n} \right\}$$

for all $n \geq 0$, since Right's remaining options 1, $\frac{1}{2}$, ..., $2^{-(n-1)}$ are all dominated.

GREEN HACKENBUSH

Recall that Either player is permitted to chop a grEen (dottEd) edge; for example, on

$$G = \qquad = \{ 0 \mid 0 \}$$
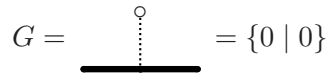
either player may move to 0. Consider the sum $G + 2^{-n}$:
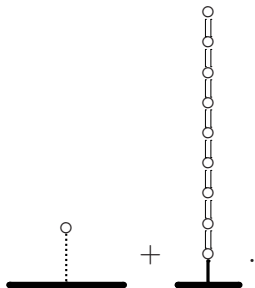
Left can win this game by following a simple strategy: chop the green edge preferentially or, if it's already gone, move to 0. In fact G is dominated by even the smallest numbers, and we might say that

$$\frac{1}{2^n} > G > -\frac{1}{2^n}$$

for all n. But $G \neq 0$, since it's a first-player win! Moreover,

is a second-player win, so that $G + G = 0$, or equivalently $G = -G$. The value of G isn't a number at all, and in fact it's **infinitesimal** with respect to numbers like 2^{-n}. It has the special name

$$*$$

(pronounced "star"), and in Chapter II we'll see that it plays a central role in the theory of combinatorial games.

Now $*$ is obviously isomorphic to a nim-heap of size 1, and if we play GREEN HACKENBUSH with long stalks, it's clear that we're just playing NIM. For example, Bouton's Theorem tells us that
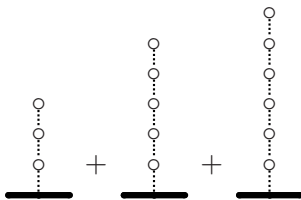
is a second-player win. We previously assigned these components **nim values** of 3, 5, and 6, respectively; but in the more general context of **game values** we need to call them something else, to distinguish them from ordinary integers. So we'll use the notation $*m$ (pronounced "star m") to denote a nim-heap of size m, and we'll write (for example)

$$*3 + *5 + *6 = 0.$$

It's clear that *any* nim-heap is dwarfed by 2^{-n}.

In BLUE–RED HACKENBUSH we could compute the value of *monochromatic* positions just by counting up the number of edges. This doesn't work in GREEN HACKENBUSH, though, and it's worth taking a moment to confirm that

$$+ \qquad = 0$$
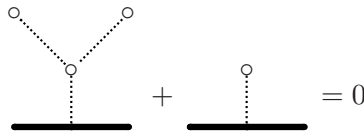
so that the value of the first component is $*1$, even though it has three edges. Fortunately there's a simple rule you can use to work out the values of arbitrary GREEN HACKENBUSH trees (see Exercise 2.1).

Tricolor HACKENBUSH

Let's consider one more example, the position

$$G = \qquad .$$
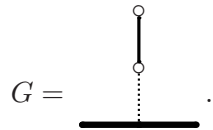
G is a first-player win, but it's clearly more favorable to Left than $*$, and in fact two copies of G make an $\mathscr{L}$-position. For on

$$G + G =$$
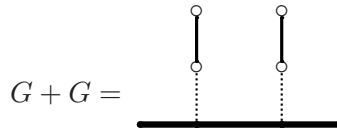
the *first* player to chop a green edge will lose, and Left can ensure this will always be Right. But $G + G$ is still infinitesimal: Right can win the sum

$$G + G - 2^{-n} =$$
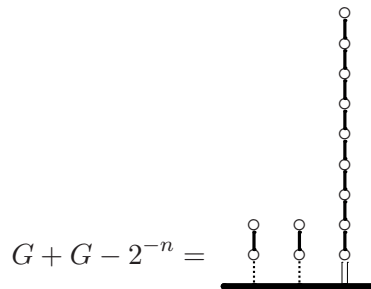
using the familiar strategy of chopping green edges preferentially.

Exercises

2.1   A HACKENBUSH position G is a **tree** if there is a unique path from the ground to each node. Let G be a tree that is completely green. Suppose that v is a vertex of G and the subtree above v consists of simple paths extending from v (i.e., every vertex above v has valence ≤ 2). Let $a_1, \ldots, a_k$ denote the lengths of these paths. Show that the subtree above v can be replaced by a single path

of length $a_1 \oplus \cdots \oplus a_k$, without changing the value of G. Show furthermore that by repeatedly applying this transformation, G can be reduced to a single stalk.

3. How to Read This Book

Combinatorial Game Theory is divided into eight chapters. Chapter I is primarily expository; each of the remaining seven chapters describes a major topic in detail.

The basic theory of short partizan games is developed in Chapter II; then Chapter III applies this theory to study the abstract algebraic and combinatorial structure of $\mathbb{G}$. Short impartial games are studied in Chapters IV (normal play) and V (misère play), with brief detours into loopy impartial (Section IV.4) and misère partizan (Section V.6) games.

The remaining chapters discuss various extensions of the partizan theory: loopy games in Chapter VI; orthodox temperature theory in Chapter VII; and transfinite games, including the surreal numbers, in Chapter VIII.

These eight chapters are further subdivided into a total of thirty-nine sections. It is not necessary to read all of this material straight through; the approximate interdependence of chapters is shown in Figure 3.1.

Chapters are numbered with Roman numerals, and sections with Arabic numerals. We'll ordinarily omit the chapter number when referring to a section *within the same chapter*. For example, you are currently reading Section 3, but in later chapters it will be known as Section I.3. The same
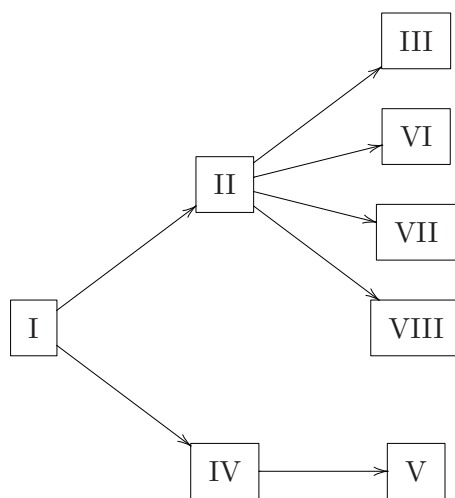
Figure 3.1. Approximate interdependence of chapters.

applies to theorems and exercises: the Fundamental Theorem (page 9) will be known as Theorem 1.5 from within Chapter I, but as Theorem I.1.5 elsewhere.

Most sections end with a list of pertinent exercises, as well as bibliographic notes summarizing results and references not covered in the body of the text. These notes are the recommended starting point for further investigation into each topic.

A full understanding of combinatorial game theory requires fluency in a wide range of mathematical subject matter—principally algebra, but also various ideas from enumerative combinatorics, set theory, and computational complexity. Some of the necessary background material is summarized in Appendix B, with an emphasis on key results used elsewhere in the text. Most results in this appendix are stated without proof and are intended primarily for review; references are included for other sources that are appropriate for more detailed study.

Notation

Figure 3.2 summarizes our naming conventions for various mathematical entitites. A particularly thorny problem is the distinction between a *game* and its *value*. Most of the time, we'll address this problem simply by suppressing the distinction—passing freely from a game to its value, or from a value to one of its representatives. This treatment requires a little extra care, but experience has shown it to be the least cumbersome approach. Therefore, we use capital Roman letters $G, H, J, \ldots$ to represent both games and game values, as needed.

Capital Roman	$G, H, J, X, Y, \ldots$	Games and values
Lowercase Roman	$x, y, z, \ldots$	Numbers
	$a, b, m, n, \ldots$	Integers
Lowercase Greek	$\alpha, \beta, \gamma, \ldots$	Ordinals
Script	$\mathscr{A}, \mathscr{B}, \mathscr{N}, \mathscr{P}, \ldots$	Sets of games
Calligraphic	$\mathcal{A}, \mathcal{B}, \ldots$	Sets of values
Capital Greek	$\Gamma, \Lambda, \ldots$	Rulesets
Boldface Roman	$\mathbf{ON}, \mathbf{PG}, \mathbf{SN}, \ldots$	Proper classes
Sans Serif	$\mathsf{P}, \mathsf{NP}, \mathsf{PSPACE}, \ldots$	Complexity classes

Figure 3.2. Default naming conventions for variables in this book.

On the other hand, we'll take care to distinguish between *sets* of games and *sets* of values, consistently using script letters $\mathscr{A}, \mathscr{B}, \ldots$ to denote sets of games, and calligraphic letters $\mathcal{A}, \mathcal{B}, \ldots$ for sets of values.

The following symbols are standard:

- $\mathbb{N}$ is the set of nonnegative integers, including 0.

- $\mathbb{N}^+$ is the set of strictly positive integers.

- $\mathbb{Z}$ is the ring of all integers.

- $\mathbb{Z}_n$ is the ring of integers modulo n (i.e., the quotient $\mathbb{Z}/n\mathbb{Z}$).

- $\mathbb{Q}$ is the field of rational numbers.

- $\mathbb{D}$ is the ring of dyadic rational numbers (those whose denominator is a power of 2).

- $\mathbb{R}$ is the field of real numbers.

- $\mathbb{R}^+$ is the set of strictly positive real numbers.

If $x \in \mathbb{R}$, then $\lfloor x \rfloor$ denotes the greatest integer $\leq x$ (the **floor** of x), and $\lceil x \rceil$ denotes the least integer $\geq x$ (the **ceiling** of x).

For an integer $n > 0$, the symbol $\mathrm{lb}(n)$ denotes the base-2 logarithm of n, rounded down—equivalently, the largest integer k such that $2^k \leq n$.

If $\mathcal{S}$ and $\mathcal{T}$ are sets, then $\mathcal{S} \subset \mathcal{T}$ *allows for* the possibility that $\mathcal{S} = \mathcal{T}$. We'll write $\mathcal{S} \subsetneq \mathcal{T}$ to mean "$\mathcal{S} \subset \mathcal{T}$ and $\mathcal{S} \neq \mathcal{T}$."

If $\mathcal{S}$ is a set, then $|\mathcal{S}|$ denotes the cardinality of $\mathcal{S}$, and $\mathrm{Pow}(\mathcal{S})$ denotes the powerset of $\mathcal{S}$ (the set of all subsets of $\mathcal{S}$).

If f is a mapping and $\mathcal{S}$ a set, then

$$f[\mathcal{S}] = \{f(x) : x \in \mathcal{S}\} \quad \text{and} \quad f^{-1}[\mathcal{S}] = \{x \in \mathrm{dom}(f) : f(x) \in \mathcal{S}\}.$$

We'll occasionally encounter sets $\mathcal{S} \in \mathrm{dom}(f)$ that also have elements in the domain of f; this is the reason for using $f[\mathcal{S}]$ instead of $f(\mathcal{S})$.

If $\mathcal{A}$ and $\mathcal{B}$ are Abelian groups, then $\mathcal{A} \times \mathcal{B}$ denotes the group of ordered pairs (a, b), with $a \in \mathcal{A}$, $b \in \mathcal{B}$, and componentwise addition. For $n \in \mathbb{N}$, we then write

$$\mathcal{A}^n = \overbrace{\mathcal{A} \times \mathcal{A} \times \ldots \times \mathcal{A}}^{n \text{ times}}.$$

We'll also write $\mathcal{A}^\omega$ for the direct *sum* of countably many copies of $\mathcal{A}$, i.e., the group of countable sequences of elements of $\mathcal{A}$ with just finitely many nonzero terms. These topics are described in more detail in Section B.1.

Standard References

The bibliography on pages 493–504 lists more than 200 references. Three of these stand out as essential companions to this volume:

- *Winning Ways for Your Mathematical Plays*, by Berlekamp, Conway, and Guy [**BCG01**], is the defining text on combinatorial games. Originally published in 1982, it is now available in a revised second edition, split into four volumes. Volume 1 is required reading. The remaining volumes are also valuable and address many topics and examples that are omitted from this book.
- Conway's *On Numbers and Games* [**Con01**] introduced the partizan theory to the world in 1976. Some of the material, particularly on the theory of surreal numbers, is still not available anywhere else.
- *Lessons in Play: An Introduction to Combinatorial Game Theory*, by Albert, Nowakowski, and Wolfe [**ANW07**], is a recent (2007) undergraduate-level textbook on combinatorial games. It is recommended as a gentle introduction and includes many examples not found elsewhere.

References to these three books will be by name, using the abbreviations *Winning Ways*, ONAG, and *Lessons in Play*. Other references will appear in square brackets: for example [**Bou01**], cross-referenced to the bibliography.

A more comprehensive combinatorial games bibliography is maintained by Aviezri Fraenkel [**Fra09**] and is recommended for references not mentioned in this book. Another valuable reference is *Games of No Chance*, a series of four volumes of papers published by Cambridge University Press [**Now96, Now02, AN09, NW13**]. The papers are freely available (search for "Games of No Chance") and are generally of very high quality.

Appendix A of this book summarizes all of the conjectures and open problems that appear throughout the text. A much longer (and evergrowing) list of unsolved problems is maintained by Guy and Nowakowski [**GN13**].

With just a few exceptions, none of the results in this book are due to the author. Attributions for specific results are generally omitted from the body of the text, in order to avoid disrupting the exposition. The historical notes at the end of each section contain an overview of the development of each topic, and an essay on the history of combinatorial game theory can be found in Appendix C.

Other Resources

Combinatorial game theory is heavily computational. Raw calculation of game values is often essential: to solve specific problems; to test or refute

conjectures; and to assemble databases that can be mined for patterns. This was recognized early in the history of the subject, leading to a succession of increasingly sophisticated computer algebra systems designed by Flanigan, Allemang, Wolfe, and many others.

The most recent and comprehensive such system, *cgsuite*, is an indispensible companion to this textbook. The author strongly recommends[2] that all readers obtain a copy (it's freely available at `www.cgsuite.org`). Experimenting with *cgsuite* is one of the best ways to become fluent in the algebra of combinatorial games, and it's an invaluable tool for working through the exercises.

4. A Survey of the Landscape

Combinatorial game theory is a vast subject. Over the past forty years it has grown to encompass a wide range of games, extending far beyond the carefully selected examples from Section 1. All of those examples were **short games**, which have just finitely many subpositions and which prohibit infinite play. The combinatorial theory of short games is central to the subject and will motivate roughly half the material in this course, but its extensions to wider classes of games are equally important.

In this section we'll introduce several further examples that are *not* short games: FOX AND GEESE; GO; CHESS; and ENTREPRENEURIAL CHESS. We'll also survey various additional topics that fall outside the scope of the disjunctive theory: nondisjunctive compounds; positional games; multiplayer games; hardness results; and exhaustive and heuristic search techniques.

FOX AND GEESE

FOX AND GEESE is an old children's game played on an 8×8 checkerboard. In the starting position, shown in Figure 4.1 on the facing page, there are four **geese** (controlled by Left) arranged against a single **fox** (the white circle, controlled by Right). On his turn a player must move exactly one of the pieces under his control. The geese move like ordinary checker pieces, one space diagonally southwest or southeast. The fox moves like a checker king, one space in any of the four diagonal directions. If the geese eventually trap the fox, so that he is completely surrounded and unable to move, then Left wins. If this never happens, then Right wins.

---

[2]Disclosure: I'm the project administrator and primary architect of *cgsuite*.
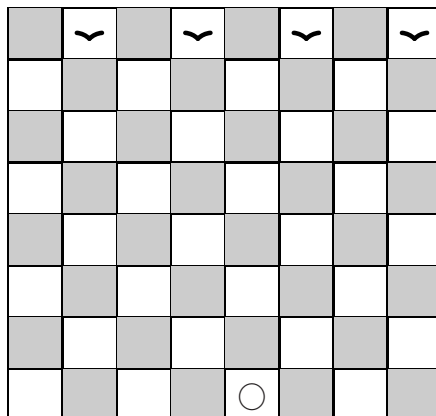
Figure 4.1. The starting position in FOX AND GEESE.

Since the geese can never retreat, the game is guaranteed to end after a finite number of moves. If the geese fail to trap the fox, then Left will eventually run out of moves, so the winning condition is really just the normal-play convention in disguise: whoever makes the last move wins. Therefore FOX AND GEESE has much in common with HACKENBUSH and DOMINEERING.

There is also a fundamental difference: FOX AND GEESE is *not* a short game, since if Right were permitted two consecutive moves, then he could return to the exact position from which he started. In the notation of Section 1, we have $G^{RR} = G$, where G is the position in Figure 4.1. So although every *play* of FOX AND GEESE is guaranteed to end (since Left can make just finitely many moves), there are infinite *runs* that involve many consecutive moves for Right.

This observation is highly relevant to the combinatorial structure of FOX AND GEESE. For consider a hypothetical sum

$$G + H,$$

in which H is a game to which *Left* can return in two consecutive moves: $H^{LL} = H$. (For example, we might take $H = -G$.) There is a four-move alternating sequence

$$G + H \ \rightarrow \ G^R + H \ \rightarrow \ G^R + H^L \ \rightarrow \ G^{RR} + H^L \ \rightarrow \ G^{RR} + H^{LL}$$

that returns to $G + H$, so in fact there exist infinite plays of $G + H$. So although *isolated* plays of G are always finite-length, its infinite runs nonetheless have combinatorial importance.

A detailed analysis of this type of disjunctive structure can sometimes yield useful insights, even for games that don't "naturally" decompose. For example, an isolated FOX AND GEESE position can never decompose in a

nontrivial way, since there is just one fox! However, the combinatorial theory can still provide a compelling window into questions like:

> Who is ahead, and by how much?

In Figure 4.1 the geese can force a win. It's not hard to work this out through a "brute force" analysis, without using any nontrivial mathematics. But it's much more difficult to quantify the precise extent of Left's advantage. What if the fox were permitted one pass move, to be used at any point during the game? Could the geese still win? How about two pass moves? Or $2\frac{1}{2}$ pass moves, in the sense of HACKENBUSH?

Combinatorial game theory provides the tools to answer these kinds of questions. It turns out that Left can still win if Right is given two pass moves, so that

$$G > 2.$$

But Right can win with just a slightly greater advantage: for all n,

$$G < 2 + \frac{1}{2^n}.$$

Moreover, if H is any *short* game with $2 < H < 2 + 2^{-n}$ for all n, then $G > H$. In particular, no short game has the same value as G. These types of quantitative comparisons are among the most powerful applications of combinatorial game theory, and they provide perhaps the clearest evidence that the cycles in G are an intrinsic part of its structure.

Games that admit infinite runs, such as FOX AND GEESE, are called **loopy**; those that don't (including all short games) are **loopfree**.

Go

Go is among the oldest and most widely played board games in the world. It originated in China as early as 2000 BCE, and today there are more than two thousand professional players throughout the world, and tens of millions of active amateurs.

Go has also been a motivating influence on combinatorial game theory throughout the history of the subject. The strong modularity of late-stage endgames inspired Conway to discover the theory of partizan games, which led in turn to the surreal numbers. In recent times Go continues to be a source of inspiration, spurring new developments in temperature theory and the study of capturing races.

Unfortunately Go serves as a poor introductory example, for several reasons. Most of our examples have simple, well-defined rules that are easily translated into an abstract mathematical description. By contrast, Go is highly complex, with all the messy peculiarities one might expect from a
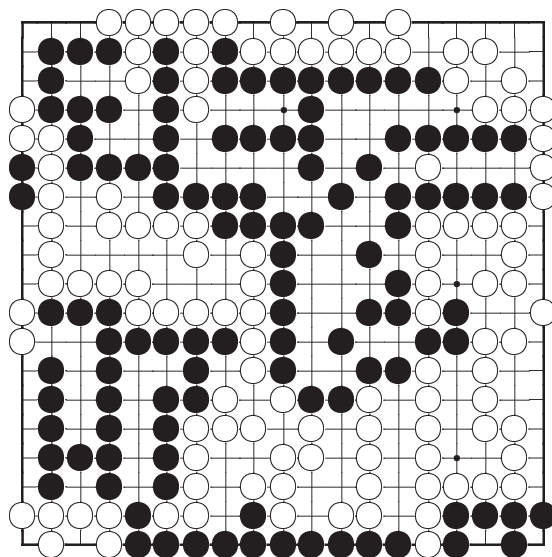
Figure 4.2. The "9-dan stumping problem" by Berlekamp and Wolfe.
White to move and win.

game that's evolved over thousands of years. A full account of its rules would
require several pages of discussion and would still fail to convey sufficient un-
derstanding to penetrate its mathematical structure even at a rudimentary
level.

For these reasons, we'll carefully avoid saying too much about Go in this
book, and the following discussion will largely be of interest to Go players.
Readers who are unfamiliar with the game can safely skip this section. Go
players, read on!

Figure 4.2 shows the famous "9-dan stumping problem" composed by
Berlekamp and Wolfe in 1994. The problem is carefully designed so that
the board decomposes into many indepdendent regions, such as the five-
point corridor in the northwest. Each such component is surrounded by
living groups of stones, in such a way that play in one component cannot
affect any other. (Some regions, such as the south-center of the board, are
composed of several weakly interacting subcomponents. In these cases the
entire region can be regarded as a single large component and can be further
broken down by quantifying these weak interactions.)

The entire problem can therefore be written as the disjunctive sum of
its various regional components. Each component has a game value (in the
sense of Section 1), and the combinatorial theory of partizan games can be
applied to find the value, and therefore the outcome, of the overall position.
We'll develop the theoretical machinery for this in Chapter II; details of its

specific applications to Go are published in Berlekamp and Wolfe's book, *Mathematical Go: Chilling Gets the Last Point* [**BW94**].

The Berlekamp–Wolfe analysis is one of combinatorial game theory's great success stories. Several 9-dan Go professionals (the top professional rank) have tried quite hard to solve Figure 4.2, without success. With sufficient knowledge of combinatorial game theory, however, it becomes a reasonably straightforward problem.

The 9-dan stumping problem is highly contrived and has little in common with positions that arise in real games. In most real-world situations, the Berlekamp–Wolfe analysis is relevant only in the very late endgame, and even then nontrivial applications are relatively rare. When *Mathematical Go* was first published in 1994, this was true of most applications of combinatorial game theory to Go.

Gradually, however, this situation is changing. Berlekamp has developed a rich temperature theory that applies to early-stage endgames and is relevant to a broader range of positions than the combinatorial theory of *Mathematical Go*. Over the past fifteen years, Berlekamp's theory has been investigated and extended by many other researchers and has been used to discover mistakes in games between top professional players. We'll explore the temperature theory in Chapter VII in an abstract mathematical context; references for specific applications to Go are given in the notes on page 49.

CHESS

On the surface, combinatorial game theory appears ill-suited to CHESS, since CHESS positions rarely decompose into disjunctive sums. In the fall of 1992, Elwyn Berlekamp made comments to this effect as part of an expository lecture on combinatorial games. Fortunately, Noam Elkies, who was in the audience at the time, interpreted Berlekamp's comments as a challenge, and he soon discovered several positions that are amenable to a disjunctive analysis.

Figure 4.3 on the next page illustrates a rare example from tournament play. The kings are in mutual zugzwang: the first player to make a king move will lose the game. Since the kings are effectively immobilized, the kingside and queenside pawn clusters cannot interact, so the board decomposes. Moreover, whoever makes the last pawn move will win, since her opponent will be forced to move his king. So we can regard Figure 4.3 as a normal-play combinatorial game with two disjoint components, one for each pawn cluster.
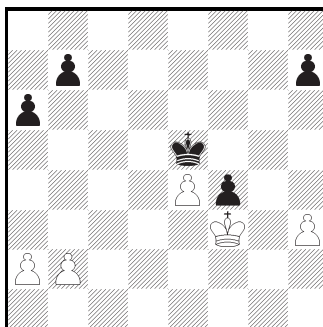
Figure 4.3. Schweda–Sika, Brno 1929.

Figure 4.3 is sufficiently simple that it can be successfully analyzed without any knowledge of combinatorial game theory. However, such a brute-force analysis provides little insight. Elkies quotes Max Euwe (a former world champion), who correctly determined that Figure 4.3 is a first-player win but had difficulty explaining why:

> Neither side appears to have any positional advantage in the normal sense .... The player with the move is able to arrange the pawn-moves to his own advantage [and win] in each case. It is difficult to say why this should be so, although the option of moving a pawn one or two squares at its first leap is a significant factor. [**EH60**]

In the language of combinatorial game theory, by contrast, one can say *exactly* why Figure 4.3 is a first-player win. We'll develop this language in Chapter II; skip ahead to Exercise II.4.9 on page 97 for the "correct" answer to Figure 4.3.

Elkies's two papers on the subject [**Elk96, Elk02**] are a nice example of how combinatorial game theory can be be applied—often in surprising contexts—to gain new insights into game situations that are otherwise baffling. The theory casts a wider net than a literal interpretation of its assumptions might suggest.

Entrepreneurial Chess

This fascinating game was invented by Elwyn Berlekamp. It is played on a quarter-infinite chessboard, with just a white king, white rook, and black king, as shown in Figure 4.4(a) on the next page. The pieces move just like their Chess counterparts, and Left (playing black) has the additional option of *cashing out*. If Left opts to cash out, then the entire position is replaced by an integer n, equal to the sum of the row and column numbers
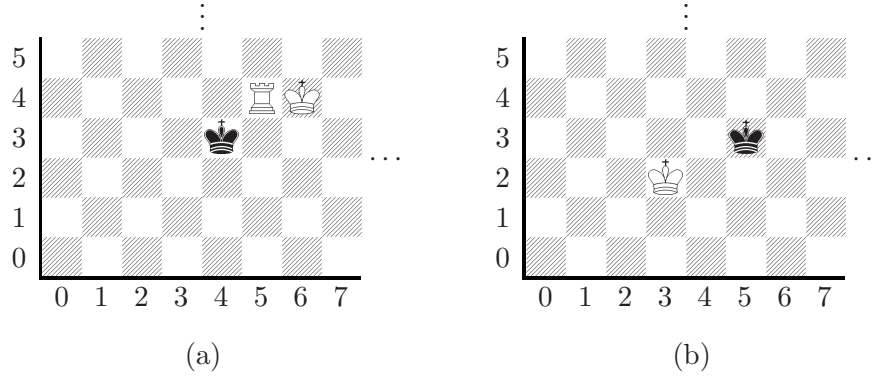
Figure 4.4. Entrepreneurial Chess. (a) A typical position; (b) a pathological position in which the rook has been captured.

for the current location of her king. For example, in Figure 4.4(a), Left may either move her king or replace the entire position by the number 7.

Entrepreneurial Chess is certainly loopy, like Fox and Geese and the other examples from this section. But there is a further difference from the classical examples from Section 1. A given Entrepreneurial Chess position might have infinitely many subpositions; we'll call games with this property **transfinite**. There exist finite games that are loopy (think Fox and Geese) and transfinite games that are loopfree (we'll see one in a moment); Entrepreneurial Chess is *neither* finite nor loopfree.

Nonetheless some positions *are* equivalent to finite values. Consider Figure 4.4(a). With best play, Left will never move her king, because there is no possibility of scoring more than 7 points (unless Right cooperates). Right can do no better than to shuttle his king around the squares adjacent to his rook. In fact it turns out that the value of Figure 4.4(a) is exactly equal to $5 + G$, where G is the Fox and Geese starting position from Figure 4.1. The machinery used to prove this (highly nonobvious!) fact is developed in Chapter VI.

The position K in Figure 4.4(b) is also quite interesting. In this pathological example, Right's rook has been captured, and Left is free to wander off indefinitely. It is clear that $K > n$ for any integer n: Left can systematically move her king to a square of value $n + 1$, then cash out; and Right is powerless to stop her advance.

But consider the difference of K with the *transfinite* Hackenbush stalk in Figure 4.5 on the facing page. The value of this Hackenbush position is the transfinite ordinal ω:

$$\omega = \{0, 1, 2, 3, \ldots \mid \ \}.$$
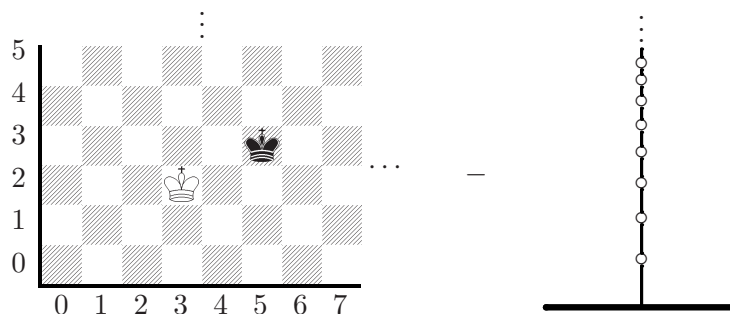
Figure 4.5. Entrepreneurial Chess meets ω.

On $K - \omega$, Right can either move his king or replace $-\omega$ by a suitable integer $-n$. It's clear that the latter is a losing move, since it reduces the position to $K - n$, which we know Left can win. But from $K - \omega$ Left can never afford to cash out, since then the position will be $n - \omega$, from which *Right* has a move to $n - (n+1)$! So on $K - \omega$ good players will simply shuttle their kings around aimlessly, and the outcome is a draw.

We conclude that $K - n$ is a win for Left for any finite n, but $K - \omega$ is drawn. It can be shown that no finite game has this property, so that the value of K is genuinely transfinite. Understanding precisely what *is* its value will require a lot of theory; peek ahead at Exercise VIII.1.12 on page 411 if you're curious about the answer.

Various Classes of Games

We've now seen several kinds of combinatorial games:

- *short impartial* games: Nim and Dawson's Kayles;
- *short partizan* games: Hackenbush and Domineering;
- *finite loopy* games: Fox and Geese, Go, and Chess;
- *transfinite* games: Entrepreneurial Chess and transfinite Hackenbush.

It turns out to be quite difficult to subsume *all* combinatorial games into a single theory. The most successful results come when we restrict ourselves to a particular subclass of games and consider the combinatorial interactions within that class.

Figure 4.6 on the following page summarizes several natural subclasses. There are three **structural constraints**, defined as follows:

Structural Constraints

Impartial	**Partizan**
Same moves for both players	Distinct moves for each player
Loopfree	**Loopy**
All runs are finite-length	Infinite runs are permitted
Finite	**Transfinite**
Finitely many subpositions	Infinitely many distinct subpositions

Play Convention

Normal Play	**Misère Play**
Last player to move wins	Last player to move loses

Figure 4.6. Various classes of combinatorial games.

Definition 4.1. Let G be a combinatorial game. G is said to be:

(a) **finite** if G has just finitely many distinct subpositions;

(b) **loopfree** if every run of G has finite length;

(c) **impartial** if Left and Right have exactly the same moves available from every subposition of G.

G is **transfinite** (resp. **loopy**, **partizan**) if it is not necessarily finite (resp. loopfree, impartial).

We say G is **short** if it is both finite and loopfree.

One can adopt any subset of these constraints (finite, loopfree, impartial) to obtain a restricted class, for a total of eight possibilities. The **play conventions**, also summarized in Figure 4.6, don't really define classes of games at all, but they are included in our table since they're a crucial ingredient in the resulting theory. Remarkably, in all cases the definition of equality is exactly the same:

$$G = H \quad \text{if} \quad o(G + X) = o(H + X) \text{ for all } X.$$

We can choose any of the eight structural classes of games for the domain of G, H, and X, and we can choose either normal or misère outcome for the meaning of $o(G)$. This gives sixteen different flavors of the fundamental equivalence, and each one yields a different theory!

In Section 1, for example, we defined $\mathbb{G}$ by choosing *short partizan* games for the domain and *normal-play* outcome for $o(G)$. $\mathbb{G}$ has a particularly rich structure that will occupy most of our attention in Chapters II through IV. Later chapters will investigate how much of this structure is retained when we vary each of these constraints (finite, loopfree, normal-play).

	Normal-Play Impartial	Normal-Play Partizan	Misère-Play Impartial	Misère-Play Partizan
Short	Sprague–Grundy IV.1–3	*Winning Ways* II, III, VII.1–2	V.1–5	V.6
Finite Loopy	Smith, Fraenkel IV.4	VI, VII.3–5		
Transfinite Loopfree	ONAG VIII.4	ONAG VIII.1–3	Uncharted Terrain	
Transfinite Loopy	Extends Finite Loopy			

Figure 4.7. A landscape of sixteen combinatorial game theories.

Figure 4.7 shows a rough roadmap of where we're headed. Short games in normal play, which are reasonably well understood, will occupy several chapters. By contrast, the theory of short partizan games in misère play was only recently isolated, and the small body of work since then has only barely scratched the surface. Finite loopy games, misère impartial games, and transfinite loopfree games fall somewhere between these extremes. You can see from Figure 4.7 that there are several classes of games (such as loopy games in misère play) about which essentially nothing is known at all!

In all these discussions, note that the *structural* classes on the right in Figure 4.6 (partizan, loopy, transfinite) are defined to include those on the left (impartial, loopfree, finite). For example, every impartial game is also *a priori* partizan. This makes the terminology far more convenient; we'll reserve the term **strictly partizan** for those rare cases when we wish to describe a game that specifically *isn't* impartial.

What's a Solution?

Much of this book is about finding *solutions* to games. Ordinarily, such a solution will take the form of a concise, closed-form algebraic description of a certain class of positions—Bouton's Theorem is the defining example. When we try to frame this sort of subjective understanding into a precise definition of *solution*, however, various subtleties arise. In Section 1 we were rather vague about our intentions:

> A **solution** to a ruleset Γ is an efficient algorithm for computing $o(G)$, for any position G of Γ.

This can be made precise using ideas from computational complexity. The rest of this discussion presupposes some background in complexity theory, which is not covered in this book; the requisite material can be found

in standard references such as [**LP97, Pap93**]. However since the focus of this book is on the algebraic and combinatorial structure of games, rather than their computational complexity, it is safe to skim or skip this discussion and to stick with the above as a working (though imprecise) definition of solution.

Here's a more formal definition of solution:

> A **solution** to a ruleset Γ is a polynomial-time algorithm for computing $o(G)$, for any position G of Γ.

Of course this revised definition is still not precise, since it leaves open the fundamental question: polynomial-time in terms of what? Typically, G is formalized as the tree (or graph) of its subpositions. But the size of this **game graph** isn't suitable to define input complexity, since the naïve brute-force algorithm (that examines each node once) is linear in the size of the graph. However there is usually some other natural choice—a "succinct encoding" for Γ that expresses its structure using the minimum necessary information. For example, in the case of DOMINEERING it's logical to use the number of squares on the board position; for HACKENBUSH the number of edges in the graph; and so on.

It's clear that the right notion of input complexity for Γ depends upon how we choose to represent its positions. This problem can be sidestepped by building a measure of input complexity directly into our definition of ruleset. Here, finally, is a precise definition of **ruleset** and **solution**:

Definition 4.2. A **ruleset** is a set of games Γ (the **positions** of the ruleset), together with a function $N : \Gamma \to \mathbb{N}$. The value $N(G)$ is called the **input complexity** of G (as an element of Γ).

A **solution** to $\Gamma = (\Gamma, N)$ is an algorithm for computing $o(G)$, for any $G \in \Gamma$, whose running time is polynomial in $N(G)$.

Fortunately it will rarely be necessary to specify N explicitly, since there are certain "natural" choices that are almost universally adopted:

- If Γ is played on a grid and G is an $m \times n$ position of Γ, then $N(G) = mn$.
- If Γ is played on a graph and the graph for G has v vertices and e edges, then $N(G) = v + e$.
- If Γ is played with heaps of tokens and G is the sum of heaps of sizes $a_1, a_2, \ldots, a_k$, then

$$N(G) = \mathrm{lb}(a_1) + \mathrm{lb}(a_2) + \cdots + \mathrm{lb}(a_k).$$

($N(G)$ is the number of bits of information necessary to represent the position.)

The definition of ruleset has useful implications for the algebraic theory of combinatorial games, since *for the most part* the main goals of the theory have now been precisely stated. The reason for qualifying this statement is that the consequences of Definition 4.2 are occasionally counterintuitive. Here are some of the common complications:

Positions of Bounded Size. Under Definition 4.2, any single (finite) position is *a priori* solved. For example, consider the starting position G in 8×8 CHESS. There is a trivial brute-force search algorithm that determines $o(G)$ by exhaustively examining every subposition. Since 8×8 CHESS has just finitely many possible positions, the running time of this algorithm is bounded by a constant and is therefore polynomial. The fact that this constant is enormously large, and hopelessly out of reach of present-day computational resources, is irrelevant. So although no one knows the outcome of 8×8 CHESS for certain, it's technically "solved" according to Definition 4.2.

For this reason, Definition 4.2 requires an infinite class of positions (with unbounded input complexity) in order to work. For abstract combinatorial games such as NIM and DAWSON'S KAYLES, this constraint is typically satisfied automatically. However, playable games such as CHESS are often assigned a fixed board size, usually with a specific starting position. Meaningful hardness results for such games depend on suitable generalizations to arbitrary board sizes.

In the case of CHESS, it's been shown that a suitable generalization of $n \times n$ CHESS is EXPTIME-complete and therefore admits no solution. This might strike CHESS players as a pedantic observation, but it actually says something deep about the specific case of 8×8 CHESS. It strongly suggests that there is intrinsic computational complexity embedded in the rules to CHESS that affects individual cases (such as 8×8) just as much as the asymptotic case. Searching for a clever algebraic solution to 8×8 CHESS is *probably* fruitless, because we know that such a solution *provably* cannot be extended to the $n \times n$ case.

Solutions to Certain Classes of Positions. FOX AND GEESE presents an interesting case. We can generalize 8×8 FOX AND GEESE to an $n \times 8$ board; the starting position G_n is an elongated version of Figure 4.1, still with four geese and one fox. Berlekamp proved that G_n has the exact value $1 + 2^{8-n}$ for all $n \geq 9$. The outcomes for $n \leq 8$ are easily computable by exhaustive search, so that the outcomes of all *starting* positions are known. Moreover, a complete strategy is known for enforcing a win by Left in polynomial time.

However, it remains unknown how to compute the outcome of an arbitrary position. In actual play of FOX AND GEESE, this is an academic observation: the "unresolved" positions can only arise if Left fails to play
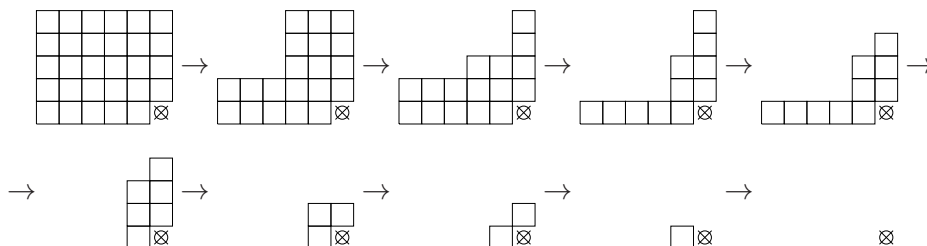
Figure 4.8. A typical game of CHOMP.

optimally or if play starts from a position other than G_n. Nonetheless, FOX AND GEESE doesn't meet our definition of "solved." This type of game is sometimes called **weakly solved**.

Another fascinating example is the impartial game CHOMP, played on an $m \times n$ grid as pictured in Figure 4.8. The players alternately remove ("chomp") a rectangle from the upper-left quadrant of the board. Denote by (i, j) the square at row i (top-to-bottom), column j (left-to-right); then a move is to select a pair (I, J) and discard all remaining squares (i, j) with $i \leq I$ and $j \leq J$. Each move must discard at least one square and *must leave at least one square remaining*. (This last rule is essential in order to prevent first player from simply removing the entire board.) The game ends when only one square remains, so that no further moves are possible, and whoever makes the last move wins.

Computing the outcome of an arbitrary CHOMP position appears to be quite difficult. However, the starting position G is easily shown to be a first-player win (except for the trivial case $m = n = 1$). The "strategy-stealing" argument goes like this: Suppose (for contradiction) that G is a second-player win. Suppose also that first player opens by removing the single square $(1, 1)$, moving to the position G'. Then second player must have a winning response, to some option G''. But G necessarily has a move directly to G'', so first player could have won by moving there initially. Contradiction!

A peculiar feature of this argument is that it gives absolutely no insight into the structure of CHOMP. We can say for certain that $m \times n$ CHOMP is a first-player win, but we can't even say what the winning moves are! Such games are sometimes called **ultra-weakly solved**. CHOMP, in particular, highlights the importance of including *all* positions in the definition of a full solution.

Enforcement of Strategies. NIMANIA is a curious game invented and studied by Fraenkel and Nešetřil [**FN85, FLN88**]. It is played with several heaps of tokens. On the k^{th} turn of the game, exactly one token is removed

from any one heap, say reducing it from size m to $m - 1$; then k *new* heaps are added to the position, each of size $m - 1$. (If $m = 1$, this last step has no effect.)

A typical game, starting with a single heap of size 3, might proceed as follows. Here a position is described as a multiset of heaps; the notation n^a abbreviates a copies of a heap of size n, and the subscript denotes the current value of k:

$$\{3\}_1 \to \{2^2\}_2 \to \{2, 1^3\}_3 \to \{2, 1^2\}_4 \to \{2, 1\}_5 \to \{2\}_6 \to \{1^7\}_7$$

after which the players alternately remove single-token heaps, and first player wins on the 13$^{\text{th}}$ move.

It's not hard to show by exhaustive analysis that starting from $\{3\}_1$, first player can force a win in 13 moves. In fact every starting position $\{n\}_1$ turns out to be a first-player win, and it's easy to describe the outcome of an arbitrary position G (with arbitrary subscript k):

- If G has the form $\{1^a\}_k$, then $o(G) = \mathcal{N}$ iff a is odd.
- If G has the form $\{2, 1^a\}_k$, then $o(G) = \mathcal{N}$ iff either a or k is odd.
- Otherwise, $o(G) = \mathcal{N}$ iff k is odd.

This gives a rather trivial strategy for NIMANIA: if G has the exact form $\{2, 1^a\}_k$, then play on a heap of size 1 if a and k are both odd, or 2 if either is even. Otherwise, play randomly!

This shows that NIMANIA is solved in the strongest sense of Definition 4.2. Amazingly, however—though first player has an easy win on $\{n\}_1$—second player can guarantee that the game lasts at least $A(n)$ moves, where A is an Ackermann function! Even $n = 4$ is effectively intractable: first player can force a win in no fewer than 2^{44} moves.

This raises the obvious question: what is the meaning of declaring a game to be a forced win, if an inordinately large number of moves is needed to consummate that win? Some authors have addressed this problem by placing additional constraints on the definition of solution—for example, by requiring that a win be enforceable in $O(c^{N(G)})$ moves, for some c. Fortunately such constraints are rarely necessary. In particular, in all of the examples we'll see in this book (except NIMANIA), each position G will have at most $O(c^{N(G)})$ subpositions.

Games Farther Afield

We'll conclude this survey with a discussion of some related topics in combinatorial game theory. Each of them is an active and important area of research; but in order to keep the focus on the disjunctive theory (and hold

this course to a manageable length) they won't be treated in detail in this book.

Nondisjunctive Compounds. Disjunctive decomposition arises naturally in games as diverse as DAWSON'S KAYLES, DOMINEERING, and GO. For this reason, disjunctive sum is a natural focal point in the study of abstract combinatorial properties of games.

Various other compounds also arise in practice, though less frequently. Aside from disjunctive sum, the **conjunctive** and **selective** compounds have historically received the most attention. In all three cases, the compound is formed by placing single copies of G and H side-by-side; but the rules for playing the compound differ:

> In the **disjunctive sum** $G + H$, move in *exactly one* of the two components.
>
> In the **conjunctive sum** $G \wedge H$, move in *both* components. (Play ends when *either* terminates.)
>
> In the **selective sum** $G \vee H$, move in *either or both* components.

These rules generalize in the obvious ways to games with three or more components; for example in

$$G_1 \wedge G_2 \wedge \cdots \wedge G_k$$

a player must make a move in *all k* of the components, and play ends as soon as *any* of them terminates.

Figure 4.9 gives precise symbolic definitions of conjunctive and selective sum, along with several additional types of compounds. All of these operations are associative, so they give rise to new variants of the fundamental equivalence and correspondingly new types of values. For example, for conjunctive sum we may write

$$G = H \quad \text{if} \quad o(G \wedge X) = o(H \wedge X) \text{ for all } X.$$

In each case, of course, $o(G)$ can mean either normal or misère outcome, and "for all" can range over any of the eight classes of games mentioned in Figure 4.6. So for each of these types of compounds we have another version of Figure 4.7, with another sixteen variants to consider!

There is evidently quite a lot of ground to cover here, and only a fraction of it has been properly explored. In this book we'll be concerned almost exclusively with the disjunctive theory; the notes on page 49 provide some references for exploring other types of compounds.

Disjunctive	Move in exactly one component.
	$$G + H \cong \left\{ G^L + H, G + H^L \mid G^R + H, G + H^R \right\}$$
Conjunctive	Move in all components. Play ends when any one of them terminates.
	$$G \wedge H \cong \left\{ G^L \wedge H^L \mid G^R \wedge H^R \right\}$$
Selective	Move in any number of components, but at least one.
	$$G \vee H \cong \left\{ G^L \vee H, G \vee H^L, G^L \vee H^L \mid G^R \vee H, G \vee H^R, G^R \vee H^R \right\}$$
Diminished Disjunctive	Move in exactly one component. Play ends immediately when any one of them terminates.
	$$G \boxplus H \cong \begin{cases} 0 & \text{if } G \cong 0 \text{ or } H \cong 0; \\ \left\{ G^L \boxplus H, G \boxplus H^L \mid G^R \boxplus H, G \boxplus H^R \right\} & \text{otherwise} \end{cases}$$
Continued Conjunctive	Move in all nonterminal components. Play ends only after all components terminate.
	$$G \triangledown H \cong \begin{cases} G + H & \text{if } G \cong 0 \text{ or } H \cong 0; \\ \left\{ G^L \triangledown H^L \mid G^R \triangledown H^R \right\} & \text{otherwise} \end{cases}$$
Shortened Selective	Move in any number of components. Play ends immediately when any one of them terminates.
	$$G \triangle H \cong \begin{cases} 0 & \text{if } G \cong 0 \text{ or } H \cong 0; \quad \text{otherwise,} \\ \left\{ G^L \triangle H, G \triangle H^L, G^L \triangle H^L \mid G^R \triangle H, G \triangle H^R, G^R \triangle H^R \right\} \end{cases}$$
Ordinal	Move in G or H; any move on G annihilates H.
	$$G : H \cong \left\{ G^L, G : H^L \mid G^R, G : H^R \right\}$$
Side	Move in G or H; Left's moves on H annihilate G, and Right's moves on G annihilate H.
	$$G \diamond H \cong \left\{ G^L \diamond H, H^L \mid G^R, G \diamond H^R \right\}$$
Sequential	Move in G unless G has terminated; in that case move in H.
	$$G \rightarrow H \cong \begin{cases} H & \text{if } G \cong 0; \\ \left\{ G^L \rightarrow H \mid G^R \rightarrow H \right\} & \text{otherwise} \end{cases}$$

Figure 4.9. Various types of compounds.

Positional combinatorial games. In a **positional game** the object is not to make the last move, but rather to fulfill some global positional objective. TIC-TAC-TOE and HEX are among the best-known examples.

The general case specifies a set X (the **board**) and a collection $\mathcal{F}$ of subsets of X (the **winning sets**). In TIC-TAC-TOE, for example, X is a nine-element set and $\mathcal{F}$ consists of all rows, columns, and diagonals of X. Each player alternately selects an element $x \in X$ and occupies it with her color; each x may only be selected once per game. The winning condition is usually defined in one of three ways:

- Under the **strong win** convention, the winner is the first player to occupy all the elements in one of the winning sets. If the board X is exhausted before either player succeeds, the game is a draw. This is how TIC-TAC-TOE is usually played.

- Under the **weak win** (or **Maker–Breaker**) convention, one player (the **Maker**) attempts to occupy a winning set, while the other (the **Breaker**) tries to prevent her. The game ends once the set X is exhausted. Maker wins just if she has occupied a winning set; otherwise Breaker wins. Draws are impossible. HEX and many other connection games fall into this category.

- Under the **reverse weak win** (or **Avoider–Enforcer**) convention, one player (the **Avoider**) strives to *avoid* occupying a winning set, while the other (the **Enforcer**) tries to ensure that she does. This is the misère version of the Maker–Breaker convention; misère HEX is a typical example.

The theory of positional games was introduced by Hales and Jewett [**HJ63**] and Erdős and Selfridge [**ES73**]. Over the past several decades it has been developed extensively, primarily through the sustained efforts of József Beck. The theory uses probabilistic methods to study the asymptotic structure of various abstract classes of games.

Surprisingly, at sufficiently *large* scales these methods can be used to deduce exact solutions! Beck cites the example of the (n, q)-CLIQUE GAME: the board X consists of all edges in the complete graph K_n on n vertices; $\mathcal{F}$ is the collection of all q-cliques of K_n. The players alternately select an edge from K_n; Maker attempts to complete a q-clique, and Breaker strives to prevent her. Remarkably, one can show that for $n = 2^{10^{10}}$, Maker has a winning strategy if and only if

$$q \le 19{,}999{,}999{,}933.$$

Similar results are obtainable for any sufficiently large n, but for smaller values of n the theory is less conclusive. For example, if n is "only" 2^{100}, then there are values of q for which the exact outcome remains unknown.

This fascinating asymptotic theory is obviously quite different from the algebraic and inductive methods that are the focus of this book, so we'll have little more to say about it. Fortunately Beck has recently published an encyclopedic and up-to-date survey of the subject, which is the natural starting point for further investigations [**Bec08**].

Multiplayer combinatorial games. Multiplayer games introduce the prospect of **coalitions**, in which several players collude to pursue a shared goal. The classical theory of economic games elegantly allows for multiplayer situations, but combinatorial settings usually require some degree of contrivance in order to exclude the possibility of collusion.

Let's consider the special case of impartial three-player games; generalizing to partizan N-player games is straightforward. Any of the impartial games discussed in this book can easily be adapted to a three-player setting. The players take turns moving on G in clockwise rotation; the game ends when the next player to move is unable to do so, in which case the entire prize goes to whoever made the last move.

The challenge is to say what is meant by $o(G)$. In some situations, one of the players can force a win no matter how the other two react, and in such cases the definition of $o(G)$ is noncontroversial. A trivial example is $G = n \cdot *$: every move is forced, and the outcome depends only on the parity of n modulo 3. But in most cases, none of the three players can force a win, because the outcome depends on how each of the other players reacts.

For example, on $*2 + *$ it is impossible for first player to win; but she can decide which of the other players wins, by moving variously to $* + *$ or $0 + *$ (say). Propp [**Pro00**] calls such positions $\mathscr{Q}$-positions and shows that the situation is quite dire: $G + *2 + *2$ is a $\mathscr{Q}$-position, for *any* game G!

Various authors have attempted to resolve this problem by making assumptions about the rationality of each player's behavior. For example:

- Li [**Li78**] supposed that if player X is unable to force a win, then she will "play for second," ensuring (if possible) that she will be the *second-to-last* player to move.

- Straffin [**Str85**] introduced a "revenge rule": if player X is faced with a kingmaker decision and chooses Y over Z, then we assume Z will favor Y over X the next time he is faced with a similar decision.

- Loeb [**Loe96**] considered probabilistic models for **stable coalitions**: minimal sets of players that can collectively guarantee a win for some player in the set.

The Li and Straffin approaches lead to elegant combinatorial theories, but they suffer from a serious defect: in practice, the resulting outcomes have little meaning, since they rely on the assumption that both *other* players will play perfectly. Loeb's approach is more resilient, but notably less "combinatorial" in nature.

Cincotti [**Cin10**] has proposed a different approach: simply leave the outcome of Propp's $\mathscr{Q}$-positions undefined, and try to identify situations under which one of the players has a true forced win. Focusing on the partizan theory, he defined multiplayer generalizations of the *recursive* characterization of $G \geq H$ and showed that certain combinations of such relations can sometimes identify the outcome.

Algorithmic combinatorial game theory.   Definition 4.2 (page 36) enables one to say precisely what is meant by a "solution" to a ruleset. Consequently, one can also say that certain particular rulesets are formally **unsolvable** (or **intractable**). For example, $n \times n$ CHESS is known to be EXPTIME-hard, so that (by standard results in complexity theory) no polynomial-time solution can exist. Many similar results are known; see Figure 4.10 on the next page for a sampling.

Algorithmic combinatorial game theory studies the computational complexity of combinatorial games. The usual technique (applied to a particular ruleset Γ) is to identify a class of Γ-positions with a special structure that can be reduced to a known, computationally difficult problem. In the case of HACKENBUSH, for example, certain positions known as **redwood beds** are equivalent to bipartite graphs, and calculating the value of a redwood bed is equivalent to finding a minimum spanning tree for its associated graph—a problem that is known to be NP-complete.

The algorithmic approach naturally complements the algebraic theory. Both approaches seek to find structure in combinatorial games: in one case, to show that this structure is simple enough to admit a complete theory; and in the other, to show that it is varied enough to preclude such possibility. The edge cases emerge as the most troubling: there are games that we "expect" to admit solutions (misère DAWSON'S KAYLES) but that have a relatively deep structure; and games that we "expect" to be intractable (DOMINEERING) but have relatively little structure. Algorithmic combinatorial game theory provides a unifying framework for attacking these problems from both directions.

A novel technique known as **constraint logic** defines a common set of problems (**constraint games**) of known complexity, which are constructed in a way that simplifies reductions from combinatorial games. The constraint logic approach has led to several new hardness results (including KONANE) and simpler proofs of existing ones (AMAZONS).

AMAZONS	PSPACE-complete	Furtak et al. [**FKUB05**]; Hearn [**Hea09**]
$n \times n$ CHESS	EXPTIME-complete	Fraenkel–Lichtenstein [**FL81**]
CHOMP	?	
CLOBBER	?	
DAWSON'S KAYLES	Solved	Guy–Smith [**GS56**]
DAWSON'S KAYLES (misère)	?	
DOMINEERING	?	
DOTS AND BOXES	NP-hard	*Winning Ways* [**BCG01**]
FLOWERS	?	
$n \times 8$ FOX AND GEESE	?	
$n \times n$ GO	EXPTIME-complete	Lichtenstein–Sipser [**LS80**]; Robson [**Rob83**]
GRUNDY'S GAME	?	
HACKENBUSH	NP-hard	*Winning Ways* [**BCG01**]
KAYLES	Solved	Guy–Smith [**GS56**]
KAYLES (misère)	Solved	Conway–Sibert [**CS92**]
KONANE	PSPACE-complete	Hearn [**Hea09**]
NIM	Solved	Bouton [**Bou01**]
PHUTBALL	PSPACE-hard	Dereniowski [**Der10**]

Figure 4.10. Complexity status of various rulesets.

Surveys by Fraenkel [**Fra96b**] and Demaine and Hearn [**DH09**] provide excellent summaries of this material. The constraint logic approach is described in detail in a new book by Hearn and Demaine [**HD09**].

Exhaustive search. Some combinatorial game positions G are sufficiently simple that they are amenable to **exhaustive search**: an attempt to determine $o(G)$ by brute force, either by exhaustively visiting every subposition of G or by visiting enough of them that the value of $o(G)$ can be definitely isolated. (For example, to determine that $o(G) = \mathcal{N}$, one need only exhibit a single winning opening move for each player; once this is done, the proof is complete and the remaining opening moves needn't be considered.)

This is *in theory* possible for every finite G, by brute-force application of the Fundamental Theorem, but *in practice* its feasibility is limited by the available computing resources. The good news is that the rapid expansion in computing power over the past several decades has dramatically increased the scope of exhaustive search techniques.

The well-known game of CHECKERS is perhaps the most dramatic illustration. Long the subject of intense interest, CHECKERS was quite recently shown to be a first-player draw: starting from the standard opening position on an 8×8 board, either player can force at least a draw with best play. This finding was announced in 2007 by a team led by Jonathan Schaeffer [**Sch96, SBB$^+$07**], after a monumental 18-year search that examined more than 10^{14} subpositions.

Figure 4.11 lists several more classic games that have succumbed to exhaustive search. Some entries are partial results. For example, DOMINEERING has an indeterminate starting boardsize, and it turns out that the

CHECKERS	Draw (on 8×8 board) Schaeffer et al. [**SBB$^+$07**]
CHESS	All pawnless six-piece endgames are solved Stiller [**Sti96**]
DOMINEERING	First-player win on $n \times n$ for $6 \leq n \leq 10$ Breuker et. al. [**BUvdH00**]; Bullock [**Bul02**]
NINE MEN'S MORRIS	Draw (on traditional board) Gasser [**Gas96**]
TIGERS AND GOATS	Draw (on traditional board) Jin and Nievergelt [**JN09**]

Figure 4.11. A sampling of games that have fallen to exhaustive search.

computational challenge is greatest the closer the board is to being square. The largest square board that has been completely solved is 10×10.

In the case of CHESS, the full solution is not remotely within the realm of what is feasible today, but exhaustive search has proven to be enormously successful in computing the exact outcomes of various endgame positions (even resulting in several official adjustments to the "50-move rule").

Exhaustive search is an undoubtedly combinatorial pursuit, but it is tangential to the main ideas of this book, which are generally directed toward solutions for infinite sequences of positions based on an understanding of the algebraic structure of games. However, it remains an active area of research, and the interested reader is enthusiastically directed to the references in Figure 4.11.

Heuristic search. In recent years, exhaustive search has led to complete solutions for many games (such as CHECKERS) that had previously been considered infeasible. However others, such as CHESS and GO, remain far (and perhaps permanently) out of reach. Therefore, computer programs that *play* such games cannot resort to exhaustive search to find the best moves, and they typically incorporate a heuristic method for approximating the value of a given position G. Designing a good heuristic is a problem of an entirely different sort from the ones considered in this book: it relies on practical optimization and high-level understanding of the structure of specific games.

Positional heuristics are often coupled with brute-force search: for example, a program might search the game tree of G as deeply as possible in the allotted time, using its heuristic to evaluate the deepest nodes when time runs out. Likewise, an algorithm might rely on heuristics to prioritize the order of its search, analyzing the heuristically best options first.

Combinations of these sorts of techniques have proven to be enormously successful in applications to certain games, including CHESS, CHECKERS, and REVERSI. The CHESS-playing program Deep Blue, which defeated world champion Garry Kasparov in a much-publicized match in 1997, relied on highly sophisticated adaptations of such techniques.

By contrast, GO has proven to be remarkably resistant to heuristic search: although the best CHESS programs now routinely defeat reigning world champions, top GO software performs no better than a talented club player. This is due largely to GO's large branching factor and the importance of larger-scale positional strategy over local tactical calculations.

In recent years, a new suite of techniques known as *Monte Carlo* methods has revolutionized computer GO. The basic Monte Carlo algorithm

evaluates a position G as follows. For each potential option G^L, the computer plays out several thousand random games all the way to the end and assigns a score to G^L based on the *probability* of a favorable outcome. The computer then selects the option that achieved the most probable win. This basic approach can be enhanced in various ways, for example, by starting random games from a deeper level of the search tree, or by introducing a heuristic bias into the random plays.

Monte Carlo algorithms are far superior to classical search algorithms at the strategic level. However they are vulnerable to tactical weakness: if the opponent has just one good reply to a particular G^L, for example, there is some risk that the randomizer will simply overlook it. Since 2006, a hybrid approach known as UCT search (Upper Confidence bounds applied to Trees) has found spectacular success and cemented the dominance of Monte Carlo methods in computer GO. In the 2011 Computer Olympiad, the top three GO programs all used a version of Monte Carlo tree search; such programs are now able to beat professional players at reasonable handicaps.

Exercises

4.1   The outcomes for NIMANIA given on page 39 are correct.

4.2   In DIVISORS a position is a finite set $\mathcal{N} \subset \mathbb{N}^+$ that is divisor-closed (if a is an element of $\mathcal{N}$, then so are all divisors of a). A move consists of selecting an integer $a \in \mathcal{N}$ and removing a from the set, together with all multiples of a. Whoever removes 1, necessarily leaving the set empty, loses.

Prove that CHOMP is isomorphic to DIVISORS played on integers of the form $2^a 3^b$.

4.3   (a) All of the operations in Figure 4.9 are associative.
       (b) Which ones are commutative? For which of them is 0 an identity?

4.4   *Selective sums.* If G and H are short impartial games, then $o(G \vee H) = \mathscr{P}$ if and only if $o(G) = \mathscr{P}$ and $o(H) = \mathscr{P}$ (normal play). Conclude that there are just two normal-play short impartial values for selective sums. How about misère play?

4.5   *Conjunctive sums.* The **remoteness** $\mathscr{R}(G)$ of a short impartial game G is defined by

$$\mathscr{R}(G) = \begin{cases} 0 & \text{if } G \cong 0; \text{ otherwise:} \\ 1 + \min\left\{\mathscr{R}(G') : \mathscr{R}(G') \text{ is even}\right\} & \text{if some } \mathscr{R}(G') \text{ is even;} \\ 1 + \max\{\mathscr{R}(G') : \mathscr{R}(G') \text{ is odd }\} & \text{if every } \mathscr{R}(G') \text{ is odd.} \end{cases}$$

       (a) $o(G) = \mathscr{P}$ if and only if $\mathscr{R}(G)$ is even.
       (b) $\mathscr{R}(G \wedge H) = \min\{\mathscr{R}(G), \mathscr{R}(H)\}$ for all G and H.
       (c) Define the **misère remoteness** $\mathscr{R}^-(G)$ of G by interchanging "odd" and "even" in the definition for $\mathscr{R}(G)$, and show that it works for determining the misère outcome of conjunctive sums.

4.6 Research the references for conjunctive and selective sums in the notes section (below), and draw the analogous tables to Figure 4.7.

4.7 TIC-TAC-TOE is a win for Maker when played under the Maker–Breaker convention.

Notes

FOX AND GEESE features prominently in *Winning Ways*. The first edition incorrectly asserted that Figure 4.1 has value $1 + \mathbf{over}$; the correct answer $2 + \mathbf{over}$ was later given by Jonathan Welton and verified by Berlekamp and Siegel. Berlekamp also proved that for $n \geq 9$, the $n \times 8$ starting position has the exact value $1 + 2^{-(n-8)}$; this analysis was incorporated into the second edition of *Winning Ways* in 2003.

The combinatorial theory of GO endgames was introduced in Wolfe's thesis [**Wol91**], building on Berlekamp's earlier discoveries in the theory of DOMINEERING [**Ber88**]. Further applications have been explored by Kim [**Kim95, BK96**], Landman [**Lan96**], Nakamura [**BN03**], Takizawa [**Tak02**], and others.

The temperature theory of GO was introduced in a seminal article by Berlekamp [**Ber96**]; many examples were subsequently found by Berlekamp, Müller, and Spight [**BMS96**]. The theory has been further pursued by Fraser [**Fras02**], Kao [**Kao97**], and Spight [**Spi99, Spi02, Spi03**].

Nakamura [**Nak09**] has recently introduced a beautiful new theory of GO capturing races, in which liberty counts are represented by partizan game values. Nakamura's theory bears a striking resemblance to the atomic weight theory (described in Section II.7 of this book).

ENTREPRENEURIAL CHESS was invented by Berlekamp and Pearson. It was inspired by the following CHESS problem, originally due to Simon Norton and mentioned in Guy's 1991 list of unsolved problems [**Guy91c**]: from the starting position in Figure 4.12, "what is the smallest board that White can win on if Black is given a win if he walks off the North or East edges of the board?"

Berlekamp and Pearson carried out a detailed temperature analysis of ENTREPRENEURIAL CHESS, showing (for example) that Figure 4.12 has mean value 17, among many other results [**BP03**]. Pearson also used related methods to solve Norton's original problem: the answer is 8×11.
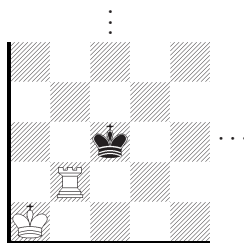

Figure 4.12. Simon Norton's CHESS problem.

Chomp is isomorphic to the game Divisors, originally introduced by Frederik Schuh [**Sch52**] (see Exercise 4.2). In the form presented here, it was invented by David Gale [**Gal74**]. Transfinite generalizations, in which the board dimensions are permitted to be transfinite ordinals, have been studied by Huddleston and Shurman [**HS02**].

These examples were chosen to illustrate the breadth and variety of combinatorial games, but in truth they barely scratch the surface. Various other games are introduced and explored throughout this text, such as Amazons (Section II.5) and Wythoff (Section IV.3), but many more were left out. Here are some prominent examples in the latter category:

- Dots and Boxes is a well-known children's game with a surprisingly rich theory. Berlekamp's 2000 treatise [**Ber00a**] is the "bible" of Dots and Boxes strategy; a later advance is described by Berlekamp and Scott [**BS02**].

- The traditional Hawaiian board game Konane meets the definition of *short combinatorial game* without any sort of contrivance. An early analysis by Ernst [**Ern95**] includes a nice account of its history; Chan and Tsai [**CT02**] later found solutions to certain sequences of $1 \times n$ positions. More recently, Santos and Silva announced the striking result that Konane contains positions of value $*m$ for every $m \geq 0$, a rare feature for a partizan game [**SS08**].

- Toads and Frogs was introduced in *Winning Ways* and later studied in depth by Erickson [**Eri96**]. More recently, these efforts have been carried much further by Thanatipanonda and Zeilberger, with the help of some exciting new techniques in automated theorem proving [**TZ09, Tha11, Tha**].

- Phutball ("Philosopher's Football") is an unusual board game invented by Conway and mentioned in *Winning Ways*. Demaine, Demaine, and Eppstein [**DDE02**] showed that it is NP-complete to determine whether a given player has an immediate winning move! More recently, Nowakowski, Ottaway, and Siegel [**Sie09b**] found that complicated loopy positions arise even on very small boards (such as 1×8).

Nondisjunctive compounds. Conjunctive and selective sums of short impartial games were studied by Smith [**Smi66**]. Selective sum has a trivial solution (Exercise 4.4). For the conjunctive case, Smith showed that the outcome of a sum can be computed from the **remoteness** of its components (Exercise 4.5), which is essentially a measure of how long the game will last with perfect play. Unlike the disjunctive theory, these results generalize easily to misère play.

Conway invented the diminished disjunctive, continued conjunctive, and shortened selective sums, and he showed that they can be analyzed using straightforward extensions of Smith's theories [**Con01**]. Remarkably, diminished disjunctive sum is also straightforward in misère play—so that ordinary disjunctive sum in misère play is the only really difficult theory of the bunch!

Winning Ways introduced a partizan theory for conjunctive and selective sums. In a fascinating twist, conjunctive sums can be analyzed using a trivial extension of the Steinhaus remoteness, whereas selective sums—which are trivial in the impartial case—require a deep and elegant theory all their own.

The theory of conjunctive and selective compounds can be extended to loopy games. The impartial cases were solved by Smith (they're really no harder than the loopfree analogues); the partizan cases were studied by Flanigan in the 1970s [**Fla79, Fla81, Fla83**].

Stromquist and Ullman found an elegant solution for sequential compounds of short impartial games [**SU93**]; this is explored in Exercise IV.1.6 on page 182. More recently, Fraser Stewart has investigated the partizan case [**Ste07**].