Surveys

Playing games with algorithms:
Algorithmic Combinatorial Game Theory

ERIK D. DEMAINE AND ROBERT A. HEARN

ABSTRACT. Combinatorial games lead to several interesting, clean problems
in algorithms and complexity theory, many of which remain open. The pur-
pose of this paper is to provide an overview of the area to encourage further
research. In particular, we begin with general background in Combinatorial
Game Theory, which analyzes ideal play in perfect-information games, and
Constraint Logic, which provides a framework for showing hardness. Then we
survey results about the complexity of determining ideal play in these games,
and the related problems of solving puzzles, in terms of both polynomial-time
algorithms and computational intractability results. Our review of background
and survey of algorithmic results are by no means complete, but should serve
as a useful primer.

1. Introduction

Many classic games are known to be computationally intractable (assuming
$P \neq NP$): one-player puzzles are often NP-complete (for instance Minesweeper)
or PSPACE-complete (Rush Hour), and two-player games are often PSPACE-
complete (Othello) or EXPTIME-complete (Checkers, Chess, and Go). Surpris-
ingly, many seemingly simple puzzles and games are also hard. Other results
are positive, proving that some games can be played optimally in polynomial
time. In some cases, particularly with one-player puzzles, the computationally
tractable games are still interesting for humans to play.

We begin by reviewing some basics of Combinatorial Game Theory in Sec-
tion 2, which gives tools for designing algorithms, followed by reviewing the

relatively new theory of Constraint Logic in Section 3, which gives tools for proving hardness. In the bulk of this paper, Sections 4–6 survey many of the algorithmic and hardness results for combinatorial games and puzzles. Section 7 concludes with a small sample of difficult open problems in algorithmic Combinatorial Game Theory.

Combinatorial Game Theory is to be distinguished from other forms of game theory arising in the context of economics. Economic game theory has many applications in computer science as well, for example, in the context of auctions [dVV03] and analyzing behavior on the Internet [Pap01].

2. Combinatorial Game Theory

A *combinatorial game* typically involves two players, often called *Left* and *Right*, alternating play in well-defined *moves*. However, in the interesting case of a *combinatorial puzzle*, there is only one player, and for *cellular automata* such as Conway's Game of Life, there are no players. In all cases, no randomness or hidden information is permitted: all players know all information about gameplay (*perfect information*). The problem is thus purely strategic: how to best play the game against an ideal opponent.

It is useful to distinguish several types of two-player perfect-information games [BCG04, pp. 14–15]. A common assumption is that the game terminates after a finite number of moves (the game is *finite* or *short*), and the result is a unique winner. Of course, there are exceptions: some games (such as Life and Chess) can be *drawn* out forever, and some games (such as tic-tac-toe and Chess) define *ties* in certain cases. However, in the combinatorial-game setting, it is useful to define the *winner* as the last player who is able to move; this is called *normal play*. If, on the other hand, the winner is the first player who cannot move, this is called *misère play*. (We will normally assume normal play.) A game is *loopy* if it is possible to return to previously seen positions (as in Chess, for example). Finally, a game is called *impartial* if the two players (Left and Right) are treated identically, that is, each player has the same moves available from the same game position; otherwise the game is called *partizan*.

A particular two-player perfect-information game without ties or draws can have one of four *outcomes* as the result of ideal play: player Left wins, player Right wins, the first player to move wins (whether it is Left or Right), or the second player to move wins. One goal in analyzing two-player games is to determine the outcome as one of these four categories, and to find a strategy for the winning player to win. Another goal is to compute a deeper structure to games described in the remainder of this section, called the *value* of the game.

A beautiful mathematical theory has been developed for analyzing two-player combinatorial games. A new introductory book on the topic is *Lessons in Play*

by Albert, Nowakowski, and Wolfe [ANW07]; the most comprehensive reference is the book *Winning Ways* by Berlekamp, Conway, and Guy [BCG04]; and a more mathematical presentation is the book *On Numbers and Games* by Conway [Con01]. See also [Con77; Fra96] for overviews and [Fra07] for a bibliography. The basic idea behind the theory is simple: a two-player game can be described by a rooted tree, where each node has zero or more *left* branches corresponding to options for player Left to move and zero or more *right* branches corresponding to options for player Right to move; leaves correspond to finished games, with the winner determined by either normal or misère play. The interesting parts of Combinatorial Game Theory are the several methods for manipulating and analyzing such games/trees. We give a brief summary of some of these methods in this section.

2.1. Conway's surreal numbers. A richly structured special class of two-player games are John H. Conway's *surreal numbers*[1] [Con01; Knu74; Gon86; All87], a vast generalization of the real and ordinal number systems. Basically, a surreal number $\{L \mid R\}$ is the "simplest" number larger than all Left options (in L) and smaller than all Right options (in R); for this to constitute a number, all Left and Right options must be numbers, defining a total order, and each Left option must be less than each Right option. See [Con01] for more formal definitions.

For example, the simplest number without any larger-than or smaller-than constraints, denoted $\{\mid\}$, is 0; the simplest number larger than 0 and without smaller-than constraints, denoted $\{0 \mid\}$, is 1; and the simplest number larger than 0 and 1 (or just 1), denoted $\{0, 1 \mid\}$, is 2. This method can be used to generate all natural numbers and indeed all ordinals. On the other hand, the simplest number less than 0, denoted $\{\mid 0\}$, is -1; similarly, all negative integers can be generated. Another example is the simplest number larger than 0 and smaller than 1, denoted $\{0 \mid 1\}$, which is $\frac{1}{2}$; similarly, all dyadic rationals can be generated. After a countably infinite number of such construction steps, all real numbers can be generated; after many more steps, the surreals are all numbers that can be generated in this way.

Surreal numbers form an ordered field, so in particular they support the operations of addition, subtraction, multiplication, division, roots, powers, and even integration in many situations. (For those familiar with ordinals, contrast with surreals which define $\omega - 1$, $1/\omega$, $\sqrt{\omega}$, etc.) As such, surreal numbers are useful in their own right for cleaner forms of analysis; see, e.g., [All87].

What is interesting about the surreals from the perspective of combinatorial game theory is that they are a subclass of all two-player perfect-information

---

[1]The name "surreal numbers" is actually due to Knuth [Knu74]; see [Con01].

Let $x = x^L | x^R$ be a game.

- $x \leq y$ precisely if every $x^L < y$ and every $y^R > x$.
- $x = y$ precisely if $x \leq y$ and $x \geq y$; otherwise $x \neq y$.
- $x < y$ precisely if $x \leq y$ and $x \neq y$, or equivalently, $x \leq y$ and $x \not\geq y$.
- $-x = -x^R | -x^L$.
- $x + y = x^L + y, x + y^L | x^R + y, x + y^R$.
- x is *impartial* precisely if x^L and x^R are identical sets and recursively every position ($\in x^L = x^R$) is impartial.
- A one-pile Nim game is defined by

$$*n = *0, \ldots, *(n-1) | *0, \ldots, *(n-1),$$

together with $*0 = 0$.

Table 1. Formal definitions of some algebra on two-player perfect-information games. In particular, all of these notions apply to surreal numbers.

games, and some of the surreal structure, such as addition and subtraction, carries over to general games. Furthermore, while games are not totally ordered, they can still be compared to some surreal numbers and, amazingly, how a game compares to the surreal number 0 determines exactly the outcome of the game. This connection is detailed in the next few paragraphs.

First we define some algebraic structure of games that carries over from surreal numbers; see Table 1 for formal definitions. Two-player combinatorial games, or trees, can simply be represented as $\{L \mid R\}$ where, in contrast to surreal numbers, no constraints are placed on L and R. The *negation* of a game is the result of reversing the roles of the players Left and Right throughout the game. The (*disjunctive*) *sum* of two (sub)games is the game in which, at each player's turn, the player has a binary choice of which subgame to play, and makes a move in precisely that subgame. A partial order is defined on games recursively: a game x is *less than or equal to* a game y if every Left option of x is less than y and every Right option of y is more than x. (Numeric) equality is defined by being both less than or equal to and more than or equal to. Strictly inequalities, as used in the definition of less than or equal to, are defined in the obvious manner.

Note that while $\{-1 \mid 1\} = 0 = \{\mid\}$ in terms of numbers, $\{-1 \mid 1\}$ and $\{\mid\}$ denote different games (lasting 1 move and 0 moves, respectively), and in this sense are *equal* in *value* but not *identical* symbolically or game-theoretically. Nonetheless, the games $\{-1 \mid 1\}$ and $\{\mid\}$ have the same outcome: the second player to move wins.

Amazingly, this holds in general: two equal numbers represent games with equal outcome (under ideal play). In particular, all games equal to 0 have the outcome that the second player to move wins. Furthermore, all games equal to a positive number have the outcome that the Left player wins; more generally, all positive games (games larger than 0) have this outcome. Symmetrically, all negative games have the outcome that the Right player wins (this follows automatically by the negation operation). Examples of zero, positive, and negative games are the surreal numbers themselves; an additional example is described below.

There is one outcome not captured by the characterization into zero, positive, and negative games: the first player to move wins. To find such a game we must obviously look beyond the surreal numbers. Furthermore, we must look for games G that are incomparable with zero (none of $G = 0$, $G < 0$, or $G > 0$ hold); such games are called *fuzzy* with 0, denoted $G \parallel 0$.

An example of a game that is not a surreal number is $\{1 \mid 0\}$; there fails to be a number strictly between 1 and 0 because $1 \geq 0$. Nonetheless, $\{1 \mid 0\}$ is a game: Left has a single move leading to game 1, from which Right cannot move, and Right has a single move leading to game 0, from which Left cannot move. Thus, in either case, the first player to move wins. The claim above implies that $\{1 \mid 0\} \parallel 0$. Indeed, $\{1 \mid 0\} \parallel x$ for all surreal numbers x, $0 \leq x \leq 1$. In contrast, $x < \{1 \mid 0\}$ for all $x < 0$ and $\{1 \mid 0\} < x$ for all $1 < x$. In general it holds that a game is fuzzy with some surreal numbers in an interval $[-n, n]$ but comparable with all surreals outside that interval. Another example of a game that is not a number is $\{2 \mid 1\}$, which is positive (> 0), and hence Right wins, but fuzzy with numbers in the range $[1, 2]$.

For brevity we omit many other useful notions in Combinatorial Game Theory, such as additional definitions of summation, superinfinitesimal games $*$ and $\uparrow$, mass, temperature, thermographs, the simplest form of a game, remoteness, and suspense; see [BCG04; Con01].

2.2. Sprague–Grundy theory. A celebrated early result in Combinatorial Game Theory is the characterization of impartial two-player perfect-information games, discovered independently in the 1930's by Sprague [Spr36] and Grundy [Gru39]. Recall that a game is *impartial* if it does not distinguish between the players Left and Right (see Table 1 for a more formal definition). The Sprague–Grundy theory [Spr36; Gru39; Con01; BCG04] states that every finite impartial game is equivalent to an instance of the game of Nim, characterized by a single natural number n. This theory has since been generalized to all impartial games by generalizing Nim to all ordinals n; see [Con01; Smi66].

Nim [Bou02] is a game played with several *heaps*, each with a certain number of tokens. A Nim game with a single heap of size n is denoted by $*n$ and is

called a *nimber*. During each move a player can pick any pile and reduce it to
any smaller nonnegative integer size. The game ends when all piles have size
0. Thus, a single pile $*n$ can be reduced to any of the smaller piles $*0$, $*1$, ...,
$*(n-1)$. Multiple piles in a game of Nim are independent, and hence any game
of Nim is a sum of single-pile games $*n$ for various values of n. In fact, a game
of Nim with k piles of sizes n_1, n_2, ..., n_k is equivalent to a one-pile Nim game
$*n$, where n is the binary XOR of n_1, n_2, ..., n_k. As a consequence, Nim can
be played optimally in polynomial time (polynomial in the encoding size of the
pile sizes).

Even more surprising is that *every* impartial two-player perfect-information
game has the same value as a single-pile Nim game, $*n$ for some n. The number
n is called the *G-value*, *Grundy-value*, or *Sprague–Grundy function* of the game.
It is easy to define: suppose that game x has k options $y_1, \ldots, y_k$ for the first
move (independent of which player goes first). By induction, we can compute
$y_1 = *n_1$, ..., $y_k = *n_k$. The theorem is that x equals $*n$ where n is the
smallest natural number not in the set $\{n_1, \ldots, n_k\}$. This number n is called the
minimum excluded value or *mex* of the set. This description has also assumed
that the game is finite, but this is easy to generalize [Con01; Smi66].

The Sprague–Grundy function can increase by at most 1 at each level of the
game tree, and hence the resulting nimber is linear in the maximum number of
moves that can be made in the game; the encoding size of the nimber is only
logarithmic in this count. Unfortunately, computing the Sprague–Grundy func-
tion for a general game by the obvious method uses time linear in the number
of possible states, which can be exponential in the nimber itself.

Nonetheless, the Sprague–Grundy theory is extremely helpful for analyzing
impartial two-player games, and for many games there is an efficient algorithm
to determine the nimber. Examples include Nim itself, Kayles, and various
generalizations [GS56b]; and Cutcake and Maundy Cake [BCG04, pp. 24–27].
In all of these examples, the Sprague–Grundy function has a succinct charac-
terization (if somewhat difficult to prove); it can also be easily computed using
dynamic programming.

The Sprague–Grundy theory seems difficult to generalize to the superficially
similar case of misère play, where the goal is to be the first player unable to
move. Certain games have been solved in this context over the years, including
Nim [Bou02]; see, e.g., [Fer74; GS56a]. Recently a general theory has emerged
for tackling misère combinatorial games, based on commutative monoids called
"misère quotients" that localize the problem to certain restricted game scenarios.
This theory was introduced by Plambeck [Pla05] and further developed by Plam-
beck and Siegel [PS07]. For good descriptions of the theory, see Plambeck's

survey [Plaa], Siegel's lecture notes [Sie06], and a webpage devoted to the topic [Plab].

2.3. Strategy stealing. Another useful technique in Combinatorial Game Theory for proving that a particular player must win is *strategy stealing*. The basic idea is to assume that one player has a winning strategy, and prove that in fact the other player has a winning strategy based on that strategy. This contradiction proves that the second player must in fact have a winning strategy. An example of such an argument is given in Section 4.1. Unfortunately, such a proof by contradiction gives no indication of what the winning strategy actually is, only that it exists. In many situations, such as the one in Section 4.1, the winner is known but no polynomial-time winning strategy is known.

2.4. Puzzles. There is little theory for analyzing combinatorial puzzles (one-player games) along the lines of the two-player theory summarized in this section. We present one such viewpoint here. In most puzzles, solutions subdivide into a sequence of moves. Thus, a puzzle can be viewed as a tree, similar to a two-player game except that edges are not distinguished between Left and Right. With the view that the game ends only when the puzzle is solved, the goal is then to reach a position from which there are no valid moves (normal play). Loopy puzzles are common; to be more explicit, repeated subtrees can be converted into self-references to form a directed graph, and losing terminal positions can be given explicit loops to themselves.

A consequence of the above view is that a puzzle is basically an impartial two-player game except that we are not interested in the outcome from two players alternating in moves. Rather, questions of interest in the context of puzzles are (a) whether a given puzzle is solvable, and (b) finding the solution with the fewest moves. An important open direction of research is to develop a general theory for resolving such questions, similar to the two-player theory.

3. Constraint logic

Combinatorial Game Theory provides a theoretical framework for giving positive algorithmic results for games, but does not naturally accommodate puzzles. In contrast, negative algorithmic results — hardness and completeness within computational complexity classes — are more uniform: puzzles and games have analogous prototypical proof structures. Furthermore, a relatively new theory called Constraint Logic attempts to tie together a wide range of hardness proofs for both puzzles and games.

Proving that a problem is hard within a particular complexity class (like NP, PSPACE, or EXPTIME) almost always involves a reduction to the problem from a known hard problem within the class. For example, the canonical problem to

reduce from for NP-hardness is Boolean Satisfiability (SAT) [Coo71]. Reducing SAT to a puzzle of interest proves that that puzzle is NP-hard. Similarly, the canonical problem to reduce from for PSPACE-hardness is Quantified Boolean Formulas (QBF) [SM73].

Constraint Logic [DH08] is a useful tool for showing hardness of games and puzzles in a variety of settings that has emerged in recent years. Indeed, many of the hardness results mentioned in this survey are based on reductions from Constraint Logic. Constraint Logic is a family of games where players reverse edges on a planar directed graph while satisfying vertex in-flow constraints. Each edge has a weight of 1 or 2. Each vertex has degree 3 and requires that the sum of the weights of inward-directed edges is at least 2. Vertices may be restricted to two types: AND vertices have incident edge weights of 1, 1, and 2; and OR vertices have incident edge weights of 2, 2, and 2. A player's goal is to eventually reverse a given edge.

This game family can be interpreted in many game-theoretic settings, ranging from zero-player automata to multiplayer games with hidden information. In particular, there are natural versions of Constraint Logic corresponding to one-player games (puzzles) and two-player games, both of bounded and unbounded length. (Here we refer to whether the length of the game is bounded by a polynomial function of the board size. Typically, bounded games are nonloopy while unbounded games are loopy.) These games have the expected complexities: one-player bounded games are NP-complete; one-player unbounded games and two-player bounded games are PSPACE-complete; and two-player unbounded games are EXPTIME-complete.

What makes Constraint Logic specially suited for game and puzzle reductions is that the problems are already in form similar to many games. In particular, the fact that the games are played on planar graphs means that the reduction does not usually need a crossover gadget, whereas historically crossover gadgets have often been the complex crux of a game hardness proof.

Historically, Constraint Logic arose as a simplification of the "Generalized Rush-Hour Logic" of Flake and Baum [FB02]. The resulting one-player unbounded setting, called *Nondeterministic Constraint Logic* [HD02; HD05], was later generalized to other game categories [Hea06b; DH08].

4. Algorithms for two-player games

Many bounded-length two-player games are PSPACE-complete. This is fairly natural because games are closely related to Boolean expressions with alternating quantifiers (for which deciding satisfiability is PSPACE-complete): there exists a move for Left such that, for all moves for Right, there exists another move for Left, etc. A PSPACE-completeness result has two consequences. First,

being in PSPACE means that the game can be played optimally, and typically all positions can be enumerated, using possibly exponential time but only polynomial space. Thus such games lend themselves to a somewhat reasonable exhaustive search for small enough sizes. Second, the games cannot be solved in polynomial time unless P = PSPACE, which is even "less likely" than P equaling NP.

On the other hand, unbounded-length two-players games are often EXPTIME-complete. Such a result is one of the few types of true lower bounds in complexity theory, implying that all algorithms require exponential time in the worst case.

In this section we briefly survey many of these complexity results and related positive results. See also [Epp] for a related survey and [Fra07] for a bibliography.

4.1. Hex. Hex [BCG04, pp. 743–744] is a game designed by Piet Hein and played on a diamond-shaped hexagonal board; see Figure 1. Players take turns filling in empty hexagons with their color. The goal of a player is to connect the opposite sides of their color with hexagons of their color. (In the figure, one player is solid and the other

Figure 1. A 5×5 Hex board.

player is dotted.) A game of Hex can never tie, because if all hexagons are colored arbitrarily, there is precisely one connecting path of an appropriate color between opposite sides of the board.

John Nash [BCG04, p. 744] proved that the first player to move can win by using a strategy-stealing argument (see Section 2.3). Suppose that the second player has a winning strategy, and assume by symmetry that Left goes first. Left selects the first hexagon arbitrarily. Now Right is to move first and Left is effectively the second player. Thus, Left can follow the winning strategy for the second player, except that Left has one additional hexagon. But this additional hexagon can only help Left: it only restricts Right's moves, and if Left's strategy suggests filling the additional hexagon, Left can instead move anywhere else. Thus, Left has a winning strategy, contradicting that Right did, and hence the first player has a winning strategy. However, it remains open to give a polynomial characterization of a winning strategy for the first player.

In perhaps the first PSPACE-hardness result for "interesting" games, Even and Tarjan [ET76] proved that a generalization of Hex to graphs is PSPACE-complete, even for maximum-degree-5 graphs. Specifically, in this graph game, two vertices are initially colored Left, and players take turns coloring uncolored vertices in their own color. Left's goal is to connect the two initially Left