

ACYCLIC CONSTRAINT LOGIC AND GAMES

Hendrik Jan Hoogeboom¹ Walter A. Kusters² Jan N. van Rijn³ Jonathan K. Vis⁴

Leiden Institute of Advanced Computer Science, Universiteit Leiden, The Netherlands

Abstract

Non-deterministic Constraint Logic is a family of graph games introduced by Demaine and Hearn that facilitates the construction of complexity proofs. It is convenient for the analysis of games, providing a uniform view. We focus on the acyclic version, apply this to Klondike, Mahjong Solitaire and Nonogram (that requires planarity), and discuss the more complicated game of Dou Shou Qi. While for the first three games we reobtain known characterizations in a simple and uniform manner, the result for Dou Shou Qi is new.

1. INTRODUCTION

Besides actually playing games, it is of great interest to know how hard these games are in the sense of computational complexity, see Kendall, Parkes, and Spoerer (2008). The games are usually generalized to allow for parameters that control board size, number of cards, etc.

In order to study the structural complexity of games, Hearn (2006) and Hearn and Demaine (2009) advocate the use of the constraint logic framework. It consists of a collection of abstract graph games. The games are played on a so-called *constraint graph*. A constraint graph is a weighted directed graph, where each edge has a weight in $\{1, 2\}$. The *inflow* of a vertex is defined to be the sum of all weights of the edges that are directed inward. A configuration (i.e., direction of the edges) of a constraint graph is legal if and only if for each vertex it holds that the inflow is at least its minimum inflow, usually 2. A move of a player is typically the reversal of one of the edges; players are only allowed to do moves that result in a legal configuration.

A notable feature of the constraint logic framework is the fact that constraint graphs can be reduced to equivalent planar versions. Many real-life games are played on a 2-dimensional board. In previous game complexity results (e.g., Culberson (1999), Flake and Baum (2002)) crossover gadgets are necessary to overcome the limitations of such a 2-dimensional game board. Crossover gadgets are in general complex and hard to construct. The generic crossover gadget for constraint logic, as presented in Figure 2 below, removes the need to devise a specific crossover gadget for every single game.

Various games based on constraint graphs are defined in Hearn and Demaine (2009). These are categorized based on the number of players and whether there is a bound on the number of moves. We will describe two of those: *Bounded Non-deterministic Constraint Logic* and *Bounded Two-Player Constraint Logic*. In particular we will also pay attention to acyclic versions and planarity issues.

The remainder of this paper is organized as follows. First we explain the different types of graph games. Next we apply these to the games Klondike, Mahjong Solitaire, Nonogram and Dou Shou Qi.

¹h.j.hoogeboom@liacs.leidenuniv.nl

²w.a.kusters@liacs.leidenuniv.nl

³j.n.van.rijn@liacs.leidenuniv.nl

⁴j.k.vis@liacs.leidenuniv.nl

1.1 Bounded Non-deterministic Constraint Logic

Bounded Non-deterministic Constraint Logic (Bounded NCL) is a one-player game (i.e., a puzzle), played on a constraint graph. A move is defined to be the reversal of one of the edges, resulting in a legal configuration, i.e., meeting the inflow condition of the vertices. Each edge may be reversed at most once. This puts an upper bound on the number of moves in this game, i.e., the number of edges in the graph. One of the edges is defined to be the *target edge*; the player wins if and only if (s)he is able to reverse the target edge.




Figure 1: Basic vertices, based on Figure 5.2 and Figure 6.2 from Hearn (2006). Edges with a weight of 2 use thick lines and have double arrows; edges with a weight of 1 use thin lines and have a single arrow. Usually these edges are referred to as “blue” and “red”, respectively.

Theorem 5.1 and Theorem 5.2 from Hearn and Demaine (2009) show (using a reduction from the Boolean satisfiability problem) that the game is NP-complete, even when the initial constraint graph only consists of AND, OR, FANOUT and CHOICE vertices as shown in Figure 1.

1.2 Bounded Two-Player Constraint Logic

Bounded Two-Player Constraint Logic (Bounded 2CL) is a two-player perfect-information game played on a constraint graph, and a partitioning of the edges in disjoint sets B and W . The players alternate turns. The white player reverses edges in W ; the black player reverses edges in B . For both players it holds that their move has to result in a legal configuration. Each edge may only be reversed once, which (as in Bounded NCL) puts an upper bound on the number of moves in the game. One of the edges in W is defined to be the target edge. The white player wins if (s)he is able to reverse this edge; if a player is unable to move, (s)he loses the game.

Theorem 6.2 from Hearn and Demaine (2009) shows that the game is PSPACE-complete, even when the constraint graph only consists of the five vertices as shown in Figure 1, where the edges from AND, OR, FANOUT and CHOICE vertices are all in the set W . Note that the black player can only play bottom edges in VARIABLE gadgets. In order to avoid clear loss for black an ample amount of additional black edges is supplied.

1.3 Acyclic graphs and crossover gadgets

In order to planarize constraint graphs, the construction shown in Figure 2(a) can be used. A pair of crossing edges can be replaced by this gadget. To obtain basic vertices as in Figure 1, each vertex with four red edges can be replaced by the so-called half-crossover gadget, which is shown in Figure 2(b). Additionally, we need to perform red-blue edge conversions, see Hearn (2006).

Edges A and B are called the *vertical external edges*. Edges C and D are called the *horizontal external edges*. It can be verified that each vertical external edge can point outward if and only if the other vertical external edge points inward. A similar property holds for the horizontal external edges. The action of reversing both vertical external edges is called *vertical propagation*; the action of reversing both horizontal external edges is called *horizontal propagation*. For example, when the edges A and B are pointing up, and the edges C and D are pointing left, the direction of all other edges follows from the inflow constraints. A sequence of, e.g., (A, F, H, G, M, O, N, B) would then perform a vertical propagation; a sequence of, e.g., $(C, K, I, L, J, E, P, R, Q, S, D)$ would perform a horizontal propagation.




Figure 2: Planar crossover gadgets, as presented in Hearn and Demaine (2009).

Although this gadget indeed simulates all the behavior of the games introduced in, e.g., Hearn and Demaine (2005) (where the same edge can be reversed multiple times) this is not the case for the constraint graphs used in Bounded NCL and Bounded 2CL. After performing a vertical propagation, due to the restriction that each edge may be reversed at most once, the gadget is in such a state that it is impossible to perform horizontal propagation, and vice versa. Although the construction in Figure 2(a) is valid, after integrating the construction of Figure 2(b) for all the vertices with four red edges (in order to restrict ourselves to the gadgets in Figure 1) it becomes clear that this is no longer the case. After, e.g., edges F and H (corresponding to a and b , respectively) are reversed, it can be verified that due to the internal state of the half-crossover, edges K and I (corresponding to c and d , respectively) cannot be reversed anymore. The only way to perform both a horizontal and vertical propagation over the same crossover gadget is when both a horizontal external edge and vertical external edge can be reversed inward at the same moment: a typical example of a *race condition*. An extensive analysis of the properties of the crossover and half-crossover gadgets as well as red-blue edge converters can be found in Hearn (2006) and van Rijn (2012).

It is clear that not all constraint graphs can be reduced to a planar equivalent using solely the gadgets presented in Figure 2. In some configurations, in particular in (initially) cyclic graphs, it is impossible to obey the additional constraint imposed by the crossover gadget that the propagation of both directions has to happen at the same moment. However, for acyclic graphs, this is never a problem. Edges can be topologically sorted, and reversed in this order. The complexity proofs of both Bounded NCL and Bounded 2CL (see Hearn and Demaine (2009)) use graphs corresponding to logical formulas, which indeed require only acyclic graphs; hence Planar Bounded NCL is NP-complete and Planar Bounded 2CL is PSPACE-complete. Note that all graphs under consideration are acyclic in their initial configuration.

Theorem 5.4 from Hearn and Demaine (2009) states that the related problem Constraint Graph Satisfiability is also NP-complete: does a given planar constraint graph, using only (initially undirected) AND and OR vertices, have a legal configuration? Note that this strictly speaking is not a game in the above sense: we only ask for a legal “final” configuration, not the sequence of moves that can be used to obtain it.

2. KLONDIKE

Klondike, also known as Patience or Solitaire, is a well-known card game, popularized by Microsoft Windows. The normal version of the game is played with a standard French card deck, without jokers. Yan *et al.* (2004) have given a formal definition of the game and provided an algorithm that plays Klondike games with a high success rate; in their version of the game, often referred to as thoughtful Solitaire, the identity of all cards is known from the beginning. Several other approaches have been proposed to deal with Klondike, see, e.g., Bjarnason, Tadevall, and Fern (2007) and Bjarnason, Fern, and Tadevall (2009).

Longpré and McKenzie (2009) have shown, amongst other complexity results, that Klondike is NP-complete even when played with two red suits and one black suit: red diamonds (\heartsuit), red hearts (\heartsuit) and black spades (\spadesuit). We will give a formal definition of the necessary subset of Klondike and confirm the NP-completeness of Klondike by a reduction from Acyclic Bounded NCL, using an argument improved upon the one from van Rijn (2012).

2.1 Definition

Generalized Klondike is played with a card deck containing m suits, each suit containing n cards ranked from 1 to n . A card with rank 1 is also referred to as an *Ace*; a card with rank n is also referred to as a *King*. The functions $rank(c)$ and $suit(c)$ return the rank and the suit of card c , respectively. Each suit is colored either red or black. The function $color(s)$ returns the color of suit s .

A Klondike game consists of m *suit stacks*, one or more *build stacks*, a *pile stack* and a *talon*. In the sequel we do not need pile stack and talon, so these will be omitted from the description. A stack is defined to be an ordered list of cards. A *configuration* describes for each card in which stack it is and on which position. For every card in a build stack it also describes whether the card is *face-up* or *face-down*. The subset of cards that are face-up on a certain build stack constitute a *card block*, and will always consist of topmost cards. In an initial configuration all cards are face-down in the build stacks (that can be of different lengths), and the suit stacks are empty.

We will define the notion of *acceptance*, which determines which moves the player can make. Each suit stack that is empty can only accept an Ace. Every suit stack that is not empty, containing card t on top, accepts card c if and only if $suit(c) = suit(t)$, and $rank(c) = rank(t) + 1$. Therefore, suit stacks accept cards of the same suit in ascending order. Each card block that is not empty, containing card t on top, accepts card c if and only if $color(suit(t)) \neq color(suit(c))$ and $rank(c) = rank(t) - 1$. Therefore, build stacks accept cards in descending order, of alternating colors. We will not employ the usual property that an empty build stack (only) accepts a King.

On each turn, the player can play cards in the following manner:

1. If all cards on a build stack are face-down, the card on top can be turned face-up, thereby creating a singleton card block.
2. A whole card block p can be moved to the top of another card block q , provided that q accepts the card at the bottom of p . (In some versions of the game a partial card block can also be moved in this manner.)
3. The top card c of a card block can be moved to a suit stack, provided that the suit stack accepts c .

The goal is to move all cards to the suit stacks, and when this is achieved the player has won.

2.2 NP-completeness

In order to prove NP-completeness, we will show that every Acyclic Bounded NCL graph can be transformed to a Klondike configuration, in such a way that the Klondike game can be won if and only if the target edge of the Acyclic Bounded NCL graph can be flipped. So we study the corresponding decision problem KLONDIKE: given a Klondike configuration, can the player win?

We will use the four gadgets from Figure 3. The gadgets consist of one, two or three build stacks, with all cards initially face-down. Each gadget gets a range of unique ranks assigned to it; for simplicity, in the figure we use the range 5–8 for all gadgets. A *lock card* represents the tail of an edge adjacent to the corresponding NCL vertex; a lock icon is displayed on these cards. A *key card* represents the head of an edge adjacent to the corresponding NCL vertex; a key icon is displayed on these cards. The rank of each key card is within the range of another gadget. The suit and rank of a key card is chosen such that once turned face-up, it accepts a lock card of the gadget from which the corresponding NCL edge is pointing (locks must be moved to their keys). The gadgets are constructed in such a way that the key card can be turned face-up if and only if the corresponding NCL edge can be flipped. Note that in the AND gadget ♠5 is also a lock card (without having a lock image).

For each lock card ℓ it is easy to see which card should be turned face-up in order to move it. If $color(rank(\ell))$ is red, this card is ♠($rank(\ell) + 1$), and otherwise it is ♡($rank(\ell) + 1$) (the card $\diamond(rank(\ell) + 1)$ will be only made available during the end play, or in the case of the OR gadget is positioned deeper within the gadget). These cards serve as key cards in other gadgets.

Now we note that the four gadgets indeed act as intended. For instance, consider the OR gadget. In order to turn the key card face-up, either the lock card ♠6 (followed by $\diamond 5$) must be moved to its corresponding key card ♡7, or the lock card ♡6 (followed by ♠5) must be moved to its corresponding key card ♠7 after which ♠6 and $\diamond 5$

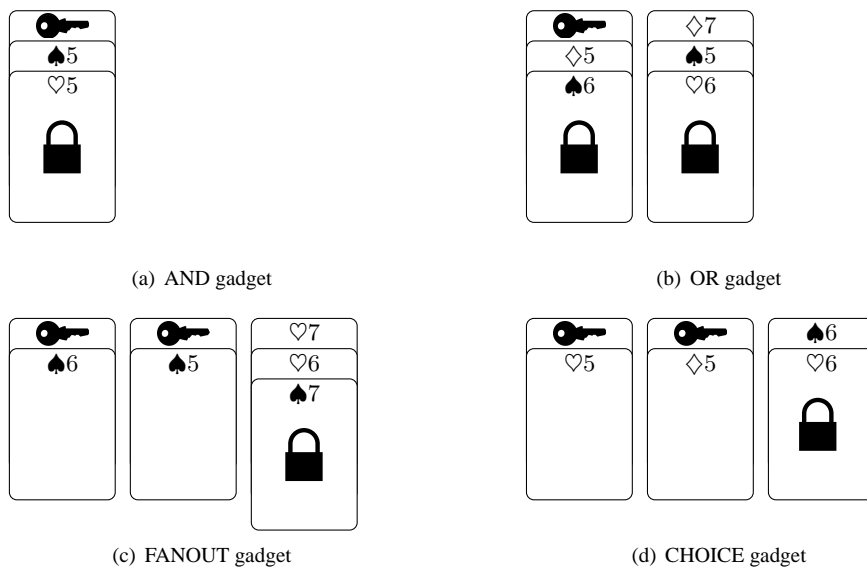


Figure 3: Klondike gadgets. Note that ♠5 in the AND gadget is also a lock card.

can be moved to $\diamond 7$. If both key cards are available, both sequences can be played. The AND gadget has a fixed order to free the key card, which is sufficient for our purpose.

Now we have:

Theorem 2.1 *KLONDIKE is NP-complete.*

Proof Reduction from Acyclic Bounded NCL. Given a constraint graph made of AND, OR, FANOUT and CHOICE vertices, we construct a corresponding Klondike configuration using the gadgets shown in Figure 3. Note that planarity is not an issue both here and for Mahjong.

We need a way to ensure that the player can move all cards to the suit stacks if and only if the key card corresponding to the target edge can be turned face-up. To this end, all cards not used in the gadgets are positioned in one big build stack (ordered by rank and within each rank in $\diamond \heartsuit \spadesuit$ order, with the three Aces at the top and ending with the three Kings at the bottom), protected by a lock card representing the target edge. Once this card is moved, all these other cards become available and allow all cards from all gadgets to be moved to the suit stacks.

Now the fact that the original NCL graph is acyclic is used. Indeed, the gadgets can be numbered, using a topological sort of the corresponding nodes, and we can take care that for every gadget the key cards used have higher rank than the cards in the gadgets. This ensures a proper order for this part of the process. In fact, even (partially) unplayed gadgets can be “discarded”, using the inductive assumption that all cards of lower ranks have already been moved to the suit stacks. Note that, if for the OR gadget $\heartsuit 5$ were used instead of $\diamond 5$, this property would not hold.

For creating the Klondike configuration, the number of cards and stacks we need are both bounded by a linear function of the number of vertices in the corresponding NCL graph. In a winning sequence there are exactly mn moves of type 1, and mn moves of type 3. As for the type 2 moves, there are at most mn of them: every card block is moved once (when focussing on its bottom card), maybe to a suit stack. Therefore Klondike is in NP, since any potential solution can be verified in polynomial time. \square

3. MAHJONG SOLITAIRE

Mahjong Solitaire, also known as Shanghai Solitaire, is a one player puzzle game mainly played on the computer in which the player is presented with a randomly arranged stack of tiles. The goal is to remove all tiles in matching pairs of two. Condon *et al.* (1997) have given a formal definition of this game, and have shown that a version of this game with imperfect information is PSPACE-complete. Eppstein (2012) has stated a proof that a version of this game with perfect information is NP-complete. In the paper by de Bondt (2012) Mahjong is proven to

be NP-complete by a reduction from 3-SAT. We will give a formal definition of this game and validate the latter result by a reduction from Acyclic Bounded NCL.

3.1 Definition

The game uses Mahjong tiles, that are divided into m disjoint *tile sets* \mathcal{T}_p of $|\mathcal{T}_p| = s_p$ matching tiles, where s_p is an even number ($p = 1, 2, \dots, m$). We define the set of all tiles to be $\mathcal{T} = \bigcup_p \mathcal{T}_p$. Two tiles a and b *match*, if and only if for some p it holds that $a, b \in \mathcal{T}_p$. Below we say that elements of the same tile set have the same color. We generalize the standard game simply by assuming that there is an arbitrarily large, finite number of tiles. A *configuration* C is a set of positions (i, j, k) , where each of i, j, k is a non-negative integer, satisfying the following constraints:

1. If $(i, j, k) \in C$ and $(i, j', k) \in C$ where $j < j'$, then for every j'' in the range $[j, j']$, $(i, j'', k) \in C$;
2. If $(i, j, k) \in C$ where $k > 0$ then $(i, j, k - 1) \in C$.

This captures the fact that tiles are arranged in three dimensions. Tiles can be stacked on top of each other; all tiles with common k are at the same height. All tiles with a common i index, form a *cross section*. Tiles at the same height, with common i index, form a *row*. The first condition ensures that there cannot be gaps in a row; the second, that a tile at height $k > 0$ must have a tile underneath it (in fact, at position $(i, j, k - 1)$).

With respect to a given configuration, a position (i, j, k) is *hidden* if in the configuration also a position $(i, j, k + 1)$ exists; the other positions are called *visible*. An *arrangement* consists of a set of tiles \mathcal{T} , a configuration C of size $|\mathcal{T}|$, and a bijective function f from the positions of C to \mathcal{T} . Here $f(i, j, k)$ denotes the tile at position (i, j, k) . If the function f maps position (i, j, k) to tile t we say $pos(t) = (i, j, k)$. The elements of \mathcal{T} will be mapped to the elements of C in such a way, that every combination is possible. With respect to a given arrangement, we say a position (i, j, k) is *available* if it is not hidden, and either position $(i, j - 1, k) \notin C$ or position $(i, j + 1, k) \notin C$ or both, i.e., we can only take tiles that are at one of the ends of a row, and that have no tiles on top of it. An arrangement is called empty if \mathcal{T} is empty. In order to avoid misunderstandings, all tiles can be seen from the beginning: the player has perfect information.

A legal move consists of the removal of two matching tiles a, b that are both available. Formally, $\mathcal{T}' = \mathcal{T} - \{a, b\}$ and $C' = C - \{pos(a), pos(b)\}$. The game is won if a series of moves results in the empty arrangement.

3.2 NP-completeness

In order to prove NP-completeness, we will show that every Acyclic Bounded NCL graph can be transformed to a Mahjong configuration, in such a way that the Mahjong game can be won if and only if the target edge of the Acyclic Bounded NCL graph can be flipped. We study the decision problem MAHJONG SOLITAIRE: given a Mahjong configuration, can the player win?

We will use the gadgets in Figure 4. The gadgets consist of one cross section containing between two and five tile stacks, with different numbers representing different tile sets. Every gadget has a unique range of numbers. Again, for simplicity, in the figure we use the range 1–7 for all gadgets. A *lock tile* represents the tail of an edge adjacent to the corresponding NCL vertex; a lock icon is displayed on these tiles. A *key tile* represents the head of an edge which is adjacent to the corresponding NCL vertex; a key icon is displayed on these tiles. Corresponding key and lock tiles from different gadgets share the same number, for obvious reasons.

When a key tile is available, it can be removed together with a lock tile from one of the other gadgets. The target edge in the corresponding NCL graph is always represented by a key tile. When this target edge is available this will initiate the end game, which is always winning for the player as we will show further on.

The four gadgets in Figure 4 have their intended behavior. For instance, consider the CHOICE gadget. To free either one of the key tiles the lock tile has to be removed. Now, the actual choice has to be made: the newly freed “3-tile” must be used to remove either the leftmost or rightmost “3-tile”. After removing both “4-tiles” precisely one of the key tiles is available. Note that the gadgets resemble those for Klondike; in fact, the AND gadget can also be modeled to look even more like its counterpart.

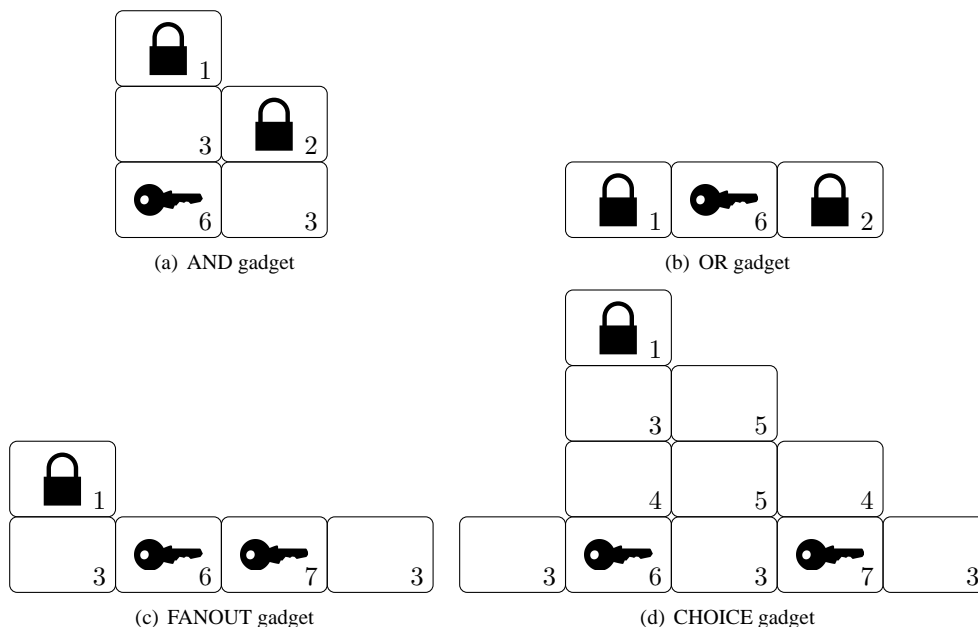


Figure 4: Mahjong gadgets as a cross section of a configuration.

Now we have:

Theorem 3.1 MAHJONG SOLITAIRE is NP-complete.

Proof Reduction from Acyclic Bounded NCL. Given a constraint graph made of AND, OR, FANOUT and CHOICE vertices, we construct a corresponding Mahjong configuration using the gadgets shown in Figure 4.

In order to have a way of clearing all remaining tiles after the target tile is removed, we supply a victory gadget consisting of one linear row of tiles with two “5-tiles” for every CHOICE gadget. The victory gadget itself is protected in a similar way as the FANOUT gadget: a pair of matching tiles is placed at both sides, and a tile matching the target tile is placed on the left one of these. This ensures that none of the tiles in the victory gadget can be used before the tiles representing the target edge are removed. Again, we will use the fact that the original NCL graph is acyclic by numbering the gadgets using the topological sort of the corresponding vertices. This ordering defines the proper order for this process allowing partially unplayed (CHOICE) gadgets to be removed by using the tiles from the victory gadget, that are placed in this same order. Note that from this acyclicity it follows that when the player is able to remove all CHOICE gadgets, all other (partially) unplayed gadgets can also be removed.

Both the number of tiles and the number of tile sets we need to use is linearly bounded by the number of vertices used in the corresponding Acyclic Bounded NCL graph. Therefore Mahjong Solitaire is in NP, since any potential solution can be verified in polynomial time. \square

4. NONOGRAM

A *Nonogram*, also referred to as a Japanese puzzle, is a logic puzzle which can be considered as an image reconstruction problem. The player is presented a rectangular grid; for each row and column a description consisting of one or more integers is provided, representing the numbers of consecutive cells that need to be black. If the player can color a subset of cells in such a way that it is consistent with the description of all rows and columns, (s)he has solved the puzzle and won the game. An example of a Nonogram and its solution is shown in Figure 5. Batenburg and Kosters (2012) have given a formal definition of Nonograms and provided an algorithm for solving many Nonograms in polynomial time. In Nagao *et al.* (1996) it is proven that the Another Solution Problem for Nonograms is NP-complete, and more in particular that the question whether a given puzzle has a solution is NP-complete. We will also give a formal definition of Nonograms and show that the latter decision problem is NP-complete, by reduction from Constraint Graph Satisfiability.




Figure 5: An example Nonogram (a) and its unique solution (b), taken from Batenburg and Kosters (2012).

4.1 Definition

A Nonogram is a puzzle in which the player is presented an $m \times n$ grid of *cells*, consisting of m rows and n columns. The state of a cell is either *white* = 0 or *black* = 1. Initially, all cells are *white*. A *line* is defined to be either a row or a column.

For each line a description d is provided, d being an ordered series of integers (d_1, d_2, \dots, d_k) . The description is adhered to, if there are exactly k *black segments* in the line, where each successive black segment s is of size d_s ($s = 1, 2, \dots, k$). A black segment is defined to be a group of consecutive cells in the line, such that all cells within the interval are *black*, and both cells adjacent to the interval, if any, are *white*. Now the puzzle is solved if the player can make a subset of the cells black, in such a way that all descriptions are adhered to. The corresponding decision problem NONOGRAM asks if a given Nonogram can be solved. For more information on Nonograms, the reader is referred to Batenburg and Kosters (2012) and the references therein.

4.2 NP-completeness

We will show that solving Nonograms is NP-complete, by reduction from Constraint Graph Satisfiability (Hearn and Demaine (2009)), only using two initially undirected gadgets: AND and OR. The global layout of the construction will be as in Figure 6. There will be several groups of D adjacent columns (or rows) where the description consists of a single element, i.e., m (or n), such that the pattern of Figure 6 is maintained. We call these lines the *separation lines*. Between each group of separation lines, there are G other lines. In the case of Figure 6, $D = 5$ and $G = 7$. The descriptions and the width of the delimiters will not interfere with those of the single elements in between. As a result of this construction we can specify disjoint *subnonograms* between the separation lines.

It is also possible to send a signal between two orthogonal adjacent subnonograms by slightly adjusting delimiters between them. This is illustrated in Figure 7. The figure shows two subnonograms separated by $D = 3$ separation lines; one subnonogram between cells $(1, 1)$ and $(4, 4)$, inclusive, and one subnonogram between cells $(1, 8)$ and $(4, 11)$, inclusive. If we were to decide that $(3, 4)$ should be *black*, this would explicitly mean that cell $(3, 8)$ cannot be *black*. (Note that this would also explicitly mean that $(2, 8)$ is *black*, and $(2, 4)$ is not *black*.) The opposite is also true. We will use this property to construct gadgets within a subnonogram, and propagate signals between them. This way we can embed a constraint graph on a Nonogram grid.

In Figure 8 a template of the gadgets is shown. From the description of every gadget follows immediately that the black cells must be *black* and the dotted cells must be *white*. The state of the other cells is dependent on the type and state of the gadget.

The gadgets are shown in Figure 9 and have $G = 7$. As it does not influence the functionality, on each side these are surrounded by only one separation line. (In the large construction we will use $D = 5$, with obvious adaptations of the descriptions.) These are already *black*. If a cell corresponding to an edge is *white*, this means that the edge is pointing away from the vertex, and vice versa. Besides AND and OR gadgets, we also provide




Figure 6: Global layout.




Figure 7: The two solutions of a Nonogram featuring two subnonograms horizontally separated by 3 separation lines.




Figure 8: Template for Nonogram gadgets. Squares that are black are black in all instances; squares with a dot are white in all instances.

two gadgets needed for wiring.

The three cells marked with a, b and c correspond with the edges; for the (initially undirected) AND gadget a and b correspond with the “red” ones. Given the values for the cells corresponding to the edges, the gadgets are within the so-called “simple” Nonogram class, and can be easily solved. For a definition and solving algorithm of the “simple” nonograms, the reader is referred to Batenburg and Kosters (2012). In any case, the solutions can be easily verified. They are uniquely characterized by $(a, b, c) \in \{(0, 0, 1), (1, 0, 1), (0, 1, 1), (1, 1, 0), (1, 1, 1)\}$ (AND gadget) and $(a, b, c) \in \{(1, 0, 0), (0, 1, 0), (0, 0, 1), (1, 0, 1), (0, 1, 1), (1, 1, 0), (1, 1, 1)\}$ (OR gadget), implying that the gadgets indeed perform as desired. If $(a, b, c) = (1, 0, 1)$, the OR gadget has two solutions, making use




Figure 9: Nonogram gadgets.

of a so-called switching component.

Now we have:

Theorem 4.1 *NONOGRAM is NP-complete.*

Proof We can simulate a planar constraint graph with only AND and OR nodes on a Nonogram grid using the global layout of Figure 6 and the two top gadgets shown in Figure 9. The two bottom gadgets from Figure 9 can be used for wires, in a straight line or as a corner. Now the arrows can be inserted in a legal way if and only if the resulting Nonogram can be solved.

Di Battista *et al.* (1994) have proven that a graph with maximal degree 3 can always be stored in a square grid of width v , where v is the number of vertices contained by the graph. This ensures that the number of rows and columns used in our reduction is bounded linearly by the number of vertices in the corresponding NCL graph.

NONOGRAM is clearly in NP, as any potential solution can be verified in polynomial time. □

5. DOU SHOU QI

Dou Shou Qi (meaning: “Game of Fighting Animals”), as described by Pritchard and Beasley (2007), is a Chinese board game. In the Western world it is often called *Jungle*, *The Jungle Game*, *Jungle Chess*, or *Animal Chess*. *Dou Shou Qi* is a two-player abstract strategy game and it contains some elements from Chess and Stratego as well as some other chess-like Chinese games (e.g., Banqi). Its origins are not entirely clear, but it seems that it

evolved rather recently (around the 1900s). It has been suggested by some that the game often ends in a draw, but preliminary results from van Rijn and Vis (2013) show a remarkably low percentage of draws.

The game Dou Shou Qi is not extensively studied in literature. In Burnett (2010), a definition of the game is given and an attempt is made to characterize certain local properties of subproblems that occur when analyzing the game. These so-called loosely coupled subproblems can be analyzed separately in contrast to analyzing the problem as a whole, resulting in a possible speed-up in the overall analysis. A first complexity result has been obtained by van Rijn and Vis (2013). Dou Shou Qi is proven PSPACE-hard by reduction from logic circuits. Here, we will prove Dou Shou Qi to be PSPACE-hard by reduction from Planar Bounded 2CL based on the reduction given by van Rijn (2012).

We cannot prove Dou Shou Qi to be PSPACE-complete; as an unbounded two-player game it is probably not in PSPACE. Dou Shou Qi is clearly in EXPTIME — like Chess.

5.1 Definition

Dou Shou Qi is played on a rectangular board consisting of 9×7 squares, see Figure 10. There are several different kinds of squares. The *dens* (D), one for each player, are located in the center of the first and last row and are protected on all sides by *traps* (T). Furthermore, there are two bodies of water (W), while the remaining squares are ordinary land squares.

Each player has eight different pieces representing different animals with a respective *strength*, according to which they can *capture* some of the opponent's pieces. Pieces can only capture a piece of equal or lower strength, with the exception of the weakest piece which is able to capture the strongest piece. The strength of the animals from weak to strong is: 1 rat, 2 cat, 3 wolf, 4 dog, 5 panther, 6 tiger, 7 lion, 8 elephant. The initial placement of the pieces is fixed, see Figure 10. Players alternate turns with white moving first. Each turn a piece must be moved either one square horizontally or vertically. Pieces are forbidden to enter their own den and are usually blocked by water. The rat is the only piece able to move through the water where it is also capable of capturing, i.e., the enemy rat, but it is forbidden to capture an elephant while attacking from the water. Lions and tigers are able to leap over water either horizontally or vertically, but they are blocked by any rat on the intermediate water squares.

7		T	D	T		6
	4		T		2	
1		5		3		8
	W	W		W	W	
	W	W		W	W	
	W	W		W	W	
8		3		5		1
	2		T		4	
6		T	D	T		7

Figure 10: A schematic Dou Shou Qi game board showing the initial position.

Pieces can be trapped by the traps surrounding the opponent's den: their strength is effectively reduced to zero, meaning that they can be captured by any enemy piece. The objective of the game is to place a piece in the opponents den or to eliminate all of the opponent's pieces. A stalemate position is declared a draw.

5.2 PSPACE-hardness

The Bounded 2CL graph will be simulated on a $m \times n$ board, where both players have k pieces. Whether a natural generalization of the game would imply that the k pieces all have a strength from the interval $[1, 8]$ or a distinct strength from the interval $[1, k]$ is open for discussion. In our reduction all pieces have a strength from the interval $[2, 5]$, excluding all pieces with special capabilities. The original game board contains several properties, i.e., clustered water squares, narrow paths between the water, traps surrounding the dens, which are symmetrical and highly regular. Which of these properties should be preserved on a generalized game board is open for debate, however in our reduction we took the liberty of using water squares and traps freely in the gadgets.

We will show a complexity proof for the decision problem DOU SHOU QI: given a Dou Shou Qi position, does the player on turn have a forced win?

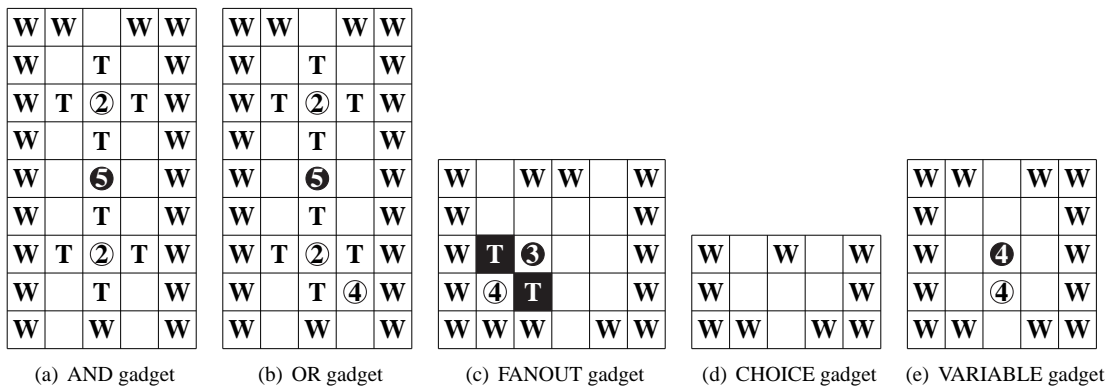


Figure 11: Dou Shou Qi gadgets.

We will reduce from Planar Bounded 2CL. The main gadgets are shown in Figure 11. The reversal of an edge in the original Bounded 2CL graph will be modeled as the movement of a white dog (strength 4) into another gadget. The VARIABLE vertex (in its initial state) can be reversed by the current player. The same is true for the VARIABLE gadget. Since both pieces are of the same strength, the current player can capture the piece of the other player, and move its piece through to the next gadget.

There are some additional issues that need to be addressed. First, white dogs that enter a gadget should not be allowed to go back into the previous gadget. Next, white dogs in a FANOUT gadget should not be allowed to move through the same exit twice. Finally, black pieces in a VARIABLE gadget should not be allowed to leave the gadget through the exit corresponding to the white edge in the graph game. In order to prevent this behavior, we have created some additional support gadgets, that will be attached to the inputs and outputs of the gadgets. These are shown in Figure 12.

The construction shown in Figure 12(a) is called a *black edge protector*. The white player can move a dog from bottom to top, but not the other way around. When a white dog enters the construction, the black piece will retreat behind either the left trap or the right trap, and the white dog can pass. When passed, the black piece moves back to its original position. The white piece cannot move back, it would be captured when entering the traps. The construction shown in Figure 12(b) is a *white edge protector*, it allows only white pieces to pass. Black pieces can be captured upon entering a trap. Note that these constraints do not apply when the opposing player attacks from both sides. We will show further on how to deal with this. The construction in Figure 12(c) is an outflow protector, with the left entrance square as input and the right entrance square as output. It ensures that upon arrival of either one or two white pieces, only one can pass.

A chain of two white edge protectors, one black edge protector and another two white edge protectors is called a *one-way channel*. First, it ensures that no black piece can move through it. It will be captured upon entering a white edge protector. After a capture, the white cat can resume its position, preventing black pieces from passing, regardless of their number. Because it is always adjacent to another white edge protector, even an attack from both sides is useless. Next, it ensures that all black pieces within are unable to move out of the construction they started in, by the same argument. Finally, linking several one-way channels to each other ensures that white pieces can move through it in only one direction. White pieces that move in the opposite direction will be stopped at the black edge protector. Indeed, when having a piece at both the input and the output, the white player can enable




Figure 12: Constructions used to support gadgets.

its piece at the output to move back through this gadget. However, in order to pass a number of subsequent black edge protectors, the white player needs an equal number of white dogs at the input to ensure such a passing. There can never be more than two white pieces at the input of a one-way channel, thus linking three one-way channels together prevents white pieces from moving into the forbidden direction. A gadget protector is a chain of three one-way channels, one outflow protector and another three one-way channels. The gadget protector is attached to every entrance of the gadgets shown in Figure 11, ensuring that these facilitate exactly the same behavior as their equivalents in the graph game.

Now we have:

Theorem 5.1 *DOU SHOU QI is PSPACE-hard.*

Proof Reduction from Bounded 2CL. Given a planar constraint graph made of AND, OR, FANOUT, CHOICE and VARIABLE vertices, we construct a corresponding Dou Shou Qi game board where the white player has a forced win if and only if (s)he has a forced win on the original Bounded 2CL graph; otherwise the black player has a forced win. Note that there are no draws in Bounded 2CL, neither are there in the reduction by optimal play.

The target edge will be represented by a gadget containing a black den, and it will have a black edge protector (Figure 12(a)) in front of it, preventing other pieces than the white dogs from entering it. The white player can move a piece into this gadget if and only if (s)he can set the corresponding Bounded 2CL graph to true. The black player is given a piece that can move straight to the white den. This will take him so many moves, that if the corresponding Bounded 2CL graph can be set to true, by the time (s)he reaches it the white player has already won the game. \square

6. CONCLUSIONS

We reduced Acyclic Bounded NCL to KLONDIKE and MAHJONG SOLITAIRE and (planar) Constraint Graph Satisfiability to NONOGRAM, proving them to be NP-complete. By using the acyclic property to our advantage, we were able to keep the reductions elegant and easy to understand. For games that require to return to an “empty” configuration (like Klondike and Mahjong Solitaire) acyclicity is even technically essential. We acknowledge the NCL framework to be well-suited for reductions for games, but it is not without drawbacks. Often the primary gadgets are relatively easy to construct, while the construction of victory gadgets is sometimes less trivial. Finally, we reduced Planar Bounded 2CL to the game of Dou Shou Qi proving it to be PSPACE-hard. The generic planarization of the NCL graphs and 2CL graphs is very useful for reductions to games played on a 2-dimensional board. As an unbounded two-player game Dou Shou Qi is expected to be EXPTIME-complete in the classification by Hearn and Demaine (2009). It is an open problem to construct the Dou Shou Qi gadgets for the special vertices, e.g., multiplayer AND, that build the relevant 2CL graph.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees, in particular for the suggested simplification regarding the Nonogram problem.

7. References

- Batenburg, K. J. and Kusters, W. A. (2012). On the difficulty of Nonograms. *ICGA Journal*, Vol. 35, pp. 195–205.
- Bjarnason, R., Fern, A., and Tadepalli, P. (2009). Lower bounding Klondike Solitaire with Monte-Carlo planning. *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS-2009)*, pp. 26–33, AAAI.
- Bjarnason, R., Tadepalli, P., and Fern, A. (2007). Searching Solitaire in real time. *ICGA Journal*, Vol. 30, pp. 131–142.
- Bondt, M. de (2012). Solving Mahjong Solitaire boards with peeking. *CoRR*, Vol. abs/1203.6559.
- Burnett, J. W. (2010). Discovering and searching loosely coupled subproblems in Dou Shou Qi. M.Sc. thesis, Tufts University.
- Condon, A., Feigenbaum, J., Lund, C., and Shor, P. (1997). Random debaters and the hardness of approximating stochastic functions. *SIAM J. Comput.*, Vol. 26, pp. 369–400.
- Culberson, J. (1999). Sokoban is PSPACE-complete. *Proceedings in Informatics*, Vol. 4, pp. 65–76.
- Di Battista, G., Eades, P., Tamassia, R., and Tollis, I. G. (1994). Algorithms for drawing graphs: An annotated bibliography. *Comput. Geom. Theory Appl.*, Vol. 4, pp. 235–282.
- Eppstein, D. (2012). Website: Computational complexity of games and puzzles. <http://www.ics.uci.edu/~eppstein/cgt/hard.html#shang>.
- Flake, G. W. and Baum, E. B. (2002). Rush Hour is PSPACE-complete, or “Why you should generously tip parking lot attendants”. *Theoretical Computer Science*, Vol. 270, pp. 895–911.
- Hearn, R. A. (2006). *Games, puzzles, and computation*. Ph.D. thesis, Massachusetts Institute of Technology.
- Hearn, R. A. and Demaine, E. D. (2005). PSPACE-completeness of sliding-block puzzles and other problems through the Nondeterministic Constraint Logic model of computation. *Theoretical Computer Science*, Vol. 343, pp. 72–96.
- Hearn, R. A. and Demaine, E. D. (2009). *Games, puzzles, and computation*. AK Peters.
- Kendall, G., Parkes, A., and Spoerer, K. (2008). A survey of NP-complete puzzles. *ICGA Journal*, Vol. 31, pp. 13–34.
- Longpré, L. and McKenzie, P. (2009). The complexity of Solitaire. *Theoretical Computer Science*, Vol. 410, pp. 5252–5260.
- Nagao, T., Ueda, N., Sato, T., and Watanabe, O. (1996). NP-completeness results for Nonogram via parsimonious reductions. Technical report, Tokyo Institute of Technology.
- Pritchard, D. B. and Beasley, J. D. (2007). *The classified encyclopedia of Chess variants*. Beasley, Harpenden.
- Rijn, J. N. van (2012). Playing games: The complexity of Klondike, Mahjong, Nonograms and Animal Chess. M.Sc. thesis, Leiden University.
- Rijn, J. N. van and Vis, J. K. (2013). Complexity and retrograde analysis of the game Dou Shou Qi. *Proceedings of the 25th Benelux Conference on Artificial Intelligence*, pp. 239–246.
- Yan, X., Diaconis, P., Rusmevichientong, P., and Van Roy, B. (2004). Solitaire: Man versus machine. *Proceedings of Advances in Neural Information Processing Systems 17 (NIPS 2004)*, pp. 1553–1560.