DIMACS Technical Report 2009-09
March 2009

On Acyclicity of Games with Cycles[1]

by

Daniel Andersson
Department of Computer Science
Aarhus University
koda@cs.au.dk

Vladimir Gurvich
RUTCOR
Rutgers University
gurvich@rutcor.rutgers.edu

Thomas Dueholm Hansen
Department of Computer Science
Aarhus University
tdh@cs.au.dk

[1]This work was supported by the Center for Algorithmic Game Theory at Aarhus University, funded by the Carlsberg Foundation. The second author was partially supported also by DIMACS and by Graduate School of Information Science and Technology, University of Tokyo.

ABSTRACT

We study restricted improvement cycles (ri-cycles) in finite positional n-person games with perfect information modeled by directed graphs (digraphs) that may contain cycles. We obtain criteria of restricted improvement acyclicity (ri-acyclicity) in two cases: for $n = 2$ and for acyclic digraphs. We provide several examples that outline the limits of these criteria and show that, essentially, there are no other ri-acyclic cases. We also discuss connections between ri-acyclicity and some open problems related to Nash-solvability.

1 Main Concepts and Results

1.1 Games in Normal Form

1.1.1 Game Forms, Strategies, and Utility Functions.

Given a set of players $I = \{1, \ldots, n\}$ and a set of strategies X_i for each $i \in I$, let $X = \prod_{i \in I} X_i$.

A vector $x = (x_i, i \in I) \in X$ is called a *strategy profile* or *situation*.

Furthermore, let A be a set of outcomes. A mapping $g : X \to A$ is called a *game form*. In this paper, we restrict ourselves to *finite* game forms, that is, we assume that sets I, A and X are finite.

Then, let $u : I \times A \to \mathbb{R}$ be a utility function. Standardly, the value $u(i, a)$ (or $u_i(a)$) is interpreted as the payoff to player $i \in I$ in case of the outcome $a \in A$. In figures, the notation $a <_i b$ means $u_i(a) < u_i(b)$.

Sometimes, it is convenient to exclude ties. Accordingly, u is called a *preference profile* if the mapping u_i is injective for each $i \in I$; in other words, u_i defines a complete order over A describing the preferences of player $i \in I$.

A pair (g, u) is called a *game in normal form*.

1.1.2 Improvement Cycles and Acyclicity.

In a game (g, u), an *improvement cycle (im-cycle)* is defined as a sequence of k strategy profiles $\{x^1, \ldots, x^k\} \subseteq X$ such that x^j and x^{j+1} coincide in all coordinates but one $i = i(j)$ and, moreover, $u_i(x^{j+1}) > u_i(x^j)$, that is, player i makes profit by substituting strategy x_i^{j+1} for x_i^j; this holds for all $j = 1, \ldots, k$ and, standardly, we assume that $k + 1 = 1$.

A game (g, u) is called *im-acyclic* if it has no im-cycles. A game form g is called *im-acyclic* if for each u the corresponding game (g, u) is im-acyclic.

We call x^{j+1} an improvement with respect to x^j for player $i = i(j)$. We call it a *best reply* (BR) improvement if player i can get no strictly better result provided all other players keep their strategies. Correspondingly, we introduce the concepts of a BR im-cycle and BR im-acyclicity. Obviously, im-acyclicity implies BR im-acyclicity but not vice versa.

1.1.3 Nash Equilibria and Acyclicity.

Given a game (g, u), a strategy profile $x \in X$ is called a *Nash equilibrium* (NE) if $u_i(x) \geq u_i(x')$ for each $i \in I$, whenever $x'_j = x_j$ for all $j \in I \setminus \{i\}$. In other words, x is a NE if no player can get a strictly better result by substituting a new strategy (x'_i for x_i) when all other players keep their old strategies. Conversely, if x is not a NE then there is a player who can improve his strategy. In particular, he can choose a best reply. Hence, a NE-free game (g, u) has a BR im-cycle. Let us remark that the last statement holds only for finite games, while the converse statement is not true at all.

A game (g, u) is called *Nash-solvable* if it has a NE. A game form g is called *Nash-solvable* if for each u the corresponding game (g, u) has a NE.

1.2 Positional Games with Perfect Information

1.2.1 Games in Positional Form.

Let $G = (V, E)$ be a finite directed graph (digraph) whose vertices $v \in V$ and directed edges $e \in E$ are called *positions* and *moves*, respectively. An edge $e = (v', v'')$ is a move from position v' to v''. Let $out(v)$ and $in(v)$ denote the sets of moves from and to v, respectively.

A position $v \in V$ is called *terminal* if $out(v) = \emptyset$, i.e., there are no moves from v. Let V_T denote the set of all terminals. Let us also fix a starting position $v_0 \in V \setminus V_T$. A directed path from v_0 to a terminal position is called a *finite play*.

Furthermore, let $D : V \setminus V_T \to I$ be a decision mapping, with $I = \{1, \ldots, n\}$ being the set of players. We say that the player $i = D(v) \in I$ makes a decision (move) in a position $v \in D^{-1}(i) = V_i$. Equivalently, D is defined by a partition $D : V = V_1 \cup \ldots \cup V_n \cup V_T$. In this paper we do not consider random moves.

A triplet $\mathcal{G} = (G, D, v_0)$ is called a *positional game form*.

1.2.2 Cycles, Outcomes, and Utility Functions.

Let C denote the set of simple (that is, not self-intersecting) directed cycles in G.

The set of outcomes A can be defined in two ways:

(i) $A = V_T \cup C$, that is, each terminal and each directed cycle is a separate outcome.

(ii) $A = V_T \cup \{C\}$, that is, each terminal is an outcome and all directed cycles (infinite plays) constitute one special outcome $c = \{C\}$.

Case (i) was considered in [2] for two-person ($n = 2$) game forms. In this paper, we analyze case (ii) for n-person games.

Remark 1. *Let us mention that as early as in 1912, Zermelo already considered case (ii) for the zero-sum two-person games in his pioneering work [11], where the game of Chess was chosen as a basic example. Obviously, the corresponding graph contains cycles: one appears whenever a position is repeated in a play. By definition, any infinite play is a draw. More precisely, Chess results in a draw whenever the same position appears three times in a play. Yet, this difference does not matter, since we are going to restrict ourselves to positional (stationary) strategies; see Remark 2 below.*

Note that players can rank outcome c arbitrarily in their preferences. In contrast, in [1] it was assumed that infinite play $c \in A$ is the worst outcome for all players.

1.2.3 Positional Games in Normal Form.

A triplet $\mathcal{G} = (G, D, v_0)$ and quadruple $(G, D, v_0, u) = (\mathcal{G}, u)$ are called a *positional form* and a *positional game*, respectively. Positional games can also be represented in *normal form*, as described below. A mapping $x : V \setminus V_T \to E$ that assigns to every non-terminal position v a move $e \in out(v)$ from this position is a *situation* or *strategy profile*.

A *strategy* of player $i \in I$ is the restriction $x_i : V_i \to E$ of x to $V_i = D^{-1}(i)$.

Remark 2. *A strategy x_i of a player $i \in I$ is interpreted as a decision plan for every position $v \in V_i$. Note that, by definition, the decision in v can depend only on v itself but not on the preceding positions and moves. In other words, we restrict the players to their pure positional strategies.*

Each strategy profile $x \in X$ uniquely defines a play $p = p(x)$ that starts in v_0 and then follows the moves prescribed by x. This play either ends in a terminal of V_T (p is finite) or results in a cycle, $a(x) = c$ (p is infinite). Thus, we obtain a game form $g(\mathcal{G}) : X \to A$ that is called the *normal form* of $\mathcal{G}$. This game form is standardly represented by an n-dimensional table whose entries are outcomes of $A = V_T \cup \{c\}$; see Figures 1, 6 and 8.

The pair $(g(\mathcal{G}), u)$ is called the *normal form* of a positional game $(\mathcal{G}, u)$.

1.3 On Nash-Solvability of Positional Game Forms

In [2], Nash-solvability of positional game forms was considered for case (i); each directed cycle is a separate outcome. An explicit characterization of Nash-solvability was obtained for the two-person ($n = 2$) game forms whose digraphs are *bidirected*: $(v', v'') \in E$ if and only if $(v'', v') \in E$. In [1], Nash-solvability of positional game forms was studied (for a more general class of payoff functions, so-called *additive or integral* payoffs, yet) with the following additional restriction:

(ii') The outcome c, infinite play, is ranked as the worst one by all players.

Under assumption (ii'), Nash-solvability was proven in three cases:

(a) Two-person games ($n = |I| = 2$);

(b) Games with at most three outcomes ($|A| \leq 3$);

(c) Play-once games, in which each player controls only one position ($|V_i| = 1 \; \forall \, i \in I$).

It was also conjectured in [1] that Nash-solvability holds in general.

Conjecture 1. *([1]) A positional game is Nash-solvable whenever (ii') holds.*

This Conjecture would be implied by the following statement: *every im-cycle $\mathcal{X} = \{x^1, \ldots, x^k\} \subseteq X$ contains a strategy profile x^j such that the corresponding play $p(x^j)$ is infinite*. Indeed, Conjecture 1 would follow, since outcome $c \in A$ being the worst for all players, belongs to no im-cycle. However, the example of Section 2.2 will show that such an approach fails. Nevertheless, Conjecture 1 is not disproved. Moreover, a stronger conjecture was recently suggested by Gimbert and Sørensen, [4]. They assumed that, in case of terminal payoffs, condition (ii') is not needed.

Conjecture 2. *A positional game is Nash-solvable if all cycles form a single outcome.*

They gave a simple and elegant proof for the two-person case. With their permission, we reproduce it in Section 5.

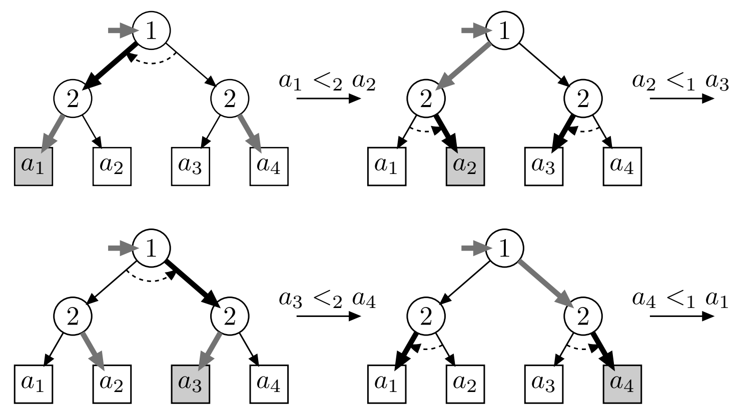

1 : $a_2 < a_3,\ a_4 < a_1$

2 : $a_1 < a_2,\ a_3 < a_4$

$$\begin{array}{cc} & 2 \\ 1 & \begin{bmatrix} a_1 & a_1 & \rightarrow & a_2 & a_2 \\ & \uparrow & & \downarrow & \\ a_3 & a_4 & \leftarrow & a_3 & a_4 \end{bmatrix} \end{array}$$

Figure 1: Im-cycle in a tree

1.4 Restricted Improvement Cycles and Acyclicity

1.4.1 Improvement Cycles in Trees.

Kukushkin [8, 9] was the first to consider im-cycles in positional games. He restricted himself to trees and observed that even in this case im-cycles can exist; see example in Figure 1.

However, it is easy to see that unnecessary changes of strategies take place in this im-cycle. For example, let us consider transition from $x^1 = (x_1^1, x_2^2)$ to $x^2 = (x_1^1, x_2^3)$. Player 1 keeps his strategy x_1^1, while 2 substitutes x_2^3 for x_2^2 and gets a profit, since $g(x_1^1, x_2^2) = a_1$, $g(x_1^1, x_2^3) = a_2$, and $u_2(a_1) < u_2(a_2)$.

Yet, player 2 switches simultaneously from a_4 to a_3. Obviously, this cannot serve any practical purpose, since the strategy is changed outside the actual play.

In [8], Kukushkin also introduced the concept of *restricted improvements* (ri). In particular, he proved that positional games on trees become ri-acyclic if players are not allowed to change their decisions outside the actual play. For completeness, we sketch his simple and elegant proof in Section 3.1, where we also mention some related results and problems.

Since we consider arbitrary finite digraphs (not only trees), let us define accurately several types of restrictions for this more general case. This restriction (introduced by Kukushkin) will be called the *inside play restriction*.

1.4.2 Inside Play Restriction.

Given a positional game form $\mathcal{G} = (G, D, v_0)$ and strategy profile $x^0 = (x_i^0,\, i \in I) \in X$, let us consider the corresponding play $p_0 = p(x^0)$ and outcome $a_0 = a(x^0) \in A$. This outcome is either a terminal, $a_0 \in V_T$, or a cycle, $a_0 = c$.

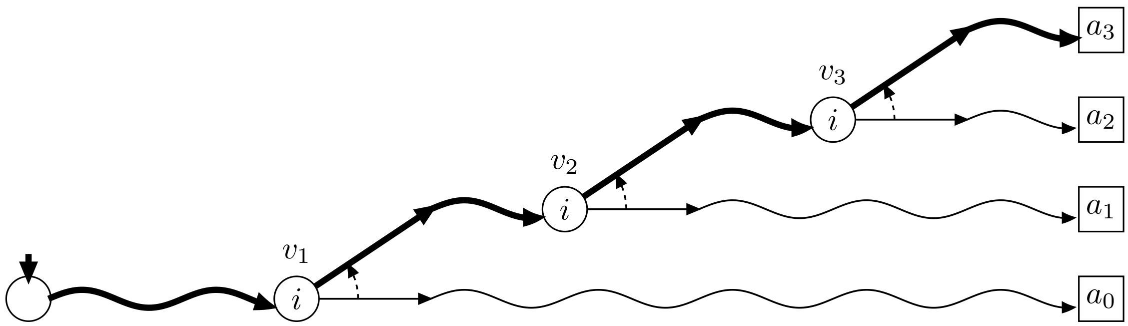

Figure 2: Inside play restriction.

Let us consider the strategy x_i^0 of a player $i \in I$. He is allowed to change his decision in any position v_1 from p_0. This change will result in a new strategy profile x^1, play $p_1 = p(x^1)$, and outcome $a_1 = a(x^1) \in A$.

Then, player i may proceed, changing his strategy further. Now, he is only allowed to change the decision in any position v_2 that is located *after* v_1 in p_1, etc., until a position v_m, strategy profile x^m, play $p_m = p(x^m)$, and outcome $a_m = a(x^m) \in A$ appear; see Figure 2, where $m = 3$.

Equivalently, we can say that all positions $v_1, \ldots, v_m$ belong to one play.

Let us note that, by construction, obtained plays $\{p_0, p_1, \ldots, p_m\}$ are pairwise distinct. In contrast, the corresponding outcomes $\{a_0, a_1, \ldots, a_m\}$ can coincide and some of them might be the infinite play outcome $c \in A$.

Whenever the acting player i substitutes the strategy x_i^m, defined above, for the original strategy x_i^0, we say that this is an *inside play deviation*, or in other words, that this change of decision in x satisfies the *inside play restriction.*

It is easy, but important, to notice that this restriction, in fact, does not limit the power of a player. More precisely, if a player i can reach an outcome a_m from x by a deviation then i can also reach a_m by an inside play deviation.

From now on, we will consider only such *inside play restricted* deviations and, in particular, only restricted improvements (ri) and talk about ri-cycles and ri-acyclicity rather than im-cycles and im-acyclicity, respectively.

1.4.3 Types of Improvements.

We define the following four types of improvements:

Standard improvement (or just improvement): $u_i(a_m) > u_i(a_0)$;

Strong improvement: $u_i(a_m) > u_i(a_j)$ for $j = 0, 1, \ldots, m-1$;

Last step improvement: $u_i(a_m) > u_i(a_{m-1})$;

Best reply (BR) improvement: a_m is the best outcome that player i can reach from x (as we already noticed above, the inside play restriction does not restrict the set of reachable outcomes).

Obviously, each best reply or strong improvement is a standard improvement. Furthermore, it is easy to verify that no other containments hold between the above four classes. We will consider ri-cycles and ri-acyclicity specifying in each case a type of improvement from the above list.

Let us note that any type of ri-acyclicity still implies Nash-solvability. Indeed, if a positional game has no NE then for every strategy profile $x \in X$ there is a player $i \in I$ who can improve x to some other strategy profile $x' \in X$. In particular, i can always choose a strong BR restricted improvement. Since we consider only finite games, such an iterative procedure will result in a strong BR ri-cycle. Equivalently, if we assume that there is no such cycle then the considered game is Nash-solvable; in other words, already strong BR ri-acyclicity implies Nash-solvability.

1.5 Sufficient Conditions for Ri-acyclicity

We start with Kukushkin's result for trees.

Theorem 1. *([8]). Positional games on trees have no restricted standard improvement cycles.*

After trees, it seems natural to consider acyclic digraphs, so we asked Kukushkin whether he has a generalization of Theorem 1 for this case. He said that he did not consider it yet. However, very soon he came up with a result that can be modified as follows.

Theorem 2. *([10]). Positional games on acyclic digraphs have no restricted last step improvement cycles.*

Let us remark that Theorem 1 does not immediately follow from Theorem 2, since a standard improvement might be not a last step improvement. Finally, for two-person positional games, which might have directed cycles, the following statement holds.

Theorem 3. *Two-person positional games have no restricted strong improvement cycles.*

Obviously, Theorem 3 implies Nash-solvability of two-person positional games; see also Section 5 for an independent proof due to Gimbert and Sørensen.

The proofs of Theorems 1, 2 and 3 are given in Sections 3.1, 3.2 and 3.3, respectively.

2 Examples of Ri-cycles; Limits of Theorems 1, 2 and 3

In this paper, we emphasize negative results showing that it is unlikely to strengthen one of the above theorems or obtain other criteria of ri-acyclicity.

2.1 Example Limiting Theorems 2 and 3

For both Theorems 2 and 3, the specified type of improvement is essential. Indeed, the example in Figure 3 shows that a two-person game on an acyclic digraph can have a ri-cycle. However, it is not difficult to see that in this ri-cycle, not all improvements are strong and some are not even last step improvements.

Thus, all conditions of Theorems 2 and 3 are essential.

Furthermore, we note that if in Theorem 3 we substitute BR improvement for strong improvement, the modified statement will not hold, see Figure 4.

By definition, every change of strategy must result in an improvement for the corresponding player. Hence, each such change implies an ordering of the two outcomes; in our figures, it appears as a label on the transition arrow between situations. An entire im-cycle implies a set of inequalities, which must be satisfiable in order to allow a consistent preference profile. Note that it is also sufficient to allow ties and have a partial ordering of the outcomes.

2.2 On c-free Ri-cycles

In Section 1.3, we demonstrated that Conjecture 1 on Nash-solvability would follow from the following statement:

(i) *There are no c-free im-cycles.*

Of course, (i) fails. As we know, im-cycles exist already in trees; see Figure 1. However, let us substitute (i) by the similar but much weaker statement:

(ii) *Every restricted strong BR ri-cycle contains a strategy profile whose outcome is infinite play.*

One can derive Conjecture 1 from (ii), as well as from (i). Yet, (ii) also fails. Indeed, let us consider the ri-cycle in Figure 5. This game is play-once; each player controls only one position. Moreover, there are only two possible moves in each position. For this reason, every ri-cycle in this game is BR and strong.

There are seven players ($n = 7$) in this example, yet, by teaming up players in coalitions we can reduce the number of players to four while the improvements remain BR and strong. Indeed, this can be done by forming three coalitions $\{1, 7\}, \{3, 5\}$, $\{4, 6\}$ and merging the preferences of the coalitionists. The required extra constraints on the preferences of the coalitions are also shown in Figure 5.

It is easy to see that inconsistent (i.e., cyclic) preferences appear whenever any three players form a coalition. Hence, the number of coalitions cannot be reduced below 4, and it is, in fact, not possible to form 4 coalitions in any other way while keeping improvements BR and strong.

Obviously, for the two-person case, **(ii)** follows from Theorem 3.

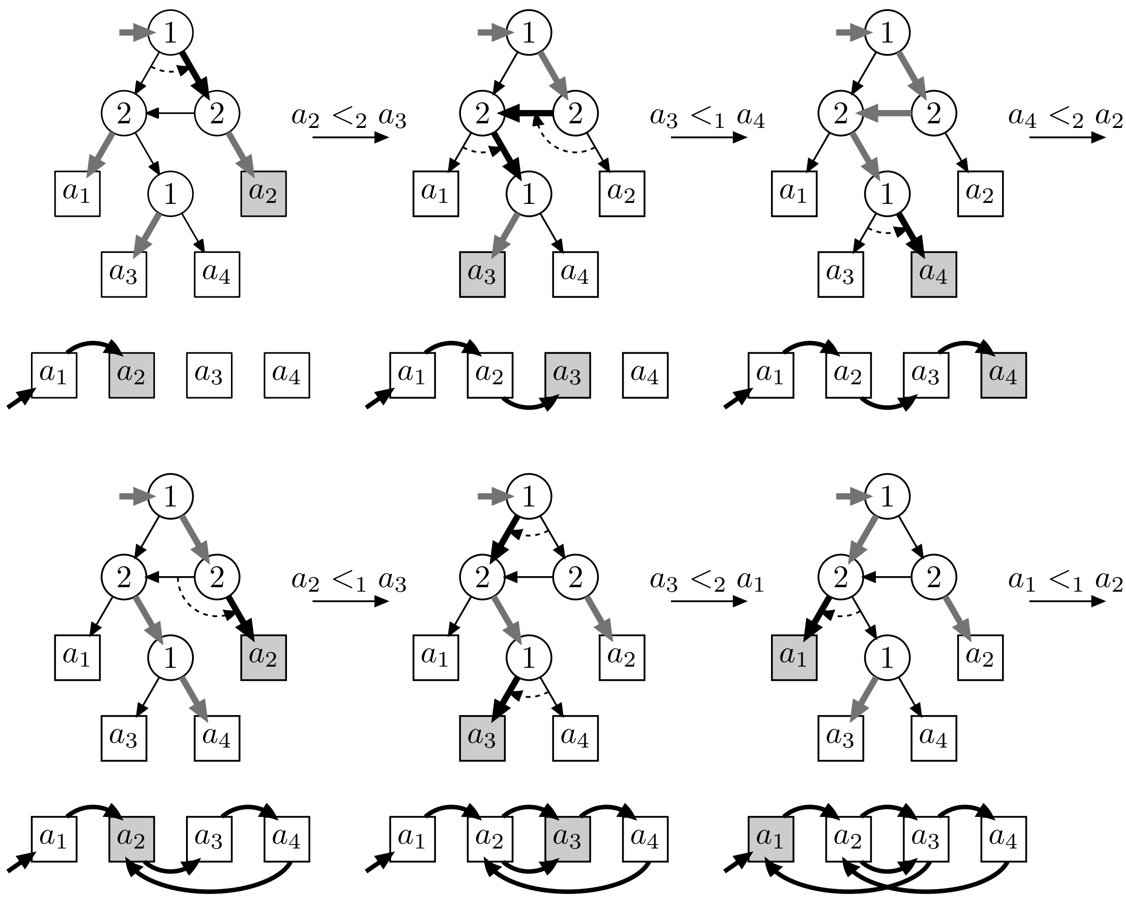

$$1: \; a_1 < a_2 < a_3 < a_4$$

$$2: \; a_4 < a_2 < a_3 < a_1$$

Figure 3: 2-person ri-cycle in an acyclic digraph. Beneath each situation is a graph of outcomes with edges defined by the previous improvement steps; these will be of illustrative importance in the proof of Theorem 3.

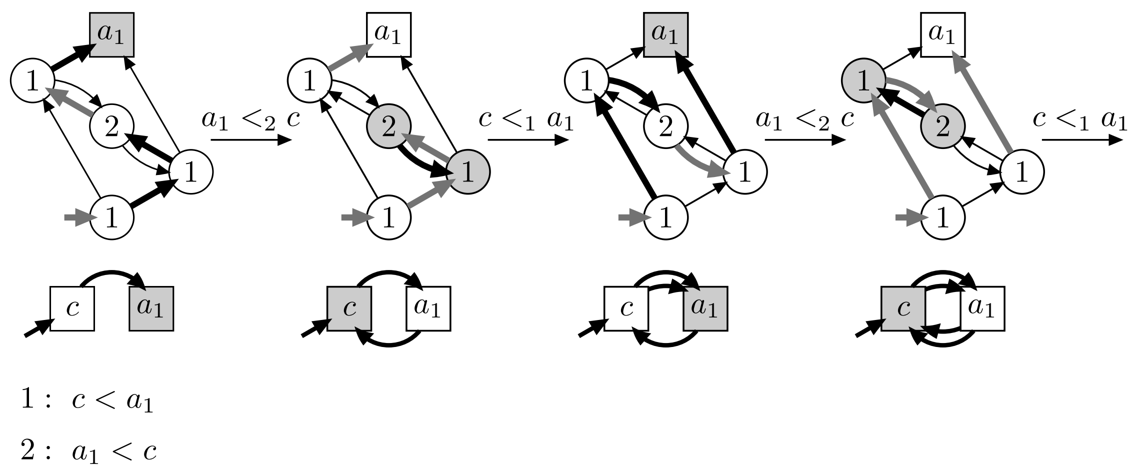

Figure 4: 2-person BR ri-cycle in graph with cycles.

Remark 3. *We should confess that our original motivation fails. It is hardly possible to derive new results on Nash-solvability from ri-acyclicity. Although, ri-acyclicity is much weaker than im-acyclicity, it is still too much stronger than Nash-solvability. In general, by Theorems 3 and 2, ri-acyclicity holds for $n = 2$ and for acyclic digraphs. Yet, for these two cases Nash-solvability is known.*

It is still possible that (ii) (and, hence, Conjecture 1) holds for $n = 3$, too. Strong BR ri-cycles in which the infinite play outcome c occurs do exist for $n = 3$, however. Such an example is provided in Section 2.3.4.

However, ri-acyclicity is of independent (of Nash-solvability) interest. In this paper, we study ri-acyclicity for the case when each terminal is a separate outcome, while all directed cycles form one special outcome. For the alternative case, when each terminal and each directed cycle is a separate outcome, Nash-solvability was considered in [2], while ri-acyclicity was never studied.

2.3 Flower Games: Ri-cycles and Nash-Solvability

2.3.1 Flower Positional Game Forms.

A positional game form $\mathcal{G} = (G, D, v_0)$ will be called a *flower* if there is a (chordless) simple cycle C in G that contains all positions, except the initial one, v_0, and the terminals, V_T; furthermore, we assume that there are only moves from v_0 to C and from C to V_T; see examples in Figures 6, 7 and 9.

By definition, C is the unique cycle in G. Nevertheless, it is enough to make flower games very different from acyclic games; see [1] (where flower games are referred to as St. George games). Here we consider several examples of ri-cycles in flower game forms of 3 and 4 players; see Figures 6, 7 and 9. Let us note that the game forms of Figures 6 and 7 are play-once: each player is in control of one position, that is, $n = |I| = |V \setminus V_T| = 3$ or 4,

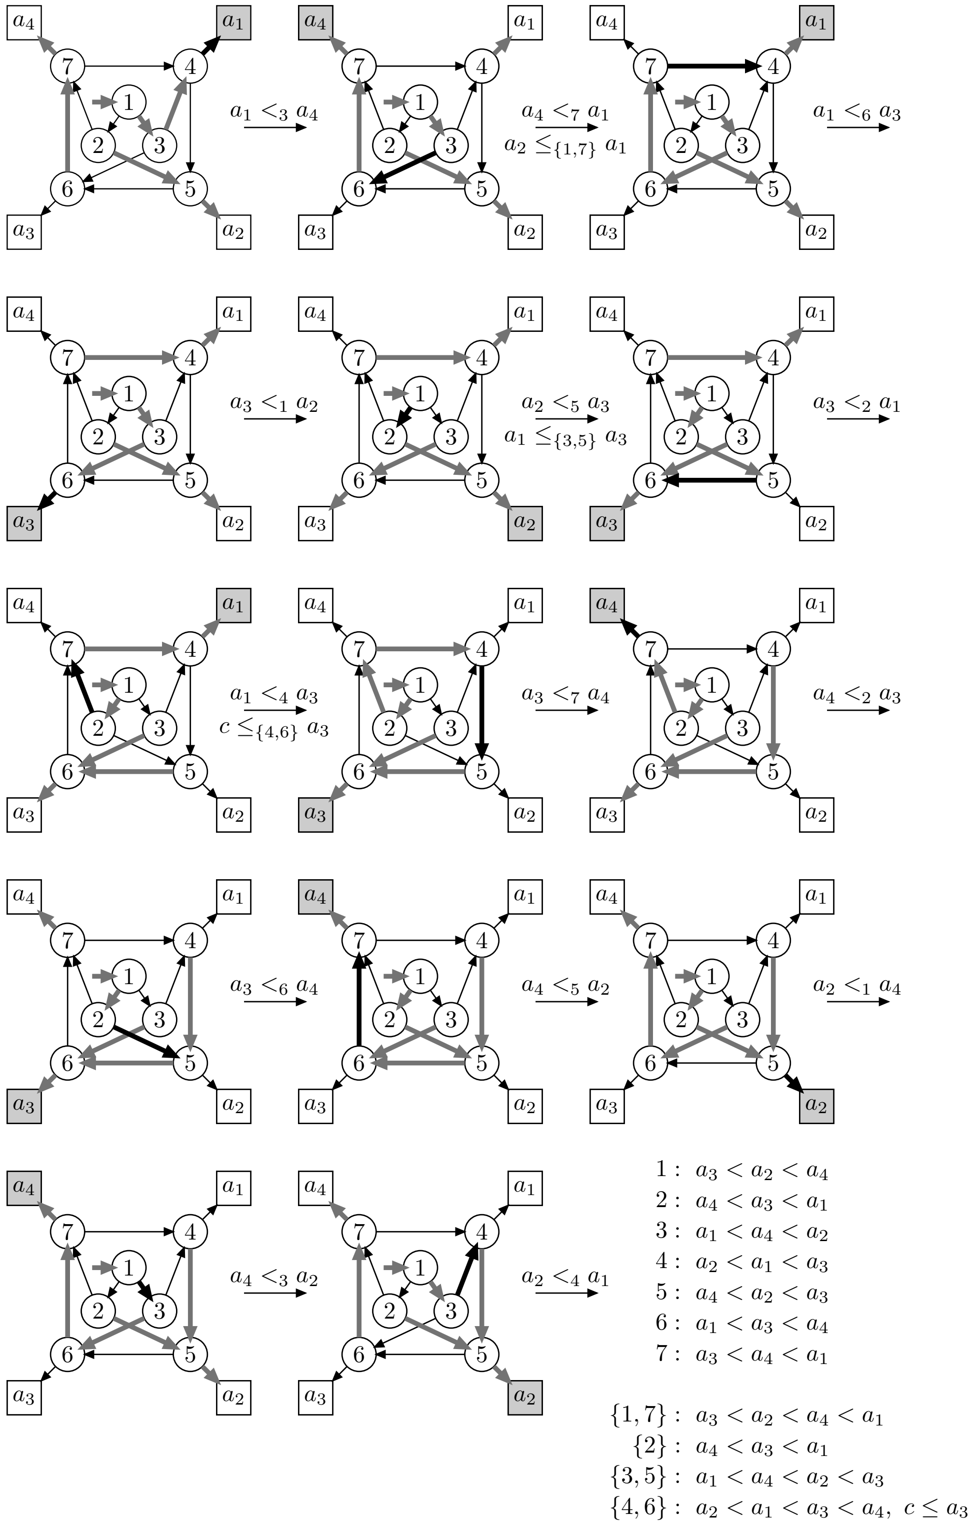

Figure 5: c-free strong BR ri-cycle.

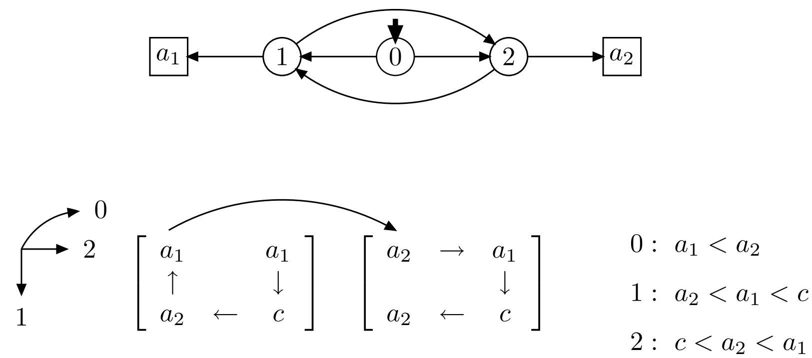

Figure 6: Ri-acyclic flower game form with 3 players.

respectively. In fact, Figure 9 can also be turned into a six-person play-once flower game.

2.3.2 Flower Three-Person Game Form.

Positional and normal forms of a three-person flower game are given in Figure 6. This game form is ri-acyclic. Indeed, it is not difficult to verify that an im-cycle in it would result in inconsistent preferences for one of the players. Yet, there is a sequence of 7 restricted improvement steps, that is, a "Hamiltonian improvement path" in the normal form.

2.3.3 Flower Four-Person Game Form.

Positional and normal forms of a four-person flower game are given in Figures 7 and 8, respectively, where a ri-cycle is shown. Obviously, it is a strong and BR ri-cycle, since there are only two possible moves in every position. However, it contains c.

The number of players can be reduced by forming either the coalition $\{1, 2\}$ or the coalition $\{1, 3\}$. However, in the first case the obtained ri-cycle is not BR, though it is strong, whereas a non-restricted improvement appears in the second case.

Moreover, no c-free ri-cycle can exist in this four-person flower game form. To see this, let us consider the graph of its normal form shown in Figure 8. It is not difficult to verify that, up to isomorphism, there is only one ri-cycle, shown above. All other ri-cycle candidates imply inconsistent preferences; see the second graph in Figure 8.

2.3.4 On BR Ri-cycles in Three-Person Flower Games.

In Section 2.2 we gave an example of a c-free strong BR ri-cycle in a four-person game. Yet, the existence of such ri-cycles in three-person games remains open. However, a strong BR ri-cycle that contains c can exist already in a three-person flower game; see Figure 9.

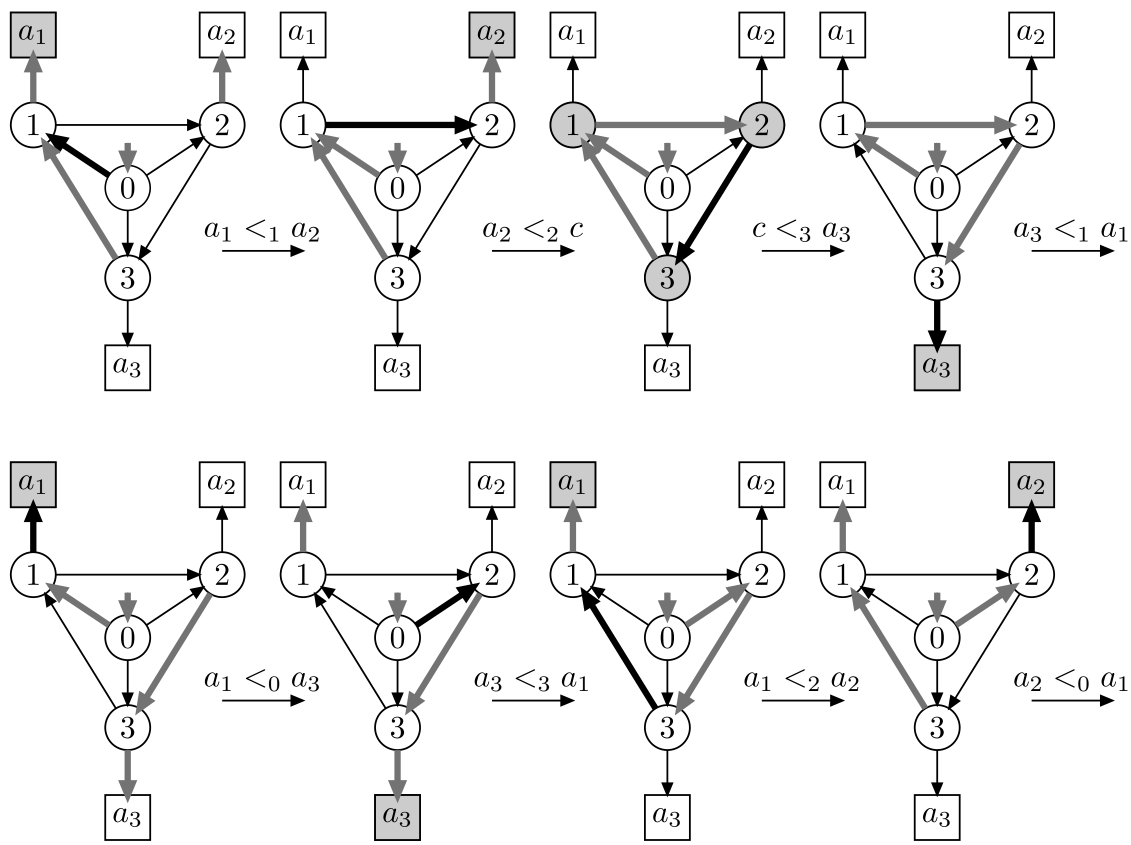

$0: \; a_2 < a_1 < a_3$
$1: \; a_3 < a_1 < a_2$
$2: \; a_1 < a_2 < c$
$3: \; c < a_3 < a_1$

Figure 7: Positional form of a ri-cycle in a flower game form with 4 players.

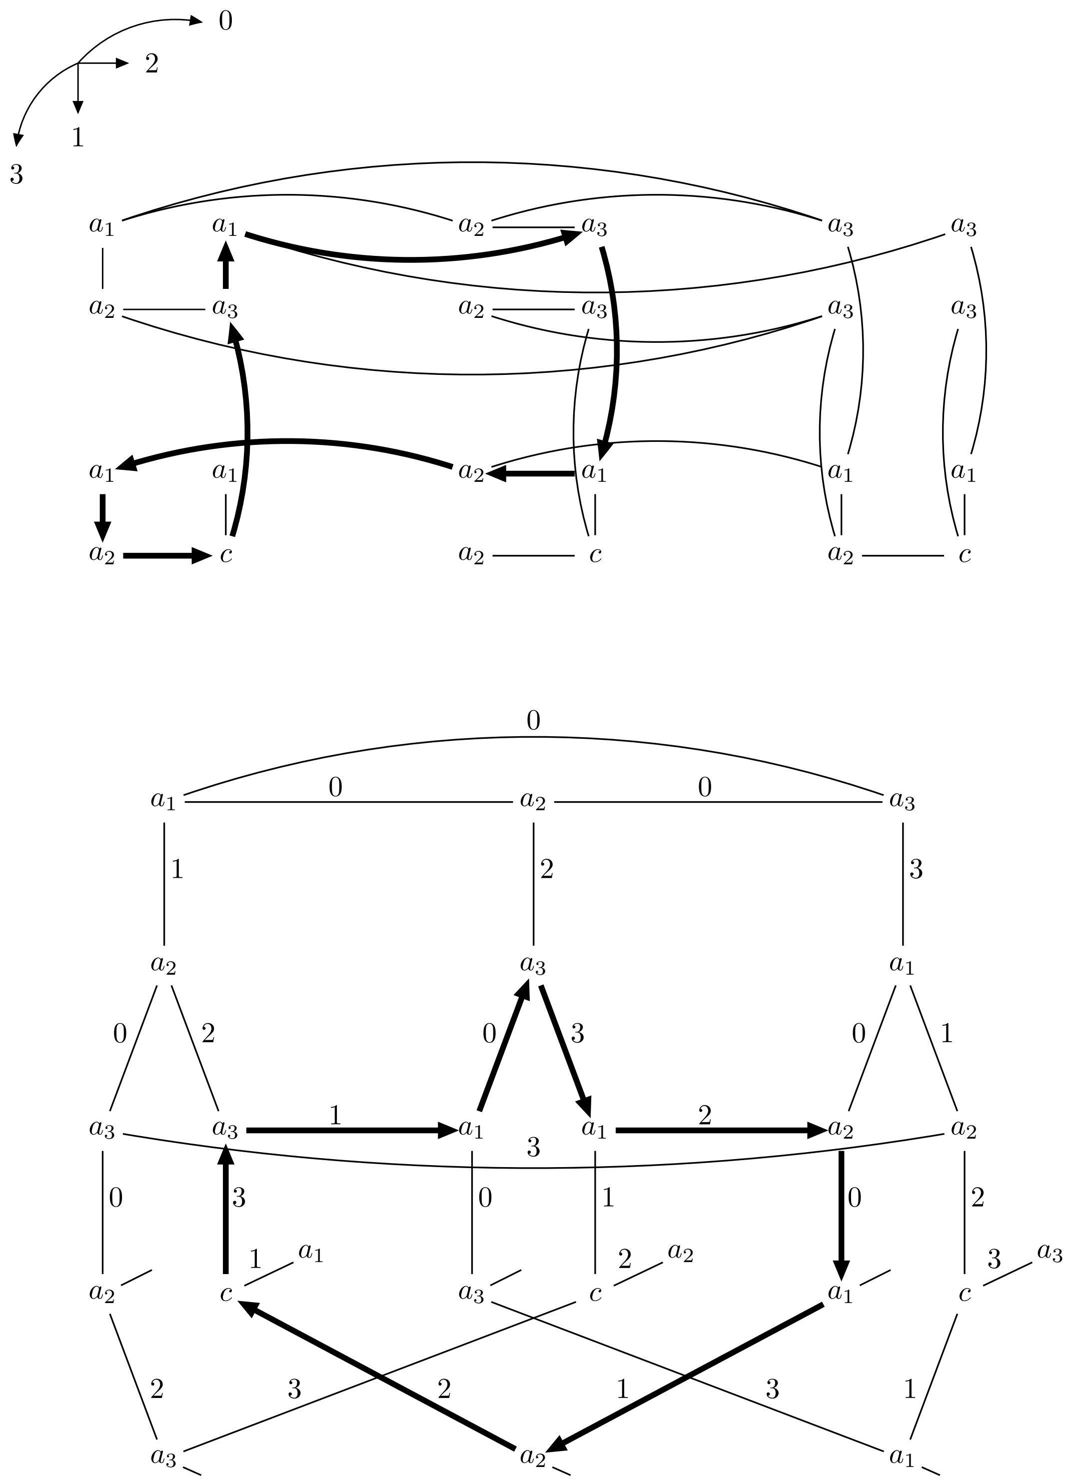

Figure 8: Normal form and unfolded normal form of a ri-cycle in a flower game form with 4 players.

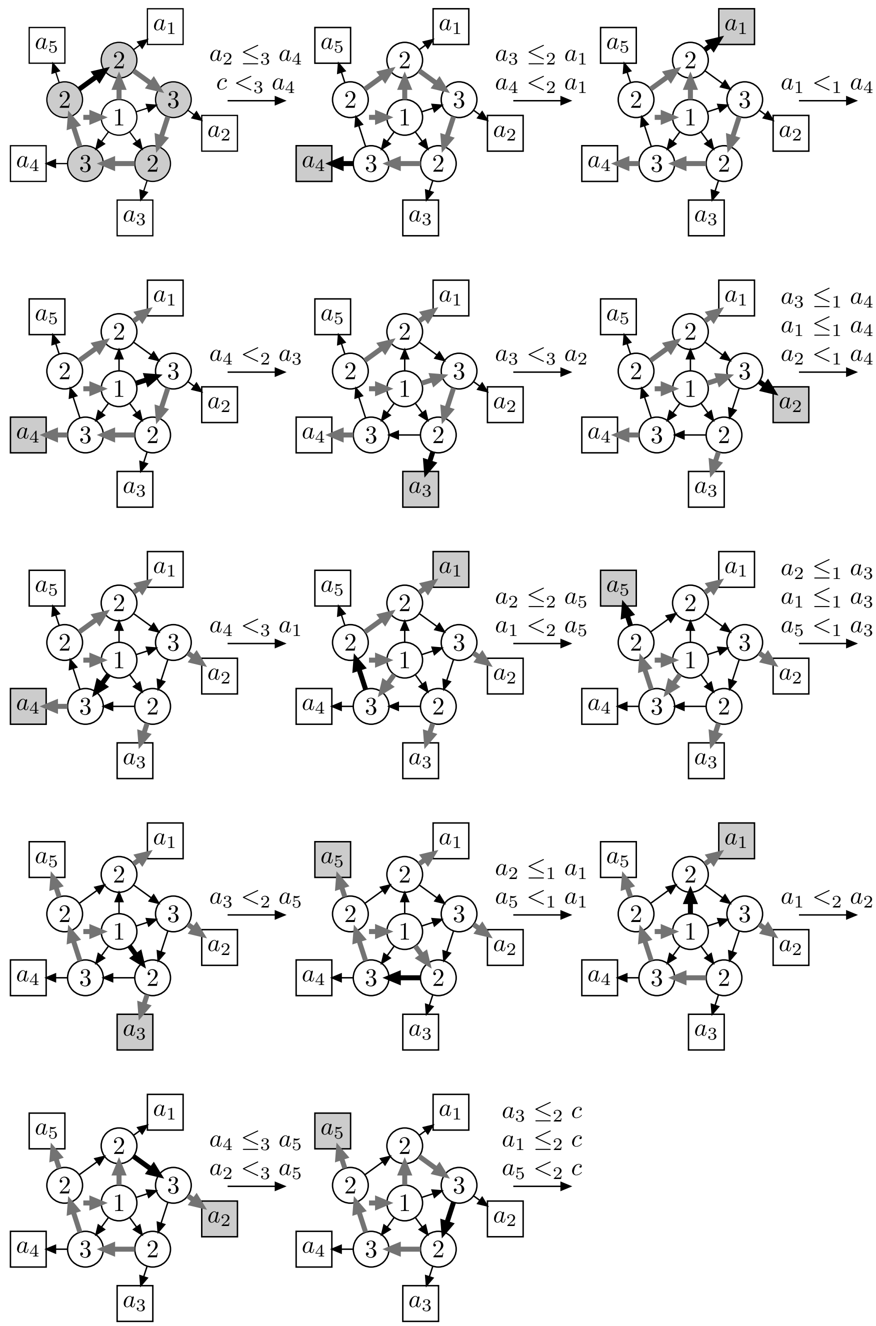

Feasible total order:

$1: \; a_5 < a_2 < a_1 < a_3 < a_4 < c$

$2: \; a_4 < a_3 < a_1 < a_2 < a_5 < c$

$3: \; c < a_3 < a_2 < a_4 < a_1 < a_5$

Figure 9: Strong BR ri-cycle in a 3-person flower game.

2.3.5 Nash-Solvability of Flower Games.

In this section we assume without loss of generality that v_0 is controlled by player 1 and that every position v in C has exactly two moves: one along C and the other to a terminal $a = a_v \in V_T$. Indeed, if a player has several terminal moves from one position then, obviously, all but one, which leads to the best terminal, can be eliminated.

We will call positions in C *gates* and, given a strategy profile x, we call gate $v \in C$ *open* (*closed*) if move (v, a_v) is chosen (not chosen) by x.

First, let us consider the simple case when player 1 controls only v_0 and later we will reduce Nash-solvability of general flower games to this special case.

Lemma 1. *Flower games in which player* 1 *controls only* v_0 *are Nash-solvable.*

Proof. Let us assume that there is a move from v_0 to each position of C. In general, the proof will remain almost the same, except for a few extra cases with a similar analysis.

Now, either: (i) for each position $v \in C$ the corresponding player $i = D(v)$ prefers c to $a = a_v$, or (ii) there is a $v' \in C$ such that $i' = D(v')$ prefers $a' = a_{v'}$ to c. If a player controls several such positions then let a' be his best outcome.

In case (i), a strategy profile such that all gates are closed is a NE. In case (ii), the strategy profile where player 1 moves from v_0 to v', the gate v' is open, and all other gates are closed, is a NE.

□

Theorem 4. *Flower games are Nash-solvable.*

Proof. We will give an indirect proof, deriving a contradiction with Lemma 1. Let $(\mathcal{G}, u)$ be a NE-free flower game. Moreover, let us assume that it is *minimal*, that is, a NE appears whenever we delete any move from $\mathcal{G}$. This assumption implies that for each gate v, there is a strategy profile x^1 such that this gate is closed but it is opened by a BR restricted improvement x^2. Since the game is NE-free, there is an infinite sequence $\mathcal{X} = \{x^1, x^2, \ldots\}$ of such BR restricted improvements. Then, it follows from Theorem 2 that gate v will be closed again by a profile $x^k \in \mathcal{X}$. Indeed, if we delete the closing edge (i.e., v remains open), the resulting graph is acyclic.

Now, let us assume that gate v is controlled by player 1. Let v' be the closest predecessor of v in C such that there is a move from v_0 to v'. Opening v, player 1 can at the same time choose the move (v_0, v').

Clearly, until v will be closed again no gate between v' and v in C, including v' itself, will be opened. Indeed, otherwise the corresponding gate could not be closed again by any sequence $\mathcal{X}$ of restricted best replies. Since player 1 already performed a BR, the next one must be performed by another player. However, these players control only the gates between v' and v in C. Hence, one of them will be opened.

Thus, we obtain a contradiction. Indeed, if a NE-free flower game has a gate of player 1, it will never be required to open. By deleting such gates repeatedly one obtains a NE-free flower game that has no gates of player 1. This contradicts Lemma 1. □

3 Proofs of Theorems 1, 2 and 3

3.1 Ri-acyclicity for Trees: Proof of Theorem 1

As we know, im-cycles can exist even for trees (see Section 1.4.1) but ri-cycles cannot. Here we sketch the proof from [8].

Given a (directed) tree $G = (V, E)$ and an n-person positional game $(\mathcal{G}, u) = (G, D, v_0, u)$, let $p_i = \sum_{v \in V_i}(|out(v)| - 1)$ for every player $i \in I = \{1, \ldots, n\}$. It is not difficult to verify that $1 + \sum_{i=1}^{n} p_i = p = |V_T|$.

Let us fix a strategy profile $x = (x_1, \ldots, x_n)$. To every move $e = (v, v')$ which is *not* chosen by x let us assign the outcome $a = a(e, x)$ which x would result in starting from v'. it is easy to see that these outcomes together with $a(x)$ form a partition of V_T.

Given a player $i \in I$, let us assign p_i numbers $u_i(a(e, x))$ for all $e = (v, v')$ not chosen by x_i, where $v \in V_i$. Let us order these numbers in monotone non-increasing order and denote the obtained p_i-dimensional vector $y_i(x)$.

Let player $i \in I$ substitute a restricted improvement x'_i for x_i; see Section 1.4.2. The new strategy profile x' results in an outcome $a_k \in A = V_T$ which is strictly better for i than the former outcome $a_0 = a(x)$. Let us consider vectors $y_j(x)$ and $y_j(x')$ for all $j \in I$. It is not difficult to verify that these two vectors are equal for each $j \in I$, except $j = i$, while $y_i(x)$ and $y_i(x')$, for the acting player i, differ by only one number: $u_i(a_k)$ in $y_i(x')$ substitutes for $u_i(a_0)$ in $y_i(x)$. The new number is strictly larger than the old one, because, by assumption of Theorem 1, x'_i is an improvement with respect to x_i for player i. Thus, vectors y_j for all $j \neq i$ remain unchanged, while y_i becomes strictly larger. Hence, no ri-cycle can appear. □

An interesting question: what is the length of the longest sequence of restricted improvements? Given $n = |I|$, $p = |A|$, and p_i such that $\sum_{i=1}^{n} p_i = p - 1 \geq n \geq 1$, the above proof of Theorem 1 implies the following upper bound: $\sum_{i=1}^{n} p_i(p - p_i)$. It would also be interesting to get an example with a high lower bound.

3.2 Last Step Ri-acyclicity for Acyclic Digraphs: Proof of Theorem 2

Given a positional game $(\mathcal{G}, u) = (G, D, v_0, u)$ whose digraph $G = (V, E)$ is acyclic, let us order positions of V so that $v < v'$ whenever there is a directed path from v to v'. To do so, let us assign to each position $v \in V$ the length of a longest path from v_0 to v and then order arbitrarily positions with equal numbers.

Given a strategy profile x, let us, for every $i \in I$, assign to each position $v \in V_i$ the outcome $a(v, x)$ which x would result in starting from v and the number $u_i(a(v, x))$. These numbers form a $|V \setminus V_T|$-dimensional vector $y(x)$ whose coordinates are assigned to positions $v \in V \setminus V_T$. Since these positions are ordered, we can introduce the inverse lexicographic order over such vectors y.

Let a player $i \in I$ choose a last step ri-deviation x'_i from x_i. Then, $y(x') > y(x)$, since the last changed coordinate increased: $u_i(a_k) > u_i(a_{k-1})$. Hence, no last step ri-cycle can exist. □

3.3 Strong Ri-acyclicity of Two-Person Games: Proof of Theorem 3

Let us consider a two-person positional game $\mathcal{G} = (G, D, v_0, u)$ and a strategy profile x such that in the resulting play $p = p(x)$ the terminal move (v, a) belongs to a player $i \in I$. Then, a strong improvement x'_i results in a terminal $a' = p(x')$ such that $u_i(a') > u_i(a)$. (This holds for n-person games, as well.)

Given a *strong* ri-cycle $\mathcal{X} = \{x^1, \ldots, x^k\} \in X$, let us assume, without any loss of generality, that the game (G, D, v_0, u) is *minimal* with respect to $\mathcal{X}$, that is, the ri-cycle $\mathcal{X}$ is broken by eliminating any move from $\mathcal{G}$. Furthermore, let $\mathcal{A}(\mathcal{X})$ denote the set of the corresponding outcomes: $A(\mathcal{X}) = \{a(x^j),\ j = 1, \ldots, k\}$. Note that several of the outcomes of the strategy profiles may be the same and that $A(\mathcal{X})$ may contain $c \in A$ (infinite play).

Let us introduce the directed multigraph $\mathcal{E} = \mathcal{E}(\mathcal{X})$ whose vertex-set is $A(\mathcal{X})$ and the directed edges are k pairs (a_j, a_{j+1}), where $a_j = a(x^j),\ \ j = 1, \ldots, k$, and $k + 1 = 1$. It is easy to see that $\mathcal{E}$ is *Eulerian*, i.e., $\mathcal{E}$ is strongly connected and for each vertex its in-degree and out-degree are equal. An example of this construction is shown in Figure 3.

Let $\mathcal{E}_1$ and $\mathcal{E}_2$ be the subgraphs of $\mathcal{E}$ induced by the edges corresponding to deviations of players 1 and 2, respectively. Then, $\mathcal{E}_1$ and $\mathcal{E}_2$ are acyclic, since a cycle would imply an inconsistent preference relation. In the example of Figure 3, the edges are partitioned accordingly (above and below the vertices), and the subgraphs are indeed acyclic.

Hence, there is a vertex a^1 whose out-degree in $\mathcal{E}_1$ and in-degree in $\mathcal{E}_2$ both equal 0. In fact, an outcome $a^1 \in A(\mathcal{X})$ most preferred by player 1 must have this property. (We do not exclude ties in preferences; if there are several best outcomes of player 1 then a^1 can be any of them.) Similarly, we define a vertex a^2 whose in-degree in $\mathcal{E}_1$ and out-degree in $\mathcal{E}_2$ both equal 0.

Let us remark that either a^1 or a^2 might be equal to c, yet, not both. Thus, without loss of generality, we can assume that a^1 is a terminal outcome.

Obviously, a player, 1 or 2, has a move to a^1. If 1 has such a move then it cannot be improved in $\mathcal{X}$, since $u_1(a_j) \leq u_1(a^1)$ for all $j = 1, \ldots, k$ and $\mathcal{X}$ is a strong ri-cycle. Let us also recall that a^1 has no incoming edges in $\mathcal{E}_2$. Hence, in $\mathcal{X}$, player 2 never makes an improvement that results in a^1. In other words, a player who has a move to a^1 will make it either always, player 1, or never, player 2. In both cases we obtain a contradiction with the minimality of digraph $\mathcal{G}$. □

4 Laziness Restriction

In addition to the inside play restriction, let us consider the following closely related but stronger restriction.

Let player i substitute strategy x'_i for x_i to get a new outcome $a' = a(x')$ instead of $a = a(x)$. We call such a deviation *lazy*, or say that it satisfies the *laziness restriction*, if it minimizes the number of positions in which player i changes the decision to reach a'.

Let us note that the corresponding strategy x'_i might not be unique.

Obviously, each lazy deviation satisfies the inside play restriction.

Furthermore, if a lazy deviation is an improvement, $u_i(a) < u_i(a')$, then this improvement is strong.

Proposition 1. *Given a strategy profile x, a target outcome $a' \in A$, and a player $i \in I$, the problem of finding a lazy deviation from x_i to x'_i such that $a(x') = a'$ (and x' is obtained from x by substituting x'_i for x_i) reduces to the shortest directed path problem.*

Proof. Let us assign a length $d(e)$ to each directed edge $e \in E$ as follows: $d(e) = 0$ if move e is prescribed by x, $d(e) = 1$ for every other possible move of the acting player i, and $d(e) = \infty$ for all other edges. Then let us consider two cases: (i) $a' \in V_T$ is a terminal and (ii) $a' = c$.

In case (i), a shortest directed path from v_0 to a' defines a desired x'_i, and vice versa. Case (ii), $a' = c$, is a little more complicated.

First, for every directed edge $e = (v, v') \in E$, let us find a shortest directed cycle C_e that contains e and its length d_e. This problem is easily reducible to the shortest directed path problem, too. The following reduction works for an arbitrary weighted digraph $G = (V, E)$. Given a directed edge $e = (v, v') \in E$, let us find a shortest directed path from v' to v. In case of non-negative weights, this can be done by Dijkstra's algorithm.

Then, it is also easy to find a shortest directed cycle C_v through a given vertex $v \in V$ and its length d_v; obviously, $d_v = \min_{v' \in V}(d_e \mid e = (v, v'))$.

Then, let us apply Dijkstra's algorithm again to find a shortest path p_v from v_0 to every vertex $v \in V$ and its length d^0_v.

Finally, let us find a vertex v^* in which $\min_{v \in V}(d^0_v + d_v)$ is reached. It is clear that the corresponding shortest directed path p_{v^*} and directed cycle C_{v^*} define the desired new strategy x'_i.

□

5 Nash-Solvability of Two-Person Game Forms

If $n = 2$ and $c \in A$ is the worst outcome for both players, Nash-solvability was proven in [1]. In fact, the last assumption is not necessary: even if outcome c is ranked by two players arbitrarily, Nash-solvability still holds. This observation was recently made by Gimbert and Sørensen.

A two-person game form g is called:

Nash-solvable if for every utility function $u : \{1, 2\} \times A \to \mathbb{R}$ the obtained game (g, u) has a Nash equilibrium.

zero-sum-solvable if for each zero-sum utility function ($u_1(a) + u_2(a) = 0$ for all $a \in A$) the obtained zero-sum game (g, u) has a Nash equilibrium, which is called a saddle point for zero-sum games.

$\pm$-*solvable* if zero-sum solvability holds for each u that takes only values: $+1$ and -1.

Necessary and sufficient conditions for zero-sum solvability were obtained by Edmonds and Fulkerson [3] in 1970; see also [5]. Somewhat surprisingly, these conditions remain necessary and sufficient for $\pm$-solvability and for Nash-solvability, as well; in other words, all three above types of solvability are equivalent, in case of two-person game forms [6]; see also [7] and Appendix 1 of [2].

Proposition 2. *([4]). Each two-person positional game form in which all cycles form one outcome is Nash-solvable.*

Proof. Let $\mathcal{G} = (G, D, v_0, u)$ be a two-person zero-sum positional game, where $u : I \times A \to \{-1, +1\}$ is a zero-sum ± 1 utility function. Let $A_i \subseteq A$ denote the outcomes winning for player $i \in I = \{1, 2\}$. Without any loss of generality we can assume that $c \in A_1$, that is, $u_1(c) = 1$, while $u_2(c) = -1$. Let $V^2 \subseteq V$ denote the set of positions in which player 2 can enforce a terminal from A_2. Then, obviously, player 2 wins whenever $v_0 \in V^2$. Let us prove that player 1 wins otherwise, when $v_0 \in V^1 = V \setminus V^2$.

Indeed, if $v \in V^1 \cap V_2$ then $v' \in V^1$ for every move (v, v') of player 2; if $v \in V^1 \cap V_1$ then player 1 has a move (v, v') such that $v' \in V_1$. Let player 1 choose such a move for every position $v \in V^1 \cap V_1$ and an arbitrary move in each remaining position $v \in V^2 \cap V_1$. This rule defines a strategy x_1. Let us fix an arbitrary strategy x_2 of player 2 and consider the profile $x = (x_1, x_2)$. Obviously, the play $p(x)$ cannot come to V_2 if $v_0 \in V_1$. Hence, for the outcome $a = a(x)$ we have: either $a \in V^1$ or $a = c$. In both cases player 1 wins. Thus, the game is Nash-solvable. $\square$

Let us recall that this result also follows immediately from Theorem 3.

Finally, let us briefly consider a refinement of the Nash equilibrium concept, the so-called *subgame perfect* equilibrium, where a strategy profile is an equilibrium regardless of the choice of starting position.

It is not difficult to see that already for two-person games a Nash equilibrium can be unique but not subgame perfect. Let us consider the example in Figure 10. There are only four different strategy profiles and for all of them there is a choice of starting position for which the profile is not an equilibrium.

6 Conclusions and Open Problems

Nash-solvability of n-person positional games (in which all directed cycles form a single outcome) holds for $n = 2$ and remains an open problem for $n > 2$.

For $n = 2$, we prove strong ri-acyclicity, which implies Nash-solvability. Computing Nash equilibria efficiently is another interesting issue for further investigation.

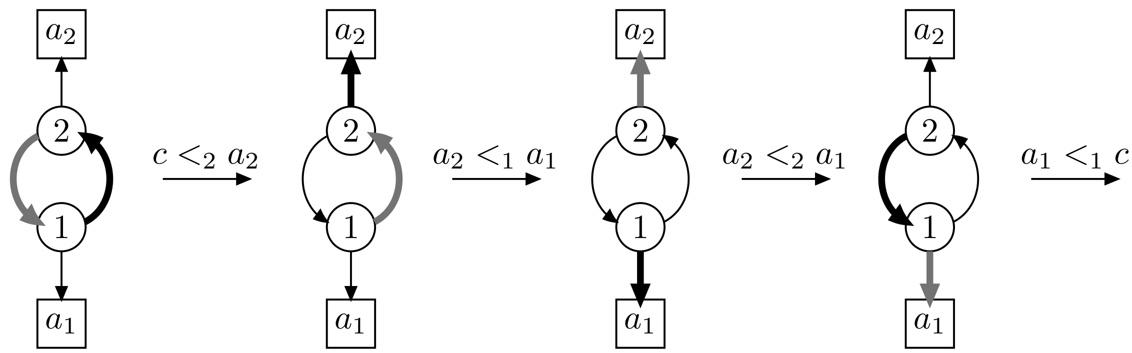

$1: \ a_2 < a_1 < c$

$2: \ c < a_2 < a_1$

Figure 10: Two-person game with no subgame perfect positional strategies. The improvements do not obey any inside play restriction, since there is no fixed starting position.

For $n \geq 4$ there are examples of best reply strong c-free ri-cycles. Yet, it remains open whether such examples exist for $n = 3$.

Acknowledgements.

We are thankful to Gimbert, Kukushkin, and Sørensen for helpful discussions.

References

[1] E. Boros and V. Gurvich, On Nash-solvability in pure strategies of finite games with perfect information which may have cycles. Math. Social Sciences 46 (2003).

[2] E. Boros, V. Gurvich, K. Makino, and Wei Shao, Nash-solvable bidirected cyclic two-person game forms, Rutcor Research Report 26-2007 and DIMACS Technical Report 2008-13, Rutgers University.

[3] J. Edmonds and D.R. Fulkerson, Bottleneck Extrema, RM-5375-PR, The Rand Corporation, Santa Monica, Ca., Jan. 1968; J. Combin. Theory, 8 (1970), 299-306.

[4] H. Gimbert and T.B. Sørensen, Private communications, July 2008.

[5] V. Gurvich, To theory of multi-step games USSR Comput. Math. and Math. Phys. 13 (6) (1973), 143-161.

[6] V. Gurvich, Solution of positional games in pure strategies, USSR Comput. Math. and Math. Phys. 15 (2) (1975), 74-87.

[7] V. Gurvich, Equilibrium in pure strategies, Soviet Mathematics Doklady 38 (3) (1988), 597-602.

[8] N.S. Kukushkin, Perfect information and potential games, Games and Economic Behavior, **38** (2002), 306–317.

[9] N.S. Kukushkin, Acyclicity of improvements in finite game forms, Manuscript, http://www.ccas.ru/mmes/mmeda/ququ/GF.pdf (2009)

[10] N. S. Kukushkin, Private communications, August 2008.

[11] E. Zermelo, Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels, *Proc. 5th Int. Cong. Math. Cambridge 1912* , Vol. II, Cambridge University Press (1913) 501–504.