More Games of No Chance
MSRI Publications
Volume **42**, 2002

A Symmetric Strategy in Graph Avoidance Games

FRANK HARARY, WOLFGANG SLANY, AND OLEG VERBITSKY

ABSTRACT. In the graph avoidance game two players alternately color the edges of a graph G in red and in blue respectively. The player who first creates a monochromatic subgraph isomorphic to a forbidden graph F loses. A *symmetric strategy* of the second player ensures that, independently of the first player's strategy, the blue and the red subgraph are isomorphic after every round of the game. We address the class of those graphs G that admit a symmetric strategy for all F and discuss relevant graph-theoretic and complexity issues. We also show examples when, though a symmetric strategy on G generally does not exist, it is still available for a particular F.

1. Introduction

In a broad class of games that have been studied in the literature, two players, A and B, alternately color the edges of a graph G in red and in blue respectively. In the achievement game the objective is to create a monochromatic subgraph isomorphic to a given graph F. In the avoidance game the objective is, on the contrary, to avoid creating such a subgraph. Both the achievement and the avoidance games have strong and weak versions. In the strong version A and B both have the same objective. In the weak version B just plays against A, that is, tries either to prevent A from creating a copy of F in the achievement game or to force such creation in the avoidance game. The weak achievement game, known also as the Maker-Breaker game, is most studied [4; 1; 13]. Our paper is motivated by the strong avoidance game [7; 5] where monochromatic F-subgraphs of G are forbidden, and the player who first creates such a subgraph loses.

Slany's research was partly supported by Austrian Science Foundation grant Z29-INF. This work was done while Verbitsky was visiting the Institut für Informationssysteme at the Technische Universität Wien, supported by a Lise Meitner Fellowship of the Austrian Science Foundation (FWF grant M 532).

The instance of a strong avoidance game with complete graphs $G = K_6$ and $F = K_3$ is well known under the name SIM [15]. Since for any edge bicoloring of K_6 there is a monochromatic K_3, a draw in this case is impossible. It is proven in [12] that a winning strategy in SIM is available for B. A few other results for small graphs are known [7]. Note that, in contrast to the weak achievement games, if B has a winning strategy in the avoidance game on G with forbidden F and if G is a subgraph of G', then it is not necessary that B also has a winning strategy on G' with forbidden F. Recognition of a winner seems generally to be a nontrivial task both from the combinatorial and the complexity-theoretic point of view (for complexity issues see, e.g., [16]).

In this paper we introduce the notion of a *symmetric strategy*[1] for B. We say that B follows a symmetric strategy on G if after every move of B the blue and the red subgraphs are isomorphic, regardless of A's strategy. As easily seen, if B plays so, he at least does not lose in the avoidance game on G with any forbidden F. There is a similarity with the *mirror-image strategy* of A in the achievement game [2]. However, the latter strategy is used on two disjoint copies of the complete graph, and therefore in our case things are much more complicated.

We address the class $\mathcal{C}_{\rm sym}$ of those graphs G on which a symmetric strategy for B exists. We observe that $\mathcal{C}_{\rm sym}$ contains all graphs having an involutory automorphism without fixed edges. This subclass of $\mathcal{C}_{\rm sym}$, denoted by $\mathcal{C}_{\rm auto}$, includes even paths and cycles, bipartite complete graphs $K_{s,t}$ with s or t even, cubes, and the Platonic graphs except the tetrahedron. We therefore obtain a lot of instances of the avoidance game with a winning strategy for B. More instances can be obtained based on closure properties of $\mathcal{C}_{\rm auto}$ that we check with respect to a few basic graph operations.

Nevertheless, recognizing a suitable automorphism and, therefore, using the corresponding symmetric strategy is not easy. Based on a related result of Lubiw [11], we show that deciding membership in $\mathcal{C}_{\rm auto}$ is NP-complete.

We then focus on games on complete graphs. We show that K_n is not in $\mathcal{C}_{\rm sym}$ for all $n \geq 4$. Moreover, for an arbitrary strategy of B, A is able to violate the isomorphism between the red and the blue subgraphs in at most $n - 1$ moves. Nevertheless, we consider the avoidance game on K_n with forbidden P_2, a path of length 2, and point out a simple symmetric strategy making B the winner. This shows an example of a graph G for which, while a symmetric strategy in the avoidance game does not exist in general, it does exist for a particular forbidden F.

The paper is organized as follows. Section 2 contains the precise definitions. In Section 3 we compile the membership list for $\mathcal{C}_{\rm sym}$ and $\mathcal{C}_{\rm auto}$. In Section 4 we investigate the closure properties of $\mathcal{C}_{\rm sym}$ and $\mathcal{C}_{\rm auto}$ with respect to various graph products. In Section 5 we prove the NP-completeness of $\mathcal{C}_{\rm auto}$. Section 6 analyses the avoidance game on K_n with forbidden P_2.

[1]Note that this term has been used also in other game-theoretic situations (see, e.g., [14]).

2. Definitions

We deal with two-person positional games of the following kind. Two players, A and B, alternately color the edges of a graph G in red and in blue respectively. Player A starts the game. In a *move* a player colors an edge that was so-far uncolored. The i-th *round* consists of the i-th move of A and the i-th move of B. Let a_i (resp. b_i) denote an edge colored by A (resp. B) in the i-th round.

A *strategy* for a player determines the edge to be colored by him at every round of the game. Formally, let ε denote the empty sequence. A strategy of A is a function S_1 that maps every possibly empty sequence of distinct edges $e_1, \dots, e_i$ into an edge different from e_1, …, and e_i and from $S_1(\varepsilon)$, $S_1(e_1)$, $S_1(e_1, e_2)$, …, and $S_1(e_1, \dots, e_{i-1})$. A strategy of B is a function S_2 that maps every nonempty sequence of distinct edges $e_1, \dots, e_i$ into an edge different from e_1, …, and e_i and from $S_2(e_1)$, $S_2(e_1, e_2)$, …, and $S_2(e_1, \dots, e_{i-1})$. If A follows a strategy S_1 and B follows a strategy S_2, then $a_i = S_1(b_1, \dots, b_{i-1})$ and $b_i = S_2(a_1, \dots, a_i)$.

Let $A_i = \{a_1, \dots, a_i\}$ (resp. $B_i = \{b_1, \dots, b_i\}$) consist of the red (resp. blue) edges colored up to the i-th round. A *symmetric strategy of* B *on* G ensures that, regardless of A's strategy, the subgraphs A_i and B_i are isomorphic for every $i \leq m/2$, where m is the number of edges of G.

The class of all graphs G on which B has a symmetric strategy will be denoted by $\mathcal{C}_{\mathrm{sym}}$.

Suppose that we are given graphs G and F and that F is a subgraph of G. The *avoidance game on* G *with a forbidden subgraph* F or, shortly, the game AVOID(G, F) is played as described above with the following ending condition: The player who first creates a monochromatic subgraph of G isomorphic to F loses.

Note that a symmetric strategy of B on G is nonlosing for B in AVOID(G, F), for every forbidden F. Really, the assumption that B creates a monochromatic copy of F implies that such a copy is already created by A earlier in the same round.

3. Automorphism-based strategy

Recall that the *order* of a graph is its number of vertices and the *size* of a graph is its number of edges. Given a graph G, we denote its vertex set by $V(G)$ and its edge set by $E(G)$. An *automorphism* of a graph G is a permutation of $V(G)$ that preserves the vertex adjacency. Recall that the *order* of a permutation is the minimal k such that the k-fold composition of the permutation is the identity permutation. In particular, a permutation of order 2, also called an *involution*, coincides with its inversion. We call an automorphism of order 2 *involutory*.

The symmetric strategy can be realized if a graph G has an involutory automorphism that moves every edge. More precisely, an automorphism $\phi : V(G) \to V(G)$ determines a permutation $\phi' : E(G) \to E(G)$ by $\phi'(\{u, v\}) = \{\phi(u), \phi(v)\}$.

We assume that ϕ is involutory and ϕ' has no fixed element. In this case, whenever A chooses an edge e, B chooses the edge $\phi'(e)$. This strategy of B is well defined because $E(G)$ is partitioned into 2-subsets of the form $\{e, \phi'(e)\}$. This strategy is really symmetric because after completion of every round ϕ induces an isomorphism between the red and the blue subgraphs. We will call such a strategy *automorphism-based*.

Definition 3.1. $\mathcal{C}_{\rm auto}$ is a subclass of $\mathcal{C}_{\rm sym}$ consisting of all those graphs G on which B has an automorphism-based symmetric strategy.

We now list some examples of graphs in $\mathcal{C}_{\rm auto}$.

Example 3.2 (Graphs in $\mathcal{C}_{\rm auto}$)**.** (i) P_n, a path of length n, if n is even.
(ii) C_n, a cycle of length n, if n is even.
(iii) Four Platonic graphs excluding the tetrahedron.
(iv) Cubes of any dimension.[2]
(v) Antipodal graphs (in the sense of [3]) of size more than 1. Those are connected graphs such that for every vertex v, there is a unique vertex $\bar{v}$ of maximum distance from v. The correspondence $\phi(v) = \bar{v}$ is an automorphism [9]. As easily seen, it is involutory and has no fixed edge. The class of antipodal graphs includes the graphs from the three preceding items.
(vi) $K_{s,t}$, a bipartite graph whose classes have s and t vertices, if st is even.
(vii) $K_{s,t} - e$, that is, $K_{s,t}$ with an edge deleted, provided st is odd.
(viii) K_n, a complete graph on n vertices, with a matching of size $\lfloor n/2 \rfloor$ deleted. Note that in this and the preceding examples, for all choices of edges to be deleted, the result of deletion is the same up to an isomorphism.

It turns out that a symmetric strategy is not necessarily automorphism-based.

Theorem 3.3. *$\mathcal{C}_{\rm auto}$ is a proper subclass of $\mathcal{C}_{\rm sym}$.*

Below is a list of a few separating examples.

Example 3.4 (Graphs in $\mathcal{C}_{\rm sym} \setminus \mathcal{C}_{\rm auto}$)**.** (i) A triangle with one more edge attached (the first graph in Figure 1). This is the only connected separating example of even size we know. In particular, none of the connected graphs of size 6 is in $\mathcal{C}_{\rm sym} \setminus \mathcal{C}_{\rm auto}$. Note that the definition of $\mathcal{C}_{\rm sym}$ does not exclude graphs of odd size, as given in the further examples.
(ii) The graphs of size 5 shown in Figure 1.
(iii) Paths P_1, P_3, and P_5.
(iv) Cycles C_3, C_5, and C_7.
(v) Stars $K_{1,n}$, if n is odd.

Note that in spite of items 4 and 5, P_7 and C_9 are not in $\mathcal{C}_{\rm sym}$.

[2] More generally, cubes are a particular case of grids, i.e., Cartesian products of paths. The central symmetry of a grid moves each edge unless exactly one of the factors is an odd path.

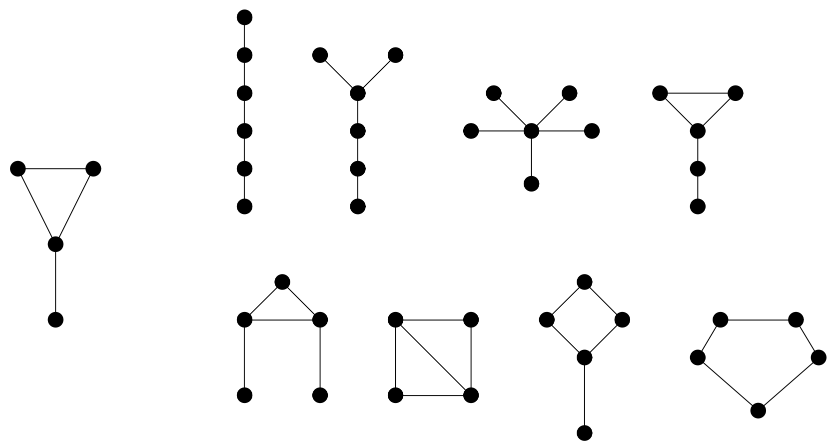

Figure 1. Graphs of size 4 and 5 that are in $\mathcal{C}_{\text{sym}}$ but not in $\mathcal{C}_{\text{auto}}$.

Question 3.5. How much larger is $\mathcal{C}_{\text{sym}}$ than $\mathcal{C}_{\text{auto}}$? Are there other connected separating examples than those listed above?

4. Closure properties of $\mathcal{C}_{\text{auto}}$

We now recall a few operations on graphs. Given two graphs G_1 and G_2, we define a product graph on the vertex set $V(G_1) \times V(G_2)$ in three ways. Two vertices (u_1, u_2) and (v_1, v_2) are adjacent in the *Cartesian product* $G_1 \times G_2$ if either $u_1 = v_1$ and $\{u_2, v_2\} \in E(G_2)$ or $u_2 = v_2$ and $\{u_1, v_1\} \in E(G_1)$; in the *lexicographic product* $G_1[G_2]$ if either $\{u_1, v_1\} \in E(G_1)$ or $u_1 = v_1$ and $\{u_2, v_2\} \in E(G_2)$; in the *categorical product* $G_1 \cdot G_2$ if $\{u_1, v_1\} \in E(G_1)$ and $\{u_2, v_2\} \in E(G_2)$.

If the vertex sets of G_1 and G_2 are disjoint, we define the *sum* (or *disjoint union*) $G_1 + G_2$ to be the graph with vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2)$.

Using these graph operations, from Example 3.2 one can obtain more examples of graphs in $\mathcal{C}_{\text{sym}}$. Note that the class of antipodal graphs itself is closed with respect to the Cartesian product [9].

Theorem 4.1. (i) $\mathcal{C}_{\text{auto}}$ *is closed with respect to the sum and with respect to the Cartesian, the lexicographic, and the categorical products.*
(ii) *Moreover,* $\mathcal{C}_{\text{auto}}$ *is an ideal with respect to the categorical product, that is, if* G *is in* $\mathcal{C}_{\text{auto}}$ *and* H *is arbitrary, then both* $G \cdot H$ *and* $H \cdot G$ *are in* $\mathcal{C}_{\text{auto}}$.

Proof. For the sum the claim 1 is obvious. Consider three auxiliary product notions. Given two graphs G_1 and G_2, we define product graphs $G_1 \otimes_1 G_2$, $G_1 \otimes_2 G_2$, and $G_1 \otimes_3 G_2$ each on the vertex set $V(G_1) \times V(G_2)$. Two vertices (u_1, u_2) and (v_1, v_2) are adjacent in $G_1 \otimes_1 G_2$ if $\{u_1, v_1\} \in E(G_1)$ and $u_2 = v_2$; in $G_1 \otimes_2 G_2$ if $\{u_1, v_1\} \in E(G_1)$ and $u_2 \neq v_2$; and in $G_1 \otimes_3 G_2$ if $u_1 = v_1$ and $\{u_2, v_2\} \in E(G_2)$.

Given two permutations, ϕ_1 of $V(G_1)$ and ϕ_2 of $V(G_2)$, we define a permutation ψ of $V(G_1) \times V(G_2)$ by $\psi(u_1, u_2) = (\phi_1(u_1), \phi_2(u_2))$. If both ϕ_1 and ϕ_2 are involutory, so is ψ. If ϕ_1 and ϕ_2 are automorphisms of G_1 and G_2 respectively, then ψ is an automorphism of each $G_1 \otimes_i G_2$, $i = 1, 2, 3$. Finally, it is not hard to see that if both ϕ_1 and ϕ_2 move all edges, so does ψ in each $G_1 \otimes_i G_2$, $i = 1, 2, 3$.

Notice now that $E(G_1 \otimes_1 G_2)$, $E(G_1 \otimes_2 G_2)$, and $E(G_1 \otimes_3 G_2)$ are pairwise disjoint. Notice also that $E(G_1 \times G_2) = E(G_1 \otimes_1 G_2) \cup E(G_1 \otimes_3 G_2)$ and $E(G_1[G_2]) = E(G_1 \otimes_1 G_2) \cup E(G_1 \otimes_2 G_2) \cup E(G_1 \otimes_3 G_2)$. It follows that if ϕ_1 and ϕ_2 are fixed-edge-free involutory automorphisms of G_1 and G_2 respectively, then ψ is a fixed-edge-free involutory automorphism of both $G_1 \times G_2$ and $G_1[G_2]$. Thus, $\mathcal{C}_{\mathrm{auto}}$ is closed with respect to the Cartesian and the lexicographic products.

To prove the claim 2, let $G \in \mathcal{C}_{\mathrm{auto}}$, ϕ be a fixed-edge-free involutory automorphism of G, and H be an arbitrary graph. Define a permutation ψ of $V(G \cdot H)$ by $\psi(u, v) = (\phi(u), v)$. It is not hard to see that ψ is a fixed-edge-free involutory automorphism of $G \cdot H$. Thus, $G \cdot H \in \mathcal{C}_{\mathrm{auto}}$. The same is true for $H \cdot G$ because $G \cdot H$ and $H \cdot G$ are isomorphic.

Example 4.2 ($\mathcal{C}_{\mathrm{sym}}$ is not closed with respect to the Cartesian, the lexicographic, and the categorical products)**.** Denote the first graph in Example 3.4 by K_3+e. The following product graphs are not in $\mathcal{C}_{\mathrm{sym}}$: $(K_3+e)\times P_2$, $P_2[K_3+e]$, and $(K_3 + e) \cdot (K_3 + e)$. To show this, for each of these graphs we will describe a strategy allowing A to destroy an isomorphism between the red and the blue subgraphs, regardless of B's strategy.

$(K_3 + e) \times P_2$ has a unique vertex v of the maximum degree 5, and v is connected to the two vertices v_1 and v_2 of degree 4 that are connected to each other. In the first move of a symmetry-breaking strategy, A chooses the edge $\{v, v_1\}$. If B chooses an edge not incident to v, A creates a star $K_{1,5}$ and thus breaks B's symmetric strategy. If B chooses an edge incident to v but not $\{v, v_2\}$, A chooses $\{v, v_2\}$ and breaks B's symmetric strategy by creating a triangle K_3 in the third move. Assume therefore that in the second round B chooses $\{v, v_2\}$. In the next moves A creates a star with center v. If B tries to create a star with the same center, the symmetry eventually breaks because A can create a $K_{1,3}$ while B can create at most a $K_{1,2}$. Assume therefore that in the first four rounds A creates a $K_{1,4}$ with center v and B creates a $K_{1,4}$ with center v_2 (see Figure 2). In the rounds 5–8 A attaches a new edge to every leaf of the red star. Player B's symmetric strategy is broken because he cannot attach any edge to v.

$P_2[K_3 + e]$ consists of three copies of $K_3 + e$ on the vertex sets $\{u_1, u_2, u_3, u_4\}$, $\{v_1, v_2, v_3, v_4\}$, and $\{w_1, w_2, w_3, w_4\}$, and of 32 edges $\{v_i, u_j\}$ and $\{v_i, w_j\}$ for all $1 \le i, j \le 4$ (see Figure 3). The vertex v_1 has the maximum degree 11, v_2 and v_3 have degree 10, v_4 has degree 9, and all other vertices have degree at most 7. In the first move of a symmetry-breaking strategy A chooses the edge $\{v_1, v_2\}$. If B in response does not choose $\{v_1, v_3\}$, A does it and breaks B's

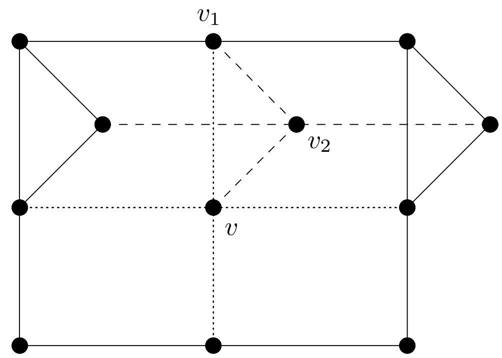

Figure 2. First four rounds of $\mathcal{A}$'s symmetry-breaking strategy on $(K_3 + e) \times P_2$ ($\mathcal{A}$'s edges dotted, $\mathcal{B}$'s edges dashed, uncolored edges continuous).

symmetric strategy by creating a star with center v_1. If B chooses $\{v_1, v_3\}$, in the second move A chooses $\{v_1, u_4\}$. If B then chooses an edge going out of v_1, A breaks B's symmetric strategy by creating a $K_{1,6}$. Assume therefore that in the second move B chooses an edge $\{v_3, x\}$. If $x = v_2$ or $x = u_4$, A chooses $\{u_4, v_2\}$ and breaks B's symmetric strategy by creating a triangle K_3. Assume therefore that x is another vertex (for example, $x = u_1$ as in Figure 3). In the third move A chooses $\{v_2, v_3\}$. If B chooses $\{x, v_2\}$, A chooses $\{u_4, v_3\}$ and breaks B's symmetric strategy by creating a quadrilateral C_4. Otherwise, in the next moves A creates a star $K_{1,10}$ with center v_2. Player B's symmetric strategy eventually is broken because he can create at most a $K_{1,9}$ with center v_1 or v_3 or at most a $K_{1,7}$ with center x.

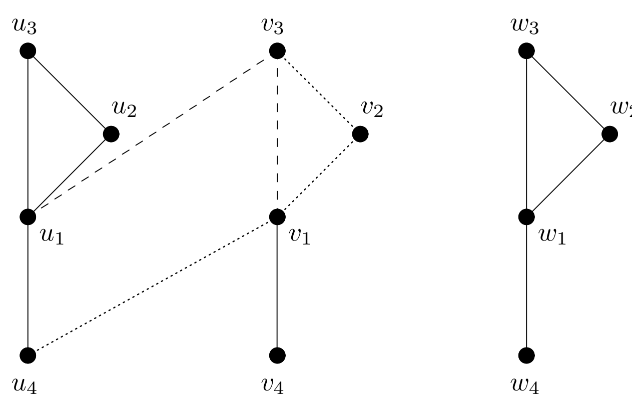

Figure 3. First three moves of $\mathcal{A}$'s symmetry-breaking strategy on $P_2[K_3 + e]$ ($\mathcal{A}$'s edges dotted, $\mathcal{B}$'s edges dashed, uncolored edges continuous, uncolored edges $\{v_i, u_j\}$, $\{v_i, w_j\}$ not shown).

$(K_3 + e) \cdot (K_3 + e)$ has a unique vertex v of the maximum degree 9, whereas all other vertices have degree at most 6. A symmetry-breaking strategy of A consists in creating a star with center v.

Remark 4.3. $\mathcal{C}_{\mathrm{sym}}$ is not closed with respect to the sum because, for example, it does not contain $K_3 + P_3$. Nevertheless, if G_1 and G_2 are in $\mathcal{C}_{\mathrm{sym}}$ and both have even size, $G_1 + G_2$ is easily seen to be in $\mathcal{C}_{\mathrm{sym}}$.

5. Complexity of $\mathcal{C}_{\mathrm{auto}}$

Though the graph classes listed in Example 3.2 have efficient membership tests, in general the existence of an involutory automorphism without fixed edges is not easy to determine.

Theorem 5.1. *Deciding membership of a given graph G in the class $\mathcal{C}_{\mathrm{auto}}$ is NP-complete.*

Proof. Consider the related problem Order 2 Fixed-Point-Free Automorphism whose NP-completeness was proven in [11]. This is the problem of recognition if a given graph has an involutory automorphism without fixed *vertices*. We describe a polynomial time reduction R from Order 2 Fixed-Point-Free Automorphism to $\mathcal{C}_{\mathrm{auto}}$.

Given a graph G, we perform two operations:

Step 1. Split every edge into two adjacent edges by inserting a new vertex, i.e., form the subdivision graph $S(G)$ (see [6, p. 80]).
Step 2. Attach a 3-star by a leaf at every non-isolated vertex of $S(G)$ which also was in G.

As a result we obtain $R(G)$ (see an example in Figure 4). We have to prove that G has an involutory automorphism without fixed vertices if and only if $R(G)$ has an involutory automorphism without fixed edges.

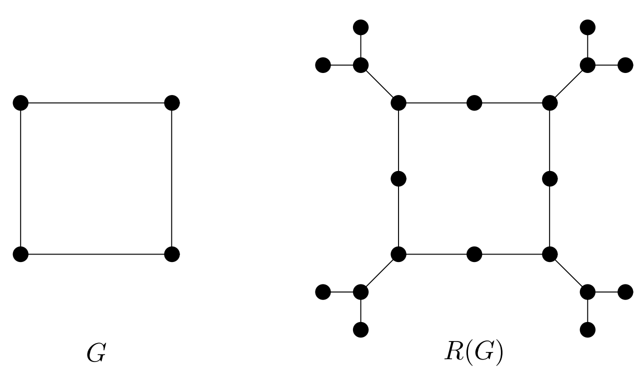

Figure 4. An example of the reduction.

Every involutory automorphism of G without fixed vertices determines an involutory automorphism of $R(G)$ that, thanks to the new vertices, has no fixed edge. On the other hand, consider an arbitrary automorphism ψ of $R(G)$. Since ψ maps the set of vertices of degree 1 in $R(G)$ onto itself, ψ maps every 3-star added in Step 2 into another such 3-star (or itself) and therefore it maps $V(G)$ onto itself. Suppose that u and v are two vertices adjacent in G and let z be the vertex inserted between u and v in Step 1. Then $\psi(z)$ is adjacent in $R(G)$ with both $\psi(u)$ and $\psi(v)$. As easily seen, $\psi(z)$ can appear in $R(G)$ only in Step 1 and therefore $\psi(u)$ and $\psi(v)$ are adjacent in G. This proves that ψ induces an automorphism of G. The latter is involutory if so is ψ. Finally, if ψ has no fixed edge, then every 3-star added in Step 2 is mapped to a different such 3-star and consequently the induced automorphism of G has no fixed vertex.

Theorem 5.1 implies that, despite the combinatorial simplicity of an automorphism-based strategy, realizing this strategy by B on $G \in \mathcal{C}_{\text{auto}}$ requires him to be at least NP powerful. The reason is that an automorphism-based strategy subsumes finding an involutory fixed-edge-free automorphism of any given $G \in \mathcal{C}_{\text{auto}}$, whereas this problem is at least as hard as testing membership in $\mathcal{C}_{\text{auto}}$.

Given the order or the size, there are natural ways of efficiently generating a graph in $\mathcal{C}_{\text{auto}}$ with respect to a certain probability distribution. Theorem 5.1 together with such a generating procedure has two imaginable applications in "real-life" situations.

Negative scenario. Player B secretly generates $G \in \mathcal{C}_{\text{auto}}$ and makes an offer to A to choose F at his discretion and play the game AVOID(G, F). If A accepts, then B, who knows a suitable automorphism of G, follows the automorphism-based strategy and at least does not lose. A is not able to observe that $G \in \mathcal{C}_{\text{auto}}$, unless he can efficiently solve NP.[3]

Positive scenario. Player A insists that before the game an impartial third person, hidden from B, permutes at random the vertices of G. Then applying the automorphism-based strategy in the worst case becomes for Player B as hard as testing isomorphism of graphs. More precisely, Player B faces the following search problem.

PAR (Permuted Automorphism Reconstruction)
Input: G, H, and β, where G and H are isomorphic graphs in $\mathcal{C}_{\text{auto}}$, and β is a fixed-edge-free involutory automorphism of H.
Find: α, a fixed-edge-free involutory automorphism of G.

We relate this problem to GI, the Graph Isomorphism problem, that is, given two graphs G_0 and G_1, to recognize if they are isomorphic. We use the notion of the Turing reducibility extended in a natural way over search problems. We

[3] We assume here that $\mathcal{A}$ fails to decide if $G \in \mathcal{C}_{\text{auto}}$ at least for some G. We could claim this failure for *most* G if $\mathcal{C}_{\text{auto}}$ would be proven to be complete for the average case [10].

say that two problems are polynomial-time equivalent if they are reducible one to another by polynomial-time Turing reductions.

Theorem 5.2. *The problems* PAR *and* GI *are polynomial-time equivalent.*

Proof. We use the well-known fact that the decision problem GI is polynomial-time equivalent with the search problem of finding an isomorphism between two given graphs [8, Section 1.2].

A reduction from PAR *to* GI. We describe a simple algorithm solving PAR under the assumption that we are able to construct a graph isomorphism. Given an input (G, H, β) of PAR, let π be an isomorphism from G to H. As easily seen, computing the composition $\alpha = \pi^{-1}\beta\pi$ gives us a solution of PAR.

A reduction from GI *to* PAR. We will describe a reduction to PAR from the problem of constructing an isomorphism between two graphs G_0 and G_1 of the same size. We assume that both G_0 and G_1 are connected and their size is odd. To ensure the odd size, one can just add an isolated edge to both of the graphs. To ensure the connectedness, one can replace the graphs with their complements. To not lose the odd size, the latter operation should be applied only for graphs of order n such that $n(n-1)/2$ is even; hence adding one or two isolated vertices may be required beforehand. If we find an isomorphism between the modified graphs, an isomorphism between the original graphs is easily reconstructed.

We form the triple (G, H, β) by setting $G = G_0 + G_1$, $H = G_0 + G_0$, and taking β to be the identity map between the two copies of G_0. If G_0 and G_1 are isomorphic, this is a legitimate instance of PAR. By the connectedness of G_0 and G_1, if $\alpha : V(G) \to V(G)$ is a solution of PAR on this instance, it either acts within the connected components $V(G_0)$ and $V(G_1)$ independently or maps $V(G_0)$ to $V(G_1)$ and vice versa. The first possibility actually cannot happen because the size of G_0 and G_1 is odd and hence α cannot be at the same time involutory and fixed-edge-free. Thus α is an isomorphism between G_0 and G_1.

Question 5.3. Is deciding membership in $\mathcal{C}_{\mathrm{sym}}$ NP-hard? A priori we can say only that $\mathcal{C}_{\mathrm{sym}}$ is in PSPACE. Of course, if the difference $\mathcal{C}_{\mathrm{sym}} \setminus \mathcal{C}_{\mathrm{auto}}$ is decidable in polynomial time, then NP-completeness of $\mathcal{C}_{\mathrm{sym}}$ would follow from Theorem 5.1.

6. Game AVOID(K_n, P_2)

Games on complete graphs are particularly interesting. Notice first of all that in this case a symmetric strategy is not available.

Theorem 6.1. *$K_n \notin \mathcal{C}_{\mathrm{sym}}$ for $n \geq 4$.*

Proof. We describe a strategy of A that violates the isomorphism between the red and the blue subgraphs at latest in the $(n-1)$-th round. In the first two rounds A chooses two adjacent edges ensuring that at least one of them is adjacent also to the first edge chosen by B. Thus, after the second round the game can be in one of five positions depicted in Figure 5.

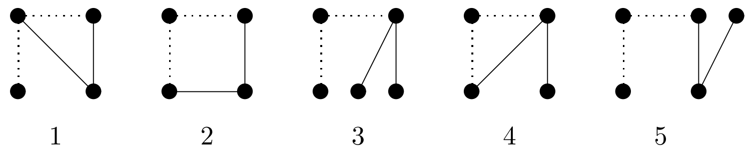

Figure 5. First two rounds of $\mathcal{A}$'s symmetry-breaking strategy ($\mathcal{A}$'s edges dotted, $\mathcal{B}$'s edges continuous).

In positions 1 and 2 A creates a triangle, which is impossible for B. In positions 3, 4, and 5 A creates an $(n-1)$-star, while B is able to create at most an $(n-2)$-star (in position 5 A first of all chooses the uncolored edge connecting two vertices of degree 2).

Let us define $\mathcal{C}_{\mathrm{II}}$ to be the class of all graphs G such that, for all F, B has a nonlosing strategy in the game AVOID(G, F). Clearly, $\mathcal{C}_{\mathrm{II}}$ contains $\mathcal{C}_{\mathrm{sym}}$. It is easy to check that K_4 is in $\mathcal{C}_{\mathrm{II}}$, and therefore $\mathcal{C}_{\mathrm{sym}}$ is a proper subclass of $\mathcal{C}_{\mathrm{II}}$.

It is an interesting question if $K_n \in \mathcal{C}_{\mathrm{II}}$ for all n. We examine the case of a forbidden subgraph $F = P_2$, a path of length 2. For all $n > 2$, we describe an efficient winning strategy for B in AVOID(K_n, P_2). Somewhat surprisingly, this strategy, in contrast to Theorem 6.1, proves to be symmetric in a weaker sense.

More precisely, we say that a strategy of B is *symmetric in AVOID*(G, F) if, independently of A's strategy, the red and the blue subgraphs are isomorphic after every move of B in the game. Let us stress the difference with the notion of a symmetric strategy on G we used so far. While a strategy symmetric on G guarantees the isomorphism until G is completely colored (except one edge if G has odd size), a strategy symmetric in AVOID(G, F) guarantees the isomorphism only as long as A does not lose in AVOID(G, F).

Theorem 6.2. *Player* B *has a symmetric strategy in the game AVOID*(K_n, P_2).

Proof. Let A_i (resp. B_i) denote the set of the edges chosen by A (resp. B) in the first i rounds. The strategy of B is, as long as A_i is a matching, to choose an edge so that the subgraph of K_n with edge set $A_i \cup B_i$ is a path. The only case when this is impossible is that n is even and $i = n/2$. Then B chooses the edge that makes $A_i \cup B_i$ a Hamiltonian cycle (see Figure 6).

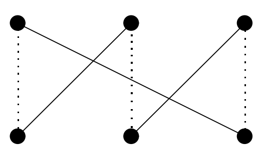

Figure 6. A game (K_6, P_2) ($\mathcal{A}$'s edges dotted, $\mathcal{B}$'s edges continuous).

Question 6.3. What is the complexity of deciding, given G, whether or not B has a winning strategy in AVOID(G, P_2)?

It is worth noting that in [16], PSPACE-completeness of the winner recognition in the avoidance game with precoloring is proven even for a fixed forbidden graph F, namely for two triangles with a common vertex called the "bowtie graph". Notice also that AVOID(G, P_2) has an equivalent vertex-coloring version: the players color the vertices of the line graph $L(G)$ and the loser is the one who creates two adjacent vertices of the same color.

Question 6.4. Does K_n belong to $\mathcal{C}_{\mathrm{II}}$? In particular, does B have winning strategies in AVOID$(K_n, K_{1,3})$, AVOID(K_n, P_3), and AVOID(K_n, K_3) for large enough n?

References

[1] J. Beck. Van der Waerden and Ramsey type games. *Combinatorica* 1:103–116 (1981).

[2] E. R. Berlekamp, J. H. Conway, R. K. Guy. *Winning ways for your mathematical plays.* Academic Press, New York (1982).

[3] A. Berman, A. Kotzig, G. Sabidussi. Antipodal graphs of diameter 4 and extremal girth. In: *Contemporary methods in graph theory,* R. Bodendiek (ed.), BI-Wiss.-Verl., Mannheim, pp. 137–150 (1990).

[4] P. Erdős, J. L. Selfridge. On a combinatorial game. *J. Combin. Theory A* 14:298–301 (1973).

[5] M. Erickson, F. Harary. Generalized Ramsey theory XV: Achievement and avoidance games for bipartite graphs. *Graph theory*, Proc. 1st Southeast Asian Colloq., Singapore 1983, Lect. Notes Math. 1073:212–216 (1984).

[6] F. Harary. *Graph theory.* Addison-Wesley, Reading MA (1969).

[7] F. Harary. Achievement and avoidance games for graphs. *Ann. Discrete Math.* 13:111–119 (1982).

[8] J. Köbler, U. Schöning, and J. Torán. *The Graph Isomorphism problem: its structural complexity.* Birkhäuser (1993).

[9] A. Kotzig On centrally symmetric graphs. (in Russian) *Czech. Math. J.* 18(93):606–615 (1968).

[10] L. A. Levin. Average case complete problems. *SIAM J. Comput.* 15:285–286 (1986).

[11] A. Lubiw. Some NP-complete problems similar to graph isomorphism. *SIAM J. Comput.* 10:11–21 (1981).

[12] E. Mead, A. Rosa, C. Huang. The game of SIM: A winning strategy for the second player. *Math. Mag.* 47:243–247 (1974).

[13] A. Pekeč. A winning strategy for the Ramsey graph game. *Comb. Probab. Comput.* 5(3):267–276 (1996).

[14] A. G. Robinson, A. J. Goldman. The Set Coincidence Game: Complexity, Attainability, and Symmetric Strategies. *J. Comput. Syst. Sci.* 39(3):376–387 (1989).

[15] G. J. Simmons. The game of SIM. *J. Recreational Mathematics*, 2(2):66 (1969).

[16] W. Slany. Endgame problems of Sim-like graph Ramsey avoidance games are PSPACE-complete. *Theoretical Computer Science*, to appear.

Frank Harary
Computer Science Department
New Mexico State University
Las Cruces, NM 88003
United States
fnh@crl.nmsu.edu

Wolfgang Slany
Institut für Informationssysteme
Technische Universität Wien
Favoritenstr. 9
A-1040 Wien
Austria
wsi@dbai.tuwien.ac.at

Oleg Verbitsky
Department of Mechanics and Mathematics
Lviv University
Universytetska 1
79000 Lviv
Ukraine
oleg@ov.litech.net