LOWER BOUNDS FOR EXISTENTIAL PEBBLE GAMES AND k -CONSISTENCY TESTS

CHRISTOPH BERKHOLZ

Lehrstuhl für Informatik 7, RWTH Aachen University
e-mail address: berkholz@informatik.rwth-aachen.de

ABSTRACT. The existential k -pebble game characterizes the expressive power of the existential-positive k -variable fragment of first-order logic on finite structures. The winner of the existential k -pebble game on two given finite structures can be determined in time $O(n^{2k})$ by dynamic programming on the graph of game configurations. We show that there is no $O(n^{(k-3)/12})$ -time algorithm that decides which player can win the existential k -pebble game on two given structures. This lower bound is unconditional and does not rely on any complexity-theoretic assumptions.

Establishing strong k -consistency is a well-known heuristic for solving the constraint satisfaction problem (CSP). By the game characterization of Kolaitis and Vardi [14] our result implies that there is no $O(n^{(k-3)/12})$ -time algorithm that decides if strong k -consistency can be established for a given CSP-instance.

1. INTRODUCTION

For two finite relational structures $\mathbf{A}$ and $\mathbf{B}$ the homomorphism problem asks if there is a mapping from the domain A of $\mathbf{A}$ to the domain B of $\mathbf{B}$ that preserves all relations. As pointed out by Feder and Vardi [7] this problem is equivalent to the constraint satisfaction problem (CSP) where the variables correspond to the domain of $\mathbf{A}$, the values correspond to the domain of $\mathbf{B}$ and the constraints are encoded in the relations of $\mathbf{A}$ and $\mathbf{B}$. Thus, every homomorphism from $\mathbf{A}$ to $\mathbf{B}$ corresponds to a solution of the CSP. Since the homomorphism problem and the CSP are NP-complete in general, there is need to look for heuristics. One well-known method introduced in the context of constraint satisfaction is the procedure of establishing strong k -consistency, which can be implemented by an $O(n^{2k})$ -time algorithm (see *e.g.* [4, 14, 2]). We will explain this concept below in the setting of the homomorphism problem.

From a logical point of view, there is a homomorphism between two finite structures $\mathbf{A}$ and $\mathbf{B}$ if and only if every existential-positive first-order sentence that is true on $\mathbf{A}$ is also true on $\mathbf{B}$. Instead of considering the full existential-positive fragment of first-order logic

2012 ACM CCS: [Theory of computation]: Computational complexity and cryptography—Complexity theory and logic.

Key words and phrases: existential pebble game, finite variable logic, first order logic, parameterized complexity theory, k -consistency, constraint propagation.

* This is the full version of the conference paper [3].

one can relax that question and ask whether every k -variable existential-positive first-order sentence true on $\mathbf{A}$ is also true on $\mathbf{B}$. Such questions can be analyzed using combinatorial games. For the k -variable existential-positive fragment of first-order logic the corresponding game is the existential k -pebble game [12], which is defined in Section 2.

Kolaitis and Vardi [14] showed that this logical relaxation of the homomorphism problem is equivalent to the k -consistency heuristic. That is, strong k -consistency can be established if and only if Duplicator wins the existential k -pebble game on $\mathbf{A}$ and $\mathbf{B}$. To sum up, the following three statements are equivalent on finite relational structures $\mathbf{A}$ and $\mathbf{B}$:

- Strong k -consistency can be established.
- For every existential-positive k -variable first-order sentence φ : $\mathbf{A} \models \varphi \implies \mathbf{B} \models \varphi$.
- Duplicator has a winning strategy in the existential k -pebble game.

Now we state our main result that provides a lower bound on the computational complexity of the statements above.

Theorem 1. For every fixed $k \geq 15$ and any $\varepsilon > 0$, the winner of the existential k -pebble game on two given finite relational structures $\mathbf{A}$ and $\mathbf{B}$ cannot be determined in time $O((\|\mathbf{A}\| + \|\mathbf{B}\|)^{\frac{k-2}{12}-\varepsilon})$ on deterministic multi-tape Turing machines.

From this theorem we directly get an $\Omega(n^{\frac{k-2}{12}-\varepsilon})$ lower bound for deciding if strong k -consistency can be established, where n is the size of the CSP-instance. As an upper bound, the query “Does Spoiler win the existential k -pebble game on $\mathbf{A}$ and $\mathbf{B}$?” is LFP^{2k} -definable for two given structures $\mathbf{A}$ and $\mathbf{B}$ [13] and is decidable by an $O(|A|^k|B|^k)$ -time algorithm. We prove Theorem 1 by a reduction from the k -pebble game of Kasai, Adachi and Iwata [10], called KAI-game, to the existential $(k+1)$ -pebble game. Our result then follows from an $n^{\Omega(k)}$ lower bound for this game [1], which in turn follows from the deterministic time hierarchy theorem.

1.1. Related Work. Kolaitis and Panttaja [11] proved that for every fixed $k \geq 2$ the problem of determining the winner of the existential k -pebble game is complete for PTIME under LOGSPACE reductions. Furthermore, they established that the problem is complete for EXPTIME when k is part of the input. It follows that there is no algorithm for this problem whose running time is polynomial in the size of the structures as well as in the number of pebbles. Parameterized by the number of pebbles k , the problem is known to be $\text{W}[1]$ -hard. This follows directly from the fact that a graph G contains a k -clique if and only if Duplicator has a winning strategy for the existential k -pebble game on the complete graph on k vertices and G . Thus, the existence of an algorithm of running time $f(k)n^c$ for some computable function f and constant c would imply $\text{W}[1] = \text{FPT}$, an unlikely event in parameterized complexity theory. However, since we do not know whether $\text{W}[1] = \text{FPT}$ it is consistent with our previous knowledge that there exists an $O(2^k n^2)$ algorithm determining the winner of the existential k -pebble game on two relational structures. Thus, for every fixed k , it was possible that there exists a quadratic time algorithm deciding if strong k -consistency can be established.

To prove the EXPTIME-completeness Kolaitis and Panttaja reduced the KAI-game to the existential pebble game. In this reduction the number of pebbles used in the existential pebble game depends on the size of the KAI-game instance and is not bounded by any function of the number of pebbles used in the KAI-game. Thus, their reduction fails to prove a lower bound for fixed k and it was left as an open question if such a lower bound

can be proven. In this paper we reduce the k -pebble KAI-game to the existential $(k + 1)$ -pebble game, and thus keep the parameter small. Some constructions are quite similar to those used by Kolaitis and Panttaja. However, the proof differs significantly at crucial points. Furthermore, to devise the winning strategy for Duplicator we use a different proof technique.

Finally, the parameterized complexity of k -consistency has also been investigated by Gaspers and Szeider [8]. We discuss their work after the introduction to k -consistency in Section 1.3.

1.2. Further Applications in Finite Model Theory. We can improve Theorem 1 in two ways. First, all partial homomorphisms in Duplicator’s winning strategy (defined below) are in fact partial isomorphisms. Thus, Duplicator has a winning strategy in the k -pebble game that corresponds to the existential k -variable fragment of first-order logic, where negation is allowed in front of atomic formulas. This implies that it requires $\Omega((\|\mathbf{A}\| + \|\mathbf{B}\|)^{\frac{k-2}{12}-\varepsilon})$ time to decide if every existential k -variable first-order sentence true on $\mathbf{A}$ is also true on $\mathbf{B}$. Second, the structures constructed in our reduction are directed graphs. Therefore, Theorem 1 holds even when we restrict ourselves to σ -structures, where σ is a relational signature containing at least one binary relation.

1.3. Establishing Strong k -Consistency. To fix the terminology, we briefly introduce the concept of establishing strong k -consistency as it is defined in [14]. Let $\mathbf{A}$ and $\mathbf{B}$ be two finite relational structures with universes A and B . A *k -partial homomorphism* is a partial homomorphism with domain size k . $\mathbf{A}$ and $\mathbf{B}$ are *k -consistent* if for every $(k - 1)$ -partial homomorphism h from $\mathbf{A}$ to $\mathbf{B}$ and every $a \in A$ there is a partial homomorphism that extends h and is defined on a . Two structures are *strongly k -consistent* if they are *i -consistent* for every $i \leq k$.

Strong k -consistency *can be established* for $\mathbf{A}$ and $\mathbf{B}$ if there are two strongly k -consistent structures $\mathbf{A}'$ and $\mathbf{B}'$ over the same universes A and B such that the following two statements hold.

- Every k -partial homomorphism from $\mathbf{A}'$ to $\mathbf{B}'$ is a k -partial homomorphism from $\mathbf{A}$ to $\mathbf{B}$.
- Every function from A to B is a homomorphism from $\mathbf{A}'$ to $\mathbf{B}'$ if and only if it is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$.

Loosely speaking, strong k -consistency can be established for two structures if they can be made strongly k -consistent by adding new relations and without changing the solution space with respect to the homomorphism problem. It is easy to see that if there is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$, then strong k -consistency can be established. Although the converse is not true in general, it holds for some classes of structures [2, 5].

All known k -consistency algorithms, as *e.g.* [4], iteratively propagate new constraints (relations in our notion) until the instance becomes k -consistent. Gasper and Szeider [8] argued that the task of checking whether the instance is already k -consistent is inherent in this procedure and thus lower bounds its complexity. This motivates them to analyze the following parameterized problem: “Given two finite relational structures and a parameter k , are the structures strongly k -consistent?” They showed that this problem is complete for the parameterized complexity class $\text{co-W}[2]$. Hence, assuming $\text{FPT} \neq \text{co-W}[2]$, the problem is not solvable in time $O(f(k)n^c)$ for some computable function f and constant

c. We analyze the complexity of a stronger statement: “Given two finite relational structures and a parameter k , can strong k -consistency be established?” The outcome of this decision problem matches the outcome of a k -consistency algorithm and thus characterizes the complexity of the k -consistency test precisely. The proof of Theorem 1 implies that establishing strong k -consistency is complete for the parameterized complexity class XP. Thus, assuming $\text{co-W}[2] \neq \text{XP}$, trying to make the instance k -consistent is a harder task than checking whether it is already k -consistent.

2. PEBBLE GAMES AND PROOF OF THEOREM 1

Figure 1: KAI-game: Moving pebble c according to rule (u, v, w, c, d) .

In this section we first introduce the KAI-game and the existential pebble game. Then we state the reduction from the KAI-game to the existential pebble game in our main lemma (Lemma 5) and use it to prove Theorem 1.

Kasai, Adachi and Iwata [10] introduced a simple combinatorial pebble game that nicely simulates Turing machines. The authors showed in [10] and [1] that playing various variants of this game is complete for different complexity classes. Here we stick to the k -pebble variant, that is restricted to a fixed set $[k] := \{1, \dots, k\}$ of pebbles. An instance of the k -pebble KAI-game is a tuple $(X, R, \mathfrak{s}, \gamma)$, where X is the set of nodes, $R = R' \times \{(c, d) \in [k]^2 \mid c \neq d\}$ with $R' \subseteq [X]^3$ the set of rules, $\mathfrak{s}: [k] \rightarrow X$ the start position and $\gamma \in X$ the goal. A rule is of the form (u, v, w, c, d) , with distinct pebbles c, d , pairwise distinct nodes u, v, w and the intended meaning that if pebble c is on u and pebble d is on v and there is no pebble on w then one player can move pebble c from u to w (see Figure 1). This is a slightly more wasteful notion than the original one used in [10], where the relation $R' \subseteq X^3$ (instead of $R' \times \{(c, d) \in [k]^2 \mid c \neq d\}$) is given as input. However, this technical modification does not affect the purpose of the game and increases the size of an instance only by a constant factor if k is fixed, and by a polynomial factor if k is part of the input. A *position* of the KAI-game is an injective mapping $\mathbf{p}: [k] \rightarrow X$. A rule $r = (u, v, w, c, d) \in R$ is *applicable* to a position $\mathbf{p}$ if $\mathbf{p}(c) = u$, $\mathbf{p}(d) = v$ and $\mathbf{p}(z) \neq w$ for all $z \in [k]$. Furthermore, if r is applicable to $\mathbf{p}$ then $r(\mathbf{p})$ denotes the position defined as $r(\mathbf{p})(c) = w$ and $r(\mathbf{p})(z) = \mathbf{p}(z)$, for all $z \in [k] \setminus \{c\}$. The set of all rules in R applicable to a position $\mathbf{p}$ is denoted by $\text{appl}(\mathbf{p})$, and $T_r(\mathbf{p}) \subseteq [k]$ is the set of KAI-pebbles i such that $\mathbf{p}(i)$ contradicts the applicability condition of rule r : $T_{(u,v,w,c,d)}(\mathbf{p}) := \{i \in [k] \mid (i = c \text{ and } \mathbf{p}(i) \neq u) \text{ or } (i = d \text{ and } \mathbf{p}(i) \neq v) \text{ or } \mathbf{p}(i) = w\}$. Thus, $r \in \text{appl}(\mathbf{p})$ iff $T_r(\mathbf{p}) = \emptyset$.

The k -pebble KAI-game is played by two players and proceeds in rounds. In the first round Player 1 starts with position $\mathfrak{s}$ and chooses a rule $r \in \text{appl}(\mathfrak{s})$. The new position is $\mathbf{p} = r(\mathfrak{s})$. In the next round Player 2 chooses a rule $r \in \text{appl}(\mathbf{p})$ and applies it to $\mathbf{p}$. Then it is Player 1's turn and so on. Player 1 wins the game if he reaches a position $\mathbf{p}$, where $\mathbf{p}(z) = \gamma$ for one $z \in [k]$ or where Player 2 is unable to move. Player 2 wins if she has a strategy ensuring that Player 1 cannot reach such a position. The next definition formalizes winning strategies for Player 2. They contain sets of positions $\mathcal{K}_i, i \in \{1, 2\}$ where it is Player i 's turn and a mapping κ that tells Player 2 for every position which rule to choose next.

Definition 2. A *winning strategy* for Player 2 in the KAI-game on $(X, \{r_1, \dots, r_m\}, \mathfrak{s}, \gamma)$ is a triple $\mathcal{K} = (\mathcal{K}_1, \mathcal{K}_2, \kappa)$ where $\mathcal{K}_1 \subseteq \{\mathbf{p} \mid \mathbf{p}: [k] \rightarrow X\}$ and $\mathcal{K}_2 \subseteq \{\mathbf{p} \mid \mathbf{p}: [k] \rightarrow X \setminus \{\gamma\}\}$ are sets of positions and $\kappa: \mathcal{K}_2 \rightarrow [m]$ is a mapping such that the following holds:

- $\mathfrak{s} \in \mathcal{K}_1$.
- For every $\mathbf{p} \in \mathcal{K}_1$ and every $r_i \in \text{appl}(\mathbf{p})$: $r_i(\mathbf{p}) \in \mathcal{K}_2$.
- For every $\mathbf{p} \in \mathcal{K}_2$: $r_{\kappa(\mathbf{p})} \in \text{appl}(\mathbf{p})$ and $r_{\kappa(\mathbf{p})}(\mathbf{p}) \in \mathcal{K}_1$.

Kasai, Adachi and Iwata [10] showed that the problem of determining the winner of the k -pebble KAI-game is PTIME-complete (for every fixed $k \geq 3$) under LOGSPACE-reductions and complete for EXPTIME when k is part of the input. Furthermore, they proved the following unconditional lower bound.

Theorem 3 ([1]). For every fixed $k \geq 6$ and any $\varepsilon > 0$, the winner of the k -pebble KAI-game on a given instance I cannot be determined in time $O(\|I\|^{\frac{k-1}{4}-\varepsilon})$ on deterministic multi-tape Turing machines, where $\|I\|$ is the size of the input I .

The proof of this theorem essentially relies on the deterministic time hierarchy theorem, which states that multi-tape Turing machines of running time n^k cannot be simulated within time $n^{k-\varepsilon}$. On the other hand, Turing machines of running time n^k can be simulated within the $(4k+1)$ -pebble KAI-game and hence the lower bound follows. In terms of parametrized complexity, their argument also leads to XP-completeness of the k -pebble KAI-game with parameter k as pointed out in [6].

The existential $(k+1)$ -pebble game [12] is played by two players Spoiler and Duplicator on two relational structures $\mathbf{A}$ and $\mathbf{B}$ with domains A and B , respectively. First, Spoiler puts pebbles $a_1, \dots, a_{k+1}$ on elements of A and Duplicator answers by putting pebbles $b_1, \dots, b_{k+1}$ on elements of B . In each further round Spoiler picks up a pebble a_i from A and places it on another element in A and Duplicator moves the corresponding pebble b_i in B . Spoiler wins the game if he can reach a position where the mapping defined by $a_i \mapsto b_i$ is not a partial homomorphism from $\mathbf{A}$ to $\mathbf{B}$. Duplicator wins the game if she has a winning strategy that tells her for every move of Spoiler how to place her pebbles such that the positions of the pebbles define a partial homomorphism. A winning strategy for Duplicator can be stated formally as a set of partial homomorphisms:

Definition 4 ([12]). A *winning strategy* for Duplicator in the existential $(k+1)$ -pebble game on structures $\mathbf{A}$ and $\mathbf{B}$ is a nonempty family $\mathcal{H}$ of partial homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ satisfying the following properties:

closure: If $h \in \mathcal{H}$ and $g \subset h$ then $g \in \mathcal{H}$.

extension: For every $g \in \mathcal{H}$, $|\text{Dom}(g)| \leq k$, and every $z \in A$ there is an $h \in \mathcal{H}$ with $g \subseteq h$ and $z \in \text{Dom}(h)$.

For a set H of partial homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ we let $\text{cl}(H) := \{g \mid g \subseteq h, h \in H\}$ be the closure of H under taking subsets and write $\text{cl}(h)$ instead of $\text{cl}(\{h\})$. It is easy to see that if h is a total homomorphism from $\mathbf{A}$ to $\mathbf{B}$, then $\text{cl}(h)$ is a winning strategy in the existential $(k+1)$ -pebble game on $\mathbf{A}$ and $\mathbf{B}$. Now we state our main lemma and prove Theorem 1. The proof of the main lemma is deferred to Section 3.5.

Lemma 5 (Main Lemma). There is a reduction from the k -pebble KAI-game to the existential $(k+1)$ -pebble game that computes for every instance $I = (X, R, \mathfrak{s}, \gamma)$ two directed graphs G_S and G_D such that the following constraints hold:

- Player 1 has a winning strategy in the k -pebble KAI-game on I if and only if Spoiler has a winning strategy in the existential $(k + 1)$ -pebble game on G_S and G_D .
- $|V(G_S)| + |V(G_D)| = O(|X| \cdot |R| \cdot k^2)$.
- $|E(G_S)| + |E(G_D)| = O(k^4(|X|^2|R| + |X| \cdot |R|^2))$.
- The reduction is computable in $\text{DTIME}(O(\|I\|^3))$ and in LOGSPACE .

Proof of Theorem 1. Let $k \geq 15$ be a fixed integer and $\varepsilon > 0$. Assume that $\mathbb{A}$ is an algorithm that determines the winner of the existential k -pebble game on structures $\mathbf{A}$ and $\mathbf{B}$ in time $O((\|\mathbf{A}\| + \|\mathbf{B}\|)^{\frac{k-2}{12}-\varepsilon})$. Let $\mathbb{B}$ be the algorithm that first applies the reduction from Lemma 5 to a given instance I of the $(k - 1)$ -pebble KAI-game and then executes $\mathbb{A}$. Since $\|G_S\| + \|G_D\| = O(\|I\|^3)$, $\mathbb{B}$ has running time $O(\|I\|^3 + \|I\|^{3(\frac{k-2}{12}-\varepsilon)})$, and thus solves the k' -pebble KAI-game in time $O(\|I\|^{\frac{k'-1}{4}-\varepsilon'})$ for $k' = k - 1$ and $\varepsilon' = \varepsilon/3$. This contradicts Theorem 3. $\square$

In addition, Lemma 5 also implies EXPTIME-completeness when k is part of the input and PTIME-completeness for every fixed $k \geq 4$, hence subsumes the result of Kolaitis and Panttaja [11]. Since the reduction is also an fpt-reduction, it follows that determining the winner in the existential k -pebble game is complete for the parameterized complexity class XP.

In our reduction we first construct two colored graphs out of smaller graphs, called gadgets. In order to prove the existence of a winning strategy for one player, we combine strategies for the gadgets to a strategy for the whole graph. The easier part is to do that for Spoiler. As in [9] and [11], we say that Spoiler *can reach* position p_j from position p_i if he has a strategy in the game such that starting from position p_i he wins the game, or position p_j occurs in the game after some finite number of rounds. Since this relation is transitive, we can combine such strategies to show that Spoiler can reach some position p from $\emptyset$; if p does not define a partial homomorphism, this gives us a winning strategy for Spoiler.

For Duplicator this is more difficult. A *critical strategy* in the existential $(k + 1)$ -pebble game is a nonempty family $\mathcal{H}$ of partial homomorphisms satisfying the closure property (Definition 4) together with a set of *critical positions* $\text{crit}(\mathcal{H}) \subset \mathcal{H}$ such that $h \in \text{crit}(\mathcal{H}) \implies |\text{Dom}(h)| = k$ and all $g \in \mathcal{H} \setminus \text{crit}(\mathcal{H})$ satisfy the extension property. A critical strategy is nearly a winning strategy in the sense that Duplicator wins unless the game reaches a critical position. Note that a critical strategy with $\text{crit}(\mathcal{H}) = \emptyset$ is a winning strategy and every critical strategy in the $(k + 1)$ -pebble game is a winning strategy in the k -pebble game. Let $\hat{\mathcal{H}} := \mathcal{H} \setminus \text{crit}(\mathcal{H})$. As for winning strategies, the union of critical strategies is also a critical strategy. The following lemma enables us to construct a winning strategy out of critical strategies.

Lemma 6. If $\mathcal{H}_1, \dots, \mathcal{H}_l$ are critical strategies on the same structures and for all $i \in [l]$ and all $p \in \text{crit}(\mathcal{H}_i)$ there exists a $j \in [l]$ such that $p \in \hat{\mathcal{H}}_j$, then $\bigcup_{i \in [l]} \mathcal{H}_i$ is a winning strategy on these structures. $\square$

Every gadget Q that we construct consists of two graphs Q_S and Q_D for Spoiler's and Duplicator's side, respectively. The graphs contain *boundary vertices* $\text{bd}(Q_S) \subseteq V(Q_S)$ and $\text{bd}(Q_D) \subseteq V(Q_D)$, which are the vertices shared with other gadgets. That is, vertices in $V(Q_S) \setminus \text{bd}(Q_S)$ ($V(Q_D) \setminus \text{bd}(Q_D)$) are only adjacent to vertices in $V(Q_S)$ ($V(Q_D)$). A *boundary function* of a strategy $\mathcal{H}$ on a gadget Q is a mapping $\beta: \text{bd}(Q_S) \rightarrow \text{bd}(Q_D)$ such

that $\beta(z) = h(z)$ for all $h \in \mathcal{H}$ and all $z \in \text{bd}(Q_S) \cap \text{Dom}(h)$. We say that two strategies $\mathcal{G}$ and $\mathcal{H}$ on gadgets Q and Q' are *connectable*, if they have boundary functions $\beta_{\mathcal{G}}$ and $\beta_{\mathcal{H}}$ and it holds that $\beta_{\mathcal{G}}(z) = \beta_{\mathcal{H}}(z)$ for all $z \in \text{bd}(Q_S) \cap \text{bd}(Q'_S)$. If $\mathcal{G}$ and $\mathcal{H}$ are two connectable strategies, we define the *composition*

$$\mathcal{G} \uplus \mathcal{H} = \{g \cup h \mid g \in \mathcal{G}, h \in \mathcal{H}\}.$$

Lemma 7. Let $\mathcal{G}$ and $\mathcal{H}$ be two connectable critical strategies on gadgets $Q = (Q_S, Q_D)$ and $Q' = (Q'_S, Q'_D)$, respectively. The composition $\mathcal{G} \uplus \mathcal{H}$ is a critical strategy on $Q_S \cup Q'_S$ and $Q_D \cup Q'_D$ with $\text{crit}(\mathcal{G} \uplus \mathcal{H}) = \text{crit}(\mathcal{G}) \cup \text{crit}(\mathcal{H})$. $\square$

Playing according to the strategy $\mathcal{G} \uplus \mathcal{H}$ on Q and Q' means that Duplicator uses strategy $\mathcal{G}$ on Q and strategy $\mathcal{H}$ on Q' . The requirements on the boundary ensure that strategy $\mathcal{G}$ equals strategy $\mathcal{H}$ on the intersection of Q and Q' . We use the operator $\uplus$ to construct global critical strategies for the whole graph out of critical strategies on the gadgets. Then we show that the union of those global critical strategies is by Lemma 6 a winning strategy for Duplicator.

3. THE REDUCTION

Let $([n], R, \mathfrak{s}, \gamma)$ be an instance of the k -pebble KAI-game and $m := |R|$. As in [11], the main idea is to simulate every play of the KAI-game within the existential pebble game such that Spoiler imitates the moves of Player 1 and Duplicator imitates the moves of Player 2. First, we construct two colored simple graphs, G_S and G_D , and then show how to omit the colors while switching to directed graphs. We use $|V(G_S)|$ colors to color every vertex of Spoiler's graph G_S differently and partition the vertices of Duplicator's graph with these colors. Thus, whenever Spoiler pebbles a vertex in G_S there is a corresponding set of vertices in G_D Duplicator can pebble.

To encode a position of the KAI-game in the existential $(k+1)$ -pebble game we introduce the vertices $\{x^1, \dots, x^k\}$ in Spoiler's graph and $\{x_l^i \mid i \in [k], 0 \leq l \leq n\}$ in Duplicator's graph. For each i , all the vertices $\{x^i\} \cup \{x_0^i, \dots, x_n^i\}$ are colored with the same unique color, denoted by c_{x^i} . The vertices x_0^i play a special role in the construction, so we draw $\circ$ for vertices with subindex 0 in the figures and $\bullet$ for all other vertices. Furthermore, we introduce vertices y^i, y_l^i in the same way.

If $\mathfrak{p}: [k] \rightarrow [n]$ is a position of the k pebbles in the KAI-game and it is Player 1's turn, then $\{(x^i, x_{\mathfrak{p}(i)}^i) \mid i \in [k]\}$ is the corresponding position in the existential pebble game. If it is Player 2's turn, then $\{(y^i, y_{\mathfrak{p}(i)}^i) \mid i \in [k]\}$ is the corresponding position. During the course of the game Spoiler pebbles some vertex x^i asking, "Where does KAI-pebble i lie?" Due to the coloring Duplicator has to answer with some vertex x_l^i meaning "KAI-pebble i lies on node l ." The vertices $\circ$ are used to handle the case when Spoiler does not play in the intended way, that is, Spoiler has a winning strategy if and only if he has a winning strategy on the $\bullet$ vertices. In order to name positions that include $\circ$ vertices, we define for positions $\mathfrak{p}$ and sets $T \subseteq [k]$ the mapping $(\mathfrak{p}, T)$ as

$$(\mathfrak{p}, T)(i) = \begin{cases} 0, & i \in T, \\ \mathfrak{p}(i), & \text{else,} \end{cases}$$

and write $\mathfrak{p}$ for $(\mathfrak{p}, \emptyset)$ and $\mathbf{0}$ for $(\mathfrak{p}, [k])$. Now we have to introduce gadgets to ensure that Spoiler can simulate a play of the KAI-game. That is, if Player 1 can reach a position $\mathfrak{p}$

Figure 2: The graph G_D . The dotted edges are only for visual purposes and need to be contracted. The gadgets $MS(r_i)$ and $MD(r_i)$ are distinct copies of the switch $M^{k,n}$. The special vertices x_j^i and y_j^i are also depicted in Figure 3.

Figure 3: Vertex blocks to encode positions in the KAI-game.

in the KAI-game, then Spoiler can reach the encoded position on the x - or y -vertices. The following list of properties ensures this.

- For the start position $\mathfrak{s}$, Spoiler can reach $\{(x^i, x_{\mathfrak{s}(i)}^i) \mid i \in [k]\}$ from $\emptyset$.
- For a position $\mathfrak{p}$ and every rule $r \in \text{appl}(\mathfrak{p})$, Spoiler can reach $\{(y^i, y_{r(\mathfrak{p})(i)}^i) \mid i \in [k]\}$ from $\{(x^i, x_{\mathfrak{p}(i)}^i) \mid i \in [k]\}$.
- Spoiler can reach $\{(x^i, x_{r(\mathfrak{p})(i)}^i) \mid i \in [k]\}$ from $\{(y^i, y_{\mathfrak{p}(i)}^i) \mid i \in [k]\}$ for a rule $r \in \text{appl}(\mathfrak{p})$ of Duplicator's choice.

- $\{(y^i, y_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ is not a partial homomorphism if $\mathbf{p}(i) = \gamma$ for some $i \in [k]$.

It follows from these properties that if Player 1 has a winning strategy in the KAI-game, then Spoiler wins the existential pebble game by simulating Player 1's winning strategy. The difficult task is to prove that this is the only way for Spoiler to win. We give a brief description of the construction and argue how Spoiler is intended to play on it. Duplicator's graph is illustrated in Figure 2. The gadgets are glued together at their boundary vertices and the vertex blocks that are glued together inherit their colors. In order to make sure that the colors partition the graphs we define a new color for every combination of colors occurring at one vertex in the graph. In Spoiler's graph we proceed the same way with Spoiler's side of the gadgets.

To implement the last condition, we simply delete the color c_{y^i} from y_γ^i for all $i \in [k]$. Since y^i and all y_l^i , $l \in [n] \setminus \{\gamma\}$, are still colored c_{y^i} , it follows that the mapping $y^i \mapsto y_l^i$ that encodes "KAI-pebble i lies on node l " is a partial homomorphism if and only if l is not the goal node γ . It follows that Spoiler wins the game if he can reach the position $\{(y^i, y_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ where $\mathbf{p}$ is a winning position for Player 1 in the KAI-game.

To make sure that Spoiler can reach the start position, we introduce the initialization gadget INIT^s , whose boundary $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph is identified with vertices $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph. The boundary vertices of the other gadgets have a similar form and can be divided into input vertices x (with certain indices) and output vertices y that are colored in the same way as the x - and y -vertices. As above, a position $\mathbf{p}$ in the KAI-game is encoded as $\{(x^i, x_{\mathbf{p}}^i) \mid i \in [k]\}$ on these vertex blocks and we call it $\mathbf{p}$ on x . The direction of the gadgets is indicated in Figure 2 by arrows. Thus, the players are intended to move clockwise in the graph.

For each rule r we define different rule gadgets $RS(r)$ and $RD(r)$ in which Spoiler can reach the position $r(\mathbf{p})$ on the output y from $\mathbf{p}$ on the input x if r is applicable to $\mathbf{p}$. Hence, from a position $\mathbf{p}$ on x Spoiler can choose an applicable rule r and reach $r(\mathbf{p})$ on the output y of some rule gadget $RS(r)$. The choice gadget C^m enables Duplicator to choose one of the m rules she wants to apply. That is, Duplicator can choose a rule r such that from $\mathbf{p}$ on y Spoiler can reach $\mathbf{p}$ on the input of $RD(r)$ and then $r(\mathbf{p})$ on the output of $RD(r)$. The most complex gadget is the multiple input one-way switch $M^{k,n}$, which is a generalization of the multiple input one-way switch defined in [11]. In our construction we put one copy of $M^{k,n}$ at the output vertices of every rule gadget. Spoiler's strategy on this gadget is nevertheless simple: he can pebble a position through the switch, that is, he can reach $\mathbf{p}$ on the output from $\mathbf{p}$ on the input. This concludes the description of how the gadgets can be used by Spoiler to ensure the four properties above.

Duplicator's strategy is to force Spoiler to play exactly this way. Especially, if the KAI-game does not stop, then Duplicator can play the existential pebble game forever by forcing Spoiler to simulate this infinite play. The main tool for Duplicator is to answer with $\circ$ vertices whenever Spoiler plays incorrectly: if Spoiler pebbles a vertex x^i he is not supposed to pebble now, then Duplicator answers with x_0^i . The strategies on the gadgets ensure that such positions extend to partial homomorphisms and thus Spoiler does not benefit from them.

3.1. Rule Gadgets. The rule gadgets $RS(r)$ and $RD(r)$ consist of input vertices $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph, and output vertices $y^1, \dots, y^k$ and

Figure 4: Rule gadget $RS(u, v, w, c, d)$. The range of i is $[k] \setminus \{c, d\}$. Figure 5: Rule gadget $RD(u, v, w, c, d)$. The range of i is $[k] \setminus \{c, d\}$.

$y_0^1, \dots, y_k^n$. For each rule $r = (u, v, w, c, d)$ we connect the vertices in the gadgets $RS(r)$ and $RD(r)$ as shown in Figure 4 and 5.

If r is applicable to $\mathbf{p}$, then Spoiler can reach the position $\{(y^i, y_{r(\mathbf{p})(i)}^i) \mid i \in [k]\}$ from $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ in both gadgets. To do so, he picks up the remaining pebble and puts it on y^1 . Then he picks up the pebble from x^1 and puts it on y^2 and so on. This fact is stated in Lemma 8(i) and 9(i). Assume that r is not applicable to $\mathbf{p}$ (then $T_r(\mathbf{p}) \neq \emptyset$) and the current position is $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$. On $RS(r)$ Duplicator can pebble some $\circ$ vertex y_0^i when Spoiler asks for y^i ($i \in T_r(\mathbf{p})$), and thus can avoid valid positions on the y -vertices. Therefore, Spoiler is penalized when he chooses a rule r not applicable to $\mathbf{p}$ and plays on $RS(r)$. This strategy is stated in Lemma 8(ii) for $T = \emptyset$. If Duplicator chooses a rule r not applicable to $\mathbf{p}$ and plays on $RD(r)$, then she will be penalized, because Spoiler wins immediately from position $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ (Lemma 9(iii)) by pebbling on y^i for some $i \in T_r(\mathbf{p})$. Furthermore, Lemma 8(ii) and 9(ii) state that for invalid positions on the x -vertices (i.e. $T \neq \emptyset$), Duplicator can avoid valid positions on the y -vertices.

Lemma 8. For every rule $r = (u, v, w, c, d)$ and position $\mathbf{p}: [k] \rightarrow [n]$ the following holds in the existential $(k+1)$ -pebble game on $RS(r)$:

- (i) If $r \in \text{appl}(\mathbf{p})$, then Spoiler can reach $\{(y^i, y_{r(\mathbf{p})(i)}^i) \mid i \in [k]\}$ from $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$.
- (ii) Duplicator has a winning strategy $\mathcal{R}_{(\mathbf{p}, T)}$ with boundary function $\{(x^i, x_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\} \cup \{(y^i, y_{(r(\mathbf{p}), T \cup T_r(\mathbf{p}))(i)}^i) \mid i \in [k]\}$, for all $T \subseteq [k]$.

Proof. For $i = 1, \dots, k$ Spoiler takes the remaining pebble and puts it on y^i . Since there is an edge $\{x^i, y^i\}$, Duplicator has to answer with $y_{r(\mathbf{p})(i)}^i$ because this is the only vertex adjacent to $x_{\mathbf{p}(i)}^i$. In the next step Spoiler picks up the pebble pair from $x^i, x_{\mathbf{p}}^i$ and proceeds with $i+1$.

The boundary function defined in (ii) preserves the vertex colors and maps edges $\{x^i, y^i\}$ to edges $\{x_{(\mathbf{p}, T)(i)}^i, y_{(r(\mathbf{p}), T \cup T_r(\mathbf{p}))(i)}^i\}$, hence defines a total homomorphism on $RS(r)$. It follows that $\mathcal{R}_{(\mathbf{p}, T)} := \{h \mid h \subseteq \beta\}$ is a winning strategy. $\square$

Lemma 9. For every rule $r = (u, v, w, c, d)$ and position $\mathbf{p}: [k] \rightarrow [n]$ the following holds in the existential $(k+1)$ -pebble game on $RD(r)$:

- (i) If $r \in \text{appl}(\mathbf{p})$, then Spoiler can reach $\{(y^i, y_{r(\mathbf{p})(i)}^i) \mid i \in [k]\}$ from $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$.
- (ii) If $r \in \text{appl}(\mathbf{p})$, then Duplicator has a winning strategy $\mathcal{R}_{(\mathbf{p}, T)}$ with boundary function $\{(x^i, x_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\} \cup \{(y^i, y_{(r(\mathbf{p}), T)(i)}^i) \mid i \in [k]\}$, for all $T \subseteq [k]$.
- (iii) If $r \notin \text{appl}(\mathbf{p})$, then Spoiler wins from $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$.

Proof. Statement (i) is analog to Lemma 8(i). The boundary function β stated in (ii) defines a total homomorphism on $RD(r)$. Thus, $\mathcal{R}_{(\mathbf{p}, T)} := \{h \mid h \subseteq \beta\}$ is a winning strategy.

In order to win from $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ for a rule r not applicable to $\mathbf{p}$ (Statement (iii)), Spoiler chooses some KAI-pebble $i \in T_r(\mathbf{p})$ (whose position $\mathbf{p}(i)$ witnesses that r is not applicable to $\mathbf{p}$) and puts the remaining pebble on y^i . Since Duplicator has to answer with a vertex y_j^i of the same color and none of them is adjacent to $x_{\mathbf{p}(i)}^i$, she loses immediately. $\square$

3.2. The Multiple Input One-Way Switch. As for the rule gadgets, the multiple input one-way switch has input vertices $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph, and output vertices $y^1, \dots, y^k$ and $y_0^1, \dots, y_n^k$, respectively. We say that a position $(\mathbf{p}, T)$ is on the input (output) to denote the positions $\{(x^i, x_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\}$ ($\{(y^i, y_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\}$).

For Spoiler the switch ensures that he can reach $\mathbf{p}$ on the output from $\mathbf{p}$ on the input (Lemma 10(i)). For Duplicator there are several strategies. She has a winning strategy called *output strategy*, where any position is on the output and $\mathbf{0}$ is on the input (Lemma 10(ii)). This ensures that Spoiler cannot move backwards and reach $\mathbf{p}$ on the input from $\mathbf{p}$ on the output. Next, for every nonempty $T \subseteq [k]$ Duplicator has a winning strategy where $(\mathbf{p}, T)$ is on the input and $\mathbf{0}$ is on the output (Lemma 10(iii)). Thus, she has a strategy such that Spoiler can reach only the $\mathbf{0}$ position on the output from invalid positions on the input. These strategies are called *restart strategies*. We will see later that Spoiler has to restart the game, that is, he has to pick up all pebbles and start playing on the initialization gadget, if he reaches a position that is contained in a restart strategy. To ensure that Spoiler picks up all pebbles when reaching $\mathbf{p}$ on the output from $\mathbf{p}$ on the input, Duplicator has a critical *input strategy* with $\mathbf{p}$ on the input and $\mathbf{0}$ on the output, whose critical positions are either contained in an output strategy (where $\mathbf{p}$ is on the output) or in a restart strategy (Lemma 10(iv)). If Duplicator plays according to this input strategy, the only way for Spoiler to bring $\mathbf{p}$ from the input to the output is to pebble a critical position inside the switch (using all the pebbles) and force Duplicator to switch to the corresponding output strategy.

In order to define the *multiple input one-way switch* $M^{k,n}$ we construct the two graphs $M_S^{k,n}$ for Spoiler's side and $M_D^{k,n}$ for Duplicator's side. Let

$$\begin{aligned} V(M_S^{k,n}) &= \{x^i, a^i, b^i, y^i \mid i \in [k]\}, \\ E(M_S^{k,n}) &= \{\{x^i, a^i\}, \{b^i, y^i\} \mid i \in [k]\} \\ &\quad \cup \{\{a^i, a^j\}, \{b^i, b^j\} \mid i, j \in [k]; i \neq j\} \\ &\quad \cup \{\{a^i, b^j\} \mid i, j \in [k]\}. \end{aligned}$$

That is, $M_S^{k,n}$ simply consists of k input vertices x^i and k output-vertices y^i , which are each connected to one vertex of a $2k$ -clique. For Duplicator's side of the graph, we define for

Figure 6: Subgraph of the multiple input one-way switch $M^{k,n}$. For the first block in Duplicator's side, all inner-block edges are drawn. Note that there is no edge between $a_{s,l}^i$ and $b_{0,l}^i$.

$i \in [k]$:

$$\begin{aligned}
 X^i &= \{x_s^i \mid 0 \leq s \leq n\}, & Y^i &= \{y_s^i \mid 0 \leq s \leq n\} \\
 A_+^i &= \{a_{s,l}^i \mid s \in [n], l \in [k]\}, & A^i &= A_+^i \cup \{a_0^i\} \\
 B_+^i &= \{b_s^i \mid s \in [n]\}, & B^i &= B_+^i \cup \{b_{0,l}^i \mid l \in [k]\}.
 \end{aligned}$$

The set of vertices of $M_D^{k,n}$ is

$$V(M_D^{k,n}) = \bigcup_{i \in [k]} (X^i \cup A^i \cup B^i \cup Y^i).$$

The graph consists of k blocks, where the i -th block contains the vertices $X^i \cup A^i \cup B^i \cup Y^i$. For every $i \in [k]$ we color $\{x_s^i\} \cup X^i$ as well as $\{a_s^i\} \cup A^i$, $\{b_s^i\} \cup B^i$ and $\{y_s^i\} \cup Y^i$ with a unique color. We first define the inner-block edges E^i , which are also shown in Figure 6, and then the inter-block edges $E^{i,j}$ (for notational convenience we always assume $l, p \in [k]$ and $s, q \in [n]$):

$$E^i = (\{x_0^i\} \times A^i) \cup \{\{x_s^i, a_{s,l}^i\}\} \cup (\{a_0^i\} \times B^i) \cup \quad (E1)-(E3)$$

$$\{\{a_{s,l}^i, b_s^i\}\} \cup \{\{a_{s,l}^i, b_{0,p}^i\} \mid l \neq p\} \cup \{\{b_s^i, y_s^i\}\} \cup \quad (E4)-(E6)$$

$$(\{b_{0,l}^i \mid l \in [k]\} \times Y^i), \quad (E7)$$

$$E^{i,j} = \{\{a_{s,l}^i, a_{q,p}^j\} \mid l \neq p\} \cup \{\{a_0^i, a_{s,l}^j\}\} \cup \{\{b_{0,l}^i, b_{0,p}^j\}\} \cup \quad (E8)-(E10)$$

$$\{\{b_s^i, b_q^j\}\} \cup \{\{a_0^i, b_{0,l}^j\}\} \cup \{\{a_0^i, b_s^j\}\} \cup \quad (E11)-(E13)$$

$$\{\{a_{s,l}^i, b_q^j\}\} \cup \{\{a_{s,l}^i, b_{0,p}^j\} \mid l \neq p\}. \quad (E14)-(E15)$$

Finally, $E(M_D^{k,n}) = \bigcup_{i \in [k]} E^i \cup \bigcup_{i,j \in [k]; i \neq j} E^{i,j}$. Our switch is a further development of the switch used by Kolaitis and Panttaja [11]. In fact, their multiple input one-way switch M^{k+1} is isomorphic to $M^{k,2}$. The next lemma states the main properties of the switch.

Lemma 10. For every position $\mathbf{p}: [k] \rightarrow [n]$, the following statements hold in the existential $(k+1)$ -pebble game on $M^{k,n}$:

- (i) Spoiler can reach $\{(y^1, y_{\mathbf{p}(1)}^1), \dots, (y^k, y_{\mathbf{p}(k)}^k)\}$ from $\{(x^1, x_{\mathbf{p}(1)}^1), \dots, (x^k, x_{\mathbf{p}(k)}^k)\}$.
- (ii) Duplicator has a winning strategy $\mathcal{H}_{(\mathbf{p}, T)}^{\text{out}}$ with boundary function $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ and $\{(y^i, y_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\}$ for all $T \subseteq [k]$.
- (iii) Duplicator has a winning strategy $\mathcal{H}_{(\mathbf{p}, T)}^{\text{restart}}$ with boundary function $\{(x^i, x_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\}$ and $\{(y^i, y_0^i) \mid i \in [k]\}$ for all nonempty $T \subseteq [k]$.
- (iv) Duplicator has a critical strategy $\mathcal{H}_{\mathbf{p}}^{\text{in}}$ with boundary function $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\} \cup \{(y^i, y_0^i) \mid i \in [k]\}$ and sets of critical positions $\mathcal{C}_{\mathbf{p}, t}^{\text{restart-crit}}$ (for $t \in [k]$) and $\mathcal{C}_{\mathbf{p}}^{\text{out-crit}}$ such that:

$$\begin{aligned} \text{a}[\mathbf{p}]^* \text{crit}(\mathcal{H}_{\mathbf{p}}^{\text{in}}) &= \bigcup_{t \in [k]} \mathcal{C}_{\mathbf{p}, t}^{\text{restart-crit}} \cup \mathcal{C}_{\mathbf{p}}^{\text{out-crit}}, \\ \text{a}[\mathbf{p}]^* \mathcal{C}_{\mathbf{p}, t}^{\text{restart-crit}} &\subseteq \mathcal{H}_{(\mathbf{p}, \{t\})}^{\text{restart}} \text{ and} \\ \text{a}[\mathbf{p}]^* \mathcal{C}_{\mathbf{p}}^{\text{out-crit}} &\subseteq \mathcal{H}_{\mathbf{p}}^{\text{out}}. \end{aligned}$$

Proof. Let $\mathbf{p}: [k] \rightarrow [n]$ be an arbitrary position. We first construct the strategy for Spoiler to prove (i). Starting from position $\{(x^1, x_{\mathbf{p}(1)}^1), \dots, (x^k, x_{\mathbf{p}(k)}^k)\}$, Spoiler places the $(k+1)$ st pebble on a^1 . Duplicator has to answer with $a_{\mathbf{p}(1), l_1}^1$ for some $l_1 \in [k]$, mapping the edge $\{x^1, a^1\}$ to some edge in (E2). Next, Spoiler picks up the pebble from x^1 and puts it on a^2 . Again, Duplicator has to answer with $a_{\mathbf{p}(2), l_2}^2$ for some $l_2 \in [k] \setminus \{l_1\}$. The index l_2 has to be different from l_1 because there is an edge between a_1 and a_2 , but none between $a_{\mathbf{p}(1), l_1}^1$ and $a_{\mathbf{p}(2), l_2}^2$ in (E8). Following that scheme, Spoiler can reach the position $\{(a^1, a_{\mathbf{p}(1), l_1}^1), \dots, (a^k, a_{\mathbf{p}(k), l_k}^k)\}$ for pairwise distinct $l_1, l_2, \dots, l_k$. Now, Spoiler pebbles b^1 with the free pebble and Duplicator has to answer with a vertex in B^1 (due to the vertex-colors) that is adjacent to all $a_{\mathbf{p}(1), l_1}^1, \dots, a_{\mathbf{p}(k), l_k}^k$. This is only the case for $b_{\mathbf{p}(1)}^1$ (due to (E4) and (E14)), since every vertex of the form b_{0, l_i}^1 is not adjacent to the vertex $a_{\mathbf{p}(i), l_i}^i$ according to (E5) and (E15). In the next step Spoiler picks up the pebble from $a_{\mathbf{p}(1), l_1}^1$ and puts it on b^2 . Duplicator has to answer with $b_{\mathbf{p}(2)}^2$, mapping the edge $\{b^1, b^2\}$ to (E11). Because every vertex of the form $b_{0, l}^2$ is not connected to $b_{\mathbf{p}(1)}^1$, this is the only choice for Duplicator. Thus, Spoiler can reach $\{(b^1, b_{\mathbf{p}(1)}^1), \dots, (b^k, b_{\mathbf{p}(k)}^k)\}$ and from there he reaches $\{(y^1, y_{\mathbf{p}(1)}^1), \dots, (y^k, y_{\mathbf{p}(k)}^k)\}$ with the same technique.

In order to derive the winning strategies for Duplicator in (ii) and (iii) we consider several total homomorphisms from Spoiler's to Duplicator's side. Consider the edges (E1), (E3) and (E7) connecting $\bullet$ vertices with $\circ$ vertices in one block of Duplicator's side. They can be used by Duplicator to pebble a $\circ$ vertex when Spoiler moves upwards. This is the crucial ingredient for Duplicator's output strategies (ii). To formalize this let $H_{\mathbf{p}}^{\text{out}}$ denote the set of total homomorphisms $h_{\mathbf{p}, \sigma}^{\text{out}} = \{(x^i, x_0^i), (a^i, a_{\mathbf{p}(i), \sigma(i)}^i), (b^i, b_{\mathbf{p}(i)}^i), (y^i, y_{\mathbf{p}(i)}^i) \mid i \in [k]\}$, where $\sigma \in S_k$ is some permutation on $[k]$. Furthermore, let $h_{(\mathbf{p}, T)}^{\text{out}} = \{(x^i, x_0^i), (a^i, a_0^i), (b^i, b_{0, i}^i), (y^i, y_{(\mathbf{p}, T)(i)}^i) \mid i \in [k]\}$. Recall that $\text{cl}(h) := \{g \mid g \subseteq h\}$. Since $h_{(\mathbf{p}, T)}^{\text{out}}$ and all $h_{\mathbf{p}, \sigma}^{\text{out}} \in H_{\mathbf{p}}^{\text{out}}$ are total,

$$\mathcal{H}_{(\mathbf{p}, T)}^{\text{out}} := \begin{cases} \text{cl}(h_{(\mathbf{p}, T)}^{\text{out}}), & T \neq \emptyset, \\ \text{cl}(\{h_{(\mathbf{p}, \emptyset)}^{\text{out}}\} \cup H_{\mathbf{p}}^{\text{out}}), & \text{else,} \end{cases}$$

is a winning strategy for Duplicator satisfying (ii).

If a homomorphism maps all the a^i vertices to A_+^i , then it has to map all b^i vertices to B_+^i . This is due to the missing edges in (E5), (E15) and has also been used in Spoiler's strategy above. On the other hand, if at least one a^i is mapped to a_0^i , then every b^i can be mapped to $b_{0,l}^i$, where l is chosen such that $a_{p(j),l}^j$ is not in the image of the homomorphism for every j . Duplicator benefits from this, because she can now map the y^i vertices arbitrarily using the edges (E7). This behavior is used in the following restart strategies. Note that a homomorphism mapping some a^i to a_0^i also maps x^i to x_0^i , hence restart strategies require invalid input positions. For all nonempty $T \subseteq [k]$, let $\mathcal{H}_{(p,T)}^{\text{restart}} := \text{cl}(H_{(p,T)}^{\text{restart}})$, where $H_{(p,T)}^{\text{restart}}$ is the set of total homomorphisms h satisfying the constraints $h(x^i) = x_{(p,T)(i)}^i$ and $h(y^i) = y_0^i$. This set clearly satisfies (iii). As an example let $g \in H_{(p,\{t\})}^{\text{restart}}$ be the following homomorphism:

$$\begin{aligned} g(x^i) &= x_{(p,\{t\})(i)}^i, & g(b^i) &= b_{0,t}^i, \\ g(a^t) &= a_0^t, & g(y^i) &= y_0^i, \\ g(a^i) &= a_{p(i),i}^i, \quad i \neq t. \end{aligned}$$

It remains to consider the critical input strategies (iv). They formalize the following behavior of Duplicator at the time when Spoiler wants to pebble a position p through the switch as in (i). If Spoiler pebbles a^i or b^i , Duplicator answers within A_+^i or $B^i \setminus B_+^i$, respectively. This allows her to answer on the boundary according to the boundary function defined in (iv). However, she may run into trouble when Spoiler places k pebbles on a^i and b^i vertices, because they extend to a $(k+1)$ -clique on Spoiler's side, but not on Duplicator's side (on the blocks A_+^i and $B^i \setminus B_+^i$). These positions form the critical positions where Duplicator switches to an output or restart strategy. If all k pebbles are on $a_1, \dots, a_k$, as in Spoiler's strategy (i), then Duplicator switches to the output strategy $\mathcal{H}_p^{\text{out}}$. In all other cases she switches to a restart strategy. For all $l \in [k]$ and permutations σ on $[k]$ we define partial homomorphisms:

$$\begin{aligned} h_{p,\sigma}^{\text{in}}(x^i) &= x_{p(i)}^i & h_{p,\sigma,l}^{\text{in}}(x^i) &= x_{p(i)}^i \\ h_{p,\sigma}^{\text{in}}(a^i) &= a_{p(i),\sigma(i)}^i & h_{p,\sigma,l}^{\text{in}}(a^i) &= a_{p(i),\sigma(i)}^i, \quad i \neq \sigma^{-1}(l) \\ h_{p,\sigma}^{\text{in}}(b^i) &= \text{undefined} & h_{p,\sigma,l}^{\text{in}}(b^i) &= b_{0,l}^i \\ h_{p,\sigma}^{\text{in}}(y^i) &= y_0^i & h_{p,\sigma,l}^{\text{in}}(y^i) &= y_0^i \end{aligned}$$

It is easy to check, that $h_{p,\sigma}^{\text{in}}$ defines a homomorphism from $M_S^{k,n} \setminus B$ to $M_D^{k,n}$, and $h_{p,\sigma,l}^{\text{in}}$ defines a homomorphism from $M_S^{k,n} \setminus \{a^{\sigma^{-1}(l)}\}$ to $M_D^{k,n}$. For all $\sigma \in S_k$

$$h_{p,\sigma}^{\text{out-crit}} := \{(a^i, a_{p(i),\sigma(i)}^i) \mid i \in [k]\}$$

and for all $\sigma \in S_k$ and $j, t \in [k]$

$$h_{p,\sigma,j,t}^{\text{restart-crit}} := \{(a^i, a_{p(i),\sigma(i)}^i) \mid i \in [k] \setminus \{t\}\} \cup \{(b^j, b_{0,\sigma(t)}^j)\}.$$

Now we can define the sets used in (iv):

$$\begin{aligned} \mathcal{H}_p^{\text{in}} &= \text{cl}(\{h_{p,\sigma}^{\text{in}} \mid \sigma \in S_k\} \cup \{h_{p,\sigma,l}^{\text{in}} \mid \sigma \in S_k, l \in [k]\}), \\ \mathcal{C}_p^{\text{out-crit}} &= \{h_{p,\sigma}^{\text{out-crit}} \mid \sigma \in S_k\}, \end{aligned}$$

$$\begin{aligned} \mathcal{C}_{p,t}^{\text{restart-crit}} &= \{h_{p,\sigma,j,t}^{\text{restart-crit}} \mid \sigma \in S_k, j \in [k]\}, \\ \text{crit}(\mathcal{H}_p^{\text{in}}) &= \bigcup_{t \in [k]} \mathcal{C}_{p,t}^{\text{restart-crit}} \cup \mathcal{C}_p^{\text{out-crit}}. \end{aligned}$$

First note that $h_{p,\sigma}^{\text{out-crit}} \subset h_{p,\sigma}^{\text{in}}$ and $h_{p,\sigma,j,t}^{\text{restart-crit}} \subset h_{p,\sigma,\sigma(t)}^{\text{in}}$. Thus, $\text{crit}(\mathcal{H}_p^{\text{in}}) \subseteq \mathcal{H}_p^{\text{in}}$. It easily follows from the definitions, that $h_{p,\sigma}^{\text{out-crit}} \subset h_{p,\sigma}^{\text{in}}$. Furthermore, every $h_{p,\sigma,l,t}^{\text{restart-crit}}$ can be extended to a homomorphism in $\mathcal{H}_{(p,\{t\})}^{\text{restart}}$ by defining the boundary as required and mapping a^t to a_0^t and b^j to $b_{0,\sigma(t)}^j$ for all $j \in [k]$. This proves statement b) and c) from (iv). It remains to show that $\mathcal{H}_p^{\text{in}}$ is a critical strategy with critical positions $\text{crit}(\mathcal{H}_p^{\text{in}})$.

Claim 11. For all $g \in \mathcal{H}_p^{\text{in}}$ with $|\text{Dom}(g)| \leq k$, either $g \in \text{crit}(\mathcal{H}_p^{\text{in}})$ or for all $z \in V(M_S^{k,n})$ there exist an $h \in \mathcal{H}_p^{\text{in}}$, such that $g \subseteq h$ and $z \in \text{Dom}(h)$.

Proof. As g is a partial homomorphism from $\mathcal{H}_p^{\text{in}}$, we can fix some $\sigma \in S_k$ and $l \in [k]$ such that $g \subset \{(x^i, x_{p(i)}^i), (a^i, a_{p(i),\sigma(i)}^i), (b^i, b_{0,l}^i), (y^i, y_0^i) \mid i \in [k]\}$.

Case 1: $|\text{Dom}(g) \cap A| = k$. In this case, $g = h_{p,\sigma}^{\text{out-crit}}$ and hence, $g \in \text{crit}(\mathcal{H}_p^{\text{in}})$.

Case 2: $|\text{Dom}(g) \cap A| = k - 1$. If $\text{Dom}(g) \cap B \neq \emptyset$, then $g = h_{p,\sigma,j,\sigma^{-1}(l)}^{\text{restart-crit}}$ for some $j \in [k]$. Thus, we can assume that $\text{Dom}(g) \cap B = \emptyset$ and show for all z that g satisfies the extension property. If z is the unique element in $A \setminus \text{Dom}(g)$, then $h_{p,\sigma}^{\text{in}}$ extends g . If $z \in X \cup B \cup Y$, then $h_{p,\sigma,l}^{\text{in}}$ extends g .

Case 3: $|\text{Dom}(g) \cap A| \leq k - 2$. Let j_1 and j_2 be two distinct indices such that $a^{j_1}, a^{j_2} \notin \text{Dom}(g)$. Furthermore, we can assume that $\sigma(j_1) = l$. For $z \neq a^{j_1}$ the homomorphism $h_{p,\sigma,l}^{\text{in}}$ extends g . If $z = a^{j_1}$, then $h_{p,\sigma',l}^{\text{in}}$ extends g , where $\sigma' := \{(i, \sigma(i)) \mid i \in [k] \setminus \{j_1, j_2\}\} \cup \{(j_1, \sigma(j_2)), (j_2, \sigma(j_1))\}$. □

□

Figure 7: The initialization gadget INIT^5 (for $k = 3$, $n = 4$, $\mathfrak{s}(1) = 2$, $\mathfrak{s}(2) = 4$, $\mathfrak{s}(3) = 3$.)

3.3. The Initialization Gadget. The initialization gadget INIT^s is built out of two multiple input one-way switches M^1 and M^2 , vertices y in Spoiler's graph and y_1, y_2 in Duplicator's graph (that are colored c_y), and the boundary vertices $x^1, \dots, x^k$ and $x_0^1, \dots, x_n^k$ (where all x^i and x_i^j are colored c_{x^i}). The x - and y -vertices are connected to M^1 and M^2 as shown in Figure 7. Lemma 12 (i)–(iii) provides the strategies on INIT^s . The main property is that Spoiler can reach the start position at the boundary (i) and Duplicator has a corresponding counter strategy (ii) in this situation. Furthermore, if an arbitrary position occurs at the boundary during the game, Duplicator has a strategy to survive (iii).

Lemma 12. For every start position $s: [k] \rightarrow [n]$ the following holds in the existential $(k+1)$ -pebble game on INIT^s :

- (i) Spoiler can reach $\{(x^i, x_{s(i)}^i) \mid i \in [k]\}$ from $\emptyset$.
- (ii) There is a winning strategy $\mathcal{I}^{\text{init}}$ for Duplicator with boundary function $\{(x^i, x_{s(i)}^i) \mid i \in [k]\}$.
- (iii) For every $p: [k] \rightarrow [n]$ and every $T \subseteq [k]$ there is a critical strategy $\mathcal{I}_{(p,T)}^{\text{init}}$ with boundary function $\{(x^i, x_{(p,T)(i)}^i) \mid i \in [k]\}$ and $\text{crit}(\mathcal{I}_{(p,T)}^{\text{init}}) \subseteq \mathcal{I}^{\text{init}}$.

Spoiler's strategy is quite simple. First he pebbles y . Duplicator has to answer with either y_1 or y_2 . Then Spoiler can reach $\{(x^i, x_{s(i)}^i) \mid i \in [k]\}$ by pebbling through either M^1 or M^2 . To construct the strategies for Duplicator, we can combine the strategies of the switches M^1 and M^2 such that she plays an input strategy on one switch and a restart or output strategy on the other switch. Assume that Spoiler reaches a critical position on the switch where Duplicator plays the input strategy, say M^1 . Duplicator can now flip the strategies such that she plays a restart or output strategy on M^1 , depending on which kind of critical position Spoiler has reached, and an input strategy on M^2 .

Proof of Lemma 12. We first develop the strategy for Spoiler (i). Spoiler first pebbles y . Duplicator has to response with either y_1 or y_2 . Depending on Duplicator's choice, Spoiler can reach either $\{(a^i, a_{s(i)}^i) \mid i \in [k]\}$ or $\{(b^i, b_{s(i)}^i) \mid i \in [k]\}$. By Lemma 10.(i) Spoiler reaches $\{(c^i, c_{s(i)}^i) \mid i \in [k]\}$ ($\{(d^i, d_{s(i)}^i) \mid i \in [k]\}$) and from there he can reach the position $\{(x^i, x_{s(i)}^i) \mid i \in [k]\}$. For Duplicator's strategies we start with a discussion of possible moves on the boundary of the switches and the x - and y -vertices. At the top of the gadget Duplicator can map y to y_1 and is then forced to answer with s at the input of M^1 and for some $R \subseteq [k]$ with (s, R) at the input of M^2 . This strategy is called $\mathcal{G}_R^1$, the dual strategy where y is mapped to y_2 is called $\mathcal{G}_R^2$.

$$\begin{aligned}\mathcal{G}_R^1 &:= \text{cl}(\{(y, y_1)\} \cup \{(a^i, a_{s(i)}^i), (b^i, b_{(s,R)(i)}^i) \mid i \in [k]\}) \\ \mathcal{G}_R^2 &:= \text{cl}(\{(y, y_2)\} \cup \{(b^i, b_{s(i)}^i), (a^i, a_{(s,R)(i)}^i) \mid i \in [k]\})\end{aligned}$$

At the bottom of the switch $\mathcal{K}_{(p,T)}$ denotes the strategy where $\mathbf{0}$ is at the output of both switches and some arbitrary (p, T) is at the x -block. In the strategy $\mathcal{K}^{\text{out-}i}$ the start position s occurs at the output of the switch M^i and on the x -block, whereas $\mathbf{0}$ is at the output of the other switch.

$$\begin{aligned}\mathcal{K}_{(p,T)} &:= \text{cl}(\{(c^i, c_0^i), (d^i, d_0^i), (x^i, x_{(p,T)(i)}^i) \mid i \in [k]\}) \\ \mathcal{K}^{\text{out-}1} &:= \text{cl}(\{(c^i, c_{s(i)}^i), (d^i, d_0^i), (x^i, x_{s(i)}^i) \mid i \in [k]\}) \\ \mathcal{K}^{\text{out-}2} &:= \text{cl}(\{(c^i, c_0^i), (d^i, d_{s(i)}^i), (x^i, x_{s(i)}^i) \mid i \in [k]\})\end{aligned}$$

Now we can combine these partial strategies with the strategies on the switches described in Lemma 10. In strategy $\mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}i}$ Duplicator plays an input strategy on switch i , a restart strategy on the other switch and an arbitrary position $(\mathbf{p}, T)$ occurs at the x -block. These strategies were combined to the critical strategy $\mathcal{I}_{(\mathbf{p},T)}^{\text{init}}$ described in (ii).

$$\begin{aligned}\mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}1} &:= \mathcal{G}_R^1 \uplus \mathcal{H}_s^{\text{in}} \langle M^1 \rangle \uplus \mathcal{H}_{(s,R)}^{\text{restart}} \langle M^2 \rangle \uplus \mathcal{K}_{(\mathbf{p},T)} \\ \mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}2} &:= \mathcal{G}_R^2 \uplus \mathcal{H}_{(s,R)}^{\text{restart}} \langle M^1 \rangle \uplus \mathcal{H}_s^{\text{in}} \langle M^2 \rangle \uplus \mathcal{K}_{(\mathbf{p},T)} \\ \mathcal{I}_{(\mathbf{p},T)}^{\text{init}} &:= \bigcup_{t \in [k]} (\mathcal{I}_{\{t\},(\mathbf{p},T)}^{\text{in-}1} \cup \mathcal{I}_{\{t\},(\mathbf{p},T)}^{\text{in-}2})\end{aligned}$$

All critical positions of $\mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}i}$ are restart or output critical positions on the switch M^i . By Lemma 10.(iv).(b) every restart critical positions of $\mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}1}$ is contained in one of the strategies $\mathcal{I}_{\{t\},(\mathbf{p},T)}^{\text{in-}2}$ as non-critical position. Hence, the only critical positions $\text{crit}(\mathcal{I}_{(\mathbf{p},T)}^{\text{init}})$ of the combined strategy are output critical positions on the switches. These output critical positions will be contained in the strategies $\mathcal{I}_s^{\text{init-}i}$ where Duplicator plays an output strategy on switch i . Together with $\mathcal{I}_s^{\text{init}}$ they form the winning strategy $\mathcal{I}^{\text{init}}$ from (ii).

$$\begin{aligned}\mathcal{I}^{\text{init-}1} &:= \mathcal{G}_{[k]}^2 \uplus \mathcal{H}_s^{\text{out}} \langle M^1 \rangle \uplus \mathcal{H}_s^{\text{in}} \langle M^2 \rangle \uplus \mathcal{K}^{\text{out-}1} \\ \mathcal{I}^{\text{init-}2} &:= \mathcal{G}_{[k]}^1 \uplus \mathcal{H}_s^{\text{in}} \langle M^1 \rangle \uplus \mathcal{H}_s^{\text{out}} \langle M^2 \rangle \uplus \mathcal{K}^{\text{out-}2} \\ \mathcal{I}^{\text{init}} &:= \mathcal{I}^{\text{init-}1} \cup \mathcal{I}^{\text{init-}2} \cup \mathcal{I}_s^{\text{init}}\end{aligned}$$

$\mathcal{I}^{\text{init}}$ is a union of critical strategies with s at the boundary. To prove that $\mathcal{I}^{\text{init}}$ is indeed a winning strategy on the gadget, we apply Lemma 6 and show that every critical position of one strategy is contained as non-critical position in another strategy. Critical positions are inside the input strategy $\mathcal{H}_s^{\text{in}}$ on one of the switches. By Lemma 10.(iv) they are either contained in an output or restart strategy on the corresponding switch. Hence, all restart critical positions on M^1 and M^2 are contained in $\mathcal{I}_s^{\text{init}}$ and all output critical positions on M^1 (M^2) are contained in $\mathcal{I}^{\text{init-}1}$ ($\mathcal{I}^{\text{init-}2}$). Recall that $\hat{\mathcal{S}} := \mathcal{S} \setminus \text{crit}(\mathcal{S})$, by Lemma 10.(iv) we get:

$$\begin{aligned}\text{crit}(\mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}2}) &= \text{crit}(\mathcal{I}^{\text{init-}1}) = \text{crit}(\mathcal{H}_s^{\text{in}} \langle M^2 \rangle) \subseteq \mathcal{H}_s^{\text{out}} \langle M^2 \rangle \cup \bigcup_{t \in [k]} \mathcal{H}_{(s,\{t\})}^{\text{restart}} \langle M^2 \rangle \\ &\subseteq \hat{\mathcal{I}}^{\text{init-}2} \cup \bigcup_{t \in [k]} \hat{\mathcal{I}}_{\{t\},s}^{\text{in-}1}, \\ \text{crit}(\mathcal{I}_{R,(\mathbf{p},T)}^{\text{in-}1}) &= \text{crit}(\mathcal{I}^{\text{init-}2}) = \text{crit}(\mathcal{H}_s^{\text{in}} \langle M^1 \rangle) \subseteq \mathcal{H}_s^{\text{out}} \langle M^1 \rangle \cup \bigcup_{t \in [k]} \mathcal{H}_{(s,\{t\})}^{\text{restart}} \langle M^1 \rangle \\ &\subseteq \hat{\mathcal{I}}^{\text{init-}1} \cup \bigcup_{t \in [k]} \hat{\mathcal{I}}_{\{t\},s}^{\text{in-}2}.\end{aligned}$$

Hence, $\text{crit}(\mathcal{I}_{(\mathbf{p},T)}^{\text{init}}) \subseteq \mathcal{I}^{\text{init}}$ and $\mathcal{I}^{\text{init}}$ is a winning strategy by Lemma 6. $\square$

3.4. The Choice Gadget. The boundary of the choice gadget consists of input vertices $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph. These vertices are identified with y -vertices in the final graph. The output vertices are of the form $(y_q)^1, \dots, (y_q)^k$ and $(y_q)_0^1, \dots, (y_q)_n^k$ for all $q \in [m]$ and are connected to the rule gadgets $RD(r_q)$. This gadget enables Spoiler to reach positions $((y_q)^i, (y_q)_{\mathbf{p}(i)}^i)$ from $(x^i, x_{\mathbf{p}(i)}^i)$ but Duplicator can choose the desired $q \in [m]$. This choice will later coincide with the rule $r_q \in R$ Player 2 chooses in the KAI-game when position $\mathbf{p}$ is pebbled. The *choice gadget* C^m is defined as follows:

$$\begin{aligned}
V(C_S^m) &= \{x^i, a^i \mid i \in [k]\} \cup \{(y_q)^i \mid i \in [k], q \in [m]\}, \\
E(C_S^m) &= \{\{x^i, a^i\} \mid i \in [k]\} \cup \\
&\quad \{\{a^i, (y_q)^i\} \mid i \in [k], q \in [m]\} \\
&\quad \cup \{\{a^i, a^j\} \mid i, j \in [k]; i \neq j\}, \\
V(C_D^m) &= \{x_l^i \mid i \in [k], 0 \leq l \leq n\} \cup \{a_0^i \mid i \in [k]\} \\
&\quad \{a_{l,q}^i \mid i \in [k], l \in [n], q \in [m]\} \\
&\quad \cup \{(y_q)_l^i \mid i \in [k], q \in [m], 0 \leq l \leq n\}, \\
E(C_D^m) &= \{\{x_0^i, a_0^i\} \mid i \in [k]\} \\
&\quad \cup \{\{x_l^i, a_{l,q}^i\} \mid i \in [k], l \in [n], q \in [m]\} \\
&\quad \cup \{\{a_0^i, (y_q)_0^i\} \mid i \in [k], q \in [m]\} \\
&\quad \cup \{\{a_{l,q}^i, (y_q)_l^i\} \mid i \in [k]; l \in [n]; q \in [m]\} \\
&\quad \cup \{\{a_{l,q}^i, (y_{q'})_0^i\} \mid i \in [k]; l \in [n]; q, q' \in [m]; q \neq q'\} \\
&\quad \cup \{\{a_{l,q}^i, a_{l',q}^j\} \mid i, j \in [k]; i \neq j; l, l' \in [n]; q \in [m]\}.
\end{aligned}$$

Furthermore, for all $i \in [k]$, all vertices a^i and $a_{l,q}^i$ are colored with the unique color c_{a^i} , and for all $i \in [k]$, $q \in [m]$ the vertices $(y_q)^i$ and $(y_q)_l^i$ are colored with the unique color $c_{(y_q)^i}$. One partition of C^m is shown in Figure 8. Recall that Spoiler can reach a position

Figure 8: The i -th partition of the gadget C^m .

if he has a strategy such that either this position occurs after some finite number of rounds or he wins the game. We extend this notion to sets of positions by saying that Spoiler can

reach one of the positions $p_1, \dots, p_m$ from p_0 if starting from position p_0 either Spoiler wins the game or one of the positions $p_1, \dots, p_m$ occurs after a finite number of rounds.

Lemma 13. In the existential $(k+1)$ -pebble game on C^m ,

- (i) for every $\mathbf{p}: [k] \rightarrow [n]$ Spoiler can reach one of the positions $\{((y_l)^i, (y_l)_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ for $l \in [m]$ from $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$, and
- (ii) for every $l \in [m]$, $\mathbf{p}: [k] \rightarrow [n]$, $T \subseteq [k]$, Duplicator has a winning strategy $\mathcal{C}_{(\mathbf{p}, T)}^l$ with boundary function $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\} \cup \{(y_q)^i, (y_q)_0^i \mid i \in [k], q \in [m] \setminus \{l\}\} \cup \{((y_l)^i, (y_l)_{\mathbf{p}(i)}^i) \mid i \in [k]\}$.

Proof. We first present Spoiler's strategy (i). Starting from position $\{(x^i, x_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ Spoiler puts the $(k+1)$ st pebble on a^1 and Duplicator has to response with $a_{\mathbf{p}(1), l}^1$ for one $l \in [m]$ she can choose. Then Spoiler picks up the pebble from x^1 and puts it on a^2 . Duplicator has to response with $a_{\mathbf{p}(2), l}^2$, since this vertex is the only one adjacent to $x_{\mathbf{p}(2)}^2$ and $a_{\mathbf{p}(1), l}^1$. Next, Spoiler picks up the pebble from x^2 and puts it on a^3 and so on. Thus, Spoiler reaches $\{(a^i, a_{\mathbf{p}(i), l}^i) \mid i \in [k]\}$. In the next step he places the $(k+1)$ st pebble on $(y_l)^1$, and Duplicator has to answer with $(y_l)_{\mathbf{p}(1)}^1$, since this vertex is the only one colored in the same color as $(y_l)^1$ and adjacent to $a_{\mathbf{p}(1), l}^1$. Now Spoiler picks up the pebble from a^1 and puts it on $(y_l)^2$, Duplicator has to answer with $(y_l)_{\mathbf{p}(2)}^2$. Following that strategy Spoiler can reach $\{((y_l)^i, (y_l)_{\mathbf{p}(i)}^i) \mid i \in [k]\}$. Duplicator's strategy (ii) is simply defined as $\mathcal{C}_{(\mathbf{p}, T)}^l = \text{cl}(h_{(\mathbf{p}, T)}^l)$, where $h_{(\mathbf{p}, T)}^l$ is the following total homomorphism on C^m :

$$h_{(\mathbf{p}, T)}^l(z) := \begin{cases} x_{(\mathbf{p}, T)(i)}^i, & \text{if } z = x^i, \\ a_{\mathbf{p}(i), l}^i, & \text{if } z = a^i, i \notin T, \\ a_0^i, & \text{if } z = a^i, i \in T, \\ (y_l)_{(\mathbf{p}, T)(i)}^i, & \text{if } z = (y_l)^i, \\ (y_q)_0^i, & \text{if } z = (y_q)^i, q \in [m] \setminus \{l\}. \end{cases}$$

Hence, if some (not necessarily valid) position $(\mathbf{p}, T)$ is on the input Duplicator can force Spoiler to bring this position to the output block (y_l) using the strategy $\mathcal{C}_{(\mathbf{p}, T)}^l$. $\square$

The strategies $\mathcal{C}_0^1 = \dots = \mathcal{C}_0^m$ have the $\mathbf{0}$ position at every vertex block and will be denoted by $\mathcal{C}_0$.

3.5. Proof of the Main Lemma.

Proof of Lemma 5. It remains to construct winning strategies for Spoiler and Duplicator on the colored graphs G_S and G_D . First, we develop the winning strategy for Spoiler. Hence assume that Player 1 has a winning strategy in the k -pebble KAI-game. By Lemma 12, Spoiler can reach position $\{(x^i, x_{\mathbf{s}(i)}^i) \mid i \in [k]\}$. Let r be the first rule Player 1 chooses in the KAI-game. By Lemma 8 Spoiler can reach $\{(y^i, y_{r(\mathbf{s})(i)}^i) \mid i \in [k]\}$, where the y vertices are the output-boundary of the rule gadget $RS(r)$. By Lemma 10 he can pebble through the switch and reach $\{(y^i, y_{r(\mathbf{s})(i)}^i) \mid i \in [k]\}$. Let $\mathbf{p} := r(\mathbf{s})$. If $\mathbf{p}$ maps some pebble to γ , then Spoiler wins the game since y^γ is colored c_{y^γ} whereas y_γ^i is not. Otherwise (by Lemma 13)

Spoiler can reach position $\{((y_q)^i, (y_q)_{\mathbf{p}(i)}^i) \mid i \in [k]\}$ for one $q \in [m]$ of Duplicator's choice at the output of the choice gadget and the input of $\text{RD}(r_q)$. If $r_q \notin \text{appl}(\mathbf{p})$, then Spoiler wins immediately by Lemma 9 (especially Spoiler wins if $\text{appl}(\mathbf{p}) = \emptyset$ and Player 2 cannot move). If $r_q \in \text{appl}(\mathbf{p})$, then Spoiler can reach $\{(x^i, x_{r_q(\mathbf{p})(i)}^i) \mid i \in [k]\}$. Spoiler chooses the next rule according to Player 1's winning strategy and so on. Since Player 1 eventually puts a pebble c on node γ , Spoiler can reach position (y^c, y_γ^c) and thus he wins the game.

Assume that $\mathcal{K} = (\mathcal{K}_1, \mathcal{K}_2, \kappa)$ is a winning strategy for Player 2 in the k -pebble KAI-game. Recall that Player 2 can play in such a way that every position occurring in the KAI-game is either contained in $\mathcal{K}_1$ or $\mathcal{K}_2$ where $\mathcal{K}_i$ is the set of position when it is Player i 's turn. We define a global critical strategy $\mathcal{S}_{\mathbf{p}}(\mathcal{D}_{\mathbf{p}})$ for every position $\mathbf{p}$ in $\mathcal{K}_1$ ($\mathcal{K}_2$). Duplicator can now simulate Player 2's winning strategy in the KAI-game by playing according to the critical strategy $\mathcal{S}_{\mathbf{p}}(\mathcal{D}_{\mathbf{p}})$ if the position $\mathbf{p}$ is the current position in the KAI-game and it is Player 1's (Player 2's) turn. If Spoiler pebbles output critical positions in these strategies, Duplicator switches the strategies in the same way as the positions in the KAI-game change. If Spoiler plays incorrectly in the sense that he pebbles a restart critical position at the switches, then Duplicator moves to a corresponding restart strategy.

Now we construct these critical strategies for the whole graph out of smaller critical strategies $\mathcal{F}$ defined on gadgets Q (denoted $\mathcal{F}\langle Q \rangle$) using the $\uplus$ -operator and Lemma 7. The global strategy $\mathcal{S}^{\text{init}}$ means "the KAI-game has just started, position $\mathbf{s}$ is on the board and it is Player 1's turn." The strategy $\mathcal{S}_{\mathbf{p}}(\mathcal{D}_{\mathbf{p}})$ denotes "position $\mathbf{p}$ is on the board and it is Player 1's (Player 2's) turn."

$$\begin{aligned} \mathcal{S}^{\text{init}} &= \mathcal{I}^{\text{init}} \uplus \mathcal{C}_0 \uplus \biguplus_{l \in \text{appl}(\mathbf{s})} \left(\mathcal{R}_{\mathbf{s}} \langle RS(r_l) \rangle \uplus \mathcal{H}_{r_l(\mathbf{s})}^{\text{in}} \langle MS(r_l) \rangle \right) \uplus \\ &\quad \biguplus_{l \in [m] \setminus \text{appl}(\mathbf{s})} \left(\mathcal{R}_{\mathbf{s}} \langle RS(r_l) \rangle \uplus \mathcal{H}_{(r_l(\mathbf{s}), T_{r_l})}^{\text{restart}} \langle MS(r_l) \rangle \right) \uplus \\ &\quad \biguplus_{l \in [m]} \left(\mathcal{R}_0 \langle RD(r_l) \rangle \uplus \mathcal{H}_{\mathbf{s}}^{\text{out}} \langle MD(r_l) \rangle \right). \end{aligned}$$

We define the global critical strategies $\mathcal{S}_{\mathbf{p}}$ and $\mathcal{S}_{(\mathbf{p}, T)}^{\text{restart}}$ for all $\mathbf{p} \in \mathcal{K}_1$ and $T \neq \emptyset$. In the strategy $\mathcal{S}_{\mathbf{p}}$ the position $\mathbf{p}$ is at the x -vertices encoding that $\mathbf{p}$ is the current position in the KAI-game and it is Player 1's turn. The strategies $\mathcal{S}_{(\mathbf{p}, T)}^{\text{restart}}$ contain the restart critical positions of $\mathcal{S}_{\mathbf{p}}$.

$$\begin{aligned} \mathcal{S}_{\mathbf{p}} &= \mathcal{I}_{\mathbf{p}}^{\text{init}} \uplus \mathcal{C}_0 \uplus \biguplus_{l \in \text{appl}(\mathbf{p})} \left(\mathcal{R}_{\mathbf{p}} \langle RS(r_l) \rangle \uplus \mathcal{H}_{r_l(\mathbf{p})}^{\text{in}} \langle MS(r_l) \rangle \right) \uplus \\ &\quad \biguplus_{l \in [m] \setminus \text{appl}(\mathbf{p})} \left(\mathcal{R}_{\mathbf{p}} \langle RS(r_l) \rangle \uplus \mathcal{H}_{(r_l(\mathbf{p}), T_{r_l})}^{\text{restart}} \langle MS(r_l) \rangle \right) \uplus \\ &\quad \biguplus_{l \in [m]} \left(\mathcal{R}_0 \langle RD(r_l) \rangle \uplus \mathcal{H}_{\mathbf{p}}^{\text{out}} \langle MD(r_l) \rangle \right), \\ \mathcal{S}_{(\mathbf{p}, T)}^{\text{restart}} &= \mathcal{I}_{(\mathbf{p}, T)}^{\text{init}} \uplus \mathcal{C}_0 \uplus \\ &\quad \biguplus_{l \in [m]} \left(\mathcal{R}_{(\mathbf{p}, T)} \langle RS(r_l) \rangle \uplus \mathcal{H}_{(r_l(\mathbf{p}), T \cup T_{r_l})}^{\text{restart}} \langle MS(r_l) \rangle \right) \uplus \\ &\quad \mathcal{R}_0 \langle RD(r_l) \rangle \uplus \mathcal{H}_{(\mathbf{p}, T)}^{\text{out}} \langle MD(r_l) \rangle. \end{aligned}$$

Furthermore, for all $\mathbf{p} \in \mathcal{K}_2$ and $T \neq \emptyset$ let $\mathcal{D}_{\mathbf{p}}$ and $\mathcal{D}_{(\mathbf{p},T)}^{\text{restart}}$ be the following global critical strategies. Similar as in the strategies above, $\mathcal{D}_{\mathbf{p}}$ puts the position $\mathbf{p}$ at the y -vertices encoding that $\mathbf{p}$ is the current position in the KAI-game and it is Player 2's turn. Again, the strategies $\mathcal{D}_{(\mathbf{p},T)}^{\text{restart}}$ contain the restart critical positions of $\mathcal{D}_{\mathbf{p}}$.

$$\begin{aligned} \mathcal{D}_{\mathbf{p}} &= \mathcal{I}_{\mathbf{0}}^{\text{init}} \uplus \mathcal{C}_{\mathbf{p}}^{\kappa(\mathbf{p})} \uplus \biguplus_{l \in [m]} (\mathcal{R}_{\mathbf{0}} \langle RS(r_l) \rangle \uplus \mathcal{H}_{\mathbf{p}}^{\text{out}} \langle MS(r_l) \rangle) \uplus \\ &\quad \biguplus_{l \in [m] \setminus \{\kappa(\mathbf{p})\}} (\mathcal{R}_{\mathbf{0}} \langle RD(r_l) \rangle \uplus \mathcal{H}_{\mathbf{0}}^{\text{restart}} \langle MD(r_l) \rangle) \uplus \\ &\quad \mathcal{R}_{\mathbf{p}} \langle RD(r_{\kappa(\mathbf{p})}) \rangle \uplus \mathcal{H}_{r_{\kappa(\mathbf{p})}(\mathbf{p})}^{\text{in}} \langle MD(r_{\kappa(\mathbf{p})}) \rangle, \\ \mathcal{D}_{(\mathbf{p},T)}^{\text{restart}} &= \mathcal{I}_{\mathbf{0}}^{\text{init}} \uplus \mathcal{C}_{(\mathbf{p},T)}^{\kappa(\mathbf{p})} \uplus \biguplus_{l \in [m]} (\mathcal{R}_{\mathbf{0}} \langle RS(r_l) \rangle \uplus \mathcal{H}_{(\mathbf{p},T)}^{\text{out}} \langle MS(r_l) \rangle) \uplus \\ &\quad \biguplus_{l \in [m] \setminus \{\kappa(\mathbf{p})\}} (\mathcal{R}_{\mathbf{0}} \langle RD(r_l) \rangle \uplus \mathcal{H}_{\mathbf{0}}^{\text{restart}} \langle MD(r_l) \rangle) \uplus \\ &\quad \mathcal{R}_{(\mathbf{p},T)} \langle RD(r_{\kappa(\mathbf{p})}) \rangle \uplus \mathcal{H}_{(r_{\kappa(\mathbf{p})}(\mathbf{p}),T)}^{\text{restart}} \langle MD(r_{\kappa(\mathbf{p})}) \rangle. \end{aligned}$$

Before we formally state Duplicator's winning strategy, we briefly describe these critical strategies. First, the only critical positions of $\mathcal{S}_{(\mathbf{p},T)}^{\text{restart}}$ and $\mathcal{D}_{(\mathbf{p},T)}^{\text{restart}}$ are inside the initialization gadget and contained in $\mathcal{S}^{\text{init}}$. Thus, this is a good situation for Duplicator, since Spoiler has to restart the game by playing on the initialization gadget. At the beginning of the game Duplicator plays according to the strategy $\mathcal{S}^{\text{init}}$, where position $\mathbf{s}$ is on the x -vertices. The only critical positions of that strategy are inside the $MS(r)$ -gadgets for rules r applicable to $\mathbf{s}$. If Spoiler pebbles some restart-critical positions there, then Duplicator can switch to $\mathcal{S}_{(\mathbf{s},T)}^{\text{restart}}$. If Spoiler pebbles an output-critical position on $MS(r)$, then Duplicator can switch to strategy $\mathcal{D}_{r(\mathbf{s})}$ where position $\mathbf{p} = r(\mathbf{s})$ is on the y -vertices. The only critical positions now are inside the switch $MD(r_{\kappa(\mathbf{p})})$ and inside the initialization gadget. If Spoiler pebbles a restart-critical position on $MD(r_{\kappa(\mathbf{p})})$, then Duplicator sticks to $\mathcal{D}_{(\mathbf{p},T)}^{\text{restart}}$. If Spoiler pebbles an output-critical position, then Duplicator chooses strategy $\mathcal{S}_{\mathbf{p}'}$, where $\mathbf{p}' = r_{\kappa(\mathbf{p})}(\mathbf{p})$. The critical positions from $\mathcal{S}_{\mathbf{p}'}$ are within the initialization gadget and the switches $MS(r)$ for rules r applicable to $\mathbf{p}'$. Thus, combining all these critical strategies allows Duplicator to play forever. Now we define the winning strategy for Duplicator:

$$\begin{aligned} \mathcal{H} &= \mathcal{S}^{\text{init}} \cup \bigcup_{\mathbf{p} \in \mathcal{K}_1, T \subseteq [k], T \neq \emptyset} (\mathcal{S}_{\mathbf{p}} \cup \mathcal{S}_{(\mathbf{p},T)}^{\text{restart}}) \\ &\quad \cup \bigcup_{\mathbf{p} \in \mathcal{K}_2, T \subseteq [k], T \neq \emptyset} (\mathcal{D}_{\mathbf{p}} \cup \mathcal{D}_{(\mathbf{p},T)}^{\text{restart}}). \end{aligned}$$

Since $\mathcal{H}$ is a union of critical strategies, it suffices by Lemma 6 to show that for each critical strategy $\mathcal{G}$ and each partial homomorphism $h \in \text{crit}(\mathcal{G})$ there is a critical strategy $\mathcal{F}$ such that $h \in \hat{\mathcal{F}}$. From the definition of the global critical strategies and the properties of the partial critical strategies they contain, it follows that

$$\begin{aligned} \text{crit}(\mathcal{S}_{(\mathbf{p},T)}^{\text{restart}}) &\subseteq \hat{\mathcal{S}}^{\text{init}}, \\ \text{crit}(\mathcal{D}_{(\mathbf{p},T)}^{\text{restart}}) &\subseteq \hat{\mathcal{S}}^{\text{init}}, \end{aligned}$$

$$\begin{aligned}
\text{crit}(\mathcal{S}_{\mathfrak{p}}) &\subseteq \hat{\mathcal{S}}^{\text{init}} \cup \bigcup_{T \neq \emptyset} \hat{\mathcal{S}}^{\text{restart}}_{(\mathfrak{p}, T)} \cup \bigcup_{l \in \text{appl}(\mathfrak{p})} \hat{\mathcal{D}}_{r_l(\mathfrak{p})}, \\
\text{crit}(\mathcal{D}_{\mathfrak{p}}) &\subseteq \hat{\mathcal{S}}^{\text{init}} \cup \hat{\mathcal{S}}_{r_{\kappa(\mathfrak{p})}(\mathfrak{p})} \cup \bigcup_{T \neq \emptyset} \hat{\mathcal{D}}^{\text{restart}}_{(\mathfrak{p}, T)}, \\
\text{crit}(\mathcal{S}^{\text{init}}) &\subseteq \bigcup_{l \in \text{appl}(\mathfrak{s})} \hat{\mathcal{D}}_{r_l(\mathfrak{s})}.
\end{aligned}$$

From the definition of $\mathcal{H}$ and the properties of the KAI-game winning strategy $\mathcal{K}$, it follows, that if a global critical strategy mentioned in the left hand side of the above inclusions is a strategy in $\mathcal{H}$, then so are all strategies on the right hand side. This concludes the proof of Lemma 5 for colored simple graphs G_S and G_D . $\square$

3.6. Getting Rid of the Colors. As in [11] our construction involves $|V(G_S)|$ unary predicates. To settle the complexity of deciding whether Spoiler has a winning strategy in the existential k -pebble game on σ -structures for fixed finite signatures σ , we use the following construction to switch from colored simple graphs to directed graphs. Let $P_1, \dots, P_w$ be the colors used in the graphs G_S and G_D , and P an additional color. We introduce $w + 1$ vertices $d_0, \dots, d_w$ that are colored P in both graphs G_S and G_D . For every $1 \leq i < w$ and each vertex x colored P_i , there are directed edges (d_i, x) and (x, d_{i+1}) . Furthermore, there are directed edges from vertices colored P_w to d_w and one from d_1 to d_0 . Now we can delete the colors $P_1, \dots, P_w$ without giving Duplicator more freedom. That is, we replace the requirement “ $x \in P_i$ ” by the statement “there exists an alternating directed path (where every second element is colored P) of length $2i - 1$ to a vertex colored P and having an out-neighbor colored P ”. This can easily be checked by Spoiler using two pebbles. To get rid of the remaining color P we add a loop $E(x, x)$ for every vertex $x \in P$.

Figure 9: Colors $P_1, \dots, P_w$ can now be deleted.

4. CONCLUSION

We proved an $\Omega(n^{\frac{(k-2)}{12}-\varepsilon})$ lower bound for deciding which player can win the existential k -pebble game on two given finite relational structures and got an $\Omega(n^{\frac{(k-2)}{12}-\varepsilon})$ lower bound for the k -consistency test as a consequence. Furthermore, the lower bound also applies to the k -pebble game that characterizes the expressive power of the existential k -variable fragment (where negation is allowed in front of atomic formulas).

The parameterized complexity of the whole k -variable first-order logic L^k and the counting logic C^k is wide open. It is not even known if it is $W[1]$ -hard to decide if two given finite relational structures can be distinguished by an L^k (C^k) sentence. Regarding the classical complexity, L^k -equivalence as well as C^k -equivalence is complete for polynomial time [9], but it is an open problem whether the problems are complete for EXPTIME when k is part of the input.

REFERENCES

- [1] Akeo Adachi, Shigeki Iwata, and Takumi Kasai. Some combinatorial game problems require $\Omega(n^k)$ time. *J. ACM*, 31, March 1984.
- [2] Albert Atserias, Andrei A. Bulatov, and Víctor Dalmau. On the power of k -consistency. In *Proc. ICALP'07*, pages 279–290, 2007.
- [3] Christoph Berkholz. Lower bounds for existential pebble games and k -consistency tests. In *Proc. LICS'12*, pages 25–34, 2012.
- [4] Martin C. and Cooper. An optimal k -consistency algorithm. *Artificial Intelligence*, 41(1):89 – 95, 1989.
- [5] Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics. In *Proc. CP'02*, pages 310–326, 2002.
- [6] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer-Verlag, 1999.
- [7] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic snp and constraint satisfaction: A study through datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.
- [8] Serge Gaspers and Stefan Szeider. The parameterized complexity of local consistency. In *Proc. CP'11*, pages 302–316, 2011.
- [9] Martin Grohe. Equivalence in finite-variable logics is complete for polynomial time. In *In Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 264–273, 1996.
- [10] Takumi Kasai, Akeo Adachi, and Shigeki Iwata. Classes of pebble games and complete problems. *SIAM J. Comput.*, 8(4):574–586, 1979.
- [11] Phokion G. Kolaitis and Jonathan Panttaja. On the complexity of existential pebble games. In *Proc. CSL'03*, pages 314–329, 2003.
- [12] Phokion G. Kolaitis and Moshe Y. Vardi. On the expressive power of datalog: Tools and a case study. *J. Comput. Syst. Sci.*, 51(1):110–134, 1995.
- [13] Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. *J. Comput. Syst. Sci.*, 61:302–332, October 2000.
- [14] Phokion G. Kolaitis and Moshe Y. Vardi. A game-theoretic approach to constraint satisfaction. In *Proc AAAI/IAAI'00*, pages 175–181, 2000.

ACKNOWLEDGEMENT

I want to thank Martin Grohe and Martin Otto for a fruitful discussion on an earlier version of this paper. Furthermore, I thank the anonymous referees for the useful comments that helped to improve this paper.