# Decidable Relationships between Consistency Notions for Constraint Satisfaction Problems

Albert Atserias[1] and Mark Weyer[2]

[1] Universitat Politècnica de Catalunya, Barcelona, Spain
[2] Humboldt Universität zu Berlin, Berlin, Germany

**Abstract.** We define an abstract pebble game that provides game interpretations for essentially all known consistency algorithms for constraint satisfaction problems including arc-consistency, $(j, k)$-consistency, $k$-consistency, $k$-minimality, and refinements of arc-consistency such as peek arc-consistency and singleton arc-consistency. Our main result is that for any two instances of the abstract pebble game where the first satisfies the additional condition of being stacked, there exists an algorithm to decide whether consistency with respect to the first implies consistency with respect to the second. In particular, there is a decidable criterion to tell whether singleton arc-consistency with respect to a given constraint language implies $k$-consistency with respect to the same constraint language, for any fixed $k$. We also offer a new decidable criterion to tell whether arc-consistency implies satisfiability which pulls from methods in Ramsey theory and looks more amenable to generalization.

## 1 Introduction

Comparing finite structures with respect to some partial order or equivalence relation is a classic theme in logic and algorithms. Notable examples include isomorphisms, embeddings, and homomorphisms, as well as preservation of formulas in various logics, and (bi-)simulation relations of various types.

Let $\leq$ and $\leq'$ be partial orders on finite structures, where $\leq$ is a *refinement* of $\leq'$ which is however *harder* than $\leq'$. More precisely, $\leq$ is a refinement of $\leq'$ in that the implication

$$\mathbf{A} \leq \mathbf{B} \Longrightarrow \mathbf{A} \leq' \mathbf{B} \tag{1}$$

holds true, but the reverse implication is not true in general. Also, $\leq$ is harder than $\leq'$ in the sense that determining whether $\mathbf{A}$ is smaller than $\mathbf{B}$ is computationally harder for $\leq$ than for $\leq'$. A question of interest in this situation is to characterize the structures $\mathbf{A}$ (resp. $\mathbf{B}$) for which the implication in (1) is actually an equivalence. Let us summarize a few known instances where this equivalence holds.

### 1.1 Some examples

For a collection of first-order formulas $L$, we write $\mathbf{A} \leq^L \mathbf{B}$ if every sentence from $L$ that is true in $\mathbf{A}$ is also true in $\mathbf{B}$. Clearly, $\leq^L$ defines a partial order, and if $L$

is closed under negation, it defines an equivalence relation $\equiv^L$. When $L$ is FO, the collection of all first-order formulas, it is well known that $\equiv^L$ agrees with isomorphism on finite structures. On the other hand, the $k$-variable fragment of first-order logic $\mathrm{FO}^k$ gives a coarsening of isomorphism which, for fixed $k$, can be decided in polynomial time. Thus, we are asking for the finite structures $\mathbf{A}$ for which the equivalence

$$\mathbf{A} \equiv^{\mathrm{FO}^k} \mathbf{B} \Longleftrightarrow \mathbf{A} \cong \mathbf{B} \tag{2}$$

holds for every finite $\mathbf{B}$.

If $\mathbf{A}$ is a colored path (a word), then (2) holds for $k \geq 3$ and every $\mathbf{B}$ [16]. More generally, for every colored tree $\mathbf{A}$, equation (2) holds with $k \geq d+1$ and every $\mathbf{B}$, where $d$ is a bound on the degree of the tree. Other fascinating examples arise for graphs embedded in surfaces: if $\mathbf{A}$ is a 3-connected planar graph, (2) holds for every large constant $k$ and every $\mathbf{B}$, and indeed $k \geq 15$ suffices [13, **?**].

When $L$ is $\exists\mathrm{FO}^+$, the existential-positive fragment of FO, the partial order $\mathbf{A} \leq^L \mathbf{B}$ coincides with the existence of a homomorphism from $\mathbf{A}$ into $\mathbf{B}$. Also its $k$-variable fragment $\exists\mathrm{FO}^{k,+}$ gives a coarsening that can be decided in polynomial time. Thus, in this case we are asking for the finite structures $\mathbf{A}$ for which the equivalence

$$\mathbf{A} \leq^{\exists\mathrm{FO}^{+,k}} \mathbf{B} \Longleftrightarrow \mathbf{A} \to \mathbf{B} \tag{3}$$

holds true for every $\mathbf{B}$, where $\mathbf{A} \to \mathbf{B}$ denotes the existence of a homomorphism.

If $\mathbf{A}$ is a colored tree, it is easy to show that (3) holds for $k \geq 2$ and every $\mathbf{B}$. More generally, if the treewidth of $\mathbf{A}$ is less than $k$, even if up to homomorphic equivalence, then (3) holds for every $\mathbf{B}$ [8]. Interestingly, this result is tight: if (3) holds for every $\mathbf{B}$, then the treewidth of $\mathbf{A}$ is less than $k$, up to homomorphic equivalence [2].

For equation (3), restrictions imposed on $\mathbf{B}$ have a different meaning than restrictions imposed on $\mathbf{A}$, and obtaining tight characterizations becomes much harder. Still, some cases are known. For example, if $\mathbf{B}$ is a bipartite graph, then (3) holds for $k \geq 3$ and every $\mathbf{A}$. And on restriction to graphs, this is one of the few instances where we get a characterization: if $\mathbf{B}$ is a graph (with at least one edge) for which (3) holds for every $\mathbf{A}$, then $k \geq 3$ and $\mathbf{B}$ is bipartite [20]. For general relational structures, the state of affairs is much more complicated, as discussed next.

### 1.2 On characterization results and CSPs

Unfortunately, full characterizations as those discussed in the previous section may be hopeless even for natural instances of $\leq$ and $\leq'$. Consider for example the problem that, given $\mathbf{A}$, asks whether the equivalence in (2) holds for every $\mathbf{B}$. For $k = 2$, an algorithm follows from the fact that the finite satisfiability problem for $\mathrm{FO}^2$ is decidable. But for $k = 3$, equivalence with respect to $\mathrm{FO}^k$ is able to encode Diophantine problems [14] and we quickly face undecidability.

For coarser partial orders, such as homomorphism, there is still some hope. For example, the results mentioned above show that the equivalence in (3) holds for every **B** if and only if the treewidth of **A** is less than $k$ up to homomorphic equivalence, which is a decidable criterion. On the other hand, the dual question of characterizing the finite structures **B** for which (3) holds for every **A** is one of the main questions in the seminal work by Feder and Vardi [10] on constraint satisfaction problems. It corresponds to the question of characterizing all constraint languages for which the so-called *k-consistency algorithm* solves any instance. In symbols:

$$k\text{-CON}(\mathbf{B}) \stackrel{?}{=} \text{CSP}(\mathbf{B}) \tag{4}$$

where $k$-CON(**B**) denotes the collection of all instances that are $k$-consistent with respect to **B**, and CSP(**B**) is the collection of all **A** such that $\mathbf{A} \rightarrow \mathbf{B}$.

### 1.3 New results

Motivated by question (4), we offer a unifying approach to the *consistency algorithms* that were considered in the literature. These include arc-consistency, $k, \ell$-consistency, $k$-consistency, $k$-minimality, and refined versions of arc-consistency such as peek arc-consistency and singleton arc-consistency. For pairs of these algorithms, which we denote as partial orders $\leq$ and $\leq'$, we want to be able to decide for which finite structures **B** the equivalence

$$\mathbf{A} \leq \mathbf{B} \Longleftrightarrow \mathbf{A} \leq' \mathbf{B} \tag{5}$$

holds true for every finite **A**. Along the lines of Kolaitis and Vardi [18], we phrase each of these algorithms as an instance of a general pebble game. This abstract setting allows us to prove that the equivalence in (5) is decidable for pairs of algorithms including arc-consistency, peek arc-consistency, singleton arc-consistency, and some others. The simple argument pivots around three components: the fact that such games enjoy treewidth duality, the identification of a subclass of games –called *stacked*– that have definitions in monadic second-order logic, and the decidability of MSO on structures of bounded treewidth. It is worth pointing out, as an interesting feature, that the MSO definitions of stacked games span different levels of the so-called "closure of monadic NP" introduced by Ajtai, Fagin, and Stockmeyer [1]. In particular, they seem to go beyond monadic NP.

One further consequence of these results is that, for a given finite structure **B**, the equality SAC(**B**) = $k$-CON(**B**) is decidable, where SAC(**B**) denotes the collection of all instances that are singleton arc-consistent with respect to **B**. To our knowledge, this sort of result was unknown before. Another remarkable consequence is that a solution to problem (4) automatically gives a solution to problem SAC(**B**) = CSP(**B**), and similarly for other pairs of algorithms.

Finally, we close the paper by offering a new decidable criterion for the problem AC(**B**) = CSP(**B**), where AC(**B**) denotes the instances that are arc-consistent with respect to **B**. Our new proof pulls from ideas in Ramsey theory and looks more amenable to generalization when compared to the previous direct proof by Feder and Vardi [10]. Indeed, our method was introduced by Kolaitis and Vardi [17] for solving a completely different problem related to the asymptotic probability of strict NP properties, which indicates its wider generality.

## 2 Preliminaries

We use standard notation and terminology in finite model theory; see [9]. All our vocabularies are finite and relational, perhaps with additional constant symbols. Homomorphisms preserve tuples and constants, strong homomorphisms preserve also non-tuples, embeddings are injective homomorphisms, and strong embeddings are injective strong homomorphisms. We write $h : \mathbf{A} \to \mathbf{B}$ to denote that $h$ is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. If $h$ does not matter, we write $\mathbf{A} \to \mathbf{B}$ to denote its existence. We use the convention that if $\mathbf{A}$ and $\mathbf{B}$ do not share the same vocabulary, automatically $\mathbf{A} \not\to \mathbf{B}$. The same conventions apply to embeddings $\xrightarrow{e}$, strong embeddings $\xrightarrow{s}$, and isomorphisms $\cong$.

The structure $\mathbf{A}$ is a substructure of $\mathbf{B}$ if $A \subseteq B$ and the identity mapping is an embedding. It is an induced substructure if the embedding is strong. In this case, $\mathbf{A}$ is the substructure of $\mathbf{B}$ induced by $A$, and we denoted it by $\mathbf{B} \upharpoonright A$. The union of $\mathbf{A}$ and $\mathbf{B}$ is the structure with universe $A \cup B$ where the relation $R$ is interpreted by $R^{\mathbf{A}} \cup R^{\mathbf{B}}$. The disjoint union of $\mathbf{A}$ and $\mathbf{B}$ is the union of two copies of $\mathbf{A}$ and $\mathbf{B}$ with disjoint universes. If $C = A \cap B$ and $\mathbf{A}$ and $\mathbf{B}$ agree on $C$ in the sense that $\mathbf{A} \upharpoonright C = \mathbf{B} \upharpoonright C$, the union of $\mathbf{A}$ and $\mathbf{B}$ is called the glued union through $\mathbf{C}$, where $\mathbf{C} = \mathbf{A} \upharpoonright C = \mathbf{B} \upharpoonright C$.

For the definitions of treewidth and tree-decompositions of graphs and relational structures we refer the reader to, say, [11]. We write $\mathrm{TW}(k)$ for the class of all finite structures of treewidth at most $k$.

For the definitions of first and second-order logic, MSO, least and greatest fixed-point logic, and Datalog see [9]. Co-Datalog stands for the negations of Datalog formulae. If $k$ is an integer, $k$-ary Datalog has all recursive predicates of arity at most $k$. An SNP formula is a formula of the form $\exists \overline{X} \forall \overline{x} \varphi$, where $\overline{X}$ is a sequence of relation variables, $\overline{x}$ is a sequence of first-order variables, and $\varphi$ is a quantifier-free formula. A $k$-ary SNP formula has all relation variables of arity at most $k$. The closure of monadic SNP stands for the collection of formulas of the form $\exists \overline{X}_1 \forall \overline{x}_1 \cdots \exists \overline{X}_m \forall \overline{x}_m \varphi$, where all relation variables are unary and $\varphi$ is quantifier-free. The closure of monadic NP was introduced in [1]. It follows from general theory that every formula of $k$-ary co-Datalog is equivalent to a $k$-ary SNP formula, and that every formula of monadic universal greatest fixed-point logic is equivalent to a formula in the closure of monadic SNP.

A *consistency notion* is just any reflexive transitive relation between structures, which is isomorphism invariant. Thus, $\to$, $\xrightarrow{e}$, $\xrightarrow{s}$, and $\cong$ are consistency notions. If $L$ is a logic and $\leq^L$ denotes preservation of $L$-formulas, then $\leq^L$ is also a consistency notion. Let $\leq$ and $\leq'$ be two consistency notions. We say that $\leq$ is a *refinement* of $\leq'$, or $\leq'$ a *coarsening* of $\leq$, if $\mathbf{A} \leq \mathbf{B}$ implies $\mathbf{A} \leq' \mathbf{B}$.

## 3 Generalized pebble game

In this section we define the abstract pebble game for which we prove our results. The methods would work for versions of the game that are even more general, but as this is on the expense of intuition, we have made the definition only as general as necessary to include the important consistency notions in the literature.

Before we start, let us introduce some necessary notation and terminology. Let $\{c_1, c_2, \ldots\}$ be a countable set of constant symbols. For a natural number $k$, a $k$-*numbered structure* is a structure $\mathbf{D}$ for a vocabulary that contains $\{c_1, \ldots, c_k\}$ such that $D = \left\{c_1^{\mathbf{D}}, \ldots, c_k^{\mathbf{D}}\right\}$. Observe, that this does not imply $|D| = k$. For a structure $\mathbf{D}$ and $d_1, \ldots, d_k \in D$, let $(\mathbf{D}, d_1, \ldots, d_k)$ denote the $k$-numbered structure, which is obtained from $\mathbf{D} \restriction \{d_1, \ldots, d_k\}$ by interpreting $c_i$ by $d_i$ for all $1 \leq i \leq k$. All $k$-numbered structures can be represented this way.

### 3.1 Definition of the game

The game comes parameterized by two sets $G$ and $S$. The *growing set* $G$ is a collection of pairs $(k, \mathbf{D})$, where $k$ is a natural number and $\mathbf{D}$ is an $\ell$-numbered structure for some $\ell \geq k$. The *shrinking set* $S$ is a collection of pairs $(k, K)$, such that $k$ is a natural number and $K \subseteq \{1, \ldots, k\}$. Further, we require $G$ to be closed under isomorphisms.

The game $\mathcal{G}^{G,S}$ is played by two players, Spoiler and Duplicator, on a board formed by two structures $\mathbf{A}$ and $\mathbf{B}$. The positions of a play of the game are sequences

$$((a_1, b_1), \ldots, (a_k, b_k)), \tag{6}$$

where $a_i \in A$ and $b_i \in B$. The initial position is the empty sequence. From a position $p$ as in (6), Spoiler has a set of options:

1. *Growing round:* Spoiler may announce a growing round in which he picks some $\ell \geq k$, some $a_{k+1}, \ldots, a_\ell$ from $A$, and an $\ell$-numbered substructure $\mathbf{S}$ of $(\mathbf{A}, a_1, \ldots, a_\ell)$, provided that the pair $(k, \mathbf{S})$ belongs to $G$. Then it is Duplicator's turn, who is required to pick some $b_{k+1}, \ldots, b_\ell$ from $B$ such that $\mathbf{S} \to (\mathbf{B}, b_1, \ldots, b_\ell)$. If she succeeds, the next position is $((a_1, b_1), \ldots, (a_\ell, b_\ell))$; if she does not, the game is over.
2. *Shrinking round:* For every $(k, K)$ in $S$, Spoiler has the option to move to $((a_i, b_i) : i \in K)$, the subsequence of $p$ induced by $K$.

Duplicator wins a play if she can play infinitely. We write $\mathbf{A} \leq^{G,S} \mathbf{B}$ if Duplicator has a winning strategy to win every play of $\mathcal{G}^{G,S}$ on the board formed by $\mathbf{A}$ and $\mathbf{B}$. If $b$ is an integer, we say that the game $\mathcal{G}^{G,S}$ is *grow-bounded* by $b$ if every pair $(k, \mathbf{D})$ in $G$ has $k \leq b$. We say that it is *fully-bounded* by $b$ if every pair $(k, \mathbf{D})$ in $G$ has $|D| \leq b$.

The first thing we need to observe is that our pebble games define relations that are always coarser than homomorphisms:

**Lemma 1.** *Let $G$ an $S$ define a pebble game. Then $\to$ is a refinement of $\leq^{G,S}$, and $\leq^{G,S}$ is reflexive.*

*Proof.* If $h : \mathbf{A} \to \mathbf{B}$, then Duplicator has a winning strategy by answering all growing rounds with $h$. In other words, in position $((a_1, h(a_1)), \ldots, (a_k, h(a_k)))$, if Spoiler picked $a_{k+1}, \ldots, a_\ell$ and a substructure $\mathbf{S}$ of $(\mathbf{A}, a_1, \ldots, a_\ell)$, Duplicator replies with $h(a_{k+1}), \ldots, h(a_\ell)$. Then $\mathbf{S} \to (\mathbf{A}, a_1, \ldots, a_\ell) \to (\mathbf{B}, h(a_1), \ldots, h(a_\ell))$

so this is a valid move. In this way, Duplicator can play infinitely to win. The second claim is immediate from considering the identity homomorphism from **A** to **A**. □

On the other hand, not all $\leq^{G,S}$ are transitive, hence not all of them induce proper consistency notions. All of our examples are transitive, though.

### 3.2 Examples

The $k$-consistency algorithm was studied by Freuder [12]. It can be defined as follows. Let **A** and **B** be two structures. Let $H$ be the collection of all partial homomorphisms $h$ from **A** to **B** such that $|\mathrm{Dom}(h)| \leq k$. For every $h$ in $H$ with $|\mathrm{Dom}(h)| \leq k$ and every $a$ in $A$, if there does not exist any $b$ in $B$ such that $g := h \cup \{(a, b)\}$ is a partial homomorphism from **A** to **B**, for which every $f \subseteq g$ with $|\mathrm{Dom}(f)| \leq k$ belongs to $H$, remove $h$ from $H$, and repeat. Whenever $H$ does not change anymore, stop. If $H$ stabilizes to a non-empty set, we say that **A** is $k$-consistent with respect to **B**. Otherwise we say that it is $k$-inconsistent.

The form we presented of the $k$-consistency algorithm is somewhat closer to the game interpretation given by Kolaitis and Vardi [18]. Our framework, of course, also captures it. We formulate a slightly more general version of it, called $k, \ell$-consistency, that appears in [10] and goes back to [12]:

*Example 1.* Let $k$ and $\ell$ be natural numbers with $0 < k < \ell$. Define

$$G = \{(i, \mathbf{D}) : i \leq k \text{ and } \mathbf{D} \text{ is } j\text{-numbered with } i \leq j \leq \ell\}$$
$$S = \{(i, K) : K \subseteq \{1, \ldots, i\}, \text{ and } |K| \leq k\},$$

Then $\leq^{G,S}$ is called $k, \ell$-consistency and we denote it by $\leq^{k,\ell}$. This game is grow-bounded by $k$ and fully-bounded by $\ell$. The special case $\leq^{k,k+1}$ is called $k$-consistency.

A variant of the $k$-consistency algorithm was introduced by Bulatov, who called it $k$-minimality. Although the differences are minor, specially when the vocabulary is finite, we show how to phrase it in our framework as the particular case of 1-minimality is a very well-known algorithm called arc-consistency.

Before we phrase the $k$-minimality algorithm in game-theoretic terms, let us present the algorithmic view of arc-consistency. Let **A** and **B** be two structures. The algorithm maintains a set $S_a \subseteq B$ for every $a \in A$, initially set to $B$. For every $a \in A$, every $b \in S_a$, every relation symbol $R$, and every $(a_1, \ldots, a_r) \in R^{\mathbf{A}}$, if there does not exist any $(b_1, \ldots, b_r) \in R^{\mathbf{B}}$ such that, for every $j \in \{1, \ldots, r\}$, it holds that $b_j \in S_{a_j}$, and $b_j = b$ whenever $a_j = a$, remove $b$ from $S_a$, and repeat. Whenever the $S_a$'s do not change anymore, stop. If the $S_a$'s stabilize to non-empty sets, we say that **A** is arc-consistent with respect to **B**. Otherwise we say that it is arc-inconsistent. The algorithmic version of $k$-minimality is a straightforward generalization of this algorithm that maintains relations of arity at most $k$.

In order to define $k$-minimality in game terms, we need the following additional concept. Let $i \leq j$. An $i, j$*-tuple structure* is a $j$-numbered structure $\mathbf{D}$ such that $R^{\mathbf{D}} \neq \emptyset$ holds for exactly one $R$ in the vocabulary, such that $R^{\mathbf{D}}$ has exactly one tuple $(d_1, \ldots, d_r)$, and such that $\{d_1, \ldots, d_r\} = \left\{c_i^{\mathbf{D}}, \ldots, c_j^{\mathbf{D}}\right\}$.

*Example 2.* Let $k$ be a natural number with $k > 0$. Define

$$
\begin{aligned}
G = \{(i, \mathbf{D}) : i \leq k \text{ and } \mathbf{D} \text{ is } j\text{-numbered with } i \leq j \leq k\} \cup \\
\{(k, \mathbf{D}) : \mathbf{D} \text{ is a } 1, j\text{-tuple structure with } j \geq k\} \\
S = \{(i, K) : \ K \subseteq \{1, \ldots, i\} \text{ and } |K| \leq k\},
\end{aligned}
$$

Then $\leq^{G,S}$ is called $k$-minimality and we denote it by $\leq^{k\text{-MIN}}$. This game is grow-bounded by $k$, and fully-bounded by the maximum arity of the relation symbols in the vocabulary. The particular case $\leq^{1\text{-MIN}}$ is called arc-consistency and we denote it by $\leq^{\text{AC}}$.

We discuss two more examples that will re-appear later in the paper. These are refinements of arc-consistency that have been studied in the literature, sometimes interchangeably. The first refinement is called peek arc-consistency in [6]. In algorithmic form, this stands for the procedure that, for every $a \in A$, checks if there exists some $b \in B$ for which the arc-consistency algorithm started with $S_a = \{b\}$, and $S_{a'} = B$ for $a' \neq a$, stabilizes with non-empty sets. As a game, this is phrased as follows:

*Example 3.* Define

$$
\begin{aligned}
G = \{(i, \mathbf{D}) : \mathbf{D} \text{ is } j\text{-numbered with } i \leq j \leq 2\} \cup \\
\{(2, \mathbf{D}) : \mathbf{D} \text{ is a } 2, j\text{-tuple structure with } j \geq 2\} \\
S = \{(i, \{1, j\}) : 2 \leq j \leq i\},
\end{aligned}
$$

Then $\leq^{G,S}$ is called peek arc-consistency and we denote it by $\leq^{\text{PAC}}$. This game is grow-bounded by 2 and fully-bounded by the maximum arity of the relation symbols in the vocabulary plus one.

The last example is singleton arc-consistency [5]. Algorithmically, we maintain sets $T_a \subseteq B$, initially set to $B$, and for every $a \in A$ and every $b \in T_a$ check whether arc-consistency started with $S_a = \{b\}$, and $S_{a'} = T_{a'}$ for $a' \neq a$, stabilizes with non-empty sets. If it does not, we remove $b$ from $T_a$, and repeat. Game-theoretically, here is how this is defined:

*Example 4.* Define

$$
\begin{aligned}
G = \{(i, \mathbf{D}) : \mathbf{D} \text{ is } j\text{-numbered with } i \leq j \leq 2\} \cup \\
\{(2, \mathbf{D}) : \mathbf{D} \text{ is a } 2, j\text{-tuple structure with } j \geq 2\} \\
S = \{(i, \{j\}) : 1 \leq j \leq i\} \cup \{(i, \{1, j\}) : 2 \leq j \leq i\},
\end{aligned}
$$

Then $\leq^{G,S}$ is called singleton-arc-consistency and we denote it by $\leq^{\text{SAC}}$. Again, this game is grow-bounded by 2 and fully-bounded by the maximum arity of the relation symbols in the vocabulary plus one.

### 3.3 Definability

We turn now to definability. We say that $\leq$ induces a consistency notion that is *definable* in some logic if, for every finite structure $\mathbf{B}$, there exists a formula $\varphi$ in the logic, such that for every finite structure $\mathbf{A}$ we have $\mathbf{A} \leq \mathbf{B}$ iff $\mathbf{A} \models \varphi$. We say that the definition is *effective* if furthermore such a $\varphi$ can be computed from $\mathbf{B}$. The following is well-known:

**Lemma 2.** *$\rightarrow$ induces a consistency notion that is is definable in monadic* SNP*. Furthermore, the definition is effective.*

For the general pebble game, it does not seem possible to stay within monadic SNP, not even monadic second-order logic. However, standard methods give the following:

**Lemma 3.** *Let $G$ and $S$ define a pebble game that is grow-bounded by $k$ and fully-bounded. Then, $\leq^{G,S}$ induces a consistency notion that is definable in $k$-ary co-Datalog and hence in $k$-ary* SNP*. Furthermore, if $G$ and $S$ are decidable, the definition is effective.*

Note, in particular, that $\leq^{\mathrm{AC}}$ induces a consistency notion that is definable in monadic co-Datalog and hence in monadic SNP. According to this lemma, $\leq^{\mathrm{PAC}}$ and $\leq^{\mathrm{SAC}}$ induce consistency notions that are definable in binary co-Datalog and binary SNP, as they are both grow-bounded by 2. We will show more as both notions are definable in a monadic fragment of second-order logic. This will follow from a general condition on pebble games that we define next.

Let $G$ and $S$ define a pebble game. The game is called *stacked* if for every $(k, K)$ in $S$ there exist $0 \leq i, j \leq k$ such that $K \setminus \{j\} = \{1, \ldots, i\}$. Arc-consistency, $1, \ell$-consistency, peek-arc-consistency, and singleton-arc-consistency are all stacked. Note also that these examples are grow-bounded by 2 but, in general, stacked pebble games need not be grow-bounded by any fixed $k$. Thus, the following result, which is the main result of this section, gets interesting when compared to Lemma 3.

**Lemma 4.** *Let $G$ and $S$ define a fully-bounded stacked pebble game. Then, $\leq^{G,S}$ induces a consistency notion that is definable in monadic universal greatest fixed-point logic and hence in the closure of monadic* SNP*. Furthermore, if $G$ and $S$ are decidable, the definition is effective.*

*Proof (rough sketch).* The construction makes heavy use of nested and simultaneous fixed points. Intuitively, the nesting levels correspond to the different $i$ for shrinking sets of the form $\{1, \ldots, i, j\}$. In this setting, only the $j$ are able to change entries in positions, which leads to monadic fixed points being sufficient.

The full proof is rather technical and can be found in Appendix A. □

### 3.4 Treewidth duality

Let $\mathcal{C}$ be a class of structures. The binary relation $\leq$ has $\mathcal{C}$-*duality*, if for every pair of structures $\mathbf{A}$ and $\mathbf{B}$ such that $\mathbf{A} \not\leq \mathbf{B}$, there exists a $\mathbf{C} \in \mathcal{C}$ such that $\mathbf{C} \leq \mathbf{A}$ and $\mathbf{C} \not\leq \mathbf{B}$. We say that $\leq$ has *treewidth-$k$-duality*, if it has $\mathcal{C}$-duality for some $\mathcal{C} \subseteq \mathrm{TW}(k)$.

Observe that this differs in formulation from the duality used in [10]. There, it is established that for a certain $\leq$, namely $k$-consistency, $\mathbf{A} \not\leq \mathbf{B}$ implies the existence of $\mathbf{C}$ in $\mathrm{TW}(k)$, such that $\mathbf{C} \to \mathbf{A}$ and $\mathbf{C} \not\to \mathbf{B}$. As for this $\leq$ it turns out that $\mathbf{C} \leq \mathbf{D}$ and $\mathbf{C} \to \mathbf{D}$ are equivalent for every $\mathbf{C}$ in $\mathrm{TW}(k)$ and every structure $\mathbf{D}$, the different formulations amount to the same. However, this equivalence does not carry over to other $\leq$. In particular, the following result does not follow directly from the fact that every bounded pebble game induces a consistency notion that is definable in co-Datalog. It requires its own proof.

**Lemma 5.** *Let $G$ and $S$ define a pebble game fully-bounded by $k$. Then $\leq^{G,S}$ has treewidth-$(k-1)$-duality.*

*Proof (rough sketch).* For the structure $\mathbf{C}$, we use an unravelling of $\mathbf{A}$. The actual notion of unravelling is game specific, but it coincides with the usual one for the $k$-consistency game. In order to obtain a finite $\mathbf{C}$, we truncate the unravelling tree at a depth which is large enough to contain all relevant moves in a game between $\mathbf{A}$ and $\mathbf{B}$. The full proof can be found in Appendix B. □

## 4 Application: Decidable relative consistency results

The *relative consistency problem* for $\leq$ and $\leq'$ is the problem of, given some finite structure $\mathbf{B}$, deciding whether the implication

$$\mathbf{A} \leq \mathbf{B} \Longrightarrow \mathbf{A} \leq' \mathbf{B} \tag{7}$$

holds for every finite $\mathbf{A}$. A simple application of the decidability of the satisfiability problem for MSO on structures bounded treewidth (see [11]) gives:

**Theorem 1.** *Let $\leq$ and $\leq'$ induce consistency notions such that:*

1. *both notions are effectively definable in* MSO.
2. *$\leq'$ has treewidth duality.*
3. *$\leq'$ is a refinement of $\leq$.*

*Then, the relative consistency problem for $\leq$ and $\leq'$ is decidable.*

*Proof.* Let $\mathbf{B}$ be given and let $k$ be such that $\leq'$ has $\mathrm{TW}(k)$ duality. Let $\varphi$ be an MSO definition of the consistency notion induced by $\leq$ and $\mathbf{B}$, and let $\varphi'$ be the one for $\leq'$ and $\mathbf{B}$. We will show that the implication (7) fails if and only if there exists some $\mathbf{C}$ in $\mathrm{TW}(k)$ such that $\mathbf{C} \leq \mathbf{B}$ and $\mathbf{C} \not\leq' \mathbf{B}$. This last condition is equivalent to the satisfiability of $\varphi \wedge \neg\varphi'$ in $\mathrm{TW}(k)$, which is decidable.

The 'if' part of the claim is immediate. For the 'only if' part, let $\mathbf{A}$ be such that $\mathbf{A} \leq \mathbf{B}$ and $\mathbf{A} \not\leq' \mathbf{B}$. Using duality, let $\mathbf{C}$ in $\mathrm{TW}(k)$ be such that $\mathbf{C} \leq' \mathbf{A}$ and $\mathbf{C} \not\leq' \mathbf{B}$. From $\mathbf{C} \leq' \mathbf{A}$ we obtain $\mathbf{C} \leq \mathbf{A}$ because $\leq'$ is a refinement of $\leq$, and then $\mathbf{C} \leq \mathbf{B}$ using transitivity. □

Even though $k$-consistency induces a consistency notion that is probably not definable in MSO when $k > 1$, we can still prove the following result:

**Theorem 2.** *Let $k \geq \mathrm{ar}(\sigma)$, and let $\leq$ induce a consistency notion such that:*

1. *the notion is effectively definable in* MSO.
2. *$\leq^k$ is a refinement of $\leq$.*

*Then, the relative consistency problem for $\leq$ and $\leq^k$ is decidable.*

*Proof.* The proof follows the same lines as in Theorem 1, using $\mathrm{TW}(k)$ duality of $\leq^k$. It remains to replace MSO definability of $\leq^k$. For this purpose, note that the previous proof only needed the formula $\varphi'$ to describe whether $\mathbf{A} \leq^k \mathbf{B}$ for $\mathbf{A}$ in $\mathrm{TW}(k)$, not for all finite structures. Recall that $\leq^k$ and $\rightarrow$ coincide on $\mathrm{TW}(k)$ (see [8]), and that $\rightarrow$ is definable in MSO. This is all we need. $\square$

It is also possible to generalize this result to any consistency notion defined by a fully-bounded game replacing $\leq^k$. This requires some additional techniques, including a notion of *game treewidth*, which will appear in the full version of the paper.

As a consequence of all the above, we have the following:

**Corollary 1.** *The following relative consistency problems are decidable:*

1. *Arc-consistency and peek arc-consistency.*
2. *Peek arc-consistency and singleton arc-consistency.*
3. *Singleton arc-consistency and $k$-consistency for $k \geq \mathrm{ar}(\sigma)$.*
4. *$1,\ell$-consistency and $1,\ell'$-consistency for $\ell \leq \ell'$.*
5. *$1,\ell$-consistency and $k$-consistency for $k \geq \max(\ell - 1, \mathrm{ar}(\sigma))$.*
6. *Transitive combinations of the above.*

As a side note we give a further result, which in particular implies that the injective variant of the relative consistency problem for $\leq^k$ and $\rightarrow$ is decidable.

**Theorem 3.** *Let $\leq$ induce a consistency notion such that:*

1. *$\leq$ is decidable.*
2. *$\xrightarrow{e}$ is a refinement of $\leq$.*

*Then, the relative consistency problem for $\leq$ and $\xrightarrow{e}$ is decidable.*

*Proof.* Although a more elementary presentation would be possible, we proceed along the lines of the previous proofs. First, $\xrightarrow{e}$ has size-plus-one duality: If $\mathbf{A} \stackrel{e}{\not\rightarrow} \mathbf{B}$, then there is some $\mathbf{C}$ such that $\mathbf{C} \xrightarrow{e} \mathbf{A}$, $\mathbf{C} \stackrel{e}{\not\rightarrow} \mathbf{B}$, and $|C| \leq |B| + 1$: If $|A| \leq |B| + 1$, we let $\mathbf{C} := \mathbf{A}$. Otherwise, let $C$ be any subset of $A$ of size $|B| + 1$ and let $\mathbf{C} = \mathbf{A} \upharpoonright C$. Then $\mathrm{id} : \mathbf{C} \xrightarrow{e} \mathbf{A}$ and $\mathbf{C} \stackrel{e}{\not\rightarrow} \mathbf{B}$ follows from injectivity.

Hence, we only need to decide whether $\mathbf{A} \leq \mathbf{B}$ implies $\mathbf{A} \xrightarrow{e} \mathbf{B}$ for all $\mathbf{A}$ with $|A| \leq |B| + 1$. As $\leq$ is decidable, this can be solved by brute force. $\square$

## 5 Decidable criterion for arc-consistency

In this section we concentrate on arc-consistency. We want to be able to decide, for a given $\mathbf{B}$, whether $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$ implies $\mathbf{A} \to \mathbf{B}$ for every finite $\mathbf{A}$. Put differently, we want to detect if there exists some *counterexample*: a finite $\mathbf{A}$ such that $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$ and yet $\mathbf{A} \not\to \mathbf{B}$. Since our goal is to build an arc-consistent instance, we start by developing the closure properties of this class of structures.

For the rest of this section, we fix a finite structure $\mathbf{B}$ with vocabulary $\sigma$. For every $b$ in $B$, let $P_b$ be a new unary relation symbol and let $\tau$ be $\sigma \cup \{P_b : b \in B\}$. For every $\mathbf{A}$ such that $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$, let $\mathcal{W}(\mathbf{A})$ be the collection of all $\tau$-expansions of $\mathbf{A}$ whose interpretations for $\{P_b : b \in B\}$ satisfy the following conditions:

i. for every $a \in A$ there exists some $b \in B$ such that $a \in P_b^{\mathbf{A}}$,
ii. for every $a \in A$, every $b \in B$ such that $a \in P_b^{\mathbf{A}}$, every $R \in \sigma$, and every $(a_1, \ldots, a_r) \in R^{\mathbf{A}}$, there exists $(b_1, \ldots, b_r) \in R^{\mathbf{B}}$ such that, for every $j \in \{1, \ldots, r\}$, it holds that $a_j \in P_{b_j}^{\mathbf{A}}$, and $b_j = b$ whenever $a_j = a$.

The condition $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$ is equivalent to the statement that $\mathcal{W}(\mathbf{A})$ is non-empty. Alternatively, we could have taken this as our definition of $\leq^{\mathrm{AC}}$. Let $\mathcal{W}$ denote the union of all $\mathcal{W}(\mathbf{A})$ as $\mathbf{A}$ ranges over all finite structures such that $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$.

**Lemma 6.** *$\mathcal{W}$ is closed under induced substructures and glued unions.*

*Proof.* Closure under induced substructures is immediate since the conditions i. and ii. defining $\mathcal{W}(\mathbf{A})$ are universal on $\mathbf{A}$. We concentrate on glued unions. Let $\mathbf{A}_1$ and $\mathbf{A}_2$ be two structures in $\mathcal{W}$ that agree on the common part: that is, for $A_0 = A_1 \cap A_2$, we have $\mathbf{A}_1 \upharpoonright A_0 = \mathbf{A}_2 \upharpoonright A_0$. Let $\mathbf{A}_3$ be the glued union of $\mathbf{A}_1$ and $\mathbf{A}_2$, and let $\mathbf{A}$ be its $\sigma$-reduct. We claim that the sets $\{P_b^{\mathbf{A}_3} : b \in B\}$ satisfy the conditions i. and ii. that define $\mathcal{W}(\mathbf{A})$. This will show that $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$ and at the same time put $\mathbf{A}_3$ in $\mathcal{W}(\mathbf{A})$ and $\mathcal{W}$.

For condition i., fix an element $a \in A_3$. If $a \in A_k$ for $k \in \{1, 2\}$, then there exists some $b \in B$ such that $a \in P_b^{\mathbf{A}_k}$ and the same $b$ serves for $\mathbf{A}_3$. The fact that $\mathbf{A}_1$ and $\mathbf{A}_2$ agree on $A_0$ guarantees that this is well-defined. For condition ii., fix $a \in A_3$, $b \in B$, $R \in \sigma$, and $(a_1, \ldots, a_r) \in R^{\mathbf{A}}$ as in its statement. Since $(a_1, \ldots, a_r)$ belongs to $R^{\mathbf{A}}$ and $\mathbf{A}$ is also the glued union of the $\sigma$-reducts of $\mathbf{A}_1$ and $\mathbf{A}_2$, it must necessarily be the case that either $\{a_1, \ldots, a_r\} \subseteq A_1$ or $\{a_1, \ldots, a_r\} \subseteq A_2$. In case $\{a_1, \ldots, a_r\} \subseteq A_k$ for $k \in \{1, 2\}$, let $b_1, \ldots, b_r \in B$ be given by condition ii. on $\mathbf{A}_k$. Again the fact that $\mathbf{A}_1$ and $\mathbf{A}_2$ agree on $A_0$ guarantees that this choice is well-defined and valid for $\mathbf{A}_3$. □

It follows from the lemma that $\mathcal{W}$ is an amalgamation class and, by Fraïssé's construction (see [15]), there exists a countably infinite structure $\mathbf{S}^+$ satisfying the following three properties:

1. every finite induced substructure of $\mathbf{S}^+$ is isomorphic to a structure in $\mathcal{W}$,
2. every structure in $\mathcal{W}$ is isomorphic to a finite induced substructure of $\mathbf{S}^+$,
3. for every two finite subsets $S_1$ and $S_2$ of $S^+$, if $\mathbf{S}^+ \upharpoonright S_1$ and $\mathbf{S}^+ \upharpoonright S_2$ are isomorphic, then there exists an automorphism of $\mathbf{S}^+$ that maps $S_1$ to $S_2$.

From now on, we write $\mathbf{S}$ for the $\sigma$-reduct of $\mathbf{S}^+$. Except for the fact that it is infinite, $\mathbf{S}$ is the candidate counterexample we are looking for. To establish this, the first and second properties of $\mathbf{S}^+$ will suffice; the third property will be discussed later on. We start showing that $\mathbf{S}$ is arc-consistent:

**Lemma 7.** $\mathbf{S} \leq^{\mathrm{AC}} \mathbf{B}$

*Proof.* If we show that the sets $\{P_b^{\mathbf{S}^+} : b \in B\}$ satisfy the conditions i. and ii. that define $\mathcal{W}(\mathbf{S})$, it will follow that the duplicator has a winning strategy witnessing that $\mathbf{S} \leq^{\mathrm{AC}} \mathbf{B}$. For condition i., fix an element $a \in S^+$. Let $\mathbf{A}$ be the finite substructure $\mathbf{S}^+ \upharpoonright \{a\}$. By the first property of $\mathbf{S}^+$, the structure $\mathbf{A}$ belongs to $\mathcal{W}$. Let then $b$ be the witness to condition i. for $\mathbf{A}$. The same $b$ works for $\mathbf{S}^+$. For condition ii., fix an element $a \in S^+$, $b \in B$, $R \in \sigma$, and $(a_1, \ldots, a_r) \in R^{\mathbf{S}}$ as in its statement. Let $\mathbf{A}$ be the finite substructure $\mathbf{S}^+ \upharpoonright \{a_1, \ldots, a_r\}$. By the first property of $\mathbf{S}^+$, the structure $\mathbf{A}$ belongs to $\mathcal{W}$. Let then $b_1, \ldots, b_r \in B$ be the witnesses to condition ii. for $\mathbf{A}$. The same witnesses work for $\mathbf{S}^+$. □

Next we show that the existence of a homomorphism $\mathbf{S} \to \mathbf{B}$ determines if arc-consistency solves CSP($\mathbf{B}$).

**Lemma 8.** *The following are equivalent:*

1. $\mathbf{S} \to \mathbf{B}$.
2. $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$ *implies* $\mathbf{A} \to \mathbf{B}$ *for every finite* $\mathbf{A}$.

*Proof.* Assume $\mathbf{S} \to \mathbf{B}$ and let $\mathbf{A}$ be a finite structure such that $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$. This means that $\mathcal{W}(\mathbf{A})$ is not empty; let $\mathbf{A}^+$ be a member of $\mathcal{W}(\mathbf{A})$ and therefore of $\mathcal{W}$. By the second property of $\mathbf{S}^+$, the structure $\mathbf{A}^+$ embeds into $\mathbf{S}^+$, and hence $\mathbf{A}$ also embeds into $\mathbf{S}$. Since $\mathbf{S} \to \mathbf{B}$, also $\mathbf{A} \to \mathbf{B}$.

The converse is proved by a standard compactness argument. As we will not really need this implication in what follows, we omit the standard proof. At any rate, it will be a consequence of the results below (of course, without falling in a circularity; see the proof of Theorem 5). □

Our next goal is to finitize $\mathbf{S}^+$. Since we cannot satisfy the three properties of $\mathbf{S}^+$ in a finite structure, we relax them significantly. This will give us a very naive first candidate to a finitized $\mathbf{S}^+$ which we will strengthen later on. Let $r$ be the maximum arity of the relations in $\sigma$. Let $\mathbf{N}^+$ be a finite $\tau$-structure satisfying the following two properties:

1. every induced substructure of $\mathbf{N}^+$ is isomorphic to some structure in $\mathcal{W}$,
2. every structure in $\mathcal{W}$ of cardinality at most $r$ is isomorphic to some induced substructure of $\mathbf{N}^+$.

Note that the disjoint union of all structures in $\mathcal{W}$ of cardinality at most $r$ does the job. We will take this canonical example as our $\mathbf{N}^+$. Note that $\mathbf{N}^+$ belongs to $\mathcal{W}$ as $\mathcal{W}$ is closed under glued unions and hence under disjoint unions. From now on, let $\mathbf{N}$ be the $\sigma$-reduct of $\mathbf{N}^+$.

By itself, $\mathbf{N}$ is way too naive. If $\mathbf{N}$ is unsatisfiable, meaning that $\mathbf{N} \not\to \mathbf{B}$, we are certainly done as we have a counterexample. But if it is satisfiable, there is not much we can say. We will ask then for a stronger condition on $\mathbf{N}$ that comes inspired by the third property of $\mathbf{S}^+$. We will say that $\mathbf{N}$ is *strongly* $\mathbf{B}$*-satisfiable* if there exists a homomorphism $f : \mathbf{N} \to \mathbf{B}$ such that $f(a_1) = f(a_2)$ for every pair of points $a_1$ and $a_2$ in $N$ for which $\mathbf{N}^+ \restriction \{a_1\}$ and $\mathbf{N}^+ \restriction \{a_2\}$ are isomorphic. The following lemma links our naive candidate $\mathbf{N}$ with our ideal candidate $\mathbf{S}$, in one direction:

**Lemma 9.** *If* $\mathbf{N}$ *is strongly* $\mathbf{B}$*-satisfiable, then* $\mathbf{S} \to \mathbf{B}$.

*Proof.* Let $f : \mathbf{N} \to \mathbf{B}$ be a homomorphism witnessing that $\mathbf{N}$ is strongly satisfiable. Let $c_1, c_2, c_3, \ldots$ be a fixed enumeration of the universe of $\mathbf{S}$. We define a sequence of mappings $h_0, h_1, h_2, h_3, \ldots$ where $h_i$ has domain $\{c_1, \ldots, c_i\}$, inductively. Let $h_0$ be the empty mapping. Let $i > 0$ and suppose that $h_{i-1}$ is already defined. Let $a_i$ be any element of $\mathbf{N}^+$ for which $\mathbf{S}^+ \restriction \{c_i\}$ and $\mathbf{N}^+ \restriction \{a_i\}$ are isomorphic. Such an $a_i$ must exist by the first property of $\mathbf{S}^+$ and the definition of $\mathbf{N}^+$. Let $h_i$ be the extension of $h_{i-1}$ that sets $h_i(c_i) = f(a_i)$. From the fact that $\mathbf{N}$ is strongly satisfied by $f$, this does not depend on the choice of $a_i$.

We claim that the map $h = \bigcup_i h_i$ is a homomorphism from $\mathbf{S}$ to $\mathbf{B}$. Fix a tuple $(c_{i_1}, \ldots, c_{i_r})$ in some relation $R^{\mathbf{S}}$. Let $d_1, \ldots, d_r$ be such that $\mathbf{S}^+ \restriction \{c_{i_1}, \ldots, c_{i_r}\}$ and $\mathbf{N}^+ \restriction \{d_1, \ldots, d_r\}$ are isomorphic with $c_{i_j}$ mapped to $d_j$. Such $d_1, \ldots, d_r$ exist by the first property of $\mathbf{S}^+$ and the definition of $\mathbf{N}^+$. Since $f$ is a homomorphism, we have $(f(d_1), \ldots, f(d_r)) \in R^{\mathbf{B}}$. On the other hand, $\mathbf{S}^+ \restriction \{c_{i_j}\}$ is isomorphic to both $\mathbf{N}^+ \restriction \{d_j\}$ and $\mathbf{N}^+ \restriction \{a_{i_j}\}$. It follows that $f(d_j) = f(a_{i_j}) = h(c_{i_j})$, and therefore also $(h(c_{i_1}), \ldots, h(c_{i_r})) \in R^{\mathbf{B}}$. Thus $h$ is a homomorphism. □

Our next goal is to reverse the implication in Lemma 9. For this we need to introduce some terminology from Ramsey theory.

Let $\mathbf{C}$ and $\mathbf{D}$ be structures and let $p \geq 1$ and $c \geq 1$ be integers. We write $\mathbf{D} \to (\mathbf{C})^p_c$ if for every mapping $f : \binom{D}{p} \to \{1, \ldots, c\}$ there exists a strong embedding $e : \mathbf{C} \xrightarrow{s} \mathbf{D}$ such that for every two sets $A \subseteq C$ and $B \subseteq D$ with $|A| = |B| = p$ and $\mathbf{C} \restriction A \cong \mathbf{D} \restriction B$, it holds that $f(e(A)) = f(e(B))$. Here, the notation $\binom{M}{p}$ stands for the collection of all subsets of $M$ of size $p$. A classic result in Ramsey theory states that for every $p$ and $c$ and every finite structure $\mathbf{C}$, there exists a finite structure $\mathbf{D}$ such that $\mathbf{D} \to (\mathbf{C})^p_c$. See [19] for a beautiful exposition and a discussion on the long history of this result.

On the one hand, we require the Ramsey result for the much simpler case of $p = 1$. On the other, we require it relative to a particular class of finite structures. If $\mathcal{K}$ is a class of finite structures and $p \geq 1$, we say that $\mathcal{K}$ is a $p$*-Ramsey class* if for every $c \geq 1$ and every $\mathbf{C}$ in $\mathcal{K}$, there exists a $\mathbf{D}$ in $\mathcal{K}$ such that $\mathbf{D} \to (\mathbf{C})^p_c$. We say that $\mathcal{K}$ is a *pigeonhole class* if it is a 1-Ramsey class. Relativized Ramsey theorems are also known and have an equally long history. The version stated below seems not to appear in the literature but can be proved by standard methods in the area. We note that the restriction to $p = 1$ is essential in all known approaches. We provide details in Appendix C.

**Theorem 4.** *Let $\mathcal{K}$ be a class of finite structures that is closed under induced substructures and glued unions. Then $\mathcal{K}$ is a pigeonhole class.*

We see how this solves our problem by reversing the implication in Lemma 9. First note that, by Lemma 6 and Theorem 4, the class $\mathcal{W}$ is a pigeonhole class. Let $\mathbf{M}^+$ be the structure $\mathbf{D}$ given by the Theorem 4 with $\mathbf{C} = \mathbf{N}^+$ and $c = |B|$. Let $\mathbf{M}$ be the $\sigma$-reduct of $\mathbf{M}^+$. These two structures will be used in the following:

**Lemma 10.** *If $\mathbf{S} \to \mathbf{B}$, then $\mathbf{N}$ is strongly $\mathbf{B}$-satisfiable.*

*Proof.* Let $h : \mathbf{S} \to \mathbf{B}$. By the second property of $\mathbf{S}^+$, there exists $f : \mathbf{M} \xrightarrow{e} \mathbf{S}$. Composing we get $h \circ f : \mathbf{M} \to \mathbf{B}$. As $\mathbf{M}^+ \to (\mathbf{N}^+)^1_{|B|}$, there exists $e : \mathbf{N}^+ \xrightarrow{s} \mathbf{M}^+$ such that if $\mathbf{N}^+ \restriction \{a_1\} \cong \mathbf{N}^+ \restriction \{a_2\}$, then $h(f(e(a_1))) = h(f(e(a_2)))$. Thus, $h \circ f \circ e$ is a homomorphism witnessing that $\mathbf{N}$ is strongly satisfiable. ◻

Finally, we obtain the characterization:

**Theorem 5.** *The following conditions are equivalent:*

1. $\mathbf{S} \to \mathbf{B}$,
2. $\mathbf{M} \to \mathbf{B}$,
3. $\mathbf{N}$ *is strongly* $\mathbf{B}$*-satisfiable,*
4. $\mathbf{A} \leq^{\mathrm{AC}} \mathbf{B}$ *implies* $\mathbf{A} \to \mathbf{B}$ *for every finite* $\mathbf{A}$.

*Proof.* Implication 1. to 2. follows from the second property of $\mathbf{S}^+$ and the fact that $\mathbf{M}^+$ belongs to $\mathcal{W}$. Implication 2. to 3. is in the proof of Lemma 10. Implication 3. to 1. is Lemma 9. This shows that 1., 2., and 3. are equivalent. The equivalence between 1. and 4. is Lemma 8. But since we proved only that 1. implies 4. in that Lemma, let us note how 4. implies 2.: $\mathbf{M}^+$ belongs to $\mathcal{W}$ and hence $\mathbf{M} \leq^{\mathrm{AC}} \mathbf{B}$, which means that if 4. holds, then 2. holds as well. ◻

Note that 3. is a perfectly decidable condition. Condition 2. is also decidable as $\mathbf{M}^+$ and $\mathbf{M}$ are explicitely defined from $\mathbf{N}^+$. To see this last claim one needs to look into the proof of Theorem 4 given in Appendix C.

## 6 Concluding remarks

Important progress on the analysis of the $k$-consistency algorithm was achieved recently through the algebraic approach to CSPs. Complete decidable classifications are now known for digraphs without sources or sinks [4] and for special triads [3]. Even for general structures a solution was announced recently. As soon as this breakthrough is confirmed, our results give also decidability for SAC and other stacked games. A natural next step would be understanding this decidability proof through some explicit algebraic condition, or perhaps by showing that $k$-consistency is no more powerful than SAC for solving CSPs. On a related note, we are not aware of algebraic conditions that allow comparing the relative strength of two different algorithms as in Corollary 1. Again, this could be because different consistency algorithms collapse after all, or because some refinement of the algebraic approach awaits for discovery.

Finally, the decidable criterion we gave for AC has an appealing combinatorial flavour that calls for generalization. An explicit question we were unable to answer and that stopped our progress is this: Is the class of all instances that are $k$-consistent with respect to a fixed $\mathbf{B}$ the collection of reducts of some amalgamation class? Results in the style of [7] indicate that this might be possible.

*Acknowledgment* We thank Elitza Maneva for useful insights and discussions.

## References

1. M. Ajtai, R. Fagin, and L. J. Stockmeyer. The closure of monadic NP. In *30th Symp. on the Theory of Computing*, 1998.
2. A. Atserias, A. A. Bulatov, and V. Dalmau. On the power of k-consistency. In *34th Intl. Colloq. on Automata, Languages and Programming*, pages 279–290, 2007.
3. L. Barto, M. Kozik, M. Maroti, and T. Niven. CSP dichotomy for special triads. *Proc. Amer. Math. Soc.*, 2009. to appear.
4. L. Barto, M. Kozik, and T. Niven. The CSP dichotomy holds for digraphs with no sources and no sinks. *SIAM Journal on Computing*, 38(5):1782–1802, 2009.
5. C. Bessiere and R. Debruyne. Theoretical analysis of singleton arc consistency and its extensions. *Artificial Intelligence*, 172(1):29–41, 2008.
6. M. Bodirsky and H. Chen. Peek arc consistency. *CoRR*, abs/0809.0788, 2008.
7. J. Covington. Homogenizable relational structures. *Illinois Journal of Mathematics*, 34(4):731–743, 1990.
8. V. Dalmau, Ph. G. Kolaitis, and M. Y. Vardi. Constraint satisfaction, bounded treewidth, and finite variable logics. In *8th Intl. Conf. on Principles and Practice of Constraint Programming*, pages 310–326, 2002.
9. H. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer-Verlag, 1995.
10. T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction. *SIAM Journal on Computing*, 28(1):57–104, 1998.
11. J. Flum and M. Grohe. *Parameterized Complexity Theory.* Springer, 2006.
12. E. C. Freuder. A sufficient condition for backtrack-free search. *Journal of the ACM*, 29(1):24–32, 1982.
13. M. Grohe. Fixed-point logics on planar graphs. In *13th IEEE Symposium on Logic in Computer Science*, pages 6–15, 1998.
14. M. Grohe. Large Finite Structures with Few Lk-Types. *Information and Computation*, 179(2):250–278, 2002.
15. W. Hodges. *Model theory.* Cambridge University Press, 1993.
16. N. Immerman and D. Kozen. Definability with bounded number of bound variables. *Information and Computation*, 83:121–139, 1989.
17. Ph. G. Kolaitis and M. Y. Vardi. The decision problem for the probabilities of higher-order properties. In *19th Symp. Theory of Comp.*, pages 425–435, 1987.
18. Ph. G. Kolaitis and M. Y. Vardi. A game-theoretic approach to constraint satisfaction. In *17th Nat. Conf. on Artificial Intelligence*, pages 175–181, 2000.
19. J. Nešetřil. Ramsey classes and homogeneous structures. *Combinatorics, Probability & Computing*, 14(1-2):171–189, 2005.
20. J. Nešetřil and Z. Zhu. On bounded treewidth duality of graphs. *Journal of Graph Theory*, 23:151–162, 1996.

## A Proof of Lemmata 3 and 4

*Proof (of Lemma 3).* Let $\mathbf{B}$ be given. We have to construct a boolean Datalog query, which is satisfied for an input database $\mathbf{A}$ iff Spoiler wins the game between $\mathbf{A}$ and $\mathbf{B}$. For convenience, we allow equality atoms in bodies of rules. These can be eliminated with standard means.

For all $0 \leq m \leq k$ and all $b_1, \dots, b_m \in B$ we use an intensional predicate $X_{b_1,\dots,b_m}$. The intended semantics is that $X_{b_1,\dots,b_m}$ contains those $a_1, \dots, a_m \in A$ for which Spoiler can win from position $((a_1, b_1), \dots, (a_m, b_m))$. Accordingly, the goal predicate is $X_{()}$.

We translate shrinking moves $(m, M) \in S$ with $m \leq k$ into Datalog rules as follows: Assume that $M = \{i_1, \dots, i_{m'}\}$ with $1 \leq i_1 < \dots < i_{m'} \leq m$. Then we add a rule

$$X_{b_1,\dots,b_m} x_1 \dots x_m \leftarrow X_{b_{i_1},\dots,b_{i_{m'}}} x_{i_1} \dots x_{i_{m'}} .$$

For growing moves $(m, \mathbf{S})$ with $\mathbf{S}$ $m'$-numbered for some $m' \leq k$, let $\varphi_1, \dots, \varphi_n$ be an enumeration of all atomic facts (including equalities) that hold in $\mathbf{S}$, with $x_i$ replaced for $c_i$ for all $1 \leq i \leq m'$. Further, let $\bar{b}_1, \dots, \bar{b}_{n'}$ be an enumeration of all $m'$-tuples $\bar{b}_i = b_{i,1} \dots b_{i,m'}$, such that $b_{i,j} = b_j$ for all $1 \leq j \leq m$ and such that $\mathbf{S} \to (\mathbf{B}, b_{i,1}, \dots, b_{i,m'})$. For each $1 \leq i \leq n'$ let $\psi_i = X_{b_{i,1},\dots,b_{i,m'}} x_1 \dots x_{m'}$. Then we add a rule

$$X_{b_1,\dots,b_m} x_1 \dots x_m \leftarrow \varphi_1, \dots, \varphi_n, \psi_1, \dots, \psi_{n'} .$$

It remains to consider the moves which concern positions of length greater than $k$. Let us call such positions long and other positions short. As $k$ is a grow bound, plays use long positions as follows: A growing move turns a short position into a long position. Then, only shrinking moves are allowed until again some position is short. If Spoiler wants to win, he has to reach a short position after a finite sequence of shrinking moves.

Without loss of generality, the game already allows to combine subsequent shrinking moves into a single shrinking move. Then, we may assume that on long positions Spoiler only plays shrinking moves which immediately result in a short position. Now let $(m, \mathbf{S}) \in G$ be such, that $\mathbf{S}$ is $m'$-numbered for some $m' > k$. Define $\varphi_1, \dots, \varphi_n$ and $\bar{b}_1, \dots, \bar{b}_{n'}$ as above. Now for every sequence $(m', M_1), \dots, (m', M_{n'})$ from $S$ such that $|M_i| \leq k$ for all $1 \leq i \leq n'$, we add a rule

$$X_{b_1,\dots,b_m} x_1 \dots x_m \leftarrow \varphi_1, \dots, \varphi_n, \psi'_1, \dots, \psi'_{n'} ,$$

where $\psi'_i$ is $\psi_i$ combined with the move $(m', M_i)$, that is if $M_i = \{j_1, \dots, j_{m''}\}$ with $1 \leq j_1 < \dots < j_{m''} \leq m'$, then $\psi'_i = X_{b_{i,j_1},\dots,b_{i,j_{m''}}} x_{j_1} \dots x_{j_{m''}}$. □

*Proof (of Lemma 4).* As Duplicator's winning condition is a co-reachability condition, it is clear that it can be expressed by a greatest fixed point. It remains to show, that the body of that fixed point is expressible by universal first-order logic. It is well-known that simultaneous fixed points can be simulated by nested fixed points and this simulation preserves monadicity and universal FO. Hence we may make free use of simultaneous greatest fixed points.

For convenience of presentation, we assume that $K \neq \emptyset$ for all $(h, K) \in S$. This is without loss of generality: As Spoiler has to fulfill a reachability condition and the game rules are positional, he does not need to repeat any position, in particular not the initial one. Together with stackedness, we conclude that each $K$ as above has the form $K = \{1, \ldots, i, j\}$ for some $0 \leq i < j \leq h$. Still, the empty position may appear as a result of a growing round.

First, we will present the encoding of the winning region. Let $k$ be a full bound on the game. For each $1 \leq \ell \leq k$ and all $b_1, \ldots, b_\ell \in B$, we use a monadic variable $X_{b_1,\ldots,b_\ell}$. It will always be used in a context, where first-order variables $x_1, \ldots x_{\ell-1}$ are available. The intended semantic of $X_{b_1,\ldots,b_\ell}$ is as follows: Let $\mathbf{A}$ be the background structure and let $a_1, \ldots, a_{\ell-1} \in A$ be the interpretations of $x_1, \ldots, x_{\ell-1}$. Then $a_\ell \in A$ should belong to (the fixed-point of) $X_{b_1,\ldots,b_\ell}$, iff $((a_1, b_1), \ldots, (a_\ell, b_\ell))$ is a winning position for Duplicator. It remains to capture the position which is the empty tuple. For this we use a nullary variable $X_{()}$ which should contain the empty tuple iff it is winning for Duplicator.

For all $0 \leq m \leq \ell \leq k$ and all $b_1, \ldots, b_\ell \in B$, we define a formula $\varphi^m_{b_1,\ldots,b_\ell}$ with free variables $x_1, \ldots, x_\ell$, which may contain $X_{()}$ and all $X_{b_1,\ldots,b_i,b'}$ for all $0 \leq i < m$ and all $b' \in B$. We intend $\varphi^m_{b_1,\ldots,b_\ell}$ to define the set of all $(a_1, \ldots, a_\ell)$ such that from position $((a_1, b_1), \ldots, (a_\ell, b_\ell))$, Duplicator can force the play into one of the following situations:

1. No shrinking round ever removes any of the first $m$ pairs of this position and ultimately, Duplicator wins.
2. At some point, some of these pairs are removed, and the succeeding position is in some $X_{b'_1,\ldots,b'_{\ell'}}$. Observe, that this implies $\ell' \leq m$ and $b'_i = b_i$ for all $i < \ell'$, because the game is stacked.

Our final formula is the greatest fixed point of $\varphi^0_{()}$.

We define the $\varphi^m_{b_1,\ldots,b_\ell}$ by induction on $k - m$. So, assume that all $\varphi^{m'}_{b'_1,\ldots,b'_{\ell'}}$ for $m' > m$ are already defined, and consider the cases $\ell = m + 1$, $\ell > m + 1$ and $\ell = m$ separately.

First the case $\ell = m + 1 \leq k$. For fixed $b_1, \ldots, b_m$, we consider the simultaneous greatest fixed point of all $\varphi^{m+1}_{b_1,\ldots,b_m,b'}$ for all $b'$. This simultaneous fixed point binds the variables $X_{b_1,\ldots,b_m,b'}$, so its components are formulae which only use $X_{b_1,\ldots,b_i,b''}$ with $i < m$. It also binds (and the reintroduces) $x_{m+1}$ but not $x_1, \ldots, x_m$. We take these formulae as the various $\varphi^m_{b_1,\ldots,b_m,b'}$.

Next, for all $\ell \leq k$ such that $\ell > m + 1$, we obtain $\varphi^m_{b_1,\ldots,b_\ell}$ from $\varphi^{m+1}_{b_1,\ldots,b_\ell}$ by substituting all $X_{b_1,\ldots,b_m,b'}$ by their respective $\varphi^m_{b_1,\ldots,b_m,b'}$ just defined.

Finally for $\ell = m$. First, for any $\ell'$-numbered structure $\mathbf{D}$, let $\varphi_{\mathbf{D}}(x_1, \ldots, x_{\ell'})$ denote the conjunction of all atomic formulae on the variables $x_1, \ldots, x_{\ell'}$ that are satisfied on $\mathbf{D}$ by the mapping $x_i \mapsto c^{\mathbf{D}}_i$. Evaluated on $\mathbf{A}$ with the assignment $x_i \mapsto a_i$, this formula states that $\mathbf{D}$ is a substructure of $(\mathbf{A}, a_1, \ldots, a_{\ell'})$. Finally we let

$$\varphi^m_{b_1,\dots,b_m} := \bigwedge_{\substack{(m,\mathbf{D})\in G\\ \mathbf{D}\text{ is } m\text{-numbered}}} \left(\varphi_{\mathbf{D}} \to \begin{cases}\bot & \text{if } \mathbf{D} \not\to (\mathbf{B}, b_1,\dots,b_m)\\ X_{b_1,\dots,b_m} x_m & \text{if } \mathbf{D} \to (\mathbf{B}, b_1,\dots,b_m)\end{cases}\right) \wedge$$

$$\bigwedge_{\substack{(m,\mathbf{D})\in G\\ \mathbf{D}\text{ is } \ell'\text{-numbered}\\ \ell'>m}} \forall x_{m+1}\dots\forall x_{\ell'} \left(\varphi_{\mathbf{D}} \to \bigvee_{\substack{b_{m+1},\dots,b_{\ell'}\in B\\ \mathbf{D}\to(\mathbf{B},b_1,\dots,b_{\ell'})}} \varphi^m_{b_1,\dots,b_{\ell'}}\right) \wedge$$

$$\bigwedge_{\substack{(m,M)\in S\\ M=\{1,\dots,i,j\}}} X_{b_1,\dots,b_i,b_j} x_j$$

in case $m > 0$. For $m = 0$ we have to replace the unary atoms $X_{b_1,\dots,b_m}x_m$ by the nullary atom $X_{()}$. Also, we can shorten the formula due to the absense of shrinking moves (recall our convention that the empty position is not reachable by shrinking moves), and the fact that $\varphi_{\mathbf{D}}$ is always true for an empty structure $\mathbf{D}$. Thus

$$\varphi^0_{()} := X_{()} \wedge \bigwedge_{\substack{(0,\mathbf{D})\in G\\ \mathbf{D}\text{ is } \ell'\text{-numbered}\\ \ell'>0}} \forall x_1\dots\forall x_{\ell'} \left(\varphi_{\mathbf{D}} \to \bigvee_{\substack{b_1,\dots,b_{\ell'}\in B\\ \mathbf{D}\to(\mathbf{B},b_1,\dots,b_{\ell'})}} \varphi^0_{b_1,\dots,b_{\ell'}}\right).$$

□

## B Proof of Lemma 5

*Proof (of Lemma 5).* To reduce notational clutter, we write $\leq$ for $\leq^{G,S}$. Let $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ denote the game played on $\mathbf{A}$ and $\mathbf{B}$. Now let $\mathbf{A} \not\leq \mathbf{B}$, where $\mathbf{A}$ and $\mathbf{B}$ are $\sigma$-structures. We take $\mathbf{C}$ to be the unravelling of $\mathbf{A}$, described next. Observe, that the legal moves for Spoiler in $\mathcal{G}$ do not depend on $\mathbf{B}$. Hence, let $\mathcal{G}_{\mathbf{A},-}$ denote Spoiler's half of that game, that is for all positions $((a_1, b_1), \ldots, (a_k, b_k))$ of $\mathcal{G}_{\mathbf{A},\mathbf{B}}$, the solitary game $\mathcal{G}_{\mathbf{A},-}$ has a position $(a_1, \ldots, a_k)$ and $\mathcal{G}_{\mathbf{A},-}$ contains only Spoiler's moves. Of course, there is no winning condition. Now, elements of $C$ are pairs $(p, a)$, where $p$ is a (finite) partial play of $\mathcal{G}_{\mathbf{A},-}$, which ends with a growing round, and $a$ was newly introduced into the position in this growing round, which means that both, $a$ was picked by Spoiler in this round, and that $a$ was not present previously. For $c = (p, a) \in C$ and a partial play $p'$ of $\mathcal{G}_{\mathbf{A},-}$, we say that $p'$ *covers* $c$, if $p$ is a prefix of $p'$, and $a$ occurs in all positions of $p'$ after $p$. Now, let $\pi_2 : C \to A$ be the projection to the second component, and for $R \in \sigma$ let

$$R^{\mathbf{C}} := \big\{ \big(c_1, \ldots, c_{\mathrm{ar}(R)}\big) \in \pi_2^{-1}\big(R^{\mathbf{A}}\big) \;\big|\; \text{some partial play } p \text{ covers all } c_i \big\}.$$

By definition of $\mathbf{C}$, we have $\pi_2 : \mathbf{C} \to \mathbf{A}$, so $\mathbf{C} \leq \mathbf{A}$. The same holds for all (finite) substructures $\mathbf{D}$ of $\mathbf{C}$ instead of $\mathbf{C}$.

For $\mathbf{C} \not\leq \mathbf{B}$, we describe a winning strategy for Spoiler in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$. For this, Spoiler privately keeps a board of $\mathcal{G}_{\mathbf{A},\mathbf{B}}$. On this board, he plays according to his winning strategy. He will ensure the following invariants:

1. If $((c_1, b_1), \ldots, (c_k, b_k))$ is the current position in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$, then the one in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ is $((\pi_2(c_1), b_1), \ldots, (\pi_2(c_k), b_k))$.
2. All elements from $C$ in the current position in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$ are covered by the partial play of $\mathcal{G}_{\mathbf{A},-}$, which is obtained from the current partial play in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ by forgetting all about $\mathbf{B}$.

For this, he plays shrinking rounds in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$ just as in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$, where he uses his winning strategy. Note, that this preserves the invariants. For growing rounds, say in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ he chooses $a_{k+1}, \ldots, a_\ell$ and a substructure $\mathbf{S}$ of

$$(\mathbf{A}, \pi_2(c_1), \ldots, \pi_2(c_k), a_{k+1}, \ldots, a_\ell)$$

after a partial play $p$. Then, let $p'$ be the respective partial play of $\mathcal{G}_{\mathbf{A},-}$, with this growing round already amended. In order to define $p'$, we do not need to wait for Duplicator's answer in any game. Let

$$(c_{k+1}, \ldots, c_\ell) := ((p', a_{k+1}), \ldots, (p', a_\ell))$$

and let $\pi$ be the restriction of $\pi_2$ to $c_1, \ldots, c_\ell$. Then $c_1, \ldots, c_\ell$ are all covered by $p'$, so $\pi : (\mathbf{C}, c_1, \ldots, c_\ell) \cong (\mathbf{A}, a_1, \ldots, a_\ell)$. Hence $\mathbf{S}' := \pi^{-1}(\mathbf{S})$ is a substructure of $(\mathbf{C}, c_1, \ldots, c_\ell)$ and $\pi : \mathbf{S}' \cong \mathbf{S}$. As $G$ is isomorphism invariant, $(k, \mathbf{S}') \in G$, so it is legal for Spoiler to choose $c_{k+1}, \ldots, c_\ell$ and $\mathbf{S}'$ in the respective growing round in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$. Let $b_{k+1}, \ldots, b_\ell$ be Duplicator's answer in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$. Spoiler uses this also as

Duplicator's answer in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$. Again, the invariants are preserved. Furthermore, in case this choice from Duplicator does not terminate $\mathcal{G}_{\mathbf{C},\mathbf{B}}$, that is if $\mathbf{S}' \to (\mathbf{B}, b_1, \ldots, b_\ell)$, then isomorphism invariance of $\to$ implies $\mathbf{S} \to (\mathbf{B}, b_1, \ldots, b_\ell)$, that is the game $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ also continues. In particular, as Spoiler's strategy for $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ in winning, at some point Duplicator cannot choose in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ and thus also not in $\mathcal{G}_{\mathbf{C},\mathbf{B}}$.

Observe further that, as Duplicator can always repeat what she did in a position already seen earlier in the game, we may assume that Spoiler's winning strategy in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$ guarantees that no position is seen twice. Now, let $k$ be a full bound on the game. Then, in $\mathcal{G}_{\mathbf{A},\mathbf{B}}$, there are at most

$$n := |A|^{k+1} \cdot |B|^{k+1}$$

positions. In a strategy where positions are not seen twice, Spoiler will win after at most $n$ rounds. So we do not need all of $\mathbf{C}$: We can restrict the unravelling accordingly, which leaves us with a finite substructure $\mathbf{D}$ of $\mathbf{C}$, for which we also have $\mathbf{D} \not\to \mathbf{B}$.

The last thing we need to show is that $\mathbf{C}$ has a tree-decomposition of width at most $k-1$. This will imply that $\mathbf{D} \in \mathrm{TW}(k-1)$. Now consider the set $V$ of partial plays of $\mathcal{G}_{\mathbf{A},-}$. We turn it into a tree, by letting $\{p, p'\} \in E$, if $p'$ extends $p$ by exactly one position. For $v \in V$, we let $B_v$ be the set of all $c \in C$ which are covered by $v$. We have $|B_v| \leq k$, because if $v$ covers $(p, a)$, then $a$ is part of the sequence which is the last position of $v$, and for each such $a$ in this sequence, there is exactly one $p$, such that $(p, a)$ is covered by $v$; namely the prefix of $v$ up to the last growing round that introduced $a$. Each $(p, a) \in C$ is covered by $p$, so it occurs in $B_p$. For each tuple $(c_1, \ldots, c_r)$ in $\mathbf{C}$, by definition of $\mathbf{C}$ there is some $v \in V$ which covers all $c_i$, so $c_1, \ldots, c_r \in B_v$. For $c = (p, a) \in C$, the set of $v \in V$ such that $c \in B_v$ is the set of all partial plays that extend $p$ for which $a$ appears in its last position and was not deleted since the last position of $p$. This set is clearly connected. Hence, the bags $B_v$ form a tree decomposition of $\mathbf{C}$ of width at most $k-1$. □

## C Proof of Theorem 4

The proof we give can be considered standard. Moreover, as we are interested only on pigeonhole classes, the arguments are way more elementary than the general case of $p$-Ramsey classes for $p > 1$. The construction we give here is very similar to the construction given in the proof in the book by Graham, Rothschild and Spencer for 2-Ramsey classes of graphs.

A singleton structure is a structure whose universe is a singleton; has cardinality one. For a singleton structure $\mathbf{T}$, we write $\mathbf{D} \to (\mathbf{C})^{\mathbf{T}}_c$ if for every $f : D \to \{1, \dots, c\}$ there exists a strong embedding $e : \mathbf{C} \xrightarrow{s} \mathbf{D}$ such that for every $a_1$ and $a_2$ in $C$ for which $\mathbf{D} \restriction \{a_1\}$ and $\mathbf{D} \restriction \{a_2\}$ are isomorphic to $\mathbf{T}$, it holds that $f(e(a_1)) = f(e(a_2))$. We say that $f$ is $\mathbf{T}$-homogeneous on the copy of $\mathbf{C}$ given by the embedding $e$. We say that $\mathcal{K}$ is a $\mathbf{T}$-pigeonhole class if for every $c \geq 1$ and every $\mathbf{C}$ in $\mathcal{K}$, there exists $\mathbf{D}$ in $\mathcal{K}$ such that $\mathbf{D} \to (\mathbf{C})^{\mathbf{T}}_c$. First we prove the following special case of what we want:

**Lemma 11.** *Let $\mathcal{K}$ be a class of finite structures that is closed under induced substructures and glued unions, and let $\mathbf{T}$ be a singleton structure. Then $\mathcal{K}$ is a $\mathbf{T}$-pigeonhole class.*

*Proof.* Fix a structure $\mathbf{C}$ in $\mathcal{K}$ and an integer $c \geq 1$. Let $C = \{a_1, \dots, a_n\}$ and let $m = n(c+1)$. We define a sequence of structures $\mathbf{D}_0, \mathbf{D}_1, \dots, \mathbf{D}_m$ inductively. The final $\mathbf{D}$ will be $\mathbf{D}_m$. Each $\mathbf{D}_i$ will consist of $m$ parts, where each part will form an *independent set* in the sense that no pair of distinct points from one part will appear together in a tuple of a relation in $\mathbf{D}_i$.

The structure $\mathbf{D}_0$ consists of $\binom{m}{n}$ disjoint copies of $\mathbf{C}$, one for each choice of $n$ distinct parts

$$1 \leq p_1 < \cdots < p_n \leq m,$$

in such a way that the copy of $a_i$ is placed in part $p_i$ of $\mathbf{D}_0$. Having defined $\mathbf{D}_i$, we define $\mathbf{D}_{i+1}$. Let $\mathbf{E}_{i+1}$ be the substructure induced by part $i+1$ of $\mathbf{D}_i$. The construction guarantees that $\mathbf{E}_{i+1}$ is an independent set. Let $\mathbf{F}_{i+1}$ be the disjoint union of $c|E_{i+1}|$ copies of $\mathbf{E}_{i+1}$. By mixing points from different copies we may have many more copies of $\mathbf{E}_{i+1}$ in $\mathbf{F}_{i+1}$ than those that form $\mathbf{F}_{i+1}$. The structure $\mathbf{D}_{i+1}$ is formed by placing $\mathbf{F}_{i+1}$ in part $i+1$ and gluing one copy of $\mathbf{D}_i$ through every induced copy of $\mathbf{E}_{i+1}$ in $\mathbf{F}_{i+1}$. All the copies of $\mathbf{D}_i$ are glued disjointly except possibly at the overlaps of the copies of $\mathbf{E}_{i+1}$ in $\mathbf{F}_{i+1}$. For $p \in \{1, \dots, m\} - \{i+1\}$, we place in part $p$ of $\mathbf{D}_{i+1}$ all the points in part $p$ of each of the copies of $\mathbf{D}_i$. These are still independent sets.

This defines $\mathbf{D}$. The first thing we need to check is that $\mathbf{D}$ is belongs to $\mathcal{K}$. For this we proceed inductively. For $\mathbf{D}_0$ this is direct because it is the disjoint union of copies of $\mathbf{C}$, and $\mathbf{C}$ belongs to $\mathcal{K}$. And if it is true for $\mathbf{D}_i$, it also true for $\mathbf{D}_{i+1}$ because it is obtained by gluing copies of $\mathbf{D}_i$ through an induced substructure of $\mathbf{F}_{i+1}$, which is a disjoint union of copies of $\mathbf{E}_{i+1}$, which is an induced substructure of $\mathbf{D}_i$.

Let us now check that the required property $\mathbf{D} \to (\mathbf{C})^{\mathbf{T}}_c$ is satisfied. Let $f : D \to \{1, \dots, c\}$ be an arbitrary coloring of $\mathbf{D} = \mathbf{D}_m$. Looking at $f$ on part

$m$ of $\mathbf{D}_m$, we want to find a copy of $\mathbf{E}_m$ on which $f$ is $\mathbf{T}$-homogeneous. This is proved in the following claim:

*Claim.* There exists a copy of $\mathbf{E}_m$ in $\mathbf{F}_m$ on which $f$ is $\mathbf{T}$-homogeneous.

*Proof.* Let $S$ be the collection of all $a$ in $E_m$ such that $\mathbf{E}_m \restriction \{a\} \cong \mathbf{T}$. If $S$ is empty, we are done; $f$ is already $\mathbf{T}$-homogeneous on every copy of $\mathbf{E}_m$ in $\mathbf{F}_m$. If $S$ is not empty, fix an element $a \in S$ and note that there are at least $c|E_m|$ copies of $a$ in $\mathbf{F}_m$. By the pigeonhole principle, there at least $|E_m|$ copies of $a$ that get the same color. As obviously $|S| \leq |E_m|$, we find our copy of $\mathbf{E}_m$ by choosing $|S|$ copies of $a$ that get the same color under $f$, and an arbitrary copy of every other $a'$ in $E_m \setminus S$. □

We continue now with the proof. By the construction of $\mathbf{D}_m$, the copy of $\mathbf{E}_m$ in $\mathbf{F}_m$ can be extended to a copy of $\mathbf{D}_{m-1}$ inside $\mathbf{D}_m$. Now we proceed inductively backwards. Looking at $f$ on part $m-1$ of this copy of $\mathbf{D}_{m-1}$, we find a copy of $\mathbf{E}_{m-1}$ on which $f$ is $\mathbf{T}$-homogeneous. Repeating, we end up finding a copy of $\mathbf{D}_0$ inside $\mathbf{D}$ where $f$ is $\mathbf{T}$-homogeneous on every part. This gives a well-defined coloring of the parts of $\mathbf{D}_0$ with $c+1$ colors: we color a part of $\mathbf{D}_0$ by $f(e)$ if there exists some $e$ in that part such that $\mathbf{D}_0 \restriction \{e\} \cong \mathbf{T}$, and we color it by 0 otherwise. Since there are $n(c+1)$ parts and only $c+1$ colors, some set of $n$ parts must be monochromatic. The monochromatic $n$ parts in $\mathbf{D}_0$ give a copy of $\mathbf{C}$ on which $f$ is $\mathbf{T}$-homogeneous. □

Now we can proceed with the proof of Theorem 4:

*Proof.* Fix a structure $\mathbf{C}$ in $\mathcal{K}$. Let $C = \{a_1, \ldots, a_n\}$, let $\mathbf{D}_0 = \mathbf{C}$, and define $\mathbf{D}_{i+1}$ by applying Lemma 11 on $\mathbf{D}_i$ and $\mathbf{T} = \mathbf{C} \restriction \{a_i\}$. The final $\mathbf{D}$ is $\mathbf{D}_n$. □