Canonical Representations for Circular-Arc Graphs Using Flip Sets

Maurice Chandoo

Leibniz Universität Hannover, Theoretical Computer Science,
Appelstr. 4, 30167 Hannover, Germany
`chandoo@thi.uni-hannover.de`

──── **Abstract** ────

We show that computing canonical representations for circular-arc (CA) graphs reduces to computing certain subsets of vertices called flip sets. For a broad class of CA graphs, which we call uniform, it suffices to compute a CA representation to find such flip sets. As a consequence canonical representations for uniform CA graphs can be obtained in polynomial-time. We then investigate what kind of CA graphs pose a challenge to this approach. This leads us to introduce the notion of restricted CA matrices and show that the canonical representation problem for CA graphs is logspace-reducible to that of restricted CA matrices. As a byproduct, we obtain the result that CA graphs without induced 4-cycles can be canonized in logspace.

1 Introduction

We consider an arc to be a connected set of points on the unit circle including the endpoints. A CA graph is a graph whose vertices can be assigned arcs such that two vertices are adjacent iff their corresponding arcs intersect. More formally, given a graph G we call it a CA graph if there exists a function ρ which maps every vertex u of G to an arc $\rho(u)$ such that u and v are adjacent iff their arcs $\rho(u)$ and $\rho(v)$ have non-empty intersection. We call such a mapping ρ a CA representation of G. CA graphs are a form of geometrical intersection graphs. Let $\mathcal{X}$ be a family of sets over some ground set. Any subset Y of $\mathcal{X}$ defines a graph G_Y which has Y as its vertex set and two vertices are adjacent if they have non-empty intersection. The graph G_Y is called intersection graph of Y. We say a (finite) graph G is an intersection graph of $\mathcal{X}$ if it is isomorphic to the intersection graph of Y for some $Y \subseteq \mathcal{X}$. In this language CA graphs are intersection graphs of arcs. The intersection graphs of intervals on the real line are called interval graphs. In this sense any set of geometrical objects defines a (geometrical intersection) graph class. CA graphs are a generalization of interval graphs since every set of intervals on the real line can be 'bent' into arcs while preserving the intersection relation. Therefore every interval graph is a CA graph.

Being a generalization of interval graphs—the archetype of geometrical intersection graphs—CA graphs are quite prominent as well and have been known for decades. Since then structural properties and algorithmic problems for this class have been thoroughly investigated with [5] and [16] being two of the earliest works in this regard. In particular, finding characterizations of CA graphs and constructing a CA representation for a given CA graph have received a great deal of attention. Remarkably, finding a forbidden induced subgraph characterization of CA graphs is still an open problem. See [12] for a survey on this line of research and [1] for one of the most recent results in that direction. It should also be

mentioned that CA graphs are of practical relevance with applications arising in disciplines such as genetics and operations research. An explanation of the connection between genetics and interval graphs in layman's terms can be found in [17]. For a specialized account on this connection emphasizing circularity see [15]. An example of how CA graphs can be used to model the problem of phasing traffic lights is given in [6].

In this work we consider the canonical representation problem for CA graphs. The representation problem for CA graphs is as follows. Given a CA graph G as input we want to output a CA representation ρ_G of G. The canonical variant of this problem imposes the additional requirement that for every pair of isomorphic CA graphs G and H their representations ρ_G and ρ_H should have identical underlying sets of arcs, i.e. $\{\rho_G(v) \mid v \in V(G)\} = \{\rho_H(v) \mid v \in V(H)\}$. Notice that solving the representation problem for CA graphs implies solving the recognition problem for CA graphs, i.e. the question given a graph G is it a CA graph. Likewise, solving the canonical representation problem for CA graphs implies solving the isomorphism problem for CA graphs, i.e. deciding whether two given CA graphs are isomorphic.

Consider the following generalization of interval graphs: 2-interval graphs are intersection graphs of two intervals on the real line. It is easy to see that this class contains CA graphs because given a set of arcs one can cut the circle at some point and straighten the arcs. The arcs which are cut can be modeled as two intervals. It is interesting to note that the isomorphism problem for interval graphs is logspace-complete [9] while the one for 2-interval graphs is already GI-complete and CA graphs lie inbetween these two classes. The GI-completeness for 2-interval graphs follows from the fact that every line graph is a 2-interval graph and line graphs are already GI-complete. To see why this inclusion holds consider a graph G and its line graph $L(G)$. Assign every vertex v of G an interval I_v on the real line such that no two intervals I_u and I_v intersect for every pair of distinct vertices u and v of G. The 2-interval model for $L(G)$ is obtained by mapping every edge $\{u, v\}$ of G to the two intervals I_u and I_v.

While a polynomial-time algorithm for deciding isomorphism of interval graphs is known since 1976 due to Booth and Lueker this question still remains open for CA graphs. There have been two claimed polynomial-time algorithms for deciding isomorphism of CA graphs in [18] and [7] which were shown to be incorrect in [4] and [3] respectively. For interval graphs even a linear-time algorithm for isomorphism is known [13]. A more recent result is that canonical interval representations for interval graphs can be computed in logspace and that this is optimal in the sense that recognition and deciding isomorphism for interval graphs is logspace-complete [9]. These two hardness results also carry over to the class of CA graphs. Furthermore, the isomorphism problem for proper CA graphs [11] and Helly CA graphs [10] have been shown to be decidable in logspace. It is also shown how to obtain canonical representations for these subclasses in logspace.

In this article we explain how the method used in [10] to obtain canonical representation for Helly CA graphs can be adapted to CA graphs in general. Following this approach, canonical representations for CA graphs can be found by computing certain subsets of vertices called flip sets in an isomorphism-invariant manner. We introduce the class of uniform CA graphs for which this method yields canonical representations in polynomial-time. We then aim to isolate the instances of CA graphs which are difficult to handle with this method. We try to capture these hard instances by what we call restricted CA matrices and show that the canonical representation problem for CA graphs is logspace-reducible to that of restricted CA matrices. During this isolation process we find a subset of uniform CA graphs, namely Δ-uniform CA graphs, for which canonical representations can be computed in

logspace. The Δ-uniform CA graphs contain Helly CA graphs and every CA graph without an induced 4-cycle. This generalizes the canonization result for Helly CA graphs given in [10]. A preliminary version of this work appeared in [2].

The paper is organized as follows. In the third section we formalize the idea of computing invariant flip sets in order to obtain canonical representations for CA graphs. This leads to the definition of invariant flip set functions. In the fourth section we investigate for what CA graphs a particular invariant flip set function is easy to compute. This leads to the class of uniform CA graphs. We also provide an alternative characterization of uniform CA graphs in terms of whether certain triangles in a CA graph have an unambiguous representation. The main result of this section is that the representation problem for uniform CA graphs, the canonical representation problem for uniform CA graphs and the non-Helly triangle representability problem (introduced in section 4) for uniform CA graphs are all logspace-equivalent. In the fifth section we consider the structure of non-uniform CA graphs, introduce restricted CA matrices and show how the canonical representation problem for CA graphs can be reduced to that of restricted CA matrices. In the process of proving this reduction the class of Δ-uniform CA graphs is defined and it is shown that this class can be canonized in logspace.

2 Preliminaries

For a number $n \in \mathbb{N}$ we write $[n]$ for $\{1, \ldots, n\}$. Given two sets A, B we say A and B intersect if $A \cap B \neq \emptyset$. We say A and B overlap, in symbols $A \between B$, if $A \cap B, A \setminus B$ and $B \setminus A$ are non-empty.

We consider graphs without self-loops which sometimes have colored vertices and colored edges. They can be seen as relational structures with the vertex set as universe and vertex colors encoded as unary relations and edge colors as binary relations. The standard notion of isomorphism for relational structures applies. We describe a graph with vertex colors as tuple (G, c) where c is a function that maps the vertices of G to the colors. We talk about a graph with edge colors as a square matrix whose entries represent the edge colors and identify the indices of the matrix and the vertices of the graph. Consequently, we identify a square matrix with the graph that it represents and talk about it in graph-theoretical terms. By a class of (relational) structures we mean a set of such structures which is closed under isomorphism.

We call a bijective function τ which maps the vertices of a graph G to some set V' a relabeling of G and $\tau(G)$ denotes the graph obtained after relabeling the vertices of G according to τ. Let G and H be two graphs and let $X \subseteq V(G)$ and $Y \subseteq V(H)$. We say X and Y are in the same orbit, in symbols $X \sim_{\mathrm{orb}} Y$, if there exists an isomorphism π from G to H such that $\pi(X) = Y$. Let f be a function which maps a graph along with a subset of its vertex set to a binary string, i.e. $f(G, X) \in \{0, 1\}^*$ and $X \subseteq V(G)$. We call f an invariant for a graph class $\mathcal{C}$ if $f(G, X) = f(H, Y)$ whenever $X \sim_{\mathrm{orb}} Y$ and $G, H \in \mathcal{C}$. Let us call a function f which maps a graph G to a family of subsets of its vertex set, i.e. $f(G) \subseteq \mathcal{P}(V(G))$, a vertex set selector. For example, the function that maps a graph to the set of its cliques is a vertex set selector. The characteristic function χ_f of a vertex set selector f is defined as $\chi_f(G, X) = 1 \Leftrightarrow X \in f(G)$. We say a vertex set selector f is invariant for a graph class $\mathcal{C}$ if its characteristic function χ_f is an invariant for $\mathcal{C}$. We call f globally invariant if χ_f is an invariant for all graphs. Intuitively, a vertex set selector f is invariant for $\mathcal{C}$ if a graph $G \in \mathcal{C}$ can be arbitrarily relabeled and f still returns the 'same' vertex sets as before w.r.t. $\sim_{\mathrm{orb}}$.

The following definitions are with respect to a graph G. Throughout the paper it will be

always clear from context with respect to what graph these expressions are to be interpreted. For a vertex v we define its open neighborhood $N(v)$ as the set of vertices which are adjacent to v and its closed neighborhood $N[v] = N(v) \cup \{v\}$. A vertex v is called universal if $N[v] = V(G)$. For two vertices u, v we say that u and v are twins if $N[u] = N[v]$. A graph G is twin-free if for every pair of distinct vertices $u \neq v$ it holds that $N[u] \neq N[v]$. A twin class is an inclusion-maximal set of vertices X such that for all $u, v \in X$ it holds that u and v are twins. For two subsets of vertices S, S' with $S' \subseteq S$ we define the exclusive neighborhood $N_S(S')$ as all vertices $v \in V(G) \setminus S$ such that v is connected to all vertices in S' and to none in $S \setminus S'$. Let A be a square matrix with entries from a set $\mathcal{E}$. For a vertex u of A and $x \in \mathcal{E}$ we define $N^x(u) = \{v \in V \mid A_{u,v} = x\}$.

Logspace Transducers and Reductions.   We assume deterministic Turing machines as default model of computation. A logspace transducer is a deterministic Turing machine M with a read-only input tape, a work tape and a write-only output tape. The work tape is only allowed to use at most $\mathcal{O}(\log n)$ cells where n denotes the input length. To write onto the output tape M has a designated state called output state with the following semantic. If M enters the output state then the symbol in the current cell of the work tape is written to the current cell of the output tape and the head on the output tape is moved one cell to the right. Other than that, M cannot write or move the head on the output tape. This means as soon as something is written to the output tape it cannot be modified afterwards. Let Σ and Γ be the input and work alphabet of M respectively. Then M computes a function $f_M \colon \Sigma^* \to \Gamma^*$. We say a (partial) function f is computed by a logspace transducer M if $f(x) = f_M(x)$ whenever $f(x)$ is defined. We call f logspace-computable if there exists a logspace transducer M which computes f. The class of logspace-computable functions is closed under composition. Let f be a function which maps words over some alphabet to words over some other alphabet. We say that the length of f is polynomially bounded if $|f(x)|$ is polynomially bounded by $|x|$. Only functions whose length is polynomially bounded can be logspace-computable since the runtime of a logspace transducer is polynomially bounded. A language $L \subseteq \Sigma^*$ is in logspace if its characteristic function is logspace-computable.

Given two functions f and g we say f is logspace-reducible to g if there exists $l \in \mathbb{N}$ and logspace-computable functions $r_1, \ldots, r_l$ such that f can be expressed as composition of g and $r_1, \ldots, r_l$. Intuitively, this means that an oracle which can compute g can be queried a constant number of times when constructing a logspace transducer for f. For two functions f and g we say that they are logspace-equivalent if f is logspace-reducible to g and vice versa. Analogously, given three functions f, g_1, g_2 we say f is logspace-reducible to g_1 and g_2 if there exists $l \in \mathbb{N}$ and logspace-computable functions $r_1, \ldots, r_l$ such that f can be expressed as composition of g_1, g_2 and $r_1, \ldots, r_l$.

Circular-Arc Graphs and Representations.   A CA model is a set of arcs $\mathcal{A} = \{A_1, \ldots, A_n\}$ on the circle. Let $p \neq p'$ be two points on the circle. Then the arc A specified by $[p, p']$ is given by the part of the circle that is traversed when starting from p going in clockwise direction until p' is reached. We say that p is the left and p' the right endpoint of A and write $l(\cdot), r(\cdot)$ to denote the left and right endpoint of an arc in general. If $A = [p, p']$ then the arc obtained by swapping the endpoints $\overline{A} = [p', p]$ covers exactly the opposite part of the circle plus the endpoints. We say $\overline{A}$ is obtained by flipping A. In our context, we are only interested in the intersection structure of a CA model and thus only the relative position of the endpoints to each other matter. All endpoints can w.l.o.g. be assumed to be pairwise different and no arc covers the full circle. Under these assumptions, a CA model $\mathcal{A}$ with n
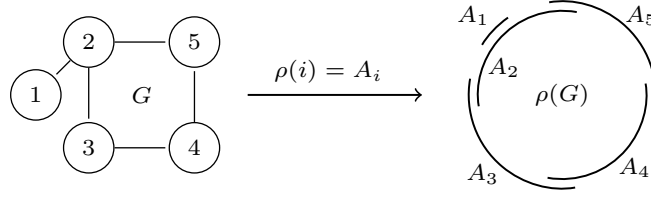
■ **Figure 1** A CA graph and a representation of it

arcs can be described as a unique string as follows. Pick an arbitrary arc $A \in \mathcal{A}$ and relabel the arcs with $1, \dots, n$ in order of appearance of their left endpoints when traversing the circle clockwise starting from the left endpoint of A. Then write down the endpoints in order of appearance when traversing the circle clockwise starting from the left endpoint of A. Do this for every arc and pick the lexicographically smallest resulting string as representation for $\mathcal{A}$. For example, the smallest such string for the CA model in Figure 1 would result from choosing A_1: $(l(1), r(1), l(2), r(5), l(3), r(2), \dots)$. Let $\text{str}(\mathcal{A})$ denote this smallest string representation. For a CA model $\mathcal{A}$ let $\mathcal{A}^{\text{r}}$ be the CA model obtained after reversing the order of its endpoints. Observe that reversing the endpoints does not affect the intersection structure of a CA model. Therefore we consider two CA models $\mathcal{A}$ and $\mathcal{B}$ to be equal if $\text{str}(\mathcal{A}) = \text{str}(\mathcal{B})$ or $\text{str}(\mathcal{A}^{\text{r}}) = \text{str}(\mathcal{B})$.

Let G be a graph and $\rho = (\mathcal{A}, f)$ consists of a CA model $\mathcal{A}$ and a bijective mapping f from the vertices of G to the arcs in $\mathcal{A}$. Then ρ is called a CA representation of G if for all $u \neq v \in V(G)$ it holds that $\{u, v\} \in E(G) \Leftrightarrow f(u) \cap f(v) \neq \emptyset$. We write $\rho(x)$ to mean the arc $f(x)$ corresponding to the vertex x, $\rho(G)$ for the CA model $\mathcal{A}$ and for a subset $V' \subseteq V(G)$ let $\rho[V'] = \{\rho(v) \mid v \in V'\}$. A graph is a CA graph if it has a CA representation.

We say a CA model $\mathcal{A}$ has a hole if there exists a point on the circle which isn't contained by any arc in $\mathcal{A}$. Every such CA model can be understood as interval model (a set of intervals on the real line) by straightening the arcs. Conversely, every interval model can be seen as CA model by bending the intervals. Therefore a graph is an interval graph iff it admits a CA representation with a hole.

A family of sets $\mathcal{F}$ over some ground set is called Helly if for all subsets $\mathcal{F}'$ of $\mathcal{F}$ such that all elements in $\mathcal{F}'$ intersect pairwise it holds that $\cap_{A \in \mathcal{F}'} A$ is non-empty. A CA graph G is called Helly (HCA graph) if it has a CA representation ρ with a Helly CA model $\rho(G)$. This is the case iff for all inclusion-maximal cliques C in G it holds that the overall intersection of C in ρ is non-empty, i.e. $\bigcap_{v \in C} \rho(v) \neq \emptyset$. Every interval model has the Helly property and therefore every interval graph is a Helly CA graph.

The intersection type of two circular arcs A and B can be one of the following five types:
- `di`: A and B are disjoint — $A \cap B = \emptyset$
- `cs`: A contains B — $B \subset A$
- `cd`: A is contained by B — $A \subset B$
- `cc`: A and B jointly cover the circle (circle cover) — $A \between B$ and $A \cup B =$ whole circle
- `ov`: A and B overlap — $A \between B$ and $A \cup B \neq$ whole circle

Using these types we can associate a matrix with every CA model. An intersection matrix is a square matrix with entries $\{\texttt{cc}, \texttt{cd}, \texttt{cs}, \texttt{di}, \texttt{ov}\}$. Given a CA model $\mathcal{A}$ we define its intersection matrix $\mu_{\mathcal{A}}$ such that $(\mu_{\mathcal{A}})_{A,B} \in \{\texttt{cc}, \texttt{cd}, \texttt{cs}, \texttt{di}, \texttt{ov}\}$ reflects the intersection type of the arcs $A \neq B \in \mathcal{A}$. An intersection matrix μ is called a CA (interval) matrix if it is the intersection matrix of some CA model (with a hole). See Figure 2 for an example of a CA model and the CA matrix which it induces. Given an intersection matrix μ and two distinct

vertices u, v of μ we sometimes write $u \, \alpha \, v$ instead of $\mu_{u,v} = \alpha$ if μ is clear from the context. Also, we sometimes talk about an intersection matrix μ as if it were a graph. In that case we consider two vertices u, v of μ to be adjacent if they do not have a `di`-entry in μ.

When trying to construct a CA representation for a CA graph G it is clear that whenever two vertices are non-adjacent their corresponding arcs must be disjoint in every CA representation of G. For two adjacent vertices the intersection type of their corresponding arcs might depend on the particular CA representation of G that one considers. Hsu has shown that this ambiguity can be removed as follows [7].

We adopt the notation of [10].

▶ **Definition 2.1.** For a graph G we define its neighborhood matrix λ_G which is an intersection matrix as

$$(\lambda_G)_{u,v} = \begin{cases} \texttt{di} & \text{, if } \{u, v\} \notin E(G) \\ \texttt{cd} & \text{, if } N[u] \subsetneq N[v] \\ \texttt{cs} & \text{, if } N[v] \subsetneq N[u] \\ \texttt{cc} & \text{, if } N[u] \between N[v] \text{ and } N[u] \cup N[v] = V(G) \\ & \quad \text{and } \forall w \in N[u] \setminus N[v] : N[w] \subset N[u] \\ & \quad \text{and } \forall w \in N[v] \setminus N[u] : N[w] \subset N[v] \\ \texttt{ov} & \text{, otherwise} \end{cases}$$

for all $u \neq v \in V(G)$.

Let μ be an intersection matrix with vertex set V and let $\rho = (\mathcal{A}, f)$ where $\mathcal{A}$ is a CA model and f is a bijective mapping from V to $\mathcal{A}$. We say ρ is a CA representation of μ if f is an isomorphism from μ to the intersection matrix $\mu_\mathcal{A}$ of $\mathcal{A}$. We denote the set of such CA representations for μ with $\mathcal{N}(\mu)$. The representation problem for CA matrices is to compute a CA representation for a given CA matrix μ. The canonical representation problem for CA matrices is defined analogously to the canonical representation problem for CA graphs. We say ρ is a normalized CA representation for a graph G if ρ is a CA representation for the neighborhood matrix λ_G of G. An example of a normalized representation can be seen in Figure 3. Let us denote the set of all normalized CA representations for G with $\mathcal{N}(G) = \mathcal{N}(\lambda_G)$.

▶ **Lemma 2.2** (Corollary 2.3. [7]). *Every twin-free CA graph G without a universal vertex has a normalized CA representation, that is $\mathcal{N}(G) \neq \emptyset$.*

▶ **Lemma 2.3.** *The canonical representation problem for CA graphs is logspace-reducible to the canonical representation problem for vertex-colored twin-free CA graphs without a universal vertex.*

Proof. For a graph G let G_0 denote the induced subgraph of G that is obtained by removing all universal vertices from G and only taking one vertex from each twin-class and deleting

$\mu_\mathcal{A}$	a	b	c	d
a	–	cs	ov	cc
b	cd	–	di	di
c	ov	di	–	ov
d	cc	di	ov	–
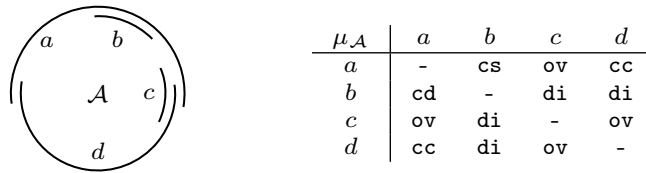
■ **Figure 2** A CA model $\mathcal{A}$ and its intersection matrix $\mu_\mathcal{A}$

the rest. Let c_0 be a coloring of G_0 which assigns each vertex the cardinality of its twin class in G. It holds that (G_0, c_0) and the number of universal vertices in G suffice to reconstruct G. Let G be a CA graph. Compute the graph (G_0, c_0). Since (G_0, c_0) is twin-free and without universal vertices we can compute a canonical representation ρ_0 for it. For a vertex v of G let v_0 denote the twin of v that occurs in G_0. A canonical representation of G is given by $v \mapsto \rho_0(v_0)$ for every non-universal vertex v of G and every universal vertex of G is represented by an arc which intersects with all other arcs. ◀

Therefore for our purposes it suffices to consider only twin-free graphs without universal vertices and a vertex-coloring.

▶ **Proviso.** From this point on we assume every graph to be twin-free and without a universal vertex unless explicitly stated otherwise. As a consequence we view CA graphs as a set of CA matrices in the sense that the neighborhood matrix of every CA graph is a CA matrix.

Flips in Intersection Matrices. McConnell [14] observed that the operation of flipping arcs in CA models has a counterpart in intersection matrices. He called this counterpart operation algebraic flips. Note that for two arcs A, B with intersection type $\alpha \in \{\mathtt{cc}, \mathtt{cd}, \mathtt{cs}, \mathtt{di}, \mathtt{ov}\}$ the intersection type of $\overline{A}$ and B is solely determined by α. More precisely, the intersection type of $\overline{A}$ and B is $Z_{10}(\alpha)$ where Z_{10} is the function defined in Table 1. Similarly, the intersection type of A and $\overline{B}$ is given by $Z_{01}(\alpha)$. Using the functions Z_{ij} we can define the operation of flipping a set of vertices in an intersection matrix.

▶ **Definition 2.4.** Let μ be an intersection matrix with vertex set V and $X \subseteq V$. We define the intersection matrix $\mu^{(X)}$ obtained after flipping the vertices X in μ as

$$\mu_{u,v}^{(X)} = Z_{ij}(\mu_{u,v}) \text{ with } i = 1 \text{ iff } u \in X \text{ and } j = 1 \text{ iff } v \in X$$

for all $u \neq v$ in V.

Since flipping the same set of arcs twice is an involution it follows that $(\mu^{(X)})^{(X)} = \mu$.

▶ **Definition 2.5.** Let V be a set of vertices, let $\mathcal{A}$ be a set of arcs and let ρ be a function that maps V to $\mathcal{A}$. Then $\rho^{(X)} \colon V \to \mathcal{A}$ for $X \subseteq V$ is defined as follows:

$$\rho^{(X)}(v) = \begin{cases} \overline{\rho(v)} & \text{, if } v \in X \\ \rho(v) & \text{, if } v \notin X \end{cases}$$

Notice that flipping vertices in an intersection matrix is equivalent to flipping arcs in a CA representation in the following sense. Given an intersection matrix λ and a subset of its vertices X it holds that $\rho \in \mathcal{N}(\lambda) \Leftrightarrow \rho^{(X)} \in \mathcal{N}(\lambda^{(X)})$. Also, it is not difficult to observe that flipping is an isomorphism-invariant operation in the sense that flipping sets of vertices which are in the same orbit lead to isomorphic intersection matrices.
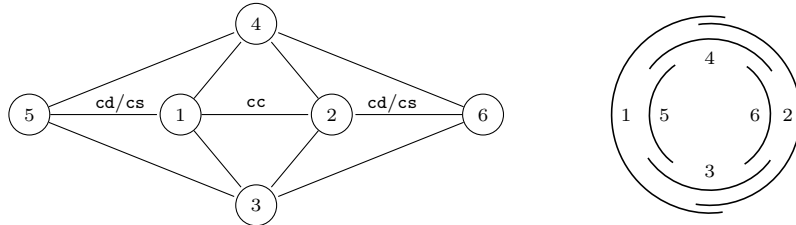
Figure 3 A CA graph and a normalized representation thereof. Every non-labeled edge corresponds to an $\mathtt{ov}$-entry in the neighborhood matrix.

3   Flip Trick

In this section we generalize the idea used by Köbler, Kuhnert and Verbitsky in [10] to compute canonical representations for Helly CA graphs. They showed that finding canonical representations for Helly CA graphs can be reduced to finding canonical representations for vertex-colored interval matrices. We show that the idea behind this reduction also works for CA matrices in general. Recall that CA graphs can be seen as special case of CA matrices since the neighborhood matrix of every CA graph is a CA matrix. The converse does not hold, i.e. there exist CA matrices which are not expressible as the neighborhood matrix of a CA graph (for instance any CA matrix with only two vertices that are not disjoint). The key result here, which is used in the subsequent sections, is that finding canonical representations for CA matrices is logspace-reducible to the task of computing what we call an invariant flip set function.

McConnell showed in [14] that CA representations for CA graphs can be computed as follows. Given a CA graph G with neighborhood matrix λ one can compute a set of vertices X of G such that $\lambda^{(X)}$ is an interval matrix. We call such a set X a flip set. Then by computing an interval representation ρ for $\lambda^{(X)}$ and flipping back the arcs X in ρ one obtains a CA representation for λ and therefore for G as well [14]. We essentially use the same argument to obtain canonical CA representations.

▶ **Definition 3.1.** Let λ be a CA matrix. A subset of vertices X of λ is called a flip set if there exists a representation $\rho \in \mathcal{N}(\lambda)$ and a point x on the circle such that $v \in X$ iff $\rho(v)$ contains the point x.

The concept of flip sets has already been implicitly defined and used in both [14] and [10]. They observed that $\lambda^{(X)}$ is an interval matrix whenever X is a flip set of a CA matrix λ. In fact, the other direction holds as well leading to the following characterization.

▶ **Lemma 3.2.** *Let λ be a CA matrix and X is a subset of vertices of λ. It holds that X is a flip set iff $\lambda^{(X)}$ is an interval matrix.*

Proof. "⇒": Let X be a flip set of λ. Let $\rho \in \mathcal{N}(\lambda)$ be a witnessing representation of the fact that X is a flip set, i.e. there exists a point x on the circle such that every arc $\rho(v)$ with $v \in X$ contains x and every arc $\rho(v)$ with $v \notin X$ does not contain x. Consider the representation $\rho^{(X)} \in \mathcal{N}(\lambda^{(X)})$. It holds that no arc $\rho^{(X)}(v)$ with $v \in V(\lambda)$ contains the point x which implies that there is a hole in $\rho^{(X)}$ and thus $\lambda^{(X)}$ is an interval matrix.

"⇐": Let X be a subset of vertices of λ such that $\lambda^{(X)}$ is an interval matrix. We argue that X must be a flip set. Let $\rho \in \mathcal{N}(\lambda^{(X)})$ be a CA representation of $\lambda^{(X)}$ containing a hole at point x on the circle. Such a representation must exist since $\lambda^{(X)}$ is an interval matrix. This means the arc $\rho(v)$ does not contain the point x for every vertex $v \in V(\lambda)$. Consider the

🟧 **Table 1** Algebraic flip functions $Z_{xy} \colon \{\mathtt{cc}, \mathtt{cd}, \mathtt{cs}, \mathtt{di}, \mathtt{ov}\} \to \{\mathtt{cc}, \mathtt{cd}, \mathtt{cs}, \mathtt{di}, \mathtt{ov}\}$

$Z_{xy}(\alpha)$	cc	cd	cs	di	ov
Z_{00}	cc	cd	cs	di	ov
Z_{01}	cs	di	cc	cd	ov
Z_{10}	cd	cc	di	cs	ov
Z_{11}	di	cs	cd	cc	ov

representation $\rho^{(X)} \in \mathcal{N}((\lambda^{(X)})^{(X)}) = \mathcal{N}(\lambda)$. Then it can be checked that $\rho^{(X)}(v)$ contains the point x iff v is in X and therefore X is a flip set with respect to λ. ◄

We already mentioned that the canonical representation problem for vertex-colored interval matrices can be solved in logspace due to [10]. However, since the theorem that we reference just states this result for uncolored interval matrices we shortly explain how to modify the proof to incorporate the coloring, which is a straightforward task for anyone familiar with the proof.

▶ **Theorem 3.3** ([10, Thm. 5.5]). *The canonical representation problem for vertex-colored interval matrices can be solved in logspace.*

Proof. In Theorem 5.5 of [10] it is stated that a canonical interval representation for an interval matrix can be found in logspace. To prove this they convert the input interval matrix λ into a colored tree $\mathbb{T}(\lambda)$ called Δ tree which is a complete invariant for interval matrices. The leafs of this tree correspond to the vertices of λ. By appending the color of a vertex from our vertex-colored interval matrix λ to the existing color of its corresponding leave node in the colored Δ tree $\mathbb{T}(\lambda)$ one obtains a complete invariant for vertex-colored interval matrices. Then by applying the same argument given in the proof of Theorem 5.5 one can also compute a canonical representation for a vertex-colored interval matrix using this slightly modified colored Δ tree. ◄

A consequence of Lemma 3.2 and Theorem 3.3 is that flip sets can be recognized in logspace. Given an intersection matrix λ and a subset of vertices X of λ it suffices to check whether $\lambda^{(X)}$ is an interval matrix by trying to compute an interval representation.

▶ **Definition 3.4.** Let $\mathcal{C}$ be a class of CA matrices and f is a vertex set selector. The function f is called an invariant flip set function for $\mathcal{C}$ if the following conditions hold:
1. For every $\lambda \in \mathcal{C}$ there exists an $X \in f(\lambda)$ such that X is a flip set of λ
2. f is invariant for $\mathcal{C}$
Recall that f is globally invariant if f is invariant for all intersection matrices.

▶ **Theorem 3.5.** *Let $\mathcal{C}$ be a class of CA matrices. The canonical representation problem for vertex-colored $\mathcal{C}$ is logspace-reducible to the problem of computing an invariant flip set function for $\mathcal{C}$.*

Proof. Let f be an invariant flip set function for $\mathcal{C}$. Given a vertex-colored CA matrix (λ, c) with $\lambda \in \mathcal{C}$ a canonical representation can be computed as follows. For every flip set $X \in f(\lambda)$ we associate it with the colored interval matrix $I_X = (\lambda^{(X)}, c_X)$ where $c_X(v) = (c(v), \text{red})$ if v is in X and $(c(v), \text{blue})$ if v is not in X for all $v \in V(\lambda)$. For a colored interval matrix I let $\hat{\rho}_I$ denote a canonical representation of I. Such a canonical representation can be computed in logspace due to Theorem 3.3. Let $\hat{X}$ denote a flip set in $f(\lambda)$ such that the interval model of $\hat{\rho}_{I_{\hat{X}}}$ is lexicographically minimal, i.e. for all flip sets X in $f(\lambda)$ it holds that the model of $\hat{\rho}_{I_X}$ is not smaller than the model of $\hat{\rho}_{I_{\hat{X}}}$. Let $\hat{\rho}$ denote the CA representation that is obtained after flipping the red arcs in $\hat{\rho}_{I_{\hat{X}}}$. Since these are the arcs that were flipped to convert λ into I_X it holds that $\hat{\rho}$ is a representation of λ. To see that this leads to a canonical representation consider two isomorphic vertex-colored CA matrices (λ, c) and (μ, d) with $\lambda, \mu \in \mathcal{C}$ and $V(\lambda)$ and $V(\mu)$ are disjoint. Let $\mathcal{I}_\lambda$ be the set of colored interval matrices I_X for all flip sets $X \in f(\lambda)$, and the set $\mathcal{I}_\mu$ is defined analogously. Let $\mathcal{M}_\lambda$ be the set of interval models M such that there exists an $I \in \mathcal{I}_\lambda$ and M is the model underlying the canonical representation $\hat{\rho}_I$ of I. The set $\mathcal{M}_\mu$ is defined analogously. Since f is invariant it follows that

for every $I \in \mathcal{I}_\lambda$ there exists an $I' \in \mathcal{I}_\mu$ such that I and I' are isomorphic, and vice versa. Since the models in $\mathcal{M}_\lambda$ and $\mathcal{M}_\mu$ only depend on the isomorphism type of the matrices in $\mathcal{I}_\lambda$ and $\mathcal{I}_\mu$ it follows that $\mathcal{M}_\lambda = \mathcal{M}_\mu$. The CA models which underlie the canonical representations of λ and μ are both derived from the smallest element in $\mathcal{M}_\lambda = \mathcal{M}_\mu$ and thus are identical.   ◄

Suppose that there is a partition of the set of CA graphs into two classes $\mathcal{C}$ and $\mathcal{D}$ such that you can efficiently compute invariant flip set functions for both classes. One might be misled into thinking that this implies canonical representations for all CA graphs can be found efficiently. However, this is not the case unless the class $\mathcal{C}$ (or $\mathcal{D}$) can be efficiently recognized, or one of the two invariant flip set functions is globally invariant.

▶ **Lemma 3.6.** *Let $\mathcal{C}$ and $\mathcal{D}$ be classes of CA matrices. The canonical representation problem for $\mathcal{C} \cup \mathcal{D}$ is logspace-reducible to the canonical representation problem for $\mathcal{C}$ and the problem of computing a globally invariant flip set function for $\mathcal{D}$.*

Proof. Let f be a globally invariant flip set function for $\mathcal{D}$. Let $\mathcal{D}'$ be the set of CA matrices λ such that $f(\lambda)$ contains a flip set. Clearly, $\mathcal{D}$ is a subset of $\mathcal{D}'$. It holds that $f(\lambda)$ contains a flip set iff $\lambda \in \mathcal{D}'$. Since f is globally invariant it follows that f is an invariant flip set function for $\mathcal{D}'$. To obtain a canonical representation for a matrix $\lambda \in \mathcal{C} \cup \mathcal{D}$ first compute $f(\lambda)$. If $f(\lambda)$ contains a flip set it holds that $\lambda \in \mathcal{D}'$ and therefore the output of f can be used to find a canonical representation for λ. If $f(\lambda)$ contains no flip set it must be the case that $\lambda \in \mathcal{C}$ and therefore the canonization algorithm for $\mathcal{C}$ can be applied.   ◄

We conclude this section by restating the invariant flip set function that was used in [10] to compute canonical representations for Helly CA graphs and explain why it is correct:

$$f_{\mathrm{HCA}}(G) = \big\{ N[u] \cap N[v] \mid u, v \in V(G) \big\}$$

In a Helly CA graph G every inclusion-maximal clique C of G is a flip set. To see why this holds let ρ be a representation of G with the Helly property. Since C is a clique this means every pair of arcs $\rho(u)$ and $\rho(v)$ with $u, v \in C$ intersects. By the Helly property it follows that the overall intersection $\bigcap_{v \in C} \rho(v)$ is non-empty. This means there exists a point x on the circle such that every arc $\rho(v)$ with $v \in C$ contains x. Assume there exists a vertex $w \in V(G) \setminus C$ such that $\rho(w)$ contains x. This means w must be adjacent to every vertex in C, which contradicts that C is inclusion-maximal. Hence C is a flip set.

In [10, Thm. 3.2] it is shown that every Helly CA graph contains at least one inclusion-maximal clique which can be expressed as the common neighborhood of two vertices. Therefore $f_{\mathrm{HCA}}(G)$ returns at least one flip set for every Helly CA graph G. Also, it is trivial to see that f_{HCA} is globally invariant.

4 Uniform Circular-Arc Graphs

We define the class of uniform CA graphs for which computing a particular invariant flip set function reduces to computing a representation. As a consequence, canonical representations for this class of CA graphs can be computed in polynomial-time. This is an interesting class for two reasons. First, it seems to capture the instances where it is easy to apply the flip trick. Secondly, its complement (within the CA graphs) is a rather exotic class of CA graphs with a quite particular structure. While the initial definition of uniformity makes it apparent why it suffices to find an arbitrary representation in order to obtain a canonical one, it is rather impractical when trying to understand what constitutes a uniform

CA graph. We provide a more pleasant characterization of uniform CA graphs in terms of how certain triangles in a CA graph can be represented. This alternative characterization also reveals that every Helly CA graph is uniform. Additionally, we show that the canonical representation problem for uniform CA graphs is logspace-equivalent to what we call the non-Helly triangle representability problem. This problem is: given a CA graph G and a set T of three pairwise overlapping vertices as input, does there exist a representation ρ of G such that T covers the whole circle in ρ?

The following kind of flip set will lead us to uniform CA graphs when trying to compute canonical representations. Given a CA matrix λ recall that X is a flip set of λ if there exists a representation $\rho \in \mathcal{N}(\lambda)$ and a point x on the circle such that $x \in \rho(v)$ iff $v \in X$ for all vertices v of λ. We impose the additional restriction that x is not allowed to be an arbitrary point on the circle but instead has to be one of the endpoints in ρ.

▶ **Definition 4.1.** Let λ be a CA matrix and $u \in V(\lambda)$. A flip set X of λ is a u-flip set if there exists a representation $\rho \in \mathcal{N}(\lambda)$ and an endpoint x of $\rho(u)$ such that $v \in X$ iff $\rho(v)$ contains the point x.

Clearly, every CA graph has a u-flip set for every vertex u. On the other hand, there are CA graphs that have flip sets which are not u-flip sets for any vertex u. For example, consider the cycle graph with $n \geq 4$ vertices. Every flip set that consists of exactly one vertex is not a u-flip set for any vertex u of the cycle graph.
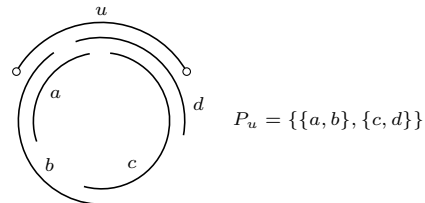
Consider the following task: given a CA graph G and a vertex u, find a u-flip set of G. Clearly, no vertex v which is disjoint from u or contained by u belongs to X since in every representation the arc of v does not contain any of the two endpoints of the arc of u. Similarly, if a vertex v contains u or forms a circle cover with u then in every representation the arc of v contains both endpoints of u and therefore must be included in X. See Figure 4 for a schematic overview.

It remains to decide for the set of vertices $N^{\mathsf{ov}}(u)$ that overlap with u whether they should be included in X. A vertex v which overlaps with u contains exactly one of the endpoints of u in any representation. Let x, y be two vertices that overlap with u. We say x and y overlap from the same side with u in ρ if $\rho(x)$ and $\rho(y)$ contain the same endpoint of $\rho(u)$. Evidently, this is an equivalence relation with respect to v and ρ which partitions $N^{\mathsf{ov}}(u)$ into two parts, namely the part which contains the left endpoint and the one which contains the right endpoint. If X is a u-flip set then $X \cap N^{\mathsf{ov}}(u)$ must be an equivalence class of the 'overlap from the same side with u in ρ'-relation for some $\rho \in \mathcal{N}(G)$.

▶ **Definition 4.2.** For a CA matrix λ and a vertex u of λ we say a partition Y of $N^{\mathsf{ov}}(u)$ into two parts is a u-ov-partition if there exists a representation $\rho \in \mathcal{N}(\lambda)$ such that two vertices $x, y \in N^{\mathsf{ov}}(u)$ are in the same part of Y iff $\rho(x)$ and $\rho(y)$ overlap from the same side with $\rho(u)$. We say ov-partition to mean an u-ov-partition for an arbitrary $u \in V(\lambda)$.
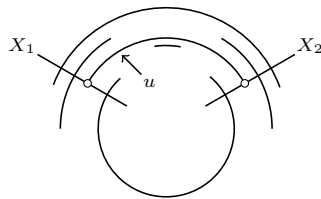
Figure 4 Exemplary u-flip sets X_1 and X_2 **Figure 5** Example of a u-overlap partition P_u

In general, for a vertex u of a CA graph G there can be multiple u-ov-partitions. In fact, there are instances with exponentially many u-ov-partitions with respect to $|N^{\mathsf{ov}}(u)|$. A trivial way of obtaining at least one u-ov-partition for every vertex u of a CA graph G is to compute an arbitrary representation $\rho \in \mathcal{N}(G)$. But the ov-partitions obtained by this method are not invariant and thus do not yield canonical representations. However, if one considers CA graphs where there is only one u-ov-partition for every vertex u then an arbitrary representation suffices.

▶ **Definition 4.3** (Uniform CA Graphs). A CA graph G is uniform if for every vertex u in G there exists exactly one u-ov-partition. This partition is denoted by $P_u = \{P_{u,1}, P_{u,2}\}$.

▶ **Lemma 4.4.** *The following mapping is an invariant flip set function for uniform CA graphs. Let G be a uniform CA graph.*

$$F_{\mathrm{uniform}}(G) = \bigcup_{\substack{u \in V(G) \\ i \in \{1,2\}}} \left\{ \{u\} \cup N^{\mathsf{cd}}(u) \cup N^{\mathsf{cc}}(u) \cup P_{u,i} \right\}$$

Proof. Let G be a uniform CA graph and X is in $F_{\mathrm{uniform}}(G)$ with $X = \{u\} \cup N^{\mathsf{cd}}(u) \cup N^{\mathsf{cc}}(u) \cup P_{u,i}$ for some $u \in V(G)$ and $i \in \{1,2\}$. It follows from Figure 4 and the definition of ov-partitions that X is a u-flip set. The invariance of $F_{\mathrm{uniform}}(G)$ follows from the fact that the intersection type of two vertices as well as the property of being an ov-partition is independent of the vertex labels.   ◀

We remark that the function F_{uniform} is undefined for non-uniform CA graphs since the sets $P_{u,1}$ and $P_{u,2}$ are not well-defined in that context.

▶ **Theorem 4.5.** *Canonical representations for uniform CA graphs can be computed in polynomial-time.*

Proof. Let G be a uniform CA graph. Compute a normalized representation ρ of G and extract the u-ov-partition for each vertex u from ρ. Then compute $F_{\mathrm{uniform}}(G)$ from Lemma 4.4 to obtain a canonical CA representation for G. Since CA representations can be computed in polynomial-time (see for instance [14]) it follows that this procedure also works in polynomial-time.   ◀

Considering that our definition of uniform CA graphs arose from the desire to compute invariant u-flip sets one might expect that these graphs are only a small special case of CA graphs. Surprisingly, quite the opposite is the case as we will see. We give an alternative definition of uniform CA graphs which gives a better intuition as to why many CA graphs are uniform.

▶ **Definition 4.6.** Let λ be a CA matrix. An ov-triangle T of λ is a set of three vertices that overlap pairwise, i.e. for all $u \neq v$ in T it holds that $u\,\mathsf{ov}\,v$. An ov-triangle T is representable as non-Helly triangle (interval triangle) if there exists a representation $\rho \in \mathcal{N}(\lambda)$ such that the set of arcs $\{\rho(x) \mid x \in T\}$ does (not) cover the whole circle. Let $\mathcal{T}_{\mathrm{NHT}}(\lambda)$ and $\mathcal{T}_{\mathrm{IT}}(\lambda)$ denote the sets of ov-triangles representable as non-Helly triangles and interval triangles respectively.

This definition also applies to CA graphs via their neighborhood matrix, i.e. $\mathcal{T}_{\mathrm{IT}}(G) = \mathcal{T}_{\mathrm{IT}}(\lambda)$ and $\mathcal{T}_{\mathrm{NHT}}(G) = \mathcal{T}_{\mathrm{NHT}}(\lambda)$ where λ is the neighborhood matrix of G. See Figure 6 for

an example where the vertices u, x, z are represented as non-Helly triangle on the left and interval triangle on the right.

Recall that a set of arcs which intersect pairwise but have overall empty intersection is called non-Helly. Since three pairwise overlapping arcs that cover the whole circle have overall empty intersection we call such a set a non-Helly triangle. In fact, one can verify that this is the only non-Helly arrangement of three arcs. A complete list of inclusion-minimal non-Helly CA models can be found in [8, Corrollary 3.1].

▶ **Theorem 4.7.** *A CA graph G is uniform iff $\mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G) = \emptyset$.*

Proof. "$\Rightarrow$": Assume there exists a uniform CA graph G with $\mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G) \neq \emptyset$. Let T be an ov-triangle in $\mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G)$ and $T = \{x, y, z\}$. This means there exist two representations $\rho_I, \rho_N \in \mathcal{N}(G)$ such that T is represented as interval triangle in ρ_I and as non-Helly triangle in ρ_N. We assume w.l.o.g. that $\rho_I(y) \subset \rho_I(x) \cup \rho_I(z)$, i.e. y is placed in-between x and z in ρ_I. This means y and z must be in the same part of the unique x-ov-partition P_x. However, y and z do not contain the same endpoint of x in the representation ρ_N, which contradicts that G is uniform.

"$\Leftarrow$": Assume there exists a CA graph G with $\mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G) = \emptyset$ that is not uniform. This means there exist a vertex u, two vertices $x, y \in N^{\mathrm{ov}}(u)$ and two representations $\rho, \rho' \in \mathcal{N}(G)$ such that x and y overlap from the same side with u in ρ but not in ρ'. This implies that x and y must overlap and therefore $T = \{u, x, y\}$ is an ov-triangle. Notice that T must be represented as interval triangle in ρ because x and y both contain the same endpoint of u. It holds that T is represented as interval triangle in ρ' as well since otherwise $T \in \mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G)$. Also, we assume w.l.o.g. that $\rho(y) \subset \rho(x) \cup \rho(u)$. Since u and y overlap it holds that $N[u] \setminus N[y] \neq \emptyset$. Due to ρ' it follows that $N[u] \setminus N[y] \subseteq N[u] \cap N[x]$. For a vertex $z \in N[u] \setminus N[y]$ to intersect with both u and x it is necessary that z overlaps with u and x due to the representation ρ. It follows that $\{u, x, z\}$ is represented as non-Helly triangle in ρ. On the other hand, $\{u, x, z\}$ must be represented as interval triangle in ρ' and therefore $\mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G) \neq \emptyset$, contradiction. See Figure 6 for a schematic overview of ρ and ρ'. ◀

Observe that if an ov-triangle T of G is representable as non-Helly triangle then this implies that T must have certain structural properties in G. For example, every vertex of G must be adjacent to at least one of the vertices in T since T covers the whole circle in some representation. Similarly, if T is representable as interval triangle this also implies some structural properties. For instance, there must be an $x \in T$ such that every vertex that is adjacent to x must also be adjacent to at least one other vertex in T. If an ov-triangle is representable as both non-Helly triangle and interval triangle then it must satisfy all of these structural properties at once. As a consequence such an ov-triangle must have a very particular structure which extends to the whole graph as we will see in the next section.

A CA graph is Helly if it has a Helly CA representation. In [8, Theorem 4.1] it is shown that every 'stable' representation of a Helly CA graph is Helly. Since every normalized
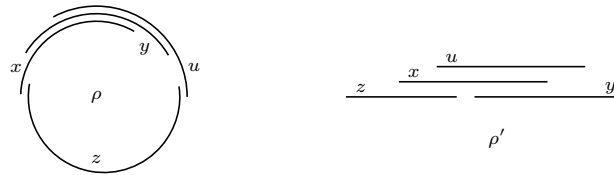
■ **Figure 6** "$\Leftarrow$"-direction in the proof of Theorem 4.7

representation has the 'stable' property it follows that a CA graph is Helly iff every normalized representation of it is Helly. If a CA graph G is Helly this implies that $\mathcal{T}_{\mathrm{NHT}}(G)$ is empty, and therefore every Helly CA graph is uniform.

A natural question to consider is the computational complexity of deciding whether an ov-triangle is representable as non-Helly triangle or interval triangle. Given a CA graph G and an ov-triangle T of G let us call the problem of deciding whether T is in $\mathcal{T}_{\mathrm{NHT}}(G)$ the non-Helly triangle representability problem. Analogously, deciding whether T is in $\mathcal{T}_{\mathrm{IT}}(G)$ is called the interval triangle representability problem. In the case of uniform CA graphs these two problems are complementary, i.e. an ov-triangle T is in $\mathcal{T}_{\mathrm{NHT}}(G)$ iff T is not in $\mathcal{T}_{\mathrm{IT}}(G)$. In the following, we show that solving either of these two problems for uniform CA graphs is logspace-equivalent to computing a canonical representation for uniform CA graphs.

▶ **Definition 4.8.** Let G be a CA graph and $T = \{u, v, w\}$ is an ov-triangle of G. We say v is amidst u and w if one of the following conditions holds:
1. $N_T(u)$ and $N_T(w)$ are non-empty
2. there exists a $z \in N_T(u, w)$ such that $\{u, w, z\} \in \mathcal{T}_{\mathrm{NHT}}(G)$

▶ **Lemma 4.9.** *Let G be a uniform CA graph and $T = \{u, v, w\}$ is an ov-triangle of G with $T \notin \mathcal{T}_{\mathrm{NHT}}(G)$. Then the following statements are equivalent:*
1. *v is amidst u and w*
2. *$\exists \rho \in \mathcal{N}(G) : \rho(v) \subset \rho(u) \cup \rho(w)$*
3. *$\forall \rho \in \mathcal{N}(G) : \rho(v) \subset \rho(u) \cup \rho(w)$*

Proof. "2 ⇒ 1": Let ρ be in $\mathcal{N}(G)$ such that $\rho(v) \subset \rho(u) \cup \rho(w)$ and assume that v is not amidst u, w. Since v overlaps with u and w it holds that $N[u] \setminus N[v]$ and $N[w] \setminus N[v]$ are non-empty. Because $N_T(u) = N_T(w) = \emptyset$ it must hold that $N_T(u, w) \neq \emptyset$. Let $z \in N_T(u, w)$. For z to intersect with u and w in ρ it must hold that $\{u, w, z\}$ is represented as non-Helly triangle in ρ. This contradicts the assumption that v is not amidst u, w.

"1 ⇒ 3": Let v be amidst u and w and assume that there exists a $\rho \in \mathcal{N}(G)$ such that $\rho(v) \not\subset \rho(u) \cup \rho(w)$. Since $T \notin \mathcal{T}_{\mathrm{NHT}}(G)$ and G is uniform it follows by Theorem 4.7 that T must be represented as interval triangle in every representation, which includes ρ. We assume w.l.o.g. that $\rho(w) \subset \rho(u) \cup \rho(v)$. From that it follows that $N_T(w)$ is empty and therefore there must be a $z \in N_T(u, w)$ such that $\{u, w, z\}$ is a non-Helly triangle in ρ, which is impossible.

"3 ⇒ 2": clear. ◀

▶ **Definition 4.10.** Let G be a CA graph and $u \in V(G)$. Let the binary relation $\sim_u$ on $N^{\mathrm{ov}}(u)$ be defined such that $x \sim_u y$ holds if one of the following holds:
1. $x = y$
2. x cd y or x cs y
3. x ov y, $\{u, x, y\} \notin \mathcal{T}_{\mathrm{NHT}}(G)$ and u is not amidst x and y

▶ **Lemma 4.11.** *For every uniform CA graph G and $u \in V(G)$ it holds that the partition induced by $\sim_u$ equals the unique u-ov-partition P_u. Stated differently, $x \sim_u y$ iff x and y are in the same part of P_u.*

Proof. "⇒": Let $x \sim_u y$ and assume for the sake of contradiction that x and y are not in the same part of the u-ov-partition. This means there exists a representation $\rho \in \mathcal{N}(G)$ such that $\rho(x)$ and $\rho(y)$ contain different endpoints of $\rho(u)$. This is only possible if x and y overlap. Since $\{u, x, y\} \notin \mathcal{T}_{\mathrm{NHT}}(G)$ this means $\{u, x, y\}$ must be represented as interval triangle in ρ. In order for $\rho(x)$ and $\rho(y)$ to contain different endpoints of $\rho(u)$ it must

hold that $\rho(u) \subset \rho(x) \cup \rho(y)$, which implies that u is amidst x and y by Lemma 4.9. This contradicts $x \sim_u y$.

"$\Leftarrow$": Let x and y be in the same part of the u-ov-partition and assume that $x \sim_u y$ does not hold. This implies that x and y must overlap and therefore $\{u, x, y\}$ form an ov-triangle. For $x \sim_u y$ to not hold it must be either the case that $\{u, x, y\}$ is only representable as non-Helly triangle or u is amidst x and y. In both cases this contradicts x and y being in the same part of the u-ov-partition. ◄

▶ **Theorem 4.12.** *The representation, canonical representation, non-Helly triangle representability and interval triangle representability problem for uniform CA graphs are logspace-equivalent.*

Proof. The non-Helly triangle representability and interval triangle representability problem for uniform CA graphs are logspace-equivalent because they are complementary in the sense that an ov-triangle is representable as non-Helly triangle iff it is not representable as interval triangle. This follows from the fact that an ov-triangle can only be either represented as non-Helly triangle or interval triangle and these two possibilities are mutually exclusive in the case of uniform CA graphs. As a consequence these two problems are trivially reducible to the representation problem for uniform CA graphs. Given a uniform CA graph G, an ov-triangle T of G and a representation $\rho \in \mathcal{N}(G)$ it holds that $T \in \mathcal{T}_{\mathrm{NHT}}(G)$ iff $T \notin \mathcal{T}_{\mathrm{IT}}(G)$ iff T is represented as non-Helly triangle in ρ.

The representation problem is obviously reducible to the canonical representation problem. Therefore it remains to show that the canonical representation problem for uniform CA graphs is reducible to the non-Helly triangle representability problem. To obtain a canonical representation for a uniform CA graph we can use the invariant flip set function given in Lemma 4.4. To compute this function we need to figure out the unique ov-partitions for each vertex. By Lemma 4.11 this can be done by computing the equivalence relation $\sim_u$ for each vertex u. It can be verified that this relation is computable in logspace using queries of the form $T \in \mathcal{T}_{\mathrm{NHT}}(G)$. ◄

The isomorphism problem for CA graphs can be reduced to the one for non-uniform CA graphs in polynomial-time due to Theorem 4.5. However, a reduction from the canonical representation problem for CA graphs to the one for non-uniform CA graphs does not immediately follow from Theorem 4.5 unless uniform CA graphs can be recognized in polynomial-time. An alternative approach to construct such a reduction is to solve the non-Helly triangle representability problem for uniform CA graphs with an additional requirement.

▶ **Definition 4.13.** The globally invariant non-Helly triangle representability problem for uniform CA graphs is defined as follows. Let A be an algorithm that correctly decides the non-Helly triangle representability problem for uniform CA graphs. Let f_A be the function computed by A, i.e. for a graph G and an ov-triangle T of G it holds that $f_A(G, T) = 1$ iff A accepts (G, T). We say A decides the globally invariant non-Helly triangle representability problem for uniform CA graphs if f_A is an invariant *for all graphs*. Stated differently, the output of A must be independent of the vertex labels.

▶ **Lemma 4.14.** *The canonical representation problem for CA graphs is logspace-reducible to the globally invariant non-Helly triangle representability problem for uniform CA graphs and the canonical representation problem for vertex-colored non-uniform CA graphs.*

Proof. Suppose we are given an algorithm A which solves the globally invariant non-Helly triangle representability problem for uniform CA graphs. We argue that A can be used to compute a globally invariant flip set function for uniform CA graphs. From Lemma 3.6 it then follows that the canonical representation problem for CA graphs reduces to that for vertex-colored non-uniform CA graphs.

Given a CA graph G let $\Delta(G, A)$ be the set of ov-triangles T of G such that A accepts (G, T). If G is a uniform CA graph then $\Delta(G, A) = \mathcal{T}_{\text{NHT}}(G)$. Consider Definition 4.8 and 4.10 and suppose that each occurrence of $\mathcal{T}_{\text{NHT}}(G)$ is replaced by $\Delta(G, A)$. Let us call the new relation $\sim_u^A$. Clearly, in the case of uniform CA graphs $\sim_u$ and $\sim_u^A$ coincide. Next, consider the following variant of F_{uniform}:

$$F_{\text{uniform}}^A(G) = \bigcup_{\substack{u \in V(G) \\ X \in (N^{\text{ov}}(u)/\sim_u^A)}} \left\{ \{u\} \cup N^{\text{cd}}(u) \cup N^{\text{cc}}(u) \cup X \right\}$$

where $(N^{\text{ov}}(u)/\sim_u^A)$ denotes the equivalence classes of $\sim_u^A$. If $\sim_u^A$ is not an equivalence relation let $(N^{\text{ov}}(u)/\sim_u^A) = \emptyset$. If G is a uniform CA graph then it follows from Lemma 4.11 that $F_{\text{uniform}}(G) = F_{\text{uniform}}^A(G)$. Therefore F_{uniform}^A is an invariant flip set function for uniform CA graphs. Additionally, it can be verified that F_{uniform}^A is globally invariant due to the fact that the answer of A is independent of the vertex labels. Also, the function F_{uniform}^A can be computed in logspace using queries of the form $T \in \Delta(G, A)$. Observe that $\Delta(G, A)$ only provides n^3 bits of information with $n = |V(G)|$ and therefore can be computed 'in a single query' by a functional oracle which outputs the n^3 bits of information.   ◀

5   Non-Uniform CA Graphs and Restricted CA Matrices

In the first part of this section we examine the structure of non-uniform CA graphs. Every such graph must have two ov-triangles which have exactly one vertex in common and both are representable as interval triangle and as non-Helly triangle. This pair of ov-triangles enforces a particular structure in non-uniform CA graphs. In the second part we introduce restricted CA matrices, which try to partly capture this structure. Roughly speaking, restricted CA matrices can be seen as a generalization of the neighborhood matrices of non-uniform CA graphs. We pay the price of considering this more general class of structures in order to provide a logspace reduction from the canonical representation problem for CA graphs to that of restricted CA matrices.

▶ **Definition 5.1.** Given a CA graph G, an induced 4-cycle $C = (u, w, w', u')$ of G and $v \in V(G) \setminus C$. We say (C, v) is a non-uniformity witness of G if $\{u, v, w\}, \{u', v, w'\} \in \mathcal{T}_{\text{IT}}(G) \cap \mathcal{T}_{\text{NHT}}(G)$. We also simply call (C, v) a witness of G.

▶ **Theorem 5.2.** *A CA graph G is non-uniform iff G has a non-uniformity witness.*

Proof. "⇒": Let G be a non-uniform CA graph. Due to Theorem 4.7 there exists an ov-triangle T of G with $T \in \mathcal{T}_{\text{IT}}(G) \cap \mathcal{T}_{\text{NHT}}(G)$. Let $T = \{u, v, w\}$ and $\rho_{\text{I}} \in \mathcal{N}(G)$ such that v is in-between u and w, i.e. $\rho_{\text{I}}(v) \subset \rho_{\text{I}}(u) \cup \rho_{\text{I}}(w)$. First, we show that there exists an induced 4-cycle $C = (u, w, w', u')$ in G.

From the non-Helly triangle representation of T it follows that $N[u] \cup N[v] \cup N[w] = V(G)$. Since v is in-between u and w this means $N[u] \cup N[w] = V(G)$. It holds that u and w overlap. Therefore one of the conditions in the definition of the neighborhood matrix for u and w to form a circle cover must be violated. Let us assume w.l.o.g. that the violated condition is that there exists a $u' \in N[u] \setminus N[w]$ such that $N[u'] \not\subseteq N[u]$. This means u' must overlap with u
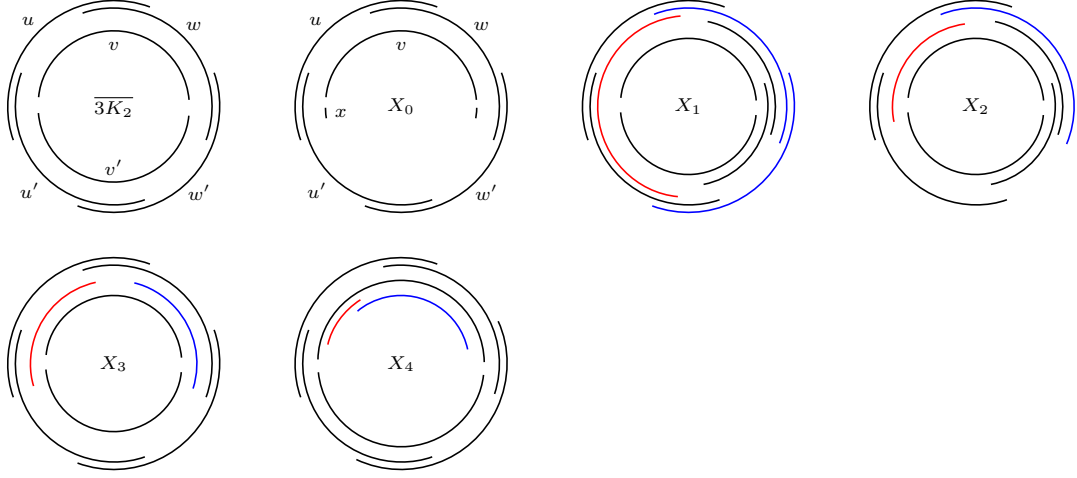
■ **Figure 7** Examples of non-uniform CA graphs and one uniform CA graph X_4

and there exists a $w' \in N[u'] \setminus N[u]$. Since $w' \notin N[u]$ it follows from $N[u] \cup N[w] = V(G)$ that $w' \in N[w]$ and because w is disjoint from u', and because w' intersects with both u' and w it follows that w' overlaps with u' and w. Therefore $C = (u, w, w', u')$ is an induced 4-cycle in G.

It remains to show that $\{u', v, w'\}$ is an ov-triangle and that it is in both $\mathcal{T}_{\mathrm{IT}}(G)$ and $\mathcal{T}_{\mathrm{NHT}}(G)$. Consider the representation ρ_{I} from before. Assume for the sake of contradiction that v does not overlap with u'. Then due to ρ_{I} it must be the case that u' is disjoint from v and thus $u' \in N_T(u)$. However, due to fact that T is representable as non-Helly triangle this would imply that u' is contained by u, which is not the case. Therefore u' overlaps with v as the other intersections types are out of question. For the same reason w' overlaps with v and hence $T' = \{u', v, w'\}$ is an ov-triangle. Now, it can be verified that in every representation of G where T is a non-Helly triangle it follows that T' must be an interval triangle and vice versa. This concludes that T' is in $\mathcal{T}_{\mathrm{IT}}(G) \cap \mathcal{T}_{\mathrm{NHT}}(G)$.

"⇐": Follows directly from Theorem 4.7. ◄

In Figure 7 five non-uniform CA graphs and one uniform CA graph (X_4) are given by their CA models. We explain how to verify this claim. First, we have to check that every CA model is normalized. This means the graphs which are induced by these models must be twin-free and without a universal vertex. Additionally, the intersection types of the arcs must match the intersection types in the induced graph (or more precisely its neighborhood matrix). A quick way to determine whether two overlapping arcs also overlap in the graph is to check if they jointly occur in an induced n-cycle for some $n \geq 4$.

To see that the first five CA graphs are non-uniform we have to find an ov-triangle that is representable as both interval and non-Helly triangle. In the case of $\overline{3K_2}$ this ov-triangle can be $\{u, v, w\}$. In the given representation $\{u, v, w\}$ is represented as interval triangle. Observe that v and v' are in the same orbit and therefore the labels v and v' can be swapped in the representation. After swapping v and v' the ov-triangle $\{u, v, w\}$ is represented as non-Helly triangle. For the graph X_0 we can also choose the ov-triangle $\{u, v, w\}$. In this case there is an automorphism which swaps u with u' and w with w' and has the other vertices as fix-points. After changing the labels in the representation according to this automorphism it

	1				2				3		4		5		6		7
u	cs	ov	cs	ov	di	di	di	di	ov	ov	ov	cs	ov	ov	di	di	ov
w	di	di	di	di	cs	ov	cs	ov	di	di	ov	ov	ov	cs	ov	ov	ov
w'	di	di	di	di	cs	cs	ov	ov	ov	ov	di	di	ov	ov	ov	cs	ov
u'	cs	cs	ov	ov	di	di	di	di	ov	cs	ov	ov	di	di	ov	ov	ov

■ **Table 2** Intersection types of restricted CA matrices with witness cycle (u, w, w', u)

holds that $\{u, v, w\}$ is represented as non-Helly triangle. We remark that $\overline{3K_2}$ and X_0 are minimal in the sense that no induced subgraph of them is a non-uniform CA graph. Next, let us consider the graphs X_1 to X_3. Observe that the black arcs in each of these graphs form an induced $\overline{3K_2}$ subgraph. We assume that the black arcs are labeled with u, u', v, v', w, w' in the same way that the representation of $\overline{3K_2}$ is labeled. It holds that v and v' are in the same orbit in all of these four graphs because they have the same open neighborhood. Therefore $\{u, v, w\}$ is representable as both interval and non-Helly triangle due to the same argument that we made for $\overline{3K_2}$.

To show that X_4 is uniform we argue that it has a unique normalized representation, i.e. $|\mathcal{N}(X_4)| = 1$. Observe that this graph has a unique CA model. Additionally, it has no non-trivial automorphism (it is rigid). Therefore X_4 has a unique CA representation.

▶ **Fact 5.3.** *Every non-uniform CA graph contains $\overline{3K_2}$ or X_0 as induced subgraph.*

Proof. Let G be a non-uniform CA graph. Due to Theorem 5.2 there exists a witness (C, v) of G with $C = (u, w, w', u')$. Since G does not contain a universal vertex it holds that $V(G) \setminus N[v]$ is non-empty. Due to the fact that $\{u, v, w\}$ and $\{u', v, w'\}$ can be represented as interval triangles it follows that $N_C(C \setminus \{x\}) \subseteq N[v]$ for all $x \in C$. Therefore $V(G) \setminus N[v] \subseteq N_C(C) \cup N_C(u, u') \cup N_C(w, w')$. Suppose there is a $v' \in N_C(C) \setminus N[v]$. Then the vertices of C along with v and v' form an induced $\overline{3K_2}$-subgraph of G. Assume that this is not the case, i.e. $N_C(C) \subseteq N[v]$. Since u and v overlap it must hold that $N[u] \setminus N[v] \neq \emptyset$. The only vertices that can be adjacent to $N[u]$ but not to $N[v]$ must be in $N_C(u, u')$ since $N_C(C) \subseteq N[v]$. Therefore there exists a vertex $x \in N_C(u, u')$ that is not adjacent to v. For the same reason there must be a vertex $y \in N_C(w, w')$ not adjacent to v because $N[w] \setminus N[v] \neq \emptyset$. The vertices of C along with v, x and y form an induced X_0-subgraph. ◀

▶ **Definition 5.4** (Restricted CA Matrix)**.** Let λ be a CA matrix. We say λ is a restricted CA matrix if it contains an induced 4-cycle $C = (u, w, w', u')$ called witness cycle such that:
1. $N_C(u, w)$, $N_C(u', w')$ and $N_C(x)$ are empty for every $x \in C$
2. For all $x \in N_C(C)$ it holds that x overlaps with all vertices in C

Observe that the intersection matrix of every CA model that is shown in Figure 7 is a restricted CA matrix.

▶ **Fact 5.5.** *Given an intersection matrix λ, vertices $x, y_1, \ldots, y_k$ of λ and intersection types $\alpha_1, \ldots, \alpha_k$, we say x is an $(\alpha_1, \ldots, \alpha_k)$-neighbor of $(y_1, \ldots, y_k)$ if $\lambda_{x,y_i} = \alpha_i$ for all $i \in [k]$. A CA matrix λ is restricted iff λ contains an induced 4-cycle $C = (u, w, w', u')$ such that for all vertices $x \in V(\lambda) \setminus C$ there exists a column $\overline{\alpha}$ in Table 2 such that x is a $\overline{\alpha}$-neighbor of C.*

Proof. We use the numbers in the table headline to refer to the different columns. For example, 2.3 refers to the third column from left in the second part of the table: (di, cs, ov, di).

"$\Rightarrow$": Let λ be a restricted CA matrix with witness cycle $C = (u, w, w', u')$. We need to show for every $x \in V(\lambda) \setminus C$ there exists a column $\overline{\alpha}$ in Table 2 such that x is a $\overline{\alpha}$-neighbor of C. Due to the definition of restricted CA matrices it must hold that x is in (exactly) one of the following seven sets: $N_C(C), N_C(u, u'), N_C(w, w')$ or $N_C(C \setminus \{z\})$ for a $z \in C$. If x is in $N_C(C)$ then x overlaps with every vertex of C by definition. This corresponds to the last column 7.1 of the table. If $x \in N_C(u, u')$ then x is disjoint from w and w'. In that case x is an $\overline{\alpha}$-neighbor of C where $\overline{\alpha}$ must be one of the four columns in part one of the table. For the same reason if $x \in N_C(w, w')$ then it is an $\overline{\alpha}$-neighbor of C where $\overline{\alpha}$ corresponds to one of the two columns in the second part of the table. If x is in $N_C(C \setminus \{w\})$ then x is disjoint from w and x overlaps with both u and w'. The intersection type between x and u' can be one of the following: x overlaps with u or x is contained by u or x contains u. The first two cases are covered by the third part of the table. However, if x contains u then there exists no corresponding column in the table since it does not have any `cd`-entries. This can be resolved by using the following observation: if x is in $N_C(C \setminus \{w\})$ and contains u' then (u, w, w', x) is a witness cycle of λ as well. As a consequence we can assume without loss of generality that a witness cycle C of λ can be chosen such that there exists no $x \in N_C(C \setminus \{w\})$ which contains u'. The same argument applies to the remaining three cases $x \in N_C(C \setminus \{z\})$ with $z \in \{u, u', w'\}$.

"$\Leftarrow$": clear. ◀

In the remainder of this section we prove that the canonical representation problem for CA graphs is logspace-reducible to the canonical representation problem for vertex-colored restricted CA matrices. The proof outline looks as follows. First, we define a subset of uniform CA graphs, namely Δ-uniform CA graphs, for which the globally invariant non-Helly triangle representability problem can be solved in logspace. Therefore the canonical representation problem for CA graphs is logspace-reducible to that of CA graphs which are not Δ-uniform. This reduction follows from a slightly modified version of Lemma 4.14. Then we show that the neighborhood matrix of a non-Δ-uniform CA graph can be converted into a vertex-colored restricted CA matrix by flipping 'long' arcs. By coloring the flipped arcs the isomorphism type is preserved.

▶ **Definition 5.6.** For a graph G we define Δ_G as the following set of `ov`-triangles (see Definition 4.6). An `ov`-triangle T of G is in Δ_G if there exist three pairwise different vertices u, v, w in T such that the following holds:
1. $N[u] \cup N[v] \cup N[w] = V(G)$
2. For all $z \in T$ it holds that if a vertex $x \in N_T(z)$ then x `cd` z
3. If there exist u', w' such that (u, w, w', u') is an induced 4-cycle and v overlaps with u' and w' then $N[v] \subseteq N[u'] \cup N[w']$

▶ **Definition 5.7.** A CA graph G is Δ-uniform if $\Delta_G \cap \mathcal{T}_{\mathrm{IT}}(G) = \emptyset$.

Let us explain the intuition behind these two definitions. The set Δ_G approximates $\mathcal{T}_{\mathrm{NHT}}(G)$. More precisely, whenever an `ov`-triangle $T = \{u, v, w\}$ is in $\mathcal{T}_{\mathrm{NHT}}(G)$ this implies that T satisfies certain constraints such as for example $N[u] \cup N[v] \cup N[w] = V(G)$. The set Δ_G consists of three such constraints. Therefore if an `ov`-triangle is representable as non-Helly triangle it must also be in Δ_G, i.e. $\mathcal{T}_{\mathrm{NHT}}(G) \subseteq \Delta_G$. The Δ-uniform CA graphs can be alternatively seen as the subset of uniform CA graphs where the constraints of Δ_G suffice to characterize $\mathcal{T}_{\mathrm{NHT}}(G)$, i.e. $\Delta_G = \mathcal{T}_{\mathrm{NHT}}(G)$.

▶ **Lemma 5.8.** *For every graph G it holds that $\mathcal{T}_{\mathrm{NHT}}(G) \subseteq \Delta_G$. If G is a Δ-uniform CA graph then $\mathcal{T}_{\mathrm{NHT}}(G) = \Delta_G$.*

Proof. For the first claim consider a graph G. If G is not a CA graph then $\mathcal{T}_{\mathrm{NHT}}(G) = \emptyset$. Therefore we can assume that G is a CA graph. Given an ov-triangle $T \in \mathcal{T}_{\mathrm{NHT}}(G)$ we show that it must be in Δ_G. Let $\rho \in \mathcal{N}(G)$ be a representation such that $T = \{u, v, w\}$ is represented as non-Helly triangle in it. Since $\rho(u) \cup \rho(v) \cup \rho(w)$ covers the whole circle it follows that $N[u] \cup N[v] \cup N[w] = V(G)$, which is the first condition of Definition 5.6. To see that the second condition holds we consider a vertex $x \in N_T(u)$ without loss of generality. Since x is not adjacent to v and w it holds that $\rho(x) \subseteq \mathbb{C} \setminus (\rho(v) \cup \rho(w))$ where $\mathbb{C}$ denotes the whole circle. Since $\mathbb{C} \setminus (\rho(v) \cup \rho(w)) \subset \rho(u)$ it follows that $\rho(x) \subset \rho(u)$. Due to the fact that ρ is a normalized representation this implies that x is contained by u. To see that the third condition of Δ_G holds let u', w' be vertices such that (u, w, w', u') is an induced 4-cycle of G. Since T is represented as non-Helly triangle in ρ it must hold that $\{u', v, w'\}$ is an interval triangle in ρ with $\rho(v) \subset \rho(u') \cup \rho(w')$ and therefore $N[v] \subseteq N[u'] \cup N[w']$.

For the second claim let G be a Δ-uniform CA graph. From the previous claim we know that $\mathcal{T}_{\mathrm{NHT}}(G) \subseteq \Delta_G$. Since every ov-triangle must be in $\mathcal{T}_{\mathrm{NHT}}(G) \cup \mathcal{T}_{\mathrm{IT}}(G)$ it follows that $\Delta_G \subseteq \mathcal{T}_{\mathrm{NHT}}(G) \cup \mathcal{T}_{\mathrm{IT}}(G)$. The definition of Δ-uniform requires $\Delta_G \cap \mathcal{T}_{\mathrm{IT}}(G) = \emptyset$ and thus $\Delta_G \subseteq \mathcal{T}_{\mathrm{NHT}}(G)$. ◀

▶ **Fact 5.9.** Δ-*uniform CA graphs are a strict subset of uniform CA graphs.*

Proof. Assume there exists a Δ-uniform CA graph G which is not uniform. This means there exists an ov-triangle $T \in \mathcal{T}_{\mathrm{NHT}}(G) \cap \mathcal{T}_{\mathrm{IT}}(G)$. Due to the previous lemma it holds that $\mathcal{T}_{\mathrm{NHT}}(G) \subseteq \Delta_G$. This implies that $T \in \Delta_G \cap \mathcal{T}_{\mathrm{IT}}(G)$ which contradicts that G is Δ-uniform. Therefore every Δ-uniform CA graph is uniform.

An example of a uniform CA graph that is not Δ-uniform is the graph X_4 in Figure 7. In the third paragraph after Theorem 5.2 we argued that X_4 is a uniform CA graph because it has a unique normalized representation. Assume that the black arcs of X_4 are labeled with u, u', v, v', w, w' in the same way that the representation of $\overline{3K_2}$ is labeled in Figure 7. To see that X_4 is not Δ-uniform it suffices to check that the ov-triangle $\{u, v, w\}$ is in Δ_{X_4} and represented as interval triangle. ◀

▶ **Corollary 5.10.** *The globally invariant non-Helly triangle representability problem for Δ-uniform CA graphs can be solved in logspace.*

Proof. Given a CA graph G and an ov-triangle T output yes iff $T \in \Delta_G$. This is correct because in the case of a Δ-uniform CA graph G it holds that $\Delta_G = \mathcal{T}_{\mathrm{NHT}}(G)$ (Lemma 5.8). Clearly, Δ_G is computable in logspace and an invariant. ◀

▶ **Lemma 5.11.** *Let G be a CA graph that is not Δ-uniform. Then there exists an induced 4-cycle $C = (u, w, w', u')$ such that $N[u] \cup N[w] = N[u'] \cup N[w'] = V(G)$ and a vertex v that overlaps with every vertex in C.*

Proof. The argument is essentially the same as the one made for the "$\Rightarrow$"-direction in the proof of Theorem 5.2. The difference is that instead of the stronger assumption that $T \in \mathcal{T}_{\mathrm{NHT}}(G)$ we only require that $T \in \Delta_G$.

Since G is not Δ-uniform there exists an ov-triangle $T = \{u, v, w\}$ of G such that $T \in \Delta_G$ and there is a representation $\rho \in \mathcal{N}(G)$ such that T is represented as interval triangle in ρ. Furthermore, let us assume w.l.o.g. that $\rho(v) \subset \rho(u) \cup \rho(w)$. Since $T \in \Delta_G$ it holds that $N[u] \cup N[v] \cup N[w] = V(G)$. Due to the interval representation of T in ρ it follows that $N[u] \cup N[w] = V(G)$. Since u and w do not form a circle cover it must hold that there exists a vertex $u' \in N[u] \setminus N[w]$ such that $N[u'] \setminus N[u]$ is non-empty. If u' is disjoint from v it follows that u' must be contained by u from the second condition in Definition 5.6 of

Δ_G. This cannot be the case and therefore $u' \in N_T(u, v)$. For u' to have a neighbor which is not adjacent to u it must hold that $\rho(u') \not\subseteq \rho(u)$. Therefore u' overlaps with u and v. Let $w' \in N[u'] \setminus N[u]$. If $w' \in N_T(w)$ then w' would be contained by w due to the second condition of Δ_G. Again, this cannot be the case and therefore $w' \in N_T(v, w)$. From the representation ρ it follows that w must overlap with u', v and w. Then $C = (u, w, w', u')$ is an induced 4-cycle of G such that v overlaps with every vertex of C. It remains to show that $N[u'] \cup N[w'] = V(G)$. Due to the third condition of Δ_G it holds that $N[v] \subseteq N[u'] \cup N[w']$. Additionally, it holds that $\rho(u) \setminus \rho(v) \subset \rho(u')$ and $\rho(w) \setminus \rho(v) \subset \rho(w')$. As a consequence $N[u'] \cup N[w'] = V(G)$. ◀

▶ **Corollary 5.12.** *Canonical representations for CA graphs without induced 4-cycle can be computed in logspace.*

Proof. By Lemma 5.11 the class of CA graphs without induced 4-cycle is a subset of Δ-uniform CA graphs and due to Corollary 5.10 and Theorem 4.12 a canonical representation for such graphs can be computed in logspace. ◀

▶ **Corollary 5.13.** *Helly CA graphs are a strict subset of Δ-uniform CA graphs.*

Proof. Assume G is a Helly CA graph which is not Δ-uniform. Then due to Lemma 5.11 there exists an induced 4-cycle C and a vertex v not in C which overlaps with every vertex in C. In any normalized representation of G it must hold that v forms a non-Helly triangle with two vertices from C. This contradicts that G is Helly. The graph ⋈ is a Δ-uniform CA graph which is not Helly. ◀

▶ **Theorem 5.14.** *The canonical representation problem for CA graphs is logspace-reducible to the canonical representation problem for vertex-colored restricted CA matrices.*

Proof. For brevity let $\mathcal{Z}$ denote the set of all CA graphs which are not Δ-uniform. Since the globally invariant non-Helly triangle representability problem for Δ-uniform CA graphs can be solved in logspace (see Corollary 5.10) it follows from a modified version of Lemma 4.14 that the canonical representation problem for CA graphs is logspace-reducible to the canonical representation problem for vertex-colored $\mathcal{Z}$. To see this replace 'uniform' with 'Δ-uniform' and 'non-uniform' with 'non-Δ-uniform' in the statement (and proof) of Lemma 4.14.

For a CA graph G let us say a subset of vertices X of G is an R-flip set if $\lambda_G^{(X)}$ is a restricted CA matrix. To find canonical representations for $\mathcal{Z}$ we construct an invariant vertex set selector f such that $f(G)$ contains at least one R-flip set for every $G \in \mathcal{Z}$. Then to obtain a canonical representation for $G \in \mathcal{Z}$ let $\hat{X}$ denote the R-flip set in $f(G)$ such that $\text{canon}(\lambda_G^{(\hat{X})}, c_{\hat{X}})$ is lexicographically minimal with c_X being the coloring which assigns every vertex $v \in X$ the color red and the other vertices are blue. Let ρ be a canonical normalized representation for $(\lambda_G^{(\hat{X})}, c_{\hat{X}})$. Then $\rho^{(\hat{X})}$ is a canonical representation for G. Notice, that $\rho^{(\hat{X})}$ can be computed in logspace by computing canonical representations for vertex-colored restricted CA matrices. The correctness of this approach follows from the same argument made in the proof of Theorem 3.5 in the flip trick section. The analogy is straightforward. The R-flip sets in this context correspond to flip sets and the invariant vertex set selector f takes the place of the invariant flip set function. Given a CA graph G and $X \subseteq V(G)$ it can be easily checked in logspace whether $\lambda_G^{(X)}$ is a restricted CA matrix.

For a CA graph G let $C(G)$ denote the set of all ordered induced 4-cycles in G. Now, we claim that the following logspace-computable function f is an invariant vertex set selector with the desired property:

$$f(G) = \bigcup_{C \in C(G)} \big\{ \{ x \in V(G) \setminus C \mid \exists y \in C : x \ \mathtt{cs} \ y \} \big\}$$

It is not difficult to check that f is invariant. It remains to argue why $f(G)$ contains at least one R-flip set for every $G \in \mathcal{Z}$. Let $G \in \mathcal{Z}$ and $C = (u, w, w', u')$ is an induced 4-cycle in G such that $N[u] \cup N[w] = N[u'] \cup N[w'] = V(G)$. The existence of such an induced 4-cycle is guaranteed by Lemma 5.11. Observe that if there exists a $u_1 \in N_C(u, w, u')$ with $u_1 \ \mathtt{cs} \ u$ then $C_1 = (u_1, w, w', u')$ also satisfies the previous condition $N[u_1] \cup N[w] = V(G)$. Therefore we can assume that there exists no $z \in C$ and $z_1 \in N_C(N[z] \cap C)$ such that $z_1 \ \mathtt{cs} \ z$. From $N[u] \cup N[w] = N[u'] \cup N[w'] = V(G)$ it immediately follows that $N_C(u, w)$, $N_C(u', w')$ and $N_C(x)$ are empty for every $x \in C$.

We prove that $\lambda^{(X)}$ is a restricted CA matrix with witness cycle C where λ is the neighborhood matrix of G and $X = \{ x \in V(G) \setminus C \mid \exists y \in C : x \ \mathtt{cs} \ y \}$. Note that $X \in f(G)$ via C. To reference the neighborhoods of G (which are the same as the ones of λ) or $\lambda^{(X)}$ we write N^G and $N^{\lambda^{(X)}}$ to distinguish between them. First, we show that $N_C^{\lambda^{(X)}}(u, w) = \emptyset$. Assume the opposite, i.e. there exists $x \in N_C^{\lambda^{(X)}}(u, w)$. If x was not flipped, i.e. $x \notin X$, then it also holds that $x \in N_C^G(u, w)$, which contradicts that $N_C^G(u, w)$ is empty. If x was flipped, i.e. $x \in X$, then it must be the case that x contains u' and w' in λ. This means $N_G[u'] \cup N_G[w'] \subseteq N_G[x]$ which implies that x is a universal vertex in G since $N_G[u'] \cup N_G[w'] = V(G)$, contradiction. For the same reason it holds that $N_C^{\lambda^{(X)}}(u', w')$ and $N_C^{\lambda^{(X)}}(z)$ are empty for all $z \in C$. It remains to show that for all $x \in N_C^{\lambda^{(X)}}(C)$ it holds that x overlaps with all vertices of C in $\lambda^{(X)}$. Notice that $\lambda_{x,z}^{(X)} \in \{ \mathtt{ov}, \mathtt{cs}, \mathtt{cc} \}$ for every $z \in C$. Otherwise x would not be in $N_C(C)$. We consider the following two cases: in the first one we assume that x contains one vertex of C in $\lambda^{(X)}$ and in the second one we assume that x forms a circle cover with one vertex of C in $\lambda^{(X)}$. We prove that neither of these cases can occur and therefore x must overlap with all vertices of C in $\lambda^{(X)}$. For the first case assume that w.l.o.g. x contains u in $\lambda^{(X)}$ and intersects with the other vertices of C in $\lambda^{(X)}$. If $x \in X$ then it was flipped. It follows that x was disjoint from u in λ and therefore $x \in N_C^G(w, w', u')$. Since $x \in X$ it also must hold that x contains at least one of the vertices w, w', u' in G. It follows that x contains w' since it cannot contain the other two in λ. However, this contradicts our choice of C which says that there exists no $w_1' \in N_C^G(w, w', u')$ such that w_1' contains w' in λ. If $x \notin X$ then it must hold that x already contained u in λ. But then x should be in X, contradiction. For the second case assume x forms a circle cover with u in $\lambda^{(X)}$. If x forms a circle cover with u then this implies that x contains w' in $\lambda^{(X)}$ and therefore this reduces to the first case. We conclude that both conditions of Definition 5.4 are satisfied and hence $\lambda^{(X)}$ is a restricted CA matrix.   ◀

6   Further Research

Finding a polynomial-time isomorphism test for CA graphs remains an open problem. We have shown that it suffices to consider only non-uniform CA graphs for this problem. This particular class of CA graphs offers quite a lot of structure, which is caused by what we named non-uniformity witnesses. It seems plausible that such witnesses can be exploited to devise an isomorphism test. Additionally, we proved that the canonical representation problem for CA graphs is logspace-reducible to that of restricted CA matrices. The central question with regard to the flip trick is how invariant flip sets for restricted CA matrices or non-uniform CA graphs can be computed. Also, we remark that CA representations for CA graphs can be computed in logspace if flip sets for restricted CA matrices can be found in

logspace. Another interesting problem is to extend Definition 5.6 of Δ_G such that it captures $\mathcal{T}_{\mathrm{NHT}}(G)$ on uniform CA graphs, i.e. $\Delta_G = \mathcal{T}_{\mathrm{NHT}}(G)$ for all uniform CA graphs G. If this can be done in such a way that Δ_G remains an invariant and computable in logspace then everything that is said about Δ-uniform CA graphs in section 5 also applies to uniform CA graphs.

―――― **References** ――――

1 Yixin Cao, Luciano N. Grippo, and Martín D. Safe. Forbidden induced subgraphs of normal Helly circular-arc graphs: Characterization and detection. *Discrete Applied Mathematics*, 216, Part 1:67 – 83, 2017. Special Graph Classes and Algorithms — in Honor of Professor Andreas Brandstädt on the Occasion of His 65th Birthday. URL: `http://www.sciencedirect.com/science/article/pii/S0166218X15004308`, `doi:http://dx.doi.org/10.1016/j.dam.2015.08.023`.

2 Maurice Chandoo. Deciding Circular-Arc Graph Isomorphism in Parameterized Logspace. In Nicolas Ollinger and Heribert Vollmer, editors, *33rd Symposium on Theoretical Aspects of Computer Science (STACS 2016)*, volume 47 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:13, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: `http://drops.dagstuhl.de/opus/volltexte/2016/5727`, `doi:http://dx.doi.org/10.4230/LIPIcs.STACS.2016.26`.

3 Andrew Curtis, Min Lin, Ross McConnell, Yahav Nussbaum, Francisco Soulignac, Jeremy Spinrad, and Jayme Szwarcfiter. Isomorphism of graph classes related to the circular-ones property. *Discrete Mathematics and Theoretical Computer Science*, 15(1), 2013. URL: `http://www.dmtcs.org/dmtcs-ojs/index.php/dmtcs/article/view/2298`.

4 Elaine Marie Eschen. *Circular-Arc Graph Recognition and Related Problems*. PhD thesis, Vanderbilt University, Nashville, TN, USA, 1998. UMI Order No. GAX98-03921.

5 F. Gavril. Algorithms on circular-arc graphs. *Networks*, 4(4):357–369, 1974. URL: `http://dx.doi.org/10.1002/net.3230040407`, `doi:10.1002/net.3230040407`.

6 Martin Charles Golumbic. *Algorithmic Graph Theory and Perfect Graphs (Annals of Discrete Mathematics, Vol 57)*. North-Holland Publishing Co., Amsterdam, The Netherlands, The Netherlands, 2004.

7 Wen-Lian Hsu. $O(M \cdot N)$ algorithms for the recognition and isomorphism problems on circular-arc graphs. *SIAM J. Comput.*, 24(3):411–439, June 1995. URL: `http://dx.doi.org/10.1137/S0097539793260726`, `doi:10.1137/S0097539793260726`.

8 Benson L. Joeris, Min Chih Lin, Ross M. McConnell, Jeremy P. Spinrad, and Jayme L. Szwarcfiter. Linear-time recognition of Helly circular-arc models and graphs. *Algorithmica*, 59(2):215–239, 2011. URL: `http://dx.doi.org/10.1007/s00453-009-9304-5`, `doi:10.1007/s00453-009-9304-5`.

9 Johannes Köbler, Sebastian Kuhnert, Bastian Laubner, and Oleg Verbitsky. Interval graphs: Canonical representations in logspace. *SIAM Journal on Computing*, 40(5):1292–1315, 2011. `doi:10.1137/10080395X`.

10 Johannes Köbler, Sebastian Kuhnert, and Oleg Verbitsky. Helly circular-arc graph isomorphism is in logspace. In *Mathematical Foundations of Computer Science 2013*, volume 8087 of *Lecture Notes in Computer Science*, pages 631–642. Springer Berlin Heidelberg, 2013. URL: `http://dx.doi.org/10.1007/978-3-642-40313-2_56`, `doi:10.1007/978-3-642-40313-2_56`.

11   Johannes Köbler, Sebastian Kuhnert, and Oleg Verbitsky.   Solving the canonical representation and star system problems for proper circular-arc graphs in logspace. *Journal of Discrete Algorithms*, 38-41(Supplement C):38 – 49, 2016.   URL: `http://www.sciencedirect.com/science/article/pii/S1570866716300016`, `doi:https://doi.org/10.1016/j.jda.2016.03.001`.

12   Min Chih Lin and Jayme L. Szwarcfiter. Characterizations and recognition of circular-arc graphs and subclasses: A survey.   *Discrete Mathematics*, 309(18):5618 – 5635, 2009.   Combinatorics 2006, A Meeting in Celebration of Pavol Hell's 60th Birthday (May 1–5, 2006).   URL: `http://www.sciencedirect.com/science/article/pii/S0012365X08002161`, `doi:http://dx.doi.org/10.1016/j.disc.2008.04.003`.

13   George S. Lueker and Kellogg S. Booth. A linear time algorithm for deciding interval graph isomorphism. *J. ACM*, 26(2):183–195, April 1979. URL: `http://doi.acm.org/10.1145/322123.322125`, `doi:10.1145/322123.322125`.

14   Ross M. McConnell.   Linear-time recognition of circular-arc graphs.   *Algorithmica*, 37(2):93–147, 2003.   URL: `http://dx.doi.org/10.1007/s00453-003-1032-7`, `doi:10.1007/s00453-003-1032-7`.

15   Franklin W. Stahl. Circular genetic maps. *Journal of Cellular Physiology*, 70(S1):1–12, 1967. URL: `http://dx.doi.org/10.1002/jcp.1040700403`, `doi:10.1002/jcp.1040700403`.

16   Alan Tucker. Characterizing circular-arc graphs. *Bull. Amer. Math. Soc.*, 76(6):1257–1260, 11 1970. URL: `http://projecteuclid.org/euclid.bams/1183532398`.

17   Michael S. Waterman and Jerrold R. Griggs.   Interval graphs and maps of dna.   *Bulletin of Mathematical Biology*, 48(2):189 – 195, 1986.   URL: `http://www.sciencedirect.com/science/article/pii/S0092824086800064`, `doi:http://dx.doi.org/10.1016/S0092-8240(86)80006-4`.

18   Tsong-Ho Wu. *An $O(n^3)$ Isomorphism Test for Circular-Arc Graphs*. PhD thesis, SUNY Stony Brook, New York, NY, USA, 1983.