Deciding Circular-Arc Graph Isomorphism in Parameterized Logspace

Maurice Chandoo

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

Abstract

We compute a canonical circular-arc representation for a given circular-arc (CA) graph which implies solving the isomorphism and recognition problem for this class. To accomplish this we split the class of CA graphs into uniform and non-uniform ones and employ a generalized version of the argument given by Köbler et al. (2013) that has been used to show that the subclass of Helly CA graphs can be canonized in logspace. For uniform CA graphs our approach works in logspace and in addition to that Helly CA graphs are a strict subset of uniform CA graphs. Thus our result is a generalization of the canonization result for Helly CA graphs. In the non-uniform case a specific set Ω of ambiguous vertices arises. By choosing the parameter k to be the cardinality of Ω this obstacle can be solved by brute force. This leads to an $\mathcal{O}(k + \log n)$ space algorithm to compute a canonical representation for non-uniform and therefore all CA graphs.

1 Introduction

An arc is a connected set of points on the circle. A graph G is called a CA graph if every vertex v can be assigned to an arc $\rho(v)$ such that two vertices u, v are adjacent iff their respective arcs intersect, i.e. $\rho(u) \cap \rho(v) \neq \emptyset$. We call such a (bijective) mapping ρ a CA representation of G and the set of arcs $\rho(G) = \{\rho(v) \mid v \in V(G)\}$ a CA model. The recognition problem for CA graphs is to decide whether a given graph G is a CA graph. The canonical representation problem for CA graphs consists of computing a CA representation ρ_G for a given CA graph G with the additional canonicity constraint that whenever two CA graphs G and H are isomorphic the CA models $\rho_G(G)$ and $\rho_H(H)$ are identical.

Similarly, a graph is an interval graph if every vertex can be assigned to an interval on a line such that two vertices share an edge iff their intervals intersect. It is easy to see that every interval graph is a CA graph since every interval model is a CA model.

The class of CA graphs started to gain attraction after a series of papers in the 1970's by Alan Tucker. However, there is still no known better upper bound for deciding CA graph isomorphism than for graph isomorphism in general even though considerable effort has been done to this end. There have been two claimed polynomial-time algorithms in [9], [3] which have been disproven in [2], [1] respectively. For the subclass of interval graphs a linear-time algorithm for isomorphism has been described in [7]. A series of newer results show that canonical representations for interval graphs, proper CA graphs and Helly CA graphs (a superset of interval graphs) can be computed in logspace [5, 4, 6]. Furthermore, recognition and isomorphism for interval graphs is logspace-hard[5] and these two hardness results carry over to the class of CA graphs; for recognition the reduction requires a little additional work.
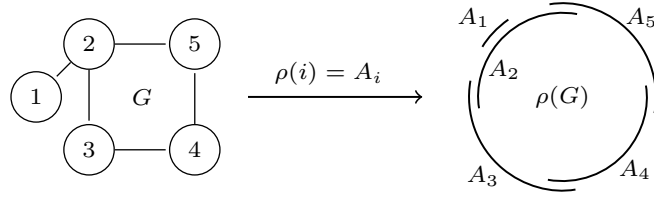
Figure 1 A CA graph and representation.

Our main contribution is that we extend the argument used in [6] to compute canonical representations for HCA graphs to all CA graphs. We split the class of CA graphs into uniform and non-uniform ones and show that the mentioned argument can be applied in both cases using only $\mathcal{O}(k + \log n)$ space. The parameter k describes the cardinality of an obstacle set Ω that occurs only in the non-uniform case. This means $k = 0$ in the uniform case and hence we also obtain a logspace algorithm for uniform CA graphs which are a superclass of HCA graphs. To the best of our knowledge this is the first non-trivial algorithm to decide isomorphism specifically for the class of CA graphs.

This paper is structured as follows. In section 2 we define CA graphs along with their representations and recall the concept of normalized representations from [3]. In section 3 we explain what we mean by flip trick and formalize this with the notions of flip sets and candidate functions. This idea has been used by [8] to compute a CA representation for CA graphs in linear time and [6] has modified it to compute canonical (Helly) CA representations for HCA graphs. In section 4 and 5 the flip trick is applied to uniform and non-uniform CA graphs respectively.

2   Preliminaries

Given two sets A, B we say A and B intersect if $A \cap B \neq \emptyset$. We say A, B overlap, in symbols $A \between B$, if $A \cap B, A \setminus B$ and $B \setminus A$ are non-empty. Let $A = (a_{u,v})_{u,v \in V(A)}, B = (b_{u,v})_{u,v \in V(B)}$ be two square matrices over vertex sets $V(A), V(B)$. We say A and B are isomorphic, in symbols $A \cong B$, if there exists a bijection $\pi \colon V(A) \to V(B)$ such that $a_{u,v} = b_{\pi(u),\pi(v)}$ for all $u, v \in V(A)$; π is called an isomorphism. For two graphs G, H with adjacency matrices A_G, A_H we say that $G \cong H$ if $A_G \cong A_H$. We consider only undirected graphs without self-loops. A graph class $\mathcal{C}$ is a subset of all graphs which is closed under isomorphism, i.e. if $G \in \mathcal{C}$ and $H \cong G$ then $H \in \mathcal{C}$. We define the graph isomorphism problem for a graph class $\mathcal{C}$ as $\mathrm{GI}(\mathcal{C}) = \{(G, H) \mid G, H \in \mathcal{C} \text{ and } G \cong H\}$. For a graph G and a vertex $v \in V(G)$ we define the open neighborhood $N(v)$ of v as the set of vertices that are adjacent to v and the closed neighborhood $N[v] = N(v) \cup \{v\}$. For a subset of vertices $V' \subseteq V(G)$ we define the common neighborhood of V' as $N[V'] = \bigcap_{v \in V'} N[v]$. We also write $N[u, v, \dots]$ instead of $N[\{u, v, \dots\}]$. A vertex v is called universal if $N[v] = V(G)$. For two vertices $u \neq v \in V(G)$ we say that u and v are twins if $N[u] = N[v]$. A twin class is an equivalence class induced by the twin relation. A graph is said to be without twins if every twin class has cardinality one. For a graph G and $S' \subseteq S \subseteq V(G)$ we define the exclusive neighborhood $N_S(S')$ as all vertices $v \in V(G) \setminus S$ such that v is adjacent to all vertices in S' and to none in $S \setminus S'$.

▶ **Definition 1** (Label-independent)**.** Let f be a function which maps graphs to a subset of subsets of vertices, i.e. $f(G) \subseteq \mathcal{P}(V(G))$. We say f is label-independent if for every pair of isomorphic graphs G, H and all isomorphisms π from G to H it holds that

$$f(H) = \{\pi(X) \mid X \in f(G)\} \text{ with } \pi(X) = \{\pi(v) \mid v \in X\} .$$

2.1 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 the opposite part of the circle. We say $\overline{A}$ is obtained by flipping A. When considering a CA model with respect to its intersection structure only the relative position of the endpoints to each other matter. W.l.o.g. all endpoints can be assumed to be pairwise different and no arc covers the full circle. Therefore a CA model $\mathcal{A}$ with n arcs can be described as a unique string as follows. Pick an arbitrary arc $A \in \mathcal{A}$ and relabel the arcs with $1, \ldots, n$ 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 the chosen arc 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 Fig. 1 would result from choosing A_1 $(l(1), r(1), l(2), r(5), l(3), r(2), \ldots)$. In the following we identify $\mathcal{A}$ with its string representation.

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 $\rho[V'] = \{\rho(v) \mid v \in V'\}$. Given a set $X \subseteq V(G)$ we write $\rho^+(X)$ to denote $\cup_{v \in X} \rho(v)$. 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 is not contained in any arc in $\mathcal{A}$. Every such CA model can be understood as interval model by straightening the arcs. Therefore a graph is an interval graph if it admits a CA representation with a hole. A CA graph G is called Helly (HCA graph) if it has a CA representation ρ such that for all maxcliques, i.e. inclusion-maximal cliques, C in G it holds that $\bigcap_{v \in C} \rho(v) \neq \emptyset$. Every interval model has the Helly property and therefore every interval graph is an HCA graph.

2.2 Normalized Representation

In [3] it was observed that the intersection type of two arcs $A \neq B$ can be one of the following five types: A and B are disjoint (`di`), A is contained in B (`cd`), A contains B (`cs`), A and B jointly cover the circle (circle cover `cc`) or A and B overlap (`ov`) but do not jointly cover the circle. Using these types we can associate a matrix with every CA model. An intersection matrix is a square matrix with entries $\{\texttt{di}, \texttt{ov}, \texttt{cs}, \texttt{cd}, \texttt{cc}\}$. Given a CA model $\mathcal{A}$ we define its intersection matrix $\mu_{\mathcal{A}}$ such that $(\mu_{\mathcal{A}})_{a,b}$ 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 (interval) model.

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. If two vertices u, v are adjacent the intersection type of their corresponding arcs might be ambiguous. It would be convenient if the intersection type for every pair of vertices would be uniquely determined by G itself. This can be achieved by associating a graph G with an intersection matrix λ_G

called the neighborhood matrix which is defined for all $u \neq v \in V(G)$ as

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

and the first case applies whose condition is satisfied.

Let μ be an intersection matrix over the vertex set V and $\rho = (\mathcal{A}, f)$ 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 the matrix μ if μ is isomorphic to the intersection matrix of $\mathcal{A}$ via f and denote the set of such CA representations with $\mathcal{N}(\mu)$. Then we say ρ is a normalized CA representation of a graph G if ρ is a CA representation of the neighborhood matrix λ_G of G. An example of a normalized representation can be seen in Fig. 1. Let us denote the set of all normalized CA representations of G with $\mathcal{N}(G) = \mathcal{N}(\lambda_G)$.

▶ **Lemma 2** ([3]). *Every CA graph G without twins and universal vertices has a normalized CA representation, that is $\mathcal{N}(G) \neq \emptyset$.*

For our purpose it suffices to consider only graphs without twins and universal vertices for the same reasons as in [6]. The point is that a universal vertex can be removed from the graph and later added as arc which covers the whole circle in the representation. For each twin class an arbitrary representative vertex can be chosen and colored with the cardinality of its twin class; the other twins are removed.

▶ **Lemma 3.** *The canonical CA representation problem for CA graphs is logspace reducible to the canonical CA representation problem for colored CA graphs without twins and universal vertices.*

Henceforth we assume every graph to be twin-free and without universal vertices and shall only consider normalized representations. For two vertices $u \neq v$ in a graph G we write $u \, \alpha \, v$ instead of $(\lambda_G)_{u,v} = \alpha$. For example, $u \, \texttt{cd} \, v$ indicates that the arc of u must be contained in the arc of v for every (normalized) representation of G.

3 Flip Trick

Let $\mathcal{A}$ be a CA model and $X \subseteq \mathcal{A}$ is a subset of arcs to be flipped. We define the resulting CA model $\mathcal{A}^{(X)} = \left\{ \overline{A} \mid A \in X \right\} \cup \mathcal{A} \setminus X$. Consider a point x on the circle and let X be the set of arcs that contain this point. Then after flipping the arcs in X no other arc contains the point x and thus $\mathcal{A}^{(X)}$ has a hole and therefore must be an interval model. Let μ and $\mu^{(X)}$ be the intersection matrices of $\mathcal{A}$ and $\mathcal{A}^{(X)}$ respectively. It was observed in [8] that the interval matrix $\mu^{(X)}$ can be easily computed using μ and X as input via Table 1. With this the problem of computing a canonical CA representation for colored CA graphs can be reduced to the canonical interval representation problem for colored interval matrices, which can be solved in logspace[6] (the colored part is not mentioned explicitly but can be easily incorporated into the proof by adding the colors to the leaves of the colored Δ tree).

Table 1 Effects of flipping arcs in the intersection matrix

$\mu_{A,B}$	di	cd	cs	cc	ov
$\mu_{\bar{A},B}$	cs	cc	di	cd	ov
$\mu_{A,\bar{B}}$	cd	di	cc	cs	ov
$\mu_{\bar{A},\bar{B}}$	cc	cs	cd	di	ov

The idea is that given a CA graph G if we can compute a set of vertices X as described above then we can obtain a canonical CA representation for G by the following argument. The neighborhood matrix λ_G of G is a CA matrix and the matrix $\lambda_G^{(X)}$ must be an interval matrix. Compute a canonical interval representation for $\lambda_G^{(X)}$ and flip the arcs in X back. This leads to a representation for λ_G and thus G. The required set X can be specified as follows.

▶ **Definition 4.** Let G be a CA graph. Then a non-empty $X \subseteq V(G)$ is a flip set iff there exists a representation $\rho \in \mathcal{N}(G)$ and a point x on the circle such that $v \in X \Leftrightarrow x \in \rho(v)$.

In fact, for the argument to obtain a canonical representation to hold it is only required that X is chosen such that $\lambda_G^{(X)}$ is an interval matrix. However, it can be shown that this is equivalent to the above definition. Now, we can reframe the argument given in [6] as follows.

▶ **Definition 5** (Candidate function). Let $\mathcal{C}$ be a subset of all CA graphs and f is a function which maps graphs to a subset of subsets of their vertices, i.e. $f(G) \subseteq \mathcal{P}(V(G))$. We call f a candidate function for $\mathcal{C}$ if the following conditions hold:
1. For every $G \in \mathcal{C}$ there exists an $X \in f(G)$ such that X is a flip set.
2. f is label-independent.

▶ **Theorem 6.** *If f is a candidate function for all CA graphs that can be computed in logspace then the canonical representation problem for CA graphs can be solved in logspace.*

Proof. Let G be a graph. To decide the recognition problem for CA graphs observe that there is a flip set in $f(G)$ iff G is a CA graph. To verify if a set $X \subseteq V(G)$ is a flip set one can check if $\lambda_G^{(X)}$ is an interval matrix by trying to compute an interval representation.

For the representation problem let G be a CA graph. Let F be the subset of $f(G)$ such that every $X \in F$ is a flip set. By the first condition it holds that F is non-empty. For every $X \in F$ a CA representation ρ_X of G can be computed by the previous argument. We return a CA representation with the lexicographically smallest underlying model

$$\underset{\{\rho_X \mid X \in F\}}{\text{argmin}} \ \rho_X(G)$$

as canonical CA representation. To see that this is indeed a canonical representation consider two isomorphic CA graphs G, H with F_G, F_H defined as F for G previously. Let π be an isomorphism from G to H and $M_G = \{\rho_X(G) \mid X \in F_G\}$ is the set of CA models induced by the flip sets F_G, similarly define M_H. Then canonicity follows by showing $M_G = M_H$. We show that $M_G \subseteq M_H$ as the argument for the other direction is analogous. Let $\mathcal{A}$ be a model in M_G. Let X be a flip set in F_G which induces $\mathcal{A}$, i.e. $\rho_X(G) = \mathcal{A}$. It must hold that $\pi(X) \in f(H)$ since f is label-independent. From $G \cong H$ it follows that $\lambda_G \cong \lambda_H$ and therefore the interval matrices $\lambda_G^{(X)}$ and $\lambda_H^{(\pi(X))}$ are isomorphic meaning that $\pi(X) \in F_H$ since it is a flip set as well. As the interval representations of both interval matrices have

identical underlying models due to canonicity it follows that they remain so after flipping X resp. $\pi(X)$. Therefore it holds that $\mathcal{A} \in M_H$.

This works in logspace since f and the representation ρ_X for every flip set X can be computed in logspace. ◀

So, we have reduced the problem of computing a canonical CA representation for CA graphs to the problem of computing a candidate function for CA graphs.

Let $\mathcal{C}$ and $\mathcal{C}'$ be two graph classes that partition all CA graphs. If $f_\mathcal{C}, f_{\mathcal{C}'}$ are candidate functions for $\mathcal{C}, \mathcal{C}'$ respectively then $f(G) = f_\mathcal{C}(G) \cup f_{\mathcal{C}'}(G)$ is a candidate function for all CA graphs. That f is label-independent follows from label-independent functions being closed under taking unions. The crux here is that we do not need to be able to distinguish if a CA graph G is in $\mathcal{C}$ or $\mathcal{C}'$. Hence, in the next two sections we consider two such classes that partition all CA graphs while avoid dealing with recognition of these two classes.

We complete this section by stating the candidate function used in [6] to canonize HCA graphs and explain why it is a candidate function for this subclass of CA graphs.

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

f_{HCA} always returns at least one flip set for an HCA graph because all maxcliques in an HCA graph are flip sets due to the Helly property and there exists at least one maxclique in every HCA graph that can be characterized as the common neighborhood of two vertices[6]. However, neither of these two properties hold for CA graphs in general.

It remains to argue that f_{HCA} can be computed in logspace and is label-independent. Since the same arguments have to be made for the two candidate functions devised in the next sections we introduce a tool that facilitates this and demonstrate it for f_{HCA}.

▶ **Definition 7.** Let φ be a first-order (FO) formula over graph structures with $k+1$ free variables. Then we define the function f_φ for a graph G as:

$$f_\varphi(G) = \bigcup_{(v_1,\ldots,v_k) \in V(G)^k} \big\{ \{u \in V(G) \mid G \models \varphi(v_1,\ldots,v_k,u)\} \big\}.$$

▶ **Lemma 8.** *For every FO formula φ over graph structures the function f_φ is computable in logspace and label-independent.*

To compute f_φ we can successively evaluate φ and take the union of the results, which can be both done in logspace. To show that f_φ is label-independent it suffices to apply structural induction to φ. Then f_{HCA} is computed by the FO formula $\varphi(u,v,x)$ that states that $x \in N[u,v]$. By Lemma 8 it follows that f_{HCA} is logspace-computable and label-independent.

4 Uniform CA graphs

The difficulty when trying to compute flip sets for CA graphs in general is that for a CA graph G there might be different normalized CA representations such that a set of vertices shares a common point in one representation but not in an other one. We show a subset of CA graphs, namely the uniform CA graphs, where this issue does not occur therefore making it easy to compute flip sets.

For a CA graph G consider an arbitrary vertex u. Looking at the neighbors of u we can try to compute the flip sets $X_{u,1}, X_{u,2}$ specified in Fig. 2. Both sets contain u and all vertices that contain u or form a circle cover with u. The vertices that overlap with u belong
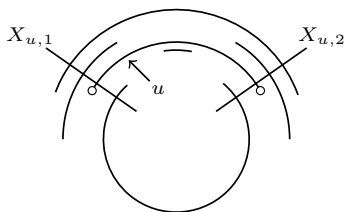
Figure 2 Uniform target flip sets.

to either $X_{u,1}$ or $X_{u,2}$ depending on the side they overlap from with u. Since we cannot determine left and right from the neighborhood matrix we want to express an equivalence relation $\sim_u$ which states that two vertices overlap from the same side with u. With the two equivalence classes induced by $\sim_u$ the two flip sets can be expressed.

Given two vertices x, y that both overlap with u it is for instance easy to see that they must overlap from different sides with u if they are disjoint. The only intersection type between x, y for which the situation is not immediately clear is `ov` as further distinctions are required. An `ov`-triangle is a set of three pairwise overlapping vertices. If x and y overlap then x, y, u form such an `ov`-triangle. Consider the possible normalized representations for an `ov`-triangle. The three vertices can either all jointly cover the circle or be a set of overlapping intervals. In the first case they overlap pairwise but their overall intersection is empty thus we call this a non-Helly triangle and the second case an interval triangle. For an interval triangle there are three different possible representations up to reflection depending on which of the three vertices is placed in-between the other two. If x, y, u is an `ov`-triangle then x and y overlap from the same side with u iff x, y, u form an interval triangle and u is not in-between x, y. We show that it is easy to derive this information in the case of a uniform CA graph G as it does not depend on a representation of G.

▶ **Definition 9.** Let G be a graph. An `ov`-triangle T is in the set Δ_G if the following holds:
1. $\bigcup_{v \in T} N[v] = V(G)$
2. For all $x \in T$ it holds that if a vertex $v \in N_T(x)$ then v `cd` x

▶ **Definition 10** (Uniform CA graph). A CA graph G is uniform if for all `ov`-triangle T in G and $\rho \in \mathcal{N}(G)$ it holds that:

$$T \in \Delta_G \Rightarrow \rho[T] \text{ is a non-Helly triangle}$$

The idea behind the definition of Δ_G is that it captures the properties that an `ov`-triangle must satisfy in the graph if it can be represented as non-Helly triangle. The definition of uniform CA graphs guarantees us that an `ov`-triangle $T \in \Delta_G$ can never be represented as interval triangle. Now, we can show that for the class of uniform CA graphs the property of being a non-Helly triangle and the in-between predicate for interval triangles is invariant across all normalized representations.

▶ **Lemma 11.** *Let G be a uniform CA graph. Then the following statements are equivalent for every* `ov`-*triangle T:*
1. *$T \in \Delta_G$*
2. *$\exists \rho \in \mathcal{N}(G) : \rho[T]$ is a non-Helly triangle*
3. *$\forall \rho \in \mathcal{N}(G) : \rho[T]$ is a non-Helly triangle*

Proof. This immediately follows from the definition of Δ_G and uniform CA graphs. ◀

As a contrasting example of a (non-uniform) CA graph for which being a non-Helly triangle depends on the representation consider the graph G obtained by taking the complement of the disjoint union of three K_2's (a triforce with a circle around the outer corners). Every edge in G is an ov-entry in λ_G and a CA model for G is given by precisely two non-Helly triangles. The possible assignments of the vertices to the arcs that follow from the automorphisms of G yield the different representations.

▶ **Definition 12.** Let G be a graph and $T = \{u, v, w\}$ is an ov-triangle with $T \notin \Delta_G$. We say v is in-between u, w if at least one of the following holds:
1. $N_T(u), N_T(w) \neq \emptyset$
2. $N_T(u, w) \neq \emptyset$ and there exists $z \in N_T(u, w)$ such that $\{u, w, z\} \in \Delta_G$

▶ **Lemma 13.** *Let G be a uniform CA graph and $T = \{u, v, w\}$ is an ov-triangle with $T \notin \Delta_G$. Then the following statements are equivalent:*
1. *v in-between u, 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. "$2 \Rightarrow 1$": There exists $\rho \in \mathcal{N}(G)$ such that $\rho(v) \subset \rho(u) \cup \rho(w)$. For all $x \neq y \in T$ it must hold that $N[x] \setminus N[y] \neq \emptyset$ due to the fact that x ov y implies $N[x] \between N[y]$. For this to be true $N_T(x)$ or $N_T(x, z)$ must be non-empty for $z \in T \setminus \{x, y\}$. It follows that $N_T(u)$ and $N_T(w)$ or $N_T(u, w)$ must be non-empty since $N_T(v) = \emptyset$ due to $\rho(v) \subset \rho(u) \cup \rho(w)$. If $z \in N_T(u, w)$ then for $\rho(z)$ to intersect with $\rho(u)$ and $\rho(w)$ but not with $\rho(v)$ implies that $\rho[u, w, z]$ forms a non-Helly triangle and therefore $\{u, w, z\} \in \Delta_G$ by Lemma 11.
"$1 \Rightarrow 3$": Assume there exists a $\rho \in \mathcal{N}(G)$ such that $\rho(u) \subset \rho(v) \cup \rho(w)$. This implies that $N_T(u) = \emptyset$ and therefore there exists a $z \in N_T(u, w)$ such that $\{u, w, z\} \in \Delta_G$. By Lemma 11 it follows that $\rho[u, w, z]$ must be a non-Helly triangle and $\rho(z)$ must be disjoint from $\rho(v)$, contradiction.
"$3 \Rightarrow 2$": is clear. ◀

Lemma 11 and 13 state a fact of the form that if a property holds in one representation then it holds in all representations hence the name uniform CA graphs.

The α-neighborhood of a vertex u in a graph G is defined as $N^\alpha(u) = \{v \in N(u) \mid u \, \alpha \, v\}$ for $\alpha \in \{\text{ov}, \text{cd}, \text{cs}, \text{cc}\}$. Now, we can define the aforementioned equivalence relation $\sim_u$ and state the candidate function for uniform CA graphs.

▶ **Definition 14.** Given a graph G and vertex $u \in V(G)$ we define the relation $\sim_u$ on $N^{\text{ov}}(u)$ such that $x \sim_u y$ holds if one of the following applies:
1. $x = y$
2. x contains or is contained in y
3. x and y overlap, $\{x, y, u\} \notin \Delta_G$ and u is not in-between x, y

▶ **Lemma 15.** *Let G be a uniform CA graph and $u \in V(G)$. Then $x \sim_u y$ holds iff x and y overlap from the same side with u in every $\rho \in \mathcal{N}(G)$.*

Proof. "$\Rightarrow$": If x, y are in a contained/contains relation this is clear. For the third condition it holds that for all $\rho \in \mathcal{N}(G)$ $\rho[x, y, u]$ is an interval triangle and x or y is in-between the other two. It follows that x, y overlap from the same side with u.
"$\Leftarrow$": If $\lambda_{x,y} \in \{\text{di}, \text{cc}\}$ this is clear. If x, y overlap then they either form a non-Helly triangle or u is in-between x, y. In both cases x, y overlap from different sides with u. ◀

▶ **Theorem 16.** *The following mapping is a candidate function for uniform CA graphs and can be computed in logspace:*

$$f_U(G) = \bigcup_{\substack{u \in V(G) \\ x \in N^{\mathsf{ov}}(u)}} \left\{ \{u\} \cup N^{\mathsf{cd}}(u) \cup N^{\mathsf{cc}}(u) \cup \{y \in N^{\mathsf{ov}}(u) | x \sim_u y\} \right\}.$$

Proof. To show that $f_{\mathrm{U}}(G)$ is a candidate function we have to prove that for every uniform CA graph G there always exists a flip set $X \in f_{\mathrm{U}}(G)$ and that f_{U} is label-independent. In fact, the even stronger claim holds that for all uniform CA graphs G every set in $f_{\mathrm{U}}(G)$ is a flip set. Let $X \in f_{\mathrm{U}}(G)$ via some $u \in V(G)$ and $x \in N^{\mathsf{ov}}(u)$. Then the set of vertices in X correspond to one of the two flip sets shown in Fig. 2. The correctness for the subset of vertices in X overlapping with u follows from Lemma 15.

To show that $f_{\mathrm{U}}(G)$ can be computed in logspace and is label-independent we apply Lemma 8. We can rewrite f_{U} as

$$f_{\mathrm{U}}(G) = \bigcup_{u,x \in V(G)} \left\{ \{z \in V(G) \mid G \models \varphi(u,x,z)\} \right\}$$

with $G \models \varphi(u,x,z)$ iff $x \in N^{\mathsf{ov}}(u)$ and $z \in \{u\} \cup N^{\mathsf{cd}}(u) \cup N^{\mathsf{cc}}(u) \cup \{y \in N^{\mathsf{ov}}(u) | x \sim_u y\}$. It remains to check that the entries in the neighborhood matrix, α-neighborhoods, exclusive neighborhoods, Δ_G, in-between and $\sim_u$ can be expressed in FO logic. ◀

▶ **Corollary 17.** *A canonical CA representation for uniform CA graphs can be computed in logspace.*

▶ **Theorem 18.** *Helly CA graphs are a strict subset of uniform CA graphs.*

Proof. First, we show "$\subseteq$" by contradiction. Assume there exists a Helly CA graph G which is non-uniform. For G to be non-uniform there must exist an ov-triangle $T \in \Delta_G$ and a representation $\rho \in \mathcal{N}(G)$ such that T is represented as interval triangle in ρ. Let $T = \{u, v, w\}$ and assume w.l.o.g. that $\rho(v) \subset \rho(u) \cup \rho(w)$ (v is in-between u, w). It follows that $N[u] \cup N[w] = V(G)$. Since $\lambda_{u,w} \neq \mathsf{cc}$ there must be a $u' \in N[u] \setminus N[w]$ such that $N[u'] \setminus N[u] \neq \emptyset$. This means there exists a $w' \in N[u', w] \setminus N[u]$. As $N[u'] \between N[w']$ it follows that u' and w' overlap. For u' it must hold that it is either in $N_T(u)$ or $N_T(u,v)$. If it is in $N_T(u)$ then by the second condition of Δ_G it follows that u' must be contained in u, contradiction. For the same reason w' is in $N_T(v,w)$. It follows that u', v, w' form an ov-triangle and must be represented as non-Helly triangle in ρ. This contradicts that G is a Helly CA graph.

To see that this inclusion is strict consider the graph G obtained by taking a triangle T and attaching a new vertex to each vertex in T (also known as net graph). In every representation $\rho \in \mathcal{N}(G)$ it must hold that T is represented as non-Helly triangle since $N_T(v) \neq \emptyset$ for all $v \in T$. For this reason G cannot be a Helly or a non-uniform CA graph. ◀

5 Non-Uniform CA graphs

From the definition of uniform CA graphs it follows that a CA graph G is non-uniform if there exists an ov-triangle T in Δ_G and a $\rho \in \mathcal{N}(G)$ such that T is represented as interval triangle in ρ. We call the pair (T, ρ) a witness for the non-uniformity of G and also say G is non-uniform via (T, ρ). Additionally, for such a witness pair (T, ρ) we call T maximal if there exists no $T' \neq T$ such that G is non-uniform via (T', ρ) and $\rho^+(T) \subset \rho^+(T')$. Such a maximal T for a given ρ must always exist. For the purpose of computing a candidate

function for this class we can assume that for a given non-uniform CA graph G we are supplied with an $\mathtt{ov}$-triangle T such that there exists a $\rho \in \mathcal{N}(G)$ with (T, ρ) being a witness for G and T is maximal. This is justified by the fact that we can iterate over all $\mathtt{ov}$-triangle $T \in \Delta_G$ trying to compute flip sets knowing that for at least one such T these conditions are met. Additionally, we write T as ordered triple (u, v, w) to indicate that v is in-between u and w in ρ.

Let G be non-uniform via (T, ρ). Consider for a vertex $x \in V(G) \setminus T$ what the possible relations between $\rho(x)$ and $\rho^+(T)$ are. For instance, $\rho(x)$ cannot be disjoint from $\rho^+(T)$ because this implies that $x \notin \cup_{t \in T} N[t]$ and therefore $T \notin \Delta_G$. Also, $\rho(x)$ cannot contain $\rho^+(T)$ as this would mean that x is a universal vertex.

▶ **Definition 19.** Let G be a non-uniform CA graph via (T, ρ). The set of normalized representations that agree with ρ on T is:

$$\mathcal{N}_\rho^T(G) = \{\rho' \in \mathcal{N}(G) \mid \rho'[T] = \rho[T]\} \ .$$

▶ **Definition 20.** Let G be a non-uniform CA graph via (T, ρ) and $\alpha \in \{\mathtt{ov}, \mathtt{cc}, \mathtt{cd}\}$. We say $x \in V(G) \setminus T$ is an α-arc in ρ' if

$$\rho'(x) \, \alpha \, \rho^+(T)$$

for some $\rho' \in \mathcal{N}_\rho^T(G)$. We call x an unambiguous α-arc if the above condition holds for all $\rho' \in \mathcal{N}_\rho^T(G)$.

▶ **Definition 21.** Let G be a graph and $T = (u, v, w) \in \Delta_G$. Then we define the following sets w.r.t. T:

$$\Gamma_{\mathtt{ov},u} = \{x \in V(G) \setminus T \mid x \, \mathtt{ov} \, u, v \, , \, x \, \mathtt{di} \, w\}$$
$$\Gamma_{\mathtt{ov},w} = \{x \in V(G) \setminus T \mid x \, \mathtt{ov} \, v, w \, , \, x \, \mathtt{di} \, u\}$$
$$\Gamma_{\mathtt{ov}} \ = \Gamma_{\mathtt{ov},u} \cup \Gamma_{\mathtt{ov},w}$$
$$\Gamma_{\mathtt{cc}} \ = \left\{x \in V(G) \setminus T \ \middle| \ \begin{array}{l} x \, \mathtt{ov} \, u, w \, , \, x \, \mathtt{di} \, v \text{ or} \\ \exists a \in T : x \, \mathtt{cc} \, a \end{array} \right\}$$
$$\Gamma_{\mathtt{cd}} \ = \left\{x \in V(G) \setminus T \ \middle| \ \begin{array}{l} x \, \mathtt{ov} \, u, w \, , \, x \, \mathtt{cs} \, v \text{ or} \\ \exists a \in T : x \, \mathtt{cd} \, a \end{array} \right\}$$
$$\Omega \ = \{x \in V(G) \setminus T \mid x \, \mathtt{ov} \, u, v, w\}$$

▶ **Lemma 22.** *If G is a non-uniform CA graph via (T, ρ) such that T is maximal then $T, \Gamma_{\mathtt{ov}}, \Gamma_{\mathtt{cc}}, \Gamma_{\mathtt{cd}}, \Omega$ partition $V(G)$.*

Proof. It is not hard to see that these sets do not overlap. To show that every vertex in G belongs to one of these sets we need to check all possible positions of the endpoints of a vertex x not in T relative to T, consider Fig. 3. We know every vertex $x \notin T$ must be represented as α-arc for some $\alpha \in \{\mathtt{ov}, \mathtt{cd}, \mathtt{cc}\}$. For $\alpha \in \{\mathtt{cd}, \mathtt{cc}\}$ it must hold that both endpoints of x must be in one of the intervals 1–5. The exemplary x depicted in Fig. 3 is a $\mathtt{cd}$-arc in the given representation and overlaps with u, v, w. If x is a $\mathtt{cd}$-arc then the number of the interval in which its left endpoint is situated must be less than or equal that of its right endpoint (in our example $2 \leq 5$). If it is a $\mathtt{cc}$-arc then the right endpoint must come before the left. The graph on the right encodes all possible placements of the two endpoints and it can be verified case-by-case that an α-arc for $\alpha \in \{\mathtt{cd}, \mathtt{cc}\}$ will occur in either Γ_α or
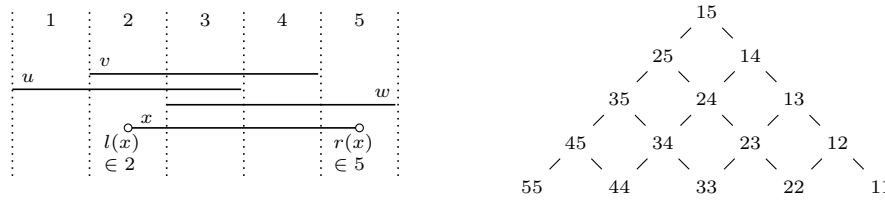
Figure 3 Possible positions of the endpoints of x relative to $T = \{u, v, w\}$

Ω. It remains to argue that an $\mathtt{ov}$-arc can only have the intersection structure denoted by $\Gamma_{\mathtt{ov}}$. W.l.o.g. assume that x is an $\mathtt{ov}$-arc that overlaps from u's side with T. It holds that the right endpoint of x is in one of the five intervals and the left endpoint must be in none of these five intervals. If $r(x) \in 1$ then $x \in N_T(u)$ and by the second condition of Δ_G it must hold that x is contained in u, contradiction. If $r(x) \in 2$ then x overlaps with u, v and is disjoint with w and therefore $x \in \Gamma_{\mathtt{ov}, u}$. If $r(x) \in \{3, 4\}$ then $T = \{u, v, w\}$ is not maximal since $\{x, v, w\} \in \Delta_G$. If $r(x) \in 5$ then it can be shown that there must be a circle cover entry in the neighborhood matrix for x, w which contradicts that they must overlap. ◀

▶ **Lemma 23.** *Let G be a non-uniform CA graph via (T, ρ) and T is maximal. All vertices in Γ_α are unambiguous α-arcs for $\alpha \in \{\mathtt{ov}, \mathtt{cd}, \mathtt{cc}\}$.*

Proof. The intersection structure of a vertex $x \in \Gamma_\alpha$ with T dictates the positioning of the endpoints relative to T in every $\rho' \in \mathcal{N}_\rho^T(G)$. It follows that this placement of the endpoints of x must hold for all $\rho' \in \mathcal{N}_\rho^T(G)$ and therefore x is an unambiguous α-arc. ◀

For a vertex $x \in \Omega$ the possible placements of its endpoints to satisfy the intersection structure with T can be one of the following four types $(\mathtt{cd}, 14), (\mathtt{cd}, 25), (\mathtt{cc}, 14), (\mathtt{cc}, 25)$. For example x in Fig. 3 is type $(\mathtt{cd}, 25)$ and flipping x would lead to type $(\mathtt{cc}, 25)$.

Let G be non-uniform via (T, ρ) and $T = (u, v, w)$ is maximal. Then the two target flip sets X_u, X_w we want to compute in the non-uniform case are immediately before the left endpoint and immediately after the right endpoint of $\rho^+(T)$ and must be of the form

$$X_i = \Gamma_{\mathtt{ov}, i} \cup \Gamma_{\mathtt{cc}} \cup \Omega'$$

for $i \in \{u, w\}$ and some subset $\Omega' \subseteq \Omega$.

▶ **Definition 24.** Let G be a non-uniform CA graph via (T, ρ) and T is maximal. We call a subset $\Omega' \subseteq \Omega$ $\mathtt{cc}$-realizable w.r.t. (T, ρ) if there exists a $\rho' \in \mathcal{N}_\rho^T(G)$ such that for all $x \in \Omega$ it holds that x is a $\mathtt{cc}$-arc in ρ' iff $x \in \Omega'$.

In other words, a $\mathtt{cc}$-realizable set is a subset of Ω such that all of its vertices can be represented as $\mathtt{cc}$-arc in a normalized representation. By finding such a set we can construct two flip sets by adding the $\mathtt{cc}$-vertices in $\Gamma_{\mathtt{cc}}$ and one side of the $\mathtt{ov}$-vertices as described above. Now, the challenge consists in finding such a $\mathtt{cc}$-realizable subset. A way to solve this is to parameterize our input by the cardinality of Ω and try all possibilities. Since Ω and its cardinality depend on the particular $\mathtt{ov}$-triangle T chosen, which we do not know a priori, we can use the following set which is a superset of every possible Ω and thus bounds the cardinality:

$$K_G = \{x \in V(G) \mid \exists \{u, v, w\} \in \Delta_G \text{ s.t. } x \mathtt{ ov } u, v, w\} .$$

▶ **Theorem 25.** *The following mapping is a candidate function for non-uniform CA graphs and can be computed in $\mathcal{O}(k + \log n)$ space for $k = |K_G|$:*

$$f_N(G) = \bigcup_{\substack{(u,v,w) \in \Delta_G \\ \Omega' \subseteq \Omega}} \{\Gamma_{\mathsf{ov},u} \cup \Gamma_{\mathsf{cc}} \cup \Omega'\}$$

where $\Gamma_{\mathsf{ov},u}, \Gamma_{\mathsf{cc}}, \Omega$ are taken w.r.t. $T = (u,v,w)$.

Proof. To show that there always exists a flip set $X \in f_N(G)$ for a non-uniform CA graph G we argue as follows. Let G be non-uniform via (T, ρ) and $T = (u,v,w)$ is maximal. Let Ω' be a cc-realizable set w.r.t. (T, ρ), which must exist since G is non-uniform. Then $X = \Gamma_{\mathsf{ov},u} \cup \Gamma_{\mathsf{cc}} \cup \Omega'$ w.r.t. T is one of the target flip sets described previously. This means there exists a $\rho \in \mathcal{N}(G)$ such that X describes the set of arcs that contain a point right before the left endpoint of $\rho^+(T)$ or a point right after the right endpoint of $\rho^+(T)$.

To see that $f_N(G)$ is label-independent a formula $\varphi(\Omega', u, v, w, x)$ can be constructed that is true iff $x \in \Gamma_{\mathsf{ov},u} \cup \Gamma_{\mathsf{cc}} \cup \Omega'$ where Ω' is a second-order set variable. Note, that $|\Omega| \leq |K_G|$. Therefore this works in $\mathcal{O}(k + \log n)$ space since one can iterate over all 2^k subsets of Ω using k bits and then apply the argument in Lemma 8 via φ which requires additional $\mathcal{O}(\log n)$ space. ◀

6   Conclusion

We showed how to canonically, or in our terms label-independently, compute flip sets for CA graphs to acquire canonical CA representations. The properties of uniform CA graphs enable us to do this easily in logspace. In the case of non-uniform CA graphs, however, it seems that the cc-realizable sets pose a non-trivial obstacle when trying to compute flip sets. The only simple remedy appears to be the proposed parameterization that enables us to use brute force. Changing the target flip sets does not seem to improve upon this situation. As a consequence, we suggest to investigate the space of cc-realizable sets. Given the restricted structure of non-uniform CA graphs this could be a reasonable first step towards deciding isomorphism for CA graphs in polynomial time.

Additionally, in [8] it was shown how to compute flip sets for CA graphs in linear time without the canonicity constraint. Can this be done in logspace as well? This would mean that recognition of CA graphs is logspace-complete.

───── **References** ─────

1 Andrew Curtis, Min Chih 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.

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

3 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. doi:10.1137/S0097539793260726.

4 Johannes Köbler, Sebastian Kuhnert, and Oleg Verbitsky. Solving the Canonical Representation and Star System Problems for Proper Circular-Arc Graphs in Logspace. In *FSTTCS*, volume 18, pages 387–399. Schloss Dagstuhl, 2012.

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

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

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

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

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