On the Isomorphism Problem for Helly Circular-Arc Graphs

Johannes Köbler, Sebastian Kuhnert[1,*], Oleg Verbitsky[2]

Humboldt-Universität zu Berlin, Institut für Informatik, Unter den Linden 6, 10099 Berlin, Germany

Abstract

The isomorphism problem is known to be efficiently solvable for interval graphs, while for the larger class of circular-arc graphs its complexity status stays open. We consider the intermediate class of intersection graphs for families of circular arcs that satisfy the Helly property. We solve the isomorphism problem for this class in logarithmic space. If an input graph has a Helly circular-arc model, then our algorithm constructs it canonically, which means that the models constructed for isomorphic graphs are equal.

Keywords: Graph isomorphism, circular-arc graphs, Helly property, graph, logarithmic space

1. Introduction

An *intersection representation* of a graph G is a mapping α from the vertex set $V(G)$ onto a family $\mathcal{A}$ of sets such that vertices u and v of G are adjacent if and only if the sets $\alpha(u)$ and $\alpha(v)$ have a nonempty intersection. The family $\mathcal{A}$ is called an *intersection model* of G. G is an *interval graph* if it admits an intersection model consisting of intervals of reals (or, equivalently, intervals of consecutive integers). The larger class of *circular-arc (CA) graphs* arises if we consider intersection models consisting of arcs on a circle. These two archetypal classes of intersection graphs have important applications, most noticeably in computational genomics, and have been intensively studied for decades in graph theory and algorithmics; for an overview see e.g. [Spi03]. In general, fixing a class of admissible intersection models, we obtain the corresponding class of intersection graphs.

In the *canonical representation problem* for a class $\mathcal{C}$ of intersection graphs, we are given a graph $G \in \mathcal{C}$ and have to compute its intersection representation α so that isomorphic graphs receive equal intersection models. This subsumes both the recognition of $\mathcal{C}$ and the isomorphism testing for graphs in $\mathcal{C}$. In their seminal work [BL76, LB79], Booth and Lueker solve both the representation and the isomorphism problems for interval graphs in linear time. Together with Laubner, we designed a canonical representation algorithm for interval graphs that takes logarithmic space [KKL$^+$11].

The case of CA graphs remains a challenge up to now. While a circular-arc intersection model can be constructed in linear time (McConnell [McC03]), no polynomial-time isomorphism test for CA graphs is currently known (though some approaches [Hsu95] have appeared in the literature; see the discussion in [CLM$^+$13]). A few natural subclasses of CA graphs have received special attention among researchers. In particular, for proper CA graphs both the recognition and the isomorphism problems are solved in linear time, respectively, in [DHH96, KN09] and [LSS08, CLM$^+$13], and in logarithmic space in [KKV12]. The latter

*Corresponding author.

Email addresses: koebler@informatik.hu-berlin.de (Johannes Köbler), kuhnert@informatik.hu-berlin.de (Sebastian Kuhnert), verbitsk@informatik.hu-berlin.de (Oleg Verbitsky)

[1]Supported by DFG grant KO 1053/7–2.

[2]Supported by DFG grant VE 652/1–1. On leave from the Institute for Applied Problems of Mechanics and Mathematics, Lviv, Ukraine.

result actually gives a logspace algorithm for computing a canonical representation of proper CA graphs, and such an algorithm is also known for unit CA graphs [Sou14]. The history of the isomorphism problem for circular-arc graphs is surveyed in more detail by Uehara [Ueh13].

Here we are interested in the class of *Helly circular-arc (HCA) graphs.* Those are graphs that admit circular-arc models having the *Helly property*, which requires that every subfamily of arcs with nonempty pairwise intersections has a nonempty overall intersection. Obeying this property is assumed in the representation problem for HCA graphs. Since any family of intervals has the Helly property and the cycles of length at least 4 are HCA graphs but not interval, the canonical representation problem for HCA graphs generalizes the canonical representation problem for interval graphs. On the other hand, not every CA model is Helly; see Fig. 1 for examples. Joeris et al. characterize HCA graphs among CA graphs by a family of forbidden induced subgraphs [JLM$^+$11].

HCA graphs were introduced by Gavril under the name of Θ circular-arc graphs [Gav74]. Gavril gave an $O(n^3)$ time representation algorithm for HCA graphs. Hsu improved this to $O(nm)$ [Hsu95]. Recently, Joeris et al. gave a linear time representation algorithm [JLM$^+$11]. The fastest known isomorphism algorithm for HCA graphs is due to Curtis et al. and works in linear time [CLM$^+$13]. Chen gave a parallel AC^2 algorithm [Che96].

We aim at designing space efficient algorithms. In [KKV13] we already presented a logspace canonical representation algorithm for HCA graphs. Our approach in [KKV13] uses techniques developed by McConnell in [McC03], and the algorithm is rather intricate. Now we suggest an alternative approach that is independent of [McC03]. The new algorithm admits a much simpler analysis and exploits some new ideas that may be of independent interest.

Theorem 1.1. *The canonical representation problem for the class of Helly circular-arc graphs is solvable in logspace.*

Note that solvability in logspace implies solvability in logarithmic time by a CRCW PRAM with polynomially many parallel processors, i.e., in AC^1. Prior to our work, no AC^1 algorithm was known for recognition and isomorphism testing of HCA graphs.

In general, solvability of the isomorphism problem for a non-trivial class of graphs in logarithmic space is an interesting result because the general graph isomorphism problem is known to be DET-hard [Tor04] and therefore also NL-hard. It is also interesting that for some classes of intersection graphs, the isomorphism problem is as hard as in general. For example, Uehara [Ueh08] shows this for intersection graphs of axis-parallel rectangles in the plane. Note that any family of such rectangles has the Helly property.

Our strategy. Recall that a hypergraph $\mathcal{H}$ is *interval* (resp. *circular-arc*) if it is isomorphic to a hypergraph whose hyperedges are intervals of integers (resp. arcs of a discrete circle). Such an isomorphism is called an interval (resp. arc) representation of $\mathcal{H}$. Like in our approach to interval graphs in [KKL$^+$11], the overall idea of our algorithm is to exploit the relationship between an input graph G and the dual of its clique hypergraph, which will be denoted by $\mathcal{B}(G)$. Fulkerson and Gross [FG65] established that G is an interval graph if and only if $\mathcal{B}(G)$ is an interval hypergraph. Moreover, represented as an interval system, $\mathcal{B}(G)$ can serve as an interval model of G. More specifically, our approach in [KKL$^+$11] consists of two steps: First,

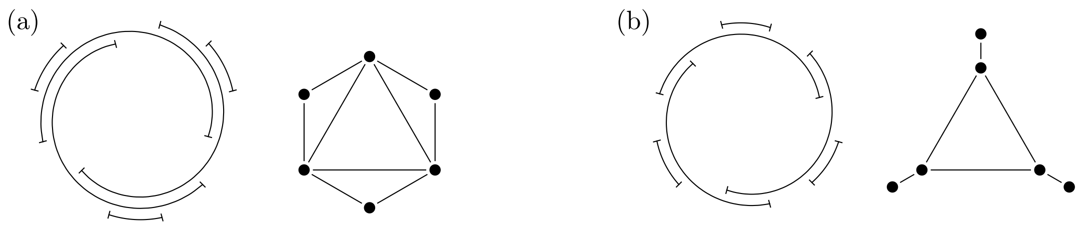

Figure 1: Two non-Helly CA models and their intersection graphs. The graph in (a) admits an HCA model, while the graph in (b) does not.

construct $\mathcal{B}(G)$ (or, equivalently, find all maxcliques in G) and, second, design a canonical representation algorithm for interval *hypergraphs* and apply it to $\mathcal{B}(G)$. The first step is implementable in logspace because all interval graphs without isolated vertices are *maximal clique irreducible* [OR81]. This means that every maxclique C contains an edge uv that is contained in no other maxclique, and implies that C is equal to the common neighborhood of u and v (cf. Lemma 7.2).

The Fulkerson-Gross theorem is extended to the class of HCA graphs by Gavril [Gav74]: G is an HCA graph if and only if $\mathcal{B}(G)$ is a CA hypergraph. Also in this case, any arc representation of $\mathcal{B}(G)$ yields a Helly arc representation for G. The canonical representation problem for CA hypergraphs is solved in logspace in [KKV12]. However, the similarity between interval and HCA graphs ends there because HCA graphs are in general not maximal clique irreducible and hence we have to use a different approach to compute the hypergraph $\mathcal{B}(G)$.

Though we are not able to find all maxcliques of an HCA graph G directly, the discussion above shows that the *canonical representation* problem for HCA graphs is logspace reducible to the *representation* problem, where we need just to construct a Helly arc representation and do not need to take care of canonicity. Indeed, once we have an arbitrary HCA model of an input graph G, we get all maxcliques of G by inspection of the sets of arcs sharing a common point. As soon as all the maxcliques are found, we form the hypergraph $\mathcal{B}(G)$ and compute its canonical representation according to [KKV12] (the details are given in Section 4).

It remains to explain how we compute a Helly arc representation α for a given HCA graph G with n vertices. It is handy to assume that the corresponding HCA model $\alpha(G)$ has $2n$ points and that no arc in $\alpha(G)$ shares extreme points with others. For a given subset $C \subset V(G)$, let $\alpha^C(G)$ denote the arc system obtained from $\alpha(G)$ by flipping the arc $\alpha(v)$ for each $v \in C$, that is, by replacing $\alpha(v)$ with the other arc on the same circle that has the same extreme points. We make use of a simple consequence of the Helly property: If C is a maxclique, then $\alpha^C(G)$ becomes an *interval* system. Note that it suffices to find only *one* maxclique C of G in order to perform the flipping operation. Such a maxclique C can be found in logspace; see Lemma 7.3. Moreover, if we can also compute the corresponding mapping α^C, we obtain the desired Helly arc representation α by performing the C-flipping for $\alpha^C(G)$ (note that $(\alpha^C)^C = \alpha$). The flipping operation is considered in detail in Section 6.

The interval system $\alpha^C(G)$ and the corresponding mapping α^C are constructed as follows. In Section 5 we combine results of Hsu [Hsu95] and Joeris et al. [JLM$^+$11] to argue that α can be supposed to be *normalized*, which means that the geometric configurations of arc pairs in the arc model $\alpha(G)$ are predetermined by natural conditions expressible in terms of the adjacency relation of G. Using these conditions, we are able to compute the *pairwise-intersection matrix* $M_\alpha = (m_{uv})$, defined by $m_{uv} = |\alpha(u) \cap \alpha(v)|$, and then also the pairwise-intersection matrix M_{α^C} for the interval system $\alpha^C(G)$. Afterwards we use another result of Fulkerson and Gross saying that an interval system is determined by its pairwise-intersection matrix up to isomorphism [FG65]. Moreover, it can be reconstructed from the pairwise-intersection matrix in logspace by an algorithm worked out in [KKW15]; see Section 3.

2. Formal definitions

Hypergraphs. Recall that a *hypergraph* is a pair $(X, \mathcal{H})$, where $X = V(\mathcal{H})$ is a set of vertices and $\mathcal{H}$ is a family of subsets of X, called *hyperedges*. We will use the same notation $\mathcal{H}$ to denote a hypergraph and its hyperedge set. A hypergraph has the *Helly property* if every set of pairwise intersecting hyperedges has a common vertex. An isomorphism from a hypergraph $\mathcal{H}$ to a hypergraph $\mathcal{K}$ is a bijection $\phi\colon V(\mathcal{H}) \to V(\mathcal{K})$ such that $H \in \mathcal{H}$ if and only if $\phi(H) \in \mathcal{K}$ for every $H \subseteq V(\mathcal{H})$. Here and below, if ϕ is a function and H is a subset of its domain, then $\phi(H) = \{\phi(x) : x \in H\}$ denotes the image of H under ϕ. In Section 4, we consider hypergraphs $\mathcal{H}$ with *multi-hyperedges*, that is, each hyperedge H is assigned a positive integer $c_{\mathcal{H}}(H)$, called the *multiplicity* of H. An isomorphism ϕ from $\mathcal{H}$ to $\mathcal{K}$ has to preserve the multiplicities, that is, it is required that $c_{\mathcal{K}}(\phi(H)) = c_{\mathcal{H}}(H)$ for every hyperedge H of $\mathcal{H}$.

Arc and interval systems. For $n \geq 3$, consider the directed cycle $\mathbb{C}$ on the vertex set $\{1, \ldots, n\}$ with arrows from i to $i + 1$ and from n to 1. The vertices of $\mathbb{C}$ will be called *points*. An *arc* $[a, b]$ consists of the points appearing in the directed path from a to b. An *arc system* $\mathcal{A}$ is a hypergraph on the vertex set $\{1,\ldots,n\}$ whose hyperedges are arcs.

Given an arc $A = [a, b]$, we refer to a and b as *extreme points* of A; in particular, a is the *start point* and b is the *end point* of A. Note that $[1, n]$ and $[i, i-1]$ for $i = 2,\ldots,n$ denote the same *complete arc*. The notion of extreme points is ambiguous in this case. Nevertheless, we will consider complete arcs with *designated* extreme points in Sections 6 and 7. An arc $[a, b]$, a complete one in particular, can alternatively be viewed as a path in $\mathbb{C}$ starting in a and ending in b.

The notions of an *interval* and an *interval system* are defined similarly with the only difference that in place of $\mathbb{C}$, the underlying structure is now the directed *path* $\mathbb{I}$ on the points $1,\ldots,n$ with arrows from i to $i+1$. Note that the extreme points of the complete interval $[1, n]$ are unambiguous.

In fact, we can consider an interval system on an arbitrary linearly ordered set. In particular, in Section 7 we consider interval systems where the underlying path is a complete arc with designated extreme points. In this way, if an arc $[a, b]$ is understood as a path from a to b, then any arc system in $\mathbb{C}$ that does not traverse some specified edge of $\mathbb{C}$ is naturally regarded as an interval system.

Arc representations of hypergraphs. An *arc representation* of a hypergraph $\mathcal{H}$ is an isomorphism ρ from $\mathcal{H}$ to an arc system $\mathcal{A}$. It can be thought of as a circular ordering of $V(\mathcal{H})$ where every hyperedge is a segment of consecutive vertices. The arc system $\mathcal{A}$ is referred to as an *arc model* of $\mathcal{H}$. The notions of an *interval representation* and an *interval model* of a hypergraph are introduced similarly. Hypergraphs having arc representations are called *circular-arc (CA) hypergraphs*, and those having interval representations are called *interval hypergraphs.*

A *representation scheme for CA hypergraphs* is a function defined on CA hypergraphs that on input $\mathcal{H}$ outputs an arc representation $\rho_{\mathcal{H}}$ of $\mathcal{H}$. Such a representation scheme is called *canonical* if isomorphic CA hypergraphs $\mathcal{H} \cong \mathcal{K}$ always receive *equal* arc models $\rho_{\mathcal{H}}(\mathcal{H}) = \rho_{\mathcal{K}}(\mathcal{K})$. In [KKV12] we designed a canonical representation scheme for CA hypergraphs computable in logarithmic space. Note that our algorithm works for hypergraphs with multi-hyperedges.

Graphs. The vertex set of a graph G is denoted by $V(G)$. The *closed neighborhood* $N[v]$ of a vertex v consists of v itself and all the vertices adjacent to it. A vertex u is *universal* if $N[u] = V(G)$. Two vertices u and v are *twins* if $N[u] = N[v]$. Note that twins are always adjacent. The *twin class* $[v]$ of a vertex v consists of v itself along with all its twins. Between two different twin classes there are either none or all of the possible edges. This allows us to consider the *quotient graph* G' on the vertex set $V(G') = \big\{[v]\big\}_{v\in V(G)}$ where two distinct twin classes $[v]$ and $[u]$ are adjacent if v and u are adjacent in G. The map $v \mapsto [v]$ from G to G' will be referred to as the *quotient map.*

The *intersection graph* $\mathbb{I}(\mathcal{H})$ of a hypergraph $\mathcal{H}$ has the hyperedges of $\mathcal{H}$ as vertices, and two such vertices $A, B \in \mathcal{H}$ are adjacent if $A \cap B \neq \emptyset$. If $\mathcal{H}$ has hyperedges of multiplicity greater than 1, they become twins in $\mathbb{I}(\mathcal{H})$.

Arc representations and arc models of graphs. An *intersection representation of a graph* G is an isomorphism $\alpha\colon V(G) \to \mathcal{A}$ from G to the intersection graph $\mathbb{I}(\mathcal{A})$ of a hypergraph $\mathcal{A}$. The hypergraph $\mathcal{A}$ is then called an *intersection model* of G. If $\mathcal{A}$ is an arc system, we call α an *arc representation* and $\mathcal{A}$ an *arc model* of G. Graphs having arc representations are called *circular-arc (CA) graphs.* In other words, those are graphs isomorphic to the intersection graph of some CA hypergraph. *Helly circular-arc (HCA) graphs* are graphs having *Helly arc representations*, i.e., representations providing arc models that obey the Helly property. The notions of *interval representations* and *interval models* of graphs as well as *interval graphs* are defined accordingly.

A *representation scheme for a class* $\mathcal{C}$ *of CA graphs* is a function that on input $G \in \mathcal{C}$ outputs an arc representation α_G of G. A representation scheme for HCA graphs must produce Helly arc representations. If a representation scheme produces equal models for isomorphic input graphs, it is called *canonical.*

3. Pairwise intersections as a complete isomorphism invariant for interval hypergraphs

In this short section, we state a few useful facts about interval systems. Let V be a set and let $\mathcal{H}$ be a hypergraph. Given a bijection $\lambda\colon V \to \mathcal{H}$, we define the *pairwise-intersection matrix* $M_\lambda = (m_{uv})_{u,v\in V}$ of λ by $m_{uv} = |\lambda(u) \cap \lambda(v)|$. Consider two bijections $\lambda\colon V \to \mathcal{H}$ and $\mu\colon V \to \mathcal{K}$ from V to hypergraphs $\mathcal{H}$ and $\mathcal{K}$, respectively. If there is an isomorphism ψ from $\mathcal{H}$ to $\mathcal{K}$ such that $\mu = \psi \circ \lambda$, then obviously $M_\lambda = M_\mu$. It turns out that the converse is also true if $\mathcal{H}$ is an interval hypergraph.

Lemma 3.1 (Fulkerson and Gross [FG65]). *Let $\mathcal{I}$ be an interval system and $\mathcal{J}$ be an arbitrary hypergraph. Suppose that $M_\lambda = M_\mu$ for bijections $\lambda\colon V \to \mathcal{I}$ and $\mu\colon V \to \mathcal{J}$. Then there is a hypergraph isomorphism ψ such that $\mu = \psi \circ \lambda$; see Fig. 2.*

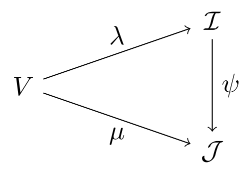

Figure 2: Lemma 3.1: If $M_\lambda = M_\mu$ and $\mathcal{I}$ is an interval hypergraph, then $\mathcal{I} \cong \mathcal{J}$.

We will use the fact that $\mathcal{I}$ and λ are efficiently reconstructible from a given matrix $M = M_\lambda$.

Lemma 3.2 (Köbler, Kuhnert, and Watanabe [KKW15]). *There is a logspace algorithm that computes for a given integer matrix $M = (m_{uv})_{u,v\in V}$ an interval system $\mathcal{I}$ and a bijection $\lambda\colon V \to \mathcal{I}$ such that $M = M_\lambda$ (if they exist).*

4. Getting canonicity for free

A *clique* in a graph G is a set of pairwise adjacent vertices. An inclusion-maximal clique will be called a *maxclique*. The *clique hypergraph* $\mathcal{C}(G)$ of a graph G has the same vertex set as G (i.e., $V(\mathcal{C}(G)) = V(G)$) and the maxcliques of G as its hyperedges. We now define the *bundle hypergraph* $\mathcal{B}(G)$, which is the dual of $\mathcal{C}(G)$. The hypergraph $\mathcal{B}(G)$ has the maxcliques of G as vertices (i.e., $V(\mathcal{B}(G)) = \mathcal{C}(G)$) and a hyperedge B_v for each vertex v of G, where B_v consists of all the maxcliques that contain v. We call B_v the *(maxclique) bundle* of v and denote the corresponding map $v \mapsto B_v$ from $V(G)$ to $\mathcal{B}(G)$ by β_G.

We begin with a well-known general fact (see, e.g., [MM99, Theorem 1.14]).

Lemma 4.1. *Let G be a graph. Then the map β_G is an intersection representation for G, that is, two vertices u and v of G are adjacent if and only if $B_u \cap B_v \neq \emptyset$. Moreover, the corresponding intersection model $\mathcal{B}(G)$ has the Helly property.*

The following classical result provides a link between HCA graphs and CA hypergraphs; it is exemplified in Fig. 3.

Lemma 4.2 (Gavril [Gav74]). *G is an HCA graph if and only if $\mathcal{B}(G)$ is a CA hypergraph. Moreover, if ρ is a CA representation of $\mathcal{B}(G)$, then $\alpha_G = \rho \circ \beta_G$ (i.e., $\alpha_G(v) = \rho(B_v)$ for all $v \in V(G)$) is a Helly arc representation of G.*

Note that if u and v are twins of G, then $B_u = B_v$ and consequently $\alpha_G(u) = \alpha_G(v)$, i.e., twins are mapped to arcs of multiplicity greater than one. Although we do not require this, we remark that α_G can be transformed in logspace into a Helly arc representation of G that maps different vertices to different arcs.

Lemma 4.3. *The canonical representation problem for HCA graphs is logspace reducible to the (not necessarily canonical) representation problem for HCA graphs that have neither twins nor universal vertices.*

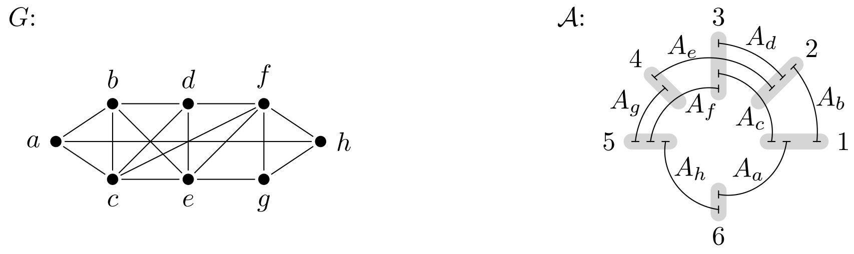

Figure 3: The graph G contains the maxcliques $C_1 = \{a, b, c\}$, $C_2 = \{b, c, d, e\}$, $C_3 = \{c, d, e, f\}$, $C_4 = \{e, f, g\}$, $C_5 = \{f, g, h\}$, and $C_6 = \{a, h\}$. Its bundle hypergraph $\mathcal{B}_G$ admits the HCA model $\mathcal{A}$ via the representation $\rho\colon \mathcal{C}(G) \to \{1, 2, 3, 4, 5, 6\}$ that maps each maxclique C_i to the point i, and thus $\rho(B_v) = A_v$ for each $v \in V(G)$. The function $\alpha\colon V(G) \to \mathcal{A}$ that maps each vertex v to the arc A_v is a Helly arc representation of G.

Proof. We first show that the canonical representation problem for HCA graphs reduces in logspace to the problem of computing $\mathcal{C}(G)$, that is, to finding all maxcliques in a given HCA graph G. Indeed, given $\mathcal{C}(G)$, we can easily construct the bundle hypergraph $\mathcal{B}(G)$ and the mapping β_G. By Lemma 4.2, we can combine β_G with an arc representation $\rho_{\mathcal{B}(G)}$ of the CA hypergraph $\mathcal{B}(G)$ and obtain an arc representation $\alpha_G = \rho_{\mathcal{B}(G)} \circ \beta_G$. If $\rho_{\mathcal{B}(G)}$ is chosen according to the logspace-computable canonical representation scheme for CA hypergraphs designed in [KKV12], then $G \mapsto \alpha_G$ will be a canonical representation scheme for HCA graphs. Indeed, if $G \cong H$, then $\mathcal{B}(G) \cong \mathcal{B}(H)$, which implies that $\alpha_G(G) = \rho_{\mathcal{B}(G)}(\mathcal{B}(G))$ is equal to $\alpha_H(H) = \rho_{\mathcal{B}(H)}(\mathcal{B}(H))$.

Note now that the problem of computing $\mathcal{C}(G)$ is equivalent to its restriction to graphs having neither twins nor universal vertices. Indeed, let G' be obtained from G by computing its quotient-graph Q with respect to the twin-relation and removing the universal vertex $[u]$ from Q (if it exists). Given $\mathcal{C}(G')$, we easily obtain $\mathcal{C}(G)$ by possibly re-inserting $[u]$ in each maxclique of G' and by converting each maxclique $\{[v_1], \ldots, [v_k]\}$ of Q to the corresponding maxclique $[v_1] \cup \ldots \cup [v_k]$ of the original graph G.

It remains to show that finding $\mathcal{C}(G')$ reduces to computing an arbitrary Helly arc representation α of G'. Given an HCA model $\alpha(G')$, for each point x of the underlying cycle we can compute the set $C_x = \big\{v \in V(G') : x \in \alpha(v)\big\}$. Obviously, C_x is a clique in G'. By the Helly property, among these cliques there are all the maxcliques of G'. Since maximality of a given clique is easy to detect in logspace, this allows us to compute $\mathcal{C}(G')$. ∎

5. Normalized arc representations of HCA graphs

A system $\mathcal{A}$ of m arcs on the $2m$-point cycle will be called *sharp* if all extreme points of the arcs in $\mathcal{A}$ are pairwise distinct; in other words, every point of the cycle is either the start or the end point of exactly one arc. An arc representation of a graph G is *sharp* if the corresponding arc model of G is sharp. It suffices to deal with arc representations of this kind because any CA graph has a sharp arc model.

Moreover, we are particularly interested in those arc representations where the mutual placement of any two arcs $\alpha(u)$ and $\alpha(v)$ is determined by conditions on u and v expressible in terms of the adjacency relation of G. Note that in any arc representation α, arcs $\alpha(u)$ and $\alpha(v)$ are disjoint exactly when the vertices u and v are non-adjacent. Furthermore, $\alpha(u) \subseteq \alpha(v)$ only if $N[u] \subseteq N[v]$ (but in general, the converse may not be true).

If two arcs A and B intersect but neither of them includes the other, then there are two possibilities: Either each arc contains exactly one extreme point of the other arc or each arc contains both extreme points of the other arc. In the former case we say that A and B *strictly overlap* and write $A \frown B$. In the latter case we say that A and B form a *circle cover* and write $A \circledcirc B$. Note that the relation $\alpha(u) \circledcirc \alpha(v)$ is only possible if the vertices u and v satisfy the following conditions:

$$N[u] \cup N[v] = V(G); \tag{1}$$
$$w \in N[u] \setminus N[v] \text{ implies } N[w] \subset N[u]; \tag{2}$$
$$w \in N[v] \setminus N[u] \text{ implies } N[w] \subset N[v]. \tag{3}$$

Definition 5.1 (cf. Hsu [Hsu95]). A sharp arc representation α of a graph G is called *normalized* if the following two conditions are met for every two vertices u and v of G:

1. $\alpha(u) \subseteq \alpha(v)$ exactly when $N[u] \subseteq N[v]$.
2. $\alpha(u) \circledcirc \alpha(v)$ exactly when all the three conditions (1)–(3) are true.

We will benefit from the following fact.

Lemma 5.2 (Hsu [Hsu95]). *Every CA graph G without twins and universal vertices has a normalized arc representation.*

When applied to an HCA graph G, Lemma 5.2 itself does not guarantee that G has a normalized *Helly* arc representation. Nevertheless, this is true and follows from a result of Joeris et al. in [JLM$^+$11].

Lemma 5.3. *Let G be an HCA graph without twins and universal vertices. Then every normalized arc representation of G is a Helly arc representation.*

Proof. Joeris et al. in [JLM$^+$11] introduce the concept of a *stable* arc model and prove that every stable arc model of an HCA graph has the Helly property [JLM$^+$11, Theorem 4.1]. Therefore, it suffices to show that any arc model $\mathcal{A}$ of G produced by a normalized arc representation $\alpha\colon V(G) \to \mathcal{A}$ is stable.

In the case that G has no universal vertex, the stability property can be defined as follows. Given an arc system in a cycle $\mathbb{C}$, by an *extreme sequence* we mean a maximal path in $\mathbb{C}$ containing extreme points of the same type, i.e., either start or end points. Let S be an extreme sequence of start points succeeding an extreme sequence E of end points in the arc model $\mathcal{A}$ of G. Suppose that an arc A has start point in S and an arc B has end point in E. Then the stability of $\mathcal{A}$ means that all such A and B are disjoint.

Assume now that there are such $A = \alpha(u)$ and $B = \alpha(v)$ with nonempty intersection. It is easy to observe that the vertices u and v satisfy the conditions (1)–(3), while A and B strictly overlap. This contradicts the assumption that α is normalized. ∎

We will use yet another property of normalized arc representations of HCA graphs. If $A \frown B$, then we will write $A \prec B$ in the case that B contains the end point of A (equivalently, A contains the start point of B).

Lemma 5.4. *Let α be a normalized Helly arc representation of a graph G and let u, v, w be vertices of G. If the arcs $\alpha(u)$, $\alpha(v)$ and $\alpha(w)$ strictly overlap each other, then*

$$\alpha(u) \cap \alpha(v) \subseteq \alpha(w) \text{ if and only if } N[u] \cap N[v] \subseteq N[w]. \tag{4}$$

Proof. To prove (4) in the forward direction, we need only the Helly property and the fact that $\alpha(u) \cap \alpha(v) \neq \emptyset$. Indeed, let $x \in N[u] \cap N[v]$. It follows that $\alpha(x)$ intersects both $\alpha(u)$ and $\alpha(v)$. By the Helly property, $\alpha(x)$ intersects even the intersection $\alpha(u) \cap \alpha(v)$. Hence, $x \in N[w]$ follows from the assumption that $\alpha(u) \cap \alpha(v) \subseteq \alpha(w)$.

Now suppose that $\alpha(u) \cap \alpha(v) \nsubseteq \alpha(w)$. We assume without loss of generality that $\alpha(u) \prec \alpha(v)$. As the Helly property precludes the case $\alpha(v) \prec \alpha(w) \prec \alpha(u)$ (see Fig. 4.a), we either have $\alpha(u) \prec \alpha(w)$ and $\alpha(v) \prec \alpha(w)$ or $\alpha(w) \prec \alpha(u)$ and $\alpha(w) \prec \alpha(v)$. We assume the former relation (see Fig. 4.b); the latter case is symmetric. Because α is normalized and $\alpha(v) \nsubseteq \alpha(w)$, there exists a vertex $x \in N[v] \setminus N[w]$. Its arc $\alpha(x)$ must contain a point in $\alpha(v) \setminus \alpha(w)$, which is a subset of $\alpha(u)$. Thus, x is also a neighbor of u and witnesses that $N[u] \cap N[v] \nsubseteq N[w]$. ∎

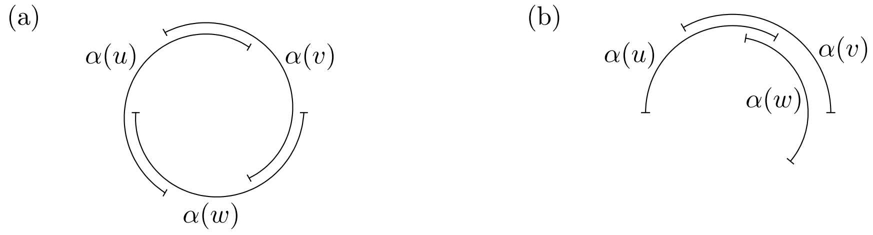

Figure 4: Proof of Lemma 5.4

6. Flipping in a sharp arc system

As we already mentioned in Section 2, the arcs $[1, n]$ and $[a, a-1]$ for $a = 2, \ldots, n$ all coincide with the complete arc $\{1, \ldots, n\}$ on the n-point cycle $\mathbb{C}$. However, in this section we differentiate between these n arcs and call them *complete arcs with designated extreme points*. Suppose that an arc $A = [a, b]$ contains more than one point, that is, $a \neq b$. In this case, we will say that the arc $\tilde{A} = [b, a]$ is obtained from A by *flipping*. This operation applies, in particular, to the n complete arcs with designated extreme points, producing the n two-point arcs $[n, 1]$ and $[a-1, a]$ for $a = 2, \ldots, n$. If applied to two-point arcs, the flipping operation produces complete arcs with designated extreme points. Note also that sharpness is preserved by flipping.

Suppose that an arc system $\mathcal{A}$ contains no one-point arc but possibly contains complete arcs with designated extreme points. Given a mapping $\nu\colon V \to \mathcal{A}$ and a set $C \subseteq V$, we define the *C-flipped mapping* $\nu^C : V \to \mathcal{A}^{\nu(C)}$ by $\nu^C(v) = \widetilde{\nu(v)}$ for $v \in C$ and $\nu^C(v) = \nu(v)$ for $v \notin C$. Here $\mathcal{A}^{\mathcal{X}} = \{\tilde{A} : A \in \mathcal{X}\} \cup \{A : A \in \mathcal{A} \setminus \mathcal{X}\}$ for a subset $\mathcal{X} \subseteq \mathcal{A}$. We will use this notation in the case when $V = V(G)$ is the vertex set of a graph G, though ν is not necessarily an arc representation. In the following lemma we denote arc systems by $\mathcal{I}$ and $\mathcal{J}$ (which we typically do for interval systems) in order to make the notation consistent with Section 7, where this lemma is used.

Lemma 6.1. *Let $\mathcal{I}$ be an arc system containing no one-point arc but possibly complete arcs with designated extreme points. Let ψ be a hypergraph isomorphism from $\mathcal{I}$ to another arc system $\mathcal{J}$ that takes the extreme points of each arc $A \in \mathcal{I}$ to the extreme points of the arc $\psi(A) \in \mathcal{J}$. Given a mapping $\lambda\colon V \to \mathcal{I}$, define the mapping $\mu\colon V \to \mathcal{J}$ by $\mu = \psi \circ \lambda$. Then, for any subset $C \subseteq V$, ψ is an isomorphism from $\mathcal{I}^{\lambda(C)}$ to $\mathcal{J}^{\mu(C)}$ and, moreover, $\mu^C = \psi \circ \lambda^C$; see Fig. 5.*

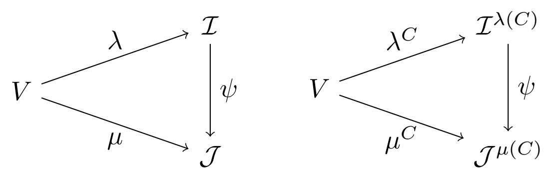

Figure 5: Lemma 6.1: Flipping preserves isomorphisms that respect extreme points.

Proof. For every $v \in V$, the isomorphism ψ maps the arc $\lambda(v) = [a^-, a^+]$ in $\mathcal{I}$ onto the arc $\mu(v) = [b^-, b^+]$ in $\mathcal{J}$. Since ψ preserves the extreme points, $\psi(\{a^-, a^+\}) = \{b^-, b^+\}$. This implies that for every $v \in V$, ψ maps the flipped arc $\widetilde{\lambda(v)}$ onto the flipped arc $\widetilde{\mu(v)}$, irrespective of whether or not $v \in C$. Therefore, ψ always maps $\lambda^C(v)$ onto $\mu^C(v)$, which exactly means that ψ is an isomorphism from $\mathcal{I}^{\lambda(C)}$ to $\mathcal{J}^{\mu(C)}$ and $\mu^C = \psi \circ \lambda^C$. ∎

Lemma 6.1 holds under the assumption that the given isomorphism ψ respects the extreme points of all arcs. It is easy to see that there are isomorphisms that violate this assumption. For example, the transposition (23), while being an automorphism of the interval system $\{[1, 3], [2, 4]\}$, exchanges the extreme points of two different intervals. However, two isomorphic sharp interval systems always admit an isomorphism that does respect extreme points. Before we prove this below in Lemma 6.3, we need to recall some general notions and facts about interval systems.

A *slot* of a hypergraph $\mathcal{H}$ is an inclusion-maximal subset S of $V(\mathcal{H})$ such that each hyperedge of $\mathcal{H}$ contains either all of S or none of it. We say that hyperedges A and B *overlap*, and write $A \between B$, if they intersect but neither of them includes the other. Similarly to the intersection graph $\mathbb{I}(\mathcal{H})$, for any hypergraph $\mathcal{H}$ we consider its *overlap graph* $\mathbb{O}(\mathcal{H})$ with the vertex set $\mathcal{H}$, where $A \in \mathcal{H}$ and $B \in \mathcal{H}$ are adjacent if $A \between B$. By *overlap-connected components* of $\mathcal{H}$ we mean the subsets of $\mathcal{H}$ spanning the connected components of the overlap graph $\mathbb{O}(\mathcal{H})$. If $\mathcal{O}$ and $\mathcal{O}'$ are different overlap-connected components, then either they are vertex-disjoint or all hyperedges of one of the two components are contained in a single slot of the other component.[3] If $\mathcal{H}$ is connected, then this containment relation determines a tree-like decomposition of $\mathcal{H}$ into its overlap-connected components.[4] The root in this tree will be referred to as the *top component*; the other components will be called *inner*.

Note that in the terminology of Section 5, two overlapping arcs either strictly overlap or form a circle cover. The latter relation is excluded in interval systems. The following fact about interval systems is due to [CY91, Theorem 2]; see also [KKL+11, Section 2.2].

Lemma 6.2 (Chen and Yesha [CY91]). *Suppose that $\mathcal{I}$ and $\mathcal{J}$ are isomorphic overlap-connected interval systems. Let $I_1, \ldots, I_k$ be all slots of $\mathcal{I}$ listed in the order as they appear in the integer line. Similarly, let $J_1, \ldots, J_k$ be the sequence of slots of $\mathcal{J}$ as they appear in the integer line. Then any isomorphism from $\mathcal{I}$ to $\mathcal{J}$ maps either each I_s onto J_s or each I_s onto J_{k+1-s}.*

Lemma 6.3. *Let $\mathcal{I}$ and $\mathcal{J}$ be isomorphic sharp interval systems. Then there is a hypergraph isomorphism ψ from $\mathcal{I}$ to $\mathcal{J}$ that respects extreme points, that is, it takes the extreme points of each interval $A \in \mathcal{I}$ to the extreme points of the interval $\psi(A) \in \mathcal{J}$. Moreover, for every isomorphism ψ from $\mathcal{I}$ to $\mathcal{J}$ there is an isomorphism ψ' that respects extreme points and maps each interval $A \in \mathcal{I}$ to $\psi'(A) = \psi(A)$.*

Proof. We proceed by induction on the number of overlap-connected components of $\mathcal{I}$. In the base case, $\mathcal{I}$ and $\mathcal{J}$ are overlap-connected. Let $I_1, \ldots, I_k$ be the slots of $\mathcal{I}$ and let $J_1, \ldots, J_k$ be the slots of $\mathcal{J}$ as described in Lemma 6.2. By Lemma 6.2, we can assume that an isomorphism ψ from $\mathcal{I}$ to $\mathcal{J}$ maps each I_s onto J_s; the other case is symmetric.

We show that for each $A \in \mathcal{I}$, the isomorphism ψ either respects the extreme points of A or can be locally modified to respect them. Let $A = [a^-, a^+]$ and $A = \bigcup_{s=p}^{q} I_s$. It follows that $\psi(A) = \bigcup_{s=p}^{q} J_s$, $a^- \in I_p$, and $a^+ \in I_q$. Moreover, if $\psi(A) = [b^-, b^+]$, then $b^- \in J_p$ and $b^+ \in J_q$.

Notice now that since $\mathcal{I}$ is sharp, every slot contains at most two points. Moreover, every two-point slot $[c^-, d^+]$ consists of the start point of some interval C and the end point of another interval D. The transposition of the points c^- and d^+ violates neither C nor D, nor any other interval.

If I_p is a one-point slot, we immediately conclude that $\psi(a^-) = b^-$. Suppose that $I_p = [a^-, x^+]$ is a two-point slot. Let $J_p = [b^-, y^+]$. If $\psi(a^-) = b^-$, we are done. Otherwise we can ensure $\psi'(a^-) = b^-$ by changing ψ only on I_p.

In order to ensure that $\psi'(a^+) = b^+$, we may need to modify ψ on I_q. In fact, we just need to inspect all two-point slots; if such a slot needs modification, this will simultaneously fix inconsistency between a pair of start points and a pair of end points. The analysis of the overlap-connected case is complete.

Suppose now that $\mathcal{I}$ and $\mathcal{J}$ have more than one overlap-connected component, that is, are not overlap-connected. If $\mathcal{I}$ and $\mathcal{J}$ are disconnected, then the claim readily follows by applying the induction assumption to the corresponding connected components of $\mathcal{I}$ and $\mathcal{J}$.

It remains to consider the case when $\mathcal{I}$ and $\mathcal{J}$ are connected but not overlap-connected. Assume that an interval $A \in \mathcal{I}$ contains an inner overlap-connected component $\mathcal{S} \subset \mathcal{I}$. Then $\psi(V(\mathcal{S})) \subset \psi(A)$ for any isomorphism ψ from $\mathcal{I}$ to $\mathcal{J}$. If we remove all points in $V(\mathcal{S})$ from $\mathcal{I}$ and all points in $\psi(V(\mathcal{S}))$ from $\mathcal{J}$, then the resulting interval systems $\mathcal{I}'$ and $\mathcal{J}'$ will still contain the extreme points of A and $\psi(A)$ respectively, and

[3] This follows from a simple observation that the conditions $B \subset A$, $B \between B'$, and $\neg(B' \between A)$ imply that $B' \subset A$.

[4] If $\mathcal{H}$ is an interval system, then this decomposition gives rise to the concept of a *PQ-tree* [BL76].

ψ will induce an isomorphism from $\mathcal{I}'$ to $\mathcal{J}'$ (recall that both interval systems are sharp). By the induction assumption, there are isomorphisms from $\mathcal{I}'$ to $\mathcal{J}'$ and from $\mathcal{S}$ to $\psi(\mathcal{S})$ that agree with ψ on hyperedges and respect extreme points. Merging them, we get the desired isomorphism ψ' from $\mathcal{I}$ to $\mathcal{J}$. ∎

When we want to apply Lemmas 6.1 and 6.3, the interval systems under consideration need to be sharp. It may happen that we deal with an isomorphic copy of a sharp interval system that itself is not sharp; consider for example, $\{[1,4],[1,2]\}$ that is isomorphic to $\{[1,4],[2,3]\}$. In such cases the following fact will be helpful.

Lemma 6.4. *Suppose that a given interval system $\mathcal{J}$ is isomorphic to a sharp interval system. Then an isomorphism from $\mathcal{J}$ to a sharp interval system $\mathcal{J}'$ can be computed in logspace.*

Proof. Suppose that $\mathcal{J}$ is isomorphic to a sharp interval system $\mathcal{S}$ and φ is an isomorphism from $\mathcal{J}$ to $\mathcal{S}$. Since $\mathcal{S}$ cannot contain any 1-point interval, the same holds true for any isomorphic system and in particular for $\mathcal{J}$. Furthermore, $\mathcal{J}$ cannot contain any point that serves simultaneously as the start point of an interval A and the end point of another interval B; otherwise the intervals $\varphi(A)$ and $\varphi(B)$ in $\mathcal{S}$ would also intersect at only one point and thus share an extreme point.

Given $\mathcal{J}$, we construct an interval system $\mathcal{J}'$ in three steps, each doable in logspace.

1. Remove all *interior points* from $\mathcal{J}$, that is, those points that are not extreme for any interval.
2. For each point x that is the start point of two or more intervals $A_1,\ldots,A_k$, do the following. W.l.o.g., assume that $A_1 \supset A_2 \supset \ldots \supset A_k$. Let $y \in A_1$ be the point next to x. We provide the intervals $A_2,\ldots,A_k$ with new pairwise distinct start points $a_2^-,\ldots,a_k^-$ that will be inserted between x and y in this order; x remains the start point of A_1.
3. Do similarly with the shared end points.

Being removed in the first step, interior points never appear later. The 2nd and the 3rd steps ensure that any two intervals in $\mathcal{J}'$ share neither the start nor the end point. Thus, $\mathcal{J}'$ is sharp. It remains to show that $\mathcal{J}'$ is isomorphic to $\mathcal{J}$.

Let V be a set of labels for the intervals in $\mathcal{S}$, $\mathcal{J}$, and $\mathcal{J}'$. Fix bijections $\lambda\colon V \to \mathcal{S}$, $\mu\colon V \to \mathcal{J}$, and $\mu'\colon V \to \mathcal{J}'$ such that for every label v, $\varphi(\mu(v)) = \lambda(v)$ and $\mu'(v)$ is obtained from $\mu(v)$ by the above transformation. In order to show that $\mathcal{J}' \cong \mathcal{J}$, we will show that $\mathcal{J}' \cong \mathcal{S}$ by using Lemma 3.1. To this end we need to check the equality $M_{\mu'} = M_\lambda$ of the corresponding pairwise-intersection matrices.

Note that the transformation of $\mathcal{J}$ into $\mathcal{J}'$ has the following features.

- If $\mu(u)$ contains the start (resp. end) point of $\mu(v)$ as an inner point, then $\mu'(u)$ contains the start (resp. end) point of $\mu'(v)$ as an inner point.
- If $\mu(u)$ does not contain the start (resp. end) point of $\mu(v)$, then $\mu'(u)$ does not contain the start (resp. end) point of $\mu'(v)$.
- If $\mu(u)$ and $\mu(v)$ share the start (resp. end) point and $\mu(u) \subset \mu(v)$, then $\mu'(u) \subset \mu'(v)$ and the start (resp. end) point of $\mu'(u)$ becomes an inner point of $\mu'(v)$.

These observations readily imply the following equivalences:

$$\mu(w) \cap \mu(v) = \emptyset \iff \mu'(w) \cap \mu'(v) = \emptyset, \tag{5}$$
$$\mu(w) \subset \mu(v) \iff \mu'(w) \subset \mu'(v), \tag{6}$$
$$\mu(w) \between \mu(v) \iff \mu'(w) \between \mu'(v). \tag{7}$$

Moreover, for every triple of pairwise overlapping intervals $\mu(w),\mu(u),\mu(v)$ we have

$$\mu(u) \cap \mu(v) \subseteq \mu(w) \iff \mu'(u) \cap \mu'(v) \subseteq \mu'(w). \tag{8}$$

Since $\lambda(v)$ and $\mu(v)$ correspond to each other under an isomorphism between $\mathcal{S}$ and $\mathcal{J}$, the equivalences (5)–(8) hold true also if μ is replaced with λ.

Suppose now that $\{\nu(v)\}_{v\in V}$ is a sharp interval system, for example, $\nu = \lambda$ or $\nu = \mu'$. Note that the pairwise-intersection matrix M_ν is completely determined by the set-theoretic relations between the intervals. Specifically,

$$|\nu(v)| = 2 + |\{w \in V : \nu(w) \between \nu(v)\}| + 2 \cdot |\{w \in V : \nu(w) \subset \nu(v)\}|.$$

Furthermore,

$$|\nu(u) \cap \nu(v)| = \begin{cases} 0 & \text{if } \nu(u) \cap \nu(v) = \emptyset, \\ |\nu(u)| & \text{if } \nu(u) \subset \nu(v), \\ |\nu(v)| & \text{if } \nu(u) \supset \nu(v). \end{cases}$$

Finally, if $\nu(u) \between \nu(v)$, then

$$\begin{aligned} |\nu(u) \cap \nu(v)| = 2 \quad &+ \quad 2 \cdot |\{w \in V : \nu(w) \subset \nu(u),\, \nu(w) \subset \nu(v)\}| \\ &+ \quad |\{w \in V : \nu(w) \subset \nu(u),\, \nu(w) \between \nu(v)\}| \\ &+ \quad |\{w \in V : \nu(w) \between \nu(u),\, \nu(w) \subset \nu(v)\}| \\ &+ \quad |\{w \in V : \nu(w) \between \nu(u),\, \nu(w) \between \nu(v),\, \nu(u) \cap \nu(v) \not\subseteq \nu(w)\}|. \end{aligned}$$

Since the set-theoretic relations are the same for λ and μ', we conclude that $M_{\mu'} = M_\lambda$. As claimed, Lemma 3.1 now implies that $\mathcal{J}' \cong \mathcal{S} \cong \mathcal{J}$.

In general, the algorithm is run on an arbitrary $\mathcal{J}$. After computing $\mathcal{J}'$ we invoke the algorithm of [KKL+11] to find a hypergraph isomorphism from $\mathcal{J}$ to $\mathcal{J}'$. In the case of failure, we conclude that the input system $\mathcal{J}$ is not isomorphic to any sharp interval system. ∎

7. A representation scheme for HCA graphs in logspace

We are now prepared to prove Theorem 1.1. By Lemma 4.3, it suffices to design a (not necessarily canonical) representation scheme for HCA graphs that have no twins and no universal vertices and to show that this scheme is computable in logspace.

Let G be an input graph on n vertices. We assume that G is HCA and has neither twins nor universal vertices. By Lemmas 5.2 and 5.3, we know that G admits a normalized Helly arc representation.

Lemma 7.1. *Let α be a normalized Helly arc representation of a graph G without twins and universal vertices. Then the pairwise-intersection matrix M_α depends on G only (being the same for all normalized Helly arc representations of G) and can be computed in logspace on input G.*

Proof. Consider first $m_{vv} = |\alpha(v)|$. The arc $\alpha(v)$ contains its two own extreme points and additionally, every vertex u adjacent to v contributes one or two extreme points of $\alpha(u)$ to $\alpha(v)$. More precisely, the following configurations are possible.

$\alpha(u) \subset \alpha(v)$: By Condition 1 in Definition 5.1, this happens exactly when $N[u] \subset N[v]$, which is verifiable in logspace. In this case, u contributes 2 points to $\alpha(v)$.

$\alpha(u) \circlearrowleft \alpha(v)$: By Condition 2 in Definition 5.1, this happens exactly when the logspace-verifiable conditions (1)–(3) are met. Also in this case, u contributes 2 points to $\alpha(v)$.

$\alpha(u) \frown \alpha(v)$: This is the remaining case: u contributes 1 point to $\alpha(v)$.

Consider now $m_{uv} = |\alpha(u) \cap \alpha(v)|$ for $u \neq v$. In the simplest case of non-adjacent u and v we have $m_{uv} = 0$. Also, $m_{uv} = m_{uu}$ if $\alpha(u) \subset \alpha(v)$ or, equivalently, $N[u] \subset N[v]$. Similarly, $m_{uv} = m_{vv}$ if $N[v] \subset N[u]$. Furthermore, $m_{uv} = m_{uu} + m_{vv} - 2n$ if $\alpha(u) \circlearrowleft \alpha(v)$, which is equivalent to (1)–(3).

It remains to compute m_{uv} if $\alpha(u) \frown \alpha(v)$. The intersection contains one extreme point of $\alpha(u)$ and one of $\alpha(v)$. Any other vertex w contributes 0, 1, or 2 extreme points of $\alpha(w)$. The contribution is 0 when

$\alpha(w)$ is disjoint from $\alpha(u)$ or $\alpha(v)$ or when it contains at least one of these arcs. Let us analyze the remaining cases (some cases symmetric up to swapping u and v are omitted). The first four conditions are verifiable in logspace similarly to the above.

$\alpha(w) \subset \alpha(u)$ and $\alpha(w) \subset \alpha(v)$: The vertex w contributes 2 to m_{uv}.

$\alpha(w) \subset \alpha(u)$ and $\alpha(w) \frown \alpha(v)$: The vertex w contributes 1 to m_{uv}.

$\alpha(w) \circlearrowleft \alpha(u)$ and $\alpha(w) \circlearrowleft \alpha(v)$: The vertex w contributes 2 to m_{uv}.

$\alpha(w) \circlearrowleft \alpha(u)$ and $\alpha(w) \frown \alpha(v)$: The vertex w contributes 1 to m_{uv}.

$\alpha(w) \frown \alpha(u)$ and $\alpha(w) \frown \alpha(v)$: This case is more complicated. Without loss of generality, suppose that $\alpha(u) \prec \alpha(v)$. Note first that the arc configuration $\alpha(v) \prec \alpha(w) \prec \alpha(u)$ cannot occur since it is non-Helly. There remain two subcases.

- $\alpha(u) \prec \alpha(w) \prec \alpha(v)$: This happens exactly when $\alpha(u) \cap \alpha(v) \subseteq \alpha(w)$, which is equivalent to the logspace-verifiable condition $N[u] \cap N[v] \subseteq N[w]$ by Lemma 5.4. In this subcase, the vertex w contributes 0 to m_{uv}.
- $\alpha(w) \prec \alpha(u)$ and $\alpha(w) \prec \alpha(v)$ or $\alpha(u) \prec \alpha(w)$ and $\alpha(v) \prec \alpha(w)$: This is the complementary subcase and w contributes 1 to m_{uv}.

The analysis is complete. The matrix entry m_{uv} is obtained by summing up the contributions of $\alpha(w)$ over all w. ∎

Our next task is to find an arbitrary maxclique $C \in \mathcal{C}(G)$. We have to argue that this is doable in logspace. An edge uv in a graph G is called *essential* if it is contained in a unique maxclique C. The following lemma implies that for each edge uv, we can check in logspace if it is essential. If so, then the corresponding maxclique C is equal to $N[u] \cap N[v]$ and, hence, can also be computed in logspace. A short proof of this simple fact is provided for the reader's convenience.

Lemma 7.2. *An edge uv is essential if and only if the intersection $N = N[u] \cap N[v]$ is a clique.*

Proof. Note first that any clique containing uv is included in N. Hence, if N is a clique then N is the only maxclique containing uv.

On the other hand, if N contains non-adjacent vertices x and y, then the two triangles $\{u, v, x\}$ and $\{u, v, y\}$ can be extended to two different maxcliques that contain uv. ∎

It is known [OR81] that if G is an interval graph without isolated vertices, then every maxclique in G contains an essential edge. This allows us to compute the bundle hypergraph $\mathcal{B}(G)$ in logspace, which is an important ingredient of our canonical representation scheme for interval graphs in [KKL+11]. However, there are HCA graphs that do not enjoy this property; an example is the Hajós (or 3-sun) graph depicted in Fig. 1(a). Fortunately, every nonempty HCA graph has at least one maxclique that can be efficiently found due to the fact that it contains an essential edge.

Lemma 7.3. *Every nonempty HCA graph G contains an essential edge uv.*

Proof. It is enough to consider the case that G has neither twins nor universal vertices. Our analysis is based on the Helly arc representation $\beta = \rho \circ \beta_G$ of G where ρ is an arc representation of the CA hypergraph $\mathcal{B}(G)$; see Lemma 4.2. Fix v to be a non-isolated vertex whose maxclique bundle B_v is minimal under inclusion. Note that $B_v \cup B_w = \mathcal{C}(G)$ for no vertex $w \in N[v]$ for else w would be universal. Thus, for every $w \in N[v]$ either $B_v \subseteq B_w$ or $B_v \between^* B_w$, where the last notation means that each of the four sets $B_v \cap B_w$, $B_v \setminus B_w$, $B_w \setminus B_v$, and $\mathcal{C}(G) \setminus (B_v \cup B_w)$ is nonempty. If $B_v \subseteq B_w$ for all $w \in N[v]$, then $N[v]$ is a clique and we are done (we can choose u arbitrarily from $N[v]$). Otherwise fix $u \in N[v]$ to be a vertex with $|B_v \cap B_u|$ being as small as possible. Note that $B_v \between^* B_u$.

It remains to argue that uv is an essential edge. By Lemma 7.2, we have to show that the intersection $N = N[u] \cap N[v]$ is a clique. Assume to the contrary that N contains non-adjacent x and y. Looking at the Helly arc representation β, we see that the arcs $\beta(x)$ and $\beta(y)$ must intersect the arc $\beta(v) \cap \beta(u)$ from different sides. Hence, one of $\beta(x)$ and $\beta(y)$ must contain the extreme point of $\beta(v)$ contained in $\beta(u)$. Without loss of generality, suppose that this is $\beta(x)$. It follows that $|\beta(v) \cap \beta(x)| < |\beta(v) \cap \beta(u)|$, which contradicts the assumption that $|B_v \cap B_u|$ is as small as possible. ∎

Lemma 7.4. *Let $\alpha\colon V(G) \to \mathcal{A}$ be a sharp Helly arc representation of a graph G without universal vertices and let $C \in \mathcal{C}(G)$ be a maxclique in G. Consider the C-flipped mapping $\alpha^C\colon V(G) \to \mathcal{A}^{\alpha(C)}$. Then there is a complete arc D with designated extreme points such that $\mathcal{I} = \mathcal{A}^{\alpha(C)}$ is an interval system on D.*

Proof. Since α is a Helly representation of G, the arcs in the set $\alpha(C)$ have a common point x. Let $A \in \alpha(C)$ be the arc that has x as one of its extreme points. Choose y to be the point of A next to x. Then the complete arc D having x and y as its designated extreme points fulfills the claimed property. ∎

We remark that the sharpness condition in Lemma 7.4 is crucial. Indeed, consider the graph G and its Helly arc representation α given in Fig. 3. The C_6-flipped mapping α^{C_6} results in a non-interval arc system.

Lemma 7.5. *Let α be a normalized Helly arc representation of a graph G without twins and universal vertices. Let $\lambda = \alpha^C$ where $C \in \mathcal{C}(G)$. Then M_λ can be computed in logspace from M_α and C.*

Proof. Let $M_\lambda = (m^\lambda_{uv})$ and $M_\alpha = (m^\alpha_{uv})$. We have $m^\lambda_{vv} = m^\alpha_{vv}$ if $v \notin C$ and $m^\lambda_{vv} = 2n + 2 - m^\alpha_{vv}$ if $v \in C$. For different u and v, m^λ_{uv} is computed by inspection of several cases. If $u \notin C$ and $v \notin C$, then $m^\lambda_{uv} = m^\alpha_{uv}$. Otherwise, m^λ_{uv} can be computed as detailed in Table 1; the case where $u \notin C$ and $v \in C$ is symmetric to that where $u \in C$ and $v \notin C$.

This completes the proof since as argued in the proof of Lemma 7.1, the relationship between $\alpha(u)$ and $\alpha(v)$ is recognizable in logspace. ∎

Now we can complete the description of our algorithm for computing a Helly arc representation of the input graph G. Suppose that $\alpha\colon V(G) \to \mathcal{A}$ is a normalized Helly arc representation of G. What follows does not depend on a particular choice of α.

Step 1. Compute the intersection matrix M_α. By Lemma 7.1, this matrix can be computed in logspace and does not depend on α.

Step 2. Compute a maxclique C of G. This is doable in logspace according to Lemmas 7.2 and 7.3.

Step 3. Compute the intersection matrix M_λ for the C-flipped mapping $\lambda = \alpha^C$. This can be done in logspace due to Lemma 7.5.

Note that by Lemma 7.4, the flipped arc system $\mathcal{I} = \mathcal{A}^{\alpha(C)}$ is actually an interval system.

Step 4. Compute an interval system $\mathcal{J}$ and a mapping $\mu\colon V(G) \to \mathcal{J}$ such that $M_\mu = M_\lambda$. For that purpose, we invoke the algorithm of Lemma 3.2.

Note that by Lemma 3.1, $\mathcal{J}$ and $\mathcal{I}$ are isomorphic hypergraphs.

Step 5. Modify μ and $\mathcal{J}$ so that $\mathcal{J}$ becomes sharp if it is not so from the very beginning. This is possible due to Lemma 6.4 because $\mathcal{J}$ is isomorphic to the sharp interval system $\mathcal{I}$.

Recall that λ is a mapping from $V(G)$ to $\mathcal{I}$. Lemma 3.1 ensures that there is a hypergraph isomorphism ψ from $\mathcal{I}$ to $\mathcal{J}$ such that

$$\mu = \psi \circ \lambda.$$

Moreover, by Lemma 6.3 we can assume that ψ respects extreme points of intervals in $\mathcal{I}$ and $\mathcal{J}$.

Step 6. Now, we "close" the interval $1, \ldots, 2n$ to the cycle where 1 succeeds $2n$ and regard $\mathcal{J}$ and $\mathcal{I}$ as arc systems, that possibly have complete arcs with designated extreme points. The mapping ψ stays a hypergraph isomorphism respecting extreme points of all arcs.

	$u \in C$ and $v \notin C$	$u \in C$ and $v \in C$
$\alpha(u) \cap \alpha(v) = \emptyset$	$m^{\lambda}_{uv} = m^{\alpha}_{vv}$	
$\alpha(u) \subset \alpha(v)$	$m^{\lambda}_{uv} = m^{\alpha}_{vv} - m^{\alpha}_{uu} + 2$	$m^{\lambda}_{uv} = 2n + 2 - m^{\alpha}_{vv}$
$\alpha(u) \supset \alpha(v)$	$m^{\lambda}_{uv} = 0$	$m^{\lambda}_{uv} = 2n + 2 - m^{\alpha}_{uu}$
$\alpha(u) \circlearrowleft \alpha(v)$	$m^{\lambda}_{uv} = 2n + 2 - m^{\alpha}_{uu}$	$m^{\lambda}_{uv} = 0$
$\alpha(u) \frown \alpha(v)$	$m^{\lambda}_{uv} = m^{\alpha}_{vv} - m^{\alpha}_{uv} + 1$	$m^{\lambda}_{uv} = 2n + 2 + m^{\alpha}_{uv} - m^{\alpha}_{uu} - m^{\alpha}_{vv}$

Table 1: The entries m^{λ}_{uv} of the matrix M_{λ} can be computed from M_{α}. Different rules apply depending on the relation of $\alpha(u)$ to $\alpha(v)$ and whether u and/or v belong to the flipping set C.

Step 7. From μ and C, compute the C-flipped mapping $\mu^C \colon V(G) \to \mathcal{J}^{\mu(C)}$.

Note that by Lemma 6.1,

$$\mu^C = \psi \circ \lambda^C = \psi \circ \alpha$$

and ψ is a hypergraph isomorphism from $\mathcal{I}^{\lambda(C)} = \mathcal{A}$ to $\mathcal{J}^{\mu(C)}$. It follows that, like α, the constructed mapping μ^C is a Helly arc representation of G.

The proof of Theorem 1.1 is complete. Note that the above argument also shows that any two normalized arc representations of an HCA graph without twins and universal vertices yield arc models that are isomorphic as hypergraphs.

Acknowledgement

We are grateful to an anonymous referee, whose detailed comments were very valuable for us in the course of preparing the revised version of the paper.

References

[BL76] K. Booth and G. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using *PQ*-tree algorithms. *J. Comput. Syst. Sci.*, 13(3):335–379, 1976.

[Che96] L. Chen. Graph isomorphism and identification matrices: Parallel algorithms. *IEEE Trans. Parallel Distrib. Syst.*, 7(3):308–319, 1996.

[CY91] L. Chen and Y. Yesha. Parallel recognition of the consecutive ones property with applications. *J. Algorithms*, 12(3):375–392, 1991.

[CLM+13] A. R. Curtis, M. C. Lin, R. M. McConnell, Y. Nussbaum, F. J. Soulignac, J. P. Spinrad, and J. L. Szwarcfiter. Isomorphism of graph classes related to the circular-ones property. *Discrete Mathematics & Theoretical Computer Science*, 15(1):157–182, 2013.

[DHH96] X. Deng, P. Hell, and J. Huang. Linear-time representation algorithms for proper circular-arc graphs and proper interval graphs. *SIAM J. Comput.*, 25(2):390–403, 1996.

[FG65] D. Fulkerson and O. Gross. Incidence matrices and interval graphs. *Pac. J. Math.*, 15:835–855, 1965.

[Gav74] F. Gavril. Algorithms on circular-arc graphs. *Networks*, 4(4):357–369, 1974.

[Hsu95] W.-L. Hsu. $O(MN)$ algorithms for the recognition and isomorphism problems on circular-arc graphs. *SIAM J. Comput.*, 24(3):411–439, 1995.

[JLM+11] B. L. Joeris, M. C. Lin, R. M. McConnell, J. P. Spinrad, and J. L. Szwarcfiter. Linear time recognition of Helly circular-arc models and graphs. *Algorithmica*, 59(2):215–239, 2011.

[KN09] H. Kaplan and Y. Nussbaum. Certifying algorithms for recognizing proper circular-arc graphs and unit circular-arc graphs. *Discrete Applied Mathematics*, 157(15):3216–3230, 2009.

[KKL+11] J. Köbler, S. Kuhnert, B. Laubner, and O. Verbitsky. Interval graphs: Canonical representations in Logspace. *SIAM J. on Computing*, 40(5):1292–1315, 2011.

[KKV12] J. Köbler, S. Kuhnert, and O. Verbitsky. Solving the canonical representation and Star System problems for proper circular-arc graphs in logspace. In *Proc. Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, number 18 in LIPIcs, pages 387–399. Leibniz-Zentrum für Informatik, 2012.

[KKW15] J. Köbler, S. Kuhnert, and O. Watanabe. Interval graph representation with given interval and intersection lengths. *Journal of Discrete Algorithms*, 34(9):108–117, 2015.

[KKV13] J. Köbler, S. Kuhnert, and O. Verbitsky. Helly circular-arc graph isomorphism is in logspace. In *Proc. 38th MFCS*, number 8087 in LNCS, pages 631–642. Springer, 2013.

[LSS08] M. C. Lin, F. J. Soulignac, and J. L. Szwarcfiter. A simple linear time algorithm for the isomorphism problem on proper circular-arc graphs. In *Proc. 11th Scandinavian Workshop on Algorithm Theory (SWAT)*, number 5124 in LNCS, pages 355–366. Springer, 2008.

[LB79] G. Lueker and K. Booth. A linear time algorithm for deciding interval graph isomorphism. *J. ACM*, 26(2):183–195, 1979.

[McC03] R. M. McConnell. Linear-time recognition of circular-arc graphs. *Algorithmica*, 37(2):93–147, 2003.

[MM99] T. A. McKee and F. R. McMorris. *Topics in intersection graph theory.* SIAM Monographs on Discrete Mathematics and Applications 2. Philadelphia: SIAM, 1999.

[OR81] R. J. Opsut and F. S. Roberts. On the fleet maintenance, mobile radio frequency, task assignment, and traffic phasing problems. In *The theory and applications of graphs*, pages 479–492. Wiley, 1981.

[Sou14] F. J. Soulignac. Minimal and short representations of unit interval and unit circular-arc graphs. `http://arxiv.org/abs/1408.3443v2`, 2014.

[Spi03] J. Spinrad. *Efficient graph representations.* Number 19 in Field Institute Monographs. AMS, 2003.

[Tor04] J. Torán. On the hardness of Graph Isomorphism. *SIAM J. Comput.* 33(5):1093–1108, 2004.

[Ueh08] R. Uehara. Simple geometrical intersection graphs. In *Proc. 2nd Int. Workshop on Algorithms and Computation (WALCOM)*, number 4921 in LNCS, pages 25–33. Springer, 2008.

[Ueh13] R. Uehara. Tractabilities and intractabilities on geometric intersection graphs. *Algorithms*, 6(1):60–83, 2013.