Induced Covering Numbers of Graphs

Bachelor's Thesis of

Yidi Zang

At the Department of Informatics
Institute of Theoretical Informatics (ITI)

Reviewer:	Dr. rer. nat. Torsten Ueckerdt
Second reviewer:	T.T.-Prof. Dr. Thomas Bläsius
Advisor:	Miriam Goetze

2024/12/19 – 2025/06/20

Karlsruher Institut für Technologie
Fakultät für Informatik
Postfach 6980
76128 Karlsruhe

I declare that I have developed and written the enclosed thesis completely by myself. I have not used any other than the aids that I have mentioned. I have marked all parts of the thesis that I have included from referenced literature, either in their original wording or paraphrasing their contents. I have followed the by-laws to implement scientific integrity at KIT.

Karlsruhe, 2025/06/20

. . . .

(Yidi Zang)

Abstract

Graph covering numbers describe the minimum number of graphs G_i of a given *guest class* $\mathcal{G}$ needed to cover a *host graph* H. While these numbers have been extensively explored, we consider a variant where the graphs G_i have to be induced subgraphs of H. This induced global $\mathcal{G}$-covering number is denoted by $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H)$. One extension of these covering numbers is the induced *union* $\mathcal{G}$-covering number $\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H)$, where each graph G_i can consist of multiple vertex-disjoint induced subgraphs of H, as long as each of them is in $\mathcal{G}$.

An induced covering number $\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(H)$ is called separable from the non-induced covering number $\mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(H)$, if there exists a guest class $\mathcal{G}$ and a host class $\mathcal{H}$, such that no function $f\colon \mathbb{N} \to \mathbb{N}$ exists, such that $\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(\mathcal{H}) \leq f(\mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(\mathcal{H}))$. In other words, the induced covering number cannot be bounded by its non-induced counterpart. In this thesis, we show that induced covering numbers can be separated even when only considering hosts H with bounded treewidth. Conversely, if we restrict the guest class $\mathcal{G}$ to only *monotone* graph classes, that is, graph classes closed under taking subgraphs, we show that for hosts H with bounded degeneracy, the induced union covering number is not separable.

Finally, we investigate the complexity of determining induced covering numbers. We present examples in which determining whether the induced union covering number is at most some integer k is NP-complete, and also show how to find a minimal induced global forest-cover of an outerplanar graph in linear time.

Zusammenfassung

Graph covering numbers beschreiben die minimale Anzahl von Graphen G_i aus einer gegebenen *Gastklasse* $\mathcal{G}$, die benötigt werden, um einen *Hostgraphen* H zu überdecken. Während diese Zahlen bereits umfassend untersucht wurden, betrachten wir eine Variante, bei der die Graphen G_i induzierte Teilgraphen von H sein müssen. Diese sogenannte induzierte globale $\mathcal{G}$-covering number wird mit $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H)$ bezeichnet. Eine Erweiterung dieses Konzepts sind die induzierten *union* $\mathcal{G}$-covering numbers $\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H)$, bei denen jeder Graph G_i aus mehreren paarweise disjunkten, induzierten Teilgraphen von H bestehen darf, sofern jeder dieser Teilgraphen zur Klasse $\mathcal{G}$ gehört.

Eine induzierte covering number $\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(H)$ heißt *separable* von der nicht-induzierten Überdeckungszahl $\mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(H)$, wenn es eine Gastklasse $\mathcal{G}$ und eine Hostklasse $\mathcal{H}$ gibt, so dass keine Funktion $f\colon \mathbb{N} \to \mathbb{N}$ existiert, mit der $\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(\mathcal{H}) \leq f(\mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(\mathcal{H}))$ gilt. Anders gesagt: Die induzierte covering number lässt sich nicht durch ihre nicht-induzierte Variante nach oben beschränken. In dieser Arbeit zeigen wir, dass induzierte covering numbers sogar dann trennbar sind, wenn man sich auf Hostgraphen H mit beschränkter Baumweite beschränkt. Umgekehrt zeigen wir, dass induzierte covering numbers nicht trennbar sind, sofern die Gastklasse $\mathcal{G}$ *monoton* ist, also unter dem Bilden von Teilgraphen abgeschlossen, und die Hostgraphen H eine beschränkte Degeneriertheit aufweisen.

Abschließend untersuchen wir die Komplexität der Bestimmung induzierter covering numbers. Wir präsentieren Beispiele, in denen es NP-vollständig ist, zu entscheiden, ob die induzierte Vereinigungsüberdeckungszahl höchstens einen gegebenen Wert k beträgt. Zudem zeigen wir, wie man in linearer Zeit ein minimal induziertes globales Wald-cover eines außenplanaren Graphen bestimmen kann.

Contents

1 Introduction

A graph cover of a *host graph* H with a *guest graph class* $\mathcal{G}$ is a multiset of graphs in $\mathcal{G}$ such that all edges in H are covered. The global *$\mathcal{G}$-covering number* of a host graph H is the minimum number of graphs of $\mathcal{G}$ needed in a cover. In one of the earliest studies of graph coverings by Erdös et al. in 1966 [EGP66], they showed that every host graph H can be covered with at most $\lfloor \frac{|V(H)|^2}{4} \rfloor$ triangles and singular edges. Since then, the global covering number has been explored extensively for different guest classes and host graphs, for example Nash-Williams proved exact bounds for the global covering number with forests [Nas64]. The covering number with forests as guest class is often referred to as *arboricity*, and many subclasses of forests have also been studied such as linear arboricity (linear forests, i.e. paths) [Alo88], star arboricity (star forests) [HMS96], and many more.

As many of these covering numbers have been looked at individually, Knauer and Ueckerdt [KU16] created a framework to unify the expression of these covering numbers. Some variants of covering numbers have shown to also be interesting to study.

They define three different $\mathcal{G}$-covers, the global, local, and folded $\mathcal{G}$-cover; each being less restrictive than the previous. The three types of $\mathcal{G}$-covers yield three covering numbers with $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}} \geq \mathrm{c}_{\mathrm{l}}^{\mathcal{G}} \geq \mathrm{c}_{\mathrm{f}}^{\mathcal{G}}$.

For example the *local* covering number, the cover does not minimize the number of graphs, but rather how often a vertex in H is covered. Fishburn et al. studied the local covering for the guest class of complete bipartite graphs [FH96].

The framework allows for more general results, for example the separability of the global and local covering number for certain guest class. The global covering number $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}$ is called separable from the local covering number $\mathrm{c}_{\mathrm{l}}^{\mathcal{G}}$, if there exists a guest class $\mathcal{G}$ and a host class $\mathcal{H}$, such that no function $f\colon \mathbb{N} \to \mathbb{N}$ exists, such that $\mathrm{c}_{\mathrm{l}}^{\mathcal{G}}(\mathcal{H}) \leq f(\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(\mathcal{H}))$. In other words, the local covering number can get arbitrarily larger than the global covering number.

Induced graph coverings are a variant of normal coverings, where there is an additional restriction on the graphs in the cover. Every graph of the guest class $\mathcal{G}$ used in an induced cover has to also be an induced subgraph of H. This requirement for induced subgraphs makes sense in applications where covers need to be more localized. When looking at the vertices of one graph G of any induced cover, all edges between these vertices have to be included in G. On the contrary, a non-induced cover may cover a dense part of the host with many sparse graphs which might be undesirable for certain applications. While induced covers do not prevent this entirely, the requirement for induced subgraphs certainly help by disallowing many sparse configurations.

Induced coverings are as far as we know largely unexplored compared to non-induced coverings. Axenovich, Ueckerdt et al. proved some results for induced forest- and star-arboricities [ADRU19]. They also look at *weak* induced arboricities, where not the entire subgraph has to be induced, but only the individual connected components.

Our goal in this thesis is transfer and expand the framework from “Three ways to cover a graph” [KU16] to induced graph coverings. We examine induced covering numbers on some well-known graph classes including linear forests and star forests. We also determine the separability of induced and non-induced covering numbers for host and guest classes with bounded degeneracy/treewidth. Finally, we establish the computational complexity for determining the induced covering numbers for certain guest classes.

2 Covering Numbers

2.1 Preliminaries

Let G and H be graphs. A subgraph H' of H is *induced* if every edge uv in $E(H)$ with $u, v \in V(H')$ is also an edge in H'.

A map $\varphi\colon V(G) \longrightarrow V(H)$ is called *graph homomorphism* if for every edge $uv \in E(G)$, the edge $\varphi(u)\varphi(v)$ exists in $E(H)$. A graph homomorphism is *edge-surjective* if for every edge $e \in E(H)$ some edge e' of G is mapped to e, i.e. $|\varphi^{-1}(uv)| \neq 0$.

The disjoint union $\dot\cup$ of graphs G_i is the vertex- and edge-disjoint union of these graphs G_i. In the resulting graph, the graphs G_i are not connected to each other.

For a graph class $\mathcal{G}$ we define $\overline{\mathcal{G}}$ as the graph class containing the disjoint union of one or more graphs in $\mathcal{G}$. For example, if $\mathcal{S}$ is the class of all star graphs, $\overline{\mathcal{S}}$ is the class of all star forest, that is the class where each graph consists of one or more disjoint stars.

2.2 Induced Covering Numbers

Let H be a graph and $\mathcal{G}$ be a graph class. We call H the *host* graph and $\mathcal{G}$ the *guest* class. A t-global $\mathcal{G}$-cover of H is an edge-surjective graph homomorphism $\varphi\colon (G_1 \dot\cup G_2 \dot\cup \ldots \dot\cup G_t) \longrightarrow H$, where for each $i \in [t], G_i \in \mathcal{G}$. The graphs G_i are called guest graphs. A cover is injective if for all $i \in [t]$, the restrictions $\varphi_{|G_i}$ are injective. In other words, $\varphi(G_i)$ is a copy of G_i in H.

A cover φ is called *induced* if for all $i \in [t]$, the graphs $\varphi(G_i)$ are induced subgraphs of H.

The induced global covering number denoted by $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H)$ is defined as

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H) := \min\{t\colon \text{ there exists an induced injective t-global } \mathcal{G}\text{-cover of } H\}.$$

To rephrase, the induced global covering number is the minimum number of graphs of $\mathcal{G}$ needed to cover all edges in the host graph H, where each graph used is an induced subgraph of H. Note that edges may be covered more than once.

An induced k-union $\mathcal{G}$-cover is an induced t-global $\mathcal{G}$-cover where the following graph is k-vertex-colorable. The graph has t vertices, each representing one of the graphs G_i in the t-global cover, and there is an edge connecting the vertices representing G_i and G_j if and only if they share a vertex in H, formally there exists $u \in V(G_i)$ and $v \in V(G_j)$ such that $\varphi(u) = \varphi(v)$.

In other words an induced k-union $\mathcal{G}$-cover is the minimum number of graphs needed to cover H, where each graph can be the union of multiple graphs G_i in $\mathcal{G}$ as long as they do not share a vertex and each G_i induces a subgraph in H. The union of multiple graphs G_i can be seen as a color class in the constructed graph on t vertices.

The induced union covering number denoted by $\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H)$ is defined as

$$\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) := \min\{k\colon \text{ there exists an induced injective k-union } \mathcal{G}\text{-cover of } H\}.$$

An induced k-local $\mathcal{G}$-cover is an induced $\mathcal{G}$-cover φ, where every vertex in H is hit at most k times, formally $\max_{v \in V(H)} |\varphi^{-1}(v)| \leq k$.

The induced local covering number denoted by $\mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H)$ is defined as

$$\mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H) := \min\{ \max_{v \in V(H)} |\varphi^{-1}(v)| : \varphi \text{ is an induced injective } \mathcal{G}\text{-cover of } H\}.$$

The local covering number does not minimize the number of guest graphs needed, but rather the maximum number of times a vertex in H is used.

The induced folded covering number denoted by $\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H)$ is defined as

$$\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H) := \min\{ \max_{v \in V(H)} |\varphi^{-1}(v)| : \varphi \text{ is an induced } \mathcal{G}\text{-cover of } H\}.$$

Note that a folded cover does not have to be injective.

2.3 Non-Induced Covering Numbers

Each of the four covering numbers $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}, \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}, \mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}, \mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}$ have a corresponding non-induced covering number $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{u}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{l}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{f}}^{\mathcal{G}}$, where the cover does not have to be induced.

$$\begin{aligned}
\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H) &:= \min\{t\colon \text{ there exists an injective } t\text{-global } \mathcal{G}\text{-cover of } H\}.\\
\mathrm{c}_{\mathrm{u}}^{\mathcal{G}}(H) &:= \min\{k\colon \text{ there exists an injective } k\text{-union } \mathcal{G}\text{-cover of } H\}.\\
\mathrm{c}_{\mathrm{l}}^{\mathcal{G}}(H) &:= \min\{\max_{v \in H} |\varphi^{-1}(v)| : \varphi \text{ is an injective } \mathcal{G}\text{-cover of } H\}.\\
\mathrm{c}_{\mathrm{f}}^{\mathcal{G}}(H) &:= \min\{\max_{v \in H} |\varphi^{-1}(v)| : \varphi \text{ is a } \mathcal{G}\text{-cover of } H\}.
\end{aligned}$$

Clearly, every induced covering number is at least the non-induced counterpart as every induced cover is also allowed as a non-induced cover.

Observation 2.1: *For any host graph H and guest class $\mathcal{G}$ we have*

$$\begin{aligned}
\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H) &\leq \mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H)\\
\mathrm{c}_{\mathrm{u}}^{\mathcal{G}}(H) &\leq \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H)\\
\mathrm{c}_{\mathrm{l}}^{\mathcal{G}}(H) &\leq \mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H)\\
\mathrm{c}_{\mathrm{f}}^{\mathcal{G}}(H) &\leq \mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H).
\end{aligned}$$

2.4 Monotonicity Of Induced Covering Numbers

Furthermore, the induced covering numbers also have an internal ordering.

Lemma 2.2: *For all guest classes $\mathcal{G}$ and all host graphs H, we have*

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H) \geq \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) \geq \mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H) \geq \mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H).$$

Clearly every t-$\mathcal{G}$-global cover is also a valid t-$\mathcal{G}$-union cover. Every k-union cover is also a valid k-local cover and every k-local cover is a valid k-folded cover. In all three cases the covering number cannot increase.

2.5 Linear Forest

We will now examine the upper bounds for induced covering numbers for some specific examples. Coloring numbers are sometimes closely related to covering numbers, the *strong chromatic index* is one of those.

The *strong chromatic index* of a graph is the minimum number of colors needed to color all edges such that each color induces a matching. Let $\mathcal{P}$ be the class of all paths.

Lemma 2.3: *For the guest class $\overline{\mathcal{P}}$, the class of disjoint unions of paths, and every host graph H we have*

$$\mathrm{ic}_{\mathrm{g}}^{\overline{\mathcal{P}}}(H) < 2 \cdot \Delta(H)^2.$$

Proof. Since the class $\overline{\mathcal{P}}$ contains all matchings, the upper bound is at most the covering number of induced matchings which corresponds to the strong chromatic index. Each color in strong chromatic coloring induces a matching, so the induced matchings of every color form an induced global $\mathcal{P}$-cover of H. As the strong chromatic index is known to be less than $2 \cdot \Delta(H)^2$ [MR97], we obtain $\mathrm{ic}_{\mathrm{g}}^{\overline{\mathcal{P}}}(H) < 2 \cdot \Delta(H)^2$. ∎

This bound is tight up to a constant factor. The complete graph K_n as the host serves as an example since every induced path is a singular edge, so the induced global covering number is equal to the number of edges.

Observation 2.4: *For every $n \in \mathbb{N}$ there exists a graph $H = K_n$ such that* $\mathrm{ic}_{\mathrm{g}}^{\overline{\mathcal{P}}}(H) > \frac{1}{2}\Delta(H)^2$.

Matchings can be used to find an upper bound for other induced covering numbers, so we start by looking at the induced matching-covering numbers.

Lemma 2.5: *For the guest class of all matchings $\mathcal{G} = \overline{\{P_2\}}$ and every host graph H we have*

$$\mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H) \leq \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) \leq 2 \cdot \left\lceil \frac{3\Delta(H) + 2}{5} \right\rceil$$
$$\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H) \leq \Delta(H) + 1.$$

Proof. Every (not necessarily induced) path in H can be split into two (non-induced) matchings by coloring the edges of the path in alternating colors, which are also union of paths. Using this construction, the induced $\mathcal{G}$-union covering number of H is at most twice the non-induced covering number with unions of paths, that is $\mathrm{ic}_{\mathrm{u}}^{\mathcal{P}}(H) \leq 2\,\mathrm{c}_{\mathrm{u}}^{\mathcal{P}}(H)$.

The non induced union $\mathcal{P}$-covering number $\mathrm{c}_{\mathrm{u}}^{\mathcal{P}}$ is also known as *linear arboricity.* Guldan shows in [Gul86] that for every graph H $\mathrm{c}_{\mathrm{u}}^{\mathcal{P}}(H) \leq \lceil \frac{3\Delta(H) + 2}{5} \rceil$.

The folded covering number $\mathrm{c}_{\mathrm{f}}^{P}(H)$ is in $\left\{\lceil \frac{\Delta(H)}{2} \rceil, \lceil \frac{\Delta(H) + 1}{2} \rceil\right\}$ [KU16]. ∎

Corollary 2.6: *For the guest class of all paths $\mathcal{P}$ and every host graph H we have*

$$\mathrm{ic}_{\mathrm{l}}^{\mathcal{P}}(H) \leq \mathrm{ic}_{\mathrm{u}}^{\mathcal{P}}(H) \leq 2 \cdot \lceil \frac{3\Delta(H) + 2}{5} \rceil$$
$$\mathrm{ic}_{\mathrm{f}}^{\mathcal{P}}(H) \leq \Delta(H) + 1.$$

As matchings are a subset of $\overline{\mathcal{P}}$, this directly follows from Lemma 2.5.

2.6 Star Forest

A tree G with at most one vertex with more than one neighbor is called a star. Let $\mathcal{S}$ be the class of all stars. The induced covering number with stars is also similar to the induced covering numbers from the previous section.

Corollary 2.7: *For the guest class $\mathcal{S}$ of all stars and every host graph H we have*

$$\mathrm{ic}_{\mathrm{l}}^{\mathcal{S}}(H) \leq \mathrm{ic}_{\mathrm{u}}^{\mathcal{S}}(H) \leq 2 \cdot \left\lceil \frac{3\Delta(H) + 2}{5} \right\rceil$$
$$\mathrm{ic}_{\mathrm{f}}^{\mathcal{S}}(H) \leq \Delta(H) + 1.$$

Since matchings are also star forests, the three induced $\mathcal{S}$-covering numbers above are bounded from above by the maximum degree of H by Lemma 2.5.

In a complete graph K_n, the only star that can be induced is the path with one edge, so the induced folded $\mathcal{S}$-covering number $\mathrm{ic}_{\mathrm{f}}^{\mathcal{S}}(K_n)$ is $\Delta(K_n)$. Therefore, the three bounds in Corollary 2.7 are asymptotically tight.

3 Separability

	Any		Hereditary		Monotone	
Host	g	u, l, f	g	u, l, f	g	u, l, f
any	Lemma 3.2	-	Lemma 3.6	Lemma 3.7	-	?
k colorable	-	-	-	?	-	?
k degenerate	-	-	-	?	Lemma 3.10	Lemma 3.9
M-minor-free	-	-	?	?	Lemma 3.14	-
treewidth k	Lemma 3.4	Lemma 3.5	?	?	Lemma 3.11	-

Figure 3.1: A table with the results of this chapter, green for separation, red for bounded, and yellow for unknown.

In this chapter, we first look at some relations between induced covering numbers and corresponding non-induced coloring numbers for all graphs and then for certain restrictions on the host or guest graph class. We already know from Observation 2.1 that the induced covering number is greater of equal to the non-induced counterpart. Here we examine by how much they can actually differ, that is, in which cases the induced covering number can actually be bounded from above by the non-induced covering number.

Definition 3.1: *For a guest class $\mathcal{G}$, and one of the four covering types* $\mathrm{x} \in \{\mathrm{g}, \mathrm{u}, \mathrm{l}, \mathrm{f}\}$ *we say a host class $\mathcal{H}$ is* $(\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{x}}^{\mathcal{G}})$-bounded, *if there exists a function f, such that*

$$\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(H) \leq f(\mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(H)) \quad \text{for all graphs } H \in \mathcal{H}.$$

If $\mathcal{H}$ is not $(\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{x}}^{\mathcal{G}})$-bounded, we call it separable.

When we talk about separation or boundedness of a type of covering number in this chapter, it will always be about the induced and non-induced counterpart of this type. This differs from the separation from one type of non-induced covering number to another type of non-induced covering number, for example, the separation of a global covering number from a local covering number.

The table Figure 3.1 show the results of this chapter. The rows are the restrictions on the host graph class $\mathcal{H}$, that is for an integer k and some graph M the host class $\mathcal{H}$ has to be k-colorable, k-degenerate, M-minor free, or treewidth at most k. The columns are restrictions on the guest class $\mathcal{G}$, that is it has to be hereditary or monotone. Each of those columns is split in two, the columns with the letter g stand for the global covering numbers, the other columns for union, local, and folded covering numbers. The green fields are separability results of this chapter, the red fields boundedness, and the yellow fields are not examined in this thesis.

3.1 General Guest Classes

We start with the most general case, where there are no restrictions on the guest and host class. No restrictions on the guest class allows for very constructed guest classes, so separability is not hard to achieve.

The global, union, and local induced covering numbers $\mathrm{ic_g}$, $\mathrm{ic_u}$ and $\mathrm{ic_l}$ cannot be bounded by their corresponding non-induced covering number.

Lemma 3.2: *There exists a host class $\mathcal{H}$ and a guest class $\mathcal{G}$ such that for every $n \geq 3$ there exists a host $H \in \mathcal{H}$ such that for $x \in \{\mathrm{g}, \mathrm{u}, \mathrm{l}\}$*

$$\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(H) \geq n - 1 \textit{ and } \mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(H) = 2.$$

Proof. Let the host class $\mathcal{H}$ be the class of all complete graphs. Let the guest class consist of K_2 and all complete graphs with one edge missing, that is we set
$\mathcal{G} = \{K_{n'} - e \mid n' \geq 3, e \in E(K_{n'})\} \cup \{K_2\}$.

For any $n \geq 3$ we consider the complete graph $H = K_n \in \mathcal{H}$. Clearly, every induced graph of a complete graph is also a complete graph, so K_2 is the only guest that is an induced subgraph of H. Therefore, there can only be matchings in any injective induced local $\mathcal{G}$-covering, so $\mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H) = n - 1$ since every matching covers at most one edge of a vertex. Using Lemma 2.2 we get $n - 1 = \mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H) \leq \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) \leq \mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H)$.

The non-induced covering number is 2 for all three covering types, as indeed for any edge $e \in E(H)$, the graphs $K_n - e$ and the graph which consists of only e form a 2-global injective $\mathcal{G}$-cover of H. In particular, we obtain $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H) = \mathrm{c}_{\mathrm{u}}^{\mathcal{G}}(H) = \mathrm{c}_{\mathrm{l}}^{\mathcal{G}}(H) = 2$. ∎

Adding restrictions to the host class $\mathcal{H}$ make the question of separability more interesting. Indeed, if there is a host class $\mathcal{H}' \subseteq \mathcal{H}$, such that $\mathrm{ic}_{\mathrm{x}}(\mathcal{H})$ and $\mathrm{c}_{\mathrm{x}}(\mathcal{H})$ can be separated, the covering numbers can also be separated for $\mathcal{H}$. Specifically, showing separability for more sparse host graphs is a stronger result than the previous Lemma 3.2. One way to limit the density of our host graphs is to limit their *treewidth*. The *treewidth* of a graph H is the minimum $k \in \mathbb{N}$, such that H is a subgraph of some k-tree. A k-tree is a graph formed by K_{k+1} and then adding more vertices, such that after each addition of a vertex v, v has exactly k neighbors and the k neighbors of v induce a clique. We use $\mathrm{tw}(G)$ as the notation for the treewidth of a graph G.

We introduce a graph class we will later use as our host class. For $n \in \mathbb{N}$ we denote the cycle on n vertices with C_n. Let C'_{2n} be a cycle graph on $2n$ vertices, with one extra chord: this one edge connects two vertices that are exactly distance n apart in the cycle. We call this graph a *split cycle*.

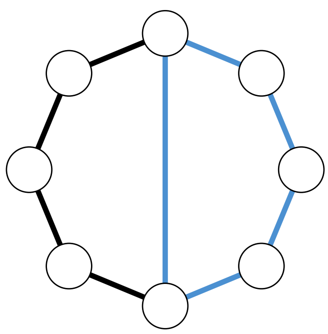

Figure 3.2: The split cycle C'_8, with one of the two smaller cycles colored blue.

Observation 3.3: *For every $n \geq 1$ there is no induced simple even cycle in C'_{4n}.*

Proof. C'_{4n} clearly only has three subgraphs which are cycles on more than two vertices. The two smaller cycles have $2n+1$ vertices, so the cycle is odd. The larger cycle uses all $4n$ vertices, which induces the entire graph C'_{4n}, which is not a cycle. ∎

Lemma 3.4: *There exists a host class $\mathcal{H}$ of treewidth 2 and a guest class $\mathcal{G}$ such that for any $n \geq 2$ there exists an $H_n \in \mathcal{H}$ with*

$$\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H_n) = 2 \text{ and } \mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H_n) \geq 4n + 1.$$

Proof. Let the guest class be the class of all even cycles $\mathcal{G} = \{C_{2k} : k \geq 2\} \cup \{K_2\}$. For every $n \geq 2$ let $H_n = C'_{4n}$ be the split cycle with $4n$ vertices. Note that this graph has treewidth 2. Let the host class $\mathcal{H}$ be $\{H_n : n \geq 2\}$.

For every $n \geq 2$ and H_n, the non-induced global $\mathcal{G}$-covering number is two, as it can be covered with a C_{4n} and the remaining chord with K_2, i.e. $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H_n) = 2$.

We now show that for every $n \geq 2$, the induced global $\mathcal{G}$-covering number $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H_n)$ is at least $4n + 1$. From Observation 3.3 we know H_n has no induced even cycle of more than two vertices. Therefore, the only induced even cycle is K_2, so any edge of C'_{4n} has to be covered individually, i.e. $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(C'_{4n}) = 4n + 1$. ∎

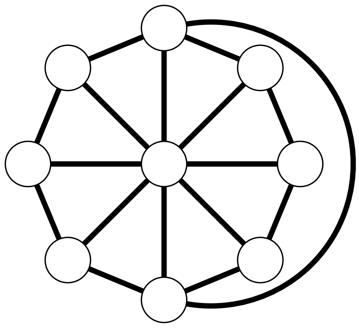

Figure 3.3: A wheel W_8, with one extra chord.

For the union covering numbers this construction is not enough for separability. As we can take many vertex disjoint edges, the induced union $\mathcal{G}$-covering number of C'_{4n} will not get large, even when $\mathcal{G}$ contains only K_2. To extend our construction to separate union covering numbers, we add one additional universal vertex to a split-cycle C'_{4n}. As we will show, this construction is enough to even separate local and folded induced covering numbers to their non-induced counterpart.

Lemma 3.5: *There exists a host class $\mathcal{H}$ of treewidth 3 and a guest class $\mathcal{G}$ such that for every $n \geq 2$ there exists a host $H \in \mathcal{H}$ such that for $x \in \{\mathrm{u}, \mathrm{l}, \mathrm{f}\}$*

$$\mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(H) = 2 \text{ and } \mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(H) \geq n$$

Proof. A cycle with one additional vertex connected to all other vertices is called a *wheel*. We denote the wheel with $k + 1$ vertices by W_k for $k \geq 3$. The size of a wheel is the number of vertices without the universal vertex, so W_k has size k. Let the guest class be the class of all even sized wheels and the complete graphs K_2 and K_3, that is $\mathcal{G} = \{W_{2k} : k \geq 2\} \cup \{K_2, K_3\}$. For an integer $n \geq 2$ we construct the graph H_n by taking a copy of W_{4k}. Denote the vertices as $v_1, \ldots, v_{4k+1}$. Assume v_1 is the universal vertex of W_{4k} and $v_2, \ldots, v_{4k+1}$ are the other vertices in cyclic order. We add the edge $v_2 v_{2k+2}$ to the graph H_n (Figure 3.3). Note that this graph is the same as a split cycle C'_{4k} with one extra universal vertex. Let the host class be $\mathcal{H} = \{H_n : n \geq 2\}$. Note that every H_n has treewidth 3 as C'_{4n} has treewidth 2 and H_n only has one additional vertex.

For every $n \geq 2$ the non-induced union $\mathcal{G}$-covering number $\mathrm{c}_{\mathrm{u}}^{\mathcal{G}}(H_n)$ is 2 as we can cover H_n with W_{4n} and a singular edge K_2.

For every $n \geq 2$ the induced union $\mathcal{G}$-covering number of H_n is at least n. Indeed, we show that there is no induced subgraph of H_n that is an even wheel on more than five vertices. Assume there is an induced subgraph G that is an even wheel on more than five vertices. In every wheel, there is one universal vertex connected to all other vertices. Denote this vertex with v_1. As there are more than four other vertices in G, this vertex has degree greater 4. There is only one vertex in H_n that has degree greater 4, that is v_1, so this vertex has to be the vertex connected to all others. The remaining vertices must induce an even cycle, but as H_n without vertex v_1 is C'_{4n}, and there is no induced even cycle in C'_{4n} by Observation 3.3. This contradicts the assumption that there is a subgraph that is an even wheel on more than 5 vertices.

As any cover uses wheels with at most 5 vertices and there is a vertex with degree $4n$, this vertex is hit at least n times, i.e. $\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H_n) \geq n$. ∎

3.2 Hereditary Guest Classes

As we have seen in the previous section, if there are no restrictions on the guest class, we may obtain many examples that lead to separability. We used some specifically constructed examples, which may feel a bit unnatural. In this section, we limit the guest class to be *hereditary*, as most commonly used graph classes are hereditary and therefore feel more natural. A graph class $\mathcal{G}$ is *hereditary* if and only if it is closed under taking induced subgraphs.

Even when only considering hereditary guest classes, the induced global covering number cannot be bounded by its corresponding non-induced global covering number.

Lemma 3.6: *There exists a host class $\mathcal{H}$ and a hereditary guest class $\mathcal{G}$ such that for every $n \geq 3$ there exists a host $H \in \mathcal{H}$ such that*

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H) \geq n \text{ and } \mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H) = 2.$$

Proof. For every $n \geq 3$, we construct the graph H_n that arises from a copy of the complete graph K_n. We denote by $v_1, v_2, \ldots, v_n$ the vertices of H_n. For each vertex v_i we add a new vertex u_i with an edge connecting v_i and u_i (Figure 3.4). We call these new edges $v_i u_i$ *hairs* of H_n. Now, we set $\mathcal{H} = \{H_n \mid n \geq 3\}$. The guest class $\mathcal{G}$ consists of the disjoint union of complete graphs. Note in particular that $\mathcal{G}$ is hereditary.

For any $n \geq 3$ we consider the host $H_n \in \mathcal{H}$. We have $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H_n) = 2$, as the edges of the complete graph in H_n can be covered with one complete graph and all other edges with one matching.

Now consider an induced injective $\mathcal{G}$-cover of H_n. We show that any guest $G \in \mathcal{G}$ covers at most one hair. Assume a guest $G \in \mathcal{G}$ contains at least two hairs $u_i v_i$ and $u_j v_j$, where $i \neq j$. The vertices u_i and u_j have degree 1, so they cannot be part of the same component of G, as each component is a complete graph. Yet, the vertices v_i and v_j induce the edge $v_i v_j$, which connects the components of u_i and u_j, a contradiction.

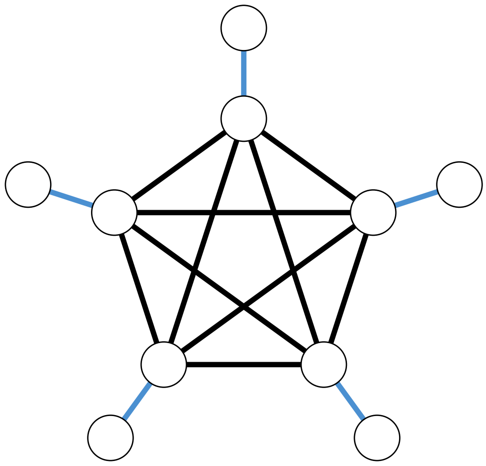

Figure 3.4: A complete graph K_5, with hairs colored blue.

We observe that the induced global $\mathcal{G}$-cover has to contain at least n guests as there are n hairs in H_n, i.e. $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H_n) \geq n$. ∎

Just as in the previous section, a slightly more complex construction is needed to separate union, local, and folded covering numbers, as Lemma 3.6 heavily relies on the fact that all guests are connected.

Lemma 3.7: *There exists a host class $\mathcal{H}$ and a hereditary guest class $\mathcal{G}$ such that for every $n \geq 2$ there is a host $H \in \mathcal{H}$ such that for $x \in \{\mathrm{u}, \mathrm{l}, \mathrm{f}\}$*

$$\mathrm{ic}_{\mathrm{x}}^{\mathcal{G}}(H) \geq n \text{ and } \mathrm{c}_{\mathrm{x}}^{\mathcal{G}}(H) = 2.$$

Proof. Let the guest class $\mathcal{G}$ be the class of all complete graphs and stars. Note that this class is hereditary. For every $n \geq 2$ construct H_n by starting with n disjoint copies of K_n. Add one extra vertex that is connected to all vertices (Figure 3.5). We call this vertex the *center vertex*. Let $\mathcal{H}$ be the class of all H_n for $n \geq 2$.

For every $n \geq 2$, H_n can be non-induced covered with one star covering all edges of the center vertex and n disjoint K_n covering the edges of the copies of K_n. Therefore, the non-induced union, local and folded $\mathcal{G}$-covering number is 2.

We now show that the induced folded $\mathcal{G}$-covering number $\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H_n)$ is at least n. We claim that no graph $G \in \mathcal{G}$ can cover more than n edges of the center vertex.

Assume a graph $G \in \mathcal{G}$ covers more than n edges of the center vertex. We denote the n copies of K_n with $K_n^1, \dots, K_n^n$. Since the center vertex has n edges to each of the n copies of K_n, according to the pigeon-hole principle, there are edges to at least two distinct K_n^i and K_n^j and there are at least two edges to the same K_n^l that are covered by G. Two vertices of distinct K_n^i and K_n^j do not have an edge between them, so G cannot be a complete graph. Two vertices of the same K_n induce the edge between them, therefore they both have degree at least 2 as they are also connected to the center vertex. This means G can also not be a star and therefore is not in $\mathcal{G}$, a contradiction to the assumption.

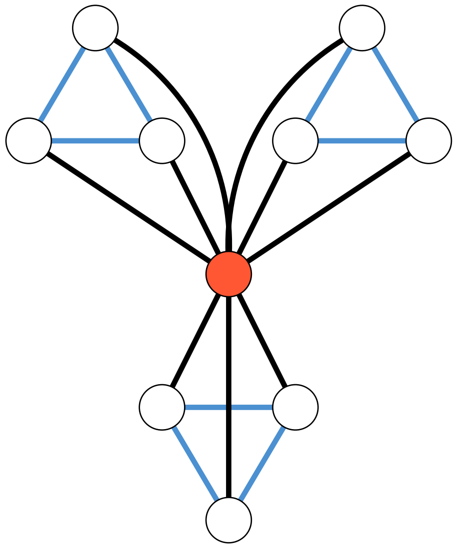

Figure 3.5: Three copies of K_3, with one universal vertex colored red.

As the center vertex is incident to n^2 edges and each $G \in \mathcal{G}$ can cover at most n of those edges, the folded covering number is at least n, i.e. $\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H_n) \geq n$. From Lemma 2.2 we obtain $\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}} \geq \mathrm{ic}_{\mathrm{l}}^{\mathcal{G}} \geq \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}$, from which the remaining inequalities of this lemma follow. ∎

For many host classes the question of separability remains open for hereditary guest classes. One interesting question might be which restrictions to the host class separate their induced union covering number from the non-induced covering number.

3.3 Monotone Guest Classes

A further natural restriction is requiring the guest classes to be *monotone*. A graph class $\mathcal{G}$ is *monotone* if and only if it is closed under taking (not necessarily induced) subgraphs. In this section, we only consider monotone guest classes. When a graph class is monotone, it often contains many useful subgraphs for covering. We will use the fact, that every monotone graph class contains all stars, or if the graph class has bounded maximum degree all stars up to that degree. To start, we show that for monotone guest classes, all host classes with bounded *degeneracy* have bounded covering numbers for union, local, and folded induced covering numbers.

A graph G has *degeneracy* d, if in every subgraph $G' \subseteq G$ there exists a vertex with degree at most d. Let $d(G)$ denote the minimum degeneracy of a graph G and $d(\mathcal{G}) = \sup_{G \in \mathcal{G}} \{d(G)\}$ denote the degeneracy of a graph class $\mathcal{G}$.

For host graphs with bounded degeneracy, we have the following useful lemma.

Lemma 3.8 ([ADRU19, Theorem 7]): *For every host graph H, we have $\mathrm{ic}_{\mathrm{u}}^{\mathcal{S}}(H) \leq 2d(H)$, where $\mathcal{S}$ is the class of all stars.*

Using this lemma, we show that when only considering monotone guest classes and host graph classes with bounded degeneracy, the union, local and folded induced covering numbers can be bounded by their non-induced counterpart.

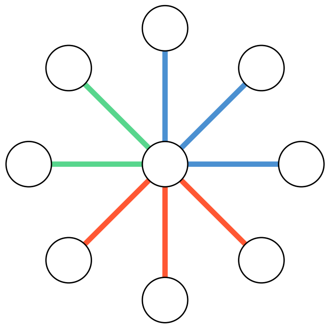

Figure 3.6: The star S_8, divided into two S_3 and one S_2.

Lemma 3.9: *For every monotone guest class $\mathcal{G}$, we have for every host H*

$$\begin{aligned} \mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) &\leq 2\, d(H) \cdot \mathrm{c}_{\mathrm{u}}^{\mathcal{G}}(H) \\ \mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(H) &\leq 2\, d(H) \cdot \mathrm{c}_{\mathrm{l}}^{\mathcal{G}}(H) \\ \mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H) &\leq 2\, d(H) \cdot \mathrm{c}_{\mathrm{f}}^{\mathcal{G}}(H). \end{aligned}$$

In particular, for every monotone guest class $\mathcal{G}$, every host class $\mathcal{H}$ with bounded degeneracy is $(\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(\mathcal{H}), \mathrm{c}_{\mathrm{u}}^{\mathcal{G}}(\mathcal{H}))$-, $(\mathrm{ic}_{\mathrm{l}}^{\mathcal{G}}(\mathcal{H}), \mathrm{c}_{\mathrm{l}}^{\mathcal{G}}(\mathcal{H}))$-, and $(\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(\mathcal{H}), \mathrm{c}_{\mathrm{f}}^{\mathcal{G}}(\mathcal{H}))$-bounded.

Proof. We will start the proof for the case that the maximum degree of every vertex in every graph $G \in \mathcal{G}$ is bounded. Let k be this maximum degree in $\mathcal{G}$ and H be any graph in $\mathcal{H}$. Since $\mathcal{G}$ is monotone, it also contains all stars with up to k leaves, that is vertices of degree 1. According to Lemma 3.8, any host H with degeneracy at most d admits a $2d$-union $\mathcal{S}$-cover. We take any valid induced $2d$-union $\mathcal{S}$-covering of H. While this induced union cover may use stars with up to $\Delta(H)$ leaves, there are only stars up to degree k in $\mathcal{G}$. Yet, every star of degree at most $\Delta(H)$ is the union of at most $\lceil \frac{\Delta(H)}{k} \rceil$ many stars in $\mathcal{G}$ (see Figure 3.6). Therefore, $\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) \leq 2d \cdot \lceil \frac{\Delta(H)}{k} \rceil$.

The non-induced folded $\mathcal{G}$-covering number is at least $\lceil \frac{\Delta(H)}{k} \rceil$ as the vertex in H with the highest degree has to be used in at least that many stars to cover all incident edges, so $\mathrm{c}_{\mathrm{f}}^{\mathcal{G}}(H) \leq \lceil \frac{\Delta(H)}{k} \rceil$.

Combining these two inequalities we get $\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}(H) \leq 2d \cdot \frac{\Delta(H)}{k} \leq 2d \cdot \mathrm{c}_{\mathrm{f}}^{\mathcal{G}}(H)$. The remaining inequalities of this lemma follow from this and Lemma 2.2.

In the case that the maximum degree of the graphs $G \in \mathcal{G}$ is not bounded, it is easy to see that from the monotony of $\mathcal{G}$ we can follow that the class of all stars $\mathcal{S}$ is contained in $\mathcal{G}$. We immediately get $\mathrm{ic}_{\mathrm{f}}^{\mathcal{G}}(H) \leq 2d(H)$ from Lemma 3.8, which also show all inequalities of this lemma. ∎

While the union, local, and folded induced covering numbers are bounded for graphs with bounded degeneracy, the same is not true for the global induced covering number. We can use the fact, that for induced global coverings, all connected components of a guest must have distance of at least two. Indeed, if there are two connected components C_1 and C_2 with distance 1, at least one edge between C_1 and C_2 would be induced. With this, we can construct an example where the induced global covering number can get arbitrarily large, while the non-induced global covering number stays the same.

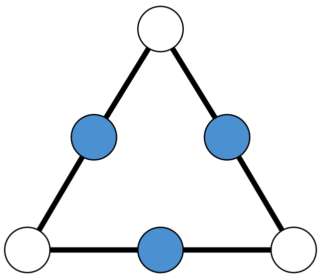

Figure 3.7: The 1-edge subdivision of the graph K_3, the new vertices are colored blue.

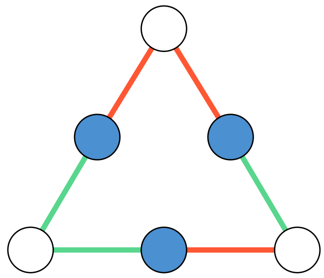

Figure 3.8: Coloring of the 1-edge subdivision of the graph K_3.

Lemma 3.10 ([AGRU18, Theorem 4(ii)(a)]): *There exists a host class $\mathcal{H}$ with bounded degeneracy and a monotone guest class $\mathcal{G}$ such that for any $n \geq 2$ there exists a host $H \in \mathcal{H}$ such that*

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H) \geq n-1 \text{ and } \mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H) \leq 2$$

Proof. For any fixed $n \geq 2$ let K_n' be the graph obtained by subdividing each edge of K_n once. That is, every edge uv in K_n is replaced by a uv-path with two edges. We denote the original vertices of K_n by $V := \{v_1, \ldots, v_n\}$ and the new vertices of the edge subdivision by V'. As all new vertices $v' \in V'$ have degree 2 and the original vertices $v \in V$ only have edges to new vertices, K_n' clearly has degeneracy 2. Let $\mathcal{H}$ be the class of all K_n' for $n \geq 2$ and the guest class $\mathcal{G}$ be the class of all star forests $\overline{\mathcal{S}}$. Note that this guest class is monotone.

First we show $\mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(K_n') \leq 2$. Let us partition the edges of K_n' into two color classes so that each of the two color classes corresponds to a star forest. Note that every edge in K_n' connects an original vertex $v \in V$ to a new vertex $v' \in V'$. Each of the new vertices $v' \in V'$ has exactly two neighbors u and w. Color the two edges uv' and $v'w$ with two different colors. As every edge is only incident to exactly one new vertex, such a coloring exists. It is easy to see that both colors form star forests with the original vertices of K_n being the centers of the stars (that is the only vertex in a star with degree greater 1) as seen in Figure 3.8.

Now we show $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(K_n') \geq n-1$. Take any induced t-global $\overline{\mathcal{S}}$-cover $\varphi\colon F_1 \dot\cup \ldots \dot\cup F_t \to K_n'$ of K_n'. Let uv' be an edge in a forest F_i and w be the only other neighbor of v' in K_n'.

If the edge $v'w$ exists in the same star forest, v' has to be a center of a star. This means u and w are leaves of F_i and cannot have any more edges in this star forest.

If the edge $v'w$ does not exist in the same star forest, no other edge with endpoint w belongs to the star forest F_i, as that would induce $v'w$.

In both cases for each edge, uv' (and $v'w$), one of the original vertices in K_n, w (and u), cannot have any more incident edges. Therefore, any star forest can only contain up to n edges and at least $\frac{2 \cdot \binom{n}{2}}{n} = n-1$ star forests have to be used. ∎

While the induced global covering number $\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}$ for monotone guest classes cannot be bounded in terms of the degeneracy of the host and the non-induced counterpart, it can be bounded when using the treewidth of the host graph.

Lemma 3.11: *For any monotone guest class $\mathcal{G}$ and every host class H*

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(H) \leq 3^{\binom{\mathrm{tw}(H)+1}{2}} \cdot \mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(H).$$

Proof. Every host graph H is the union of at most $3^{\binom{\mathrm{tw}(H)+1}{2}}$ induced star forests [ADRU19, Theorem 8]. The remaining proof is identical to the previous lemma. ∎

Another type of host classes, where the induced global covering number can be bounded in terms of the non-induced global covering number is the class of M-minor free graphs. Here we can use some coloring numbers to prove a bound.

The *t-coloring $col_t(G)$* of a graph G is defined as the minimum integer k, such that an ordering of the vertices in G exists, where from every vertex $v \in V(G)$ at most $k-1$ smaller vertices of the ordering can be reached using paths with at most t edges.

Lemma 3.12 ([Van+17, Corollary 1.3])**:** *For every graph G that excludes the complete graph K_n as a minor, we have*

$$col_t(G) \leq \binom{n-1}{2} \cdot (2t+1).$$

We use this upper bound to determine a bound for the *star chromatic number*, which is the number of colors needed to color a graph G, such that any two colors induce a star forest.

Lemma 3.13 ([JNM23, Page 188])**:** *For every graph G, $col_2(G)$ is equal to the star chromatic number.*

Using the star chromatic number, we can easily find an upper bound for the induced global covering number.

Lemma 3.14: *For every K_n-minor free host class $\mathcal{H}$ and every monotone guest class $\mathcal{G}$, we have*

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}(\mathcal{H}) \leq \binom{5\binom{n-1}{2}}{2} \cdot \mathrm{c}_{\mathrm{g}}^{\mathcal{G}}(\mathcal{H}).$$

Proof. For any K_n-minor free host graph H we know from Lemma 3.12 that the 2-coloring number is at most $5 \cdot \binom{n-1}{2}$. Furthermore, Lemma 3.13 states that the star chromatic number is equal to the 2 strong coloring number, so also $5 \cdot \binom{n-1}{2}$. From the definition of star chromatic number, any two colors of a star chromatic coloring induces a star forest. The induced star forest of all pairs of colors clearly cover the entire graph G, so the induced global covering number is at most $\binom{5\binom{n-1}{2}}{2}$. ∎

Corollary 3.15: *For every monotone guest class $\mathcal{G}$, every M-minor free host class is $(\mathrm{ic}_{\mathrm{g}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{g}}^{\mathcal{G}})$-bounded.*

With this section, the question of separability for monotone guest classes is solved for most classes of hosts. What remains open is whether hosts with unbounded degeneracy are separable for union, local, and folded induced covering numbers.

Question 3.16: *Are all hosts classes $\mathcal{H}$ $(\mathrm{ic}_{\mathrm{u}}^{\mathcal{G}}, \mathrm{c}_{\mathrm{u}}^{\mathcal{G}})$-bounded for all monotone guest classes $\mathcal{G}$?*

4 Complexity of Determining Induced Covering Numbers

4.1 NP-complete examples

Lemma 4.1 ([KSW78]): *Determining whether the non-induced global clique-cover number of a given graph H is at most a given integer k is NP-complete.*

As every subgraph in a graph H that is a clique is also an induced clique in H, this is also trivially true for the induced global clique-cover number.

Corollary 4.2: *Determining whether the induced global clique-cover number of a given graph H is at most a given integer k is NP-complete.*

Lemma 4.3: *In any triangle-free graph H, we have*

$$\mathrm{ic}_{\mathrm{u}}^{\mathcal{S}}(H) = \mathrm{c}_{\mathrm{u}}^{\mathcal{S}}(H).$$

where $\mathcal{S}$ is the class of all stars.

Proof. We claim that every non-induced union $\mathcal{S}$-cover of every triangle-free graph H is also an induced union $\mathcal{S}$-cover. Assume a non-induced union $\mathcal{S}$-cover $\varphi\colon S_1, \dots, S_t \to H$ of a triangle-free graph H is not a valid induced union $\mathcal{S}$-cover. This means that at least one star S_i in the union $\mathcal{S}$-cover is not induced. Let uv be one edge in H, where uv is not in $E(S_i)$, but is induced by S_i in H. As this edge is not in the star, neither of the two vertices u and v is the center vertex of the star. Let us denote the center vertex of the star by c. The three edges cu, cv, and uv are all in H, and form a triangle, contradicting the assumption that H is triangle-free. Therefore, every non-induced $\mathcal{S}$-cover of a triangle-free graph H is also a valid induced union $\mathcal{S}$-cover of H. Since the induced covering number is always at least as large as the non-induced counterpart, this proves the equality $\mathrm{ic}_{\mathrm{u}}^{\mathcal{S}}(H) = \mathrm{c}_{\mathrm{u}}^{\mathcal{S}}(H)$. ∎

Lemma 4.4: *For any graph H if there exists a union $\mathcal{S}$-cover of size k, then there also exists an union $\mathcal{S}$-cover of size at most k such that every edge in H is covered exactly once.*

Proof. Clearly, all subgraphs of a star S are also stars. Given a non-induced union $\mathcal{S}$-cover $\varphi\colon S_1, \dots, S_t \to H$, as long as there is an edge in H that is covered more than once, we can remove one instance of that edge in one star S_i of the cover. We end up with a cover of size at most k, where no edge in H is covered more than once. ∎

In fact, a covering as described in Lemma 4.4 corresponds to an edge-coloring of H, where every color forms a star forest. We call an edge-coloring, where every color forms a star forest, a *star-forest-coloring*.

Corollary 4.5: *For any graph H if there exists a union $\mathcal{S}$-cover of size k, then there also exist a star-forest-coloring into k colors.*

Lemma 4.6: *If there exist a star-forest-coloring of a graph H with k colors, there also exist a union $\mathcal{S}$-cover of size k.*

Proof. Take any star-forest-coloring with k colors, and for every color, the star forest of that color. The stars in these k star forest form a union $\mathcal{S}$-cover. Since the stars of a star forest are vertex-disjoint, this union $\mathcal{S}$-cover has size at most k. ∎

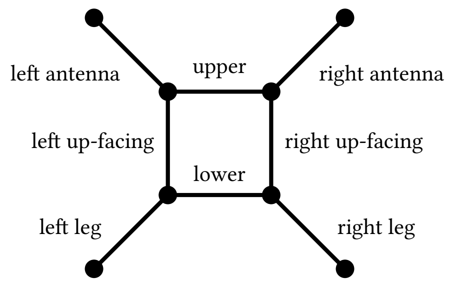

Figure 4.1: Variable gadget with its edge names

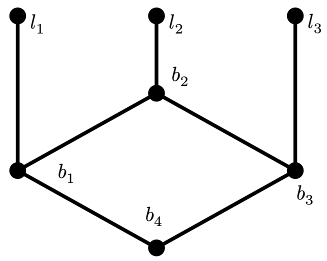

Figure 4.2: Clause gadget

Theorem 4.7 ([HMS96, Theorem 4])**:** *Determining whether* $\mathrm{c}_\mathrm{u}^{\mathcal{S}}(H) \leq 2$ *is NP-complete for even* 2*-degenerate graphs H.*

A Not-All-Equal 3SAT instance, consists of variables V and clauses C. In every clause $c \in C$, there are exactly three literals of variables in V. A clause c denotes, that the three literal cannot all three be true, or all three be false. Deciding whether a Not-All-Equal 3SAT instance has a valid assignment of variables to true or false is NP-complete [Sch78].

Theorem 4.8 ([Sch78])**:** *Not-All-Equal* 3*SAT is NP-complete.*

Our goal is to prove Theorem 4.7 by reducing Not-All-Equal 3SAT to 2-Star-Forest-Cover, the problem described in Theorem 4.7. Given any Not-All-Equal 3SAT instance I with variables V and clauses C, the auxiliary graph H contains one copy of the variable gadget (Figure 4.1) for every variable $v \in V$ and a copy of the clause gadget (Figure 4.2) for every clause $c \in C$. In every variable gadget, we call the left vertex of the lower edge the "true" vertex and the right one the "false" vertex. We add an edge from every literal vertex l_1, l_2, and l_3 of every clause gadget to the corresponding variable gadget, either the "true" vertex or the "false" vertex depending on the literal type. We call these added edges the *connecting edges*, and together with the left (or right) leg, they form the left (or right) *down-facing* edges of a variable gadget.

In the following we will prove that $\mathrm{ic}_\mathrm{u}^{\mathcal{S}}(H) \leq 2$ is true if and only if there exists a valid variable assignment in I.

Within this section, we refer to the two colors of any 2-coloring as blue and red.

Lemma 4.9: *If there exists a star-forest-coloring of H with* 2 *colors, then for every variable gadget, all its left down-facing edges have the same color c_1 and all its right down-facing edges also have the same color c_2. Note that (for this lemma) c_1 and c_2 are not necessarily distinct.*

Proof. Take any star-forest-coloring of edges with 2 colors. There clearly cannot be any monochromatic paths of length 3, that is a path 3 edges.

We use the same notation as in Figure 4.1. Assume there is a variable gadget that has left down-facing edges in both colors. Without loss of generality there is at least one red right down-facing edge in the same gadget. The lower edge has to be blue, otherwise there would be a red path of length 3. The right up-facing edges has to be red as otherwise it would extend the blue path with the lower edge and a blue left down-facing edge to length 3. The upper edge and the right antenna have to be blue to not extend the red path with the right up-facing edge. The left up-facing edge has to be red to not extend the blue path with the upper edge with the right antenna. The left antenna cannot be either color since being red would extend the red path with the left up-facing edge, and being blue would extend the path with the upper edge together with the right antenna. This contradicts there being a coloring with two colors, so all left down-facing edges of each variable gadget have to be the same color.

The proof that all right down-facing edges have the same color is identical with left and right swapped. ∎

Lemma 4.10: *If there exists a star-forest-coloring of H with 2 colors, then the left down-facing edges of all variable gadgets G_v have to be colored differently than the right down-facing edges in G_v.*

Proof. From Lemma 4.9 we know that the same side down-facing edges all have the same color. Assume left and right down-facing edges have the same color. Without loss of generality let all down-facing edges be red. The lower edge has to be blue to avoid the red path of length 3. We then choose the color c_1 for the left up-facing edge. The antenna and upper edge have to be in the other color c_2, as the up-facing edge forms a monochromatic path of length 2 with either a down-facing edge or the lower edge. The same is true for the right up-facing edge. Since the upper edge has to be different from both up-facing edges those two must have the same color c_3. The left and right antenna also have to be a different color from their corresponding up-facing edge c_3, so they also have to be in the same color c_4. The left antenna, upper edge, and right antenna all have the same color now and form a monochromatic path of length 3, so this connected monochromatic component is not a star, this contradicts the assumption. Therefore, the left down-facing and right down-facing edges have distinct colors. ∎

Lemma 4.11: *For every variable gadget G_v, if all left down-facing edges have the same color, all right down-facing edges have the same color, and these two colors are distinct, then it is possible to star-forest-color the edges of a variable gadget with 2 colors.*

Proof. Without loss of generality let the left down-facing edges be red and the right down-facing edges be blue. Color the lower edge red, the both up-facing blue, the left antenna blue, and the upper edge and the right antenna red. As there every monochromatic connected component is a star this construction proves the lemma. ∎

Lemma 4.12: *If there exists a star-forest-coloring of H with 2 colors, then every edge $l_j b_j$ from every clause gadget is colored differently from the other incident edge from a variable gadget into l_j.*

Proof. Let the other incident edge from the variable gadget be $l_j a$ and the color of that edge be c. As proven in Lemma 4.9, all down-facing edges incident to a have the same color c. The vertex a also has an incident leg in that same color c, so that leg together with the edge $l_j a$ already form a monochromatic path on two edges. Assume the edge $l_j b_j$ would also be in the same color. This would extend the monochromatic path to three edges, which means this monochromatic component does not form a star and therefore cannot be part of a star. This contradiction proves that each edges $l_j a_i$ has a different color from the other incident edge of $l_j b_j$. ∎

Lemma 4.13: *If there exists a star-forest-coloring of H with 2 colors, then in every clause gadget the three incoming connecting edges do not all have the same color.*

Proof. As a reminder, *connecting edges* are the edges connecting a vertex gadget with a clause gadget. Assume all three of the connecting edges of a clause gadget G_v have the same color, let that color be blue. We know from Lemma 4.12 that all three edges $l_j b_j$ for $j \in [3]$ have to be blue. The two edges $b_1 b_2$ and $b_2 b_3$ must be blue as otherwise they would connect the two edges $l_1 b_1$ and $l_3 b_3$ and form a path of length 3. The two other edges $b_1 b_4$ and $b_3 b_4$ must be red as otherwise they would extend the blue path with $b_1 b_2$ and $b_2 b_3$. Now the path (l_1, b_1, b_4, b_3, b_1) is red and therefore cannot be part of a star. This contradicts the claim that all three connecting edges of the same clause gadget have the same color. ∎

Lemma 4.14: *If not all three connecting edges of a clause gadget have the same color, then the edges of the clause gadget are star-forest-colorable with two colors.*

Proof. In the case that the two edges $l_1 b_1$ and $l_3 b_3$ are colored into the same colors, we color the remaining edges as follows. Without loss of generality let $l_1 b_1$ and $l_3 b_3$ be red and $l_2 b_2$ be blue. We color $b_1 b_2$ and $b_3 b_4$ blue, and $b_1 b_4$ and $b_2 b_3$ red.

In case that the two edges $l_1 b_1$ and $l_3 b_3$ have the different color, we color the remaining edges as follows. Without loss of generality let $l_1 b_1$ and $l_2 b_2$ be red and $l_3 b_3$ be blue. We color $b_1 b_2$ and $b_3 b_4$ blue, and $b_1 b_4$ and $b_2 b_3$ red.

In both cases both colors form a star forest. ∎

Lemma 4.15: *If the graph H admits a union $\mathcal{S}$-cover of size 2, then the Not-All-Equal 3SAT instance I is a yes-instance.*

Proof. If there exists a $\mathcal{S}$-cover of size 2, we know from Lemma 4.6 that there also exist a star-forest-coloring with 2 colors. Take any such coloring. From Lemma 4.9 we know that for each variable gadget all down-facing edges from the “true” vertex are in the same color and all down-facing edges from the “false” vertex are in the other. Assign every variable with red colored down-facing edges from the “true” vertex to true, and every other variable to false. From Lemma 4.9 and Lemma 4.10 we know that every variable gadget has exactly one “true” or “false” vertex where all down-facing edges are all red, and the other one has only blue down-facing edges. We claim that this a valid assignment.

Assume there is a clause that is violated. Then there exists a clause where all three clause-variables are assigned to the same value. From our assignment we know all three connecting edges to the clause gadget all have the same color. After Lemma 4.13 we know that if all three connecting edges have the same color, there is no coloring of edges into two colors, such that all connected components of one color is a star. Furthermore, according to Lemma 4.6 this means there is no 2 union $\mathcal{S}$-cover of H. This violates the assumption. Therefore, all clauses have to be valid, so the assignment is valid. ∎

Lemma 4.16: *If the Not-All-Equal 3SAT instance I is a yes-instance, then the graph is union $\mathcal{S}$-coverable with size 2.*

Proof. Take any valid assignment of variables of the NOT-ALL-EQUAL 3SAT instance. If a variable v is assigned to true, we color all down-facing edges of the variable gadget G_v from the "true" vertex red, and any down-facing edges from the "false" vertex blue. We color these edges exactly the other way if the corresponding variable is assigned false. From Lemma 4.11 we extend the star-forest-coloring with two colors to the vertex gadget G_v. As we started with a valid instance of the instance I, no three connecting edges into a clause gadget have the same color, so by Lemma 4.14 we extend the star-forest coloring with two colors to the clause gadget. As this colors all edges in a valid way, according to Lemma 4.6 H has a union $\mathcal{S}$-cover of size 2, which proves the lemma. ∎

The construction of the graph H from the Not-All-Equal 3SAT instance I is clearly polynomial, and as Lemma 4.15 and Lemma 4.16 have shown, $\mathrm{c}_{\mathrm{u}}^{\mathcal{S}}(H) \leq 2$ is true if and only if there exists a valid variable assignment of the Not-All-Equal 3SAT instance I. This proves that deciding $\mathrm{c}_{\mathrm{u}}^{\mathcal{S}}(H) \leq 2$ is NP-complete, and as the construction of H is clearly 2-degenerate, Theorem 4.7.

Corollary 4.17: *Determining* $\mathrm{ic}_{\mathrm{u}}^{\mathcal{S}} \leq 2$ *is NP-complete for the class of* 2*-degenerate graphs.*

This follows directly from Lemma 4.3 and Theorem 4.7, as the construction in Theorem 4.7 does not contain any triangle.

We can do the same proof idea to extend a result for a non-induced star-covering number to the induced star-covering number.

Lemma 4.18 ([GO09, Theorem 5])**:** *Determining* $\mathrm{c}_{\mathrm{u}}^{\mathcal{S}} \leq 3$ *is NP-complete for bipartite planar graphs.*

Corollary 4.19: *Determining* $\mathrm{ic}_{\mathrm{u}}^{\mathcal{S}} \leq 3$ *is NP-complete for bipartite planar graphs.*

As there are no triangles in a bipartite graph, this follows directly from Lemma 4.3 and Lemma 4.18

4.2 Induced Forest-Covering of Outerplanar Graphs

After looking at some cases where induced coverings are NP-complete, we look at one example where a minimal induced global cover can be found in linear time in regard to the number of vertices. Here we introduce an algorithm to determine the global induced forest-covering number of an outerplanar graph.

Theorem 4.20: *For every outerplanar graph H with n vertices, we can determine the induced forest-global covering number* $\mathrm{ic}_{\mathrm{g}}^{\mathcal{F}}(H)$ *in linear time* $\mathcal{O}(n)$*, where $\mathcal{F}$ is the class of all forests. We can also find such a cover in* $\mathcal{O}(n)$.

Axenovich et al showed that every outerplanar graph can be induced global covered with at most three forests.

Lemma 4.21 ([ADRU19, Theroem 9])**:** *The global induced forest-covering number of any outerplanar graph H is at most* 3*, that is*

$$\mathrm{ic}_{\mathrm{g}}^{\mathcal{F}}(H) \leq 3$$

where $\mathcal{F}$ is the class of all forests.

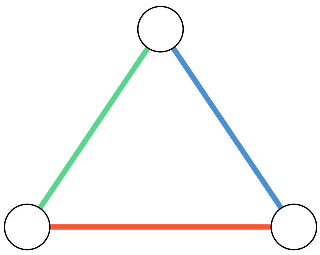

Figure 4.3: A Triangle requires three induced forest to cover all edges.

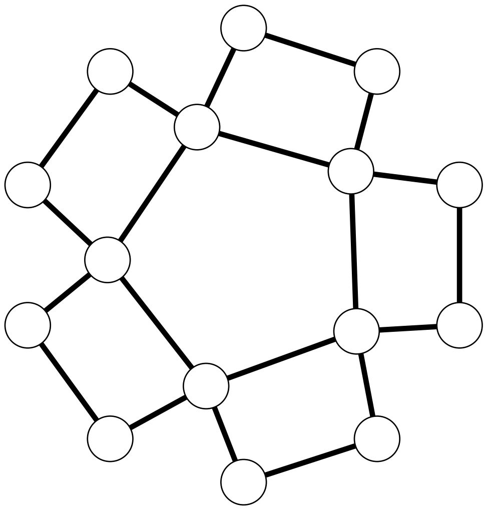

Figure 4.4: A triangle free graph that requires three induced forests to cover all edges.

One easy example where an outerplanar graph requires three forest in an induced cover is a triangle. In fact, as soon as an outerplanar graph H contains a triangle, the induced global forest-covering number is 3. One might assume that two forests are enough in any triangle-free outerplanar graph H, but this is not true for the graph in Figure 4.4. Finding an induced global cover for any outerplanar graph H with three forest is also not difficult.

As every outerplanar graph H has degeneracy of at most 2, it is trivially proper vertex-colorable with tree colors.

Observation 4.22: *Every outerplanar graph is proper vertex-colorable with three colors.*

Lemma 4.23: *For every outerplanar graph H with n vertices, we can find an induced global $\mathcal{F}$-cover of size 3 in linear time $\mathcal{O}(n)$.*

Proof. This follows from part of the proof of [HMS96, Theorem 2]. We add edges to H until H is maximal outerplanar (inner triangulated). It is easy to see that these added edges cannot decrease the induced global covering number, as the class of all forests is monotone, that is close under taking subgraphs. We finding any proper vertex-coloring of H with three colors, as this is always possible Observation 4.22. We now show the subgraph induced by any two colors is a forest. Assume the induced subgraph of two colors is not a forest, that is, it contains at least one cycle. Take the smallest subgraph that is a cycle $C \subseteq F$. Since our chosen 3-coloring is proper, C is bipartite, so the cycle C must contain at an even amount of vertices. But since H is maximal, there must exist edges in H, that form a triangle with vertices of C. This contradicts that the chosen cycle C is minimal, therefore there is no cycle in the induced subgraph with two colors F. As there are only three pairs of colors, these three induced forests form an induced global forest-cover of H of size 3. ∎

Determining whether the induced global forest-covering number $\mathrm{ic}_g^{\mathcal{F}}$ of an outerplanar graph H is 1 is trivial, as it is exactly true if H is a forest itself. Therefore, we only have to determine whether the induced global forest-covering number is 2 or 3, and in the case it is 2, find any two induced forests that cover H.

We will now describe how to determine whether two induced forests can cover H and if yes, how to find one such cover.

Theorem 4.24: *For every outerplanar graph H with n vertices, we can determine in linear time $\mathcal{O}(n)$, whether there exists two induced forests in H, that cover all edges. If so, we can also find one such pair in $\mathcal{O}(n)$.*

We start by trying to color the vertices of H in two colors, let us call the two colors blue and red. Every vertex has to be colored, but we also allow a vertex to be colored in both colors. The property of this coloring is that the blue and red vertices each induce a forest, and every edge in H is covered by at least one of the two forests. We call such a coloring a *two-forest-coloring*.

For this to be true, the following two things need to be true:

- For every edge $uv \in E(H)$, u and v both share at least one color.
- The red (blue) vertices do not induce any cycle.

It is easy to see that any vertex coloring fulfilling these criteria is a two-forest-coloring. If a two-forest-coloring exist, the global induced forest-covering number is at most 2, as the two forests induced by the colors form such a cover. Clearly, any two induced forests which cover all edges can also be represented by such a coloring.

Observation 4.25: *For every outerplanar graph H, the existence of a two-forest-coloring is equivalent to $\mathrm{ic}_{\mathrm{g}}^{\mathcal{F}}(H) \leq 2$.*

To simplify the process, let us initially assume the outerplanar graph H is 2-vertex connected.

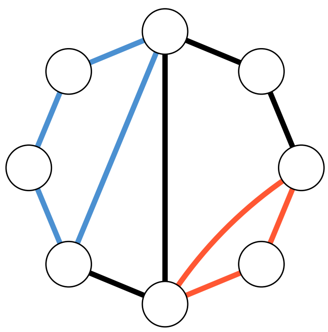

Figure 4.5: Two outermost faces of an outerplanar graph colored blue and red.

We call a face C of an outerplanar graph outermost, if all but one edge are on the outer face.

Observation 4.26: *Every 2-vertex connected outerplanar graph with at least one chord has at least two outermost faces.*

Clearly any 2-vertex connected outerplanar has at least one outermost face. If the outerplanar graph has at least one chord, take that and divide the outerplanar graph into two. Clearly, both halves have an outermost face.

Lemma 4.27: *For every 2-vertex connected outerplanar graph H with n vertices, we can determine in $\mathcal{O}(n)$, whether there exists two induced forests in H, that cover all edges. If there exists two such forests, then we can also find one such pair in linear time.*

Proof. The outer face of every 2-vertex connected outerplanar graph H is a cycle on all n vertices. We use a dynamic programming approach to find a valid coloring of the vertices efficiently. In one state, we maintain two vertices u and v, the colors of each of the two vertices, and whether they are connected in each of the two colors. This state represents whether it is possible to color all vertices in the cyclic order from u to v inclusive, such that u and v each are colored red, blue, or in both colors, and are connected in red, blue, or both colors. We initialize all adjacent pairs of vertices of the cyclic ordering, with all possible ways to assign colors to the two vertices such that they share at least one color. The connectivity in each color is automatically determined by whether the colors of the two vertices induce this color in the edge uv, in other words, they are connected in the shared colors of u and v.

Then as long as H contains at least one chord, take any outermost face F. Let v_1v_k be the chord of this outermost face F and denote the vertices of this face in cyclic order by $v_1, v_2, \ldots, v_k$. We now merge adjacent states of this face F, that is we start with the states with vertices v_1 and v_2, and merge them with the states with vertices v_2 and v_3. For this merge, we iterate over all states v_1 and v_2, let us call these states s_i, and for each such state s_i we also iterate over all states with v_2 and v_3, let us call those s'_j. When the color of v_2 is the same in both states s_i and s'_j, we create a new state with v_1 and v_3, where v_1 has the color it has in state s_i and v_3 has the color it has in state s'_j. In this state, v_1 and v_3 are connected in each color c exactly if in both state s_i and s'_j the two vertices are connected in color c. After this merge, we merge the new states of vertices v_1 and v_3 with the states of vertices v_3 and v_4 and so on, until we are only left with states of vertices v_1 and v_k.

Here we for each state s_i with v_1 and v_k, the color (red, blue, or both) of the chord between these two vertices is determined by the colors of v_1 and v_k. In case the two vertices do not share a color, this state s_i is invalid, as both endpoints of all edges must contain a common color. If this chord is colored in a color c and in this state s_i v_1 and v_k are already connected in that color c, the chord would close a monochromatic cycle, and is therefore invalid, otherwise, we create a new state with added the colors of the chord added to connectivity. Only the newly created states of v_1 and v_k are valid, as the old states do not consider the chord. We can now remove all edges of the outerplanar face F except for the chord. Clearly, the remaining graph remains a 2-vertex connected outerplanar graph. In this remaining graph, only the states of two now adjacent vertices v_1 and v_k have been added, so all states still contain only pairs of vertices adjacent in the cyclic ordering.

We repeat this process with any remaining outermost face until only one face remains. This last face can be handled exactly the same, only that we can choose one edge as the "chord", and it does not matter which edge this is. If at the end, any valid state remains, H is induced global 2-forest coverable. We can construct this cover by taking any remaining state of the remaining pair uv, coloring these two vertices in these two colors, and then backtracking from which two states this state was created. Then we color the vertices of these states in those colors and so on, until all vertices are colored. Since we maintained the previously mentioned criteria, that each vertex has a color, for all edges the two endpoints u and v share at least one common color, and no color induces a cycle, this induced cover is a valid induced global $\mathcal{F}$-cover of H. In case there are no states left, there exists no valid vertex coloring fulfilling the criteria, and therefore $\mathrm{ic}_{\mathrm{g}}^{\mathcal{F}}(H) = 3$.

For one fixed pair of vertices u and v, there can clearly only be a constant amount of states containing these two vertices u and v. As we only merge as often as there are edges in H, and every outerplanar graph G has at most $2|V(G)| - 3$ edges, this entire process can be done in linear time with regard to the number of vertices. ∎

Note that in the coloring of Lemma 4.27 we can also determine all possible colors of a fixed vertex v by always choosing an outermost face that does not contain v and finish the final face with an edge containing v.

Until now, we assumed H is connected and 2-vertex connected. If H is not connected, we can solve each connected component separately, so this does not matter. In case it is not 2-vertex connected, we can look at the block-cut-tree of H. This block-cut-tree contains all 2-vertex connected components, and two adjacent components share exactly one vertex. We can take any leaf component of the block-cut-tree and use Lemma 4.27 to find valid colorings in that 2-vertx connected component C. Here we also determine the possible colors of the vertex v that component C shares with at least one other component. When we initialize the states of a component C' that also contains vertex v, we need to make sure we allow no state where vertex v has a color it cannot have in component C. We can repeatedly remove leaf components of the block-cut-tree until nothing remains, then we either have no valid states left and there is no induced global forest-cover of H of size 2, or we can construct this cover again by backtracking our states. The block-cut-tree and the algorithm in Lemma 4.27 both run in linear time with regard to the number of vertices in H, so this proves Theorem 4.24.

In total, we can determine whether the induced global forest-covering number of any outerplanar graph H with n vertices is 1, 2, or 3 in $\mathcal{O}(n)$. We are also always able to find such a cover, so this proves Theorem 4.20.

Bibliography

[ADRU19] Maria Axenovich, Philip Dörr, Jonathan Rollin, and Torsten Ueckerdt. "Induced and weak induced arboricities". In: *Discrete Mathematics* Volume 342 (2019), pp. 511–519.

[AGRU18] Maria Axenovich, Daniel Gonçalves, Jonathan Rollin, and Torsten Ueckerdt. "The k-strong induced arboricity of a graph". In: *European Journal of Combinatorics* Volume 67 (2018), pp. 1–20.

[Alo88] Noga Alon. "The linear arboricity of graphs". In: *Israel Journal of Mathematics* Volume 62 (1988), pp. 311–325.

[EGP66] Paul Erdös, Adolph W Goodman, and Louis Pósa. "The representation of a graph by set intersections". In: *Canadian Journal of Mathematics* Volume 18 (1966), pp. 106–112.

[FH96] Peter C Fishburn and Peter L Hammer. "Bipartite dimensions and bipartite degrees of graphs". In: *Discrete Mathematics* Volume 160 (1996), pp. 127–148.

[GO09] Daniel Gonçalves and Pascal Ochem. "On star and caterpillar arboricity". In: *Discrete Mathematics* Volume 309 (2009), pp. 3694–3702.

[Gul86] Filip Guldan. "Some results on linear arboricity". In: *Journal of graph theory* Volume 10 (1986), pp. 505–509.

[HMS96] S Louis Hakimi, John Mitchem, and Edward Schmeichel. "Star arboricity of graphs". In: *Discrete Mathematics* Volume 149 (1996), pp. 93–98.

[JNM23] Yiting Jiang, Jaroslav Nešetřil, and Patrice Ossona de Mendez. "From χ-to χp-bounded classes". In: *Journal of Combinatorial Theory, Series B* Volume 158 (2023), pp. 186–209.

[KSW78] Lawrence T. Kou, Larry J. Stockmeyer, and Chak-Kuen Wong. "Covering edges by cliques with regard to keyword conflicts and intersection graphs". In: *Communications of the ACM* Volume 21 (1978), pp. 135–139.

[KU16] Kolja Knauer and Torsten Ueckerdt. "Three ways to cover a graph". In: *Discrete Mathematics* Volume 339 (2016), pp. 745–758.

[MR97] Michael Molloy and Bruce Reed. "A bound on the strong chromatic index of a graph". In: *Journal of Combinatorial Theory, Series B* Volume 69 (1997), pp. 103–109.

[Nas64] C St JA Nash-Williams. "Decomposition of finite graphs into forests". In: *Journal of the London Mathematical Society* Volume 1 (1964), pp. 12–12.

[Sch78] Thomas J Schaefer. "The complexity of satisfiability problems". In: *Proceedings of the tenth annual ACM symposium on Theory of computing*. 1978, pp. 216–226.

[Van+17] Jan Van Den Heuvel, Patrice Ossona de Mendez, Daniel Quiroz, Roman Rabinovich, and Sebastian Siebertz. “On the generalised colouring numbers of graphs that exclude a fixed minor”. In: *European Journal of Combinatorics* Volume 66 (2017), pp. 129–144.