

Labeling Schemes for Bounded Degree Graphs

David Adjiashvili

Institute for Operations Research, ETH Zürich
Rämistrasse 101, 8092 Zürich, Switzerland
email: addavid@ethz.ch

Noy Rotbart

Department of Computer Science, University Of Copenhagen
Universitetsparken 5, 2200 Copenhagen
email: noyro@diku.dk

October 29, 2018

Abstract

We investigate adjacency labeling schemes for graphs of bounded degree $\Delta = O(1)$. In particular, we present an optimal (up to an additive constant) $\log n + O(1)$ adjacency labeling scheme for bounded degree trees. The latter scheme is derived from a labeling scheme for bounded degree outerplanar graphs. Our results complement a similar bound recently obtained for bounded depth trees [Fraigniaud and Korman, SODA 10], and may provide new insights for closing the long standing gap for adjacency in trees [Alstrup and Rauhe, FOCS 02]. We also provide improved labeling schemes for bounded degree planar graphs. Finally, we use combinatorial number systems and present an improved adjacency labeling schemes for graphs of bounded degree Δ with $(e+1)\sqrt{n} < \Delta \leq n/5$.

1 Introduction

A labeling scheme is a method of distributing the information about the structure of a graph among its vertices by assigning short *labels*, such that a selected function on pairs of vertices can be computed using only their labels. The quality of a labeling scheme is mostly measured by its *size*: that is, the maximum number of bits used in a label. Additional important attributes of labeling schemes are the running times of the label generation algorithm (*encoder*), and the decoding algorithm (*decoder*), which replies to a query given a pair of labels.

Among all labeling schemes, that of *adjacency* is perhaps the most fundamental, as it directly comprises an implicit representation of the graph. For a graph G and any two of its vertices u, v , the decoder of an adjacency labeling scheme is required to deduce whether u and v are adjacent in G directly from their labels. Adjacency queries for bounded degree graphs appear naturally in networks of small dilation [4], peer-to-peer (P2P) [18] and wireless ad-hoc networks [22].

Our main contribution is an optimal (up to an additive constant) $\log n + O(1)$ size adjacency labeling scheme for bounded-degree outerplanar graphs. As a special case thereof, we obtain an optimal labeling scheme for bounded degree trees. We summarize this result in the following theorem.

Theorem 1. *For every fixed $\Delta \geq 1$, the class $\mathcal{O}(n, \Delta)$ of bounded-degree outerplanar graphs admits an adjacency labeling scheme of size $\log n + O(1)$, with encoding complexity $O(n \log n)$ and decoding complexity $O(\log \log n)$.*

Our labeling scheme utilizes a novel technique based on *edge*-universal graphs¹ for bounded degree outerplanar graphs. Unlike other results in the field which rely on a tight connection to induced-universal graphs² [1, 3, 14, 17], our technique embeds the input graph into a small edge-universal graph. Moreover, to the best of our knowledge, our labeling scheme is the first to use the total label size to separate the different components of the label. In contrast, other labeling schemes, such as [21, 2], introduce an extra overhead to support such separation.

Kannan, Naor and Rudich [15] showed that if a graph class \mathcal{R} admits an adjacency labeling scheme with maximum label length $g(n)$, then there exists an induced-universal graph with $2^{g(n)}$ vertices for \mathcal{R} , efficiently constructible from the labeling scheme. The opposite relation holds in a weaker sense. The existence of an induced-universal graph with $2^{g(n)}$ vertices for a family \mathcal{R} of graphs implies the existence of labeling scheme with size $g(n)$. This transformation is however not efficient, namely the resulting scheme has exponential running time. In light of the existing linear size induced-universal graphs for bounded degree trees [9], our contribution is in devising an *efficient* labeling scheme of optimal size.

As a corollary of Theorem 1, we also obtain an efficient $(\lfloor \Delta/2 \rfloor + 1) \log n$ size labeling scheme for graphs with maximum degree Δ . For the case of bounded degree planar graphs we construct a $\frac{\lfloor \Delta/2 \rfloor + 1}{2} \log n$ size labeling scheme with average label size of $(1 + o(1)) \log n + O(\log \log n)$, improving the best known construction for $\Delta \leq 4$. Finally, we observe that a simple application of combinatorial number systems [16] gives an adjacency labeling scheme for all graphs with maximum degree $\Delta(n)$, improving the known bounds for $\Delta(n) \in [(e+1)\sqrt{n}, n/5]$.

We summarize all known results for adjacency labeling schemes in Table 1, and our contributions in Table 2. Our results for outerplanar graphs, planar graphs and general graphs are presented in Section 2, Appendix B and Section 3, respectively. *We defer all proofs to appendix A.*

1.1 Previous Work

Alstrup and Rauhe [3] proved that forests (and trees) in $\mathcal{G}(n)$ have an adjacency labeling scheme of $\log n + O(\log^* n)$.³ Their technique uses a recursive decomposition of the tree which yields the same $\log n + O(\log^* n)$ label size for bounded degree trees. Fraigniaud and Korman [11] showed that bounded depth trees have a labeling scheme of size $\log n + O(1)$. Bonichon et al. [5] proved

¹A graph G is *edge-universal* for a class \mathcal{R} of graphs, if every graph in \mathcal{R} appears as a subgraph in G (not necessarily induced).

²A graph G is *induced-universal* for a class \mathcal{R} of graphs, if every graph in \mathcal{R} appears as an induced subgraph in G .

³ $\log^* n$ denotes the iterated logarithm of n .

Table 1: Best known adjacency labeling schemes for graphs with at most n vertices.

Family	Upper bound	Lower bound	Encoding	Decoding	Ref.
Trees	$\log n + O(\log^* n)$	$\log n + \Omega(1)$	$O(n \log^* n)$	$O(\log^* n)$	[3]
Binary trees	$\log n + O(1)$	$\log n + \Omega(1)$	$O(n)$	$O(1)$	[5]
Bd. depth δ trees	$\log n + O(1)$	$\log n + \Omega(1)$	$O(n)$	$O(1)$	[11]
Bd. deg. Δ trees	$\log n + O(\log^* n)$	$\log n + \Omega(1)$	$O(n \log^* n)$	$O(\log^* n)$	[3]
Planar graphs	$2 \log n + O(\log \log n)$	$\log n + \Omega(1)$	$O(n)$	$O(1)$	[13]
Outerplanar graphs	$\log n + O(\log \log n)$	$\log n + \Omega(1)$	$O(n)$	$O(1)$	[13]
Bd. deg. $\Delta(n)$ graphs	$\lceil \frac{\Delta(n)}{2} \rceil \log n + O(\log^* n)$	$\frac{\Delta(n)}{2} \log n$	$O(n)$	$O(\log^* n)$	[3]
Graphs	$\lfloor \frac{1}{2}n \rfloor + \lceil \log n \rceil$	$\lfloor \frac{1}{2}n \rfloor - 1$	$O(n)$	$O(1)$	[19]

Table 2: Our contribution for families of graphs with bounded degree Δ .

Family	Upper bound	Tight	Encoding	Decoding	Ref.
Trees	$\log n + O(1)$	yes	$O(n \log n)$	$O(\log \log n)$	Cor. 10
Outerplanar	$\log n + O(1)$	yes	$O(n \log n)$	$O(\log \log n)$	Thm. 1
Planar ($\Delta \leq 4$)	$\frac{3}{2} \log n + O(\log \log n)$	no	$O(n \log n)$	$O(\log \log n)$	Thm. 18
Graphs, $\Delta(n)$	$\log \binom{n}{\lceil \Delta(n)/2 \rceil} + 2 \log n$	no	$O(n)$	$O(\Delta(n) \log n)$	Thm. 14
Graphs (unbounded)	$\lceil \frac{\Delta}{2} \rceil \log n + O(1)$	no	$O(n \log n)$	$O(\log \log n)$	Cor. 11

that caterpillars and binary trees enjoy a labeling scheme of size $\log(n) + O(1)$ using a method called “Traversal and Jumping”. In a follow up paper, Bonichon et al. [6] claimed without proof that the aforementioned methods can be used to achieve the same bound for bounded degree trees. Chung [9] showed the existence of an induced-universal graph with $O(n)$ vertices for bounded degree trees.

Graphs with maximum degree Δ have *arboricity*⁴ $k(\Delta) = \lceil \Delta/2 \rceil$ [15] thus, by the theorem of Nash-Williams [20], they can be decomposed into $\lceil \Delta/2 \rceil$ forests. Alstrup and Rauhe [3] combined this result with their labeling scheme for forests to obtain a labeling scheme of size $(\lceil \Delta/2 \rceil) \log n + O(\log^* n)$ for bounded degree graphs. They also proved a matching lower bound of $\Omega(k(\Delta) \log n)$.

Butler [7] constructed an induced-universal graph for graphs with maximum degree Δ with $O(n^{\lceil \frac{\Delta+1}{2} \rceil})$ vertices. The author notes that any induced-universal graph must have at least $cn^{\lfloor \Delta/2 \rfloor + 1}$ vertices for some c depending only on Δ , which implies that the bounds are optimal when Δ is even. For odd Δ , Esperet et al. [10] showed a smaller induced-universal graph with $O(n^{\lceil \Delta/2 \rceil - 1/\Delta} \log^{2+2/\Delta} n)$ vertices. It follows that there exists a labeling scheme for $\mathcal{G}(n, \Delta)$ of size $\frac{\Delta}{2} \log n$ bits for even Δ , and $(\lceil \frac{\Delta}{2} \rceil - 1/\Delta) \log n + \log(\log^{2+2/\Delta} n)$ bits for an odd Δ (but that is not necessary efficient). We summarize the best known bounds for adjacency labeling schemes in Table 1.

1.2 Preliminaries

For two integers $0 \leq k_1 \leq k_2$ we denote $[k_1] = \{1, \dots, k_1\}$ and $[k_1, k_2] = \{k_1, \dots, k_2\}$. A binary string x is a member of the set $\{0, 1\}^*$, and we denote its length by $|x|$. We denote the concatenation of two binary strings x, y by $x \circ y$.

For a graph G we denote its set of vertices and edges by $V(G)$ and $E(G)$, respectively. The family of all graphs is denoted \mathcal{G} . For any graph family \mathcal{R} , let $\mathcal{R}(n) \subseteq \mathcal{R}$ denote the subfamily containing the graphs of at most n vertices. The collection of graphs with bounded degree Δ in $\mathcal{R}(n)$ is denoted $\mathcal{R}(n, \Delta)$. The collection of planar graphs, outerplanar graphs, and trees, in $\mathcal{G}(n)$ is denoted $\mathcal{P}(n)$, $\mathcal{O}(n)$ and $\mathcal{T}(n)$, resp. Unless otherwise stated, we assume hereafter Δ to be constant. To simplify the presentation we suppress all dependencies on Δ in all our bounds and running time estimations. All these dependencies can be computed and shown to be at most a multiplicative factor of $O(\Delta \log \Delta)$ times the claimed bounds. We defer the exact details to the journal version of the paper. Non constant bounds on the degree are denoted by $\Delta(n)$. We note that all results work for disconnected graphs. We assume trees to be rooted, and denote $\log_2 n$ as $\log n$. For a set of vertices $S \subset V(G)$ we define $G - S$ to be the graph obtained from G by removing the vertices in S and all incident edges. The set of edges in $G = (V, E)$ incident to a vertex $v \in V$ is denoted E_v .

Let $G = (V, E) \in \mathcal{G}(n)$, and let $u, v \in V$. $adjacency(v, u)$ is the boolean function over vertices in

⁴The *arboricity* of a graph G is the minimum number of edge-disjoint acyclic subgraphs whose union is G .

$G \in \mathcal{G}$ that returns **true** if and only if u and v are adjacent in G . A *label assignment* for $G \in \mathcal{G}$ is a mapping of each $v \in V(G)$ to a bit string $\mathcal{L}(v)$, called the *label* of v . An *adjacency labeling scheme* for \mathcal{G} consists of an encoder and decoder. The *encoder* is an algorithm that receives $G \in \mathcal{G}$ as input and computes the label assignment e_G . The *decoder* is an algorithm that receives any two labels $\mathcal{L}(v), \mathcal{L}(u)$ and computes the query $d(\mathcal{L}(v), \mathcal{L}(u))$, such that $d(\mathcal{L}(v), \mathcal{L}(u)) = \text{adjacency}(v, u)$. The *size* of the labeling scheme is the maximum label length. Hereafter, we refer to adjacency labeling schemes simply as labeling schemes. For the encoding and decoding algorithms, we assume a $\Omega(\log n)$ word size RAM model.

2 $\log n + O(1)$ Labeling Scheme for Bounded-Degree Outerplanar Graphs

In this section we describe a labeling scheme for outerplanar graphs with bounded degree Δ . Our method relies on an embedding technique of Bhatt, Chung, Leighton and Rosenberg [4] for bounded degree outerplanar graphs. In their paper, the authors were concerned with *edge-universal graphs* for various families of bounded degree graphs. In particular, they show that for every $n \in \mathbb{N}$ there exists a graph H_n with $O(n)$ vertices and $O(n)$ edges that contains every bounded degree outerplanar graph $G \in \mathcal{O}(n, \Delta)$ as a subgraph (not necessarily induced).

2.1 Our Methods

Our main tool is an embedding technique due to Bhatt et al. [4] of outerplanar graphs into H_n . On the one hand, the embedding is simple to compute. This fact will lead to an efficient $O(n \log n)$ time encoder. On the other hand, the embedding satisfies a useful locality property. This property allows our labels to contain both unique vertex identifiers of the graph H_n and edge identifiers, without exceeding the desired label size $\log n + O(1)$.

To obtain the latter label size via an embedding into H_n we need to overcome several difficulties. Although H_n has a linear number of edges, its maximum degree is $\Omega(\log n)$, thus, unique edge identifiers require $\Omega(\log \log n)$ bits, in general. Since also $|V(H_n)| = \Omega(n)$, it follows that a label cannot contain an arbitrary combination of vertex identifiers in V_n and edge identifiers at the same time, as it would lead to labels with size $\log n + \Omega(\log \log n)$. This difficulty is overcome by exploiting the structure of H_n further and constructing unique vertex identifiers in a particular way that allows reducing the encoding length. This solution creates an additional difficulty of separating the different parts of the label in the decoding phase. This difficulty is overcome by designing careful encoding lengths that minimize the ambiguity, and storing an additional constant amount of information to eliminate it altogether.

2.2 A Compact Edge-Universal Graph for Bounded-Degree Outerplanar Graphs

We describe next the edge-universal graph $H_n = (V_n, E_n)$ constructed by Bhatt et al. [4] for $\mathcal{O}(n, \Delta)$. We let $k = \min\{s \in \mathbb{N} : 2^s - 1 \geq n\}$ and set $N = 2^k - 1$. The construction uses two constants $c = c(\Delta), g = g(\Delta)$, that depend only on Δ .

The graph H_n is constructed from the complete binary tree T on N vertices as follows. To obtain the vertex set V_n , split every vertex $v \in V(T)$ at level t in T into $\gamma_t = c \log(N/2^t)$ vertices $w_1(v), \dots, w_{\gamma_t}(v)$. The latter set of vertices is called the *cluster* of v . For $w \in V_n$ we denote by $t(w)$ the *level* of w , that is the level in the binary tree T of the cluster to which w belongs.

The edge set E_n is defined as follows. Two vertices $w_i(v), w_j(u) \in V_n$ are adjacent if and only if the clusters they belong to in T are at distance at most g in T . Note that $w_1 w_2 \in E_n$ implies

$|t(w_1) - t(w_2)| \leq g$. This completes the construction of the graph. One can easily check that $|V_n| = O(n)$. $|E_n| = O(n)$ also holds, but we do not this fact directly. The graph H_n is illustrated in Figure 1. Our labeling scheme relies on the following result of Bhatt et al. [4].

Theorem 2 (Bhatt et al. [4]). *H_n is edge-universal for the class of bounded degree outerplanar graphs $\mathcal{O}(n, \Delta)$.*




Figure 1: An illustration of the graph H_n .

2.3 Warm-Up: a $\log n + O(\log \log n)$ Labeling Scheme

We briefly describe a simple $\log n + O(\log \log n)$ labeling scheme. First, assign unique identifiers Id to the vertices in H_n . Since $|V_n| = O(n)$ we can assume that $|Id(v)| = \log n + O(1)$ for every $v \in V_n$. Next, for every $v \in V_n$ assign unique identifiers to the edges incident to v . Since every vertex in H_n has $O(\log n)$ neighbours, each edge can be encoded using $\log \log n + O(1)$ bits.

To obtain labels for a given outerplanar graph $G = (V, E)$, compute first an embedding $\phi : V \rightarrow V_n$ of G into H_n . Next, define the label of vertex $v \in V$ to be the concatenation of $Id(\phi(v))$ with the identifiers of all edges in E_n leading to images under the embedding ϕ of neighbouring vertices, namely all $\phi(u)\phi(v)$ such that $uv \in E$. Since the maximum degree in G is bounded by the constant Δ , this results in a $\log n + O(\log \log n)$ label size. It is not difficult to see that encoding and decoding can be performed efficiently (we elaborate on it later on).

2.4 The Encoder

To reduce the size of the labels to $\log n + O(1)$ we need to refine the latter technique significantly. As a first step we employ *differential sizing*, a technique first used in the context of labeling schemes

in [21]. In differential sizing some parts of the label do not have a fixed number of allocated bits across all labels. Concretely, we use differential sizing for both vertex and edge identifiers.

The resulting labels will have the desired length, but will also contain an undesired ambiguity, that will prohibit correct decoding. We will then append a short prefix to the label that will resolve this ambiguity.

Differential Sizing - The Suffix of a Label.

Let us first formally define our naming scheme for vertex and edge identifiers in H_n .

Definition 3. A *naming* of H_n is an injective function $Id : V_n \rightarrow \mathbb{N}$ and a collection of injective functions $EId_v : E_v \rightarrow \mathbb{N}$ for every $v \in V_n$. A naming is *coherent* if for every $v, v_1, v_2 \in V_n$

1. $Id(v_1) > Id(v_2)$ implies that $t(v_1) > t(v_2)$, or $t(v_1) = t(v_2)$ and the cluster of v_2 appears to the left of the cluster of v_1 in T .
2. $EId_v(vv_1) > EId_v(vv_2)$ implies that $Id(v_1) > Id(v_2)$.

We compute a coherent naming by assigning the identifiers 1 through $|V_n|$ to V_n level by level, traversing the clusters in any single level in T from left to right, and then naming the edges incident to $v \in V_n$ from 1 to $|E_v|$ in a way that is consistent with the vertex naming.

For $v, v' \in V$ define $\alpha(v) := \lceil \log Id(v) \rceil$ and $\beta(v) := \lceil \log EId_v(vv') \rceil$ and let

$$\alpha_t = \max_{v: t(v) \leq t} \alpha(v) \quad \text{and} \quad \beta_t = \max_{vv': t(v) \leq t} \beta(v')$$

be the maximal number of bits required to encode a vertex name and an edge name for vertices with level at most t . The simple labels described in the beginning of this section store $\log n + O(1)$ and $\log \log n + O(1)$ bits for *every* vertex name, and every edge name, respectively. In contrast, our label for a vertex in level t stores α_t bits for a vertex name, and β_t for edge names. In the following lemma we prove that new labels have the desired size.

Lemma 4. *For every $t \leq \log N$ it holds that $\alpha_t \leq t + \lceil \log(\log N - t) \rceil + O(1)$, $\beta_t \leq \lceil \log(\log N - t) \rceil + O(1)$ and $\alpha_t + \Delta\beta_t = \log n + O(1)$.*

We henceforth denote the part of the label containing the vertex name and all its edge identifiers as the *suffix*.

Resolving Ambiguity.

Since the vertex name does not occupy a fixed number of bits across all labels, it is a priori unclear which part of the label contains it. To resolve this ambiguity we analyze the following function, which represents the final length of our labels for vertices in level t (up to a fixed constant). Let $D = \lceil \log N \rceil$ and $f : D \rightarrow \mathbb{N}$ be defined as

$$f(t) = \alpha_t + \Delta\beta_t = t + (\Delta + 1)\lceil \log(\log N - t) \rceil.$$

The following lemma states that all but a constant number (depending on Δ) of values in D have at most two pre-images under f . This observation is useful, since it implies that the knowledge of the level $t(v)$ of the vertex v can resolve all remaining ambiguities in its label, as the vertex name occupies exactly $\alpha_{t(v)}$ bits.

Lemma 5. Let $r(\Delta) = \lceil 8(\Delta+1) \log(\Delta+1) \rceil$. For every $t \in [\lceil \log N \rceil - r(\Delta)]$ the number of integers $t' \in D \setminus \{t\}$ that satisfy $f(t) = f(t')$, is at most one.

Remark 6. It is natural to ask if having equal label lengths for vertices in different levels can be avoided altogether (thus making Lemma 5 unnecessary). This seems not to be the case for the following reason. The number of vertices in every level is at least $\Omega(\log N)$, thus, with label size $\ell = o(\log \log N)$ one can not uniquely represent all vertices in *any* level. Furthermore, the label length is also restricted to $\log n + O(1)$, and the number of levels is $\lceil \log N \rceil$. Thus, a function assigning levels to label lengths would need to have domain $[\lceil \log N \rceil]$ and range $[g(N), \lceil \log N \rceil]$ for $g(N) = \Omega(\log \log N)$, implying that it cannot be one-to-one.

Recall that the length of the suffix of vertex $v \in V$ is exactly $\alpha_{t(v)} + \Delta \beta_{t(v)}$. We next show how the structural property proved here allows to construct a constant size *prefix*, that will eliminate the ambiguity caused by differential sizing.

Constructing the Prefix.

For a formal description of the prefix we need the following definition. We let $r(\Delta) = \lceil 8(\Delta+1) \log(\Delta+1) \rceil$, as in Lemma 5.

Definition 7. A vertex $v \in V$ is called *shallow* if its level $t(v)$ is at most $\lceil \log N \rceil - r(\Delta)$. We call a shallow vertex *early* if $t(v)$ is the smallest pre-image of $f(t(v))$. A shallow vertex that is not early is called *late*.

A vertex $v \in V$ that is not shallow is called *deep*. A deep vertex is of *type* τ , if its level satisfies $t(v) = \lceil \log N \rceil - \tau$.

It is easy to verify the following properties. Lemma 5 guarantees that if v is shallow, then $f(t(v))$ has at most two pre-images under f . If v is shallow and there is only one pre-image for $f(t(v))$, then v is early. Finally, observe that the type of deep vertices ranges in the interval $[1, r(\Delta)]$.

We are now ready to define the prefix of a label $\mathcal{L}(v)$ for a vertex $v \in V$. Every prefix starts with a single bit $D(v)$ that is set to 0 if v is shallow, and to 1 if v is deep. The second bit $R(v)$ in every prefix indicates whether a shallow vertex is early, in which case it is set to 0, or late, in which case it is set to 1. The bit $R(v)$ is always set to 0 in labels of deep vertices. The next part $Type(v)$ of the prefix contains $\lceil \log r(\Delta) \rceil$ bits representing the type of the vertex v , in case v is deep. If v is shallow this field is set to zero. This concludes the definition of the prefix. Observe that the prefix contains $O(\log \Delta) = O(1)$ bits. We stress that length s_p of the prefix is identical across all labels.

It is evident that the prefix of a label eliminates any remaining ambiguity. This follows from the fact that the level $t(v)$ of a vertex $v \in V$ can be computed from the length of the suffix and the additional information stored in the prefix. The level, in turn, allows to decompose the suffix into the vertex name and the incident edge names, which can then be used for decoding. We elaborate on the decoding algorithm later on.

The Final Labels.

The final label is obtained by concatenating the suffix to the prefix, namely $\mathcal{L}(v)$ is defined as follows.

$$\mathcal{L}(v) = \underbrace{D(v) \circ R(v) \circ Type(v)}_{\text{prefix}} \circ \underbrace{Id(\phi(v)) \circ EId_{\phi(v)}(e_1) \circ \dots \circ EId_{\phi(v)}(e_\Delta)}_{\text{suffix}}.$$

Figure ?? illustrates the label structure as a function of the level of the vertex. Note that $|\mathcal{L}(v)| = s_p + f(t(v))$, thus the level $t(v)$ of v determines $|\mathcal{L}(v)|$. Note that Lemma 4 and the fact that the prefix has constant size guarantees that $|\mathcal{L}(v)| = \log n + O(1)$, as desired. We also pad each label with sufficiently many 0's and a single '1', to arrive at a uniform length. The latter simple modification allows the decoder to work without knowing n in advance (see [12] for details).




Figure 2: The composition of labels in our labeling scheme for vertices in different levels $t = 1, \dots, \log N$. The size of the prefix is seen to be constant in every level, while the fields of lengths α_t and β_t , used to store vertex and edge identifiers, respectively, have variable sizes. The levels i and j comprise a collision with respect to the function f , thus labels of vertices in these levels have the same length.

Although it is not necessary for the correctness of our labeling scheme, we prove here uniqueness of the labels. In other words, we show that two different vertices in G necessarily get different labels.

Lemma 8. *For every two distinct $u, v \in V(G)$ it holds that $\mathcal{L}(u) \neq \mathcal{L}(v)$.*

2.5 Decoding

Consider two labels $\mathcal{L}(u)$ and $\mathcal{L}(v)$ for vertices $u, v \in V$. The decoder first extracts the levels $t(u)$ and $t(v)$ of u and v respectively, using the following simple procedure, which we describe for v . If $D(v) = 0$, v is shallow. To this end, the decoder computes all pre-images of the length of the suffix, $|\mathcal{L}(v)| - s_p$, under f . Recall, that by Lemma 5, the number of pre-images is at most two. Let $t_1 \leq t_2$ be the computed pre-images. Next, the decoder inspects the bit $R(v)$. According to the definition of the labels, $t(v) = t_1$ if $R(v) = 0$, and $t(v) = t_2$, otherwise. Consider next the case $D(v) = 1$, namely that v is deep. In this case, the decoder inspects $Type(v)$. The level of v is $t(v) = \lceil \log N \rceil - Type(v)$, by definition of a type of a deep vertex. It is obvious by the definition of the labels that the decoder extracts $t(v)$ correctly. Assume next that the decoder extracted $t(u)$ and $t(v)$. The decoder can now extract $Id(\phi(u))$ and $Id(\phi(v))$, by inspecting the first $\alpha_{t(u)}$ and $\alpha_{t(v)}$ bits of the suffix of $\mathcal{L}(u)$ and $\mathcal{L}(v)$, respectively. Next, the decoder checks if $\phi(u)\phi(v) \notin E_n$,

in which case it reports **false**. Finally, if $\phi(u)\phi(v) \in E_n$ the decoder scans all Δ blocks of $\beta_{t(u)}$ bits each, succeeding $Id(\phi(u))$ in the suffix of $\mathcal{L}(u)$, checking if one of them contains the edge-identifier $EId_{\phi(u)}(\phi(u)\phi(v))$. If this identifier is found the decoder reports **true**. Otherwise, it reports **false**. The correctness of the decoding is clear from the label definition and Lemma 5.

Lemma 9. *The decoding of the labels can be performed in time $O(\log \log n)$.*

2.6 Computing the Embedding ϕ

All the labels can clearly be computed from the graph G , the embedding ϕ and the graph H_n in time $O(n \log n)$. It is also straightforward to compute H_n in $O(n)$ time. It remains to discuss how to compute an embedding ϕ , for which we provide a high-level overview. For a detailed description, see Bhatt et al. [4].

The algorithm uses a subroutine for computing bisectors of a graph. A *bisector* of a graph $G = (V, E)$ is a set $S \subset V$ of vertices with the property that the connected components of the graph $G - S$ can be partitioned into two parts, such that the sum of the vertices in each part is the same, and no edge connects two vertices in different parts. If S is a bisector in G we say that S *bisects* $V \setminus S$.

Given a k -coloring $V = V_1 \cup \dots \cup V_k$ of V (for some $k \in \mathbb{N}$), one can define a k -*bisector* of G as a set $S \subset V$ that bisects *every* color class, namely one that bisects $V_i \setminus S$ for all $i \in [k]$. An important property of outerplanar graphs is that they admit $O(\log n)$ size k -bisectors, for every fixed k .

The algorithm works by assigning vertices in the graph G to clusters in T . The root of T is assigned up to $c \log n$ vertices that form a bisector of G with parts G_1, G_2 . In the next iteration, vertices adjacent to vertices in the bisector are given a new color. Next, two new 2-bisectors are found, one in each part G_1, G_2 , and they are assigned to the corresponding decedents of the root of T .

Let $k(\Delta) = \log \Delta + 1$. In general, the vertices stored at a vertex of T at level t correspond to a $k(\Delta)$ -bisector. The colors of this bisector correspond to the neighbors of vertex-sets stored at $k(\Delta) - 1$ nearest ancestors of the current vertex in T . The last color is reserved to the remaining vertices. Also stored in this vertex are all neighbors of the vertex-set stored in the ancestor of the current vertex at distance exactly $k(\Delta)$, that were not yet assigned to some other cluster. We refer the reader to [4] for an analysis of the sizes of clusters.

Let $T(n)$ be the running time of the latter algorithm in a graph with n vertices. $T(n)$ clearly satisfies $T(n) \leq 2T(n/2) + O(h(n))$, where $h(n)$ is the complexity of finding an $O(1)$ -bisector of $O(\log n)$ size in an n -vertex graph. For outerplanar graphs the latter can be done in linear time [8, 4], thus the labels of our labeling scheme can be computed in $O(n \log n)$ time.

2.7 Improvements and Special Cases

Several well-known techniques can be easily applied on top of our construction to reduce the additive constant in the label size. First, since graphs of maximum degree Δ have arboricity $\lfloor \frac{\Delta}{2} \rfloor + 1$, one can reduce the number of edge identifiers stored in each label to the latter number (see Kannan et al. [15]). We later show a simpler procedure that works for bounded degree graphs.

Finally, for bounded-degree trees $\mathcal{T}(n, \Delta)$, it suffices to store a single edge identifier (corresponding to the edge connecting a vertex to its parent in G). We summarize this result in the following corollary of Theorem 1.

Corollary 10. *For every fixed $\Delta \geq 1$, the class $\mathcal{T}(n, \Delta)$ of bounded-degree trees admits a labeling scheme of size $\log n + O(1)$, with encoding complexity $O(n \log n)$ and decoding complexity $O(\log \log n)$.*

3 Labeling schemes for $\mathcal{G}(n, \Delta)$ and $\mathcal{G}(n, \Delta(n))$

First we note that Theorem 1 implies almost directly a $\lceil \Delta/2 \rceil \log n$ labeling scheme for graphs of fixed bounded degree Δ . The result follows from the technique of Alstrup and Rauhe [3], Lemma 8 and the fact that any subtree of a bounded degree graph has bounded degree.

Corollary 11. *For every $\Delta \geq 1$, the class $\mathcal{G}(n, \Delta)$ of bounded-degree graphs admits labeling schemes of size $\lceil \Delta/2 \rceil \log n$, with encoding complexity $O(n \log n)$ and decoding complexity $O(\log \log n)$.*

From here on, we discuss labeling schemes for graphs of non-constant bounded degree $k = \Delta(n)$. Adjacency relation between any two vertices may be reported in only one of the labels representing them. For bounded degree graphs, the method of Kannan et al. [15] of decomposition into forests can be replaced with a simpler procedure, using Eulerian circuits, as we prove in the following.

Lemma 12. *Let $G = (V, E)$ be a graph with degree bounded by k . There exist an adjacency labeling scheme for G of size $(\lceil \frac{k}{2} \rceil + 1) \log n$.*

The current best labeling schemes for graphs works in two modes, according to the range of k . If $k \leq n/\log n$, a $k/2 \log n$ labeling scheme can be achieved [15], essentially by encoding an adjacency list. For larger k , labels defined through the adjacency matrix of the graph have size $n/2 + \log n$ [19]. Our improved labels use the well-known *combinatorial number system* (see e.g. [16]).

Lemma 13. *Let $L = \binom{n}{k}$. There is a bijective mapping $\sigma : S_k \rightarrow [0, L - 1]$ between the set of strictly increasing sequences S_k of the form $0 \leq t_1 < t_2 < \dots < t_k < n$ and $[0, L - 1]$ given by*

$$\sigma(t_1, \dots, t_k) = \sum_{i=1}^k \binom{t_i}{i}.$$

We use Lemma 13 to prove the following theorem. For this purpose we assume that the labeling scheme presented in Lemma 12 returns a subset of vertices for every vertex $v \in V$ according to the partition instead of a final label.

Theorem 14. *For $1 \leq k \leq n$, there exist an adjacency labeling scheme for $\mathcal{G}(n, k)$ with size: $\binom{n}{\lceil k/2 \rceil} + \lceil \log n \rceil + \lceil \log k \rceil$.*

The labeling scheme suggested in Theorem 14 implies a label size of approximately $(k + 2) \log n$ bits, when k is small and $\Theta(n)$ when $k = \Theta(n)$. The following lemma identifies the range of k for which our labeling scheme improves on the best known bounds.

Lemma 15. *For $(e+1)\sqrt{n} \leq k \leq \frac{n}{5}$, and $f(n, k) = \binom{n}{\lceil k/2 \rceil} + \log k + \log n$ it holds that a) $f(n, k) < \frac{n}{2}$; and b) $f(n, k) < \lceil k/2 \rceil + 2 \log(n)$.*

We conclude from Lemma 15 that our labeling scheme is preferable to [19, 7] for graphs of bounded degree k for $(e + 1)\sqrt{n} \leq k \leq \frac{n}{5}$.

References

- [1] Alstrup, S., Bille, P., Rauhe, T.: Labeling schemes for small distances in trees. SIAM J. Disc. Math. 19(2), 448–462 (2005)

- [2] Alstrup, S., Gavoille, C., Kaplan, H., Rauhe, T.: Nearest common ancestors: A survey and a new distributed algorithm. In: Proc. of the fourteenth annual ACM symposium on Parallel algorithms and architectures. pp. 258–264. SPAA '02 (2002)
- [3] Alstrup, S., Rauhe, T.: Small induced-universal graphs and compact implicit graph representations. In: Proc. The 43rd Annual IEEE Symposium on Foundations of Computer Science. pp. 53–62. IEEE (2002)
- [4] Bhatt, S., Chung, F.R.K., Leighton, T., Rosenberg, A.: Universal Graphs for Bounded-Degree Trees and Planar Graphs. *SIAM J. Disc. Math.* 2(2), 145–155 (1989)
- [5] Bonichon, N., Gavoille, C., Labourel, A.: Short labels by traversal and jumping. In: Structural Information and Communication Complexity, pp. 143–156 (2006)
- [6] Bonichon, N., Gavoille, C., Labourel, A.: Short labels by traversal and jumping. *Electr. Notes in Discrete Math.* 28, 153–160 (2007)
- [7] Butler, S.: Induced-universal graphs for graphs with bounded maximum degree. *Graphs Combinator.* 25(4), 461–468 (2009)
- [8] Chung, F.R.K.: Separator theorems and their applications. *Forschungsinst. für Diskrete Mathematik* (1989)
- [9] Chung, F.R.K.: Universal graphs and induced-universal graphs. *J. Graph Theor.* 14(4), 443–454 (1990)
- [10] Esperet, L., Labourel, A., Ochem, P.: On induced-universal graphs for the class of bounded-degree graphs. *Inform. Process. Lett.* 108(5), 255–260 (2008)
- [11] Fraigniaud, P., Korman, A.: Compact ancestry labeling schemes for xml trees. In: Proc. of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 458–466. SODA '10 (2010)
- [12] Fraigniaud, P., Korman, A.: Compact ancestry labeling schemes for xml trees. In: In Proc. 21st ACM-SIAM Symp. on Discrete Algorithms. SODA '10 (2010)
- [13] Gavoille, C., Labourel, A.: Shorter implicit representation for planar graphs and bounded treewidth graphs. In: Algorithms–ESA 2007, pp. 582–593 (2007)
- [14] Gavoille, C., Paul, C.: Distance labeling scheme and split decomposition. *Discrete Math.* 273, 115 – 130 (2003)
- [15] Kannan, S., Naor, M., Rudich, S.: Implicit representation of graphs. In: Proc. of the 20th ACM symposium on Theory of computing. pp. 334–343. STOC '88 (1988)
- [16] Knuth, D.: Combinatorial algorithms. the art of computer programming, vol. 4a (2011)
- [17] Korman, A., Peleg, D., Rodeh, Y.: Constructing labeling schemes through universal matrices. In: Algorithms and Computation, pp. 409–418 (2006)
- [18] Laoutaris, N., Rajaraman, R., Sundaram, R., Teng, S.H.: A bounded-degree network formation game. arXiv preprint cs/0701071 (2007)
- [19] Moon, J.W.: On minimal n-universal graphs. In: Proc. of the Glasgow Math. Assoc. vol. 7, pp. 32–33 (1965)

- [20] Nash-Williams, C.: Edge-disjoint spanning trees of finite graphs. J. London Math. Soc. 1(1), 445–450 (1961)
- [21] Thorup, M., Zwick, U.: Compact routing schemes. In: Proc. of the thirteenth annual ACM symposium on Parallel algorithms and architectures. pp. 1–10. SPAA '01, New York (2001)
- [22] Wang, Y., Li, X.Y.: Localized construction of bounded degree and planar spanner for wireless ad hoc networks. Mobile Networks and Applications 11(2), 161–175 (2006)

A Missing Proofs from Section 2

A.1 Proof of Lemma 4

We start with the following simple observation stating a couple of facts about the size of different parts of H_n . We let $V_n^t \subset V_n$ denote the set of vertices in H_n with level at most t . This properties will provide bounds on α_t and β_t .

Property 16. *The following properties hold for every level t and every vertex v in level t .*

- i. $|V_n^t| \leq c2^t(\log N - t + 1)$.
- ii. $|\Gamma(v)| \leq 6c\Delta^2(\log N - t + g(\Delta))$.

Proof. i. Note that the number of vertices in level i is $c2^{i-1}(\log N - i)$, thus we obtain

$$|V_n^t| = \sum_{i=1}^t 2^{i-1} c(\log N - i) \leq c2^t \log N - c \sum_{i=1}^t 2^{i-1} i.$$

Using $\sum_{i=1}^t 2^{i-1} i = t2^{t+1} - (t+1)2^t + 1 \geq 2^t(t-1)$ and substituting in the expression above we obtain the desired result.

ii. We notice that, by definition of H_n , the set of neighbors of v are exactly all vertices whose cluster is at distance at most $g(\Delta)$ from the cluster of v . It follows that every such cluster has at most $c \log(N/2^{t-g(\Delta)}) = O(\log N - t + g(\Delta))$ vertices. Finally, notice that the number of clusters that are at distance at most $g(\Delta)$ away from a given cluster is at most $3 \cdot 2^{g(\Delta)-1} = 6\Delta^2$. Putting things together we obtain a bound of $6\Delta^2(\log N - t + g(\Delta))$, as desired. \square

We can now use Property 16 to prove the bounds

$$\alpha_t \leq t + \lceil \log(\log N - t) \rceil + \log c + 2.$$

and

$$\beta_t \leq \lceil \log(\log N - t) \rceil + 2 \log \Delta + \log c + 4.$$

We turn to proving the bound on $\alpha_t + \Delta\beta_t$ for an arbitrary $t \leq \log N$. We drop the additive terms $\log c + 2$ and $2 \log \Delta + \log c + 4$ as they do not depend on n , and hence only contribute a constant additive term. Substituting the bounds above for α_t and β_t and defining $Q = \log N - t(v)$ we obtain

$$\alpha_t + \Delta\beta_t = t + (\Delta + 1) \log(\log N - t) = \log N - Q + (\Delta + 1) \log \log Q.$$

Using the fact that $(\Delta + 1) \log \log Q - Q \leq 0$ whenever $Q \geq (\Delta + 1) \log(\Delta + 1)$ and $(\Delta + 1) \log \log Q - Q \leq (\Delta + 1) \log \log(\Delta + 1)$ whenever $Q < (\Delta + 1) \log(\Delta + 1)$, we have $\alpha_t + \Delta\beta_t \leq \log N + (\Delta + 1) \log \log(\Delta + 1) = \log n + O(1)$. This concludes the proof of the lemma.

A.2 Proof of Lemma 5

Consider first the behavior of the term $h(t) = (\Delta + 1)\lceil \log(\log N - t) \rceil$ in the definition of f . Let $t_1 \in [D]$ be the smallest integer such that $\log(\log N - t_1)$ is integer. It follows that $\log N - t_1 = 2^s$ for some integer s . Consider the smallest integer $t_2 \in D$ with $t_2 > t_1$ such that $\log(\log N - t_2)$ is integer. This value clearly satisfies $\log N - t_2 = 2^{s-1}$. The smallest integer $t_3 \in D$ with $t_3 > t_2$ such that $\log(\log N - t_3)$ is integer is $\log N - t_3 = 2^{s-2}$ and so on. Let $t_1 < t_2 < \dots < t_p$ be the obtained sequence of $p = O(\log \log N)$ points that satisfies $t_i = \log N - 2^{s-i+1}$ for $i \in [p]$. We call such points *jumps*. See Figure 3 for an illustration of f . Observe next that for every $k \in [p-1]$ the term $h(t)$ is constant within the interval $I_k = [t_k + 1, t_{k+1}]$, while at the point $t_{k+1} + 1$ it decreases by $\Delta + 1$.

We rewrite f as

$$f(t) = t - \sum_{k: t_k \leq t-1} (\Delta + 1) + (\Delta + 1)\lceil \log(\log N - 1) \rceil,$$

The last term does not depend on t , thus it suffices to prove the claim for $\hat{f}(t) = t - \sum_{k: t_k \leq t-1} (\Delta + 1)$. Consider next some $t \in D$ and assume that $\hat{f}(r) \neq \hat{f}(t)$ for every $r < t$. Define $K = \lceil \log N \rceil - t$, and let $t_i \geq t$ be the smallest jump larger or equal to t . Define $\gamma = t_i - t$. Observe that the number of jumps within the interval $[t_i, \lceil \log N \rceil]$, whose length is $K - \gamma$ is at most $\log(K - \gamma) + 1$, while the next jump occurring after t_i appears exactly in the middle of the latter interval, namely $t_{i+1} - t_i = \frac{K - \gamma}{2}$. To this end assume that $\hat{f}(t)$ has at least two other pre-images $t'' > t' > t$ under \hat{f} . Since \hat{f} is monotonic between jumps, there must be at least one jump between every consecutive pair of pre-images of \hat{f} . It follows that $t \leq t_i \leq t'$ and $t' \leq t_j \leq t''$ for some $j \geq i + 1$. In particular, it follows that $[t_i, t_{i+1}] \subset [t_i, t'']$, which implies $t'' - t_i \geq \frac{K - \gamma}{2}$. Now, using the fact that within the interval $[t_i, \lceil \log N \rceil] \supset [t_i, t'']$ there are at most $\log(K - \gamma)$ jumps, we can write

$$\begin{aligned} \hat{f}(t'') &= t'' - \sum_{k: t_k \leq t''-1} (\Delta + 1) \\ &\geq t - \sum_{k: t_k \leq t-1} (\Delta + 1) + (t'' - t) - \sum_{k: t_k \in [t_i, t''-1]} (\Delta + 1) \\ &\geq \hat{f}(t) + (\gamma + \frac{K - \gamma}{2}) - (\Delta + 1)(\log(K - \gamma) + 1) \\ &\geq \hat{f}(t) + \frac{K}{2} - (\Delta + 1)(\log K + 1), \end{aligned}$$

which, using $\hat{f}(t) = \hat{f}(t'')$, implies

$$\frac{K}{2} - (\Delta + 1)(\log K + 1) \leq 0.$$

Rearranging the terms we obtain $\frac{K}{2(\log K + 1)} \leq \Delta + 1$, which clearly implies $K \leq r(\Delta)$ and $t \in [\lceil \log N \rceil - r(\Delta)]$. This concludes the proof.

A.3 Proof of Lemma 8

Consider two vertices $u, v \in V(G)$ and assume $\mathcal{L}(u) = \mathcal{L}(v)$. Since the prefix of every label has the same length and the same parts for every vertex in $V(G)$, it must hold that $D(u) = D(v)$, $R(u) = R(v)$ and $Type(u) = Type(v)$.




Figure 3: An illustration of Lemma 5

Assume first that both u and v are shallow. In this case $\mathcal{L}(u) = \mathcal{L}(v)$ implies that $f(t(u)) = f(t(v))$. Now, since $R(u) = R(v)$, both v and u are either early, or late, implying by Lemma 5 that $t(u) = t(v) =: t$. It follows by definition of the labels, that the first α_t bits of the suffices of $\mathcal{L}(u)$ and $\mathcal{L}(v)$ contain $Id(\phi(u))$ and $Id(\phi(v))$, respectively, implying $Id(\phi(u)) = Id(\phi(v))$. From the correctness of the embedding ϕ and uniqueness of the identifiers Id it follows that $u = v$.

Finally, assume that both u and v are deep. In this case $Type(u) = Type(v)$ implies again that $t(u) = t(v)$. We now reach the conclusion $u = v$ as in the previous case.

A.4 Proof of Lemma 9

All operations performed by the decoder take $O(1)$ time except for the computation of pre-images of an integer under f , and steps that inspect properties in H_n . The former can obviously be performed in $O(\log \log n)$ time as follows. If the label corresponds to a deep label, then the decoder checks all $r(\Delta) = O(1)$ possible preimages, so assume next that the vertex is shallow. Notice that $t - f(t) = O(\log \log n)$ for every $t \in D$. Thus, given $f(t)$, one can compute in $O(\log \log n)$ time all $f(y)$ for all integers y with $|y - f(t)| = O(\log \log n)$. All preimages of $f(t)$ under f will be found this way.

We focus hereafter on the complexity of deciding if $\phi(u)\phi(v) \in E_n$ and the computation of the edge-identifier corresponding to two vertices in H_n . To this end we need to elaborate on the way labeling is performed for the graph H_n . Recall that H_n is built from the complete binary tree T with $\log N$ levels by splitting each vertex at level t into $\gamma_t = c \log(N/2^t)$ vertices, thus forming clusters, and then connecting two vertices if they belong to clusters at distance at most $g(\Delta) = 2 \log \Delta + 2$ apart in T . The labeling of vertices is performed by assigning a *range* R_C of size γ_t to every cluster at level t in a single breadth-first search traversal, and then ordering the vertices in the cluster arbitrarily using unique values in the range R_C . More precisely, the ranges are constructed starting from the root of T , processing T level by level as follows. The Range of the root cluster is simply $[1, \gamma_1]$. The range of the left child of the root is assigned the range $[\gamma_1 + 1, \gamma_1 + \gamma_2]$, while the right child is assigned the range $[\gamma_1 + \gamma_2 + 1, \gamma_1 + 2\gamma_2]$ and so on. At level t , the assignment of ranges starts from the left descendant of the parent at level $t - 1$ with the lowest range (the range with the smallest lower bound), followed by the right one, and so on. This process defines an ordering

of the clusters, in each level. We denote by $POS(C)$ the position of the cluster C in this order.

Given a label L of a vertex $w \in V_n$ at level $t = t(w)$, its clusters range (which identifies the cluster) can be computed as follows. By definition of the labeling, we have that L is contained in the range $I(w) = [A + 1, A + \gamma_t]$ with

$$A = \sum_{i=1}^{t-1} 2^{i-1} \gamma_i + (POS(C(w)) - 1) \gamma_t,$$

where $C(w)$ is the cluster to which w belongs. We now notice that

$$\sum_{i=1}^{t-1} 2^{i-1} \gamma_i = \sum_{i=1}^{t-1} 2^{i-1} (\log n - i) = c \log n \sum_{i=1}^{t-1} 2^{i-1} - \frac{c}{2} \sum_{i=1}^{t-1} 2^i,$$

where both sums in the last formula have simple closed form expressions. It follows that $POS(C(w))$ can be easily computed from t and L in $O(1)$ time.

Finally, the decoder can compute both the parent and the descendants of a cluster C at level t whose range is $[A + 1, A + \gamma_t]$. The parent is computed using the fact that it is at level $t - 1$ and its position in this level is $\lceil POS(C)/2 \rceil$. The descendants can be computed using the fact that their level is $t + 1$ and their positions in this level are $2POS(C) - 1$ and $2POS(C)$. Both computations can clearly be performed in $O(1)$ time, once presented with the values γ_t . We conclude that it is possible to obtain in time $O(1)$ the ranges of all neighbours of a given cluster C in T .

We turn to the problem of deciding whether $\phi(u)\phi(v) \in E_n$. The decoder starts by computing the clusters C_u and C_v of u and v , respectively in time $O(\log n \log n)$. Then the decoder computes paths P_u and P_v of length at most $g(\Delta)$ from C_u and C_v , respectively in the tree T . Each path is obtained by successively moving at most $g(\Delta)$ steps from a cluster to its parent in T . The computation of these paths takes $O(1)$ time, using the aforementioned method for computing a parent of a cluster. By definition of H_n , we have that $\phi(u)\phi(v) \in E_n$ if and only if the clusters of $\phi(u)$ and $\phi(v)$ are at distance at most $g(\Delta)$ from one another in T . This can be now easily tested by inspecting the paths P_u, P_v , which must contain the least common ancestor (LCA) of $\phi(u)$ and $\phi(v)$, in case they are neighbours in H_n (see Fig. 1 for an illustration).

With a similar argument one shows that each edge identifier can be decoded in $O(1)$ time. The details are similar, and thus omitted. Performing the decoding to all Δ edge identifiers gives a total running time of $O(1)$. We conclude that the decoder can be implemented to have running time $O(\log \log n)$.

A.5 Proof of Lemma 12

We first create an Eulerian multigraph $G' = (V, E')$ from G by adding at most $|V|/2$ new edges, comprising the a matching that connects pairs of vertices in G with odd degree. The degrees of all vertices in G' are even, hence it is Eulerian. We proceed by finding an Eulerian circuit P in G' and directing every edge according to the direction along P . Every vertex in G' with degree d has now exactly $\frac{d}{2}$ incoming and outgoing edges. The label of a vertex will only correspond to the vertices which are adjacent through outgoing edges in P . This number is at most $\frac{d}{2}$. The identifiers of the edges in M are not included in the resulting label. The labeling scheme obtained will assign unique labels to each of the vertices, and concatenate the selected labels to each vertex. since $d \leq k$ this yields a label of size $(\lceil \frac{k}{2} \rceil + 1) \log n$.

A.6 Proof of Theorem 14

We assume the vertices of the graph $G = (V, E)$ to be numbered from 0 to $n - 1$. We call a binary vector C of length n an *adjacency vector* of a vertex $v \in V$ if it satisfies that $C_i = 1$ if and only if v_i is a neighbor of v in G . We denote by \bar{C} the binary vector with $\bar{C}_i = 1$ if and only if $C_i = 0$. We interpret the vectors C and \bar{C} and sets $S \subset V$ as sequences of integers in the range $[0, n - 1]$, as in Lemma 13. Let $v \in V$ be a vertex with $k' \leq k$ neighbors, and let C be its adjacency vector. Our labeling scheme assigns every $v \in V$ to an appropriate $\mathcal{L}(v)$ as follows.

We first use the encoder from the labeling scheme described in Lemma 12 to obtain a temporary subset S_v of neighbors of v . If $k' < \frac{2n}{3}$ we set the last bit of $\mathcal{L}(v)$ to 0, and append it to the number mapped to the sequence of integers corresponding to S_v under the bijection in Lemma 13. According to Lemma 12, we are assured that $|S_v| \leq \lceil \frac{k}{2} \rceil \leq \lceil \frac{n}{3} \rceil$. If $k' \geq \frac{2n}{3}$, we set the last bit of $\mathcal{L}(v)$ to 1, and append to it the number corresponding to \bar{C} under the bijection in Lemma 13. Since $k' > \frac{2n}{3}$, the number of 1's in \bar{C} is at most $\frac{n}{3}$. Finally, we attach to every label a binary representation of k' and a unique vertex identifier using exactly $\log n$ bits each. The encoder performs the operation in polynomial time. It is straightforward to verify the claimed label length.

The decoder receives $\mathcal{L}(u)$ and $\mathcal{L}(v)$ and extracts the corresponding vectors C_u and C_v , the adjacency vectors of u and v , resp., using Lemma 13, and possibly a bit inversion operation. The decoder returns **true** if and only if either $C_u(v) = 1$ or $C_v(u) = 1$. We note that the decoding can be implemented in $O(k \log \binom{n}{k})$ time.

A.7 Proof of Lemma 15

Stirling's approximation yields the following asymptotic approximation.

$$\log \binom{cn}{n} \sim \log \left(\frac{c^c}{(c-1)^{c-1}} \right) \cdot n.$$

Accordingly, $\log \binom{n}{\frac{1}{3}n} \sim \frac{2.75}{3}n$. Since $\log(\frac{10^{10}}{9^9})/10 < 0.5$, and since the function $\binom{n}{k}$ is increasing for $\sqrt{n} \leq k \leq \frac{n}{3}$ our labeling scheme will incur strictly less than $0.5n$ label size for the proposed labeling scheme over the range specified.

Let x_n be defined by $x_n = \log(n!) - (n \log n - n + 1/2(\log 2\pi n))$, and by its definition $x_n \rightarrow 0$ as $n \rightarrow \infty$. In addition:

$$\log \binom{k}{n} = \log(n!) - \log((n-k)!) - \log(k!)$$

$$= (x_n - x_{n-k} - x_k) + n \log n - (n-k) \log(n-k) - k(\log k) + 1/2 \log(n/k(n-k)2\pi).$$

Define $f(n, k)$ by $f(n, k) = k/2 \log n - \log \binom{k}{n}$. We are now interested in the smallest k for which $k_n = f(n, k) \leq 0$ for fixed n . $f(n, \lfloor \sqrt{n} \rfloor) < 0$, so $k_n > \sqrt{n}$, so we can assume $\sqrt{n} \leq k \leq n/2$, which implies that $1/2 \log(n/(k(n-k)2\pi)) = O(\log n)$. Furthermore $x_n - x_{n-k} - x_k = O(1)$ so

$$f(n, k) = k/2 \log n - n \log n + (n-k) \log(n-k) + k \log k + O(\log n)$$

$$= k/2 \log n + n \log((n-k)/n) + k \log(k/(n-k)) + O(\log n).$$

Note that as $n \log((n-k)/n) = O(k)$ this means that

$$\frac{f(n, k)}{k} = \frac{1}{2} \log n - \log((n-k)/k) + O(1) = \log \left(\frac{k\sqrt{n}}{n-k} \right) + O(1)$$

This means that there exists some constant $c > 0$ such that if $\log\left(\frac{k\sqrt{n}}{n-k}\right) \geq c$ then $f(n, k) \geq 0$. But

$$\log\left(\frac{k\sqrt{n}}{n-k}\right) \geq c \iff k \geq c \frac{n}{\sqrt{n} + c}.$$

If $k > c\sqrt{n}$ then $f(n, k) \geq 0$. We can conclude $k_n \leq c\sqrt{n}$, and we may assume that $c\sqrt{n} \geq k \geq \sqrt{n}$.

Now we can conclude that $n \log((n-k)/n) = -k + O(k^2/n) = -k + O(1)$. Hence

$$\frac{f(n, k)}{k} = \frac{1}{2} \log n - 1 + k \log(k/(n-k)) + O\left(\frac{\log n}{\sqrt{n}}\right),$$

and thus

$$\frac{f(n, k)}{k} = \log\left(\frac{k\sqrt{n}}{e(n-k)}\right) + O\left(\frac{\log n}{\sqrt{n}}\right).$$

We can now conclude with $c_n = \exp(O(\frac{\log n}{\sqrt{n}}))$

$$k_n = c_n e \frac{n}{\sqrt{n} + c_n e} + O(1) = e\sqrt{n} + O(\log n).$$

B Labeling Schemes for Bounded-degree Planar Graphs

We present here our labeling scheme for the class of bounded degree planar graphs $\mathcal{P}(n, \Delta)$. Again, we rely on an embedding of the given planar graph G into a graph $PL_n = (W_n, A_n)$, obtained from the complete binary tree T . The only difference between PL_n and H_n is that in PL_n , every vertex v in level t of T is divided into a cluster of $\delta_t = c' \sqrt{\frac{N}{2^t}}$ vertices $z_1(v), \dots, z_{\delta_t}(v)$ (instead of γ_t in H_n), for some constant c' . As for H_n , two vertices in W_n are connected if and only if the distance between their clusters in T is at most $\hat{g}(\Delta) = O(\log \Delta)$. Bhatt et al. [4] showed the following.

Theorem 17 (Bhatt et al. [4]). *PL_n is edge-universal for the class of bounded degree planar graphs $\mathcal{P}(n, \Delta)$. Furthermore, PL_n has $O(n)$ vertices and $O(n \log n)$ edges.*

We use the latter theorem to prove the following result, that is proved in Appendix B.

Theorem 18. *For every fixed $\Delta \geq 1$, the class $\mathcal{P}(n, \Delta)$ of bounded-degree planar graphs admits a labeling scheme of size $\frac{[\Delta/2]+1}{2} \log n + O(\log \log n)$, with encoding complexity $O(n \log n)$ and decoding complexity $O(\log n \log n)$. The average label size of the latter scheme is $(1+o(1)) \log n + O(\log \log n)$.*

Proof. We use a simplified version of the method used in Section 2, thus many details are omitted. Our labeling scheme for planar graphs works with an embedding $\psi : V \rightarrow W_n$ of the given graph G into PL_n . Levels of vertices $w \in W_n$ and $u \in V$ are defined as before. Analogously to our previous scheme, we assign unique identifiers $Id(v)$ to vertices $v \in W_n$ level by level, starting from the highest level, and use them to define edge identifiers. Let $W_n^t \subset W_n$ denote the set of vertices in PL_n with level at most t . Then we have the following lemma.

Lemma 19. *The following properties hold for every level t and every vertex v in level t .*

- i. $|W_n^t| = O(\sqrt{N} 2^t)$.
- ii. $|\Gamma(v)| = O(\sqrt{N} 2^{-t})$.

Proof. *i.* Note that the number of vertices in level i is $O(2^i \sqrt{\frac{N}{2^i}})$, thus for some constant d ,

$$|W_n^t| = c' \sum_{i=1}^t 2^{i-1} \sqrt{\frac{N}{2^i}} \leq c' \sqrt{N} \sum_{i=1}^t \sqrt{2^{i-1}} = O(\sqrt{N 2^t}).$$

ii. Since v is connected to vertices in clusters at distance at most $O(\log \Delta) = O(1)$ from its cluster, there are $O(1)$ such clusters. Furthermore, the highest level of such a cluster is at least $t - O(1)$, thus every such cluster contains $O(\delta_t) = O(\sqrt{N 2^{-t}})$ vertices, which concludes the proof. \square \square

Lemma 19 implies that a the encoding length of label containing a vertex identifier $Id(\psi(v))$ together with $s \in \mathbb{N}$ edge identifiers corresponding to $\psi(v)$ is

$$\log(\sqrt{N 2^t}) + s \log(\sqrt{N 2^{-t}}) + O(1) = \log n + \frac{s-1}{2}(\log n - t(v)) + O(1). \quad (1)$$

Additionally storing the level $t(v)$ using $O(\log \log n)$ bits allows to perform decoding easily, without needing the complications arising in our scheme for outerplanar graphs. Finally, combined with the label splitting technique of Kannan et al. [15] the statement of the theorem follows.

To prove the bound on the average label length we perform the following simple computation. Consider a constant $B \gg \Delta$ and let $t^* = (1 - B^{-1}) \log n$. By choice of t^* and using Lemma 19 it follows that $|W_n^{t^*}| = c' n^{1-\epsilon}$ and $|W_n \setminus W_n^{t^*}| = n - c' n^{1-\epsilon}$ for some constants $c', \epsilon > 0$. Now, the length of every label $\mathcal{L}(v)$ for $v \in W_n^{t^*}$ can be trivially bounded by $\Delta \log n$, while for vertices $v \in W \setminus W_n^{t^*}$, whose level is at least t^* , we use (1) to upper-bound the label size by

$$\log n + \frac{\Delta-1}{2}(\log n - t^*) + O(\log \log n) = \left(1 + B^{-1} \frac{\Delta-1}{2}\right) \log n + O(\log \log n).$$

Finally, combining the bounds one obtains the following bound on the sum of the label sizes.

$$\sum_{v \in V} |\mathcal{L}(v)| \leq c' n^{1-\epsilon} \Delta \log n + (n - c' n^{1-\epsilon}) \left(1 + B^{-1} \frac{\Delta-1}{2}\right) \log n + O(n \log \log n).$$

The term $dn^{1-\epsilon} \Delta \log n$ is clearly negligible, since Δ is fixed. We conclude that the average label length is at most $\left(1 + B^{-1} \frac{\Delta-1}{2}\right) (1 + h(n)) \log n + O(\log \log n)$ for some function $h(n) = o(1)$. Repeating the argument with $B = O(\log \log n)$ leads to a bound of

$$(1 + o(1)) \log n + O(\log \log n)$$

on the average label size, which is asymptotically the best possible average label size. \square

In light of the $2 \log n + O(\log \log n)$ labeling scheme of Gavaille and Labourel [13] for general planar graphs, our result improves the best known bounds for $\Delta \leq 4$, and matches the bound in [13] for $\Delta \leq 6$.