

Near Optimal Adjacency Labeling Schemes for Power-Law Graphs*

Casper Petersen¹, Noy Rotbart², Jakob Grue Simonsen^{†3}, and Christian Wulff-Nilsen⁴

1 University of Copenhagen, Copenhagen, Denmark
cazz@di.ku.dk

2 University of Copenhagen, Copenhagen, Denmark
noyro@di.ku.dk

3 University of Copenhagen, Copenhagen, Denmark
simonsen@di.ku.dk

4 University of Copenhagen, Copenhagen, Denmark
koolooz@di.ku.dk

Abstract

An adjacency labeling scheme labels the n nodes of a graph with bit strings in a way that allows, given the labels of two nodes, to determine adjacency based only on those bit strings. Though many graph families have been meticulously studied for this problem, a non-trivial labeling scheme for the important family of power-law graphs has yet to be obtained. This family is particularly useful for social and web networks as their underlying graphs are typically modelled as power-law graphs. Using simple strategies and a careful selection of a parameter, we show upper bounds for such labeling schemes of $\tilde{O}(\sqrt[n]{n})$ for power law graphs with coefficient α , as well as nearly matching lower bounds. We also show two relaxations that allow for a label of logarithmic size, and extend the upper-bound technique to produce an improved distance labeling scheme for power-law graphs.

1998 ACM Subject Classification E.1 Distributed data structures

Keywords and phrases Labeling schemes, Power-law graphs

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.133

1 Introduction

A body of work on large, real-world networks deals with the difficulties of storing them and effectively resolving queries on them; examples of techniques are compression [14, 13] and dissemination of the underlying graphs of these networks over several machines [35, 52, 54]. A different approach to storing information about the graph is to disseminate the structural information of the graph to its vertices and store it locally. This *peer-to-peer* strategy allows inferring the graph's local topology using only local information stored in each vertex without using costly access to large, global data structures. In particular, it can be useful to address privacy concerns and ensure a high survivability rate [18].

We posit that a useful tool for such a peer-to-peer strategy is the notion of a *labeling scheme*: an algorithm that assigns a bit string—a *label*—to each vertex so that a query between

* The full version of this paper is available at <http://arxiv.org/abs/1502.03971>.

† Jakob Grue Simonsen partially supported by the Danish Council for Independent Research Sapere Aude grant “Complexity via Logic and Algebra” (COLA).



any two vertices can be deduced solely from their respective labels. Labeling schemes are extremely well-studied in the algorithmic literature [8, 17, 20, 27, 32, 30, 37, 38, 40, 41, 49]; the main objective is to minimize the *maximum label size*: the maximum number of bits used in a label of any vertex. Among their applications are XML search engines [25], mapping services [1], and internet routing [43].

One class of graphs extensively used for modelling real-world networks is *power-law graphs*: roughly, n -vertex graphs where the number of vertices of degree k is proportional to n/k^α for some positive α . Power-law graphs (also called scale-free graphs in the literature) have been used to model the Internet AS-level graph [50, 4], and many other types of network (see, e.g., [24, 47] for overviews). The adequacy of fit of power-law graph models to actual data, as well as the empirical correctness of the conjectured mechanisms giving rise to power-law behaviour, have been subject to criticism (see, e.g., [2, 24]). In spite of such criticism, and because their degree distribution affords a reasonable approximation of the degree distribution of many networks, the class of power-law graphs remains a popular tool in network modelling. In this paper, we perform the first theoretical and practical study of adjacency labeling schemes for classes of graphs whose statistical properties—in particular their *degree distribution*—more closely resemble that of real-world networks.

1.1 Our contributions

Our contributions are:

A discrete and simple characterisation of power-law graphs

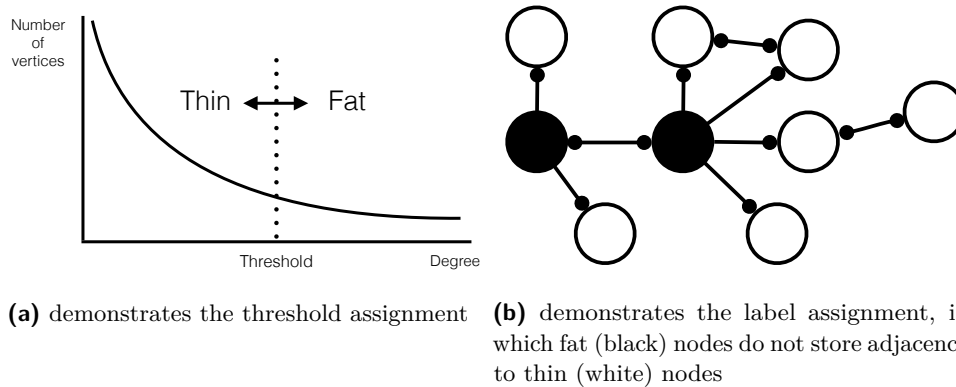
An n -vertex graph is power-law if the number of its vertices of degree k is proportional to n/k^α for some positive α . To solidify this somewhat vague definition, numerous probabilistic and deterministic definitions of power-law graphs are given in the literature. In Sec. 3, we define and prove useful properties for two simple families of graphs, \mathcal{P}_h and \mathcal{P}_l , where \mathcal{P}_h contains and \mathcal{P}_l is contained by the standard definitions of power-law graphs in the literature, including recent ones [16]. We use \mathcal{P}_h and \mathcal{P}_l to study upper and lower bounds respectively.

An $O(\sqrt[\alpha]{n}(\log n)^{1-1/\alpha})$ adjacency labeling scheme

In Sec. 4, we describe our labeling scheme, which is based on two ideas: (i) a labeling *strategy* that partitions the vertices of G into high (“fat”) and low degree (“thin”) vertices based on a threshold degree, and (ii) a threshold *prediction* that depends only on the coefficient α of a power-law curve fitted to the degree distribution of G . These ideas are illustrated in Figure 1. Using the same ideas, we get an asymptotically near-tight $O(\sqrt{n \log n})$ adjacency labeling scheme for sparse graphs. As real-world power-law graphs have $2 \leq \alpha \leq 3$ and rarely exceed 10^{10} vertices, this implies labels of the order of $10^4 - 10^5$ bits. That, and the simplicity of our labeling scheme suggests that our labeling schemes may be appealing in practice. To stress this point, we offer an experimental evaluation of our labeling scheme in the full version of the paper.

A lower bound of $\Omega(\sqrt[\alpha]{n})$ for any adjacency labeling scheme

In Sec. 5, We use our restrictive subclass of power-law graphs and show that it requires label size $\Omega(\sqrt[\alpha]{n})$ for n -vertex graphs. This lower bound shows that our upper bound above is asymptotically optimal, bar a $(\log n)^{1-1/\alpha}$ factor. By the connections between adjacency labeling schemes and universal graphs, we also obtain upper and lower bounds for induced



■ **Figure 1** Two illustrations of the main idea.

universal graphs for power-law graphs. We also show, in Sec. 6, two scenarios in which this lower bound can be bypassed.

A $o(n)$ distance labeling scheme

In Sec. 7, we demonstrate the usefulness of our strategy to arrive at a $o(n)$ distance labeling scheme for power-law graphs. Our labeling scheme is designed to outperform competing labeling schemes for small distances, in accordance with Chung and Lu’s findings [22] on the small expected diameter of power-law graphs.

1.2 Related work

Adjacency labeling schemes for key graph families are by now well understood. General graphs require a label size of $n/2 + O(1)$ [48, 8], while trees, planar graphs, and bounded degree graphs enjoy labels of logarithmic size [9, 30, 3]. Adjacency labeling schemes are also tightly coupled with the graph-theoretic concept of induced universal graphs, in which one aims to find the smallest N where there exist a graph of N vertices which contains all graphs of a particular graph family \mathcal{F}_n of n vertices as induced subgraphs. It was shown [36] that an $f(n) \log n$ adjacency labeling scheme for \mathcal{F}_n constructs an induced universal graph for this family of $2^{f(n)}$ vertices. In the context of sparse graphs, a body of work on universal graphs¹ for this family was investigated both by Babai et al. [11] and by Alon and Asodi [5].

Routing labeling schemes for power-law graphs have been investigated by Brady and Cowen [17], and by Chen et al. [21]. Labeling schemes for properties other than adjacency have been investigated for various classes of graphs, e.g., distance [32], and flow [37]. Dynamic labeling schemes were studied by Korman and Peleg [38, 40, 41] and recently by Dahlgaard et al [27]. Experimental evaluations for some labeling schemes for various properties on general graphs have been performed by Caminiti et. al [20], Fischer [29] and Rotbart et. al [49].

In the context of distributed graph computing systems, a somewhat related paradigm of computation is the vertex centric computing model “think like a vertex”. In this model each vertex exchanges messages only with nearby vertices, to improve locality and simplify the design and implementation of such systems. Among the numerous systems proposed are Pregel [45], Power-Graph [35] and GraphLab [44]. For a recent survey on the topic see [46].

¹ A graph that contains each graph from the graph family as a subgraph, not necessarily induced.

2 Preliminaries

Throughout the paper we consider n -vertex, undirected graphs. For a real $c > 0$, a graph is c -sparse if it has at most cn edges and *sparse* if it is c -sparse for some constant c . For $0 < c \leq n - 1$, the set of c -sparse graphs with n vertices is denoted by $\mathcal{S}_{c,n}$. If \mathcal{F} is a set of graphs, \mathcal{F}_n denotes the subset of graphs in \mathcal{F} having exactly n vertices. The degree of a vertex v in a graph is denoted by $\Delta(v)$, and for non-negative integers k , the set of vertices in a graph G of degree k is denoted by V_k . The length of a binary string $x \in \{0, 1\}^*$ is denoted by $|x|$.

Let \mathcal{F} be a set of graphs. An *adjacency labeling scheme* (from hereon just *labeling scheme*) for \mathcal{F} is a pair consisting of an *encoder* and a *decoder*. The encoder is an algorithm that receives $G \in \mathcal{F}$ as input and outputs a bit string $\mathcal{L}(v) \in \{0, 1\}^*$, called the *label* of v , for each vertex v in G . The decoder is an algorithm that receives any two labels $\mathcal{L}(v), \mathcal{L}(u)$ as input and outputs **true** if u and v are adjacent in G and **false** otherwise. Note that the graph G is not an input to the decoder. The *size* of a labeling scheme is the map $\text{size} : \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{size}(n)$ is the maximum length of any label assigned by the encoder to any vertex in any graph $G \in \mathcal{F}_n$. The *degree distribution* of a graph $G = (V, E)$ is the mapping $\text{ddist}_G(k) : \mathbb{N}_0 \rightarrow \mathbb{Q}$ defined by $\text{ddist}_G(k) := \frac{|V_k|}{n}$.

3 Defining Power-Law Graphs

In the literature *power-law* graphs are usually defined as the class of n vertex graphs G such that $\text{ddist}_G(k)$ is proportional to $k^{-\alpha}$ for some real number $\alpha > 1$. Ideally, and ignoring rounding, $\text{ddist}_G(k) = Ck^{-\alpha}$ for all k for constant C . As the degree distribution of a graph must be a probability distribution, we have $\sum_{k=1}^{\infty} Ck^{-\alpha} = C \sum_{k=1}^{\infty} k^{-\alpha} = 1$, hence $C = 1/\zeta(\alpha)$ where ζ is the Riemann zeta function. However, in the literature, concessions are usually made that relax the restrictions on $\text{ddist}_G(k)$, for example that the power-law property need only hold for high-degree vertices (“above a cutoff”), or that $\text{ddist}_G(k)$ is only *approximately* equal to $Ck^{-\alpha}$, with some approximation error that falls off with n . To ensure that our results hold for all these variations of power-law graphs, we define two families of graphs \mathcal{P}_h and \mathcal{P}_l with $\mathcal{P}_l \subsetneq \mathcal{P}_h$. Family \mathcal{P}_h is rich enough to contain the graphs whose degree distribution is approximately, or perfectly, power-law distributed, and our upper bound on the label size for our labeling scheme holds for any graph in \mathcal{P}_h . Family \mathcal{P}_l is used to show our lower bound and is restrictive enough that most definitions of power-law graph occurring in the literature will contain it.

In the following, let $i_1 = \Theta(\sqrt[n]{n})$ be the smallest integer such that $\lfloor Cn/i_1^\alpha \rfloor \leq 1$, and let $C' \geq (\frac{C}{\alpha-1} + \frac{i_1}{\sqrt[n]{n}} + 5)^\alpha + \frac{C}{\alpha-1}$ be a constant; we shall use C' in the remainder of the paper.

► **Definition 1.** Let $\alpha > 1$ be a real number and let $\chi : \mathbb{N} \rightarrow \mathbb{N}$ be a function. $\mathcal{P}_{h,\chi,\alpha}$ is the family of graphs G such that if $n = |V(G)|$ then for all integers k between $\chi(n)$ and $n - 1$, $\sum_{i=k}^{n-1} |V_i| \leq C'(\frac{n}{k^{\alpha-1}})$. We shall usually suppress χ and α , writing merely \mathcal{P}_h .

The function χ captures the notion of a cutoff as defined in [24, Sec. 3.1]; the intuition is that for an n -vertex graph the power-law distribution need only apply for nodes of degree higher than $\chi(n)$, rather than for all degrees. Setting $\chi(n) = 1$ corresponds to the case where the entire range of degrees follows a power-law distribution, hence even for small values of $\chi(n)$, \mathcal{P}_h morally contains all graphs with power-law degree distribution. We will later prove upper bounds that hold for *all* χ bounded from above by some function; in particular for the upper bound for adjacency labeling schemes, the bound holds for $\chi(n)$ as high as $\sqrt[n]{n/\log n}$.

The class \mathcal{P}_l contains graphs where the number of vertices of degree k must be $C \frac{n}{k^\alpha}$ rounded either up or down and the number of vertices of degree k is non-increasing with k . Note that the function $k \mapsto C \frac{1}{k^\alpha}$ is strictly decreasing.

► **Definition 2.** Let $\alpha > 1$ be a real number and let $C = 1/\zeta(\alpha)$ where ζ is the Riemann zeta function. $\mathcal{P}_{l,\alpha}$ is the set of graphs $G = (V, E)$ such that

1. $\lfloor Cn \rfloor - i_1 - 1 \leq |V_1| \leq \lceil Cn \rceil$,
2. $\lfloor C \frac{n}{2^\alpha} \rfloor \leq |V_2| \leq \lceil C \frac{n}{2^\alpha} \rceil + 1$,
3. for every i with $3 \leq i \leq n$: $|V_i| \in \{\lfloor C \frac{n}{i^\alpha} \rfloor, \lceil C \frac{n}{i^\alpha} \rceil\}$, and
4. for every i with $2 \leq i \leq n - 1$: $|V_i| \geq |V_{i+1}|$.

We usually suppress α , writing just \mathcal{P}_l .

Note that we allow slightly more noise in the sizes of V_1 and V_2 than in the remaining sets; without it, it seems tricky to prove a better lower bound than $\Omega(n^{\frac{1}{\alpha+1}})$ on label sizes.

We show the following properties of \mathcal{P}_l .

► **Proposition 1.** *The maximum degree in an n -vertex graph in \mathcal{P}_l is at most $(\frac{C}{\alpha-1} + 2) \sqrt[\alpha]{n} + i_1 + 3 = \Theta(\sqrt[\alpha]{n})$.*

Proof. Let $n > 0$ be an integer and let $k' = \lfloor \sqrt[\alpha]{n} \rfloor$. Furthermore, let $S_{k'} = \sum_{i=1}^{k'} |V_i|$, that is $S_{k'}$ is the number of vertices of degree at most k' . Let $S_{k'}^- = (\sum_{i=1}^{k'} \lfloor Cni^{-\alpha} \rfloor) - i_1 - 1$. Then $S_{k'} \geq S_{k'}^-$. We now bound $S_{k'}^-$ from below. For every i with $1 \leq i \leq k'$,

$$\begin{aligned} S_{k'}^- + k' &= -i_1 - 1 + \sum_{i=1}^{k'} (\lfloor Cni^{-\alpha} \rfloor + 1) \geq -i_1 - 1 + \sum_{i=1}^{k'} Cni^{-\alpha} = -i_1 - 1 + Cn \sum_{i=1}^{k'} i^{-\alpha} \\ &\geq n \left(1 - C \sum_{i=k'+1}^{\infty} i^{-\alpha} \right) - i_1 - 1 \geq n \left(1 - C \int_{k'}^{\infty} x^{-\alpha} dx \right) - i_1 - 1 \\ &= n \left(1 - C \left[\frac{1}{\alpha-1} x^{-\alpha+1} \right]_{k'}^{k'} \right) - i_1 - 1 = n \left(1 - \frac{C}{\alpha-1} (\lceil \sqrt[\alpha]{n} \rceil)^{-\alpha+1} \right) - i_1 - 1 \\ &\geq n \left(1 - \frac{C}{\alpha-1} (\sqrt[\alpha]{n})^{-\alpha+1} \right) - i_1 - 1 = n - \frac{Cn}{\alpha-1} n^{-1+\frac{1}{\alpha}} - i_1 - 1 \\ &= n - \frac{C}{\alpha-1} \sqrt[\alpha]{n} - i_1 - 1, \end{aligned}$$

giving $S_{k'} \geq S_{k'}^- \geq n - \frac{C}{\alpha-1} \sqrt[\alpha]{n} - \lceil \sqrt[\alpha]{n} \rceil - i_1 - 1$. There are thus at most $\frac{C}{\alpha-1} \sqrt[\alpha]{n} + \lceil \sqrt[\alpha]{n} \rceil + i_1 + 1$ vertices of degree strictly more than $k' = \lfloor \sqrt[\alpha]{n} \rfloor$. Since for every $1 \leq i \leq n - 1$: $|V_i| \geq |V_{i+1}|$, it follows that the maximum degree of any graph in \mathcal{P}_l is at most $(\frac{C}{\alpha-1} + 2) \sqrt[\alpha]{n} + i_1 + 3$. ◀

► **Proposition 2.** *For $\alpha > 2$, all graphs in \mathcal{P}_l are sparse.*

Proof. By Proposition 1, the maximum degree of an n -vertex graph in \mathcal{P}_l graph is at most $k' \triangleq (\frac{C}{\alpha-1} + 2) \sqrt[\alpha]{n} + i_1 + 3$, whence the total number of edges is at most $\frac{1}{2} \sum_{k=1}^{k'} k|V_k|$. By definition, $|V_k| \leq \lceil \frac{Cn}{k^\alpha} \rceil \leq \frac{Cn}{k^\alpha} + 1$ for $k \neq 2$ and $|V_2| \leq \lceil \frac{Cn}{2^\alpha} \rceil + 1$, and thus

$$\begin{aligned} \frac{1}{2} \sum_{k=1}^{k'} k|V_k| &\leq 1 + \frac{1}{2} \sum_{k=1}^{k'} k \left(\frac{Cn}{k^\alpha} + 1 \right) \leq 1 + \frac{k'(k'+1)}{4} + Cn \sum_{k=1}^{\infty} k^{-\alpha+1} \\ &= O(n^{2/\alpha}) + Cn\zeta(\alpha-1) = O(n). \end{aligned}$$

► **Proposition 3.** For any χ and $\alpha > 1$, $\mathcal{P}_{l,\alpha} \subseteq \mathcal{P}_{h,\chi,\alpha}$.

Proof. Let $d = \lfloor (\frac{C}{\alpha-1} + 2) \sqrt[\alpha]{n} + i_1 + 3 \rfloor$. For any graph in \mathcal{P}_l with n vertices and for any k , $|V_k| \leq Ck^{-\alpha}n + 1$ and by Proposition 1, $|V_k| = 0$ when $k > d$.

Let k be an arbitrary integer between $\chi(n)$ and $n - 1$. We need to show that $\sum_{i=k}^{n-1} |V_i| \leq C'(\frac{n}{k^{\alpha-1}})$. It suffices to show this for $k \leq d$. We have:

$$\begin{aligned} \sum_{i=k}^{n-1} |V_i| &\leq \sum_{i=k}^d (Ci^{-\alpha}n + 1) = d - k + 1 + Cn \sum_{i=k}^d i^{-\alpha} \\ &\leq \left(\frac{C}{\alpha-1} + \frac{i_1}{\sqrt[\alpha]{n}} + 5 \right) \sqrt[\alpha]{n} + Cn \int_k^d x^{-\alpha} dx \\ &\leq \left(\frac{C}{\alpha-1} + \frac{i_1}{\sqrt[\alpha]{n}} + 5 \right) \sqrt[\alpha]{n} + Cn \left[\frac{1}{\alpha-1} x^{-\alpha+1} \right]_k^\infty \\ &\leq \left(\left(\frac{C}{\alpha-1} + \frac{i_1}{\sqrt[\alpha]{n}} + 5 \right) \left(\frac{\sqrt[\alpha]{n}d^{\alpha-1}}{n} \right) + \frac{C}{\alpha-1} \right) nk^{-\alpha+1} \\ &\leq \left(\frac{C}{\alpha-1} + \frac{i_1}{\sqrt[\alpha]{n}} + 5 \right) \left(\frac{C}{\alpha-1} + \frac{i_1}{\sqrt[\alpha]{n}} + 5 \right)^{\alpha-1} nk^{-\alpha+1} + \left(\frac{C}{\alpha-1} \right) nk^{-\alpha+1} \\ &\leq C'nk^{-\alpha+1}, \end{aligned}$$

as desired. ◀

3.1 Comparison to other deterministic models

Numerous probabilistic and deterministic definitions of power-law graphs are given in the literature. A recent deterministic model, called shifted power-law distribution [28] has recently proven to capture a vast number of such definitions, both in theory and experimentally in [16]. We show that our definition of \mathcal{P}_h contains graphs that adhere to the model, which is defined as follows. Let $c_1 > 0$ be a constant. A graph G is *power-law bounded* for parameters $\alpha > 1$ and $t \geq 0$ if for every integer $d \geq 0$, the number of vertices of G of degree in $[2^d, 2^{d+1})$ is at most

$$c_1 n(t+1)^{\alpha-1} \sum_{i=2^d}^{2^{d+1}-1} (i+t)^{-\alpha}.$$

As experimentally verified in [16], the value of t is typically very small. If $t = O(1)$, the bound above becomes $O(n \sum_{i=2^d}^{2^{d+1}-1} i^{-\alpha})$. In this case, our family $\mathcal{P}_h(\chi, \alpha)$ is rich enough to contain these power-law bounded graphs for sufficiently large C' and any choice of χ and α . This follows since for any power-law bounded graph with n vertices and any integer k between 1 and $n - 1$, $\sum_{i=k}^{n-1} |V_i| = O(\sum_{d=\lceil \lg k \rceil}^{\lceil \lg(n-1) \rceil} n \sum_{i=2^d}^{2^{d+1}-1} i^{-\alpha}) = O(\frac{n}{k^{\alpha-1}})$. Thus our upper bound also applies to power-law bounded graphs. It is possible to extend our upper bound to super-constant t where the bound is stronger the smaller t is; we omit the details. Conversely, our family \mathcal{P}_l is restrictive enough that \mathcal{P}_l is contained in the family of power-law bounded graphs when $t = O(1)$, and the lower bound we derive thus also holds in that setting.

4 The Labeling Schemes

We now construct algorithms for labeling schemes for c -sparse graphs and for the family \mathcal{P}_h . Both labeling schemes partition vertices into *thin* vertices which are of low degree and *fat*

vertices of high degree. The *degree threshold* for the scheme is the lowest possible degree of a fat vertex. We start with c -sparse graphs.

► **Theorem 3.** *There is a $\sqrt{2cn \log n} + 2 \log n + 1$ labeling scheme for $\mathcal{S}_{c,n}$.*

Proof. Let $G = (V, E)$ be an n -vertex c -sparse graph. Let $\tau(n)$ be the degree threshold for n -vertex graphs; we choose $\tau(n)$ below. Let k denote the number of fat vertices of G , and assign each fat vertex a unique identifier between 1 and k . Each thin vertex is given a unique identifier between $k + 1$ and n .

For a $v \in V$, the first part of the label $\mathcal{L}(v)$ is a single bit indicating whether v is thin or fat followed by a string of $\log n$ bits representing its identifier. If v is thin, the last part of $\mathcal{L}(v)$ is the concatenation of the identifiers of the neighbors of v . If v is fat, the last part of $\mathcal{L}(v)$ is a *fat bit string* of length k where the i th bit is 1 iff v is incident to the (fat) vertex with identifier i .

Decoding a pair $(\mathcal{L}(u), \mathcal{L}(v))$ is straightforward: if one of the vertices, say u , is thin, u and v are adjacent iff the identifier of v is part of the label of u . If both u and v are fat then they are adjacent iff the i th bit of the fat bit string of $\mathcal{L}(u)$ is 1 where i is the identifier of v . Both decoding processes can be computed in $O(\log n)$ time using standard assumptions.

Since $|E| \leq cn$, we have $k \leq 2cn/\tau(n)$. A fat vertex thus has label size $1 + \log n + k \leq 1 + \log n + 2cn/\tau(n)$ and a thin vertex has label size at most $1 + \log n + \tau(n) \log n$. To minimize the maximum possible label size, we solve $2cn/x = x \log n$. Solving this gives $x = \sqrt{2cn/\log n}$ and setting $\tau(n) = \lceil x \rceil$ gives a label size of at most $1 + \log n + (\sqrt{2cn/\log n} + 1) \log n \leq 1 + 2 \log n + \sqrt{2cn \log n}$. ◀

By Proposition 2, graphs in \mathcal{P}_l are sparse for $\alpha > 2$. This gives a label size of $O(\sqrt{n \log n})$ with the labeling scheme in Theorem 3. We now show that this label can be significantly improved, by constructing a labeling scheme for \mathcal{P}_h which contains \mathcal{P}_l .

► **Theorem 4.** *There is a $\sqrt[3]{C'n}(\log n)^{1-1/\alpha} + 2 \log n + 1$ labeling scheme for \mathcal{P}_h .*

Proof. The proof is very similar to that of Theorem 3. We let $\tau(n)$ denote the degree threshold. If we pick $\tau(n) \geq \sqrt[3]{n/\log n}$ then by Definition 1 there are at most $C'n/\tau(n)^{\alpha-1}$ fat vertices. Defining labels in the same way as in Theorem 3 gives a label size for thin vertices of at most $1 + \log n + \tau(n) \log n$ and a label size for fat vertices of at most $1 + \log n + C'n/\tau(n)^{\alpha-1}$. We minimize by solving $x \log n = C'n/x^{\alpha-1}$, giving $x = \sqrt[3]{C'n/\log n}$. Setting $\tau(n) = \lceil x \rceil$ gives a label size of at most $\sqrt[3]{C'n}(\log n)^{1-1/\alpha} + 2 \log n + 1$. ◀

4.1 A labeling scheme for random graphs

There are schemes using randomness to “grow” graphs that, with high probability, have an approximate power-law degree distribution for a range of degrees (see e.g. [23]). For graphs obtained from such models, their degree sequences are instead probability distributions. We now show that applying our labeling scheme for \mathcal{P}_h to random graphs with the power-law distribution results in a small expected worst-case label size.

Using the definition of Mitzenmacher [47], a random variable X is said to have the *power-law* distribution (w.r.t. $\alpha > 1$) if

$$\Pr[X \geq x] \sim cx^{-\alpha+1},$$

for a constant $c > 0$, i.e., $\lim_{x \rightarrow \infty} \Pr[X \geq x]/cx^{-\alpha+1} = 1$.

Let $\epsilon > 0$ be fixed. Consider a graph G picked from a family \mathcal{F} of random graphs whose degree sequences have the power-law distribution. Order the vertices of G arbitrarily

as v_1, \dots, v_n . For $i = 1, \dots, n$, let indicator variable X_i be 1 iff v_i has degree at least $d = \sqrt[\alpha]{n/\log n}$. There is a constant $N_0 \in \mathbb{N}$ (depending on ϵ) such that if $n \geq N_0$ then for all i ,

$$E[X_i] = \Pr[X_i = 1] \leq (1 + \epsilon)cd^{-\alpha+1}.$$

With the same labeling scheme as for \mathcal{P}_h with degree threshold $\tau(n) = d$, denote by E_n the expected label size of an n -vertex graph from \mathcal{F} . Then for all $n \geq N_0$,

$$\begin{aligned} E_n &= \sum_{x=0}^n \Pr \left[\sum_{i=1}^n X_i = x \right] O((x + d \log n)) = O \left(d \log n + E \left[\sum_{i=1}^n X_i \right] \right) \\ &= O \left(d \log n + \sum_{i=1}^n E[X_i] \right) = O(d \log n + nd^{-\alpha+1}) = O \left(\sqrt[\alpha]{n}(\log n)^{1-1/\alpha} \right). \end{aligned}$$

Thus, we have:

► **Theorem 5.** *Let \mathcal{F} be a family of graphs with degree sequences having the power-law distribution w.r.t. $\alpha > 1$. Then there is a labeling scheme for \mathcal{F} such that the expected worst-case label size of any graph $G \in \mathcal{F}$ is $O(\sqrt[\alpha]{n}(\log n)^{1-1/\alpha})$ where n is the number of vertices of G .*

5 Lower Bounds

We now derive lower bounds for the label size of any labeling schemes for both $\mathcal{S}_{c,n}$ and \mathcal{P}_l . Our proofs rely on Moon's [48] lower bound of $\lfloor n/2 \rfloor$ bits for labeling scheme for general graphs. We first show that the upper bound achieved for sparse graphs is close to the best possible. The following proposition is essentially a more precise version of the lower bound suggested by Spinrad [51].

► **Proposition 4.** *Any labeling scheme for $\mathcal{S}_{c,n}$ requires labels of size at least $\lfloor \frac{\sqrt{cn}}{2} \rfloor$ bits.*

Proof. Assume for contradiction that there exists a labeling scheme assigning labels of size strictly less than $\lfloor \frac{\sqrt{cn}}{2} \rfloor$. Let G be an n -vertex graph. Let G' be the graph resulting by adding $\lfloor \frac{n^2}{c} \rfloor - n$ isolated vertices to G , and note that now G' is c -sparse. The graph G is an induced subgraph of G' . It now follows that the vertices of G have labels of size strictly less than $\lfloor \frac{\sqrt{c\lfloor n^2/c \rfloor}}{2} \rfloor \leq n/2$ bits. As G was arbitrary, we obtain a contradiction. ◀

In the remainder of this section we are assuming that $\alpha > 2$ and prove the following:

► **Theorem 6.** *For any n , any labeling scheme for n -vertex graphs of $\mathcal{P}_{h,\chi,\alpha}$ requires label size $\Omega(\sqrt[\alpha]{n})$.*

More precisely, we present a lower bound for \mathcal{P}_l which is contained in \mathcal{P}_h . Let $n \in \mathbb{N}$ be given and let $H = (V(H), E(H))$ be an arbitrary graph with i_1 vertices where $i_1 = \Theta(\sqrt[\alpha]{n})$ is defined as in Section 3. We show how to construct a graph $G = (V, E)$ in \mathcal{P}_l with n vertices that contains H as an induced subgraph. Observe that a labeling of G induces a labeling of H . As H was chosen arbitrarily and as any labeling scheme for k -vertex graphs requires $\lfloor i_1/2 \rfloor$ label size in the worst case, Theorem 6 follows if we can show the existence of G .

We construct G incrementally where initially $E = \emptyset$. Partition V into subsets V_1, \dots, V_n as follows. The set V_1 has size $\lfloor Cn \rfloor - i_1$. For $i = 2, \dots, i_1 - 1$, V_i has size $\lfloor Cn/i^\alpha \rfloor$. Letting

$n' = \sum_{i=1}^{i_1-1} |V_i|$, we set the size of V_i to 1 for $i = i_1, \dots, i_1 + n - n' - 1$ and the size of V_i to 0 for $i = i_1 + n - n', \dots, n$, thereby ensuring that the sum of sizes of all sets is n . Observe that $\sum_{i=1}^{i_1} \lfloor Cn/i^\alpha \rfloor \leq n$ so that $n' \leq n - i_1$, implying that $n - n' \geq i_1$. Hence we have at least i_1 size 1 subsets $V_{i_1}, \dots, V_{i_1+n-n'-1}$ in each of which the vertex degree allowed by Definition 2 is at least i_1 .

Let v_1, \dots, v_{i_1} be an ordering of $V(H)$, form a set $V_H \subseteq V$ of i_1 arbitrary vertices from the sets $V_{i_1}, \dots, V_{i_1+n-n'-1}$, and choose an ordering v'_1, \dots, v'_{i_1} of V_H . For all $i, j \in \{1, \dots, i_1\}$, add edge (v'_i, v'_j) to E iff $(v_i, v_j) \in E(H)$. Now, H is an induced subgraph of G and since the maximum degree of H is $i_1 - 1$, no vertex of V_i exceeds the degree bound allowed by Definition 2 for $i = 1, \dots, n$.

We next add additional edges to G in three phases to ensure that it is an element of \mathcal{P}_l while maintaining the property that H is an induced subgraph of G . For $i = 1, \dots, n$, during the construction of G we say that a vertex $v \in V_i$ is *unprocessed* if its degree in the current graph G is strictly less than i . If the degree of v is exactly i , v is *processed*.

Phase 1

Let $V' = V \setminus (V_1 \cup V_H)$. Phase 1 is as follows: while there exists a pair of unprocessed vertices $(u, v) \in V' \times V_H$, add (u, v) to E .

When Phase 1 terminates, H is clearly still an induced subgraph of G . Furthermore, all vertices of V_H are processed. To see this, note that the sum of degrees of vertices of V_H when they are all processed is $O(i_1^2) = O(n^{2/\alpha})$ which is $o(n)$ since $\alpha > 2$. Furthermore, prior to Phase 1, each of the $\Theta(n)$ vertices of V' have degree 0 and can thus have their degrees increased by at least 1 before being processed.

Phase 2

While there exists a pair of unprocessed vertices $(u, v) \in V' \times V'$, add (u, v) to E . At termination, at most one vertex of V' remains unprocessed. If such a vertex exists we process it by connecting it to $O(\sqrt[n]{n})$ vertices of V_1 ; as $|V_1| = \Theta(n)$ there are enough vertices of V_1 to accomodate this. Furthermore, prior to adding these edges, all vertices of V_1 have degree 0, and hence the bound allowed for vertices of this set is not exceeded.

Phase 3

We add edges between pairs of unprocessed vertices of V_1 until no such pair exists. If no unprocessed vertices remain we have the desired graph G . Otherwise, let $w \in V_1$ be the unprocessed vertex of degree 0. We add a single edge from w to another vertex w' of V_1 , thereby processing w and moving w' from V_1 to V_2 . Note that the sizes of V_1 and V_2 are kept in their allowed ranges due to the first two conditions in Definition 2. This proves Theorem 6.

6 $O(\log n)$ adjacency labeling schemes for some power-law graphs

The lower bound presented can be avoided in two interesting cases. The first, for random graphs generated by a popular model, and the second using an extension of the concept of labeling schemes from the literature.

BA model

As discussed in Sec. 4.1, generative models play an important role in the study of power-law graphs. Perhaps the most well-known generative model is the Barabási-Albert (BA) which, roughly, grows a graph in a sequence of time steps by inserting a single vertex at each step and attaching it to m existing vertices with probability weighted by the degree of each existing vertex [12]. The BA model generates graphs that asymptotically have a power-law degree distribution ($\alpha = 3$) for low-degree nodes [15]. However, graphs created by the BA model have low arboricity² [34]. We use this fact to devise the following highly efficient labeling scheme for such graphs.

► **Proposition 5.** *The family of graphs generated by the BA model has an $O(m \log n)$ adjacency labeling scheme.*

Proof. Let $G = (V, E)$ be an n -vertex graph resulting by the construction by the BA model with some parameter m (starting from some graph $G_0 = (V_0, E_0)$ with $|V_0| \ll n$). While it is not known how to compute the arboricity of a graph efficiently, it is possible in near-linear time to compute a partition of G with at most twice³ the number of forests in comparison to the optimal [10]. We can thus decompose the graph to $2m$ forests in near linear time and label each forest using the recent $\log n + O(1)$ labeling scheme for trees [6], and achieve a $2m(\log n + O(1))$ labeling scheme for G . ◀

If the encoder operates at the same time as the creation of the graph, Proposition 5 can be tightened to yield a $m \log n$ labeling scheme, by storing the identifiers of the vertices to the node introduced. Theorem 6 and Proposition 5 strongly suggest that local properties of power-law graphs are very different from those of a randomly generated graph using the BA model. In contrast, other generative models such as Waxman's [53], N-level Hierarchical [19], and Chung and Liu's [23] (Chapter 3) do not seem to have an obvious smaller label size than the one in Proposition 4.

Labeling schemes with a query

The concept of labeling scheme limits the number of nodes participating in a query severely. A relaxed variant thereof, called 1-query labeling scheme [39], assumes that the decoder receives both labels queried, and may access the label of a third node in order to answer the query. If this is allowed, we can construct an $O(\log n)$ 1-query adjacency labeling scheme for sparse (and power-law) graphs as follows: We assign each node v with an identifier $ID(v)$, then produce a classic [26] chaining perfect hash-function⁴ from $\{1 \dots cn\}$ to $\{1 \dots n\}$, with the guarantee that the worst case number of collisions is constant. We then compute the hash function for all edges (u, v) and store the tuple $\langle ID(v), ID(u) \rangle$ in the label of the corresponding vertex. The decoder first computes the hash value resulting from $ID(v)$ and $ID(u)$ and proceed to examine if on the label corresponding to the result of the function the tuple appears. The decoder needs only to know the primary and secondary hash functions used, description thereof amount to logarithmic number of bits, which can be concatenated to each label.

² the arboricity of a graph is the minimum number of spanning forests needed to cover its edges.

³ More precisely, for any $\epsilon \in (0, 1)$ there exist an $O(|E(G)|/\epsilon)$ algorithm [42] that computes such partition using at most $(1 + \epsilon)$ times more forests than the optimal one.

⁴ To this end, we may for example first partition the domain into c parts.

7 A distance labeling scheme

In this section we extend the usefulness of our strategy by showing a labeling scheme for small distances in power-law graphs.

For sparse graphs, Alstrup et al. [7] obtain a distance labeling scheme with maximum label size $O(\frac{n}{D} \log^2 D)$ where $D = (\log n)/(\log \frac{m+n}{n})$ and m is the number of edges in the graph. Gawrychowski et al. obtain an upper bound of [33] $O(\frac{n}{D} \log D)$ with sub-linear decoding time. Few general results on lower bounds exist. The lower bound of $\Omega(\sqrt{n})$ for adjacency given in the present paper is trivially also a lower bound for distance; for total label size, the best known lower bound remains $\Omega(n^{3/2})$ as proved by Gavoille et al. [31].

► **Lemma 7.** *For any computable $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $f(n) \leq n - 1$ for all n , and for any $\chi(n) \geq n^{1/(\alpha-1+f(n))}$ there is an $f(n)$ -distance labeling scheme for $\mathcal{P}_{h,\chi,\alpha}$ that assigns labels of length at most $O(n^{f(n)/(f(n)+1)} \log f(n))$.*

Proof. Let G be a graph in $\mathcal{P}_{h,\chi,\alpha}$. A node of G is *fat* if it has degree at least $n^{1/(\alpha-1+f(n))}$ and *thin* otherwise. The label of each node v contains (i) a table of distances to all fat nodes (if the distance is more than $f(n)$, it is simply ignored), (ii) a table of distances to all thin nodes w that are at most distance $f(n)$ away from v where the shortest path between v and w does not pass through any fat node, and (iii) a single bit signifying whether the node is fat or thin. Clearly, as $f(n)$ is computable and distances in G are computable, there is a computable encoder assigning labels. A decoder can now compute the distance between any two nodes u, v as follows: If both u or v are fat, the distance can be directly read off part (i) of the label of any node. If at least one of u and v is fat, the distance can be read off part (i) of the label of the thin node. If both nodes are thin, the decoder can check if the distance is in part (ii) of the label of either node; if the distance is not present, either the distance is strictly greater than $f(n)$, or the shortest path between u and v passes through a fat node; in this case, the decoder may brute-force check the distances from u and v to each fat node, and output the smallest sum of these two distances.

Furthermore, as all nodes of G are either thin or fat, it is clearly possible for an encoder to compute all distances less than or equal to $f(n)$ between any pair of nodes. Note that as all distances we care for are bounded above by $f(n)$, each such distance can be stored using at most $\log f(n)$ bits.

As $G = G(V, E)$ is in $\mathcal{P}_{h,\chi,\alpha}$, we have

$$\begin{aligned} \sum_{i=\chi(n)}^{n-1} |V_i| &\leq \sum_{i=n^{\frac{1}{\alpha-1+f(n)}}}^{n-1} |V_i| \leq C' \left(\frac{n}{\left(n^{\frac{1}{\alpha-1+f(n)}}\right)^{\alpha-1}} \right) \\ &\leq C' n^{1-(\alpha-1)/\alpha-1+f(n)} = C' n^{f(n)/(\alpha-1+f(n))} \end{aligned}$$

Thus, a table of distances to all fat nodes takes up at most $O\left(n^{\frac{f(n)}{\alpha-1+f(n)}} \log f(n)\right)$ bits.

Similarly, for each node v there are at most $\left(n^{1/(\alpha-1+f(n))}\right)^{f(n)} = n^{f(n)/(\alpha-1+f(n))}$ nodes at distance at most $f(n)$ away from v where the shortest path consists only of thin nodes. Hence, the associated table of distances takes up at most $O(n^{f(n)/(\alpha-1+f(n))} \log n)$ bits.

In total, each label thus has size at most $O(n^{f(n)/(f(n)+1)} \log n)$ bits. ◀

For $f(n) = \log n$, Lemma 7 yields labels of size $O\left(n^{(\log n)/(\alpha-1+\log n)} \log \log n\right)$. Unsurprisingly, as we are only considering distances up to $f(n)$, this label size is asymptotically

smaller than for the labeling schemes working for all distances in *sparse* graphs, e.g. the largest label sizes of [33] for sparse graphs is $O(n \frac{\log \log n}{\log n})$. For power-law random graphs, Chung and Lu show in [22] that, subject to mild conditions, the diameter of power-law graphs with $\alpha > 2$ is almost surely $\Theta(\log n)$. We thus expect our labeling scheme to have superior performance for such graphs.

8 Conclusion and Future Work

We have devised adjacency and distance labeling schemes for sparse graphs and graphs whose degree distribution approximately follows a power-law distribution. We have proven lower bounds for the class of power-law graphs showing that our strategy for adjacency labeling scheme is almost optimal, and showed two relaxations that allow for logarithmic size labels. In the full version of the paper we also validate experimentally that the labeling scheme for power-law graphs obtains results in practice requiring little space, and that the theoretical threshold we use in our strategy is reasonably close to the optimum threshold.

8.1 Future work

We propose the following directions:

- Our labeling schemes are designed for static networks, and while it seems not difficult to extend our idea to dynamic networks, an analysis is required to account for the communication and number of re-labels incurred by such an extension.
- Labeling schemes for power-law graphs can likely be devised for the realistic case where the scheme only has incomplete knowledge of the graph, for example when the expected frequency of vertices of each degree is known, but not the exact frequency of each vertex.
- Closing the gap of the multiplicative logarithmic factor may be of interest to the theory community. A more interesting gap exists for distance labeling schemes. As we have seen, there is a large gap between labeling schemes for short distance and adjacency for power-law (and sparse) graphs. This gap effectively deemed the distance labels uninteresting for practical applications.
- Finally, while power-law distributions may model the degree distribution of real-world networks, other distributions may fit better (see, e.g., [24]); it is interesting to see whether refinements of our labeling scheme that utilize knowledge about such distributions would result in superior labeling schemes for real-world data.

References

- 1 Ittai Abraham, Daniel Delling, Andrew V Goldberg, and Renato F Werneck. A hub-based labeling algorithm for shortest paths in road networks. In *Experimental Algorithms*, pages 230–241. Springer, 2011.
- 2 Dimitris Achlioptas, Aaron Clauset, David Kempe, and Cristopher Moore. On the bias of traceroute sampling: Or, power-law degree distributions in regular graphs. *J. ACM*, 56(4), 2009. doi:10.1145/1538902.1538905.
- 3 David Adjashvili and Noy Rotbart. Labeling schemes for bounded degree graphs. In *Automata, Languages, and Programming*, pages 375–386. Springer, 2014.
- 4 Aditya Akella, Shuchi Chawla, Arvind Kannan, and Srinivasan Seshan. Scaling properties of the internet graph. In *Proceedings of the Twenty-Second ACM Symposium on Principles of Distributed Computing, PODC 2003*, pages 337–346, 2003. doi:10.1145/872035.872087.
- 5 Noga Alon and Vera Asodi. Sparse universal graphs. *J. Comput. Appl. Math.*, 142(1):1–11, May 2002. doi:10.1016/S0377-0427(01)00455-1.

- 6 Stephen Alstrup, Søren Dahlgaard, and Mathias Bæk Tejs Knudsen. Optimal induced universal graphs and adjacency labeling for trees. In *Proceedings of the 58th Symposium on Foundations of Computer Science, FOCS'15*, Washington, DC, USA, 2015. IEEE Computer Society.
- 7 Stephen Alstrup, Søren Dahlgaard, Mathias Bæk Tejs Knudsen, and Ely Porat. Sublinear distance labeling for sparse graphs. *CoRR*, abs/1507.02618, 2015. URL: <http://arxiv.org/abs/1507.02618>.
- 8 Stephen Alstrup, Haim Kaplan, Mikkel Thorup, and Uri Zwick. Adjacency labeling schemes and induced-universal graphs. *To appear in the 47th symposium on Theory of computing (STOC)*, 2015.
- 9 Stephen Alstrup and Theis Rauhe. Small induced-universal graphs and compact implicit graph representations. In *Proceedings of the 43rd Symposium on Foundations of Computer Science, FOCS'02*, pages 53–62, Washington, DC, USA, 2002. IEEE Computer Society. URL: <http://dl.acm.org/citation.cfm?id=645413.652151>.
- 10 Srinivasa R Arikati, Anil Maheshwari, and Christos D Zaroliagis. Efficient computation of implicit representations of sparse graphs. *Discrete Applied Mathematics*, 78(1):1–16, 1997.
- 11 Laszlo Babai, Fan RK Chung, Pál Erdős, Ronald L Graham, and J Spencer. On graphs which contain all sparse graphs. *Ann. Discrete Math*, 12:21–26, 1982.
- 12 Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- 13 Paolo Boldi, Marco Rosa, Massimo Santini, and Sebastiano Vigna. Layered label propagation: A multiresolution coordinate-free ordering for compressing social networks. In *Proceedings of the 20th international conference on World Wide Web*, pages 587–596. ACM, 2011.
- 14 Paolo Boldi and Sebastiano Vigna. The webgraph framework I: compression techniques. In *Proceedings of the 13th international conference on World Wide Web*, pages 595–602. ACM, 2004.
- 15 Béla Bollobás, Oliver Riordan, Joel Spencer, and Gábor E. Tusnády. The degree sequence of a scale-free random graph process. *Random Struct. Algorithms*, 18(3):279–290, 2001. doi:10.1002/rsa.1009.
- 16 Pawek Brach, Marek Cygan, Jakub Lacki, and Piotr Sankowski. Algorithmic complexity of power law networks. In *To appear in Proceedings of the thirty third annual ACM-SIAM symposium on Discrete algorithms, SODA'16*, 2014. URL: <http://arxiv.org/abs/1507.02426>.
- 17 Arthur Brady and Lenore J Cowen. Compact routing on power law graphs with additive stretch. In *ALLENEX*, volume 6, pages 119–128. SIAM, 2006.
- 18 Sonja Buchegger, Doris Schiöberg, Le-Hung Vu, and Anwitaman Datta. Peerson: P2p social networking: early experiences and insights. In *Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*, pages 46–52. ACM, 2009.
- 19 Kenneth L Calvert, Matthew B Doar, and Ellen W Zegura. Modeling internet topology. *Communications Magazine, IEEE*, 35(6):160–163, 1997.
- 20 Saverio Caminiti, Irene Finocchi, and Rossella Petreschi. Engineering tree labeling schemes: A case study on least common ancestors. In *Algorithms-ESA 2008*, pages 234–245. Springer, 2008.
- 21 Wei Chen, Christian Sommer, Shang-Hua Teng, and Yajun Wang. A compact routing scheme and approximate distance oracle for power-law graphs. *ACM Transactions on Algorithms*, 9(1):4, 2012.
- 22 Fan Chung and Linyuan Lu. The average distance in a random graph with given expected degrees. *Internet Mathematics*, 1(1):91–113, 2004.

- 23 Fan RK Chung and Linyuan Lu. *Complex Graphs and Networks*, volume 107. American mathematical society Providence, 2006.
- 24 Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
- 25 Edith Cohen, Haim Kaplan, and Tova Milo. Labeling dynamic xml trees. *SIAM Journal on Computing*, 39(5):2048–2074, 2010.
- 26 Thomas H. Cormen, Clifford Stein, Ronald L. Rivest, and Charles E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, 2nd edition, 2001.
- 27 Søren Dahlgaard, Mathias Bæk Tejs Knudsen, and Noy Rotbart. Dynamic and multi-functional labeling schemes. In *Algorithms and Computation*, pages 141–153. Springer, 2014.
- 28 YH Eom, S Fortunato, and Matjaz Perc. Characterizing and modeling citation dynamics. *PLoS ONE*, 6(9):e24926, 2011.
- 29 Johannes Fischer. Short labels for lowest common ancestors in trees. In *Algorithms-ESA 2009*, pages 752–763. Springer, 2009.
- 30 Cyril Gavoille and Arnaud Labourel. Shorter implicit representation for planar graphs and bounded treewidth graphs. In *Algorithms-ESA 2007*, pages 582–593. Springer, 2007.
- 31 Cyril Gavoille, David Peleg, Stéphane Pérennes, and Ran Raz. Distance labeling in graphs. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms, SODA'01*, pages 210–219, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=365411.365447>.
- 32 Cyril Gavoille, David Peleg, Stéphane Pérennes, and Ran Raz. Distance labeling in graphs. *Journal of Algorithms*, 53:85–112, 2004.
- 33 Pawel Gawrychowski, Adrian Kosowski, and Przemyslaw Uznanski. Even simpler distance labeling for (sparse) graphs. *CoRR*, abs/1507.06240, 2015. URL: <http://arxiv.org/abs/1507.06240>.
- 34 Gaurav Goel and Jens Gustedt. Bounded arboricity to determine the local structure of sparse graphs. In *Graph-Theoretic Concepts in Computer Science*, pages 159–167. Springer, 2006.
- 35 Joseph E Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. Power-graph: Distributed graph-parallel computation on natural graphs. *OSDI*, 12(1):2, 2012.
- 36 Sampath Kannan, Moni Naor, and Steven Rudich. Implicit representation of graphs. In *SIAM Journal On Discrete Mathematics*, pages 334–343, 1992.
- 37 Michal Katz, Nir A Katz, Amos Korman, and David Peleg. Labeling schemes for flow and connectivity. *SIAM Journal on Computing*, 34(1):23–40, 2004.
- 38 Amos Korman. General compact labeling schemes for dynamic trees. *Distributed Computing*, 20(3):179–193, 2007.
- 39 Amos Korman and Shay Kutten. Labeling schemes with queries. In *Structural Information and Communication Complexity*, pages 109–123. Springer, 2007.
- 40 Amos Korman and David Peleg. Compact separator decompositions in dynamic trees and applications to labeling schemes. In *Distributed Computing*, pages 313–327. Springer, 2007.
- 41 Amos Korman and David Peleg. Labeling schemes for weighted dynamic trees. *Inf. Comput.*, 205(12):1721–1740, December 2007. doi:10.1016/j.ic.2007.08.004.
- 42 Łukasz Kowalik. Approximation scheme for lowest outdegree orientation and graph density measures. In *Algorithms and computation*, pages 557–566. Springer, 2006.
- 43 Dmitri Krioukov, Kevin Fall, and Xiaowei Yang. Compact routing on internet-like graphs. In *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1. IEEE, 2004.

- 44 Yucheng Low, Danny Bickson, Joseph Gonzalez, Carlos Guestrin, Aapo Kyrola, and Joseph M Hellerstein. Distributed graphlab: a framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment*, 5(8):716–727, 2012.
- 45 Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 135–146. ACM, 2010.
- 46 Robert Ryan McCune, Tim Weninger, and Greg Madey. Thinking like a vertex: a survey of vertex-centric frameworks for large-scale distributed graph processing. *CoRR*, abs/1507.04405, 2015. URL: <http://arxiv.org/abs/1507.04405>.
- 47 Michael Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1(2):226–251, 2004.
- 48 JW Moon. On minimal n-universal graphs. *Proceedings of the Glasgow Mathematical Association*, 7(01):32–33, 1965.
- 49 Noy Rotbart, Marcos Vaz Salles, and Iasonas Zotos. An evaluation of dynamic labeling schemes for tree networks. In *Experimental Algorithms*, pages 199–210. Springer, 2014.
- 50 Georgos Siganos, Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. Power laws and the as-level internet topology. *IEEE/ACM Trans. Netw.*, 11(4):514–524, 2003. doi:10.1109/TNET.2003.815300.
- 51 Jeremy P Spinrad. *Efficient graph representations*. American mathematical society, 2003.
- 52 Isabelle Stanton and Gabriel Kliot. Streaming graph partitioning for large distributed graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1222–1230. ACM, 2012.
- 53 Bernard M Waxman. Routing of multipoint connections. *Selected Areas in Communications, IEEE Journal on*, 6(9):1617–1622, 1988.
- 54 Cong Xie, Ling Yan, Wu-Jun Li, and Zhihua Zhang. Distributed power-law graph computing: Theoretical and empirical analysis. In *Advances in Neural Information Processing Systems*, pages 1673–1681, 2014.