Models and Algorithms for Graph Watermarking

David Eppstein[1]   Michael T. Goodrich[1]   Jenny Lam[2]
Nil Mamano[1]   Michael Mitzenmacher[3]   Manuel Torres[1]

[1]Dept. of Computer Science, University of California, Irvine, CA USA
[2]Dept. of Computer Science, San José State University, San José, CA USA
[3]Dept. of Computer Science, Harvard University, Cambridge, MA USA

Abstract

We introduce models and algorithmic foundations for graph watermarking. Our frameworks include security definitions and proofs, as well as characterizations when graph watermarking is algorithmically feasible, in spite of the fact that the general problem is NP-complete by simple reductions from the subgraph isomorphism or graph edit distance problems. In the digital watermarking of many types of files, an implicit step in the recovery of a watermark is the mapping of individual pieces of data, such as image pixels or movie frames, from one object to another. In graphs, this step corresponds to approximately matching vertices of one graph to another based on graph invariants such as vertex degree. Our approach is based on characterizing the feasibility of graph watermarking in terms of keygen, marking, and identification functions defined over graph families with known distributions. We demonstrate the strength of this approach with exemplary watermarking schemes for two random graph models, the classic Erdős-Rényi model and a random power-law graph model, both of which are used to model real-world networks.

1 Introduction

In the classic media watermarking problem, we are given a digital representation, R, for some media object, O, such as a piece of music, a video, or an image, such that there is a rich space, $\mathcal{R}$, of possible representations for O besides R that are all more-or-less equivalent. Informally, a *digital watermarking* scheme for O is a function that maps R and a reasonably short random message, m, to an alternative representation, R', for O in $\mathcal{R}$. The verification of such a marking scheme takes R and a presumably-marked representation, R'' (which was possibly altered by an adversary), along with the set of messages previously used for marking, and it either identifies the message from this set that was assigned to R'' or it indicates a failure. Ideally, it should difficult for an adversary to transform a representation, R' (which he was given), into another representation R'' in $\mathcal{R}$, that causes the identification function to fail. Some example applications of such digital watermarking schemes include steganographic communication and marking digital works for copyright protection (e.g., see [16, 25, 50]).

With respect to digital representations of media objects that are intended to be rendered for human performances, such as music, videos, and images, there is a well-established literature on digital watermarking schemes and even well-developed models for such schemes (e.g., see Hopper *et al.* [24]). Typically, such watermarking schemes take advantage of the fact that rendered works have many possible representations with almost imperceptibly different renderings from the perspective of a human viewer or listener.

In this paper, we are inspired by recent systems work on *graph watermarking* by Zhao *et al.* [55,56], who propose a digital watermarking scheme for graphs, such as social networks, protein-interaction graphs, etc., which are to be used for commercial, entertainment, or scientific purposes. This work by Zhao *et al.* presents a system and experimental results for their particular method for performing graph watermarking, but it is lacking in formal security and algorithmic foundations. For example, Zhao *et al.* do not provide formal proofs for circumstances under which graph watermarking is undetectable or when it is computationally feasible. Thus, as complementary work to the systems results of Zhao *et al.*, we are interested in the present paper in providing models and algorithms for graph watermarking, in the spirit of the watermarking model provided by Hopper *et al.* [24] for media files. In particular, we are interested in providing a framework for identifying when graph watermarking is secure and computationally feasible.

1.1 Additional Related Work

Under the term "graph watermarking," there is some additional work, although it is not actually for the problem of graph watermarking as we are defining it. For instance, there is a line of research involving software watermarking using graph-theoretic concepts and encodings. In this case, the object being marked is a piece of software and the goal of a "graph watermarking" scheme is to create a graph, G, from a message, m, and then embed G into the control flow of a piece of software, S, to mark S. Examples of such work include pioneering work by Collberg and Thomborson [15], as well as subsequent work by Venkatesan, Vazirani, and Sinha [52] and Collberg *et al.* [14]. (See also Chen *et al.* [7] and Bento *et al.* [4], as well as a survey by Hamilton and Danicic [22].) This work on software watermarking differs from the graph watermarking problem we study in the present paper, however, because in the graph watermarking problem we study an input graph is provided and we want to alter it to add a mark. In the graph-based software watermarking problem, a graph is instead created from a message to have a specific, known structure, such as being a permutation graph, and then that graph is embedded into the control flow of the piece of software.

A line of research that is more related to the graph watermarking problem we study is anonymization and de-anonymization for social networks (e.g., see [3, 23, 26, 37, 43, 53, 57]). One of the closest examples of such prior work is by Backstrom, Dwork, and Kleinberg [3], who show how to introduce a small set of "rogue" vertices into a social network and connect them to each other and to other vertices so that if that same network is approximately replicated in another setting it is easy to match the two copies. Such work differs from graph watermarking, however, because the set of rogue vertices are designed to "stand out" from the rest of the graph rather than "blend in," and it may in some cases be relatively easy for an adversary to identify and remove such rogue vertices. Also, we would ideally prefer graph watermarking schemes that make small changes to the adjacencies of existing vertices rather than mark a graph by introducing new vertices, since in some applications it may not be possible to introduce new vertices into a graph that we wish to watermark. In addition to this work, also of note is work by Narayanan and Shmatikov [43], who study the problem of approximately matching two social networks without marking, as well as the work on Khanna and Zane [28] for watermarking road networks by perterbing vertex positions (which is a marking method outside the scope of our approach).

Our approach to graph watermarking is also necessarily related to the problem of graph isomorphism and its approximation (e.g., see [1, 2, 17, 27, 30, 46]). In the graph isomorphism problem, we are given two n-vertex graphs, G and H, and asked if there is a mapping, χ, of vertices in G to vertices in H such that (v, w) is an edge in G if and only if $(\chi(v), \chi(w))$ is an edge in H. While the graph isomorphism problem is "famous" for having an uncertain, but unlikely [1], with respect to being NP-complete, extensions to subgraph isomorphism and graph edit distance are known to be

NP-complete (e.g., see [20]).

There is, of course, also prior work on digital watermarking in general. For background on such work, we refer the interested reader to any of the existing surveys, framework papers, or books (e.g., see [16, 24, 25, 50]).

1.2 Our Results

In this paper, we introduce a general graph watermarking framework that is based on the use of key generation, marking, and identification functions, as well as a hypothetical watermarking security experiment (which would be performed by an adversary). We define these functions in terms of graphs taken over random families of graphs, which allows us to quantify situations in which graph watermarking is provably feasible.

We also provide some graph watermarking schemes as examples of our framework, defined in terms of the classic Erdős-Rényi random-graph model and a random power-law graph model. Our schemes extend and build upon previous results on graph isomorphism for these graph families, which may be of independent interest. In particular, we design simple marking schemes for these random graph families based on simple edge-flipping strategies involving high- and medium-degree vertices. Analyzing the correctness of our schemes is quite nontrivial, however, and our analysis and proofs involve intricate probabilistic arguments. We provide an analysis of our scheme against adversaries that can themselves flip edges in order to defeat our mark identification algorithms. In addition, we provide experimental validation of our algorithms, showing that our edge-flipping scheme can succeed for a graph without specific knowledge of the parameters of its deriving graph family. We also conducted experiments to fit real-world networks to the random power-law graph model, which gave results that showed that the model was generally a good fit for the networks tested but the learned values did not fall into the range needed for our scheme.

2 Our Watermarking Framework

We begin by presenting a general framework for graph watermarking, which differs from the general model of Hopper *et al.* [24], but is similar in spirit.

Suppose we are given an undirected graph, $G = (V, E)$, that we wish to mark. To define the security of a watermarking scheme for G, G must come from a family of graphs with some degree of entropy [56]. We formalize this by assuming a probability distribution $\mathcal{D}$ over the family $\mathcal{G}$ of graphs from which G is taken.

Definition 1. A *graph watermarking scheme* is a tuple (keygen, mark, identify) over a set, $\mathcal{G}$, of graphs where

- keygen : $\mathbb{N} \times \mathbb{N} \to$ Aux is a private key generation function, such that keygen(ℓ, n) is a list of ℓ (pseudo-)random graph elements, such as vertices and/or vertex pairs, defined over a graph of n vertices. These candidate locations for marking are defined independent of a specific graph; that is, vertices in Aux are identified simply by the numbering from 1 to n. For example, keygen(ℓ, n) could be a small random graph, H, and some random edges to connect H to a larger input graph [56], or keygen(ℓ, n) could be a set of vertex pairs in an input graph that form candidate locations for marking.
- mark : Aux $\times \mathcal{G} \to \mathbb{N} \times \mathcal{G}$ takes a private key z generated by keygen, and a specific graph G from $\mathcal{G}$, and returns a pair, $S = (\text{id}, H)$, such that id is a unique identifier for H and H is the graph obtained by adding the mark determined by id to G in the location determined determined by the private key z. mark is called every time a different marked copy needs to be

produced, with the i-th copy being denoted by $S_i = (\mathsf{id}_i, H_i)$. Therefore, the unique identifiers should be thought of as being generated randomly. To associate a marked graph H_i with the user who receives it, the watermarking scheme can be augmented with a table storing user name and unique identifiers. Alternatively, the identifiers can be generated pseudo-randomly as a hash of a private key provided by the user.

- $\mathsf{identify} : \mathsf{Aux} \times \mathcal{G} \times \mathbb{N}^k \times \mathcal{G} \to \mathbb{N} \cup \{\bot\}$ takes a private key from Aux, the original graph, G, k identifiers of previously-marked copies of G, and a test graph, G', and it returns the identifier, id_i, of the watermarked graph that it is identifying as a match for G'. It may also return $\bot$, as an indication of failure, if it does not identify any of the graphs H_i as a match for G'.

In addition, in order for a watermarking scheme to be effective, we require that *with high probability*[1] over the graphs from $\mathcal{G}$ and k output pairs, $S_1, \dots, S_k$ of $\mathsf{mark}(z, G)$, for any $(\mathsf{id}, G') = S_i$, we have $\mathsf{identify}(z, G, \mathsf{id}_1, \dots, \mathsf{id}_k, G') = \mathsf{id}$.

Algorithm 1 shows a hypothetical security experiment for a watermarking scheme with respect to an adversary, $A : \mathcal{G} \to \mathcal{G}$, who is trying to defeat the scheme. Intuitively, in the hypothetical experiment, we generate a key z, choose a graph G, from family $\mathcal{G}$ according to distribution $\mathcal{D}$ (as discussed above), and then generate k marked graphs according to our scheme (for some set of k messages). Next, we randomly choose one of the marked graphs, G', and communicate it to an adversary. The adversary then outputs a graph G_A that is similar to G' where his goal is to cause our identification algorithm to fail on G_A.

Algorithm 1: Hypothetical Watermarking Security Experiment

$\mathsf{experiment}(A, k, \ell, n)$:
    1. $z \leftarrow \mathsf{keygen}(\ell, n)$
    2. $G \leftarrow_{\mathcal{D}} \mathcal{G}$
    3. $S_i \leftarrow \mathsf{mark}(z, G)$, for $i = 1, \dots, k$
    4. randomly choose $S_i = (\mathsf{id}, G')$ from $\{S_1, \dots, S_k\}$
    5. $G_A \leftarrow A(G')$

In order to characterize differences between graphs, we assume a similarity measure $\mathsf{dist} : \mathcal{G} \times \mathcal{G} \to \mathbb{R}$, defining the distance between graphs in family $\mathcal{G}$. We also include a similarity threshold θ, that defines the advantage of an adversary performing the experiment in Algorithm 1. Specifically, the *advantage* of an adversary, $A : \mathcal{G} \to \mathcal{G}$ who is trying to defeat our watermarking scheme is

$$\mathbb{P}\left[\mathsf{dist}(G, G_A) < \theta \text{ and } \mathsf{identify}(z, G, \mathsf{id}_1, \dots, \mathsf{id}_k, G_A) \neq \mathsf{id}\right].$$

The watermarking scheme is $(\mathcal{D}, \mathsf{dist}, \theta, k, \ell)$-secure against adversary A if the similarity threshold is θ and A's advantage is *polynomially negligible* (i.e., is $O(n^{-a})$ for some $a > 0$).

Examples of adversaries could include the following:

- *Arbitrary edge-flipping adversary*: a malicious adversary who can arbitrarily flip edges in the graph. That is, the adversary adds an edge if it is not already there, and removes it otherwise.
- *Random edge-flipping adversary*: an adversary who independently flips each edge with a given probability.
- *Arbitrary adversary*: a malicious adversary who can arbitrarily add and/or remove vertices and flip edges in the graph.
- *Random adversary*: an adversary who independently adds and/or removes vertices with a given probability and independently flips each edge with a given probability.

---
[1]Or "*whp*," that is, with probability at least $1 - O(n^{-a})$, for some $a > 0$.

4

One could also imagine other types of adversaries, as well, such as a random adversary who is limited in terms of the numbers or types of edges or vertices that he can change.

2.1 Random graph models

As defined above, a graph watermarking scheme requires that graphs to be marked come from some distribution. In this paper, we consider two families of random graphs—the classic Erdős-Rényi model and a random power-law graph model—which should capture large classes of applications where graph watermarking would be of interest.

Definition 2 (The Erdős-Rényi model). A *random graph* $G(n, p)$ is a graph with n vertices, where each of the $\binom{n}{2}$ possible edges appears in the graph independently with probability p.

Definition 3 (The random power-law graph model, §5.3 of [9]). Given a sequence $\mathbf{w} = (w_1, w_2, \ldots, w_n)$, such that $\max_i w_i^2 < \sum_k w_k$, the *general random graph* $G(\mathbf{w})$ is defined by labeling the vertices 1 through n and choosing each edge (i, j) independently from the others with probability $p[i, j] = \rho w_i w_j$, where $\rho = 1/\sum_j w_j$.

We define a *random power-law graph* $G(\mathbf{w}^\gamma)$ parameterized by the maximum degree m and average degree w. Let $w_i = ci^{-1/(\gamma-1)}$ for values of i in the range between i_0 and $i_0 + n$, where

$$c = \frac{\gamma - 2}{\gamma - 1} w n^{\frac{1}{\gamma-1}}, \qquad i_0 = n \left(\frac{w(\gamma - 2)}{m(\gamma - 1)} \right)^{\gamma - 1}. \tag{1}$$

This definition implies that each edge (i, j) appears with probability

$$P[i, j] = K_0 \left(n^{\gamma-3} ij \right)^{-\frac{1}{\gamma-1}}, \quad \text{where } K_0 \stackrel{\text{def}}{=} \left(\frac{\gamma - 2}{\gamma - 1} \right)^2 w. \tag{2}$$

As we show in the following proposition, this model does indeed have a power-law degree distribution.

Proposition 4. In the random power-law graph $G(\mathbf{w}^\gamma)$, the expected number of vertices with degree k is between Cn/k^γ and $Cn/(k + 1)^\gamma$ where $C = (w(\gamma - 2))^{\gamma-1} / (\gamma - 1)^{\gamma-2}$.

Proof. The function $i(k) = (c/k)^{\gamma-1}$ relating the index of a vertex to its expected degree $k = w_i$ is convex and decreasing. By the mean value theorem, the number Δi of indices j such that $k \leq w_j < k + 1$ satisfies

$$|i'(k)| \leq \frac{\Delta i}{(k + 1) - k} = \Delta i \leq |i'(k + 1)|.$$

Now the derivative of $i(k)$ is $-Cn/k^\gamma$. Noting that Δi is the expected number of vertices of degree k, the result is proven. $\square$

2.2 Graph watermarking algorithms

We discuss some instantiations of the graph watermarking framework defined above. Unlike previous watermarking or de-anonymization schemes that add vertices [3, 56], we describe an effective and efficient scheme based solely on edge flipping. Such an approach would be especially useful for applications where it could be infeasible to add vertices as part of a watermark.

Our scheme does not require adding labels to the vertices or additional objects stored in the graph for identification purposes. Instead, we simply rely on the structural properties of graphs for the purposes of marking. In particular, we focus on the use of vertex degrees, that is, the number of

edges incident on each vertex. We identify high and medium degree vertices as candidates for finding edges that can be flipped in the course of marking. The specific degree thresholds for what we mean by "high-degree" and "medium-degree" depend on the graph family, however, so we postpone defining these notions precisely until our analysis sections.

Algorithms providing an example implementation of our graph watermarking scheme are shown in Algorithm 2. The keygen algorithm randomly selects a set of candidate vertex pairs for flipping, from among the high- and medium-degree vertices, with no vertex being incident to more than a parameter t of candidate pairs. We introduce a procedure, label(G), which labels high-degree vertices by their degree ranks and each medium-degree vertex, w, by a bit vector identifying its high-degree adjacencies. This bit vector has a bit for each high-degree vertex, which is 1 for neighbors of w and 0 for non-neighbors. The algorithm mark(z, G), takes a random set of candidate edges and a graph, G, and it flips the corresponding edges in G according to a resampling of the edges using the distribution $\mathcal{D}$. The algorithm, approximate-isomorphism(G, H), returns a mapping of the high- and medium-degree vertices in G to matching high- and medium-degree vertices in H, if possible. The algorithm, identify$(z, G, \mathsf{id}_1, \ldots, \mathsf{id}_k, H)$, uses the approximate isomorphism algorithm to match up high- and medium-degree vertices in G and H, and then it extracts the bit-vector from this matching using z.

As mentioned above, we also need a notion of distance for graphs. We use two different such notions. The first is the graph edit distance, which is the minimum number of edges needed to flip to go from one graph to another. The second is vertex distance, which intuitively is an edge-flipping metric localized to vertices.

Definition 5 (Graph distances). Let $\mathcal{G}$ be the set of graphs on n vertices. If $G, H \in \mathcal{G}$, define Π as the set of bijections between the vertex sets $V(G)$ and $V(H)$. Define the *graph edit distance* $\mathsf{dist}_e : \mathcal{G} \times \mathcal{G} \to \mathbb{N}$ as

$$\mathsf{dist}_e(G, H) = \min_{\pi \in \Pi} |E(G) \oplus_\pi E(H)|,$$

where $\oplus_\pi$ is the symmetric difference of the two edge sets under correspondence π. Define the *vertex distance* $\mathsf{dist}_v : \mathcal{G} \times \mathcal{G} \to \mathbb{N}$ as

$$\mathsf{dist}_v(G, H) = \min_{\pi \in \Pi} \max_{v \in V(G)} |E(v) \oplus_\pi E(\pi(v))|,$$

where $E(v)$ is the set of edges incident to v.

3 Identifying High- and Medium-Degree Vertices

We begin analyzing our proposed graph watermarking scheme by showing how high- and medium-degree vertices can be identified under our two random graph distributions. We begin with some technical results related to graph isomorphism that form the basis of our watermarking approach, with the goal of determining the conditions under which a vertex of a random graph can be identified with high probability, either by its degree (if the degree is high) or by its set of high-degree neighbors (if it has medium degree). We ignore low-degree vertices: their information content and distinguishability are low, and they are not used by our example scheme. Because our results on vertex identifiability are used in our graph watermarking scheme, we also determine how robust these identifications are, based on how well-separated the vertices are by their degrees.

We first find a threshold number k such that the k vertices with highest degree are likely to have distinct and well-separated degree values. We call these k vertices the *high-degree* vertices. Next, we look among the remaining vertices for those that are well-separated in terms of their high-degree

Algorithm 2: Watermarking scheme for random graphs.

---

t: the maximum number of flipped edges that can be adjacent to the same vertex.

keygen(ℓ, n):
1. Let x denote the total number of high- and medium-degree vertices
2. $X = \{(u, v) \mid 1 \leq u < v \leq x\}$
3. Let z be a list of ℓ pairs randomly sampled (without replacement) from X such that no end vertex appears more than t times
4. return z

label(G):
1. sort the vertices in decreasing order by degree and identify the high- and medium-degree vertices
2. if the degrees of high-degree vertices are not unique, return failure
3. label each high-degree vertex with its position in the vertex sequence
4. label each medium-degree vertex with a bit vector encoding its high-degree adjacencies
5. if the bit vectors are not unique, return failure
6. otherwise, return the labelings

mark(z, G):
1. $S = \varnothing$
2. V is the set of high- and medium-degree vertices of G, sorted lexicographically by their labels given by $L = \mathsf{label}(G)$
3. generate an ℓ-bit string id where each bit i is independently set to 1 with probability $p_{z[i]}$, where $p_{z[i]}$ is the probability of the edge $z[i]$ in $\mathcal{D}$
4. let H be a copy of G
5. for j from 1 to ℓ:
6. $(u, v) = z[j]$
7.   if id$[j]$ is 1:
8.   insert edge $(V[u], V[v])$ in H
9.   else:
10.   remove edge $(V[u], V[v])$ from H
11. return (id, H)

approximate-isomorphism(G, H):
1. call $\mathsf{label}(G)$ and $\mathsf{label}(H)$, returning failure if either of these fail.
2. match each of G's high-degree vertices with the vertex in H with the same label.
3. match each of G's medium-degree vertices with the vertex in H whose label is closest in Hamming distance.
4. if H has a vertex that is matched more than once, return failure.
5. otherwise, return the (partial) vertex assignments between G and H.

identify($z, G, \mathsf{id}_1, \ldots, \mathsf{id}_k, H$):
1. find an $\mathsf{approximate\text{-}isomorphism}(G, H)$, returning $\perp$ if failure occurred at any step.
2. V is the set of high- and medium-degree vertices of G, sorted lexicographically by their labels given by $L = \mathsf{label}(G)$
3. V' is the set of vertices of H identified as corresponding to those in V, in that same order.
4. id is an empty bit string
5. for (u, v) in z (from left to right):
6. $b = 1$ iff there is an edge between $V'[u]$ and $V'[v]$ in H.
7.   append b to id
8. return among the id_i's the one closest to id

---

neighbors. Specifically, the (high-degree) *neighborhood distance* between two vertices is the number of high-degree vertices which are connected to exactly one of the two vertices. Note that we will omit the term "high-degree" in "high-degree neighborhood distance" from now on, as it will always be implied.

In the Erdős-Rényi model, we show that all vertices that are not high-degree nevertheless have well-separated high-degree neighborhoods whp. In the random power-law graph model, however, there will be many lower-degree vertices whose high-degree neighborhoods cannot be separated. Those that have well-separated high-degree neighborhoods with high probability form the medium-degree vertices, and the rest are the low-degree vertices.

For completeness, we include the following well-known Chernoff concentration bound, which we will refer to time and again.

Lemma 6 (Chernoff inequality [9])**.** Let $X_1, \ldots, X_n$ be independent random variables with

$$\mathbb{P}[X_i = 1] = p_i, \qquad \mathbb{P}[X_i = 0] = 1 - p_i.$$

We consider the sum $X = \sum_{i=1}^n X_i$, with expectation $\mathbb{E}[X] = \sum_{i=1}^n p_i$. Then

$$\mathbb{P}[X \leq \mathbb{E}[X] - \lambda] \leq e^{-\frac{\lambda^2}{2\mathbb{E}[X]}},$$

$$\mathbb{P}[X \geq \mathbb{E}[X] + \lambda] \leq e^{-\frac{\lambda^2}{2\mathbb{E}[X] + \lambda/3}}.$$

3.1 Vertex separation in the Erdős-Rényi model

Let us next consider vertex separation results for the classic Erdős-Rényi random-graph model. Recall that in this model, each edge is chosen independently with probability p.

Definition 7. Index vertices in non-increasing order by degree. Let d_i represent the i-th highest degree in the graph. Given $h = O(n)$, we say that a vertex is *high-degree* with respect to d_h if it has degree at least d_h. Otherwise, we say that the vertex is *medium-degree*. We just say high-degree when the value of h is understood from context.

Note that in this random-graph model, there are no low-degree vertices.

Definition 8. A graph is (d, d')-*separated* if all high-degree vertices differ in their degree by at least d and all medium-degree vertices are neighborhood distance d' apart.

Note: this definition depends on how high-degree or medium-degree vertices are defined and will therefore be different for the random power-law graph model.

Lemma 9 (Extension of Theorem 3.15 in [5])**.** Suppose $m = o(pqn/\log n)^{1/4}$, $m \to \infty$, and $\alpha(n) \to 0$. Then with probability

$$1 - m\alpha(n) - 1/\left[m\left(\log(n/m)\right)^2\right],$$

$G(n, p)$ is such that

$$d_i - d_{i+1} \geq \frac{\alpha(n)}{m^2}\left(\frac{pqn}{\log n}\right)^{1/2} \quad \text{for every } i < m,$$

where $q = 1 - p$.

Proof. We quantify and extend the probability analysis of a proof from [5]. Let

$$K = pn + (x - \varepsilon)(pnq)^{1/2}, \quad \varepsilon = (\log(n/m))^{1/2}.$$

The event of the result fails if $d_m < K$ or if there is $i < m$ such that $d_i - d_{i+1} < K$.

The statement of theorem 3.12 of [5] still holds when the words "a.e. G_p satisfies" are replaced by "G_p satisfies with probability greater than $1 - 1/\omega(n)^2$". This can be seen directly from the part of the proof where Chebychev's inequality is applied.

By this result, the probability that $d_m < K$ is $1/\left[m\left(\log(n/m)\right)^2\right]$. The probability that $d_i - d_{i+1} < K$ for a given $i < m$ is $O(\alpha(n))$. $\qquad \square$

Lemma 10 (Vertex separation in the Erdős-Rényi model)**.** Let $0 < \varepsilon < 1/9$, $d \geq 3$, $C \geq 3$, $h = n^{(1-\varepsilon)/8}$. Suppose $0 < p = p(n) \leq \frac{1}{2}$ is such that $p = \omega(n^{-\varepsilon} \log n)$. Then $G(n,p)$ is $(d, C \log n)$-separated with probability $1 - O(n^{-(1-\varepsilon)/8})$.

Proof. We prove the theorem with probability at least $1 - (d+2)n^{-(1-\varepsilon)/8}$. Let $\alpha(n) = dn^{-(1-\varepsilon)/4}$ and $m = h$. By Lemma 9, the probability that $d_i - d_{i+1} < d$ for some $i < h$ is at most $(d+1)n^{-(1-\varepsilon)/8}$.

Let X_{ij} be the expected neighborhood distance between two vertices $i, j \geq h$. We have

$$\mathbb{E}\left[X_{ij}\right] = mp(1-p) \geq (2C+1)n^{(1-9\varepsilon)/8} \log n \geq (2C+1) \log n,$$

so that, if $d' = C \log n$,

$$\frac{(\mathbb{E}\left[X_{ij}\right] - d')^2}{\mathbb{E}\left[X_{ij}\right]} \geq \mathbb{E}\left[X_{ij}\right] - 2d' \geq C \log n.$$

Since the high-degree vertices are separated by more than two degrees, the fact that they are high-degree vertices is independent of whether they are neighbors of i and j. Consequently, we can apply a Chernoff bound (Lemma 6.) Then, by the union bound, the probability that $X_{ij} < C \log n$ for some medium-degree i, j is less than $n^{-C+2} \leq n^{-(1-\varepsilon)/8}$. $\qquad \square$

Thus, high-degree vertices are well-separated with high probability in the Erdős-Rényi model, and the medium-degree vertices are distinguished with high probability by their high-degree neighborhoods.

3.2 Vertex separation in the random power-law graph model

We next study vertex separation for a random power-law graph model, which can match the degree distributions of many graphs that naturally occur in social networking and science. For more information about power-law graphs and their applications, see e.g. [6, 40, 44].

In the random power-law graph model, vertex indices are used to define edge weights and therefore do not necessarily start at 1. The lowest index that corresponds to an actual vertex is denoted i_0. So vertex indices range from i_0 to $i_0 + n$. Additionally, there are two other special indices i_H and i_M, which we define in this section, that separate the three classes of vertices.

Definition 11. The vertices ranging from i_0 to i_H are the *high-degree* vertices, those that range from $i_H + 1$ to i_M are the *medium-degree* vertices, and those beyond i_M are the *low-degree* vertices.

In this model, the value of i_0 is constrained by the requirement that $P[i_0, i_0] < 1$. When $\gamma \geq 3$, this constraint is not actually restrictive. However, when $\gamma < 3$, i_0 must be asymptotically greater than $n^{-(\gamma-3)/2}$. The constraints on i_0 also constrain the value of the maximal and average degree of the graph.

We define i_H and i_M to be independent of i_0, but dependent on parameters that control the amount and probability of separation at each level. The constraints that $i_0 < i_H$ and $i_H < i_M$ translate into corresponding restrictions on the valid values of γ, namely that $\gamma > 5/2$ and $\gamma < 3$. We define i_H in the following lemma.

Lemma 12 (Separation of high-degree vertices). *In the $G(\mathbf{w}^\gamma)$ model, let $\delta_i = |w_{i+1} - w_i|/2$. Then,*

$$\frac{c}{2(\gamma - 1)}(i + 1)^{-\frac{\gamma}{\gamma-1}} \le \delta_i \le \frac{c}{2(\gamma - 1)}i^{-\frac{\gamma}{\gamma-1}}. \tag{3}$$

Moreover, for all ε_1 satisfying $0 < \varepsilon_1 \le 1$ and $C_1 > 0$, the probability that

$$|\deg(i) - w_i| < \varepsilon_1\delta_i \qquad \text{for all } i \le i_H \stackrel{\text{def}}{=} \left(\frac{c\varepsilon_1^2}{16(\gamma-1)^2 C_1 \log n}\right)^{\frac{\gamma-1}{2\gamma-1}}$$

is at least $1 - n^{-C_1}$.

Proof. The first statement follows from the fact that w_i is a convex function of i and from taking its derivative at i and $i + 1$.

For the second statement, let $C > 0$ and let $i'_H \stackrel{\text{def}}{=} \left(\frac{c\varepsilon_1^2}{8(\gamma-1)^2 C \log n}\right)^{\frac{\gamma-1}{2\gamma-1}}$. We will show that if $i \le i'_H$, then

$$\mathbb{P}\left[|\deg(i) - w_i| \ge \varepsilon_1\delta_i\right] < n^{-C}. \tag{4}$$

Now we choose C such that $C_1 + \log i_H / \log n < C \le 2C_1$. The inequality $C \le 2C_1$ implies that $i_H \le i'_H$ and (4) holds for all $i \le i_H$. By the union bound applied to (4)

$$\mathbb{P}\left[\exists i \le i_H, |\deg(i) - w_i| \ge \varepsilon_1\delta_i\right] \le i_H n^{-C}.$$

Since $C_1 + \log i_H / \log n < C$, the right hand side is bounded above by n^{-C_1}. This proves the result.

Now, we prove (4). Clearly, since $\delta_i = (w_i - w_{i+1})/2$, we have that $w_i \ge \delta_i$. So if $\varepsilon_1 \le 1$ and $\lambda_i = \varepsilon_1\delta_i$, then $w_i \ge \lambda_i/3$. This implies that

$$\frac{\lambda_i^2}{w_i + \lambda_i/3} \ge \frac{\lambda_i^2}{2w_i} \ge \frac{c\varepsilon_1^2}{8(\gamma-1)^2}i^{-\frac{2\gamma-1}{\gamma-1}},$$

where the second inequality follows from (3) and the definition of w_i given in Definition 3. If $i \le i'_H$, the right hand side is lower-bounded by $C \log n$. The result follows by applying a Chernoff bound (Lemma 6). $\qquad\square$

For simplicity, we often use the following observation.

Observation 13. *Rewriting i_H to show its dependence on n, we have*

$$i_H(\varepsilon_1, C_1) = K_1(\varepsilon_1, C_1)\, n^{\frac{1}{2\gamma-1}} (\log n)^{-\frac{\gamma-1}{2\gamma-1}}, \quad K_1(\varepsilon_1, C_1) \stackrel{\text{def}}{=} \left(\frac{\gamma - 2}{(\gamma - 1)^3}\frac{w\varepsilon_1^2}{16C_1}\right)^{\frac{\gamma-1}{2\gamma-1}}. \tag{5}$$

For the graph model to make sense, the high-degree threshold must be asymptotically greater than the lowest index. In other words, we must have that $i_0 = o(i_H)$. Since $i_0 = \Omega(n^{-(\gamma-3)/2})$, this implies that $\gamma > 5/2$.

We next define i_M, the degree threshold for medium-degree vertices, in the following lemma.

Lemma 14 (Separation of medium-degree vertices)**.** Let K_0 be defined as in Definition 3, $K_1(\varepsilon_1, C_1)$ be defined as in (5), and

$$K_2(\varepsilon_1, C_1, \varepsilon_2, C_2) \stackrel{\text{def}}{=} \frac{K_0^{\gamma-1} K_1^{\gamma-2}(\varepsilon_1, C_1)}{(C_2 + 2\Gamma + 2\log(K_0^{\gamma-1} K_1^{\gamma-2}(\varepsilon_1, C_1)) + 2\varepsilon_2)^{\gamma-1}}. \tag{6}$$

Let X_{ij} denote the neighborhood distance between two vertices i and j in $G(\mathbf{w}^\gamma)$. If $5/2 < \gamma < 3$, for every $\varepsilon_2 > 0$ and $C_2 > 0$, the probability that

$$X_{ij} > \varepsilon_2 \log n, \quad \text{for all } i_H \leq i, j \leq i_M$$

where

$$i_M(\varepsilon_1, C_1, \varepsilon_2, C_2) \stackrel{\text{def}}{=} K_2(\varepsilon_1, C_1, \varepsilon_2, C_2) \, n^\Gamma \, (\log n)^{-\frac{3(\gamma-1)^2}{2\gamma-1}}, \quad \Gamma \stackrel{\text{def}}{=} -\frac{2\gamma^2 - 8\gamma + 5}{2\gamma - 1}, \tag{7}$$

is at least $1 - n^{-C_2}$ for sufficiently large n.

Proof. Let $C > 0$ and let

$$i'_M \stackrel{\text{def}}{=} \left(\frac{C_2 + 2\Gamma + 2\log(K_0^{\gamma-1} K_1^{\gamma-2}) + 2\varepsilon_2}{C + 2\varepsilon_2} \right)^{\gamma-1} i_M.$$

We claim that if $i_H \leq i, j \leq i'_M$, then

$$\mathbb{P}\left[X_{ij} \leq \varepsilon_2 \log n\right] \leq n^{-C}. \tag{8}$$

If we choose $C = C_2 + 2\Gamma + 2\log K_0^{\gamma-1} K_1^{\gamma-2}$, we have that $i_M = i'_M$, so that (8) applies to all i, j such that $i, j \leq i_M$. Moreover, since

$$i_M \leq K_0^{\gamma-1} K_1^{\gamma-2} n^\Gamma \leq n^{\log(K_0^{\gamma-1} K_1^{\gamma-2})} n^\Gamma,$$

our choice of C implies that $i_M^2 \, n^{-C} \leq n^{-C_2}$. By applying the union bound to (8), we have

$$\mathbb{P}\left[\exists i, j \text{ s.t. } i_H \leq i, j \leq i_M, \ X_{ij} \leq \varepsilon_2 \log n\right] \leq i_M^2 n^{-C} \leq n^{-C_2},$$

which establishes the lemma.

Let us now prove the claim. Observe that X_{ij} is the sum over the high-degree vertices k, of indicator variables X_{ij}^k for the event that vertex k is connected to exactly one of the vertices i and j. It i For fixed i and j, these are independent random variables. Therefore, we can apply a Chernoff bound. The probability that $X_{ij}^k = 1$ is

$$P[i, k](1 - P[j, k]) + P[j, k](1 - P[i, k]) \geq 2P[i_M, i_H](1 - P[i_0, i_H]).$$

Since $P[i_0, i_H] \to 0$, for sufficiently large n, this expression is bounded below by $P[i_M, i_H]$, and

$$\mathbb{E}\left[X_{ij}\right] \geq i_H P[i_M, i_H] \geq (C + 2\varepsilon_2) \log n,$$

by (2), (5) and (7), as can be shown by a straightforward but lengthy computation. Let $d = \varepsilon_2 \log n$. This implies that

$$\frac{(\mathbb{E}\left[X_{ij}\right] - d)^2}{\mathbb{E}\left[X_{ij}\right]} \geq \mathbb{E}\left[X_{ij}\right] - 2d \geq C \log n.$$

Therefore, applying the Chernoff bound (Lemma 6) to the X_{ij}^k for fixed i and j and all high-degree vertices k proves the claim. $\square$

Figure 1: Degree breakpoints for the random power-law graph model.

Observation 15. We would have the undesirable situation that $i_M = o(1)$ whenever $\frac{2\gamma^2 - 8\gamma + 5}{2\gamma - 1} > 0$, or equivalently when $\gamma > 2 + \sqrt{3/2} > 3$. In fact, in order for $i_H = o(i_M)$, we must have $\gamma < 3$.

We illustrate the breakpoints for high-, medium-, and low-degree vertices in Fig. 1.

The next lemma summarizes the above discussion and provides the forms of i_H and i_M that we use in our analysis.

Lemma 16 (Vertex separation in the power-law model)**.** Let $5/2 < \gamma < 3$. Fix $\varepsilon > 0, C_1 > 0, C_2 > 0$. Let $i_H = i_H(\varepsilon_1, C_1)$ and $i_M = i_M(\varepsilon_1, C_1, \varepsilon_2, C_2)$ where $\varepsilon_1 = 1$ and $\varepsilon_2 = \varepsilon$. Let

$$d = n^{\frac{1}{2\gamma - 1}} \quad \text{and} \quad d' = \log n.$$

For sufficiently large n, the probability that a graph $G(\mathbf{w}^\gamma)$ is not $(\varepsilon d, \varepsilon d')$-separated is at most $n^{-C_1} + n^{-C_2}$.

Proof. Let δ_i be defined as in Lemma 12. A straightforward computation using (1), (3), and (5) shows that

$$\delta_{i_H} \geq \text{constant} \cdot n^{\frac{1}{2\gamma - 1}} (\log n)^{\frac{\gamma}{2\gamma - 1}}.$$

So for sufficiently large n, we have $\delta_{i_H} \geq 3\varepsilon d/2$. For all $i \leq i_H$, the average degrees w_i of consecutive vertices are at least $3\varepsilon d/2$ apart. So for two high-degree vertices to be within εd of each other, at least one of the two must have degree at least $(3\varepsilon/2 - \varepsilon/2)d$ away from its expected degree. By Lemma 12, the probability that some high-degree vertex i satisfies $|\deg(i) - w_i| > \delta_{i_H}$ is at most n^{-C_1}.

By Lemma 14, the probability that there are two medium-degree vertices with neighborhood distance less than $\varepsilon d'$ is at most n^{-C_2}. $\qquad\square$

Thus, our marking scheme for the random power-law graph model is effective.

4 Adversary Tolerance

In this section, we study the degree to which our exemplary graph watermarking scheme can tolerate an arbitrary edge-flipping adversary. To measure success, we use the notion of security and adversary advantage which are formally defined in 2. We quantify the number of edge flips that can be tolerated under the Erdős-Rényi model and the random power-law graph model.

Theorem 17 (Security against an arbitrary edge-flipping adversary in the Erdős-Rényi model)**.** Let $0 < \varepsilon < 1/9$, $d \geq 3$, $h = n^{(1-\varepsilon)/8}$ and $p \leq 1/2$ such that $p = \omega(n^{-\varepsilon} \log n)$. Let d be sufficiently large so that

$$\varepsilon \frac{d+1}{d-1} < 1. \tag{9}$$

Suppose the similarity measure is the vertex distance dist_v, the similarity threshold is $\theta = d$, we have a number $k = n^C$ of watermarked copies, and their identifiers are generated using $\ell = 8(2C + C')n^\varepsilon$

bits. Suppose also that the identifiers map to sets of edges of a graph constrained by the fact that no more than $t = d$ edges can be incident to any vertex. The watermarking scheme defined in Algorithm 2 is $(G(n,p), \mathsf{dist}_v, \theta, k, \ell)$-secure against any deterministic adversary.

The proof of this theorem relies on two lemmas. Lemma 18 identifies conditions under which a set of bit vectors with bits independently set to 1 is unlikely to have two close bit vectors. Lemma 19 states that a deterministic adversary's ability to guess the location of the watermark is limited. Informally, this is because the watermarked graph was obtained through a random process, so that there are many likely original graphs that could have produced it.

Lemma 18 (Separation of IDs)**.** Consider $k = n^C$ random bit strings of length ℓ, where each bit is independently set to 1, and the i-th bit is 1 with probability q_i satisfying $p \leq q_i \leq 1/2$ for a fixed value p. The probability that at least two of these strings are within Hamming distance $D = 4(2C + C') \log n$ of each other is at most $n^{-C'}$ if $\ell p \geq 2D$.

Proof. The expected distance between two such strings is at least $2\ell p(1-p) \geq \ell p$. Applying Lemma 6 with $\lambda = \ell p/2$, we have that the probability that their Hamming distance is less than $\ell p/2$ is at most $e^{-\ell p/8} \leq n^{-(2C+C')}$. Therefore, the probability that at least two out of k strings are within Hamming distance $D \leq \ell p/2$ of each other is at most $k^2 n^{-(2C+C')} = n^{-C'}$. $\qquad\square$

Lemma 19 (Guessing power of adversary)**.** Consider a complete graph on N vertices, and let r of its edges be red. Let s be a sample of ℓ edges chosen uniformly at random among those that satisfy the constraint that no more than t edges of the sample can be incident to any one vertex. Suppose also that ℓ, N and t are non-decreasing functions of n such that

$$\frac{\ell^{t+1}}{N^{t-1}} \to 0 \text{ as } n \to \infty. \tag{10}$$

For sufficiently large N, the probability that s contains at least $R = 8\ell r/N^2$ red edges is bounded by $4 \exp\left(-12\ell r/(7N^2)\right)$. Moreover, if $\ell r/N^2 \to 0$, then the probability that s contains at least $R = 1$ red edge is bounded by $4 \exp\left(-cN^2/(\ell r)\right)$, for some $c > 0$ and for sufficiently large N.

Proof. In the process of selecting ℓ edges without replacement, let A be the event that the sample contain at least R red edges, and let B be the event that the sample satisfies the degree constraint. The event whose probability we want to bound is equal to

$$\mathbb{P}\left[A|B\right] \leq \frac{\mathbb{P}\left[A\right]}{\mathbb{P}\left[B\right]}.$$

Let us first show that $\mathbb{P}\left[B\right]$ can be lower bounded by a constant. To prove this, we select 2ℓ vertices with replacement uniformly at random, and pair consecutive vertices to obtain ℓ edges. Choosing vertices uniformly in this way will simplify showing that the degree constraint is satisfied. Of course we want to avoid "self-loops", or edges where both end vertices are the same. Let C denote the event that there is a vertex that is incident to more than t edges of the sample. Also, let D denote the event that the sample contains no self-loops and no duplicate edges. Then

$$\mathbb{P}[\bar{B}] = \mathbb{P}\left[C|D\right] \leq \frac{\mathbb{P}\left[C\right]}{\mathbb{P}\left[D\right]}.$$

Now, the probability of encountering a self-loop is $1/N$ and the probability of an edge being a duplicate of another is at most $2/N^2$. Therefore,

$$\mathbb{P}[\bar{D}] \leq \frac{\ell}{N} + \binom{\ell}{2}\frac{2}{N^2} \leq \frac{2\ell}{N}.$$

13

By (10), $\ell/N \to 0$. So $\mathbb{P}[D]$ is bounded away from 0. Moreover, since the edges now consist of pairs of independently chosen vertices, we can approximate the number of edges incident to each vertex by N independent Poisson random variables with parameter $2\ell/N$ thusly:

$$\mathbb{P}[C] \leq N \left(\frac{e^{-2\ell/N}(2\ell e/N)^t}{t^t} \right) (e\sqrt{2\ell}),$$

where the middle factor is a bound on the probability that one Poisson variable is at least t (Theorem 5.4 of [41]), and the last factor is an adjustment factor for this approximation (Corollary 5.9 of [41]). This expression is bounded by a constant factor times the expression on the left-hand side of (10). Consequently, $\mathbb{P}[\bar{B}]$ converges to 0, and for sufficiently large N, $\mathbb{P}[B] \geq 1/2$, as was to be shown.

Now we find an upper bound for $\mathbb{P}[A]$. To do this, we select ℓ edges with replacement uniformly at random. Because ℓ is relatively small when compared to N, it is unlikely that the sample will contain any duplicates. Formally, let E be the event that the sample contains at least R red edges, and F be the event that the sample consists of distinct edges. We have

$$\mathbb{P}[A] = \mathbb{P}[E|F] \leq \frac{\mathbb{P}[E]}{\mathbb{P}[F]}.$$

The probability that two selected edges are the same edge is $1/\binom{N}{2}$. So

$$\mathbb{P}[\bar{F}] \leq \binom{\ell}{2} \Big/ \binom{N}{2} \leq \frac{\ell^2}{2} \frac{4}{N^2} = 2 \left(\frac{\ell}{N} \right)^2.$$

So for large enough N, $\mathbb{P}[F]$ is bounded below by $\frac{1}{2}$.

Finally, we bound $\mathbb{P}[E]$. The expected number X of red edges in this sample is $\mathbb{E}[X] = \ell r/\binom{N}{2}$ which is bounded below by $2\ell r/N^2$ and bounded above by $4\ell r/N^2 = R/2$. So using these bounds and a Chernoff bound (Lemma 6), where we set λ equal to $\mathbb{E}[X]$, we have that

$$\mathbb{P}[E] \leq \exp\left(-\frac{6}{7} \frac{2\ell r}{N^2} \right).$$

If $\ell r/N^2 \to 0$ as $n \to \infty$, set λ equal to $1 - \mathbb{E}[X] = 1 - \Theta(\ell r/N^2)$:

$$\mathbb{P}[E] \leq \exp\left(-\frac{c}{\mathbb{E}[X]} \right)$$

for some constant $c > 0$. Putting it all together, we have that for large enough N, $\mathbb{P}[B] \geq 1/2$ and $\mathbb{P}[A]$ is bounded above by 2 times one of the two bounds for $\mathbb{P}[E]$. This proves the result. $\square$

Theorem 17. An upper bound on the advantage of any deterministic adversary $A : \mathcal{G} \to \mathcal{G}$ on graphs on n vertices is given by the conditional probability

$$\mathbb{P}\left[\mathsf{identify}(z, G, \mathsf{id}_1, \dots, \mathsf{id}_k, G_A) \neq \mathsf{id} | \mathsf{dist}_v(G, G_A) < \theta\right],$$

where the parameters passed to $\mathsf{identify}$ are defined according to the experiment in Algorithm 1. We show that this quantity is polynomially negligible.

For G_A to be successfully identified, it is sufficient for the following three conditions to hold:
1. the original graph $G = G(n, p)$ is $(4d, 4d)$-separated;
2. the Hamming distance between any two id and id' involved in a pair in S is at least $D = 4(2C + C') \log n$;

14

3. *A* changes no edges of the watermark.

These are sufficient conditions because we only test graphs whose vertices had at most d incident edges modified by the adversary, and another d incident edges modified by the watermarking. So for original graphs that are $(4d, 4d)$-separated, the labeling of the vertices can be successfully recovered. Finally, if the adversary does not modify any potential edge that is part of the watermark, the id of the graph is intact and can be recovered from the labeling.

Now, by Lemma 10, the probability that $G(n, p)$ is not $(4d, 4d)$-separated is less than $O(n^{-(1-\varepsilon)/8})$. Moreover, since $\ell p \geq 2D$, by Lemma 18, the probability that there are two identifiers in S that are within D of each other is at most $n^{-C'}$.

Finally, for graphs in which an adversary makes fewer than d modifications per vertex, the total number of edges the adversary can modify is $r \leq dn/2$. Since all vertices are high- and medium-degree vertices in this model, $N = n$. Therefore, $\ell r / N^2 = O(1/n^{(1-\varepsilon)}) \to 0$. Equation (9) guarantees that the hypothesis given by (10) of Lemma 19 is satisfied. Consequently, the probability that A changes one or more adversary edges is $O(\exp[cn^{1-\varepsilon}])$ for some constant c.

This proves that each of the three conditions listed above fails with polynomially negligible probability, which implies that the conditional probability is also polynomially negligible. $\qquad\square$

Theorem 20 (Security against an arbitrary edge-flipping adversary in the random power-law graph model). Let $5/2 < \gamma < 3$, $C > 0$, $i_H = i_H(\varepsilon_1, C_1)$ and $i_M = i_M(\varepsilon_1, C_1, \varepsilon_2, C_2)$ where $\varepsilon_1 = 1$, $\varepsilon_2 = 8(C + 1)$ and $C_1 = C_2 = C$.

Let $p = P[i_M, i_M]$. Suppose the similarity measure is a vector of distances $\mathsf{dist} = (\mathsf{dist}_e, \mathsf{dist}_v)$, that the corresponding similarity threshold is the vector $\theta = (r, \log n)$ where $r = p(i_M)^2/32$ is the maximum number of edges the adversary can flip in total, and $\log n$ the maximum number number of edges it can flip per vertex. Suppose that we have $k = n^{C''}$ watermarked copies of the graph, that we use $\ell = 8(2C'' + C')(\log n)/p$ to watermark a graph.

Suppose also that the identifiers map to sets of edges of a graph constrained by the fact that no more than $t = \log n$ edges can be incident to any vertex. Then the watermarking scheme defined in Algorithm 2 is $(G(\mathbf{w}^\gamma), \mathsf{dist} = (\mathsf{dist}_e, \mathsf{dist}_v), \theta = (r, \log n), k, \ell)$-secure against any deterministic adversary.

Proof. The proof is similar to the proof of Theorem 17. An upper bound on the advantage of any deterministic adversary $A : \mathcal{G} \to \mathcal{G}$ on graphs on n vertices is given by the conditional probability

$$\mathbb{P}\left[\mathsf{identify}(z, G, \mathsf{id}_1, \ldots, \mathsf{id}_k, G_A) \neq \mathsf{id} | \mathsf{dist}(G, G_A) < \theta\right],$$

where the parameters passed to $\mathsf{identify}$ are defined according to the experiment in Algorithm 1. We show that this quantity is polynomially negligible.

For G_A to be successfully identified, it is sufficient for the following three conditions to hold:
1. the original graph $G = G(\mathbf{w}^\gamma)$ is $(4 \log n, 4 \log n)$-separated;
2. the Hamming distance between any two id and id$'$ involved in a pair in S is at least $D = 4(2C'' + C') \log n$;
3. *A* changes fewer than $D/2$ edges of the watermark.

The proof is similar to the proof of Theorem 17. To be able to apply Lemma 19, we need to show that (10) holds with $N = i_M$. Recall that

$$p = K_0 \left(n^{\gamma-3} i_M^2\right)^{-\frac{1}{\gamma-1}}, \qquad i_M = K_2\, n^\Gamma (\log n)^{\Gamma'}$$

where

$$\Gamma = -\frac{2\gamma^2 - 8\gamma + 5}{2\gamma - 1}, \qquad \Gamma' = -\frac{3(\gamma - 1)^2}{2\gamma - 1}.$$

By the definition of ℓs given in the statement of this theorem, we have

$$\ell = a\frac{1}{K_0}K_2^{2(\frac{1}{\gamma-1})}n^{-\frac{2\gamma-7}{2\gamma-1}}(\log n)^{\frac{5-4\gamma}{2\gamma-1}}$$

for a constant a. Therefore

$$\frac{\ell^{t+1}}{i_M^{t-1}} = b(c^{L_1(t)})(n^{L_2(t)})(\log n)^{L_3(t)}, \tag{11}$$

where b and c are constants and L_1, L_2 and L_3 are linear functions of t that are parameterized by γ. In particular,

$$L_2(t) = \frac{2(\gamma-2)(\gamma-3)}{2\gamma-1}t + \frac{2(-\gamma^2+3\gamma+1)}{2\gamma-1}.$$

For the range of values of γ we are concerned with (i.e., $5/2 < \gamma < 3$), the first factor is negative. Since $t = \log n$ is positive, this shows that (11) converges to 0 as $n \to \infty$.

Now we prove that each of the three listed conditions fails with polynomially negligible probability. We invoke Lemma 16 to show that this is the case for the first condition. For the second condition, we use the fact that each bit is independently set to 1 with some probability $P[i,j]$ where $i \leq i_M$ and $j \leq i_M$. Thus $p = P[i_M, i_M]$ is a lower bound on these probabilities. This, together with the definition of ℓ given in this theorem, allow for the hypotheses of Lemma 18 to be met. Thus, we can apply this lemma and show that condition 2 fails with polynomially negligible probability.

Finally, we have shown in our earlier discussion leading up to the(10) that the conditions of Lemma 19 are met. Let us identify the edges chosen by the adversary as the red edges, and let the ℓ-sampled edges be the ones that the watermarking algorithm selected in procedure keygen. By Lemma 19, the number of edges that are common to both selections is at least

$$R = \frac{8\ell r}{N^2} = \frac{16Dr}{pN^2} = \frac{D}{2}$$

where we used the fact that $\ell = 2D/p$, is at most

$$4\exp\left(-\frac{12}{7}\frac{\ell r}{N^2}\right) = 4\exp\left(-\frac{12}{7}\frac{D}{16}\right) = 4n^{-\frac{3}{7}(2C''+C')}.$$

So we have our result. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

4.1 Discussion

It is interesting to note how the differences in the two random graph models translate into differences in their watermarking schemes. The Erdős-Rényi model, with its uniform edge probability, allows for constant separation of high-degree vertices, at best. But all the vertices tend to be well-separated. On the other hand, the skewed edge distribution that is characteristic of the random power-law model allows high-degree vertices to be very well-separated, but a significant number of vertices—the low-degree ones, will not be easily distinguished.

These differences lead to the intuition that virtually all edges in the Erdős-Rényi model are candidates for use in a watermark, as long as only a constant number of selected edges are incident to any single vertex. Therefore, both our watermarking function and the adversary are allowed an approximately linear number of changes to the graph. Theorem 17 confirms this intuition with a scheme that proposes $O(n^\varepsilon)$ bits for the watermark, and a nearly linear number $O(n)$ bits that the adversary may modify.

In contrast, the number of edges that can be used as part of a watermark in the random power-law graph model is limited by the number of distinguishable vertices, which is on the order of i_M or $O(n^\varepsilon)$, where $\varepsilon = -\frac{2\gamma^2-8\gamma+5}{2\gamma-1}$.

5 Experiments

Although our paper is a foundational complement to the systems work of Zhao *et al.* [55, 56], we nevertheless provide in this section the results of a small set of empirical tests of our methods, so as to experimentally reproduce the hypothetical watermarking security experiment from Algorithm 1. Our experiments are performed on two large social network graphs, Youtube [54] from the SNAP library [36], and Flickr [39], as well as a randomly generated graph drawn from the random power-law graph model distribution. Table 1 illustrates the basic properties of the networks. To generate the random power-law graph, we set the number of nodes to $n = 10000$, the maximum degree to $m = 1000$, the average degree to $w = 20$, and $\gamma = 2.75$.

5.1 Adaptations from the theoretical scheme

To adapt our theoretical framework to the rough-and-tumble world of empirical realities, we made three modifications to our framework for the sake of our empirical tests.

First, instead of using the high-degree and medium-degree thresholds derived from Lemmas 12 and 14, for the power-law distribution, to define the cutoffs for high-degree and medium-degree vertices, we used these and the other lemmas given above as justifications for the existence of such distinguishing sets of vertices and we then optimized the number of high- and medium-degree vertices to be values that work best in practice. The column, "Unique degree," from Table 1 shows, for each network, the number of consecutive nodes with unique degree when considering the nodes in descending order of degree. This is, in theory, the maximum number of high-degree nodes that could be distinguished. Since this value is too small in most cases, we applied the principles of Lemmas 10 and 12 again, in a second-order fashion, to distinguish and order the high-degree nodes. In particular, in addition to the degree of each high-degree vertex, we also label each vertex with the list of degrees of its neighbors, sorted in decreasing order. With this change, we are not restricted in our choice of number of high-degree nodes as required by applying these lemmas only in a first-order fashion. Table 2 shows the values used in our experiments based on this second-order application. As medium-degree vertices, we picked the maximum number such that there are no collisions among their bit vectors of high-degree node adjacencies.

Second, instead of returning failure if (a) two high-degree nodes have the same degree and list of degrees of their neighbors, (b) two medium-degree nodes have the same bit vector, or (c) the approximate isomorphism is not injective, we instead proceed with the algorithm. Despite the existence of collisions, the remaining nodes often provide enough information to conclude successfully.

Table 1: Network statistics

Network	# nodes	# edges	Max. degree	Avg. degree	Unique degree	Estimated γ
Power-law	10,000	94,431	960	18.89	14	—
Youtube	1,134,890	2,987,624	28,754	5.27	29	1.48
Flickr	1,715,256	15,554,181	27,203	18.14	130	1.62

Table 2: Experiment Parameters

Network	# high-degree	# medium-degree	Key size	Marking dK-2 deviation
Power law	64	374	219	0.065
Youtube	256	113	184	0.033
Flickr	300	5901	3250	0.002

Finally, we simplified how we resampled (and flipped) edges in order to create a graph watermark, using our approach for the Erdős-Rényi model even for power-law graphs, since resampling uniformly among our small set of marked edges is likely not to cause major deviations in the graph's distribution and, in any case, it is empirically difficult to determine the value of γ for real-world social networks. Therefore, we set the resampling probability to 0.5 so that it is consistent with the Erdős-Rényi model and so that each bit in the message is represented uniformly and independently.

5.2 Fitting real-world networks to the random power-law graph model

Note that during the marking step, resampling the edges of the key requires knowledge of the distribution from which the network is drawn. For this reason, we tried to fit the real-world networks to the random power-law graph model distribution. The main task was to find the exponent γ of a power law function that would best fit the degree distribution.

First, we give definitions to help introduce the problem. The degree distribution of an undirected graph G with n vertices is a probability distribution such that the probability mass function P_G is given by $P_G(k) = n_k/n$ for $k \in \{0, \ldots, n\}$ where n_k is the number of vertices with degree k. A random variable is said to follow a power-law distribution if the probability density function is given by $f(x) = cx^{-\gamma}$ for some $c > 0$ and $\gamma > 0$. As stated before, it has been found empirically that the degree distributions of many naturally-occurring graphs in social networking and science follow a power-law distribution. Note that for observed data that is believed to follow a power-law distribution, the power-law behavior of the data often only holds for values larger than some $x_{\min}$ (see [12]).

In finding γ, there are three primary focuses: obtaining γ itself, finding the value for $x_{\min}$ for which the power-law behavior holds after, and finding the associated p-value indicating how good of a fit the power-law distribution is to the degree distribution of G. The methods to obtain all three values are taken from Clauset *et al.* [12] and some of the code used to find all three values can be found at [10, 11].

To obtain γ, we use the method of maximum likelihood. The maximum-likelihood estimator for γ is given by

$$\gamma = 1 + k \left(\sum_{i=1}^{k} \ln \frac{x_i}{x_{\min}} \right)^{-1} \tag{12}$$

where x_i for $i = 1, \ldots, k$ are the values of the degree distribution of G such that $x_i \geq x_{\min}$ (see [12, 42]).

Next, we discuss finding $x_{\min}$. $x_{\min}$ is the lower bound for which the data above it follow a power-law distribution. The idea used to find $x_{\min}$ is to choose the value in the observed data (some x such that $x \in \{P_G(k) : k = 0, \ldots, n\}$) such that the empirical cumulative distribution function (CDF) of the observed data above $x_{\min}$ is most similar to the estimated CDF of the power-law distribution of the observed data obtained by using Equation 12 (see [13]). We measure similarity between distributions using the Kolmogorov-Smirnov (KS) statistic. So in order to find $x_{\min}$, for each x_i we set $x_{\min} = x_i$ and compute the KS statistic for the empirical CDF of the observed data larger than $x_{\min}$ and the CDF of the power-law distribution obtained by using Equation 12. The x_i that minimizes the KS statistic will be used as the value for $x_{\min}$ [12].

Last, we discuss finding the p-value based off of the methods in [12]. To find the p-value indicating how well the power-law distribution fits the degree distribution of G, we use the methods described above to find γ and $x_{\min}$. After doing so, we compute the KS statistic for the empirical CDF of the data and the CDF of the power-law distribution with exponent γ and $x \geq x_{\min}$. We then generate new data sets from our observed data. Let n_{tail} be the number of observed data larger than

$x_{\min}$. Then with probability n_{tail}/n, we randomly sample a point from the power-law distribution with exponent γ and $x \geq x_{\min}$ then add it to the new data set. With probability $1 - n_{\text{tail}}/n$, we randomly sample an observed data point in the interval $x < x_{\min}$ and add it to the new data set. This process is continued until we have added n total data points. Then we compute the KS statistic for the empirical CDF of the newly generated data set and the CDF of the estimated power-law distribution using this newly generated data set. We generate 10,000 data sets and compute the KS statistic for each, where 10,000 is a good rule of thumb in order to have high precision in the p-value. The p-value is then determined by the fraction of the number of generated KS statistic values that are larger than the KS statistic obtained from the original data. A p-value larger than 0.1 implies that the power-law with exponent γ and $x \geq x_{\min}$ is a good fit to our data [12].

We gathered graphs from many different domains to test if the power-law distribution is a good fit to the degree distribution. We tested social networks from Google+ [38], LiveJournal [39], Slashdot [35], Epinions [47], Pokec [51], and Twitter [38]. We also tested citation networks of U.S. patents [33] and ArXiv [21,33], collaboration networks of ArXiv [34] and the DBLP computer science bibliography [54], email communication networks from the Enron corpus [29,35] and a European research institution [34], a communication network from Wikipedia's talk pages [32], location-based online social networks (OSN) [8], Internet autonomous systems (AS) networks [18,19,33,45], snapshots of the graph of the peer-to-peer file sharing service Gnutella [34,48], road networks from Pennsylvania, Texas, and California [35], and product co-purchasing networks from Amazon [31,54]. Additionally, we tested web graphs where the nodes represent web pages and the edges are hyperlinks connecting the pages for part of Stanford's website [35], Notre Dame's website [35], Stanford and Berkeley's websites [35], and Google [35]. Our results, shown in Table 3, display that the power-law distribution is a good fit for the degree distribution for 29 of the 40 graphs tested. Moreover, 23 of the 29 cases where it is a good fit, the estimated γ is less than 2. This prevented us from using Equation 2 for the resampling probabilities, which requires $\gamma > 2$.

5.3 Experiment parameters

For the experiment parameters other than the original network and the number of high- and medium-degree nodes, we set the following values.

Maximum flips adjacent to any given node during marking: 1.

Key size: We set this to the maximum possible value (i.e., the number of high- and medium-degree vertices divided by two, as shown in Table 2), because the numbers of high- and medium-degree nodes are not large. This effectively means that every high- and medium-degree node has exactly one edge added or removed.

Number of marked graphs: 10.

Adversary: We used a time-efficient variation of the *arbitrary edge-flipping adversary*. This adversary selects a set of pairs of nodes randomly, and flips the potential edge among each pair.

5.4 Results

We evaluated how much distortion the adversary can introduce before our method fails to identify the leaked network correctly. For this purpose, we compared the identification success rate to the amount of distortion under different fractions of modified edges by the adversary. To estimate the success rate, we ran the experiment 10 times and reported the fraction of times that the leaked

Table 3: Estimated γ and associated p-values

Network Type	Network	# nodes	# edges	Estimated γ	p-value
Social Networks	Google+	107,614	12,238,285	1.8129	0.0096
	LiveJournal	5,203,764	48,709,773	1.3901	0.0171
	Slashdot 08/11	77,360	469,180	1.5947	0.6425
	Slashdot 02/09	82,168	504,230	1.5906	0.3517
	Epinions	75,879	405,740	1.6893	0.9979
	Pokec	1,632,803	22,301,964	2.4144	0.8383
	Twitter	81,306	1,342,296	1.4844	0.0000
Citation Networks	ArXiV (Physics)	34,546	420,877	1.3931	0.0000
	ArXiV (Phy. Theory)	27,769	352,285	1.4309	0.0000
	Patents	3,774,768	16,518,947	1.2776	0.2102
Collaboration Networks	ArXiV (Astro.)	18,771	198,050	2.0062	0.3105
	ArXiV (Condensed)	23,133	93,439	1.3896	0.0715
	ArXiV (Physics)	12,006	118,489	1.8967	0.5440
	DBLP	317,080	1,049,866	1.3405	0.0530
Communication Networks	Email - Enron	36,692	183,831	1.6143	0.9884
	Email - Europe	265,009	364,481	1.4703	0.9998
	Wikipedia	2,394,385	4,659,565	1.5232	0.9933
Location-based OSNs	Brightkite	58,228	214,078	1.5403	0.4212
	Gowalla	196,591	950,327	1.4542	0.3132
AS Networks	CAIDA	26,475	53,381	1.5446	0.9191
	Skitter	1,696,415	11,095,298	1.4484	0.7464
	Oregon - 1	11,174	23,409	1.5523	0.9113
	Oregon - 2	11,461	32,730	1.6024	0.8396
Peer-to-peer Networks	Gnutella 08/04/02	10,876	39,994	2.5995	0.6346
	Gnutella 08/24/02	26,518	65,369	2.0109	0.7580
	Gnutella 08/25/02	22,687	54,705	2.0495	0.6715
	Gnutella 08/30/02	36,682	88,328	1.8832	0.8464
	Gnutella 08/31/02	62,586	147,892	1.9131	0.9071
Co-purchasing Networks	Amazon 03/02/03	262,111	899,792	1.3310	0.6295
	Amazon 03/12/03	400,727	2,349,869	1.3680	0.5020
	Amazon 05/05/03	410,236	2,439,437	1.3580	0.5129
	Amazon 06/01/03	403,394	2,443,408	1.3684	0.6091
	Amazon (2012)	334,863	925,872	1.3382	0.5789
Road Networks	California	1,965,206	2,766,607	1.1415	0.0785
	Pennsylvania	1,088,092	1,541,898	3.1585	0.6290
	Texas	1,379,917	1,921,660	1.1521	0.0642
Web Graphs	Berkeley & Stanford	685,230	6,649,470	1.4591	0.0087
	Google	875,713	4,322,051	1.3890	0.3828
	Notre Dame	325,729	1,090,108	1.6577	0.8434
	Stanford	281,903	1,992,636	1.4052	0.0036

network was identified correctly. As a measure of distortion, we used the dK-2 deviation [55] between the original network and the version modified by the adversary. The dK-2 deviation is the euclidean distance between the dK-2 series [49] of the two graphs, normalized by the number of tuples in the dK-2 series. The dK-2 deviation captures the differences between the joint degree distributions of the networks, that is, the probability that a randomly selected edge has as endpoints nodes with certain degrees. We average the dK-2 deviation among the 10 runs. Figure 2 shows the outcome of our experiments. Moreover, Table 2 shows the dK-2 deviation introduced by the marking alone.

Power-law
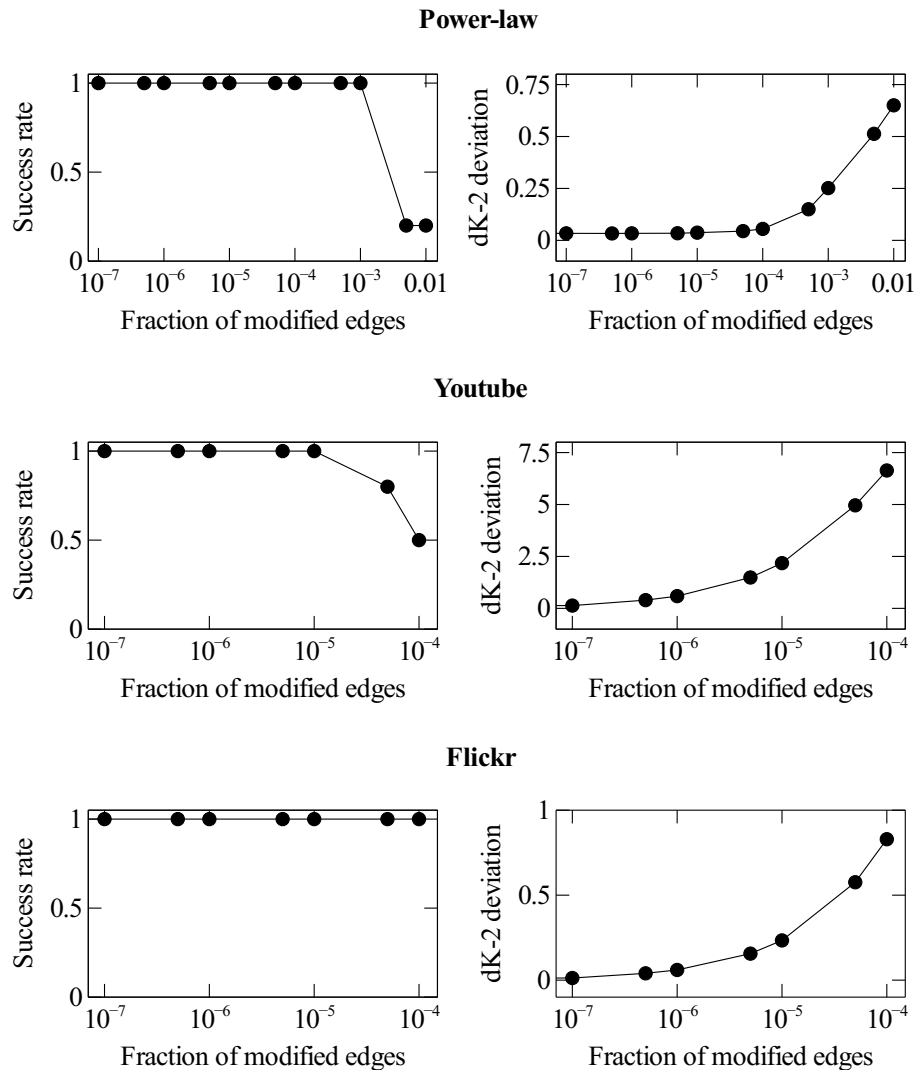
Youtube

Flickr

Figure 2: Success rate and dK-2 deviation under different fractions of modified potential edges by the adversary, for the Power law, Youtube, and Flickr networks.

5.5   Discussion

Based on our experiments, the success rate of our scheme is high but it drops after a certain threshold. This demonstrates that there is a distinct range of adversarial edge flips that can be tolerated by our scheme. Specifically, our scheme worked well when the fraction of potential edges

flipped by the adversary is up to 10^{-3} and 10^{-5} for the random power-law and Youtube networks, respectively. For these graphs, this number of flipped potential edges corresponds to 52.9% and 215.6% of the number of edges in the original graphs, respectively. For the Flickr network, the runtime of the adversary modification became excessive before the success rate could decrease, at a fraction of 10^{-4} of potential edges flipped.

The distortion introduced by the watermark is negligible compared to the distortion caused by the number of flips that the scheme can tolerate. On average, the marking modifies half of the edges on the key, which corresponds to $1.1 \cdot 10^{-3}, 3 \cdot 10^{-5}$, and 10^{-4} of the number of edges in the original random power-law, Youtube, and Flickr networks, respectively.

For the same number of flips, the dK-2 deviation in the Youtube network was much larger than in the Flickr network, which in turn was larger than that of the random power-law network. A possible explanation for this is that any set of uniform edge flips has a bigger effect on the dk2-deviation of a skewed graph than on the dK-2 deviation of a less skewed graph. Note that the Youtube network has the largest skew, as the maximum degree is on the same order as the Flickr network, but the average degree is less.

6   Conclusion

We defined a watermarking framework and a notion of security meant to capture the difficulty in removing a watermark from a graph. We studied two random graph models and showed that watermarking in these models could be achieved in such a way that no adversary could remove the watermark whp and still have a graph that is "close" to his original graph. A vital feature of our approach is that we watermark graphs to look like typical graphs from the distribution the original graph was issued from, while also making them look similar to the original. In addition, we provided an exemplary implementation that works effectively for marks consisting only of edge flips.

For future work, it would be interesting to consider solutions that can tolerate some degree of collusion. In the exemplary schemes we presented, an adversary is likely to detect many edges of the watermark, if he has access to multiple watermarked graphs produced from the same original graph.

Acknowledgments

References

[1] L. Babai. Graph isomorphism in quasipolynomial time. *ArXiv ePrint*, abs/1512.03547, 2015.

[2] L. Babai, P. Erdös, and S. M. Selkow. Random graph isomorphism. *SIAM Journal on Computing*, 9(3):628–635, 1980.

[3] L. Backstrom, C. Dwork, and J. Kleinberg. Wherefore art thou r3579x?: Anonymized social networks, hidden patterns, and structural steganography. *Commun. ACM*, 54(12):133–141, Dec. 2011.

[4] L. M. S. Bento, D. Boccardo, R. C. S. Machado, V. Pereira de Sá, and J. Szwarcfiter. Towards a provably resilient scheme for graph-based watermarking. In A. Brandstädt, K. Jansen, and R. Reischuk, editors, *Graph-Theoretic Concepts in Computer Science*, volume 8165 of *LNCS*, pages 50–63. Springer, 2013.

[5] B. Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.

[6] G. Caldarelli. *Scale-Free Networks: Complex Webs in Nature and Technology*. Oxford University Press, 2013.

[7] X. Chen, D. Fang, J. Shen, F. Chen, W. Wang, and L. He. A dynamic graph watermark scheme of tamper resistance. In *5th Int. Conf. on Information Assurance and Security (IAS)*, volume 1, pages 3–6, 2009.

[8] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090. ACM, 2011.

[9] F. Chung and L. Lu. *Complex graphs and networks*, volume 107 of *CBMS Regional Conf. Series in Math.* American Math. Society, 2006.

[10] A. Clauset and J. Ornstein. Calculating p-value for fitted power-law model. `http://tuvalu.santafe.edu/~aaronc/powerlaws/plpva.py`, July 2011.

[11] A. Clauset and J. Ornstein. Fitting a power-law distribution. `http://tuvalu.santafe.edu/~aaronc/powerlaws/plfit.py`, July 2011.

[12] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009.

[13] A. Clauset, M. Young, and K. S. Gleditsch. On the frequency of severe terrorist events. *Journal of Conflict Resolution*, 51(1):58–87, 2007.

[14] C. Collberg, S. Kobourov, E. Carter, and C. Thomborson. Graph-based approaches to software watermarking. In H. Bodlaender, editor, *Graph-Theoretic Concepts in Computer Science*, volume 2880 of *LNCS*, pages 156–167. Springer, 2003.

[15] C. Collberg and C. Thomborson. Software watermarking: Models and dynamic embeddings. In *ACM Symp. Princ. of Prog. Lang. (POPL)*, pages 311–324, 1999.

[16] I. Cox, M. Miller, J. Bloom, J. Fridrich, and T. Kalker. *Digital watermarking and steganography*. Morgan Kaufmann, 2007.

[17] T. Czajka and G. Pandurangan. Improved random graph isomorphism. *Journal of Discrete Algorithms*, 6(1):85–92, 2008.

[18] C. for Applied Internet Data Analysis. As relationships. `http://www.caida.org/data/as-relationships/`.

[19] C. for Applied Internet Data Analysis. Skitter. `http://www.caida.org/tools/measurement/skitter/`.

[20] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman, 1979.

[21] J. Gehrke, P. Ginsparg, and J. Kleinberg. Overview of the 2003 kdd cup. *ACM SIGKDD Explorations Newsletter*, 5(2):149–151, 2003.

[22] J. Hamilton and S. Danicic. A survey of graph based software watermarking. Technical report, Department of Computing, Goldsmiths, University of London, 2010.

[23] M. Hay, G. Miklau, D. Jensen, D. Towsley, and P. Weis. Resisting structural re-identification in anonymized social networks. *Proc. VLDB Endow.*, 1(1):102–114, Aug. 2008.

[24] N. Hopper, D. Molnar, and D. Wagner. From weak to strong watermarking. In *Theory of Cryptography*, volume 4392 of *LNCS*, pages 362–382. Springer, 2007.

[25] S. Katzenbeisser and F. Petitcolas. *Information hiding techniques for steganography and digital watermarking.* Artech house, 2000.

[26] A. Kayem, A. Deshai, and S. Hammer. On anonymizing social network graphs. In *Information Security for South Africa (ISSA)*, pages 1–8, Aug 2012.

[27] E. Kazemi, S. H. Hassani, and M. Grossglauser. Growing a Graph Matching from a Handful of Seeds. *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases*, 8(10):1010–1021, 2015.

[28] S. Khanna and F. Zane. Watermarking maps: Hiding information in structured data. In *11th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 596–605, 2000.

[29] B. Klimt and Y. Yang. Introducing the enron corpus. In *CEAS*, 2004.

[30] J. Kobler, U. Schöning, and J. Torán. *The graph isomorphism problem: its structural complexity.* Springer, 2012.

[31] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.

[32] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1361–1370. ACM, 2010.

[33] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 177–187. ACM, 2005.

[34] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2, 2007.

[35] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1):29–123, 2009.

[36] J. Leskovec and R. Sosič. SNAP: A general purpose network analysis and graph mining library in C++. http://snap.stanford.edu/snap, June 2014.

[37] K. Liu and E. Terzi. Towards identity anonymization on graphs. In *ACM SIGMOD Int. Conf. on Management of Data*, pages 93–106, 2008.

[38] J. J. McAuley and J. Leskovec. Learning to discover social circles in ego networks. In *NIPS*, volume 2012, pages 548–56, 2012.

[39] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *5th ACM/Usenix Internet Measurement Conference (IMC)*, 2007.

[40] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1(2):226–251, 2004.

[41] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge Univ. Press, New York, NY, 2005.

[42] A. Muniruzzaman. On measures of location and dispersion and tests of hypotheses on a pareto population. *Bulletin of the Calcuta Statistical Association*, 7:115–123, 1957.

[43] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *IEEE Symp. on Security and Privacy (SP)*, pages 173–187, 2009.

[44] M. Newman, A.-L. Barabasi, and D. J. Watts. *The Structure and Dynamics of Networks.* Princeton Studies in Complexity. Princeton Univ. Press, 2006.

[45] U. of Oregon. University of Oregon route views project. `http://www.routeviews.org/`.

[46] R. C. Read and D. G. Corneil. The graph isomorphism disease. *Journal of Graph Theory*, 1(4):339–363, 1977.

[47] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *The Semantic Web-ISWC 2003*, pages 351–368. Springer, 2003.

[48] M. Ripeanu, I. Foster, and A. Iamnitchi. Mapping the gnutella network: Properties of large-scale peer-to-peer systems and implications for system design. *ArXiv ePrint cs/0209028*, 2002.

[49] A. Sala, L. Cao, C. Wilson, R. Zablit, H. Zheng, and B. Y. Zhao. Measurement-calibrated graph models for social network experiments. In *19th Int. Conf. on the World Wide Web (WWW)*, pages 861–870, 2010.

[50] F. Y. Shih. *Digital watermarking and steganography: fundamentals and techniques.* CRC Press, 2007.

[51] L. Takac and M. Zabovsky. Data analysis in public social networks. In *International Scientific Conference and International Workshop Present Day Trends of Innovations*, pages 1–6, 2012.

[52] R. Venkatesan, V. Vazirani, and S. Sinha. A graph theoretic approach to software watermarking. In I. Moskowitz, editor, *Information Hiding*, volume 2137 of *LNCS*, pages 157–168. Springer, 2001.

[53] G. Wondracek, T. Holz, E. Kirda, and C. Kruegel. A practical attack to de-anonymize social network users. In *IEEE Symp. on Security and Privacy (SP)*, pages 223–238, 2010.

[54] J. Yang and J. Leskovec. Defining and evaluating network communities based on ground-truth. *CoRR*, abs/1205.6233, 2012.

[55] X. Zhao, Q. Liu, H. Zheng, and B. Y. Zhao. Towards graph watermarks. In *2015 ACM Conf. on Online Social Networks (COSN)*, pages 101–112, 2015.

[56] X. Zhao, Q. Liu, L. Zhou, H. Zheng, and B. Y. Zhao. Graph watermarks. *ArXiv ePrint*, abs/1506.00022, 2015.

[57] B. Zhou and J. Pei. Preserving privacy in social networks against neighborhood attacks. In *IEEE Int. Conf. on Data Engineering (ICDE)*, pages 506–515, April 2008.