First-Order Model-Checking in Random Graphs and Complex Networks

Jan Dreier
Department of Computer Science, RWTH Aachen University, Germany
https://tcs.rwth-aachen.de/~dreier
dreier@cs.rwth-aachen.de

Philipp Kuinke
Department of Computer Science, RWTH Aachen University, Germany
https://tcs.rwth-aachen.de/~kuinke
kuinke@cs.rwth-aachen.de

Peter Rossmanith
Department of Computer Science, RWTH Aachen University, Germany
https://tcs.rwth-aachen.de
rossmani@cs.rwth-aachen.de

Abstract

Complex networks are everywhere. They appear for example in the form of biological networks, social networks, or computer networks and have been studied extensively. Efficient algorithms to solve problems on complex networks play a central role in today's society. Algorithmic meta-theorems show that many problems can be solved efficiently. Since logic is a powerful tool to model problems, it has been used to obtain very general meta-theorems. In this work, we consider all problems definable in first-order logic and analyze which properties of complex networks allow them to be solved efficiently.

The mathematical tool to describe complex networks are random graph models. We define a property of random graph models called α-power-law-boundedness. Roughly speaking, a random graph is α-power-law-bounded if it does not admit strong clustering and its degree sequence is bounded by a power-law distribution with exponent at least α (i.e. the fraction of vertices with degree k is roughly $O(k^{-\alpha})$).

We solve the first-order model-checking problem (parameterized by the length of the formula) in almost linear FPT time on random graph models satisfying this property with $\alpha \geq 3$. This means in particular that one can solve every problem expressible in first-order logic in almost linear expected time on these random graph models. This includes for example preferential attachment graphs, Chung–Lu graphs, configuration graphs, and sparse Erdős–Rényi graphs. Our results match known hardness results and generalize previous tractability results on this topic.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases random graphs, average case analysis, first-order model-checking

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.40

Related Version A full version of the paper is available at https://arxiv.org/abs/2006.14488.

Funding This research has been partially supported by the German Science Foundation (DFG) under grant DFG RO 927/15-1.

1 Introduction

Complex networks, as they occur in society, biology and technology, play a central role in our everyday lives. Even though these networks occur in vastly different contexts, they are structured and evolve according to a common set of underlying principles. Over the last two

decades, with the emergence of the field of network science, there has been an explosion in research to understand these fundamental laws. One well observed property is the *small-world phenomenon*, which means that distances between vertices are very small. This has been verified for the internet and many other networks [1, 56]. Furthermore, many real networks tend to be *clustered*. They contain groups of vertices that are densely connected [67]. If two vertices share a common neighbor, then there is a high chance that there is also an edge between them. A network can be considered clustered if the ratio between the number of triangles and the number of paths with three vertices is non-vanishing. This is formalized by the clustering coefficient, which is high for many networks [72]. A third important property is a *heavy tailed degree distribution*. While most vertices have a low number of connections, there are a few hubs with a high degree. Experiments show that the degrees follow for example a power-law or log-normal distribution. In a power-law distribution, the fraction of vertices with degree k is proportional to $k^{-\alpha}$ (usually with α between 2 and 3). This behavior makes complex networks highly inhomogeneous [65, 58, 10, 15].

One important goal of theoretical computer science has always been to explore what kinds of inputs allow or forbid us to construct efficient algorithms. In this context, *algorithmic meta-theorems* [52] are of particular interest. They are usually theorems stating that problems definable in a certain logic can be solved efficiently on graph classes that satisfy certain properties. Logic is a powerful tool to model problems and therefore has been used to obtain very general meta-theorems. A well-known example is Courcelle's theorem [16], which states that every problem expressible in counting monadic second-order logic can be solved in linear time on graph classes with bounded treewidth. It has been further generalized to graph classes with bounded cliquewidth [17]. To obtain results for larger graph classes one has to consider weaker logics. The languages of relational database systems are based on first-order logic. In this logic, one is allowed to quantify over vertices and to test equality and adjacency of vertices. With k existential quantifiers, one may ask for the existence of a fixed graph with k vertices (k-subgraph isomorphism), a problem relevant to motif-counting [57, 29]. On the other hand, connectivity properties cannot be expressed in first-order logic. We define for every graph class $\mathcal{G}$ the parameterized first-order model-checking problem $p\text{-MC}(\text{FO}, \mathcal{G})$ [45].

$p\text{-MC}(\text{FO}, \mathcal{G})$	
Input:	A graph $G \in \mathcal{G}$ and a first-order sentence φ
Parameter:	The number of symbols in φ, denoted by $\lvert\varphi\rvert$
Problem:	Does φ hold on G (i.e. $G \models \varphi$)?

The aim is to show for a given graph class $\mathcal{G}$ that $p\text{-MC}(\text{FO}, \mathcal{G})$ is *fixed parameter tractable* (FPT), i.e., can be decided in time $f(\lvert\varphi\rvert)n^{O(1)}$ for some function f (see for example [18] for an introduction to fixed parameter tractability). Since input graphs may be large, a linear dependence on n is desirable. If one is successful, then every problem expressible in first-order logic can be solved on $\mathcal{G}$ in linear time.

For the class of all graphs $\mathfrak{G}$, $p\text{-MC}(\text{FO}, \mathfrak{G})$ is AW[$*$]-complete [25] and therefore most likely not fpt. Over time, tractability of $p\text{-MC}(\text{FO}, \mathcal{G})$ has been shown for more and more sparse graph classes $\mathcal{G}$: bounded vertex degree [69], forbidden minors [35], bounded local treewidth [34], and further generalizations [19, 30, 68]. Grohe, Kreutzer and Siebertz prove that $p\text{-MC}(\text{FO}, \mathcal{G})$ can be solved in almost linear FPT time $f(\lvert\varphi\rvert, \varepsilon)n^{1+\varepsilon}$ for all $\varepsilon > 0$ if $\mathcal{G}$ is a nowhere dense graph class [46]. On the other hand if $\mathcal{G}$ is a monotone somewhere dense graph class, $p\text{-MC}(\text{FO}, \mathcal{G})$ is AW[$*$]-hard [46]. Nowhere dense graph classes were introduced by Nešetřil and Ossona de Mendez as those graph classes where for every $r \in \mathbf{N}$ the size of all r-shallow clique minors of all graphs in the graph class is bounded by a function of r

(Section 4.3). A graph class is somewhere dense if it is not nowhere dense. The tractability of the model-checking problem on monotone graph classes is completely characterized with a dichotomy between nowhere dense and somewhere dense graph classes. These very general results come at a cost: Frick and Grohe showed that the dependence of the run time on φ is non-elementary [38]. We want to transfer this rich algorithmic theory to complex networks. But what is the right abstraction to describe complex networks?

Network scientists observed that the chaotic and unordered structure of real networks can by captured using *randomness*. There is a vast body of research using random processes to create graphs that mimic the fundamental properties of complex networks. The most prominent ones are the preferential attachment model [3, 64], Chung–Lu model [12, 13], configuration model [60, 59], Kleinberg model [50, 51], hyperbolic graph model [53], and random intersection graph model [47, 66]. All these are random models. It has been thoroughly analyzed how well they predict various properties of complex networks [43].

When it comes to algorithmic meta-theorems on random graph models "even the most basic questions are wide open," as Grohe puts it [45]. By analyzing which models of complex networks and which values of the model-parameters allow for efficient algorithms, we aim to develop an understanding how the different properties of complex networks control their algorithmic tractability.

In this work we show for a wide range of models, including the well known preferential attachment model, that one can solve the parameterized first-order model-checking problem in almost linear FPT time. This means in particular that one can solve every problem expressible in first-order logic efficiently on these models. Our original goal was to obtain efficient algorithms only for preferential attachment graphs, but we found an abstraction that transfers these results to many other random graph models. Roughly speaking, the following two criteria are sufficient for efficiently solving first-order definable problems on a random graph model:

- The model needs to be unclustered. In particular the expected number of triangles needs to be subpolynomial.
- For every k, the fraction of vertices with degree k is roughly $O(k^{-3})$. In other words, the degree sequence needs to be bounded by a power-law distribution with exponent 3 or higher.

Models satisfying these properties include sparse Erdős–Rényi graphs, preferential attachment graphs as well as certain Chung–Lu and configuration graphs. On the other hand, the Kleinberg model, the hyperbolic random graph model, or the random intersection graph model do not satisfy these properties. Our results generalize previous results [44, 22] and match known hardness results: The model-checking problem has been proven to be hard on power-law distributions with exponent smaller than 3 [28]. We therefore identify the threshold for tractability to be a power-law coefficient of 3. It is also a big open question whether the model-checking problem can also be solved on clustered random graph models, especially since real networks tend to be clustered. Furthermore, significant engineering challenges need to be overcome to make our algorithms applicable in practice.

1.1 Average Case Complexity

Average-case complexity analyzes the typical run time of algorithms on random instances (see [6] for a survey), based on the idea that a worst-case analysis often is too pessimistic as for many problems hard instances occur rarely in the real world. Since models of complex networks are probability distributions over graphs, we analyze the run time of algorithms under average-case complexity. However, there are multiple notions and one needs to be careful which one to choose.

Assume a random graph model is asymptotically almost surely (a.a.s.) nowhere dense, i.e., a random graph from the model with n vertices belongs with probability $1-\delta(n)$ to a nowhere dense graph class, where $\lim_{n\to\infty}\delta(n)=0$ (Section 4.3). Then the first-order model-checking problem can be efficiently solved with a probability converging to one [46]. However, with probability $\delta(n)$ the run time can be arbitrarily high and the rate of convergence of $\delta(n)$ to zero can be arbitrarily slow. These two missing bounds are undesirable from an algorithmic standpoint and the field of average-case complexity has established a theory on how the run time needs to be bounded with respect to the fraction of inputs that lead to this run time.

This is formalized by the well-established notion of *average polynomial run time*, introduced by Levin [54]. An algorithm has average polynomial run time with respect to a random graph model if there is an $\varepsilon > 0$ and a polynomial p such that for every n, t the probability that the algorithm runs longer than t steps on an input of size n is at most $p(n)/t^{\varepsilon}$. This means there is a polynomial trade-off between run time and fraction of inputs. This notion has been widely studied [6, 2] and is considered from a complexity theoretic standpoint the right notion of polynomial run time on random inputs. It is closed under invoking polynomial subroutines.

In our work, however, we wish to explicitly distinguish linear time. While Levin's complexity class is a good analogy to the class P, it is not suited to capture algorithms with average linear run time. For this reason, we turn to the expected value of the run time, a stronger notion than average polynomial time. In fact, using Markov's inequality we see that if an algorithm has expected linear run time, all previous measures of average tractability are also bounded. Their relationship is as follows.

$$\text{expected linear} \;\Rightarrow\; \text{expected polynomial} \;\Rightarrow\; \text{average polynomial} \;\Rightarrow\; \text{a.a.s. polynomial}$$

With this in mind we can present our notion of algorithmic tractability. A labeled graph is a graph where every vertex can have (multiple) labels. First-order formulas can have unary predicates for each type of label. These predicates test whether a vertex has a label of a certain type. We define $\mathfrak{G}$ to be the class of all graphs, and $\mathfrak{G}_{lb}$ to be the class of all vertex-labeled graphs. A function $L\colon \mathfrak{G} \to \mathfrak{G}_{lb}$ is an *l-labeling function* for $l \in \mathbf{N}$ if for every $G \in \mathfrak{G}$, $L(G)$ is a labeling of G with up to l classes of labels (see Section 4 for details). Furthermore, a *random graph model* is a sequence $\mathcal{G} = (\mathcal{G}_n)_{n\in\mathbf{N}}$, where $\mathcal{G}_n$ is a probability distribution over unlabeled simple graphs with n vertices.

▶ **Definition 1.** *We say* p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ *can be decided on a random graph model* $(\mathcal{G}_n)_{n\in\mathbf{N}}$ *in* expected time $f(|\varphi|, n)$ *if there exists a deterministic algorithm* $\mathcal{A}$ *which decides* p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ *on input* G, φ *in time* $t_{\mathcal{A}}(G, \varphi)$ *and if for all* $n \in \mathbf{N}$, *all first-order sentences* φ *and all* $|\varphi|$*-labeling functions* L, $\mathrm{E}_{G\sim\mathcal{G}_n}\big[t_{\mathcal{A}}(L(G), \varphi)\big] \leq f(|\varphi|, n)$. *We say* p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ *on a random graph model can be decided in* expected FPT time *if it can be decided in expected time* $g(|\varphi|)n^{O(1)}$ *for some function* g.

In particular, this definition implies efficient average run time according to Levin's notion (which is closed under polynomial subroutines). We choose to include labels into our notion of average-case hardness for two reasons: First, it makes our algorithmic results stronger, as the expected run time is small, even in the presence of an adversary that labels the vertices of the graph. Secondly, it matches known hardness results that require adversary labeling.

1.2 Previous Work

There have been efforts to transfer the results for classical graph classes to random graph models by showing that a graph sampled from some random graph model belongs with high probability to a certain algorithmically tractable graph class. For most random graph

models the treewidth is polynomial in the size of the graph [41, 4]. Therefore, people have considered more permissive graph measures than treewidth, such as low degree [44], or bounded expansion [22, 33]. Demaine et al. showed that some Chung–Lu and configuration graphs have bounded expansion and provided empirical evidence that some real-world networks do, too [22]. However, this technique is still limited, as many random graph models (such as the preferential attachment model [22, 27]) are not known to be contained in any of the well-known tractable graph classes.

The previous tractability results presented in this section all use the following technique: Assume we have a formula φ and sample a graph of size n from a random graph model. If the sampled graph belongs to the tractable graph class, an efficient model-checking algorithm for the graph class can solve the instance in FPT time. If the graph does not belong to the graph class, the naive model-checking algorithm can still solve the instance in time $O(n^{|\varphi|})$. Assume we can show that the second case only happens with probability $\delta(n)$ converging to zero faster than any polynomial. Then $\delta(n)O(n^{|\varphi|})$ converges to zero and the expected run time remains bounded by an FPT function.

Let $p(n)$ be a function with $p(n) = O(n^\varepsilon/n)$ for all $\varepsilon > 0$. Grohe showed that one can solve p-MC(FO, $\mathfrak{G}_{lb}$) on Erdős–Rényi graphs $G(n, p(n))$ in expected time $f(|\varphi|, \varepsilon)n^{1+\varepsilon}$ for every $\varepsilon > 0$ [44]. This result was obtained by showing that with high probability the maximum degree of the random graph model is $O(n^\varepsilon)$ for every $\varepsilon > 0$ and then using a model-checking algorithm for low degree graphs. Later Demaine et al. and Farrell et al. showed that certain Chung–Lu and configuration graphs whose degrees follow a power-law distribution with exponent $\alpha > 3$ [22] as well as certain random intersection graphs [33] belong with high probability to a graph class with bounded expansion. While they do not mention it explicitly, the previous argument implies that one can solve p-MC(FO, $\mathfrak{G}_{lb}$) in expected time $f(|\varphi|)n$ on these random graph models.

There further exist some average-case hardness results for the model-checking problem. It has been shown that one cannot decide p-MC(FO, $\mathfrak{G}_{lb}$) on Erdős–Rényi graphs $G(n, 1/2)$ or $G(n, p(n))$ with $p(n) = n^\varepsilon/n$ for some $0 < \varepsilon < 1$, $\varepsilon \in \mathbf{Q}$, in expected FPT time unless AW[$*$] $\subseteq$ FPT/poly [28]. The same holds for Chung–Lu graphs with exponent $2.5 < \alpha < 3$, $\alpha \in \mathbf{Q}$. These hardness results fundamentally require the adversary labeling of Definition 1. It is a big open question whether they can be transferred to model-checking without labels.

Another thing to keep in mind when considering logic and random graphs [70] are *zero-one laws*. They state that in many Erdős–Rényi graphs every first-order formula holds in the limit either with probability zero or one [70, 42, 32]. Not all random graph models satisfy a zero-one law for first-order logic (e.g. the limit probability of the existence of a K_4 in a Chung–Lu graph with weights $w_i = \sqrt{n/i}$ is neither zero nor one).

2 Our Results

We define a property called *α-power-law-boundedness*. This property depends on a parameter α and captures many unclustered random graph models for which the fraction of vertices with expected degree $d \in \mathbf{N}$ is roughly $O(d^{-\alpha})$. Our main contribution is solving the model-checking problem efficiently on all α-power-law-bounded random graph models with $\alpha \geq 3$. This includes preferential attachment graphs, Chung–Lu graphs, Erdős–Rényi graphs, and other random graph models. Note that graphs do not need to have a power-law degree distribution to be α-power-law-bounded. Our results hold for arbitrary labelings of the random graph and are based on a novel decomposition technique for local regions of random graphs. While all previous algorithms work by placing the random graph model with high probability in a sparse graph class, our technique also works for some a.a.s. somewhere dense random graphs (e.g. preferential attachment graphs [27]).

2.1 Power-Law-Boundedness

We start by formalizing our property. Since it generalizes the Chung–Lu model, we define this model first. A Chung–Lu graph with exponent α and vertices $v_1, \dots, v_n$ is defined such that two vertices v_i and v_j are adjacent with probability $\Theta(w_i w_j/n)$ where $w_i = (n/i)^{1/(\alpha-1)}$ [12]. Furthermore all edges are independent, which means that the probability that a set of edges occurs equals the product over the probabilities of each individual edge. In our model the probability of a set of edges can be a certain factor larger than the product of the individual probabilities, which allows edges to be moderately dependent.

▶ **Definition 2.** *Let $\alpha > 2$. We say a random graph model $(\mathcal{G}_n)_{n\in\mathbf{N}}$ is α-power-law-bounded if for every $n \in \mathbf{N}$ there exists an ordering $v_1, \dots, v_n$ of $V(\mathcal{G}_n)$ such that for all $E \subseteq \binom{\{v_1,\dots,v_n\}}{2}$*

$$\Pr\big[E \subseteq E(\mathcal{G}_n)\big] \leq \prod_{v_i v_j \in E} \frac{(n/i)^{1/(\alpha-1)}(n/j)^{1/(\alpha-1)}}{n} \cdot \begin{cases} 2^{O(|E|^2)} & \text{if } \alpha > 3 \\ \log(n)^{O(|E|^2)} & \text{if } \alpha = 3 \\ O(n^\varepsilon)^{|E|^2} \text{ for every } \varepsilon > 0 & \text{if } \alpha < 3. \end{cases}$$

The probability that a set of edges E occurs may be up to a factor $2^{O(|E|^2)}$ or $\log(n)^{O(|E|^2)}$ or $O(n^\varepsilon)^{|E|^2}$ (depending on α) larger than the probability in the corresponding Chung–Lu graph. For conditional probabilities this means the following: The probability bound for an edge under the condition that some set of l edges is already present may be up to a factor $2^{O(l)}$ or $\log(n)^{O(l)}$ or $O(n^\varepsilon)^l$ larger than the unconditional probability. This lets power-law-bounded random graphs capture moderate dependence between edges. The factor undergoes a phase transition at $\alpha = 3$. The smaller factor $2^{O(|E|^2)}$ for $\alpha > 3$ was chosen to guarantee linear FPT run time of our model-checking algorithm (Theorem 4) if $\alpha > 3$. The slightly larger factor of $\log(n)^{O(|E|^2)}$ for $\alpha = 3$ was chosen to capture preferential attachment graphs while still maintaining a quasilinear FPT run time of our algorithm.

The parameter α of an α-power-law-bounded random graph model controls the degree distribution. Note that if a graph class is α-power-law-bounded it is also α'-power-law-bounded for all $2 < \alpha' < \alpha$. It can be easily seen that a vertex v_i has expected degree at most $O(n^\varepsilon)(n/i)^{1/(\alpha-1)}$ for every $\varepsilon > 0$. This means the expected degree sequence of an α-power-law-bounded random graph model is not power-law distributed with exponent smaller than α. The gap is often tight: For example, Chung–Lu graphs with a power-law degree distribution exponent α are α-power-law-bounded and preferential attachment graphs have a power-law degree distribution with exponent 3 and are 3-power-law-bounded. For the interesting case $\alpha = 3$, the inequality in Definition 2 simplifies to

$$\Pr\big[E \subseteq E(\mathcal{G}_n)\big] \leq \log(n)^{O(|E|^2)} \prod_{v_i v_j \in E} \frac{1}{\sqrt{ij}}.$$

2.2 Model Checking

We now present our model-checking algorithm for α-power-law-bounded graphs. We express its run time relative to the term

$$\tilde{d}_\alpha(n) = \begin{cases} O(1) & \alpha > 3 \\ \log(n)^{O(1)} & \alpha = 3 \\ O(n^{3-\alpha}) & \alpha < 3. \end{cases}$$

This term is related to an established property of degree distributions, namely the *second order average degree* [12]. If a random graph with n vertices has expected degrees $w_1, \dots, w_n$ then the second order average degree is defined as $\sum_{i=1}^{n} w_i^2 / \sum_{k=1}^{n} w_k$. In graphs with a power-law degree distribution α we have $w_i = \Theta((n/i)^{1/(\alpha-1)})$. The second order average degree then is $\Theta\big(\sum_{i=1}^{n}(n/i)^{2/(\alpha-1)} / \sum_{k=1}^{n}(n/k)^{1/(\alpha-1)}\big)$. For $\alpha > 3$, this term is constant, for $\alpha = 3$ it is logarithmic, and for $\alpha < 3$ it is polynomial in n [12]. Thus, we can interpret $\tilde{d}_\alpha(n)$ as an estimate of the second order average degree. We prove that the model-checking problem can be solved efficiently if $\tilde{d}_\alpha(n)$ is small.

▶ **Theorem 3.** *There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on every α-power-law-bounded random graph model in expected time $\tilde{d}_\alpha(n)^{f(|\varphi|)}n$.*

The term $\tilde{d}_\alpha(n)$ naturally arises in our proofs and is not a consequence of how we defined the multiplicative factor (i.e., $2^{O(|E|^2)}$, $\log(n)^{O(|E|^2)}$, $O(n^\varepsilon)^{|E|^2}$) in Definition 2. In fact the dependence goes the other way: We defined the factor for each α as large as possible such that it does not dominate the run time of the algorithm. Next we specify exactly those values of α where the previous theorem leads to FPT run times. (In the third case $\varepsilon > 0$ can be chosen arbitrarily small since we require α to be arbitrarily close to 3.)

▶ **Theorem 4.** *Let $(\mathcal{G}_n)_{n\in\mathbf{N}}$ be a random graph model and $\varepsilon > 0$. There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ in expected time*

- $f(|\varphi|)n$ *if $(\mathcal{G}_n)_{n\in\mathbf{N}}$ is α-power-law-bounded for some $\alpha > 3$,*
- $\log(n)^{f(|\varphi|)}n$ *if $(\mathcal{G}_n)_{n\in\mathbf{N}}$ is α-power-law-bounded for $\alpha = 3$,*
- $f(|\varphi|, \varepsilon)n^{1+\varepsilon}$ *for all $\varepsilon > 0$* *if $(\mathcal{G}_n)_{n\in\mathbf{N}}$ is α-power-law-bounded for every $2 < \alpha < 3$.*

This solves the model-checking problem efficiently on a wide range of random graph models. These tractability results are matched by previous intractability results. (Note that the third case of Theorem 4 requires power-law-boundedness for *every* $2 < \alpha < 3$ and thus does not contradict Proposition 5.)

▶ **Proposition 5** ([28] and [26, Lemma 10.3])**.** *For every $2 < \alpha < 3$ there exists an α-power-law-bounded random graph model $(\mathcal{G}_n)_{n\in\mathbf{N}}$ such that one cannot solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on $(\mathcal{G}_n)_{n\in\mathbf{N}}$ in expected* FPT *time unless* AW[$*$] $\subseteq$ FPT/poly*.*

We observe a phase transition in tractability at power-law exponent $\alpha = 3$. Also the run time of our algorithm cannot be linear in n for $\alpha \leq 3$ as a 3-power-law-bounded random graph can have for example $n\log(n)$ edges in expectation. We discuss some of the algorithmic implications of our result for some well-known random graph models in Section 9. More details can be found in Section 10 of [26].

2.3 Structure

Many algorithmic results are based on *structural decompositions*. For example, bidimensionality theory introduced by Demaine et al. [20, 21] is based on the grid minor theorem, which is itself based on a structural decomposition into a clique-sum of almost-embeddable graphs developed by Robertson and Seymour [61]. The model-checking algorithm for graph classes with bounded expansion by Dvořak, Král, and Thomas [30] relies on a structural decomposition of bounded expansion graph classes by Nešetřil and Ossona de Mendez called low tree-depth colorings [63]. Our algorithms are based on a structural decomposition of α-power-law-bounded random graph models.

All algorithms prior to this work rely on showing that a certain graph model is with high probability contained in a certain well-known tractable graph class (for example bounded expansion) and then use the structural decompositions [63] of said graph class. However,

these decompositions were not originally designed with random graphs in mind and therefore may not provide the optimal level of abstraction for random graphs. Our algorithms are based on a specially defined structural decomposition. This direct approach helps us capture random graph models that could otherwise not be captured such as the a.a.s. somewhere dense preferential attachment model. By focusing on α-power-law-bounded random graph models, we obtain structural decompositions for a wide range of models.

We observe that α-power-law-bounded random graphs have mostly an extremely sparse structure with the exception of a part whose size is bounded by the second order average degree. However, this denser part can be separated well from the remaining graph. We show that local regions consist of a core part, bounded in size by the second order average degree, to which trees and graphs of constant size are attached by a constant number of edges. This decomposition is similar to so called *protrusion decompositions*, which have been used by Bodlaender et al. to obtain meta-theorems on kernelization [5]. Our structural decomposition is valid for all graphs that fit into the framework of α-power-law-boundedness, such as preferential attachment graphs or Chung–Lu graphs. We define an approximation of the second order average degree of the degree distribution as $\hat{d}_\alpha(n) = 2$ for $\alpha > 3$, $\hat{d}_\alpha(n) = \log(n)$ for $\alpha = 3$ and $\hat{d}_\alpha(n) = n^{3-\alpha}$ for $\alpha < 3$ (similarly to $\tilde{d}_\alpha(n)$ without O-notation).

▶ **Theorem 6.** *Let $(\mathcal{G}_n)_{n\in\mathbf{N}}$ be an α-power-law-bounded random graph model. There exist constants c, r_0 such that for every $r \geq r_0$ a.a.s. for every r-neighborhood H of $\mathcal{G}_n$ one can partition $V(H)$ into three (possibly empty) sets X, Y, Z with the following properties.*

- *$|X| \leq \hat{d}_\alpha(n)^{cr^2}$.*
- *Every connected component of $H[Y]$ has size at most cr and at most c neighbors in X.*
- *Every connected component of $H[Z]$ is a tree with at most one edge to $H[X \cup Y]$.*

Removing a few vertices makes the local neighborhoods even sparser:

▶ **Corollary 7.** *Let $(\mathcal{G}_n)_{n\in\mathbf{N}}$ be an α-power-law-bounded random graph model. There exist constants c, r_0 such that for every $r \geq r_0$ a.a.s. one can remove $\hat{d}_\alpha(n)^{cr^2}$ vertices from $\mathcal{G}_n$ such that every r-neighborhood has treewidth at most 26.*

Further structural results that may be interesting beyond the purpose of model-checking as well as proofs of the results outlined here can be found in Section 9 of [26]. We now discuss how we use the decomposition of Theorem 6 for our algorithms and why decompositions similar to Corollary 7 are not sufficient for our purposes.

3 Techniques and Outline

A first building block of our algorithm is Gaifman's locality theorem [39]. It implies that in order to solve the first-order model-checking problem on a graph, it is sufficient to solve the problem on all r-neighborhoods of the graph for some small r. We can therefore restrict ourselves to the model-checking problem on the neighborhoods of random graphs. With this in mind, we want to obtain structural decompositions of these neighborhoods.

One important thing to note is that a decomposition according to Corollary 7 is not sufficient. Let us focus on the interesting case $\alpha = 3$ where efficient model-checking is still possible. Corollary 7 then states that the removal of polylogarithmically many vertices yields neighborhoods with treewidth at most 26. While we could easily solve the model-checking problem on graphs with treewidth at most 26 via Courcelle's theorem [16], we cannot solve it on graphs where we need to remove a set X of $\log(n)$ vertices to obtain a treewidth of at most 26. Every vertex not in X may have an arbitrary subset of X as neighborhood. Since

there are $2^{|X|} = n$ possible neighborhoods, we can encode a large complicated structure into this graph by stating that two vertices $i, j \in \mathbf{N}$ are adjacent if and only if there is a vertex whose neighborhood in X represents a binary encoding of the edge ij (omitting some details). Because of this, the model-checking problem on this graph class is as hard as on general graphs. We need the additional requirement that X is only loosely connected to the remaining graph. The decomposition in Theorem 6 fulfills this requirement. Every component of $H \setminus X$ has at most a constant number of neighbors in X.

Let us assume we have decompositions of the neighborhoods of a graph according to Theorem 6 where the sets X are chosen as small as possible. We can now use a variant of the Feferman–Vaught theorem [48] for each r-neighborhood to prune the protrusions and thereby construct a smaller graph that satisfies the same (short) first-order formulas as the original graph, We call this smaller graph the *kernel*. The size of this kernel will be some function of $|X|$. We then use the brute-force model-checking algorithm on the kernel.

For the first steps of the algorithm (decomposition into neighborhoods, kernelization using Feferman–Vaught) one can easily show that they always take FPT time. However, the run time of the last step requires a careful analysis. One can check a formula φ on a graph of size x in time $O(x^{|\varphi|})$ by brute force. Thus, checking the formula on the kernel of all n many r-neighborhoods of a random graph takes expected time at most $n \sum_{x=1}^{n} p_x O(x^{|\varphi|})$, where p_x is the probability that the kernelization procedure on an r-neighborhood of a random graph yields a kernel of size x. In order to guarantee a run time of the form $\log(n)^{f(|\varphi|)} n$ for some function f, p_x should be of order $\log(n)^{f(|\varphi|)} x^{-|\varphi|}$.

Earlier, we discussed that the size of the kernel will be some function of $|X|$ and that we choose X as small as possible. It is therefore sufficient to bound the probability that the set X of the decomposition of a neighborhood exceeds a certain size. Parameterizing the decomposition by two values (denoted by b and μ later on) gives us enough control to guarantee such a bound on p_x. A large part of this work is devoted to proving a good trade-off between the size of the set X of the decomposition and the probability that X is of minimal size. Furthermore, computing the set X is computationally hard, so the whole procedure has to work without knowing the set X, but only its existence.

Our proofs are structured as follows. First, we show in Section 5 that α-power-law-bounded random graph models have the following structure with high probability: They can be partitioned into sets A, B, C, where $A \cup B$ is small, $B \cup C$ is sparse and A and C locally share only few edges. This is done by characterizing this structure by a collection of small forbidden edge-sets and then excluding these edge-sets using the union bound and Definition 2. Then in Section 6 we show that the partition into A, B, C implies the protrusion decomposition of Theorem 6. In Section 7, we partially recover the protrusion decomposition from a given input, and use it to kernelize each r-neighborhood into an equivalent smaller graph. At last, in Section 8, we combine Gaifman's locality theorem with the previous algorithms and probability bounds to obtain our algorithm and bound its run time. Some proofs are quite tedious, but the nature of this problem seems to stop us from using simpler methods.

3.1 Missing Proofs

Many proofs of the results presented in this paper have been omitted. In particular Section 2.3, 5 – 9 sketch only the main ideas behind our results. All missing proofs can be found in the corresponding full version of this paper [26].

4 Notations and Definitions

4.1 Graph Notation

We use common graph theory notation [23]. The *length* of a path equals its number of edges. The *distance* between two vertices u and v ($\mathrm{dist}(u,v)$) equals the length of a shortest path between u and v. For a vertex v let $N_r^G(v)$ be the set of vertices that have distance at most r to v in G. The *radius* of a graph is the minimum among all maximum distances from one vertex to all other vertices. An *r-neighborhood* in G is an induced subgraph of G with radius at most r. The *order* of a graph is $|G| = |V(G)|$. The *size* of a graph is $\|G\| = |V(G) + E(G)|$. The *edge-excess* of a graph G is $|E(G)| - |V(G)|$.

In this work we obtain results for *labeled graphs* [45]. A labeled graph is a tuple $G = (V(G), E(G), P_1(G), \ldots, P_l(G))$ with $P_i(G) \subseteq V(G)$. We call $P_1(G), \ldots, P_l(G)$ the *labels* of G. We say a vertex v is labeled with label $P_i(G)$ if $v \in P_i(G)$. A vertex may have multiple labels. We say the unlabeled simple graph $G' = (V(G), E(G))$ is the *underlying graph* of G and G is a *labeling* of G'. All notations for graphs extend to labeled graphs as expected. The union of two labeled graphs G and H, $(G \cup H)$, is obtained by setting $V(G \cup H) = V(G) \cup V(H)$, $E(G \cup H) = E(G) \cup E(H)$ and for each label $P_i(G \cup H) = P_i(G) \cup P_i(H)$.

For a graph class $\mathcal{G}$, we define $\mathcal{G}_{lb}$ to be the class of all labelings of $\mathcal{G}$. We define $\mathfrak{G}$ to be the class of all simple graphs and $\mathfrak{G}_{lb}$ to be the class of all labeled simple graphs.

4.2 Probabilities and Random Graph Models

We denote probabilities by $\Pr[*]$ and expectation by $\mathrm{E}[*]$. We consider a random graph model to be a sequence of probability distributions. For every $n \in \mathbf{N}$ a random graph model describes a probability distribution on unlabeled simple graphs with n vertices. In order to speak of probability distributions over graphs we fix a sequence of vertices $(v_i)_{i \geq 1}$ and require that a graph with n vertices has the vertex set $\{v_1, \ldots, v_n\}$. A random graph model is a sequence $\mathcal{G} = (\mathcal{G}_n)_{n \in \mathbf{N}}$, where $\mathcal{G}_n$ is a probability distribution over all unlabeled simple graphs G with $V(G) = \{v_1, \ldots, v_n\}$. Even though some random processes naturally lead to graphs with multi-edges or self-loops, we interpret them as simple graphs by removing all self-loops and replacing multiple edges with one single edge. In slight abuse of notation, we also write $\mathcal{G}_n$ for the random variable which is distributed according to $\mathcal{G}_n$. This way, we can lift graph notation to notation for random variables of graphs: For example edge sets and neighborhoods of a random graph $\mathcal{G}_n$ are represented by random variables $E(\mathcal{G}_n)$ and $N_r^{\mathcal{G}_n}(v)$.

4.3 Sparsity

At first, we define nowhere and somewhere density as a property of *graph classes* and then lift the notation to *random graph models*. There are various equivalent definitions and we use the most common definition based on shallow topological minors.

▶ **Definition 8** (Shallow topological minor [63]). *A graph H is an r-shallow topological minor of G if a graph obtained from H by subdividing every edge up to $2r$ times is isomorphic to a subgraph of G. The set of all r-shallow topological minors of a graph G is denoted by $G \widetilde{\nabla} r$. We define the maximum clique size over all shallow topological minors of G as*

$$\omega(G \widetilde{\nabla} r) = \max_{H \in G \widetilde{\nabla} r} \omega(H).$$

▶ **Definition 9** (Nowhere dense [62]). *A graph class $\mathcal{G}$ is nowhere dense if there exists a function f, such that for all $r \in \mathbf{N}$ and all $G \in \mathcal{G}$, $\omega(G \widetilde{\nabla} r) \leq f(r)$.*

▶ **Definition 10** (Somewhere dense [62]). *A graph class $\mathcal{G}$ is somewhere dense if for all functions f there exists an $r \in \mathbf{N}$ and a $G \in \mathcal{G}$, such that $\omega(G \widetilde{\nabla} r) > f(r)$.*

Observe that a graph class is somewhere dense if and only if it is not nowhere dense. We lift these notions to random graph models using the following two definitions.

▶ **Definition 11** (a.a.s. nowhere dense). *A random graph model $\mathcal{G}$ is a.a.s. nowhere dense if there exists a function f such that for all $r \in \mathbf{N}$*

$$\lim_{n\to\infty} \Pr[\omega(\mathcal{G}_n \widetilde{\nabla} r) \leq f(r)] = 1.$$

▶ **Definition 12** (a.a.s. somewhere dense). *A random graph model $\mathcal{G}$ is a.a.s. somewhere dense if for all functions f there is an $r \in N$ such that*

$$\lim_{n\to\infty} \Pr[\omega(\mathcal{G}_n \widetilde{\nabla} r) > f(r)] = 1.$$

While for graph classes the concepts are complementary, a random graph model can both be *neither* a.a.s. somewhere dense *nor* a.a.s. nowhere dense (e.g., if the random graph model is either the empty or the complete graph, both with a probability of 1/2).

4.4 First-Order Logic

We consider only first-order logic over labeled graphs. We interpret a labeled graph $G = (V, E, P_1, \ldots, P_l)$, as a structure with universe V and signature $(E, P_1, \ldots, P_l)$. The binary relation E expresses adjacency between vertices and the unary relations $P_1, \ldots, P_l$ indicate the labels of the vertices. Other structures can be easily converted into labeled graphs. We write $\varphi(x_1, \ldots, x_k)$ to indicate that a formula φ has *free variables* $x_1, \ldots, x_k$. The *quantifier rank* of a formula is the maximum nesting depth of quantifiers in the formula. Two labeled graphs G_1, G_2 with the same signature are *q-equivalent* ($G_1 \equiv_q G_2$) if for every first-order sentence φ with quantifier rank at most q and matching signature holds $G_1 \models \varphi$ if and only if $G_2 \models \varphi$. Furthermore, $|\varphi|$ is the number of symbols in φ. There exists a simple algorithm which decides whether $G \models \varphi$ in time $O(|G|^{|\varphi|})$.

4.5 Model-Checking

With all definitions in place, we can now properly restate the model-checking problem and what it means to solve it efficiently on a random graph model. The model-checking problem on labeled graphs is defined as follows.

p-MC(FO, $\mathfrak{G}_{lb}$)	
Input:	A graph $G \in \mathfrak{G}_{lb}$ and a first-order sentence φ
Parameter:	$\lvert\varphi\rvert$
Problem:	$G \models \varphi$?

Under worst-case complexity, p-MC(FO, $\mathfrak{G}_{lb}$) is AW[$*$]-complete [25] (and PSPACE-complete when unparameterized [71]). We want average case algorithms for p-MC(FO, $\mathfrak{G}_{lb}$) to be efficient for all possible labelings of a random graph model. A function $L : \mathfrak{G} \to \mathfrak{G}_{lb}$ is called a *l-labeling function* for $l \in \mathbf{N}$ if for every $G \in \mathfrak{G}$, $L(G)$ is a labeling of G with up to l labels.

▶ **Definition 1.** *We say* p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ *can be decided on a random graph model* $(\mathcal{G}_n)_{n\in\mathbf{N}}$ *in* expected time $f(|\varphi|, n)$ *if there exists a deterministic algorithm* $\mathcal{A}$ *which decides* p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ *on input* G, φ *in time* $t_{\mathcal{A}}(G, \varphi)$ *and if for all* $n \in \mathbf{N}$, *all first-order sentences* φ *and all* $|\varphi|$*-labeling functions* L, $\mathrm{E}_{G\sim\mathcal{G}_n}\left[t_{\mathcal{A}}(L(G), \varphi)\right] \leq f(|\varphi|, n)$. *We say* p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ *on a random graph model can be decided in* expected FPT time *if it can be decided in expected time* $g(|\varphi|)n^{O(1)}$ *for some function* g.

5 Structure Theorem for Power-Law-Bounded Random Graph Models

The goal of this section is to partition α-power-law-bounded random graph models. We show in Theorem 14 that their vertices can with high probability be partitioned into sets A, B, C with the following properties: The sets A and B are small, the graph $G[B \cup C]$ is locally almost a tree, i.e., has locally only a small edge-excess, and the set B almost separates A from C, i.e., every neighborhood in $G[C]$ has only a small number of edges to A. We call (A, B, C) a b-r-μ-partition. We state the formal definition.

▶ **Definition 13** (b-r-μ-partition)**.** *Let* $b, r, \mu \in \mathbf{N}^+$. *Let* G *be a graph. A tuple* (A, B, C) *is called a* b-r-μ*-partition of* G *if*

1. *the sets* A, B, C *are pairwise disjoint and their union is* $V(G)$,
2. $|A| \leq b$ *and* $|B| \leq b^{\mu}$,
3. *every* $40\mu r$*-neighborhood in* $G[B \cup C]$ *has an edge-excess of at most* μ^2, *and*
4. *for every* $20\mu r$*-neighborhood in* $G[C]$ *there are at most* μ *edges incident to both the neighborhood and to* A.

A graph for which a b-r-μ*-partition exists is called* b-r-μ*-partitionable.*

In summary, B and C are well behaved and the large set C is almost separated from A. Note that the properties of a b-r-μ-partition depend on three parameters b, r, μ. The results of this section imply that our random graphs are asymptotically almost surely b-r-μ-partitionable for $b = \tilde{d}_{\alpha}(n)^{\Omega(1)}$ and constant r, μ. It therefore helps to assume that b is a slowly growing function in n, such as $\log(n)$ and r, μ are constants. Higher values of μ boost the probability of a random graph being b-r-μ-partitionable. The parameter μ is therefore crucial for the design of efficient algorithms. The main result of this section is the following.

▶ **Theorem 14.** *Let* $(\mathcal{G}_n)_{n\in\mathbf{N}}$ *be an* α*-power-law-bounded random graph model and let* $b, r, \mu, n \in \mathbf{N}^+$ *with* $\mu \geq 5$. *The probability that* $\mathcal{G}_n$ *is not* b-r-μ*-partitionable is at most* $\tilde{d}_{\alpha}(n)^{O(\mu^6 r^2)} b^{-\mu^2/10}$.

For proofs of the results in this section, we refer the reader to Section 6 of the full version of this paper [26]. In the following we only sketch the main ideas.

For an α-power-law-bounded random graph model $(\mathcal{G}_n)_{n\in\mathbf{N}}$, we always assume the vertices of $\mathcal{G}_n$ to be $v_1, \ldots, v_n$, ordered as in Definition 2. We will choose $A = \{v_1, \ldots, v_b\}$, $B = \{v_{b+1}, \ldots, v_{b^{\mu}}\}$, $C = \{v_{b^{\mu}+1}, \ldots, v_n\}$ and show that the probability is low that (A, B, C) does not form a b-r-μ-partition. To do so, we first define $\mathcal{H}_n(b, r, \mu)$ to be a set of graphs over the vertex set $\{v_1, \ldots, v_n\}$.

▶ **Definition 15.** *Let* $b, r, \mu, n \in \mathbf{N}^+$. *We define* $\mathcal{H}_n(b, r, \mu)$ *to be the set of*

- *all graphs with vertex set* $V \subseteq \{v_{b+1}, \ldots, v_n\}$ *such that* $|V| \leq 200r\mu^3$, *all vertices have degree at least two, and the graph has an edge-excess of* μ^2, *and*

- *all graphs* $(V_1 \cup V_2, E)$ *such that* $V_1 \subseteq \{v_1, \dots, v_b\}$, $V_2 \subseteq \{v_{b^\mu+1}, \dots, v_n\}$, $|V_1 \cup V_2| \leq 25r\mu^2$, $|E| \leq 25r\mu^2$, $|V_1| \leq \mu$, *all vertices in* V_2 *have degree at least two, and the summed degree of* V_2 *is* $2|V_2| - 2 + \mu$.

We show that if (A, B, C) is not a b-r-μ-partition then the complete edge-set of some graph in $\mathcal{H}_n(b, r, \mu)$ is present in the graph.

▶ **Definition 16.** *Let* G *be a graph and* $\mathcal{H}$ *be a set of graphs over* $V(G)$. *We say* $\mathcal{H} \sqsubseteq G$ *if for some* $H \in \mathcal{H}$, $E(H) \subseteq E(G)$.

▶ **Lemma 17.** *Let* $b, r, \mu, n \in \mathbf{N}^+$. *If a graph* G *with vertex set* $\{v_1, \dots, v_n\}$ *is not* b-r-μ-*partitionable, then* $\mathcal{H}_n(b, r, \mu) \sqsubseteq G$.

For a fixed graph $H \in \mathcal{H}_n(b, r, \mu)$, α-power-law-boundedness (Definition 2) immediately gives us a good bound on the probability $\Pr\big[E(H) \subseteq E(\mathcal{G}_n)\big]$ that its edge-set occurs. Using the union bound over all graphs in $\mathcal{H}_n(b, r, \mu)$ and some tedious calculations (Lemmas 5.6, 5.7, 5.8, 5.9 in [26]) we can bound the probability of the edge-set of any graph from $\mathcal{H}_n(b, r, \mu)$ being present in the random graph model, proving the main result Theorem 14.

At last, we present one more result in which we bound the sum of the expected sizes of all r-neighborhoods in an α-power-law-bounded graph class. This is needed in Section 8 to bound the expected run time of an algorithm that iterates over all r-neighborhoods of a graph.

▶ **Lemma 18.** *Let* $(\mathcal{G}_n)_{n \in \mathbf{N}}$ *be an* α-*power-law-bounded random graph model. Let* $r, \mu, n \in \mathbf{N}^+$ *with* $\mu \geq 5$. *Let* A_b *be the event that* $b \in \mathbf{N}^+$ *is the minimal value such that* $\mathcal{G}_n$ *is* b-r-μ-*partitionable. Then*

$$\mathrm{E}\Big[\sum_{v \in V(\mathcal{G}_n)} \|\mathcal{G}_n[N_r^{\mathcal{G}_n}(v)]\| \;\Big|\; A_b\Big] \Pr[A_b] \leq (r\mu)^{O(r)} \tilde{d}_\alpha(n)^{O(\mu^6 r^2)} b^{-\mu^2/10} n.$$

6 Protrusion Decompositions of Neighborhoods

In this section, we show that local neighborhoods of power-law-bounded graph classes are likely to have the following nice structure: They consist of a (small) core graph to which so called *protrusions* are attached. Protrusions are (possibly large) subgraphs with small treewidth and boundary. The *boundary* of a subgraph is the size of its neighborhood in the remaining graph. Protrusions were introduced by Bodlaender et al. for very general kernelization results in graph classes with bounded genus [5].

Earlier, (Theorem 14) we showed that α-power-law-bounded random graph models are (for certain values of α, b, r, μ) likely to be b-r-μ-partitionable. It is therefore sufficient to show that r-neighborhoods of b-r-μ-partitionable graphs have such a nice protrusion structure.

However, in general it is not easy to find protrusions in a graph [49]. As we later need to be able to find them, we define special protrusion decompositions, called b-r-μ-local-protrusion-partitions in which (most of) the protrusions can be efficiently identified. The main and only result of this section is the following theorem.

▶ **Theorem 19.** *Let* $b, r, \mu \in \mathbf{N}^+$ *and let* G *be a* b-r-μ-*partitionable graph. Let* G^r *be an* r-*neighborhood in* G. *Then* G^r *is* $O(\mu^{17} r^3 b)$-r-$O(\mu)$-*locally-protrusion-partitionable.*

It remains to define what a b-r-μ-local-protrusion-partition of a graph G^r with radius at most r is. The definition has to strike the right balance: It needs to be permissive enough such that neighborhoods of power-law-bounded graph classes are likely to have this structure

and it needs to be restrictive enough to admit efficient algorithms. Informally speaking, a b-r-μ-local-protrusion-partition of a graph G^r is a partition (X, Y, Z) of the vertices of G^r such that X has small size and the connected components of $G^r[Y \cup Z]$ are protrusions. In order to be able to efficiently identify the protrusions, we further require that the components of $G^r[Y]$ have bounded size and the components of $G^r[Z]$ are trees. This is formalized in the following definition.

▶ **Definition 20** (b-r-μ-local-protrusion-partition)**.** *Let $b, r, \mu \in \mathbf{N}^+$. Let G^r be a graph with radius at most r. A tuple (X, Y, Z) is called a b-r-μ-local-protrusion-partition of G^r if*

1. *the sets X, Y, Z are pairwise disjoint and their union is $V(G^r)$,*
2. *$|X| \leq b^\mu$,*
3. *every connected component of $G^r[Y]$ has size at most $r\mu^7$ and at most μ neighbors in X,*
4. *every connected component of $G^r[Z]$ is a tree with at most one edge to $G^r[X \cup Y]$,*
5. *for a subgraph H of $G^r[Y \cup Z]$ we say $N^{G^r}(V(H)) \cap X$ is the boundary of H. The connected components of $G^r[Y]$ may have at most b^μ distinct boundaries, i.e., $|\{N^{G^r}(V(H)) \cap X \mid H \text{ connected component of } G^r[Y \cup Z]\}| \leq b^\mu$.*

A graph for which a b-r-μ-local-protrusion-partition exists is called b-r-μ-locally-protrusion-partitionable.

Property 3 and 4 enforce that the components of $G^r[Y \cup Z]$ are protrusions. Later, we will transform b-r-μ-local-protrusion-partitions into equivalent graphs of bounded size by replacing the protrusions with small graphs. Thus, Property 2 and 5 are there to ensure that the resulting kernelized graph will have size roughly b^μ (without Property 5 we could only guarantee a size of roughly b^{μ^2} which is too large for our purposes).

Proving Theorem 19 involves multiple pages of proofs for which we refer to the full version of this paper [26]. Here, we only sketch the construction of an $O(\mu^{17}r^3b)$-r-$O(\mu)$-local-protrusion-partition. Let a graph G and $b, r, \mu \in \mathbf{N}^+$ be fixed. We further assume G to be b-r-μ-partitionable and we fix a b-r-μ-partition (A, B, C) of G. Let further G^r be an r-neighborhood in G and let $A^r = A \cap V(G^r)$, $B^r = B \cap V(G^r)$, $C^r = C \cap V(G^r)$. We construct an $O(\mu^{17}r^3b)$-r-$O(\mu)$-local-protrusion-partition (X, Y, Z) of G^r by placing all vertices from $A^r \cup B^r$ into X. In order to distribute the vertices C^r to the sets X, Y, and Z, we define so called *ties*.

▶ **Definition 21** (Tie)**.** *Let $W \subseteq B^r \cup C^r$. We say (u_1, u_2, v) is a W-*tie *if $u_1, u_2 \in W$ and v lies on a walk p with the following properties: Every inner vertex of p is contained in C^r and has at least two neighbors in p; u_1 and u_2 are contained only as endpoints of p; and p is contained in a $20\mu r$-neighborhood in $G[B^r \cup C^r]$.*

We use this notion to partition the set C^r. We distinguish vertices connected to A^r, vertices connected to B^r (but not to A^r), those which are connected to neither but lie on a tie, and the rest. We set

- $C^r_A = N(A^r) \cap C^r$,
- $C^r_B = (N(B^r) \setminus N(A^r)) \cap C^r$,
- $C^r_Y = \{v \mid v \in C^r \setminus (C^r_A \cup C^r_B)$ and there exist $u_1, u_2 \in C^r_A \cup C^r_B$ such that (u_1, u_2, v) is a $(C^r_A \cup C^r_B)$-tie $\}$,
- $C^r_Z = C^r \setminus (C^r_A \cup C^r_B \cup C^r_Y)$.

Finally, we define X to be the union of A^r, B^r and all vertices from $C^r_A \cup C^r_B \cup C^r_Y$ which are in a connected component of $G[C^r]$ with more than one edge to B^r. We define Y to be the vertices from $C^r_A \cup C^r_B \cup C^r_Y$ which are in a connected component of $G[C^r]$ with at most one edge to B^r, and we define $Z = C^r_Z$. The fact that (X, Y, Z) is an $O(\mu^{17}r^3b)$-r-$O(\mu)$-local-protrusion-partition is proved in Section 6 of [26].

7 Compressing Neighborhoods

In this section, we kernelize b-r-μ-partitionable graphs. This means we replace the protrusions with subgraphs of bounded size that retain the same boundary. This yields a smaller graph which is q-equivalent to the original graph. The same technique has been used for obtaining small kernels in larger graph classes, e.g., in graphs that exclude a fixed minor [36]. The main result of this section is the following theorem.

▶ **Theorem 22.** *There exists an algorithm that takes $q, r, \mu \in \mathbf{N}^+$ and a connected labeled graph G with radius at most r and at most q labels as input, runs in time at most $f(q,r,\mu)\|G\|$ for some function $f(q,r,\mu)$, and computes a labeled graph $G^* \equiv_q G$. If G is b-r-μ-locally-protrusion-partitionable for some $b \in \mathbf{N}^+$ then $|G^*| \leq f(q,r,\mu)b^\mu$.*

This kernelization procedure and its run time bound is independent in b but the size of the output kernel is not: If b is small, then the output is small. The result is obtained by replacing protrusions with the help of the Feferman–Vaught theorem [48]. However, in order to replace the protrusions, one first has to identify them. The main complication in this section lies in partitioning a graph such that the relevant protrusions can be easily identified. It is crucial that we obtain the size bound $|G^*| \leq f(q,r,\mu)b^\mu$ in Theorem 22. Weaker bounds are easier to obtain but would not be sufficient for our purposes.

We use a variant of the Feferman–Vaught theorem [48] to replace a protrusion by a q-equivalent boundaried graph of minimal size. This size depends only on q and the size of the boundary. The original Feferman–Vaught theorem states that the validity of FO-formulas on the disjoint union or Cartesian product of two graphs is uniquely determined by the value of FO-formulas on the individual graphs. Makowsky adjusted the theorem for algorithmic use [55] in the context of MSO model-checking. The following proposition contains the Feferman–Vaught theorem in a very accessible form. There is also a nice and short proof in [45]. The notation is borrowed from [45], too. At first, we need to define so called q-types.

▶ **Definition 23** ([45]). *Let G be a labeled graph and $\bar{v} = (v_1, \ldots, v_k) \in V(G)^k$, for some nonnegative integer k. The* first-order q-type of $\bar{v}$ in G *is the set $\operatorname{tp}_q^{\mathrm{FO}}(G, \bar{v})$ of all first-order formulas $\psi(x_1, \ldots x_k)$ of rank at most q such that $G \models \psi(v_1, \ldots, v_k)$.*

A q-type could be an infinite set, but one can reduce them to a finite set by syntactically normalizing formulas, so that there are only finitely many normalized formulas of fixed quantifier rank and with a fixed set of free variables. These finitely many formulas can be enumerated. For a tuple $\bar{u} = (u_1, \ldots, u_k)$, we write $\{\bar{u}\}$ for the set $\{u_1, \ldots, u_k\}$. The following is a variant of the Feferman–Vaught theorem [48].

▶ **Proposition 24** ([45, Lemma 2.3]). *Let G, H be labeled graphs and $\bar{u} \in V(G)^k$, such that $V(G) \cap V(H) = \{\bar{u}\}$. Then for all $q \geq 0$, $\operatorname{tp}_q^{\mathrm{FO}}(G \cup H, \bar{u})$ is determined by $\operatorname{tp}_q^{\mathrm{FO}}(G, \bar{u})$ and $\operatorname{tp}_q^{\mathrm{FO}}(H, \bar{u})$.*

We use this proposition to define a q-type preserving protrusion replacement procedure. Assume we identify a protrusion H of a graph with boundary $\bar{u}$. Using Courcelle's theorem [17], we can compute $\operatorname{tp}_q^{\mathrm{FO}}(H, \bar{u})$ by checking all representing formulas. By then enumerating all graphs in ascending order by their size we can find a small graph H' with $\operatorname{tp}_q^{\mathrm{FO}}(H', \bar{u}) = \operatorname{tp}_q^{\mathrm{FO}}(H', \bar{u})$. Proposition 24 now states that we can replace H with H' to obtain a smaller q-equivalent graph. We repeat this for every protrusion we identify. In the full version of this paper [26], this procedure and the proof of Theorem 22 is presented in detail (Section 7).

8 Model-Checking

In this section, we finally obtain the main result of this paper, namely that for certain values of α one can perform model-checking on α-power-law-bounded random graph models in efficient expected time.

An important tool in this section is Gaifman's locality theorem [39]. It states that first-order formulas can express only local properties of graphs. It is a well established tool for the design of model-checking algorithms (e.g. [44, 45, 37]). We use it to reduce the model-checking problem on a graph to the model-checking problem on neighborhoods of said graph [26, Lemma 8.2]. This technique is described well by Grohe [45, section 5].

To illustrate our approach, consider the following thought experiment: Let X be a non-negative random variable with $\Pr[X = b] = \Theta(b^{-10})$ for all $b \in \mathbf{N}$. Assume an algorithm that gets an integer $b \in \mathbf{N}$ as input and runs in time $t(b)$. Its expected run time on input X is $\sum_{b\in\mathbf{N}} \Theta(b^{-10})t(X)$. If $t(b) = b^{10}$ then the expected run time is infinite. If $t(b) = b^8$ then the expected run time is $\Theta(1)$. Thus, small polynomial differences in the run time can have a huge impact on the expected run time. We notice that the run time on an input has to grow slower than the inverse of the probability that the input occurs.

Let us fix a formula φ and let r and μ be constants depending on φ. In this section we provide a model-checking algorithm whose run time on a graph G depends on the minimal value $b \in \mathbf{N}$ such that G is b-r-μ-partitionable. This means, we need to solve the model-checking problem on b-r-μ-partitionable graphs faster than the inverse of the probability that b is minimal.

Section 5 states that a graph from power-law-bounded graph classes is for some b not b-r-μ-partitionable with probability approximately $b^{-\mu^2}$ (we ignore the terms in r, μ and $\tilde{d}_\alpha(n)$ for now). Thus, the probability that a value b is minimal is approximately $b^{-\mu^2}$.

Let G be a graph and a be the minimal value such that G is b-r-μ-partitionable. In Section 6 we showed that all its r-neighborhoods are $O(\mu^{17}r^3b)$-r-$O(\mu)$-locally-protrusion-partitionable. The kernelization result from Section 7 states that such r-neighborhoods can be converted in linear time into $|\varphi|$-equivalent graphs of size approximately b^μ (we again ignore the factors independent of b for now). This means, using the naive model-checking algorithm, one can decide for an r-neighborhood G^r of G whether $G^r \models \varphi$ in time approximately $\|G\|b^{\mu|\varphi|}$. Thus, one can perform model-checking on all r-neighborhoods of G in time approximately $b^{\mu|\varphi|}\sum_v \|N_r^G(v)\|$. Using Gaifman's locality theorem, this (more or less) yields the answer to the model checking problem in the *whole* graph. Let G be a graph from a power-law-bounded random graph model. In summary, we have for every $b \in \mathbf{N}$:

- $b \in \mathbf{N}$ is the minimal value such that a graph is b-r-μ-partitionable with probability approximately $b^{-\mu^2}$.
- If $b \in \mathbf{N}$ is the minimal value such that G is b-r-μ-partitionable then we can decide whether $G \models \varphi$ in time approximately $b^{\mu|\varphi|}\sum_v \|N_r^G(v)\|$.

In this example one may choose $\mu = |\varphi|^2$ such that the run time grows slower than the inverse of the probability. We changed some numbers in these examples to simplify our arguments. Thus, in reality, μ needs to be chosen slightly differently.

The proofs of this section proceed as follows. We first use slightly nonstandard version of Gaifman locality [26, Lemma 8.2] and the kernelization result in Theorem 22 to solve the model-checking in b-r-μ-partitionable graphs, obtaining the following result (see [26, Lemma 8.4] for the proof).

▶ **Lemma 25.** *Let $\mu \in \mathbf{N}^+$. There exist functions ρ and f such that for every first-order sentence φ and every labeled graph G with at most $|\varphi|$ labels one can decide whether $G \models \varphi$ in time $f(\rho(|\varphi|), \mu) b^{\mu\rho(|\varphi|)} \sum_{v \in V(G)} \|G[N^G_{\rho(|\varphi|)}(v)]\|$, where $b \in \mathbf{N}^+$ is the minimal value such that G is b-$\rho(r)$-μ-partitionable.*

The run time of this algorithm depends not only on b but also on the sum of the sizes of all neighborhoods in a graph, which might be quadratic in the worst case. In order to get almost linear expected run time, we bound the expectation of this value in Lemma 18. We can now prove our main result.

▶ **Theorem 3.** *There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on every α-power-law-bounded random graph model in expected time $\tilde{d}_\alpha(n)^{f(|\varphi|)} n$.*

Proof. Let $(\mathcal{G}_n)_{n \in \mathbf{N}}$ be an α-power-law-bounded random graph model and φ be a first-order formula. We fix a $|\varphi|$-labeling function L and $n \in \mathbf{N}$. We consider labeled graphs with vertices $V(\mathcal{G}_n)$ whose underlying graph is distributed according to $\mathcal{G}_n$, and analyze the expected run time of the model-checking algorithm from Lemma 25 on these graphs.

Let ρ be the function from Lemma 25 and let $r = \rho(|\varphi|)$ and $\mu = \rho(|\varphi|)^2 + 100$. For every graph G there exists a value $b \in \mathbf{N}^+$ such that G is b-r-μ-partitionable (i.e., by setting $b = |V(G)|$, $A = V(G)$). Let A_b be the event that $b \in \mathbf{N}^+$ is the minimal value such that $\mathcal{G}_n$ is b-r-μ-partitionable and let R be the expected run time of the model-checking algorithm from Lemma 25. The expected run time of the algorithm is exactly $\sum_{b=1}^{\infty} \mathrm{E}[R \mid A_b] \Pr[A_b]$. We use Lemma 25 and 18 to bound

$$
\begin{aligned}
&\sum_{b=1}^{\infty} \mathrm{E}[R \mid A_b] \Pr[A_b] \\
\leq &\sum_{b=1}^{\infty} \mathrm{E}\Big[f'(r,\mu) b^{r\mu} \sum_{v \in V(\mathcal{G}_n)} \|\mathcal{G}_n[N^{\mathcal{G}_n}_r(v)]\| \mid A_b\Big] \Pr[A_b] \\
= &\sum_{b=1}^{\infty} f'(r,\mu) b^{r\mu} \,\mathrm{E}\Big[\sum_{v \in V(\mathcal{G}_n)} \|\mathcal{G}_n[N^{\mathcal{G}_n}_r(v)]\| \mid A_b\Big] \Pr[A_b] \\
\leq &\sum_{b=1}^{\infty} f'(r,\mu) b^{r\mu} (200 r \mu^3)^{O(r)} \tilde{d}_\alpha(n)^{O(\mu^6 r^2)} b^{-\mu^2/10} n \\
= &f'(r,\mu)(200 r \mu^3)^{O(r)} \tilde{d}_\alpha(n)^{O(\mu^6 r^2)} n \sum_{b=1}^{\infty} b^{-\mu^2/10 + r\mu}.
\end{aligned}
$$

Note that for $\mu = \rho(|\varphi|)^2 + 100$ and $r = \rho(|\varphi|)$ we have $\sum_{b=1}^{\infty} b^{-\mu^2/10 + r\mu} \leq \sum_{b=1}^{\infty} b^{-2} = O(1)$. This yields a run time of $\tilde{d}_\alpha(n)^{f(|\varphi|)} n$ for some function f. ◀

By substituting the values of $\tilde{d}_\alpha(n)$ and distinguishing three cases we obtain the following alternative form of our main result.

▶ **Theorem 4.** *Let $(\mathcal{G}_n)_{n \in \mathbf{N}}$ be a random graph model and $\varepsilon > 0$. There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ in expected time*

- $f(|\varphi|) n$ *if $(\mathcal{G}_n)_{n \in \mathbf{N}}$ is α-power-law-bounded for some $\alpha > 3$,*
- $\log(n)^{f(|\varphi|)} n$ *if $(\mathcal{G}_n)_{n \in \mathbf{N}}$ is α-power-law-bounded for $\alpha = 3$,*
- $f(|\varphi|, \varepsilon) n^{1+\varepsilon}$ *for all $\varepsilon > 0$ if $(\mathcal{G}_n)_{n \in \mathbf{N}}$ is α-power-law-bounded for every $2 < \alpha < 3$.*

9 Implications for Various Graph Models

A wide range of unclustered random graph models are α-power-law-bounded. In this section, we discuss the implications of our result for preferential attachment, Chung–Lu, and Erdős–Rényi graphs. A more detailed discussion with proofs can be found in Section 10 in the full version [26].

9.1 Preferential Attachment Model

The preferential attachment model [3, 64] may be the best-known model for complex networks. In this model, graphs are created by a random process that iteratively adds new vertices and randomly connects them to already existing ones, where the attachment probability is proportional to the current degree of a vertex. The preferential attachment process exhibits small world behavior [24] and has been widely recognized as a reasonable explanation of the heavy tailed degree distribution of complex networks [7]. It has a vanishing clustering coefficient [8], but there exist extensions of the model that admit clustering [73].

Recent efficient model-checking algorithms on random graph models only worked on random graph models that asymptotically almost surely (a.a.s.) are nowhere dense [44, 22]. It is known that preferential attachment graphs are not a.a.s. nowhere dense [22] and even a.a.s. somewhere dense [27], thus previous techniques do not work.

We define G^n_m to be the preferential attachment graph with n vertices and m edges per vertex. Usually, the parameter m is considered to be constant. We obtain efficient algorithms even if we allow m to be a function of the size of the network. For a function $m(n)\colon \mathbf{N} \to \mathbf{N}$ we define $(G^n_{m(n)})_{n\in\mathbf{N}}$ to be the random graph model where the number of edges per vertex grows according to $m(n)$. Previous work [29] implies that this model is α-power-law-bounded ([26, Lemma 10.1]), which immediately implies the following model checking result.

▶ **Corollary 26.** *Let $m\colon \mathbf{N} \to \mathbf{N}$. There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on the preferential attachment model $(G^n_{m(n)})_{n\in\mathbf{N}}$ in expected time*

- $\log(n)^{f(|\varphi|)}n$ *if* $m(n) = \log(n)^{O(1)}$,
- $f(|\varphi|,\varepsilon)n^{1+\varepsilon}$ *for every* $\varepsilon > 0$ *if* $m(n) = O(n^\varepsilon)$ *for every* $\varepsilon > 0$.

9.2 Chung–Lu Model

This model generates random graphs that fit a certain degree sequence and has been studied extensively [12, 13, 14]. The degree sequence is defined by a power-law distribution with exponent α. One can easily show that this model is α-power-law-bounded [26, Lemma 10.3]. We can therefore characterize the tractability of the labeled model-checking problem on Chung–Lu graphs based on α.

▶ **Corollary 27.** *Let $\mathcal{G}$ be the Chung–Lu random graph model with exponent α. There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on $\mathcal{G}$ in expected time*

- $f(|\varphi|)n$ *if* $\alpha > 3$,
- $\log(n)^{f(|\varphi|)}n$ *if* $\alpha = 3$.

Furthermore, if $2.5 \leq \alpha < 3$, $\alpha \in \mathbf{Q}$ then one cannot solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on $\mathcal{G}$ in expected FPT *time unless* AW$[*] \subseteq$ FPT/poly.

Previously, the model-checking problem has been known to be tractable on Chung–Lu graphs with exponent $\alpha > 3$, and hard on Chung–Lu graphs with exponent $2.5 \leq \alpha < 3$. The important case $\alpha = 3$ was open. Furthermore, the previous tractability result assumes

the maximum expected degree of a Chung–Lu graph with exponent α to be at most $O(n^{1/\alpha})$, while in the canonical definition of Chung–Lu graphs it is $\Theta(n^{1/(\alpha-1)})$. Our results hold for the canonical definition. The missing case $\alpha < 2.5$ is still open. We believe it can be proven to be hard with similar techniques as for $2.5 \leq \alpha < 3$.

9.3 Erdős–Rényi Model

One of the earliest and most intensively studied random graph models is the Erdős–Rényi model [31]. We say $G(n, p(n))$ is a random graph with n vertices where each pair of vertices is connected independently uniformly at random with probability $p(n)$. Many properties of Erdős–Rényi graphs are well studied, including but not limited to, threshold phenomena, the sizes of components, diameter, and length of paths [9]. With a three-line argument [26, Lemma 10.8], we obtain a fine grained picture over the tractability of the model-checking problem on sparse Erdős–Rényi graphs.

▶ **Corollary 28.** *There exists a function f such that one can solve p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ on $G(n, p(n))$ in expected time*

- $f(|\varphi|)n$ *if* $p(n) = O(1/n)$,
- $\log(n)^{f(|\varphi|)}n$ *if* $p(n) = \log(n)^{O(1)}/n$,
- $f(|\varphi|, \varepsilon)n^{1+\varepsilon}$ *for every* $\varepsilon > 0$ *if* $p(n) = O(n^{\varepsilon}/n)$ *for every* $\varepsilon > 0$.

The third case has been shown previously by Grohe [44]. Furthermore, under reasonable assumptions (AW[$*$] $\not\subseteq$ FPT/poly) we know that p-MC$(\mathrm{FO}, \mathfrak{G}_{lb})$ cannot be decided in expected FPT time on denser Erdős–Rényi graphs with $p(n) = n^{\delta}/n$ for some $0 < \delta < 1$, $\delta \in \mathbf{Q}$ [28].

10 Conclusion

We define α-power-law-bounded random graphs which generalize many unclustered random graphs models. We provide a structural decomposition of neighborhoods of these graphs and use it to obtain a meta-algorithm for deciding first-order properties in the the preferential attachment-, Erdős–Rényi-, Chung–Lu- and configuration random graph model.

There are various factors to consider when evaluating the practical implications of this result. The degree distribution of most real world networks is similar to a power-law distribution with exponent between two and three [15], but our algorithm is only fast for exponents at least three. This leaves many real world networks where our algorithm is slow. However, it has been shown that the model-checking problem (with labels) becomes hard on these graphs if we assume independently distributed edges [28].

So far, we do not know whether the model-checking problem is hard or tractable on clustered random graphs. If a random graph model is 3-power-law-bounded then one can show that the expected number of triangles is polylogarithmic (via union bound of all possible embeddings of a triangle). Therefore, random models with clustering, such as the Kleinberg model [50], the hyperbolic random graph model [53, 11], or the random intersection graph model [47], which have a high number of triangles currently do not fit into our framework (see [26, Section 10.5] for a proof that random intersection graphs are not α-power-law-bounded for any α). This is unfortunate, since clustering is a key aspect of real networks [72]. In the future, we hope to extend our results to clustered random graph models. We observe that some clustered random graph models can be expressed as *first-order transductions* of α-power-law-bounded random graph models. For example the random intersection graph model is a transduction of a sparse Erdős–Rényi graph. We believe this connection can be

used to transfer tractability results to clustered random graphs. If we can efficiently compute for a clustered random graph model $\mathcal{G}$ a pre-image of a transduction that is distributed like an α-power-law-bounded random graph then we can efficiently solve p-MC(FO, $\mathfrak{G}_{lb}$) on $\mathcal{G}$. The same idea is currently being considered for solving the model checking problem for transductions of sparse graph classes (e.g. structurally bounded expansion classes) [40].

In our algorithm, we use Gaifman's locality theorem to reduce our problem to r-neighborhoods of the input graph. In this construction the value of r can be exponential in the length of the formula [39]. On the other hand, the small world property states that the radius of real networks is rather small. This means, even for short formulas our neighborhood-based approach may practically be working on the whole graph instead of neighborhoods. It would be interesting to analyze for which values of r practical protrusion decompositions according to Theorem 6 exist in the real world.

At last, a big problem with all parameterized model-checking algorithms is their large run time dependence on the length of the formula. Grohe and Frick showed that already on trees every first-order model-checking algorithm takes worst-case time at least $f(|\varphi|)n$ where f is a non-elementary tower function [38]. So far, it is unclear whether this also holds in the average-case setting. The results presented in this paper have a non-elementary dependence on the length of the formula. We are curious whether one can find average-case model-checking algorithms with elementary expected FPT run time. In summary, many more obstacles need to be overcome to obtain a truly practical general purpose meta-algorithm for complex networks.

References

1 Réka Albert, Hawoong Jeong, and Albert-László Barabási. Internet: Diameter of the world-wide web. *Nature*, 401(6749):130, 1999.

2 Sanjeev Arora and Boaz Barak. *Computational complexity: A modern approach*. Cambridge University Press, 2009.

3 Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

4 Thomas Bläsius, Tobias Friedrich, and Anton Krohmer. Hyperbolic random graphs: Separators and treewidth. In *24th Annual European Symposium on Algorithms (ESA 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.

5 Hans L Bodlaender, Fedor V Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M Thilikos. (Meta) kernelization. *Journal of the ACM (JACM)*, 63(5):44, 2016.

6 Andrej Bogdanov and Luca Trevisan. Average-Case Complexity. *Foundations and Trends in Theoretical Computer Science*, 2(1):1–106, 2006.

7 Béla Bollobás, Oliver Riordan, Joel Spencer, and Gábor Tusnády. The degree sequence of a scale-free random graph process. *Random Structures & Algorithms*, 18(3):279–290, May 2001.

8 Béla Bollobás and Oliver M Riordan. Mathematical results on scale-free random graphs. *Handbook of graphs and networks: from the genome to the internet*, pages 1–34, 2003.

9 Béla Bollobás. *Random Graphs*. Cambridge University Press, 2nd edition, 2001.

10 Anna D. Broido and Aaron Clauset. Scale-free networks are rare. *Nature communications*, 10(1):1017, 2019.

11 Elisabetta Candellero and Nikolaos Fountoulakis. Clustering and the hyperbolic geometry of complex networks. *Internet Mathematics*, 12(1-2):2–53, 2016.

12 Fan Chung and Linyuan Lu. The average distances in random graphs with given expected degrees. *Proc. of the National Academy of Sciences*, 99(25):15879–15882, 2002.

13 Fan Chung and Linyuan Lu. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics*, 6(2):125–145, 2002.

14 Fan Chung and Linyuan Lu. *Complex graphs and networks*, volume 107. American Math. Soc., 2006.

15 Aaron Clauset, Cosma Rohilla Shalizi, and Mark E. J. Newman. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4):661–703, 2009.

16 Bruno Courcelle. The monadic second-order logic of graphs I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990.

17 Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000. doi:10.1007/s002249910009.

18 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.

19 Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally Excluding a Minor. In *Proceedings of the 22nd Symposium on Logic in Computer Science*, pages 270–279, 2007.

20 Erik D. Demaine, Fedor V. Fomin, Mohammadtaghi Hajiaghayi, and Dimitrios M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs. *J. ACM*, 52(6):866–893, November 2005. doi:10.1145/1101821.1101823.

21 Erik D. Demaine and M. Hajiaghayi. The bidimensionality theory and its algorithmic applications. *Comput. J.*, 51(3):292–302, 2008.

22 Erik D. Demaine, Felix Reidl, Peter Rossmanith, Fernando Sánchez Villaamil, Somnath Sikdar, and Blair D. Sullivan. Structural sparsity of complex networks: Bounded expansion in random models and real-world graphs. *J. Comput. Syst. Sci.*, 105:199–241, 2019. doi:10.1016/j.jcss.2019.05.004.

23 R. Diestel. *Graph Theory*. Springer, Heidelberg, 2010.

24 Sander Dommers, Remco van der Hofstad, and Gerard Hooghiemstra. Diameters in preferential attachment models. *Journal of Statistical Physics*, 139(1):72–107, 2010.

25 Rod G. Downey, Michael R. Fellows, and Udayan Taylor. The Parameterized Complexity of Relational Database Queries and an Improved Characterization of W[1]. *DMTCS*, 96:194–213, 1996.

26 Jan Dreier, Philipp Kuinke, and Peter Rossmanith. First-order model-checking in random graphs and complex networks, 2020. arXiv:2006.14488.

27 Jan Dreier, Philipp Kuinke, and Peter Rossmanith. Maximum shallow clique minors in preferential attachment graphs have polylogarithmic size. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 176 of *LIPIcs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

28 Jan Dreier and Peter Rossmanith. Hardness of FO model-checking on random graphs. In *14th International Symposium on Parameterized and Exact Computation, IPEC 2019, September 11-13, 2019, Munich, Germany*, volume 148 of *LIPIcs*, pages 11:1–11:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.IPEC.2019.11.

29 Jan Dreier and Peter Rossmanith. Motif counting in preferential attachment graphs. In *39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2019, December 11-13, 2019, Bombay, India*, volume 150 of *LIPIcs*, pages 13:1–13:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.FSTTCS.2019.13.

30 Zdenek Dvořak, Daniel Král, and Robin Thomas. Deciding First-Order Properties for Sparse Graphs. In *Proceedings of the 51st Conference on Foundations of Computer Science*, pages 133–142, 2010.

31 P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.

32 Ronald Fagin. Probabilities on finite models 1. *The Journal of Symbolic Logic*, 41(1):50–58, 1976.

33 Matthew Farrell, Timothy D Goodrich, Nathan Lemons, Felix Reidl, Fernando Sánchez Villaamil, and Blair D Sullivan. Hyperbolicity, degeneracy, and expansion of random intersection graphs. In *International Workshop on Algorithms and Models for the Web-Graph*, pages 29–41. Springer, 2015.

34 Jörg Flum, Markus Frick, and Martin Grohe. Query Evaluation via Tree-Decompositions. *Journal of the ACM (JACM)*, 49(6):716–752, 2002.

35 Jörg Flum and Martin Grohe. Fixed-Parameter Tractability, Definability, and Model-Checking. *SIAM Journal on Computing*, 31(1):113–145, 2001.

36 Fedor V Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M Thilikos. Bidimensionality and kernels. In *Proc. of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 503–510, 2010.

37 Markus Frick and Martin Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM (JACM)*, 48(6):1184–1206, 2001.

38 Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. *Annals of pure and applied logic*, 130(1-3):3–31, 2004.

39 Haim Gaifman. On local and non-local properties. In *Studies in Logic and the Foundations of Mathematics*, volume 107, pages 105–135. Elsevier, 1982.

40 Jakub Gajarský, Petr Hliněný, Jan Obdrzálek, Daniel Lokshtanov, and M. S. Ramanujan. A new perspective on FO model checking of dense graph classes. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 176–184, 2016. doi:10.1145/2933575.2935314.

41 Yong Gao. Treewidth of Erdős–Rényi random graphs, random intersection graphs, and scale-free random graphs. *Discrete Applied Mathematics*, 160(4-5):566–578, 2012.

42 Yu V Glebskii, DI Kogan, MI Liogon'kii, and VA Talanov. Range and degree of realizability of formulas in the restricted predicate calculus. *Cybernetics and Systems Analysis*, 5(2):142–154, 1969.

43 Anna Goldenberg, Alice X. Zheng, Stephen E. Fienberg, Edoardo M. Airoldi, et al. A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2):129–233, 2010.

44 Martin Grohe. Generalized model-checking problems for first-order logic. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 12–26. Springer, 2001.

45 Martin Grohe. Logic, graphs, and algorithms. *Logic and Automata*, 2:357–422, 2008.

46 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3), 2017.

47 Michał Karoński, Edward R. Scheinerman, and Karen B. Singer-Cohen. On random intersection graphs: The subgraph problem. *Combinatorics, Probability and Computing*, 8(1-2):131–159, 1999.

48 Carol Karp. The first order properties of products of algebraic systems. fundamenta mathematicae. *Journal of Symbolic Logic*, 32(2):276–276, 1967. doi:10.2307/2271704.

49 Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau, and Somnath Sikdar. Linear kernels and single-exponential algorithms via protrusion decompositions. *ACM Transactions on Algorithms (TALG)*, 12(2):21, 2016.

50 Jon Kleinberg. The Small-World Phenomenon: An Algorithmic Perspective. In *Proceedings of the 32nd Symposium on Theory of Computing*, pages 163–170, 2000.

51 Jon M. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845–845, 2000.

52 Stephan Kreutzer. Algorithmic meta-theorems. In *International Workshop on Parameterized and Exact Computation*, pages 10–12. Springer, 2008.

53 Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.

54 Leonid A. Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.

55 Johann A. Makowsky. Algorithmic uses of the feferman–vaught theorem. *Annals of Pure and Applied Logic*, 126(1-3):159–213, 2004.

56 Stanley Milgram. The small world problem. *Psychology Today*, 2(1):60–67, 1967.

57 Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

58 Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proc. of the 7th ACM SIGCOMM Conference on Internet Measurement*, pages 29–42. ACM, 2007.

59 M. Molloy and B. A. Reed. The size of the giant component of a random graph with a given degree sequence. *Combin., Probab. Comput.*, 7(3):295–305, 1998.

60 Michael Molloy and Bruce Reed. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms*, 6(2-3):161–180, 1995.

61 Paul D. Seymour N. Robertson. Graph minors XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89:43–76, 2003.

62 Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity*. Springer, 2012.

63 Jaroslav Nešetřil and Patrice Ossona de Mendez. Grad and classes with bounded expansion I. Decompositions. *European Journal of Combinatorics*, 29(3):760–776, 2008.

64 Derek de Solla Price. A general theory of bibliometric and other cumulative advantage processes. *Journal of the American society for Information science*, 27(5):292–306, 1976.

65 Nataša Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007.

66 Katarzyna Rybarczyk. Diameter, connectivity, and phase transition of the uniform random intersection graph. *Discrete Mathematics*, 311(17):1998–2019, 2011.

67 Satu Elisa Schaeffer. Graph clustering. *Computer Science Review*, 1(1):27–64, 2007.

68 Nicole Schweikardt, Luc Segoufin, and Alexandre Vigny. Enumeration for FO queries over nowhere dense graphs. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10-15, 2018*, pages 151–163. ACM, 2018. doi:10.1145/3196959.3196971.

69 Detlef Seese. Linear time computable problems and first-order descriptions. *Math. Struct. in Comp. Science*, 6:505–526, 1996.

70 Joel Spencer. *The strange logic of random graphs*, volume 22. Springer Science & Business Media, 2013.

71 Larry J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

72 Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440, 1998.

73 Konstantin Zuev, Marián Boguná, Ginestra Bianconi, and Dmitri Krioukov. Emergence of soft communities from geometric preferential attachment. *Scientific reports*, 5:9421, 2015.