

Efficiently Generating Geometric Inhomogeneous and Hyperbolic Random Graphs

Thomas Bläsius

Hasso Plattner Institute, Potsdam, Germany
thomas.blaesius@hpi.de

Tobias Friedrich

Hasso Plattner Institute, Potsdam, Germany
tobias.friedrich@hpi.de

Maximilian Katzmann

Hasso Plattner Institute, Potsdam, Germany
maximilian.katzmann@hpi.de

Ulrich Meyer

Goethe University, Frankfurt, Germany
umeyer@ae.cs.uni-frankfurt.de

Manuel Penschuck

Goethe University, Frankfurt, Germany
mpenschuck@ae.cs.uni-frankfurt.de

Christopher Weyand

Hasso Plattner Institute, Potsdam, Germany
christopher.weyand@hpi.de

Abstract

Hyperbolic random graphs (HRG) and geometric inhomogeneous random graphs (GIRG) are two similar generative network models that were designed to resemble complex real world networks. In particular, they have a power-law degree distribution with controllable exponent β , and high clustering that can be controlled via the temperature T .

We present the first implementation of an efficient GIRG generator running in expected linear time. Besides varying temperatures, it also supports underlying geometries of higher dimensions. It is capable of generating graphs with ten million edges in under a second on commodity hardware. The algorithm can be adapted to HRGs. Our resulting implementation is the fastest sequential HRG generator, despite the fact that we support non-zero temperatures. Though non-zero temperatures are crucial for many applications, most existing generators are restricted to $T = 0$. We also support parallelization, although this is not the focus of this paper. Moreover, we note that our generators draw from the correct probability distribution, i.e., they involve no approximation.

Besides the generators themselves, we also provide an efficient algorithm to determine the non-trivial dependency between the average degree of the resulting graph and the input parameters of the GIRG model. This makes it possible to specify the desired expected average degree as input.

Moreover, we investigate the differences between HRGs and GIRGs, shedding new light on the nature of the relation between the two models. Although HRGs represent, in a certain sense, a special case of the GIRG model, we find that a straight-forward inclusion does not hold in practice. However, the difference is negligible for most use cases.

2012 ACM Subject Classification Theory of computation → Generating random combinatorial structures

Keywords and phrases hyperbolic random graphs, geometric inhomogeneous random graph, efficient network generation

Digital Object Identifier 10.4230/LIPIcs.ESA.2019.19

Related Version A full version of the paper is available at [3], <http://arxiv.org/abs/1905.06706>.



© Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, Ulrich Meyer, Manuel Penschuck, Christopher Weyand;

licensed under Creative Commons License CC-BY

27th Annual European Symposium on Algorithms (ESA 2019).

Editors: Michael A. Bender, Ola Svensson, and Grzegorz Herman; Article No. 19; pp. 19:1–19:14

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Network models play an important role in different fields of science [8]. From the perspective of network science, models can be used to explain observed behavior in the real world. To mention one example, Watts and Strogatz [26] observed that few random long-range connections suffice to guarantee a small diameter. This explains why many real-world networks exhibit the small-world property despite heavily favoring local over long-range connections. From the perspective of computer science, and specifically algorithmics, realistic random networks can provide input instances for graph algorithms. This facilitates theoretical approaches (e.g., average-case analysis), as well as extensive empirical evaluations by providing an abundance of benchmark instances, solving the pervasive scarcity of real-world instances.

There are some crucial features that make a network model useful. The generated instances have to resemble real-world networks. The model should be as simple and natural as possible to facilitate theoretical analysis, and to prevent untypical artifacts. And it should be possible to efficiently draw networks from the model. This is particularly important for the empirical analysis of model properties and for generating benchmark instances.

A model that has proven itself useful in recent years is the *hyperbolic random graph (HRG)* model [17]. HRGs are generated by drawing vertex positions uniformly at random from a disk in the hyperbolic plane. Two vertices are joined by an edge if and only if their distance lies below a certain threshold; see Section 2.2. HRGs resemble real-world networks with respect to crucial properties. Most notable are the *power-law degree distribution* [15] (i.e., the number of vertices of degree k is roughly proportional to $k^{-\beta}$ with $\beta \in (2, 3)$), the high *clustering coefficient* [15] (i.e., two vertices are more likely to be connected if they have a common neighbor), and the small diameter [11, 19]. Moreover, HRGs are accessible for theoretical analysis (see, e.g., [15, 11, 19, 4]). Finally there is a multitude of efficient generators with different emphases [2, 24, 23, 25, 21, 13, 12]; see Section 1.2 for a discussion.

Closely related to HRGs is the *geometric inhomogeneous random graph (GIRG)* model [7]. Here every vertex has a position on the d -dimensional torus and a weight following a power law. Two vertices are then connected if and only if their distance on the torus is smaller than a threshold based on the product of their weights. When using positions on the circle ($d = 1$), GIRGs approximate HRGs in the following sense: the processes of generating a HRG and a GIRG can be coupled such that it suffices to decrease and increase the average degree of the GIRG by only a constant factor to obtain a subgraph and a supergraph of the corresponding HRG, respectively. Compared to HRGs, GIRGs are potentially easier to analyze, generalize nicely to higher dimensions, and the weights allow to directly adjust the degree distribution.

Above, we described the idealized *threshold variants* of the models, where two vertices are connected if and only if their distance is small enough. Arguably more realistic are the *binomial variants*, which allow longer edges and shorter non-edges with a small probability. This is achieved with an additional parameter T , called *temperature*. For $T \rightarrow 0$, the binomial and threshold variants coincide. Many publications focus on the threshold case, as it is typically simpler. This is particularly true for generation algorithms: in the threshold variants one can ignore all vertex pairs with sufficient distance, which can be done using geometric data structures. In the binomial case, any pair of vertices could be adjacent, and the search space cannot be reduced as easily. For practical purposes, however, a non-zero temperature is crucial as real-world networks are generally assumed to have positive temperature. Moreover, from an algorithmic perspective, the threshold variants typically produce particularly well-behaved instances, while a higher temperature leads to difficult problem inputs. Thus, to obtain benchmark instances of varying difficulty, generators for the binomial variants are key.

1.1 Contribution & Outline

Based on the algorithm by Bringmann, Keusch, and Lengler [7], we provide an efficient and flexible GIRG generator. It includes the binomial case and allows higher dimensions. Its expected running time is linear. To the best of our knowledge, this is the first efficient generator for the GIRG model. Moreover, we adapt the algorithm to the HRG model, including the binomial variant. Compared to existing HRG generators (most of which only support the threshold variant), our implementation is the fastest sequential HRG generator.

A refactoring of the original GIRG algorithm [7] allows us to parallelize our generators. They do not use multiple processors as effectively as the threshold-HRG generator by Penschuck [21], which was specifically tailored towards parallelism. However, in a setting realistic for commodity hardware (8 cores, 16 threads), we still achieve comparable run times.

Our generators come as an open-source C++ library¹ with documentation, command-line interface, unit tests, micro benchmarks, and OpenMP [6] parallelization using shared memory. An integration into NetworkKit [22] is available.

Besides the efficient generators, we have three secondary contributions. (I) We provide a comprehensible description of the sampling algorithm that should make it easy to understand how the algorithm works, why it works, and how it can be implemented. Although the core idea of the algorithm is not new [7], the previous description is somewhat technical. (II) The expected average degree can be controlled via an input parameter. However, the dependence of the average degree on the actual parameter is non-trivial. In fact, given the average degree, there is no closed formula to determine the parameter. We provide a linear-time algorithm to estimate it. (III) We investigate how GIRGs and HRGs actually relate to each other by measuring how much the average degree of the GIRG has to be decreased and increased to obtain a subgraph and supergraph of the HRG, respectively.

In the following we first discuss our main contribution in the context of existing HRG generators. In Section 2, we formally define the GIRG and HRG models. Afterwards we describe the sampling algorithm in Section 3. In Section 4 we discuss implementation details, including the parameter estimation for the average degree (Section 4.1) as well as multiple performance improvements. Section 5 contains our experiments: we investigate the scaling behavior of our generator in Section 5.1, compare our HRG generator to existing ones in Section 5.2, and compare GIRGs to HRGs in Section 5.3.

1.2 Comparison with Existing Generators

We are not aware of previous GIRG implementations. Concerning HRGs, most algorithms only support the threshold case; see Table 1. The only published exceptions are the trivial quadratic algorithm [2], and an $O((n^{3/2} + m) \log n)$ algorithm [23] based on a quad-tree data structure [24]. The latter is part of NetworkKit; we call it NKQUAD. Moreover, the code for a hyperbolic embedding algorithm [5] includes an HRG generator implemented by Bringmann based on the GIRG algorithm [7]; we call it EMBEDDER in the following. EMBEDDER has been widely ignored as a high performance generator. This is because it was somewhat hidden, and it is heavily outperformed by other threshold generators. Experiments show that our generator HYPERGIRGs is much faster than NKQUAD, which is to be expected considering the asymptotic running time. Moreover, on a single processor, we outperform EMBEDDER by an order of magnitude for $T = 0$ and by a factor of 4 for higher temperatures. As EMBEDDER does not support parallelization, this speed-up increases for multiple processors.

¹ <https://github.com/chistopher/girgs>

■ **Table 1** Existing hyperbolic random graph generators. The columns show the names used throughout the paper; the conference appearance; a reference (journal if available); whether the generator supports the binomial model; and the asymptotic running time. The time bounds hold in the worst case (wc), with high probability (whp), in expectation (exp), or empirically (emp).

Name	First Published	Ref.	Binom.	Running Time
PAIRWISE	CPC'15	[2]	✓	$\Theta(n^2)$ (wc)
QUADTREE	ISAAC'15	[24]		$O((n^{3/2} + m) \log n)$ (wc)
NKQUAD	IWOCA'16	[23]	✓	$O((n^{3/2} + m) \log n)$ (wc)
NKGEN, NKOPT	HPEC'16	[25]		$O(n \log n + m)$ (emp)
EMBEDDER	ESA'16	[5]	✓	$\Theta(n + m)$ (exp)
HYPERGEN	SEA'17	[21]		$O(n \log \log n + m)$ (whp)
RHG	IPDPS'18	[13]		$\Theta(n + m)$ (exp)
SRHG	JPDC'19	[12]		$\Theta(n + m)$ (exp)
HYPERGIRGS	this paper		✓	$\Theta(n + m)$ (exp)

For the threshold variant of HRGs, there are the following generators. The quad-tree data structure mentioned above was initially used for a threshold generator (QUADTREE) [24]. It was later improved leading to the algorithm currently implemented in NetworKit (NKGEN) [25]. A later re-implementation by Penschuck [21] improves it by about a factor of 2 (NKOPT). However, the main contribution of [21] was a new generator that features sublinear memory and near optimal parallelization (HYPERGEN). Up to date, HYPERGEN was the fastest threshold-HRG generator on a single processor. Our generator, HYPERGIRGS, improves by a factor of 1.3 – 2 (depending on the parameters) but scales worse for more processors. Finally, Funke et al. [13] provide a generator designed for a distributed setting to generate enormous instances (RHG). Its run time was later further reduced (SRHG) [12].

2 Models

2.1 Geometric Inhomogeneous Random Graphs

GIRGs [7] combine elements from random geometric graphs [14] and Chung-Lu graphs [9, 10]. Let $V = \{1, \dots, n\}$ be a set of vertices with positive weights w_1, \dots, w_n following a power law with exponent $\beta > 2$. Let W be their sum. Let \mathbb{T}^d be the d -dimensional torus for a fixed dimension $d \geq 1$ represented by the d -dimensional cube $[0, 1]^d$ where opposite boundaries are identified. For each vertex $v \in V$, let $x_v \in \mathbb{T}^d$ be a point drawn uniformly and independently at random. For $x, y \in \mathbb{T}^d$ let $\|x - y\|$ denote the L_∞ -norm on the torus, i.e. $\|x - y\| = \max_{1 \leq i \leq d} \min\{|x_i - y_i|, 1 - |x_i - y_i|\}$. Two vertices $u \neq v$ are independently connected with probability p_{uv} . For a positive temperature $0 < T < 1$,

$$p_{uv} = \min \left\{ 1, c \left(\frac{w_u w_v / W}{\|x_u - x_v\|^d} \right)^{1/T} \right\} \quad (1)$$

while for $T = 0$ a threshold variant of the model is obtained with

$$p_{uv} = \begin{cases} 1 & \text{if } \|x_u - x_v\| \leq c(w_u w_v / W)^{1/d}, \\ 0 & \text{else.} \end{cases}$$

The constant $c > 0$ controls the expected average degree. We note that the above formulation slightly deviates from the original definition; see Section 2.3 for more details.

2.2 Hyperbolic Random Graphs

HRGs [17] are generated by sampling random positions in the hyperbolic plane and connecting vertices that are close. More formally, let $V = \{1, \dots, n\}$ be a set of vertices. Let $\alpha > 1/2$ and $C \in \mathbb{R}$ be two constants, where α controls the power-law degree distribution with exponent $\beta = 2\alpha + 1 > 2$, and C determines the average degree \bar{d} . For each vertex $v \in V$, we sample a random point $p_v = (r_v, \theta_v)$ in the hyperbolic plane, using polar coordinates. Its angular coordinate θ_v is chosen uniformly from $[0, 2\pi)$ while its radius $0 \leq r_v < R$ with $R = 2 \log(n) + C$ is drawn according to the density function $f(r) = \frac{\alpha \sinh(\alpha r)}{\cosh(\alpha R) - 1}$. In the threshold case of HRGs two vertices $u \neq v$ are connected if and only if their distance is below R . The hyperbolic distance $d(p_u, p_v)$ is defined via $\cosh(d(p_u, p_v)) = \cosh(r_u) \cosh(r_v) - \sinh(r_u) \sinh(r_v) \cos(\theta_u - \theta_v)$.

The binomial variant adds a temperature $T \in [0, 1]$ to control the clustering, with lower temperatures leading to higher clustering. Two nodes $u, v \in V$ are then connected with probability $p_T(d(p_u, p_v))$ where $p_T(d) = (\exp[(d - R)/(2T)] + 1)^{-1}$. For $T \rightarrow 0$, the two definitions (threshold and binomial) coincide.

2.3 Comparison of GIRGs and HRGs

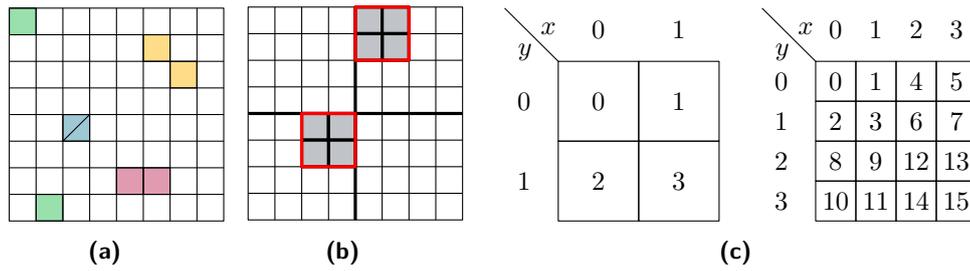
Bringmann et al. [7] show that the HRG model can be seen as a special case of the GIRG model in the following sense. Let d_{HRG} be the average degree of a HRG. Then there exist GIRGs with average degree d_{GIRG} and D_{GIRG} with $d_{\text{GIRG}} \leq d_{\text{HRG}} \leq D_{\text{GIRG}}$ such that they are sub- and supergraphs of the HRG, respectively. Moreover, d_{GIRG} and D_{GIRG} differ only by a constant factor. Formally, this is achieved by using the big-O notation instead of a single constant c for the connection probability. We call this the *generic GIRG framework*. It basically captures any specific model whose connection probabilities differ from Equation (1) by only a constant factor. From a theoretical point of view this is useful as proving something for the generic GIRG framework also proves it for any manifestation, including HRGs.

To see how HRGs fit into the generic GIRG framework, consider the following mapping [7]. Radii are mapped to weights $w_v = e^{(R-r_v)/2}$, and angles are scaled to fit on a 1-dimensional torus $x_v = \theta_v/(2\pi)$. One can then see that the hyperbolic connection probability $p_T(d)$ under the provided mapping deviates from Equation (1) by only a constant. Thus, c in Equation (1) can be chosen such that all GIRG probabilities are larger or smaller than the corresponding HRG probabilities, leading to the two average degrees d_{GIRG} and D_{GIRG} mentioned above. Bringmann et al. [7] note that the two constants, which they hide in the big-O notation, do not have to match. They leave it open if they match, converge asymptotically, or how large the interval between them is in practice. We investigate this empirically in Section 5.3.

3 Sampling Algorithm

As mentioned in the introduction, the core of our sampling algorithm is based on the algorithm by Bringmann et al. [7]. In the following, we first give a description of the core ideas and then work out the details that lead to an efficient implementation.

To explain the idea, we make two temporary assumptions and relax them in Section 3.1 and Section 3.2, respectively. For now, assume that all weights are equal and consider only the threshold variant $T = 0$. The task is to find all vertex pairs that form an edge, i.e., their distance is below the threshold $c(w_u w_v / W)^{1/d}$. Since all weights are equal, the threshold in this restricted scenario is the same for all vertex pairs. One approach to quickly identify adjacent vertices is to partition the ground space into a grid of cells. The size of the cells should be chosen, such that (I) the cells are as small as possible and (II) the diameter



■ **Figure 1** (a),(b) The grid used by weight bucket pairs with a connection probability threshold between 2^{-3} and 2^{-4} in two dimensions. (a) Each pair of colored cells represent neighbors. Note that the ground space is a torus and a cell is also a neighbor to itself. (b) The eight gray cells represent multiple distant cell pairs, which are replaced by one pair consisting of the red outlined parent cell pair. (c) Linearization of the cells on level 1 (left) and 2 (right) for $d = 2$.

of cells is larger than the threshold $c(w_u w_v / W)^{1/d}$. The latter implies that only vertices in neighboring cells can be connected thus narrowing down the search space. The former ensures that neighboring cells contain as few vertex pairs as possible reducing the number of comparisons. Figure 1a shows an example of such a grid for a 2-dimensional ground space.

3.1 Inhomogeneous Weights

Assume that we have vertices with two different weights w_1, w_2 , rather than one. As before, the cells should still be as small as possible while having a diameter larger than the connection threshold. However, there are three different thresholds now, one for each combination of weights. To resolve this, we can group the vertices by weight and use three differently sized grids to find the edges between them.

As GIRGs require not only two but many weights, considering one grid for every weight pair is infeasible. The solution is to discretize the weights by grouping ranges of weights into *weight buckets*. When searching for edges between vertices in two weight buckets, the pair of largest weights in these buckets provides the threshold for the cell diameter. This choice of the cell diameter satisfies property (II). Property (I) is violated only slightly, if the weight range within the bucket is not too large. Thus, each combination of two weight buckets uses a grid of cells, whose granularity is based on the maximum weight in the respective buckets.

As a tradeoff, we choose $\lceil \log_2 n \rceil$ many buckets which yields a sublinear number of grids. Moreover, the largest and smallest weight in a bucket are at most a factor two apart. Thus, the diameter of a cell is too large by at most a factor of four.

With this approach, a single vertex has to appear in grids of different granularity. To do this in an efficient manner, we recursively divide the space into ever smaller grid cells, leading to a hierarchical subdivision of the space. This hierarchy is naturally described by a tree. For a 2-dimensional ground space, each node has four children, which is why we call it *quadtree*. Note that each level of the quadtree represents a grid of different granularity. Moreover, the side length of a grid cell on level ℓ is $2^{-\ell}$. For a pair (i, j) of weight buckets, we then choose the level that fits best for the corresponding weights, i.e., the deepest level such that the diameter of each grid cell is above the connection threshold for the largest weights in bucket i and j , respectively. We call this level the *comparison level*, denoted by $CL(i, j)$. It suffices to insert vertices of a bucket into the deepest level among all its comparison levels. This level is called the *insertion level* and we denote it by $I(i)$. In Section 3.4, we discuss in detail how to efficiently access all vertices in a given grid cell belonging to a given weight bucket.

3.2 Binomial Variant of the Model

For $T > 0$, neighboring cell pairs are still easy to handle: a constant fraction of vertex pairs will have an edge and one can sample them by explicitly checking every pair. For distant cell pairs and a fixed pair of weight buckets, the distance between the cells yields an upper bound on the connection probability of included vertices; see Equation (1). The probability bound depends on both, the weight buckets and the cell pair distance, using the maximum weight within the buckets and the minimum distance between points in the cells. We note that the individual connection probabilities are only a constant factor smaller than the upper bound.

Knowing this, we can use geometric jumps to skip most vertex pairs [1]. The approach works as follows. Assume that we want to create an edge with probability \bar{p} for each vertex pair. For this process, we define the random variable X to be the number of vertex pairs we see until we add the next edge. Then X follows a geometric distribution. Thus, instead of throwing a coin for each vertex pair, we can do a single experiment that samples X from the geometric distribution and then skip X vertex pairs ahead. Since not all vertex pairs reach the upper bound \bar{p} , we accept encountered pairs with probability p_{uv}/\bar{p} to get correct results.

Although distant cell pairs are handled efficiently, their number is still quadratic, most of which yield no edges. To circumvent this problem, the sampling algorithm, yet again, uses a quadtree. In the quadratic set of cell pairs to compare for one weight bucket pair, non-neighboring cells are grouped together along the quadtree hierarchy. They are replaced by their parents as shown in Figure 1b until their parents become neighbors.

In conclusion, for each pair of weight buckets (i, j) the following two types of cell pairs have to be processed. Any two neighboring cell pairs on the comparison level $CL(i, j)$; and any distant cell pair with level larger or equal $CL(i, j)$ that has neighboring parents. The resulting set of distant and neighboring cell pairs for a fixed bucket pair partitions $\mathbb{T}^d \times \mathbb{T}^d$.

3.3 Efficiently Iterating Over Cell Pairs

The previous description sketches the algorithm as originally published. Here, we propose a refactoring that greatly simplifies the implementation and enables parallelization. We attribute a significant amount of HYPERGIRGS' speed up over EMBEDDER to this change.

Instead of first iterating over all bucket pairs and then over all corresponding cell pairs, we reverse this order. This removes the need to repeatedly determine the cell pairs to process for a given bucket pair. Instead it suffices to find the bucket pairs that process a given cell pair. This only depends on the level of the two cells and their type (neighboring or distant). Inverting the mapping from bucket pairs to cell pairs in the previous section yields the following. A neighboring cell pair on level ℓ is processed for bucket pairs with a comparison level of exactly ℓ . A distant cell pair on level ℓ (with neighboring parents) is processed for bucket pairs with a comparison level larger than or equal to ℓ . Thus, for each level of the quadtree we must enumerate all neighboring cell pairs, as well as distant cell pairs with neighboring parents. Algorithm 1 recursively enumerates exactly these cell pairs.

3.4 Efficient Access to Vertices by Bucket and Cell

A crucial part of the algorithm is to quickly access the set of vertices restricted to a weight bucket i and a cell A , which we denote by V_i^A . To this end, we linearize the cells of each level as illustrated in Figure 1c. This linearization is called Morton code [18] or z-order curve [20]. It has the nice properties that (I) for each cell in level ℓ , its descendants in level $\ell' > \ell$ in the quadtree appear consecutively; and (II) it is easy to convert between a cell's position in the linear order and its d -dimensional coordinates (see Section 4.2).

Algorithm 1: Sample GIRG by Recursive Iteration of Cell Pairs

Input: cell pair (A,B) ; initially called with A,B set to the root of the quadtree

- 1 **forall** bucket pairs (i, j) that process the cell pair (A, B) **do**
- 2 **if** A and B are neighbors **then**
- 3 emit each edge $(u, v) \in V_i^A \times V_j^B$ with probability p_{uv}
- 4 **else**
- 5 choose candidates $S \subseteq V_i^A \times V_j^B$ using geometric jumps and \bar{p}
- 6 emit each edge $(u, v) \in S$ with probability p_{uv}/\bar{p}
- 7 **if** A and B are neighbors **and not** maximum depth reached **then**
- 8 **forall** children X of A **do**
- 9 **forall** children Y of B **do**
- 10 recur(X, Y)

We sort the vertices of a fixed weight bucket i by the Morton code of their containing cell on the insertion level $I(i)$, using arbitrary tie-breaking for vertices in the same cell. This has the effect that for any cell A with $\text{level}(A) \leq I(i)$, the vertices of V_i^A appear consecutive. Thus, to efficiently enumerate them, it suffices to know for each cell A the index of the first vertex in V_i^A . This can be precomputed using prefix sums leading to the following lemma.

► **Lemma 1.** *After linear preprocessing, for all cells A and weight buckets i with $\text{level}(A) \leq I(i)$, vertices in the set V_i^A can be enumerated in $\mathcal{O}(|V_i^A|)$.*

4

 Implementation Details

The description in the previous section is an idealized version of the algorithm. For an actual implementation, there are some gaps to fill in. Omitting many minor tweaks, we want to sketch optimizations that are crucial to achieve a good practical run time in the following. More details on the sketched approaches can be found in the long version of this paper [3].

4.1 Estimating the Average Degree Parameter

Here, we sketch how to estimate the parameter c in Eq. (1) to achieve a given expected average degree. We estimate the constant based on the actual weights, not on their probability distribution. This leads to lower variance and allows user-defined weights.

We start with an arbitrary constant c , calculate the resulting expected average degree $\mathbb{E}[\bar{d}]$ and adjust c accordingly, using a modified binary search. This is possible, as $\mathbb{E}[\bar{d}]$ is monotone in c . We derive an exact formula for $\mathbb{E}[\bar{d}]$, depending on c and the weights. It cannot simply be solved for c , which is why we use binary search instead of a closed expression.

For the binary search, we need to evaluate $\mathbb{E}[\bar{d}]$ for different values of c . This is potentially problematic, as the formula for $\mathbb{E}[\bar{d}]$ sums over all vertex pairs. The issue preventing us from simplifying this formula is the minimum in the connection probability. Therefore, we first ignore the minimum and subtract an error term for those vertex pairs, where the minimum takes effect. The remaining hard part is to calculate this error term. Let E_S be the set of vertex pairs appearing in the error term and let S be the set of vertices with at least one partner in E_S . Although $|E_S|$ itself is sufficiently small, S is too large to determine E_S by iterating over all pairs in $S \times S$. We solve this by iterating over the vertices in S , sorted by

weight. Then, for each vertex we encounter, the set of partners in E_S is a superset of the partners of the previous vertex with smaller weight, allowing us to compute $\mathbb{E}[\bar{d}]$ in $O(S)$.

4.2 Efficiently Encoding and Decoding Morton Codes

Recall from Section 3.4 that we linearize the d -dimensional grid of cells using Morton code. As vertex positions are given as d -dimensional coordinates, we have to convert the coordinates to Morton codes (i.e., the index in the linearization) and vice versa. This is done by bitwise interleaving the coordinates. For example, the 2-dimensional Morton code of the four-bit coordinates $a = a_3a_2a_1a_0$ and $b = b_3b_2b_1b_0$ is $a_3b_3a_2b_2a_1b_1a_0b_0$. We evaluated different encoding and decoding approaches via micro benchmarks. The fastest approach, at least on Intel processors, was an assembler instruction from BMI2 proposed by Intel in 2013 [16].

4.3 Generating HRGs Avoiding Expensive Mathematical Operations

The algorithm from Section 3 can be used to generate HRGs. The algorithm works conceptually the same, except that most formulas change. This has for example the effect that we no longer get a closed formula to determine the insertion level of a weight bucket or the comparison level of a bucket pair. Instead, one has to search them, by iterating over the levels of the quadtree. Furthermore, HRGs introduce many computationally expensive mathematical operations like the hyperbolic cosine. This can be mitigated as follows.

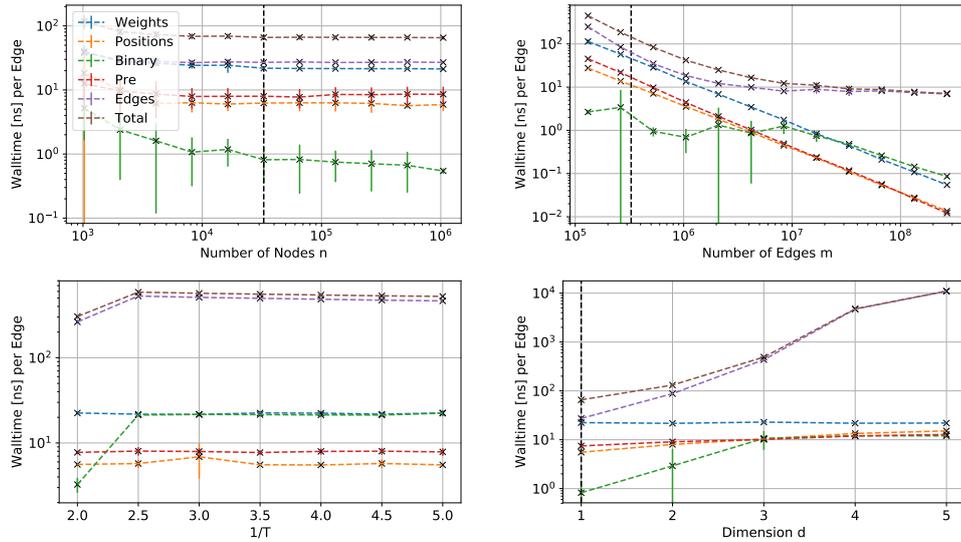
For the threshold model, an edge exists if the distance d is smaller than R . Considering how the hyperbolic distance is defined (Section 2.2), reformulating it to $\cosh(d) < \cosh(R)$ avoids the expensive $\operatorname{arccosh}$, while $\cosh(R)$ remains constant during execution and can thus be precomputed. Similar to recent threshold HRG generators, we compute intermediate values per vertex such that $\cosh(d)$ can be computed using only multiplication and addition [12, 21].

For the binomial model, evaluating the connection probability is a performance bottleneck. The straightforward way to sample edges is: compute the connection probability $p_T(d)$ depending on the distance, sample a uniform random value $u \in [0, 1]$, and create the edge if and only if $u < p_T(d)$. We can improve this by precomputing the inverse of $p_T(d)$ for equidistant values in $[0, 1]$. This lets us, for small ranges in $[0, 1]$, quickly access the corresponding range of distances. Changing the order, we first sample $u \in [0, 1]$, which falls in a range between two precomputed values, which in turn yields a range of distances. If the actual distance lies below that range, there has to be an edge and if it lies above, there is no edge. Only if it lies in the range, we actually have to compute the probability $p_T(d)$.

4.4 Parallelization

The algorithm has five steps: generate weights, generate positions, estimate the average degree constant, precompute the geometric data structure, and sample edges. The first two are trivial to parallelize. For estimating the constants, we parallelize the dominant computations with linear running time. To sample the edges, we make use of the fact that we iterate over cell pairs in a recursive manner. This can be parallelized by cutting the recursion tree at a certain level and distributing the loose ends among multiple processors.

For the preprocessing we have to do three subtasks: compute for each vertex its containing cells on its insertion level, sort the vertices according to their Morton code index, and compute the prefix sum for all cells. We parallelize all three tasks and optimize them by handling all weight buckets together, sorting by weight bucket first and Morton code second. This is done by encoding this criterion into integers that are sorted with parallel radix sort.



■ **Figure 2** Sequential run time for the steps of the GIRG sampling algorithm averaged over 10 iterations. Each plot varies a different model parameter deviating from a base configuration $d = 1$, $n = 2^{15}$, $T = 0$, $\beta = 2.5$, and $\bar{d} = 10$. The base configuration is indicated by a dashed vertical line.

5 Experimental Evaluation

We perform three types of experiments. In Section 5.1 we investigate the scaling behavior of our GIRG generator, broken down into the different tasks performed by the algorithm. In Section 5.2 we compare our HRG generator with existing generators. In Section 5.3 we experimentally investigate the difference between HRGs and their GIRG counterpart. Whenever a data point represents the mean over multiple iterations, our plots include error bars that indicate the standard deviation. Besides the implementation itself, all benchmarks and analysis scripts are also accessible in our source repository.

5.1 Scaling of the GIRG Generator

We investigate the scaling of the generator, broken down into five steps. 1. **(Weights)** Generate power-law weights. 2. **(Positions)** Generate points on \mathbb{T}^d . 3. **(Binary)** Estimate the constant controlling the average degree. 4. **(Pre)** Preprocess the geometric data structure (Section 3.4). 5. **(Edges)** Sample edges between all vertex pairs as described in Algorithm 1.

Figure 2 shows the sequential run time over the number of nodes n (top left), number of edges m (top right), temperature T (bottom right), and dimension d (bottom right). The performance is measured in nanoseconds per edge. Each data point represents the mean over 10 iterations. To make the measurements independent of the graph representation, we do not save the edges into RAM, but accumulate a checksum instead. Note that the top right plot increases the average degree, resulting in a decreased time per edge.

The empirical run times match the theoretical bounds: it is linear in n and m , grows exponentially in the dimension d , and is unaffected by the temperature T . The overall time is dominated by the edge sampling. Generating the weights includes expensive exponential functions, making it the slowest step after edge sampling. Generating the positions is significantly faster even for higher dimensions. For the parameter estimation using binary search, one can see that the run time never exceeds the time to generate the weights. For

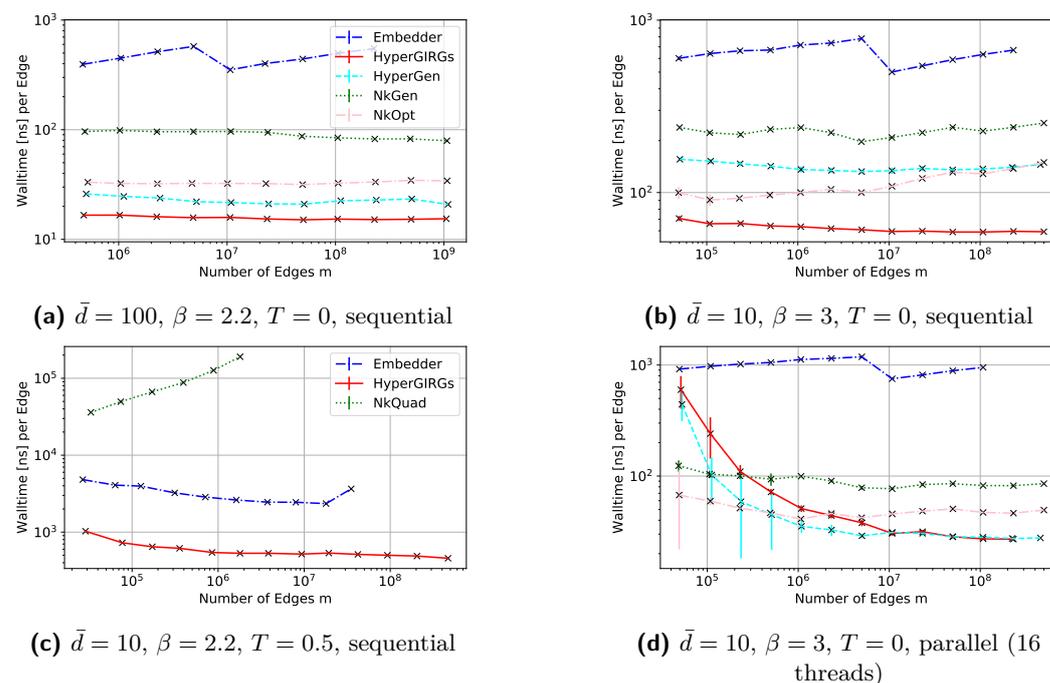


Figure 3 Comparison of HRG generators averaged over 5 iterations. (a), (b) Threshold variant for different average degrees \bar{d} and power-law exponents β . (c) Binomial variant with temperature $T = 0.5$. (d) The same configuration as (b) but utilizing multiple cores.

non-zero temperature T the performance of the binary search is similar to the generation of the weights, as it also requires exponential functions. The lower run times per edge for the increasing number of edges (top right) show that the run time is dominated by the number of nodes n . Only for very high average degrees, the cost per edge outgrows the cost per vertex.

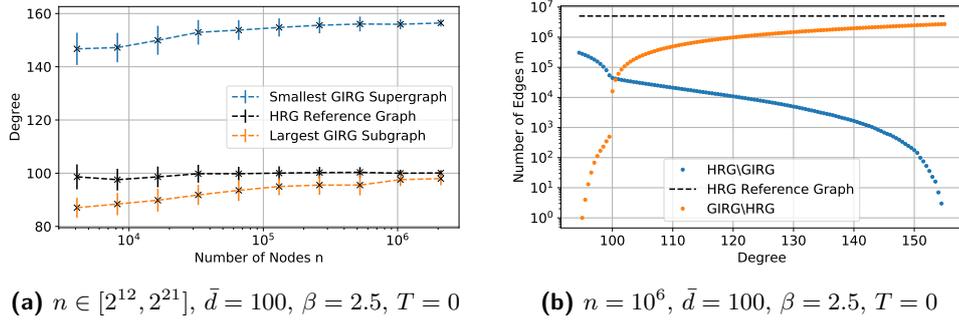
5.2 HRG Run Time Comparison

We evaluate the run time performance of HYPERGIRGs compared to the generators in Table 1, excluding the generators with high asymptotic run time as well as RHG and SRHG. RHG and SRHG are designed for distributed machines. Executed on a single compute node, the performance of the faster SRHG is comparable to HYPERGEN [12]. To avoid systematic biases between different graph representations, the implementations are modified² not to store the resulting graph. Instead, only the number of edges produced is counted and we ensure that the computation of incident nodes is not optimized away by the compiler.

We used different machines for our sequential and parallel experiments. The former are done on an *Intel Core i7-8700K* with 16 GB RAM, the latter on an *Intel Xeon CPU E5-2630 v3* with 8 cores (16 threads) and 64 GB RAM.

Our generator HYPERGIRGs is consistently faster than the competitors, independent of the parameter choices; see Figure 3a and 3b. Only for unrealistic average degrees (1k), HYPERGEN slightly outperforms HYPERGIRGs. Moreover, HYPERGIRGs beats EMBEDDER, the only other efficient generator supporting non-zero temperature, by an order of magnitude.

² The modifications are publicly available and referenced in our GitHub repository.



■ **Figure 4** Relation between the HRG and the GIRG model. (a) The values for d_{HRG} , d_{GIRG} , D_{GIRG} averaged over 50 iterations. (b) The number of missing ($\text{HRG} \setminus \text{GIRG}$) and additional ($\text{GIRG} \setminus \text{HRG}$) edges depending on the expected degree of the corresponding GIRG. It can be interpreted as a cross-section of one iteration in (a).

For higher temperatures, we compare our algorithm with the two other non-quadratic generators NKQUAD (included in NetworkKit) and EMBEDDER; see Figure 3c. We note that EMBEDDER uses a different estimation for R , which leads to an insignificant left-shift of the corresponding curve. In Figure 3c, one can clearly see the worse asymptotic running time of NKQUAD. Compared to EMBEDDER, HYPERGIRGS is consistently 4 times faster.

Figure 3d shows measurements for parallel experiments using 16 threads. The parameters coincide with Figure 3b. EMBEDDER does not support parallelization and is outperformed even more. For sufficiently large graphs, the fastest generator in this multi-core setting is HYPERGEN, which is specifically tailored towards parallel execution. Nonetheless, HYPERGIRGS shows comparable performance and overtakes the other two generators NKGGEN and NKOPT. We note that even on parallel machines, the sequential performance is of high importance: one often needs a large collection of graphs rather than a single huge instance. In this case, it is more efficient to run multiple instances of a sequential generator in parallel.

5.3 Difference Between HRGs and GIRGs

Recall from Section 2.3 that a HRG with average degree d_{HRG} has a corresponding GIRG sub- and supergraphs with average degrees d_{GIRG} and D_{GIRG} , respectively.

We experimentally determine, for given HRGs, the values for d_{GIRG} by decreasing the average degree of the corresponding GIRGs until it is a subgraph of the HRG. Analogously, we determine the value for D_{GIRG} . We focus on the threshold variant of the models, as this makes the coupling between HRGs and GIRGs much simpler (the graph is uniquely determined by the coordinates). Figure 4a shows d_{GIRG} and D_{GIRG} , compared to d_{HRG} for growing n . One can see that d_{GIRG} and D_{GIRG} are actually quite far apart. They in particular do not converge to the same value for growing n . However, at least d_{GIRG} seems to approach d_{HRG} . This indicates that every HRG corresponds to a GIRG subgraph that is missing only a sublinear fraction of edges. On the other hand, the average degree of the GIRG has to be increased by a lot to actually contain all edges also contained in the HRG.

Figure 4b gives a more detailed view for a single HRG. Depending on the average degree of the GIRG, it shows how many edges the GIRG lacks and how many edges the GIRG has in addition to the HRG. For degree 100, the GIRG contains about 38 k additional and lacks about 42 k edges. These are rather small numbers compared to the 50 M edges of the graphs.

References

- 1 Joachim H. Ahrens and Ulrich Dieter. Sequential random sampling. *ACM Transactions on Mathematical Software*, 11(2):157–169, 1985. doi:10.1145/214392.214402.
- 2 Rodrigo Aldecoa, Chiara Orsini, and Dmitri Krioukov. Hyperbolic graph generator. *Computer Physics Communications*, 196:492–496, 2015. doi:10.1016/j.cpc.2015.05.028.
- 3 Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, Ulrich Meyer, Manuel Penschuck, and Christopher Weyand. Efficiently generating geometric inhomogeneous and hyperbolic random graphs. *CoRR*, abs/1905.06706, 2019. URL: <http://arxiv.org/abs/1905.06706>, arXiv:1905.06706.
- 4 Thomas Bläsius, Cedric Freiberger, Tobias Friedrich, Maximilian Katzmann, Felix Montenegro-Retana, and Marianne Thieffry. Efficient shortest paths in scale-free networks with underlying hyperbolic geometry. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107, pages 20:1–20:14, 2018. doi:10.4230/LIPIcs.ICALP.2018.20.
- 5 Thomas Bläsius, Tobias Friedrich, Anton Krohmer, and Sören Laue. Efficient embedding of scale-free graphs in the hyperbolic plane. *IEEE/ACM Transactions on Networking*, 26(2):920–933, 2018. doi:10.1109/TNET.2018.2810186.
- 6 OpenMP Architecture Review Board. OpenMP application program interface version 5.0, 2018. URL: <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-5.0.pdf>.
- 7 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Geometric inhomogeneous random graphs. *Theoretical Computer Science*, 760:35–54, 2019. doi:10.1016/j.tcs.2018.08.014.
- 8 Deepayan Chakrabarti and Christos Faloutsos. Graph mining: Laws, generators, and algorithms. *ACM Comput. Surv.*, 38(1), 2006. doi:10.1145/1132952.1132954.
- 9 Fan Chung and Linyuan Lu. The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences*, 99(25):15879–15882, 2002. doi:10.1073/pnas.252631999.
- 10 Fan Chung and Linyuan Lu. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics*, 6(2):125–145, 2002. doi:10.1007/PL00012580.
- 11 Tobias Friedrich and Anton Krohmer. On the diameter of hyperbolic random graphs. *SIAM Journal on Discrete Mathematics*, 32(2):1314–1334, 2018. doi:10.1137/17M1123961.
- 12 Daniel Funke, Sebastian Lamm, Ulrich Meyer, Manuel Penschuck, Peter Sanders, Christian Schulz, Darren Strash, and Moritz von Looz. Communication-free massively distributed graph generation. *Journal of Parallel and Distributed Computing*, 131:200–217, 2019. doi:10.1016/j.jpdc.2019.03.011.
- 13 Daniel Funke, Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Moritz von Looz. Communication-free massively distributed graph generation. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 336–347, 2018. doi:10.1109/IPDPS.2018.00043.
- 14 Edgar N. Gilbert. Random plane networks. *Journal of the Society for Industrial and Applied Mathematics*, 9(4):533–543, 1961. doi:10.1137/0109045.
- 15 Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. Random hyperbolic graphs: Degree sequence and clustering. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 573–585, 2012. doi:10.1007/978-3-642-31585-5_51.
- 16 Intel Corporation. *Intel 64 and IA-32 Architectures Developer’s Manual*, 2019.
- 17 Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Physical Review E*, 82:036106, 2010. doi:10.1103/PhysRevE.82.036106.
- 18 Guy M Morton. A computer oriented geodetic data base and a new technique in file sequencing. Technical report, International Business Machines Company New York, 1966. URL: <https://domino.research.ibm.com/library/cyberdig.nsf/0/0dabf9473b9c86d48525779800566a39?OpenDocument>.

- 19 Tobias Müller and Merlijn Staps. The diameter of KPKVB random graphs. *CoRR*, abs/1707.09555, 2017. URL: <http://arxiv.org/abs/1707.09555>, arXiv:1707.09555.
- 20 J. A. Orenstein and T. H. Merrett. A class of data structures for associative searching. In *ACM SIGACT-SIGMOD Symposium on Principles of Database Systems (PODS)*, pages 181–190, 1984. doi:10.1145/588011.588037.
- 21 Manuel Penschuck. Generating practical random hyperbolic graphs in near-linear time and with sub-linear memory. In *International Symposium on Experimental Algorithms (SEA)*, volume 75, pages 26:1–26:21, 2017. doi:10.4230/LIPIcs.SEA.2017.26.
- 22 Christian L. Staudt, Aleksejs Sazonovs, and Henning Meyerhenke. NetworKit: A tool suite for large-scale complex network analysis. *Network Science*, 4(4):508–530, 2016. doi:10.1017/nws.2016.20.
- 23 Moritz von Looz and Henning Meyerhenke. Querying probabilistic neighborhoods in spatial data sets efficiently. In *International Workshop on Combinatorial Algorithms (IWOCA)*, pages 449–460, 2016. doi:10.1007/978-3-319-44543-4_35.
- 24 Moritz von Looz, Henning Meyerhenke, and Roman Prutkin. Generating random hyperbolic graphs in subquadratic time. In *International Symposium on Algorithms and Computation (ISAAC)*, pages 467–478, 2015. doi:10.1007/978-3-662-48971-0_40.
- 25 Moritz von Looz, Mustafa Safa Özdayi, Sören Laue, and Henning Meyerhenke. Generating massive complex networks with hyperbolic geometry faster in practice. In *IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6, 2016. doi:10.1109/HPEC.2016.7761644.
- 26 Duncan J. Watts and Steven H. Strogatz. Collective dynamics of “small-world” networks. *Nature*, 393:440–442, 1998. doi:10.1038/30918.