

Navigability is a Robust Property

Dimitris Achlioptas *
Department of Computer Science
University of California, Santa Cruz
optas@cs.ucsc.edu

Paris Siminelakis †
Department of Electrical Engineering
Stanford University
psimin@stanford.edu

Abstract

The Small World phenomenon has inspired researchers across a number of fields. A breakthrough in its understanding was made by Kleinberg who introduced Rank Based Augmentation (RBA): add to each vertex independently an arc to a random destination selected from a carefully crafted probability distribution. Kleinberg proved that RBA makes many networks *navigable*, i.e., it allows greedy routing to successfully deliver messages between any two vertices in a polylogarithmic number of steps. We prove that navigability is an inherent property of many random networks, arising without coordination, or even independence assumptions.

*Research supported by a European Research Council (ERC) Starting Grant (StG-210743) and an Alfred P. Sloan Fellowship.

†Supported in part by an Onassis Foundation Scholarship.

1 Introduction

The Small World phenomenon refers to the fact that there exist short chains of acquaintances between most pairs of people in the world, popularly known as Six Degrees of Separation [17]. Milgram’s famous 1967 experiment [16] showed that not only such chains exist, but they can also be found in a decentralized manner. Specifically, each participant in the experiment was handed a letter addressed to a certain person and was told of some general characteristics of the person, including their occupation and location. They were then asked to forward the letter to the individual they knew on a first-name basis who was most likely to know the recipient. Based on the premise that similar individuals have higher chance of knowing each other (homophily), the participants typically forwarded the message to their contact most similar to the target, a strategy that yielded remarkably short paths for most letters that reached their target (many did not).

Kleinberg’s groundbreaking work, formulated mathematically the property of finding short-paths in a decentralized manner as *navigability* [8, 10]. Since then, much progress has been made [9] and the concept of navigability has found applications in the design of P2P networks [5, 19], data-structures [4, 15] and search algorithms [14, 18]. Key to decentralization is shared knowledge in the form of geometry, i.e., shared knowledge of a (distance) function on pairs of vertices (not necessarily satisfying the triangle inequality).

Geometry. A geometry (V, d) consists of a set of vertices V and a distance function $d : V \times V \rightarrow \mathbb{R}_+$, where $d(x, y) \geq 0$, $d(x, y) = 0$ iff $x = y$, and $d(x, y) = d(y, x)$, i.e., d must be a semi-metric.

Given a graph $G(V, E)$ on a geometry (V, d) , a *decentralized search algorithm* is any algorithm that given a target vertex t and current vertex v selects the next edge $\{v, u\} \in E$ to cross by only considering the distance of each neighbor u of v to the target t , i.e., $d(u, t)$. The allowance of paths of polylogarithmic length in the definition of navigability, below, is motivated by the fact that in any graph with constant degree the diameter is $\Omega(\log(n))$, reflecting an allowance for polynomial loss due to the lack of global information.

Navigability. A graph $G(V, E)$ on geometry (V, d) is d -navigable if there exists a decentralized search algorithm which given any two $s, t \in V$ will find a path from s to t of length $O(\text{poly}(\log |V|))$.

In his original work on navigability [8, 10], Kleinberg showed that if G is the 2-dimensional grid then adding a single random edge independently to each $v \in V$ results in a navigable graph (with d being the L1 distance on the grid). The distribution for selecting the other endpoint u of each added edge is crucial. Indeed, if it can only depend on $d(v, u)$ and distinct vertices are augmented independently, Kleinberg showed that there is a *unique* suitable distribution, the one in which the probability is proportional to $d(v, u)^{-2}$ (and, more generally, $d(v, u)^{-r}$ for r -dimensional lattices). The underlying principle behind Kleinberg’s augmentation scheme has by now become known as *Rank Based Augmentation* (RBA) [11, 13].

Rank Based Augmentation. Given a geometry (V, d) , a vertex $v \in V$, and $\ell \geq 0$, let $N_v(\ell)$ be the number of vertices within distance ℓ from v . In RBA, the probability of augmenting v with an edge to any $u \in V$ is

$$P(v, u) \propto \frac{1}{N_u(d(v, u))} . \quad (1)$$

The intuition behind RBA is that navigability is attained because the added edges provide connectivity *across all distance scales*. Concretely, observe that for any partition of the range of the distance function d into intervals, the probability that the (distance of the) other endpoint of an added edge lies in a given interval is independent of the interval. Therefore, by partitioning the range of d into $O(\log n)$ intervals we see that whatever the current vertex v is, there is always $\Omega((\log n)^{-1})$ probability that its long-range edge

is to a vertex at a distance at the same “scale” as the target. Of course, that alone is not enough. In order to shrink the distance to the target by a constant factor, we also need the long-range edge to have reasonable probability to go “in the right direction”, something which is effortlessly true in regular lattices for any finite dimension. In subsequent work [11], aiming to provide rigorous results for graphs beyond lattices, Kleinberg showed that the geometric conditions needed for RBA to achieve navigability are satisfied by the geometries induced by *set-systems* satisfying certain conditions when the distance between two vertices is the size of the smallest set (homophily) containing both (see Definition 1 in Section 5).

Another canonical setting for achieving navigability by RBA is when the distance function d is the shortest-path metric of a connected graph $G_0(V, E_0)$ with large diameter $\Theta(\text{poly}(n))$. In that setting, if E_d is the random set of edges added through RBA, the question is whether the (random) graph $G(V, E_0 \cup E_d)$ is d -navigable. Works of Slivkins [15] and Fraigniaud et al. [7] have shown the existence of a threshold, below which navigability is attainable and above which (in the worst case) it is not attainable, in terms of the *doubling dimension* of the shortest path metric of G_0 . Roughly speaking, the doubling dimension corresponds to the logarithm of the possible directions that one might need to search, and the threshold occurs when it crosses $\Theta(\log \log n)$. Thus, we see that even when d is a (shortest path) metric, very significant additional constraints on d need to be imposed.

The remarkable success of RBA in conferring navigability rests crucially on its *perfect adaptation* to the underlying geometry. This adaptation, though, not only requires perfect independence and identical behavior of all vertices, but also a very specific, indeed unique, functional form for the probability distribution of edge formation. This exact fine tuning renders RBA unnatural greatly weakening its plausibility. Our goal in this paper is to demonstrate that navigability is in fact a robust property of networks that emerges from very basic considerations without adaptation, coordination, or even independence assumptions.

2 Our Contribution

As mentioned, at the foundation of navigability lies shared knowledge in the form of geometry. Our starting premise is that geometry imposes *global* constraints on the set of feasible networks. Most obviously, in a physical network where edges (wire, roads) correspond to a resource (copper, concrete) there is typically an upper bound on how much can be invested to create the network. More generally, cost may represent a number of different notions that distinguish between edges. We formalize this intuition by (i) allowing different edges to have arbitrary costs, i.e., without imposing any constraints on the cost structure, and (ii) taking as input an upper bound on the *total cost* of feasible graphs, i.e., a budget. We remain fully agnostic in all other respects, i.e., we study the *uniform* measure on all graphs satisfying the budget constraint. So, for example, if all edges have unit cost we recover the classic Erdős-Rényi $G(n, m)$ random graphs (except now m is a random variable, sharply concentrated just below the budget.)

As one can imagine, the set of all graphs feasible within a given budget may contain wildly different elements. Our capacity to study the uniform measure on such graphs comes from a very recent general theorem we developed in [1], of which this work is the first application. At a high level, the main theorem of [1] asserts that if a subset S of the set of all undirected simple graphs \mathcal{G}_n on n vertices is sufficiently symmetric, then the uniform measure on S can be well-approximated by a product measure on the edges, i.e., a measure where each edge is included independently with different edges potentially having different probabilities. Formally, a product measure on \mathcal{G}_n is specified succinctly by a symmetric matrix $\mathbf{Q} \in [0, 1]^{n \times n}$ of probabilities where $Q_{ii} = 0$ for $i \in [n]$. We denote by $G(n, \mathbf{Q})$ the measure in which possible edge $\{i, j\}$ is included independently with probability $Q_{ij} = Q_{ji}$. The main result of [1] then allows one to approximate the uniform measure by a product measure in the following very strong sense.

Sandwichability. The uniform measure $U(S)$ over an arbitrary set of graphs $S \subseteq \mathcal{G}_n$ is (ϵ, δ) -sandwichable if there exists a $n \times n$ symmetric matrix \mathbf{Q} such that the two distributions $G^\pm \sim G(n, (1 \pm \epsilon)\mathbf{Q})$, and the distribution $G \sim U(S)$ can be coupled so that $G^- \subseteq G \subseteq G^+$ with probability at least $1 - \delta$.

As discussed above, navigability requires some degree of structure in the underlying geometry. It is from this structure that we will extract the symmetry needed to apply the theorem of [1] and derive a product form approximation for graphs with a bounded total cost. Armed with such an approximation, establishing navigability becomes dramatically easier, allowing us to demonstrate its robustness and ubiquity. Roughly speaking, we isolate three ingredients that suffice for navigability on a geometry (V, d) :

- A *substrate* of connections between nearby points on V , allowing the walk to never get stuck.
- Some degree of *coherence* of the distance function d .
- Sufficient, and sufficiently uniform, edge density across all distance scales.

The first two ingredients are generalizations of existing work and, as we will see, fully compatible with RBA. The third ingredient is also motivated by the RBA viewpoint, but we will prove that it can be achieved in far more-light handed, and thus natural, manner than RBA. Moreover, in the course of doing so, we will give it a very natural *economical* interpretation as the *cost of indexing*.

2.1 The two Basic Requirements and a Unifying Framework for RBA

Substrate. A set of edges E_0 forms a substrate for a geometry (V, d) , if for every $(s, t) \in V \times V$ with $s \neq t$, there is at least one vertex v such that $\{s, v\} \in E_0$ and $d(v, t) \leq d(s, t) - 1$.

The existence of the substrate implies that (very slow) travel between any two vertices is possible, so that a decentralized algorithm never gets trivially stuck.

Coherence is a notion that comes with an associated scale factor $\gamma > 1$. Specifically, given a geometry (V, d) we will refer to the vertices whose distance from a given vertex $v \in V$ lie in the interval $(\gamma^{k-1}, \gamma^k]$ as the vertices in the k -th (distance) γ -scale from v . Also, for a fixed $\lambda < 1$ and any target vertex $t \neq v$, we will say that a vertex u is t -helpful to v if $d(v, u) \leq \gamma^k$ (u is within the same γ -scale as t), and $d(u, t) < \lambda d(v, t)$ (reduces the distance by a constant).

Coherence. Let $K = \lceil \log_\gamma |V| \rceil$. A geometry (V, d) is γ -coherent if there is $\lambda < 1$ such that for all $v \in V$:

- For all $k \in [K]$, the number of vertices in the k -th distance scale from v is $P_k(v) = \Theta(\gamma^k)$.
- For all $t \neq v$, a constant fraction of the vertices whose distance scale from v is no greater than the distance scale of t are t -helpful to v .

The two conditions above endow the, otherwise arbitrary, semi-metric d with sufficient regularity and consistency to guide the search. Although our definition of coherence is far more general, in order to convey intuition about the two conditions, think for a moment of V as a set of points in Euclidean space. The first condition guarantees that there are no “holes”, as the variance in the density of points is bounded in every distance scale. The second condition guarantees that around any vertex v the density of points does not change much depending on the direction (target vertex t) and distance scale. Besides those two conditions, we make *no further* assumptions on d and, in particular, we do *not* assume the triangle inequality.

Coherent geometries allow us to provide a unified treatment of navigability since they encompass finite-dimensional lattices, hierarchical models, any vertex transitive graph with bounded doubling dimension and, as we prove in Section 5, Kleinberg’s set systems [11] (see Definition 1).

Theorem 1. *Every set system satisfying the conditions of [11] is a γ -coherent geometry for some $\gamma > 1$.*

Theorem 2. *Let (V, d) be any γ -coherent geometry and let E_0 be any substrate for it. If E_d is the (random) set of edges obtained by applying RBA to (V, d) , then the graph $G(V, E_0 \cup E_d)$ is d -navigable w.h.p.¹*

Theorem 2 subsumes and unifies a number of previous positive results on RBA-induced navigability. Our main contribution, though, lies in showing that given a substrate and coherence, navigability can emerge without coordination from the interplay of cost and geometry.

2.2 Navigability from Organic Growth

As mentioned earlier, the success of RBA stems from the fact that the edge-creation mechanism is *perfectly* adapted to the underlying geometry so as to induce navigability. In contrast, we will not even specify an edge-creation mechanism, but rather focus only on the set of graphs feasible with a given budget. Our requirement is merely that the cost function is *informed* by the geometry.

γ -consistency. *Given a γ -coherent geometry (V, d) , a cost function $c : V \times V \rightarrow \mathbb{R}$ is γ -consistent if c takes the same value c_k for every edge $\{u, v\}$ such that $d(u, v) \in (\gamma^{k-1}, \gamma^k]$.*

In particular, note that we make no requirements of the values $\{c_k\}$, not even a rudimentary one, such as being increasing in k . All that γ -consistency entails is that the partition of edges according to cost is a coarsening of the partition of the edges by γ -scale. In fact, even this requirement can be weakened significantly, as long as some correlation between the two partitions remains, but it is technically much simpler to assume γ -consistency as it simplifies the exposition greatly. One can think of consistency as limited sensitivity with respect to distance. As an example, it means that making friends with the people next door might be more likely than making friends with other people on the same floor, and that making friends with people on the same floor is more likely than making friends with people in a different floor, but it does not really matter which floor.

Cost-geometries. *We say that $\Gamma = \Gamma(V, d, c)$ is a coherent cost-geometry if there exists $\gamma > 1$ such that (V, d) is a γ -coherent geometry and c is γ -consistent cost function.*

Random Graphs of Bounded Cost. *Given a coherent cost-geometry $\Gamma(V, d, c)$ and a real number $B \geq 0$, let $G_\Gamma(B) = \{E \subseteq V \times V : \frac{1}{n} \sum_{e \in E} c(e) \leq B\}$, i.e., $G_\Gamma(B)$ is the set of all graphs (edge sets) on V with total cost at most Bn . A uniformly random element of $G_\Gamma(B)$ will be denoted as $E_\Gamma = E_\Gamma(B)$.*

Applying the main theorem of [1], in Section 3 we will prove that random graphs of bounded cost (on a consistent cost-geometry) have a product measure approximation, in the following sense.

Theorem 3. *Given a coherent cost-geometry Γ , there exist a unique function $\lambda(B) \geq 0$ and constant $B_0(\Gamma) > 0$ such that for every $B \geq B_0(\Gamma)$ the uniform measure on $G_\Gamma(B)$ is (δ, ϵ) -sandwichable by the product measure in which the probability of every edge with cost c_k is*

$$\frac{1}{1 + \exp(\lambda(B)c_k)} \quad , \quad (2)$$

and $(\delta, \epsilon) = \left(\sqrt{\frac{24}{\log |V|}}, 2|V|^{-5K} \right)$. The number $\lambda(B) > 0$ can be explicitly defined in terms of $\{c_k\}$.

¹Throughout the paper, all asymptotics will be with respect to the number of vertices $|V| = n$. Thus, with high probability will always mean “with probability that tends to 1 as $n \rightarrow \infty$.”

The regularizer $\lambda = \lambda(B)$ in Theorem 3 corresponds to the derivative of entropy with respect to the budget B (energy), i.e., is an inverse temperature, and depends on B in a smooth one-to-one manner.

Theorem 3 will give us a great amount of access to the uniform measure on $G_\Gamma(B)$. In particular, the upper approximation $G(|V|, (1 + \epsilon)\mathbf{Q})$ will allow us to bound the total number of edges present in a typical graph, establishing sparsity for all sufficiently small budgets. On the other hand, the lower approximation $G(|V|, (1 - \epsilon)\mathbf{Q})$ will allow us to establish a lower bound on the number of edges incident to each vertex of each distance scale. Combined with the spatial uniformity afforded by independence, this will allow us to prove that navigability emerges as soon as the total number of edges within each scale is large enough, establishing navigability for all sufficiently large budgets.

Theorem 4. *For every coherent cost-geometry $\Gamma(V, d, c)$, where $|V| = n$, there exist numbers B^\pm such that if E_Γ is a uniformly random element of $G_\Gamma(B)$ then:*

- For all $B \leq B^+$, w.h.p. $|E_\Gamma| = O(n \cdot \text{poly}(\log n))$. (Sparsity)
- For all $B \geq B^-$, for any substrate E_0 , w.h.p. the graph $G(V, E_0 \cup E_\Gamma)$ is d -navigable. (Navigability)

Note that Theorem 4 shows that navigability arises eventually, i.e., for all $B \geq B^-$, without *any* further assumptions on the cost function or geometry. The caveat, if we think of B as increasing from 0, is that by the time there are enough edges across all distance scales, i.e., $B \geq B^-$, the total number of edges may be much greater than linear. This is because for an arbitrary cost structure $\{c_k\}$, by the time the “slowest growing” distance scale has the required number of edges, the other scales may be replete with edges, possibly many more than the required $\Omega(n/\text{poly} \log(n))$. This is reflected in the ordering between B^- and B^+ that determines whether the sparsity and navigability regimes are overlapping. In particular, we would like $B^- \leq B^+$ and, ideally, the ratio $R = B^+/B^- > 0$ to be large. Whether this is the case or not depends precisely on the degree of adaptation of the cost-structure $\{c_k\}$ to the geometry as we discuss next.

2.3 Navigability as a Reflection of the Cost of Indexing

Recall that for every vertex v in a γ -coherent geometry and for every distance scale $k \in [K]$, the number of vertices whose distance from v is in the k -th γ -distance scale is $P_k(v) = \Theta(\gamma^k)$. At the same time, (2) asserts that the probability of each edge is exponentially small in its cost. Thus, reconciling sparsity with navigability boils down to balancing these two factors. We will exhibit a class of cost functions that (i) have an intuitive interpretation as the *cost of indexing*, (ii) achieve a ratio $R = B^+/B^- > 0$ that *grows* with n , i.e., a very wide range of budgets for which we have both navigability and sparsity, and (iii) recover RBA as a special case corresponding to a particular budget choice.

Consider a vertex v that needs to forward a message to a neighbor u at the k -th distance scale. To do so, v needs to distinguish u among all other $P_k(v)$ vertices at the k -th distance scale, i.e., v needs to be able to *index* into that scale. To do so, it is natural to assert that v must incur a cost of $\Theta(\log_2 P_k(v)) = \Theta(k)$ (due to coherence) bits to store the unique ID of u among the other members of its equivalence class (in the eyes of v). Motivated by this we consider cost functions where for some $\alpha > 0$,

$$c_k^* = \frac{1}{\alpha} k .$$

Theorem 5. *For any coherent cost-geometry $\Gamma(V, d, c^*)$, where $|V| = n$, there exist $B^- \leq B^+$ such that:*

1. $B^+/B^- = \omega(\text{poly} \log n)$.

2. For all $B \in [B^-, B^+]$, w.h.p.:

- $|E_\Gamma(B)| = O(n \text{ poly } \log n)$.
- The graph $G(V, E_0 \cup E_\Gamma(B))$ is d -navigable.

3. There exists $B_a \in [B^-, B^+]$ such that in the approximation of E_Γ by $G(|V|, \mathbf{Q})$, for every $\{u, v\} \in E$, $\mathbf{Q}_{ij}^* = \Theta(N_u(d(u, v))^{-1})$, i.e, Rank Based Augmentation is approximately recovered.

Part 1 of Theorem 5 is equivalent to a scaling window of $\Theta(\frac{\log \log n}{\log n})$ for the exponent λ , within which navigability holds with poly-logarithmic average degree. This corroborates Kleinberg's work that gave a unique exponent of $\beta = -1$ in the context of RBA for the scaling (1) of probability. Nevertheless, under our framework this vanishing window for the highly sensitive parameter λ produces a *diverging* range for values of B , explaining the purported fragility of RBA to looking at perturbations in the *wrong scale*. In fact, we can use this feature of our model to provide the first theoretical explanation for the discrepancy between theoretical results and empirical evidence [2, 13, 6] showing that real networks exhibit an exponent $\hat{\beta} \approx 0.8 < 1$. In our setting, exponents smaller than 1 correspond to *higher* average degree and thus we can attribute this discrepancy to finite size effects (finite n) and the densification [12] of networks.

3 Deriving a Product Measure Approximation: Proof of Theorem 3

We start with some definitions that will allow us to state the main theorem of [1]. A set of graphs $S \subseteq \mathcal{G}_n$ is symmetric with respect to a partition \mathcal{P} of the set of all possible $\binom{n}{2}$ edges, if the characteristic function of S depends only on the number of edges from each part of \mathcal{P} but not on which edges.

Edge Profile. Given a partition $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_K)$ of the set of all possible $\binom{n}{2}$ edges, for a set of edges $E \in \mathcal{G}_n$ and for each $k \in [K]$, let $m_k(E)$ denote the number of edges in E from \mathcal{P}_k . The edge profile of E is $\mathbf{m}(E) := (m_1(E), \dots, m_K(E))$.

We denote the image of a symmetric set S under the edge-profile as $\mathbf{m}(S)$. As before let $P_k = |\mathcal{P}_k| = \frac{1}{2} \sum_{u \in V} P_k(u)$ be the total number of edges in part k of partition \mathcal{P} .

Edge Profile Entropy. Given an edge profile $\mathbf{v} = (v_1, \dots, v_k)$ the entropy of \mathbf{v} is $\text{ENT}(\mathbf{v}) = \sum_{k=1}^K \log \binom{P_k}{v_k}$.

The edge-profile entropy is used to express the number of graphs with a particular edge profile \mathbf{v} as $\exp(\text{ENT}(\mathbf{v}))$. Given any symmetric set $S \subseteq \mathcal{G}_n$, the probability of observing an edge profile \mathbf{v} when sampling an element uniformly at random from S is then given by $\mathbf{P}_S(\mathbf{v}) = \frac{1}{|S|} e^{\text{ENT}(\mathbf{v})}$. Thus, in order to analyze the distribution of a random edge-profile, and consequently of a random element of $G_c(n, B)$, we are going to exploit analytic properties of the entropy on the set of feasible edge profiles $\mathbf{m}(S)$.

Convexity. Let $\text{Conv}(A)$ denote the convex hull of a set A . Say that a \mathcal{P} -symmetric set $S \subseteq \mathcal{G}_n$ is convex iff the convex hull of $\mathbf{m}(S)$ contains no new integer points, i.e., if $\text{Conv}(\mathbf{m}(S)) \cap \mathbb{N}^k = \mathbf{m}(S)$.

Entropic Optimizer. Given a symmetric set S , let $\mathbf{m}^* = \mathbf{m}^*(S) \in \mathbb{R}^k$ be the unique solution to

$$\max_{\mathbf{v} \in \text{Conv}(\mathbf{m}(S))} - \sum_{k=1}^K \left[v_k \log \left(\frac{v_k}{P_k} \right) + (P_k - v_k) \log \left(\frac{P_k - v_k}{P_k} \right) \right]. \quad (3)$$

Given the maximizer $\mathbf{m}^*(S)$, the matrix $\mathbf{Q}^* = \mathbf{Q}^*(S)$ is given by letting for all $k \in [K]$ the probability of an edge $e \in \mathcal{P}_k$ be $Q_e^* := m_k^*/P_k$. To state the theorem, we need the following parameters that quantify the concentration of the uniform measure around its mode.

Thickness and Condition Number. *Given a partition \mathcal{P} and a \mathcal{P} -symmetric set S , we define*

$$\text{Thickness:} \quad \mu = \mu(S) = \min_{k \in [K]} \min\{m_k^*, P_k - m_k^*\} \quad (4)$$

$$\text{Condition number:} \quad \tau = \tau(S) = \frac{5K \log n}{\mu(S)} \quad (5)$$

We now state the main theorem employed in the proof.

Theorem 6 ([1]). *Let \mathcal{P} be any edge-partition and let S be any \mathcal{P} -symmetric convex set. For every $\epsilon > \sqrt{12\tau(S)}$, the uniform measure over S is (ϵ, δ) -sandwichable for $\delta = 2 \exp\left[-\mu(S) \left(\frac{\epsilon^2}{12} - \tau(S)\right)\right]$.*

In our setting, S is the set $G_\Gamma(B) := \{E \subset V \times V : \frac{1}{n} \sum_{e \in E} c_e \leq B\}$ of graphs with bounded average cost and \mathcal{P} is the partition induced by the coherent cost function c . The set $\mathbf{m}(S)$ is then given by $\mathbf{m}(S) = \{\mathbf{v} \in \mathbb{N}^k : \frac{1}{n} \sum_{k=1}^K c_k v_k \leq B\}$. Hence, it is easy to see that $G_\Gamma(B)$ is convex and symmetric, according to the previous definition, for all values of B . To prove Theorem 3, we need to find:

- (i) an analytic expression for the vector \mathbf{m}^* as a function of B
- (ii) the range of values of B for which applying Theorem 6 gives high probability bounds.

3.1 Finding the Entropic Optimal Edge Profile

We start by introducing a slight reparametrization in terms of the average-degree profile. For an edge set E , define the vector $\mathbf{a}(E) := \mathbf{m}(E)/n$, where as before \mathbf{m} is the edge-profile. In the same spirit, let $p_k = P_k/n$ denote the average number of edges in part (scale) k . Using this parametrization and by explicitly writing $\text{Conv}(\mathbf{a}(S))$, we can equivalently express the optimization problem (3) as:

$$\begin{aligned} \max_{\mathbf{a}} H(\mathbf{a}) &= - \sum_{k=1}^K [(p_k - a_k) \log(p_k - a_k) + a_k \log(a_k)] \\ \text{subject to} & \sum_{k=1}^K a_k c_k \leq B \\ & 0 \leq a_k \leq p_k, \quad \forall k \in [K] . \end{aligned}$$

We will refer to the above optimization problem as (Λ) and to its solution as $\mathbf{a}^* = \mathbf{a}^*(B)$. Towards obtaining an analytic expression for \mathbf{a}^* , we first show that no coordinate $k \in [K]$ lies on the natural boundary $\{0, p_k\}$.

Lemma 1. *The optimal profile $\mathbf{a}^* \in \mathcal{D}(B) := \{\mathbf{a} \in (0, p_1) \times \dots \times (0, p_K) : \sum_k^K c_k a_k \leq B\}$.*

Proof. We prove the lemma by contradiction. We show that if \mathbf{a}^* is a solution of (Λ) such that $\mathbf{a}^* \notin \mathcal{D}$, then there is an $\hat{\mathbf{a}}^* \in \mathcal{D}$ for which objective function f takes a higher value. Specifically, for $\epsilon > 0$ assume that there are indices $1 \leq i, j \leq K$ such that $a_i^* = 0$ and $a_j^* > \delta(\epsilon)^2$, where $\delta(\epsilon) = \epsilon c_i / c_j$. Define

²For any nontrivial values of B such an index can always be found.

$\hat{\mathbf{a}}^*(\epsilon) = (a_1^*, \dots, a_i^* + \epsilon, \dots, a_j^* - \delta(\epsilon), \dots, a_K^*)$. If $h(\epsilon) = H(\hat{\mathbf{a}}^*) - H(\mathbf{a}^*)$ is the difference in the objective function between the assumed optimal \mathbf{a}^* and the perturbation $\hat{\mathbf{a}}^*$, then

$$h'(\epsilon) = -\log(\epsilon) + \log(p_i - a_i - \epsilon) + \frac{c_i}{c_j} (\log(a_j - \delta(\epsilon)) - \log(p_j - a_j + \delta(\epsilon))) .$$

Observe, that $\lim_{\epsilon \rightarrow 0} h'(\epsilon) = +\infty$, since we have assumed that $a_j^* > 0$. This shows that every maximizer satisfies $\mathbf{a}^* > 0$. The same argument establishes that $a_k^* < p_k$ for all $k \in [K]$. Combining the two statements we get that any maximizer belongs in \mathcal{D} . \square

As a consequence, since they are inactive at the optimum, we can omit separable inequalities from the formulation. Further, define $\bar{B} := \frac{1}{2} \sum_{k=1}^K p_k c_k$ the average cost of the solution to the unconstrained version of (Λ) , i.e., where $\bar{a}_k := p_k/2$. If $B > \bar{B}$ then the absolute maximum entropic point $\bar{\mathbf{a}}$ is still in $\mathcal{D}(B)$ and thus the solution will be always $a_k^* = \bar{a}_k$ for every such B .

Lemma 2. *There is a unique function $\lambda(B)$ that is one-to-one for all $0 \leq B \leq \bar{B}$ and $\lambda(B) = 0$ for all $B \geq \bar{B}$, such that the unique solution of (Λ) is given by:*

$$a_k^*(B) = \frac{p_k}{1 + \exp[\lambda(B) \cdot c_k]}, \quad \forall k \in [K] . \quad (6)$$

Proof. Uniqueness of the solution follows easily from convexity of the domain and concavity of the objective function. Further, by Lemma 1, we can reduce the optimization problem (Λ) to the following:

$$\begin{aligned} \max_{\mathbf{a}} \quad & - \sum_{k=1}^K [(p_k - a_k) \log(p_k - a_k) + a_k \log(a_k)] \\ \text{subject to} \quad & \sum_{k=1}^K a_k c_k \leq B . \end{aligned}$$

To obtain an analytical solution, we form the Lagrangian of the reduced problem

$$L(\mathbf{a}, \lambda) = - \sum_{k=1}^K [(p_k - a_k) \log(p_k - a_k) + a_k \log(a_k)] + \lambda \left(B - \sum_{k=1}^K a_k c_k \right) .$$

with the additional constraint that $\lambda \geq 0$. The Karush-Kuhn-Tacker conditions read

$$\frac{\partial L}{\partial a_k} = 0 \iff \log\left(\frac{a_k}{p_k - a_k}\right) = -\lambda c_k \quad (7)$$

$$\frac{\partial L}{\partial \lambda} = 0 \iff \sum_{k=1}^K a_k c_k = B . \quad (8)$$

Solving the first equation for $a_k(\lambda)$ we get

$$a_k^* = \frac{p_k}{1 + \exp(\lambda c_k)} ,$$

Substituting this expression in (8), we get the following function of λ :

$$g(\lambda) = \sum_{k=1}^K c_k \cdot \frac{p_k}{1 + \exp(\lambda c_k)} \quad (9)$$

and the second constraint can now be written as $g(\lambda) = B$. The domain of g is the set of non-negative numbers on which g is continuous and infinitely differentiable. Under positive costs $\{c_k\}$, it is easy to see that $g'(\lambda) < 0$ for all $B < \bar{B}$, hence, g is strictly decreasing in the interval $[0, \infty)$ and $g(0) = \bar{B}$. Thus, $g : [0, \infty) \rightarrow [0, \bar{B}]$ is 1-to-1 and thus invertible. This means that every budget in $[0, \bar{B}]$ is feasible and that for each such budget there is a unique $\lambda(B) := g^{-1}(B)$. For $B \geq \bar{B}$, $\lambda(B) = 0$. Therefore, we conclude that the maximizer is always unique for any feasible B and implicitly given by $g(\lambda) = B$. \square

3.2 Thickness $\mu(B)$ of $G_\Gamma(B)$ and Sandwiching

Our next step is to use the analytical solution to the optimization problem to instantiate the thickness parameter μ defined in (4). Using (6), we can write:

$$\mu(B) = \min_{k \in [K]} m_k^* = n \cdot \min_{k \in [K]} \frac{p_k}{1 + \exp[\lambda(B)c_k]} \quad (10)$$

where we have used the facts that $a_k^* = m_k^*/n$ and $a_k^*(B) \leq 1/2 \Rightarrow m_k^* \leq P_k - m_k^*$. To get a more convenient expression, since $0 < c_k < \infty$ we can write the cost as $c_k = \frac{1}{\beta_k} \log(p_k)$ where $0 < \beta_k < \infty$ when $p_k \geq 1$. Thus, approximately³ for large p_k (eq. k) we have $\mu(B) \approx n \cdot \min_{k \in [K]} \left[p_k^{1-\lambda(B)/\beta_k} \right]$. Theorem 6, gives strong (non-constant) probability bounds as long as $\tau(B) \ll 1$. For our purposes we are going to consider that the maximum $\tau(B)$ (respectively minimum B) that we allow is $\tau_0 = \log^{-1}(n)$ (respectively B_0). Substituting the above expression for $\mu(B)$ in (5), we get that the condition $\tau \leq \tau_0$ can be rewritten as $\lambda(B) \leq \lambda_0$, where

$$\lambda_0 = \lambda_0(\{p_k\}, \{\beta_k\}) := \min_{k \in [K]} \left[\log \left(\frac{n \log p_k}{5K \log^2(n)} \right) \frac{\beta_k}{\log p_k} \right]. \quad (11)$$

Using the function $g(\lambda)$ defined in (9), we can express this constraint as $B \geq B_0 := g(\lambda_0)$.

To conclude the proof of Theorem 3 we see that $\mu(B) \geq 5K \log^2(n)$ and $\tau(B) \leq \frac{1}{\log(n)}$, for all $B \geq B_0$. Applying Theorem 6, for $\epsilon_0 = \sqrt{\frac{24}{\log n}}$ that is greater than $\sqrt{12\tau_0}$, we get that $\delta \leq 2 \exp \left[\mu(B) \left(\frac{\epsilon_0^2}{12} - \tau(B) \right) \right]$. The proof is concluded by substituting the bounds in the last expression.

4 Navigability via Reducibility

In this section we prove our results about navigability on coherent geometries. We start by giving a slightly more formal definition of coherence. Recall that given a geometry (V, d) and a fixed (scale factor) $\gamma > 1$, $P_k(v)$ denotes the number of vertices in V at “distance” $(\gamma^{k-1}, \gamma^k]$ from v . Further, for fixed $\lambda < 1$ and all $t \neq v \in V$, let k_{vt} be the non-negative integer such that $d(v, t) \in (\gamma^{k_{vt}-1}, \gamma^{k_{vt}}]$ and $D_\lambda(v, t)$ be the vertices in V whose distance from v is at most $\gamma^{k_{vt}}$ and whose distance from t is at most $\lambda \cdot d(v, t)$. Thus, $|D_\lambda(v, t)|$ is the number of nodes that could facilitate greedy routing (t -helpful), i.e., reduce the distance to t by a constant factor $\lambda < 1$.

Coherence. Fix $\gamma > 1$ and let $K = \lceil \log_\gamma(|V|) \rceil$. A geometry (V, d) is γ -coherent if:

(H1) Bounded Growth: $\exists A > 1, \alpha > 0$ such that $P_k(v) \in \gamma^k[\alpha, A]$, for all $v \in V$ and $k \in [K]$.

(H2) Isotropy: $\exists \phi > 0, 1 > \lambda > 0$ such that $|D_\lambda(v, t)| \geq \phi \gamma^{k_{vt}}$, for all $s \neq t \in V$.

³When the approximation does not hold it means that $\mu(B) = \Omega(n)$ which trivially satisfies all the requirements we need for “sandwiching” and navigability.

For graphs on coherent geometries there are two requirements for navigability. The first basic requirement is deterministic and amounts to the ability to move slowly (linear rate) towards the target. In the graph augmentation setting this was given by the fact that the initial set of edges formed a connected graph. On the other hand in Kleinberg’s work on set systems, the degree of vertices is set to $\Theta(\log^2(n))$, so that the probability of ever being stuck at a vertex is polynomially small. As mentioned in the introduction, we opt to adopt the more natural approach of assuming a *substrate*.

Substrate. *A set of edges E_0 forms a substrate for a geometry (V, d) , if for every $(s, t) \in V \times V$ with $s \neq t$, there is at least one vertex v such that $\{s, v\} \in E_0$ and $d(v, t) \leq d(s, t) - 1$. If there are multiple such vertices, we distinguish one arbitrarily and call it the local t -connection of s . A path starting from s and ending to t using only local t -connections is called a local (s, t) -path.*

The second requirement is probabilistic and expresses the fact that for all distance scales and “directions” there should be significant probability of observing an edge. This property is satisfied by Rank Based Augmentation and is essentially what was actually used to prove navigability originally.

Uniform Richness. *Given a γ -coherent geometry (V, d) with parameters $\alpha, \phi > 0$ define $k_\theta := \frac{\theta \log \log n - \log a}{\log \gamma}$ to be the distance scale of edges having distance $\Theta(\log^\theta(n))$. A product measure $G(n, \mathbf{Q})$ is then called θ -uniformly rich for (V, d) if there is a constant $M > 0$ such that for every $k \geq k_\theta$ every edge (i, j) with $d(i, j) \in (\gamma^{k-1}, \gamma^k]$ satisfies $Q_{ij} \geq \frac{1}{M \log^\theta(n)} \frac{1}{\gamma^k}$.*

In other words, since we are interested in routing in poly-logarithmic time and slow traveling can be done through the substrate (connected base graph), the probabilistic requirements concern only edges of longer distance. As we show next these two requirements are sufficient for navigability to arise in the general setting of random graphs of bounded cost.

4.1 Reducibility via Uniform Richness

We start by introducing a deterministic property of graphs that implies navigability, that of *reducibility*. The main advantage of reducibility is that it allows us to separate the construction of the random graph from the analysis of the algorithm.

Reducibility. *Given a graph $G(V, E)$, we will say that a pair $(s, t) \in V \times V$ is p -reducible if $\exists C > 0$ such that among the first $C(\log |V|)^p$ vertices of the local (s, t) -path there is at least one vertex u such that $(u, v) \in E$ and $d(v, t) \leq \lambda d(s, t)$. If every pair $(s, t) \in V \times V$ is p -reducible we will say that G is p -reducible.*

Proposition 1. *If G is p -reducible, greedy routing on G takes at most $1 + C(\log n)^{1+p}$ steps.*

Proof of Proposition 1. Given any arbitrary pair of vertices (s, t) with distance at most n , the reducibility property of G guarantees us that after at most $C \log^p n$ steps we will obtain a new pair (s', t) with distance reduced by a constant factor. Since, the new pair is also p -reducible, we can repeat the process until we reduce the distance again by a constant. After at most $\log_{1/\lambda} n$ iterations we will reach the target. Since, the pairs were arbitrary, this holds for all pairs and thus the graph is navigable in $1 + C(\log n)^{1+p}$ steps. \square

Lemma 3. *Given a γ -coherent geometry (V, d) with a substrate E_0 and a random edge set E_q sampled from a θ -uniformly rich product measure $G(n, \mathbf{Q})$, the graph $G(V, E_0 \cup E_q)$ is $(\theta + 1)$ -reducible with high probability.*

Proof. To prove that the graph is $(\theta + 1)$ -reducible we will (i) prove that the event B_{st} that any fixed source-destination pair (s, t) is not $(\theta + 1)$ -reducible has very small probability under $G(n, \mathbf{Q})$, and (ii) use union-bound to argue that the probability that any pair is not $(\theta + 1)$ -reducible is small as well. To simplify the proof, we first distinguish between pairs (s, t) where within the first $C \log^{\theta+1}(n)$ steps of the t -local path there is a vertex with distance smaller than $d(s, t)$ by a constant factor $\lambda < 1$ and where there is no such vertex. Pairs (s, t) that belong in the first case, are $(\theta + 1)$ -reducible with probability 1. Hence, we only need to focus on the latter case, where all vertices on the first $C \log^{\theta+1}(n)$ steps are within the same distance scale $k_{st} := \lceil \log_\gamma d(s, t) \rceil$ as s from t . We will refer to k_{st} as k to ease the notation. For each such vertex v on the t -local path, property (H2) of coherent geometries tells us that there are at least $\phi \gamma^k$ candidate edges that would reduce the distance from t by a constant factor $\lambda < 1$. The probability Q_{vz} of each such good edge (v, z) is lower bounded by $\frac{1}{M \log^{\theta+1}(n) \gamma^k}$, since the measure $G(n, \mathbf{Q})$ is θ -uniformly rich. Let $T(s, t)$ be the set of all such good edges. We can write the probability of the event B_{st} as:

$$\mathbb{P}_{\mathbf{Q}}(B_{st}) = \prod_{e \in T(s,t)} (1 - Q_e) \leq \left(1 - \frac{1}{M \log^{\theta+1}(n) \gamma^k}\right)^{|T(s,t)|} \leq e^{-\frac{C \log^{\theta+1}(n) \phi \gamma^k}{M \log^{\theta+1}(n) \gamma^k}} \leq n^{-\frac{C \phi}{M}}$$

where we used that $|T(s, t)| \geq C \log^{\theta+1}(n) \cdot \phi \gamma^k$ due to (H2) and the definition of reducibility. For any $\ell > 0$ and $C \geq (2 + \ell) \frac{M}{\phi}$ we get that $\mathbb{P}(B_{st}) \leq n^{-(2+\ell)}$. To finish the proof, we perform a Union Bound over all possible sets (s, t) . Let B be the even that the graph $G(V, E_0 \cup E_d)$ is not $(\theta + 1)$ -reducible, then:

$$\mathbb{P}_{\mathbf{Q}}(B) = \mathbb{P}_{\mathbf{Q}}\left(\bigcup_{st} B_{st}\right) \leq \sum_{st} \mathbb{P}_{\mathbf{Q}}(B_{st}) \leq n^2 n^{-(2+\ell)} = n^{-\ell}$$

for any $\ell > 0$. Thus, the graph $G(V, E_0 \cup E_d)$ is d -navigable with high probability. \square

4.2 Analyzing the Product Measure $G(n, \mathbf{Q}^*(B))$

Our next step will be to show that for a range of values of B , the product measure defined through (6) is θ -uniformly rich for some $\theta > 0$. In doing so, our previous result shows that such a product measure leads to navigable graphs. Recall that $\mathbf{Q}^*(B)$ is the matrix where for all $k \in [K]$ and $ij \in \mathcal{P}_k$ it holds that $Q_{ij}^* = (1 + \exp(\lambda(B)c_k))^{-1}$ and $g(\lambda(B))$ is the expected budget corresponding to an element generated according to the product measure $\mathbf{Q}^*(B)$.

Proposition 2. *For $B \geq B_\theta^+ := \max\{B_0, g(\lambda_\theta)\}$, the product measure $G(n, \mathbf{Q}^*(B))$ is θ -uniformly rich. The number λ_θ is explicitly defined as $\lambda_\theta(\{p_k\}, \{c_k\}) := \min_{k_\theta \leq k \leq K} \left[\frac{\log p_k}{c_k} \left(1 + \frac{\theta \log \log n}{\log p_k}\right) \right]$.*

Proof. This follows easily by the definition of λ_θ . In particular, consider an edge (i, j) of scale $k \geq h_\theta$:

$$Q_{ij}^*(B) = [1 + \exp(c_k \lambda(B))]^{-1} \geq [p_k \log^\theta(n)]^{-1} \geq \frac{1}{A \log^\theta(n) \gamma^k}$$

where the last inequality follows from (H1). \square

Proposition 3. *For $B \leq B_\theta^- := g(\Lambda_\theta)$ the product measure $G(n, \mathbf{Q}^*(B))$ has $O(n \cdot \log^{\theta+1}(n))$ edges with high probability. The number Λ_θ is explicitly defined as $\Lambda_\theta(\{p_k\}, \{c_k\}) := \max_{k_\theta \leq k \leq K} \left[\frac{\log p_k}{c_k} \left(1 - \frac{\theta \log \log n}{\log p_k}\right) \right]$*

Proof. For all $B_0 \leq B \leq B_+$, by definition of Λ_θ we have that for all $k \geq k_\theta$:

$$Q_{ij}^*(B) = [1 + \exp(c_k \lambda(B))]^{-1} \leq [p_k \log^{-\theta}(n)]^{-1}$$

Thus, the expected number of edges is upper bounded by:

$$n \cdot \left[Ak_\theta \cdot p_{k_\theta} + (K - k_\theta) \max_{k \geq k_\theta} p_k \frac{\log^\theta(n)}{p_k} \right] = n \cdot O\left(\log \log(n) \log^\theta(n) + \log(n) \log^\theta(n)\right)$$

as $k_\theta = O(\log \log n)$, $p_{k_\theta} = O(\log^\theta(n))$ by (H1) and $K = O(\log n)$. Applying standard Chernoff bounds [3] we get the required conclusion, as by definition for $B \geq B_0$ each class has at least a poly-logarithmic number of edges at the maximizer and thus the expected value (under the product measure) of the edges is tightly concentrated around the mean. \square

4.3 Analyzing Graphs of Bounded Cost

Proof of Theorem 4. For any $B \geq B_0$, consider $\mathbf{Q}^*(B)$ the matrix corresponding to the optimal profile (Lemma 2) and two random elements $E^\pm \sim G(n, (1 \pm \epsilon)\mathbf{Q}^*(B))$. By Theorem 3, we get that for $\epsilon = \sqrt{24/\log(n)}$ the probability of the event W , i.e. that $E^- \subseteq E_\Gamma \subseteq E^+$, is at least $1 - n^{-5K}$. To prove Theorem 4 we will condition on the above event and then use our analysis of the product measure. To prove Navigability we will use the relation $E^- \subset E_\Gamma$ and the fact that Navigability is monotone property. Let $N_d(E)$ be the event that that the graph $G(V, E_0 \cup E)$ is not d -navigable, then:

$$\mathbb{P}(N_d(E_c)) = \mathbb{P}(N_d(E_c) \cap W) + \mathbb{P}(N_d(E_c) \cap \bar{W}) \quad (12)$$

$$\leq \mathbb{P}(N_d(E_c)|W) + \mathbb{P}(\bar{W}) \quad (13)$$

$$\leq \mathbb{P}_{\mathbf{Q}^*}(N_d(E^-)) + n^{-5K} \quad (14)$$

$$\leq n^{-\ell} + n^{-5K} \quad (15)$$

where we used the law of total probability in the first equality, Bayes Theorem and monotonicity of the probability measure in the second inequality, Theorem 3 and monotonicity in the third, and Lemma 3 and Proposition 2 in the last. This proves part (a) of the theorem. To prove part (b) we follow the same method but for the event $\{|E_\Gamma| = \omega(n \text{poly}(\log(n)))\}$ and exploit the inequality $E_\Gamma \subset E^+$. Using Proposition 3 and Theorem 3 we get the required conclusion. \square

4.4 Analysis of Indexing

Proof of Theorem 5. We first start with the proof of part 3 of the Theorem. Instead of considering $c_k^* \propto k$ we can equivalently consider, due to (H1), $c_k^* \propto \log p_k$. Thus, for simplicity $c_k^* = \frac{1}{\alpha} \log p_k$. Set $B_a = g(a)$, for such B and an edge (u, v) of scale k , we have

$$Q_{uv}^* = \frac{1}{1 + \exp(\lambda(B_a)c_k^*)} = \frac{1}{1 + \exp(a \frac{\log p_k}{a})} = \frac{1}{1 + p_k}$$

Now, by property (H1) we know that for any vertex u and every vertex v within distance scale k from u , $N_u(d(u, v)) \in [a, A]\gamma^k$, thus we get that:

$$\left(\frac{a}{2A}\right) \frac{1}{N_u(d(u, v))} \leq \frac{1}{2A\gamma^k} \leq Q_{uv}^*(B_a) \leq \frac{1}{a\gamma^k} \leq \left(\frac{A}{a}\right) \frac{1}{N_u(d(u, v))} \quad (16)$$

Setting $r = 2A/a$ proves part (b). To further see the correspondence between Random Graphs of Bounded Cost when the cost corresponds to indexing and Rank based augmentation, consider the $a_k^*(B_a)$ the average number of edges of scale k per vertex. We have:

$$a_k^*(B_a) = \frac{p_k}{1 + p_k} \approx 1, \forall k \in [K]$$

Thus, we see that the scale invariance property of RBA is recovered. Furthermore, we have that in this case $B_a = \sum_{k=1}^K a_k^*(B_a) c_k^* = \Theta(\log^2(n))$ and the average degree of a random graph of bounded cost for B_a is $\Theta(\log(n))$.

To show the first two parts of the theorem we essentially obtain estimates for B^\pm given in Theorem 4 for the special case where the cost is the cost of indexing as above. We have:

$$\lambda_\theta^* = \alpha \left(1 + \theta \frac{\log \log n}{\log p_K} \right) \quad (17)$$

$$\Lambda_\theta^* = \alpha \left(1 - \theta \frac{\log \log n}{\log p_K} \right) \quad (18)$$

By property (H1) we know that $\log p_K = \Theta(\log n)$. Define as before $B^+ = g(\Lambda_\theta^*)$ and $B^- = g(\lambda_\theta^*)$. Then for every $B^- \leq B \leq B^+$ or equivalently for $\Lambda_\theta \leq \lambda(B) \leq \lambda_\theta$, we have that for some $C > 0$:

$$\Omega \left(\left[\log n^{-\frac{C\theta}{\log n}} \right]^k \right) = \frac{p_k}{1 + \exp(\lambda(B)c_k^*)} = O \left(\left[\log n^{\frac{C\theta}{\log n}} \right]^k \right)$$

where $a_k^*(B) = \frac{p_k}{1 + \exp(\lambda(B)c_k^*)}$ expresses the average number of edges of scale k per vertex. Thus, by (H1) we get that:

$$B^+ = \frac{1}{\alpha} \sum_{k=1}^K a_k^*(B^+) \log p_k \geq \frac{1}{\alpha} \log p_K a_K^*(B^+) = \Omega(\log(n)^{1+C'\theta})$$

Further, $B^- \leq B_a = \Theta(\log^2(n))$. Hence, we obtain that $B^+/B^- = \Omega(\text{poly}(\log n))$. The proof is concluded by invoking Theorem 3. \square

4.5 Rank Based Augmentation for Coherent Geometries

Recall that in RBA a single link is added for each vertex u to a random vertex v with probability given by

$$P_{\text{RBA}}(u, v) = \frac{1}{Z} \frac{1}{|N_u(d(u, v))|} \quad (19)$$

where $N_u(\ell) := \{t \in V : d(u, t) \leq \ell\}$ is the set of vertices that are within distance ℓ from u . Here we show that the Kleinberg's original proof can be applied with ease when instead of the semi-metric induced by set-system, we have a semi-metric corresponding to a coherent geometry. There are basically two steps. We first upper bound the normalizing constant Z and then lower bound the probability that for a given pair (s, t) we find an edge in the first $C \log^2(n)$ steps of a path along the substrate that reduces the distance to t by a constant factor.

Proposition 4 (Bounded Growth). *For a coherent geometry (V, d) , $\exists C < \infty$ such that $Z(1) \leq C \log(n)$.*

Proof. For a given vertex u , we divide vertices depending on their distance scale $k \in \{0, \dots, \log_\gamma(n)\}$ from u . For $k \geq 0$, we know from property (H1) that there are at most $A\gamma^k$ such vertices. Further, we also know that $|B_{k-1}(u)| = \sum_{i=0}^{k-1} P_k(u) \geq a \frac{\gamma^k - 1}{\gamma - 1}$. Using these two facts we have:

$$Z(1) = \sum_{v \in V} P_\alpha(u, v) \leq \frac{A}{a} + \sum_{k=1}^{\log(n)} P_k(u) \frac{1}{|B_{k-1}(u)|} \leq \frac{A}{a} + \frac{A}{a} \sum_{k=1}^{\log(n)} \gamma^k \frac{\gamma - 1}{\gamma^k - 1} \leq \frac{A}{a} (1 + \gamma \log_\gamma(n))$$

□

Finally, to complete the proof, we are going to employ once again reducibility.

Proof of Theorem 2. Fix any two vertices s, t , the probability of finding a long-range edge at s reducing the distance by a constant factor is at least:

$$\frac{|D(s, t)|}{Z} \frac{1}{P_k(s)} \geq \frac{1}{C \log n} \frac{\phi \gamma^k}{A \gamma^k} = \frac{\phi}{AC} \frac{1}{\log n}$$

Thus, the probability of the event B_{st} that no such edge exists after $C' \log^2(n)$ trials is at most:

$$\mathbb{P}(B_{st}) \leq \left(1 - \frac{\phi}{AC} \frac{1}{\log n}\right)^{C' \log^2(n)} \leq e^{-\frac{\phi C'}{AC} \log n} \leq n^{-\frac{\phi}{AC} C'}$$

For C' large enough and a union bound over the $\Theta(n^2)$ possible pairs of vertices, we get that if E_d is the random set of edges added through RBA and E_0 is a substrate for the coherent geometry (V, d) , then the graph $G(V, E_0 \cup E_d)$ is d -navigable with high probability. □

5 Set-Systems are Coherent Geometries

We begin by recalling the definitions of set-systems from [11].

Definition 1 (Set System). *Let V be a finite set of vertices and let $\Sigma = \{S_1, \dots, S_m\}$ be a collection of subsets of V . If a set S contains a vertex t we will say that S is t -bound.*

Fix $0 < \lambda < 1$ and $\beta > 1$. We say that Σ is a (λ, β) -set system if all the following hold:

(K1) $V \in \Sigma$.

(K2) If $|S| > 1$, then for every $t \in S$, there is a t -bound $S' \subset S$ of size $|S'| \geq \min\{\lambda|S|, |S| - 1\}$.

(K3) If $S_L(v)$ is the union of sets that contain v and have size at most $L \geq 2$, then $|S_L(v)| \leq \beta L$.

Given a set system Σ on a set of vertices V , we define the distance (semi-metric) between two vertices.

Definition 2. *For any two vertices $u, v \in V$, their distance in Σ , denoted by $d_\Sigma(u, v)$, is the size of the smallest set in Σ containing both vertices minus 1, i.e. $d_\Sigma(u, v) = \min_{S \in \Sigma} \{|S| - 1 : u, v \in S\}$.*

The goal of this section is to show that the geometry (V, d_Σ) is coherent for any (λ, β) -set system, i.e., prove that the semi-metric d_Σ satisfies properties (H1) and (H2) for a suitable $\gamma > 1$. Towards that direction, the main hurdle is obtaining for all $v \in V$ upper and lower bounds on $P_k(v)$, the number of

vertices at distance in $(\gamma^{k-1}, \gamma^k]$ from v . The basic observation that guides the proof is that for all v and $k \geq 1$

$$P_k(v) = |B_k(v)| - |B_{k-1}(v)| \quad (20)$$

where $B_k(v)$ is the set of all vertices having distance from v at most γ^k . This representation is very convenient because the properties of set systems are directly related to $|B_k(v)|$. In particular, if we get good upper and lower bound for $|B_k(v)|$ then we can obtain upper and lower bounds for $P_k(v)$ and prove (H1), which comprises the main challenge.

Obtaining the upper bound is trivial, since it is directly given by (K3). However, the lower bound on $B_k(v)$ requires more thought as it needs to be tight enough so that when substituting both bounds in (4) (in order to obtain a lower bound on $P_k(v)$) the difference is strictly positive. It turns out that the last property depends on the particular values of the parameters λ, β . We show that it is always possible to select $\gamma = \gamma(\beta, \lambda) > 1$ such that the last property holds. The main observation that will provide a lower bound on $|B_k(v)|$ is that the existence of a set S with size in $(\gamma^{k-1}, \gamma^k]$ implies that $|B_k(v)| \geq |S|$ for all $v \in S$. This is because all vertices in S have distance at most $|S| - 1$ from v . Thus, what remains is to show the existence of such set S for all $v \in V$ and k . To that end, we need the following axillary lemma that was implicitly stated and used in Kleinberg's original work [11].

Proposition 5 (Shrinkage). *For every $S \in \Sigma$ with $|S| \geq 1/(\lambda - \lambda^2)$ and for every $t \in S$, there exists a t -bound set $S' \in \Sigma$ with $\lambda^2|S| \leq |S'| \leq \lambda|S|$.*

Proof of Proposition 5. Assume, for the sake of contradiction, that there exists a set S and a vertex $t \in S$ such that the proposition does not hold. If we start with S and invoke (K2) iteratively until we reach t , we get a sequence $S = S_1 \supset S_2 \cdots \supset S_k = t$ of t -bound subsets of S . Since $|S| > \lambda|S|$, there is a largest index i such that $|S_i| > \lambda|S|$, and $|S_i| \geq 2$ since $\lambda|S| > 1$. Therefore, we can apply (K2) to S_i yielding a t -bound set of size at least $z = \min\{\lambda|S_i|, |S_i| - 1\}$. For the hypothesis to hold it must be that $z < \lambda^2|S|$, for if $z > \lambda|S|$ we contradict the maximality of i . But having $z < \lambda^2|S|$ is impossible since the fact $|S_i| > \lambda|S|$ implies $\lambda|S_i| > \lambda^2|S|$, while combined with the fact $|S| \geq 1/(\lambda - \lambda^2)$ it implies $|S_i| - 1 \geq \lambda^2|S|$. \square

This lemma will be used to show that for all vertices v one can start from the set V , that belongs in Σ by (K1), and inductively apply Lemma 5 to deduce the existence of sets S containing v at all scales. More specifically, given a (λ, β) -set system Σ , let M be the smallest integer such that $\lambda^{-2M} \geq |V|$. We partition the range of possible set-sizes in Σ as $\mathcal{I} = (I_1, \dots, I_M)$ by letting $I_k = (\lambda^{-2(k-1)}, \lambda^{-2k}]$, for $k \in [M]$. The partition \mathcal{I} implicitly partitions all pairs of vertices into groups, such that all pairs in a group have roughly the same distance in Σ , i.e., up to a factor of λ^2 . We show that for every vertex and for every interval of the partition, there is a set with size in that interval that contains the vertex.

Proposition 6. *For every $t \in V$, for every $k \in [M]$, there exists a t -bound set $S \in \Sigma$ with $|S| \in I_k$.*

Proof of Proposition 6. Assume, for the sake of contradiction, that there exists a vertex t for which the proposition does not hold. Let $k_0 \in [M]$ be the largest integer such that there is no t -bound set $S' \in \Sigma$ with $|S'| \in I_{k_0}$. If we start with V and invoke (K2) iteratively until we reach t , we get a sequence $V = S_1 \supset S_2 \cdots \supset S_k = t$ of t -bound sets. Let i_{k_0} be the largest index i such that $|S_i| \in I_{k_0+1}$. The maximality of k_0 implies $|S_{i_{k_0}+1}| \in I_{k_0-1}$. But invoking Proposition 5 for $S_{i_{k_0}}$ implies $|S_{i_{k_0}+1}| \in I_{k_0}$, a contradiction. \square

Treating \mathcal{I} as a distance scale, our next goal is to obtain for each vertex t , upper and lower bounds on the number of vertices that lie at each distance-scale from t . To achieve this we need to consider a coarser partition of the set sizes than \mathcal{I} . To do that it will be beneficial to use a partition built out of blocks of \mathcal{I} , thus

allowing us to utilize Proposition 6, proven for \mathcal{I} . In particular, the existence of a t -bound set of each size will be the basis for obtaining lower bounds on the number of vertices at each new distance scale from t .

We let $r = r(\beta, \lambda) \geq 2$ denote the smallest integer such that $\lambda^{-2(r-1)} > \beta$ and consider the partition that results by grouping together every r consecutive intervals of \mathcal{I} . That is, for $\gamma(\beta, \lambda) = \lambda^{-2r(\beta, \lambda)}$, we define the partition $\mathcal{A} = \mathcal{A}(\gamma)$ consisting of the intervals $A_k = (\gamma^{k-1}, \gamma^k]$, $k \in [K]$, where K is the smallest integer such that $\gamma^K \geq |V| - 1$. Having defined \mathcal{A} , we now let $P_k(v)$ denote the number of vertices whose distance from v lies in the set A_k and we let $P_k = \frac{1}{2} \sum_{v \in V} P_k(v)$ denote the total number of pairs of vertices whose distance lies in A_k .

Lemma 4 (Bounded Growth). *Let $\alpha = (\lambda^2 - \beta/\gamma) > 0$ and $A = (\beta - \lambda^2/\gamma)$. For all $k \in [K]$ and $v \in V$,*

$$\alpha \cdot \gamma^k \leq P_k(v) \leq A \cdot \gamma^k .$$

Proof of Lemma 4. First observe that \mathcal{A} is a coarsening of \mathcal{I} since $\gamma = \lambda^{-2r}$ and $r \geq 2$ is an integer. Next, let $B_k(v) = \sum_{i \leq k} P_i(v)$ be the number vertices in V whose distance from v lies in $A_1 \cup \dots \cup A_k$, i.e., is no more than γ^k . Condition (K3) asserts that $B_k(v) \leq \beta \gamma^k$. On the other hand, by Proposition 6, we know that for any $v \in V$ there is a v -bound set $S \in I_{rk} \subset A_k$. Since, all vertices in S have distance at most $|S| \leq \gamma^k$ from v , we get that $B_k(v) \geq |S| \geq \lambda^{-2(rk-1)} = \gamma^k \lambda^2$. Therefore, for all $k \in [K]$,

$$\lambda^2 \gamma^k \leq B_k(v) \leq \beta \gamma^k . \quad (21)$$

Using the representation (20) and invoking (21), we get

$$\lambda^2 \gamma^k - \beta \gamma^{k-1} \leq P_k(v) \leq \beta \gamma^k - \lambda^2 \gamma^{k-1}$$

which is equivalent to the claimed statement. The fact $\alpha > 0$ is implied by our choice of γ . \square

Thus we have shown property (H1). Proceeding further, we need to show that the semi-metric d_Σ satisfies also the isotropy property (Section 4), i.e. that the size of the set $D_\lambda(s, t) = \{v \in V : d(s, v) \leq \gamma^{k_{st}} \text{ and } d(v, t) \leq \lambda d(s, t)\}$ is proportional to $\gamma^{k_{st}}$, where k_{st} is the scale of $d(s, t)$. To do that we are going to show something stronger. Given any two vertices $s \neq t \in V$, consider a $S_{st} \in \Sigma$ of minimal size such that both $s, t \in S$. Then for all $k \leq k_{st}$ define the following set $G_k(s, t) = \{v \in S_{st} : d(s, v) \in A_k \text{ and } d(v, t) \leq \lambda |S_{st}|\}$ of vertices in S_{st} whose distance from s lies in the interval A_k (scale k) and whose distance from t is no more than $\lambda |S_{st}|$.

Lemma 5 (Isotropy). *For every $s \neq t \in V$ with $|S_{st}| \geq 1/(\lambda - \lambda^2)$, we have that*

$$|G_{k_{st}}(s, t) \cup G_{k_{st}-1}(s, t)| \geq \left(\frac{\alpha}{\gamma}\right) \gamma^{k_{st}} .$$

Proof of Lemma 5. Proposition 5 implies that there is a t -bound set $S' \in \Sigma$ with $\lambda^2 |S_{st}| \leq |S'| \leq \lambda |S_{st}|$. Thus, a λ^2 fraction of the vertices in S_{st} have distance from t at least a factor λ less than $|S_{st}|$. Having established an abundance of “good” vertices in S_{st} , we are left to show that a constant fraction of them are in the top two distance scales $k_{st}, k_{st} - 1$ from s (recall that $|S_{st}| \in A_{k_{st}}$). We start by noting that $Z = \sum_{i \leq k} |G_i(s, t)| \geq |S'|$, as the sum must count the vertices in S' . Since $S_{st} \in |A_{k_{st}}|$ and $|S'| \geq \lambda^2 |S_{st}|$, we get $Z \geq \lambda^2 \gamma^{k_{st}-1}$. On the other hand, the good vertices in the bottom $k_{st} - 2$ distance scales from s are a subset of all vertices containing s at those distance scales, a quantity bounded by (K3) as $\sum_{i \leq k-2} |G_i(s, t)| \leq \beta \gamma^{k_{st}-2}$. Therefore, $|G_{k_{st}}(s, t) \cup G_{k_{st}-1}(s, t)| \geq \lambda^2 \gamma^{k_{st}-1} - \beta \gamma^{k_{st}-2}$. \square

Proof of Theorem 1. In order to prove that the set system defines a coherent geometry, we need to show that properties (H1) and (H2) hold for some $\gamma > 1$. Our two lemmas achieve exactly that. The first property follows from Lemma 4 and the second property follows from Lemma 5 since $G_{k_{st}}(s, t) \cup G_{k_{st}-1}(s, t) \subset D_\lambda(s, t)$. \square

References

- [1] Dimitris Achlioptas and Paris Siminelakis. Product measure approximation of symmetric graph properties. *submitted*, 2014.
- [2] Lada A. Adamic and Eytan Adar. How to search a social network. *Social Networks*, 27:2005, 2005.
- [3] Dana Angluin and Leslie G Valiant. Fast probabilistic algorithms for hamiltonian circuits and matchings. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 30–41. ACM, 1977.
- [4] James Aspnes, Zoë Diamadi, and Gauri Shah. Fault-tolerant routing in peer-to-peer systems. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, pages 223–232. ACM, 2002.
- [5] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Workshop on Design Issues in Anonymity and Unobservability*, pages 46–66, 2000.
- [6] Aaron Clauset and Cristopher Moore. How do networks become navigable?, 2003.
- [7] P. Fraigniaud, E. Lebhar, and Z. Lotker. A lower bound for network navigability. *SIAM Journal on Discrete Mathematics*, 24(1):72–81, 2010.
- [8] Jon Kleinberg. The small-world phenomenon: an algorithm perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, STOC '00, pages 163–170, New York, NY, USA, 2000. ACM.
- [9] Jon Kleinberg. Complex networks and decentralized search algorithms. In *Proceedings of the International Congress of Mathematicians: Madrid, August 22-30, 2006: invited lectures*, pages 1019–1044, 2006.
- [10] Jon M. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845, August 2000.
- [11] Jon M. Kleinberg. Small-world phenomena and the dynamics of information. In *NIPS*, pages 431–438, 2001.
- [12] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 177–187. ACM, 2005.
- [13] David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Geographic routing in social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(33):11623–11628, 2005.

- [14] Gurmeet Singh Manku, Moni Naor, and Udi Wieder. Know thy neighbor's neighbor: the power of lookahead in randomized p2p networks. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, STOC '04, pages 54–63, New York, NY, USA, 2004. ACM.
- [15] Aleksandrs Slivkins. Distance estimation and object location via rings of neighbors. In *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, PODC '05, pages 41–50, New York, NY, USA, 2005. ACM.
- [16] Jeffrey Travers and Stanley Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1969.
- [17] D.J. Watts. *Six Degrees: The Science of a Connected Age*. W. W. Norton, 2003.
- [18] Jianyang Zeng, Wen-Jing Hsu, and Jiangdian Wang. Near optimal routing in a small-world network with augmented local awareness. In *Parallel and Distributed Processing and Applications*, pages 503–513. Springer, 2005.
- [19] Hui Zhang, Ashish Goel, and Ramesh Govindan. Using the small-world model to improve freenet performance. *Computer Networks*, 46(4):555 – 574, 2004.