Object Location Using Path Separators

Ittai Abraham
School of Computer Science and Engineering
Hebrew University of Jerusalem
Jerusalem, Israel
ittaia@cs.huji.ac.il

Cyril Gavoille*
Laboratoire Bordelais de Recherche en Informatique
University of Bordeaux
Bordeaux, France
gavoille@labri.fr

ABSTRACT

We study a novel separator property called *k-path separable*. Roughly speaking, a *k-path separable* graph can be recursively separated into smaller components by sequentially removing k shortest paths. Our main result is that every minor free weighted graph is k-path separable. We then show that k-path separable graphs can be used to solve several object location problems: (1) a small-worldization with an average poly-logarithmic number of hops; (2) an $(1+\varepsilon)$-approximate distance labeling scheme with $O(\log n)$ space labels; (3) a stretch-$(1+\varepsilon)$ compact routing scheme with tables of poly-logarithmic space; (4) an $(1+\varepsilon)$-approximate distance oracle with $O(n \log n)$ space and $O(\log n)$ query time. Our results generalizes to much wider classes of weighted graphs, namely to bounded-dimension isometric sparable graphs.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design – *Distributed networks*; G.2.2 [Discrete Mathematics]: Graph Theory – *Network problems, Graph labeling*.

General Terms: Algorithms, Theory.

Keywords: Compact Routing, Excluded Minor, Separator.

1. INTRODUCTION

Divide-and-conquer is a widely used paradigm in Computer Science. Typically, finding a good separator is a challenging algorithmic problem. A celebrated example is the Lipton and Tarjan $O(\sqrt{n})$-separator for planar graphs [33]. In this paper we study a novel separator property called *k-path separable*. Unlike most separators that bound the number of vertices, our new separator property only guarantees that the separator is composed, roughly speaking, of a sequence of at most k shortest paths (see Definition 1). We show that this property is sufficient to efficiently solve many object location problems like small-woldization, distance oracles, distance labeling and compact routing.

A graph H is a *minor* of G if H is a subgraph of a graph obtained by a series of edge contractions[1] of G. The main theorem of this paper states that every H-minor free weighted graph is k-path separable for some $k = k(H)$.

1.1 Related work

The study of graphs excluding a fixed minor has lead to fundamental graph theory results. In the context of routing, decentralized search, and object location, several natural classes of networks can be defined by their forbidden minor. Among them are *trees* [20, 32] (excluding K_3), *outerplanar* [21] (excluding K_4 and $K_{2,3}$) and *series-parallel* networks [18] (excluding K_4) that capture many network backbone structures, and *planar graphs* [22, 28] (excluding K_5 and $K_{3,3}$) that capture the structure of two dimensional maps.

Most relevant to out work is that of Thorup [44]. Motivated by the problem of reachability in directed graphs, the author studies object location problems, namely reachability and $(1+\varepsilon)$-approximate distance oracles, $(1+\varepsilon)$-approximate distance labels, and stretch-$(1+\varepsilon)$ labeled routing schemes, for planar graphs (both directed and undirected). His $(1+\varepsilon)$-approximate distance oracles require $O(1/\varepsilon \cdot n \log n)$ space and can be distributed into $O(1/\varepsilon \cdot \log n)$ space $(1+\varepsilon)$-approximate distance labels[2]. The labels can be transformed into a stretch-$(1+\varepsilon)$ labeled routing scheme. For distance and routing label based solutions, it is required that the distance or the first edge of the route must be answered by inspecting only the labels of the two endpoints vertices. For a formal definition of all the above problems, we refer to [44].

Motivated by the "small world" phenomena in social networks, Kleinberg [29], suggested a new probabilistic network model. Specifically, Kleinberg studies the random graph obtained by taking a two-dimensional grid and augmenting it by adding for each vertex a random edge. This model leads to new algorithmic and graph theoretic questions. One such question studies the complexity of greedy routing on these random graphs. Kleinberg gives a certain distribution of long range contacts that augments the grid and obtains

*Supported by the projects "PairAPair" and "GeoComp" of the ACI Masses de Données.

[1]The *contraction* of the edge e with endpoints u, v is the replacement of u and v with a single vertex whose incident edges are the edges other than e that were incident to u or v.

[2]In this model, one unit of space consists of a block of $\Omega(\omega + \log n)$ bits of memory, where ω is the number of bits to represent an edge-weight. In particular, for integral edge-weights, a block has length $\Omega(\log \Delta)$ bits where Δ is the weighted diameter of the graph.

$O(\log^2 n)$ expected greedy hop count. Fraigniaud [19] and Duchon et al. [15, 16] suggest to consider the generalized question of finding distributions that augment larger families of graphs and obtain poly-logarithmic expected greedy hop count. Specifically, Fraigniaud obtains $O(k \log^2 n)$ complexity for graphs with treewidth k and Duchon et al. obtain poly-logarithmic complexity for growth bounded graphs. Slivkins [41], considers graphs with doubling dimension α and obtains $O(2^{O(\alpha)} \log \Delta \log n)$ expected greedy hop complexity. Obtaining poly-logarithmic results for larger families of graphs is an open question recently highlighted by Kleinberg [30].

Finding small separators of a graph (or computing its treewidth) is an active field of research. The best polynomial time algorithm to date yields an $O(\sqrt{\log k})$-approximation for treewidth-k graphs [17]. Bounds on the size of the smallest separators are known for planar and bounded genus graphs [33, 26, 7, 14] (see [5, Table 1 p. 813] for the best current bounds), and more generally for graphs excluding a fixed minor [6, 12, 27, 35, 36]. See also [23] for separators in geometric objects. However, the smallest separator in these graphs may be large, namely of size $\Omega(\sqrt{n})$, for a regular n-vertex mesh for instance, whereas object location problems in regular meshes are obvious.

Object location solutions were studied originally by Awerbuch and Peleg for arbitrary graphs [8, 9]. The general approach was based on sparse covers. However, for every stretch $s < 3$, there are unweighted n-vertex graphs for which every stretch-s routing scheme requires $\Omega(n)$-bit labels [24], or for which every stretch-s distance labeling scheme requires $\Omega(n)$-bit labels [25, 45]. Upper bounds can be greatly improved to $1+\varepsilon$ stretch for any $\varepsilon > 0$ for graphs whose induced metric space has constant doubling dimension [10, 41, 2]. However even binary trees have unbounded doubling dimension. Moreover the techniques used in all the above papers are not separator based.

Recently [3] proposed a poly-log memory routing scheme for graphs excluding a minor. The scheme is name-independent, has constant stretch, and is limited to unweighted graphs. This limitation is inherent to name-independent schemes since there is a polynomial space lower bound for name-independent routing for weighted trees [4]. We stress that the stretch of the name-independent scheme is at least 3 and depends on the excluded minor, so it cannot be fixed arbitrary close to 1.

For graphs excluding a fixed minor, the structure theorem of Robertson and Seymour [37, 39] already has several important algorithmic applications by Demaine et al. [11, 12]. Specifically, they develop sub-exponential fixed-parameter algorithms for dominating set, vertex cover, and set cover in any class of graphs excluding a fixed minor.

1.2 Our contributions

Our algorithmic results are a non-trivial generalization of [44] to all graphs excluding a fixed minor. This solves an open question raised by Thorup [44]. These results are based on a new separator theorem whose proof uses the structure theorem of Robertson and Seymour [37, 39]. Our small word result significantly extends the result of [19] to weighted planar graphs and, in general, to all weighted graphs excluding a fixed minor.

Given a weighted graph G and $u, v \in V(G)$, let $d_G(u,v)$ denote the distance between u and v in G, i.e., the cost of a minimum cost path in G between u and v, the cost of a path being the sum of its edge weights. We omit the subscript G when it is clear from the context. We extended sets operators like $\cup$, $\cap$, or $\setminus$ to graphs in a natural way as follows: for two graphs (or subgraphs) A and B, $A \star B$, for any operator $\star \in \{\cup, \cap, \setminus\}$, is the graph G defined by $V(G) = V(A) \star V(B)$ and $E(G) = E(A) \star E(B)$, being clear that an edge e is kept in G only if both endpoints are in $V(G)$. In particular, the union $Q_1 \cup \cdots \cup Q_k$ of k paths of G represents a subgraph of G composed of k subgraphs of G.

The following new definition is central to our paper.

Definition 1. *A weighted graph G with n vertices is k-path separable is there exists a subgraph S, called k-path separator, such that:*

- *(P1) $S = P_0 \cup P_1 \cup \cdots$, where each subgraph P_i is the union of k_i minimum cost paths in $G \setminus \bigcup_{j<i} P_j$;*
- *(P2) $\sum_i k_i \leqslant k$; and*
- *(P3) either $G \setminus S$ is empty, or each connected component of $G \setminus S$ is k-path separable and has at most $n/2$ vertices.*

Intuitively, a k-path separable graph G can be halved by repetitively removing a small set of minimum cost paths (actually a sequence of unions of minimum cost paths). The total number of paths in order to halve G is at most k, however there is no restriction on the number of vertices contained in each path.

We observe that P_i's are pairwise disjoint, however two paths in the same P_i may intersect. Each path used to define S is not necessarily a minimum cost path in G (except all of those in P_0), but in some subgraphs of G.

For instance, any unweighted rectangular mesh is 1-path separable taking S as the middle row. Trees (excluding K_3) are 1-path separable as well, taking S as the center vertex of the tree — a single vertex being a trivial minimum cost path. Series-parallel graphs (excluding K_4) are 3-path separable because these graphs are of treewidth two and thus have a set of three vertices whose removal halves the graph. More generally, Thorup [44] showed that the union of three minimum cost paths suffices to halve any weighted planar graph (excluding K_5), and thus all these graphs are 3-path separable.

One may conjecture that graphs excluding K_r can be halved by a *single* union of $f(r)$ minimum cost paths for some suitable function f. Unfortunately, this is wrong for K_6-free-minor graphs (cf. Section 5.2) for which a *sequence* of unions of minimum costs paths is required to form a k-path separator. The following is our main result.

Theorem 1 (Main). *Every H-minor-free weighted connected graph is k-path separable, for $k = k(H)$, and a k-path separator can be computed in polynomial time.*

One may wonder how this result can be used to solve "object location problems". While the number of paths is bounded, the number of vertices of each path is unbounded. The key fact is that the paths are shortest paths in some subgraph. Specifically, given any shortest path R from s to t in G, the above result implies that there exists some path Q of some P_i in the k-path separator S and some subgraph $G' = G \setminus \bigcup_{j<i} P_j$ such that R and Q are shortest paths in G' and the two paths intersect. Informally, we show that s can choose a small set of landmarks from Q so that one of

these landmarks will be close to $R \cap Q$. Indeed, using the k-path separability property one can show:

THEOREM 2. *For every n-vertex k-path separable weighted graph G, and for every $\varepsilon > 0$, there exists an $(1+\varepsilon)$-approximate distance labeling scheme with $O(k/\varepsilon \cdot \log n)$ space labels. Moreover, the labels are constructible in polynomial time, and form a distance oracle of $O(k/\varepsilon \cdot n \log n)$ space supporting $(1+\varepsilon)$-approximate distance queries in $O(k/\varepsilon \cdot \log n)$ time.*

The *aspect ratio* of G is the value $\Delta = \max_{u \neq v} d(u,v) / \min d(u,v)$. To avoid dragging a normalization constant we assume $\min_{u \neq v} d(u,v) = 1$.

THEOREM 3. *For every connected n-vertex k-path separable weighted graph G with aspect ratio Δ there exists a distribution of long-range edges, computable in polynomial time, such that the greedy routing performs in $O(k^2 \log^2 n \log^2 \Delta)$ expected number of hops.*

1.3 Outline of techniques

Theorem 1 is the main result of this paper. Indeed, once the decomposition of the k-path separable graph is given, then many object location problems can be solved based on a generalization of the arguments of [44]. For the small-world problem we use a novel potential function argument.

Let us now sketch the heart of this paper, the decomposition theorem for H-minor-free graphs stated by Theorem 1. The main ingredient is the use of the tree structure of H-minor-free graphs by Robertson and Seymour [37, 39].

Roughly speaking a H-minor-free graph has a tree-decomposition in "almost" embeddable subgraphs on some surfaces where H cannot be embedded. "Almost" means that each subgraph of the tree is embedded up to a constant number of vertices (called apices), and up to a constant number of disjoint non-embeddable parts (called vortices) that have bounded pathwidth and are associated with a single face. The constants depend only on H.

A general paradigm for solving problems on these graphs, for example see [11], is to first solve the problem on planar graphs, then to extend it on bounded genus, and then to extend to bounded genus graphs having vortices and apices.

There are however many technical problems due to our specific k-path property. For instance, we cannot concentrate our attention only to one subgraph of the tree-decomposition without paying attention to the whole graph, because we are looking for a set of shortest paths that potentially expand to everywhere in the (weighted) graph, even if we can effectively control the extremities of the paths. Another difficulty is to adapt known planar techniques to bounded genus graphs having vortices, remarking that the underlying surface can be orientable (a sphere, a torus, a double-torus, etc) or nonorientable (for instance a Klein bottle), and that vortices may create nontrivial path crossings on the surface (see Fig. 1(a) for instance). We overcome the problems with a novel generalization of surface curves for embedded graphs with vortices: *vortex-paths* (cf. Definition 2).

To identify our k-path separator in G, we proceed in three steps. In Step one, a large separating subgraph (center) of the graph is identified and some vertices (called apices) are removed. The remaining subgraph will be the place where all the shortest paths start from. In Step two, the subgraph obtained from step one, is processed in an iterative manner, each time extracting vortex-paths, reducing the Euler genus of the embedded part, until the resulting subgraph has an embedded part that is planar. To bound the number of paths, technical care is taken to use vortex-paths that intersect each vortex in a bounded number of vertices. Step three takes the nearly planar subgraph (planar with constant number of vortices) and gives a *clique-weight* that captures the separation of the subgraph in the original graph. This weighting scheme generalizes the regular vertex weighting to capture the connectivity between the center subgraph and the remaining of the graph. Using this clique-weighting, it finds vortex-paths that separate the graph into components of size less than half. The two last parts of the proof are technically complicated by the requirement to deal with vortices, which are non-embeddable regions in the graph that break the traditional tools used for embeddable graphs, e.g. the Jordan's Curve Theorem. We believe one of our main technical contributions is our notion of *vortex-path* which plays a central role in the proof. We believe this notion has applications to other problems in H-minor-free graphs. We note that the above description is an oversimplification that hides several subtle difficulties.

The plan of this paper is as follows. Section 2 presents preliminaries to the Robertson and Seymour Theorem about the structure of H-minor-free graphs. The main proof is in Section 3 and the small world result is in Section 4. For the presentation of the main theorem and due to space limitation most of the proofs appear in [1]. In Section 5 we extended our result on path separators to doubling dimension isometric separators, and show some lower bounds. In particular, we show that the approach of [44], the tree-separator of planar graphs, cannot be extended to K_6-minor-free graphs. Distance labels and oracles are discussed in [1]. We conclude in Section 6 where we leave open the object location problem for bounded degree graphs.

2. PRELIMINARIES

2.1 The tree structure of H-minor-free graphs

A *tree-decomposition* of a graph G is a tree $\mathcal{T}$ whose vertices, called *bags*, are subsets of vertices of G such that:

1. for every vertex u of G, there exists a bag X of $\mathcal{T}$ such that $u \in X$;
2. for every edge $\{u,v\}$ of G, there exists a bag X of $\mathcal{T}$ such that $u, v \in X$; and
3. for every vertex u of G, the set of bags containing u induces a subtree of $\mathcal{T}$.

An important property following from the last two points is that every path between $u \in X$ and $v \in Y$ in G has to intersect all the bags on the path from X and Y in $\mathcal{T}$. Therefore, every bag X disconnects G provided that $\mathcal{T} \setminus X$ is composed of more than one subtree.

The *width* of a tree-decomposition $\mathcal{T}$ is $\max_{X \in \mathcal{T}} |X| - 1$. A *path-decomposition* of G is a tree-decomposition $\mathcal{T}$ where $\mathcal{T}$ is a path. The *treewidth of G* (resp. *pathwidth*) is the minimum of the width, over all tree-decompositions (resp. path-decompositions) of G.

A *joint set* in a graph having a tree-decomposition $\mathcal{T}$ is $X \cap Y$ for some bags X, Y of $\mathcal{T}$. The *torso* of G w.r.t. $\mathcal{T}$ is the graph denoted by $\tilde{G}$ where each joint set of $\mathcal{T}$ is filled-in by a complete graph.

For a graph G and a subset of vertices X, $G[X]$ denotes the subgraph of G induced by X. We denote by $\mathcal{T} \cap X$, the graph obtained by intersecting each bag of $\mathcal{T}$ with X, and by taking the subgraph of $\mathcal{T}$ induced by the resulting nonempty bags. If $G[X]$ is connected, then $\mathcal{T} \cap X$ is a tree-decomposition of $G[X]$.

A graph G is a *vortex* if there is a sequence of distinct vertices $u_1, \ldots, u_t$ of G, called *perimeter*, and a path-decomposition of G whose bags are $X_1, \ldots, X_t$ ordered such that the edges of the path are the pairs $\{X_i, X_{i+1}\}$ and $u_i \in X_i$ for all i. The width of a vortex is the width of its path-decomposition.

The following are standard terminologies about graphs on surfaces, we refer to [34] for an introduction. A *surface* is a non-null compact connected 2-manifold without boundary. A face of a graph embedded on a surface is *cellular* if its is homeomorphic to an open disc.

A graph G is h-*almost embeddable* on a surface Σ if there exists a set X of at most h vertices, called *apices*, such that $G \setminus X = G_\Sigma \cup W_1 \cup \cdots \cup W_t$, $t \leqslant h$ such that:

1. the graph G_Σ has an embedding on Σ;
2. the graphs W_i's are pairwise disjoint vortices of width at most h;
3. the perimeter of each vortex is the border of some cellular face of G_Σ.

An h-*nearly planar* graph is a graph that is h-almost embeddable on the sphere with no apices.

It is known that every graph excluding a planar minor H has bounded treewidth [37]. Actually, a much more general result, due to [39], gives us the structure of H-minor-free graphs for any graph H, that can be expressed as follows (see also [13, 27]):

THEOREM 4 (ROBERTSON & SEYMOUR [37, 39]). *For every H-minor-free graph G, there is a number $h = h(H)$ and a tree-decomposition $\mathcal{T}$ such that for every bag X of $\mathcal{T}$, either $|X| \leqslant h$, or $\tilde{G}[X]$ is h-almost embeddable on some surface on which H is not embeddable.*

A tree-decomposition $\mathcal{T}$ and the embeddings satisfying Theorem 4 can be constructed in polynomial time for fixed H, where the exponent depends on the number h. The algorithm comes directly from [38] (cf. [27, Lemma 15]). More recently, another algorithm has been presented in [11]. Note that the so constructed tree-decompositions have size linear in G, and the faces of all the embeddings are cellular.

3. FINDING A K-PATH SEPARATOR

In this section we assume that G is a weighted connected H-minor-free graph with n vertices, and that $\mathcal{T}$ is a linear size tree-decomposition for G satisfying Theorem 4. Let h be the constant, depending on the fixed graph H, involved in Theorem 4.

The family of graphs excluding a fixed minor is closed under minor taking, and thus under induced subgraph. Therefore, to prove Theorem 1 it suffices to prove that G has a k-path separator S without proving that the small ($\leqslant n/2$ vertices) connected components of $G \setminus S$ are k-path separable.

To find such a separator $S = P_0 \cup P_1 \cup \ldots$, we proceed iteratively in several steps. At a given step s, we compute a set of minimum cost paths P_s in the current graph G_s, and we keep the largest connected component G_{s+1} of $G_s \setminus P_s$, until all of them have $\leqslant n/2$ vertices.

The process is split into three main steps detailed below. Initially, we set $s = 0$, and $G_0 = G$, $\mathcal{T}_0 = \mathcal{T}$.

Step 1: Remove the center apices

The next classical result holds for every tree-decomposition of any graph G and any tree-decomposition $\mathcal{T}$ of G.

LEMMA 1. *There is a bag C of $\mathcal{T}$, called* center*, such that the connected components of $G \setminus C$ have at most $n/2$ vertices.*

Let C_0 be the bag given by Lemma 1 applied on G_0 and $\mathcal{T}_0$. C_0 can be found in polynomial time, since $\mathcal{T}_0$ has linear size. If $|C_0| \leqslant h$, then we set $S = P_0 = C_0$, and we have proved that S is a h-path separator for G_0. Otherwise, let P_0 be the set of apices of C_0, and let G_1 be a largest connected component of $G_0 \setminus P_0$. Since only G_1 can have more than $n/2$ vertices, it suffices to concentrate our attention on G_1, and to complete P_0 with some sequence of minimum cost paths $P_1, P_2, \ldots$ to form the wanted separator S for G_0.

Let $\mathcal{T}_1 = \mathcal{T}_0 \cap G_1$, and $C_1 = G_1 \cap C_0$. $\mathcal{T}_1$ is a tree-decomposition for G_1, and C_1 is a bag of $\mathcal{T}_1$ such that its torso, $\tilde{C}_1$, is h-almost embeddable with no apices.

Step 2: Make the center nearly planar

The goal of this step is to transform $\tilde{C}_1$ into a nearly planar graph by removing iteratively some set of minimum cost paths. For that we need some terminologies and definitions. For an introduction to topological graph theory we refer to [34].

Let Σ_0 be the surface on which $\tilde{C}_0$ is h-almost embedded, denote by $C_{\Sigma_0} = \tilde{C}_0 \cap \Sigma_0$ the part of $\tilde{C}_0$ embedded on Σ_0, and let g be the *Euler genus* of Σ_0 we define hereafter. The Surface Classification Theorem states that every surface is homeomorphic to a space obtained from the sphere by adding handles or crosscaps. The *Euler genus* of Σ_0 is either λ if Σ_0 was obtained by adding $\lambda \geqslant 0$ handles (orientable surface), or 2μ if Σ_0 was obtained by adding $\mu \geqslant 1$ crosscaps (nonorientable surface).

Note that in general there is no upper bound of the Euler genus of Σ_0 in terms of h. However, an important observation is that $g < |E(H)|$, because every graph X with a cellular embedding on Σ satisfies $g \leqslant |E(X)| - |V(X)| + 1$ from Euler's Formula.

A graph is (h,p,g)-*embeddable* if it is h-almost embeddable with no apices on a surface Σ of Euler genus at most g, and if the width of its vortices is at most $p \leqslant h$. So $\tilde{C}_1$ is (h,h,g)-embeddable, and h-nearly planar graphs are exactly the $(h,h,0)$-embeddable graphs.

DEFINITION 2. *Let G be an (h,p,g)-embeddable graph on Σ, and let G_Σ denote its embedded part. A* vortex-path *is a subgraph $\mathcal{V}$ of G that can be decomposed in $\mathcal{V} = Q_0 \cup X_1 \cup Y_1 \cup Q_1 \cup \cdots \cup X_t \cup Y_t \cup Q_t$ such that for every $i \in \{0, \ldots, t\}$:*

1. *Q_i, called* segment *of $\mathcal{V}$, is a path wholly contained in G_Σ;*
2. *there are t distinct vortices $W_1, \ldots, W_t$ in G such that X_i and Y_i are bags of W_i; and*
3. *one extremity of Q_i is the perimeter vertex of Y_i (for $i > 0$) and the other one is the perimeter vertex of X_{i+1} (for $i < t$), and no other vertices of Q_i is a perimeter vertex.*

The *projection* of a vortex-path $\mathcal{V}$, denoted by $\bar{\mathcal{V}}$, is the path formed by $Q_0 \cup e_1 \cup Q_1 \cup \cdots \cup e_t \cup Q_t$ where e_i, $i > 0$, is an extra edge added to G between the perimeter vertex of X_i and the perimeter vertex of Y_i, and embedded on the face (which is cellular) of the vortex W_i. We observe that $\mathcal{V}$ is not necessarily a path of G, but its projection $\bar{\mathcal{V}}$ is a curve of Σ as depicted on Fig. 1(b).

With every path P of G whose extremities are in G_Σ, one can associate a vortex-path $Q_0 \cup X_1 \cup Y_1 \cup \cdots \cup Q_t$ defined as follows (see Fig. 1): Start a walk on P from one extremity until encountering the first perimeter vertex, say x_1. This first part of P forms the segment Q_0, and the vortex bag corresponding to x_1 is X_1. Then, continue the walk on P and select y_1 to be the last perimeter vertex on P belonging to the same vortex that x_1 belongs to. This forms Y_1, the bag whose perimeter vertex is y_1. Then, continue along P up to the next perimeter vertex, forming segment Q_1, and so on. Observe that the part of P between x_1 and y_1 may enter and leave many vortices, each one several times. However, by construction, the vortex-path of P enters and leaves pairwise distinct vortices.

The path from u to v in a rooted tree T is denoted by $T(u,v)$, and $T(u,v)$ is said to be *monotone* if u and v are relatives, i.e., u is ancestor of v or v is ancestor of u.

We have the following result which is the key of Step 2:

LEMMA 2. *Let G be a (h,p,g)-embedded graph for $g > 0$, and let T be a spanning tree of G rooted in the embedded part of G. Then, there exist two vortex-paths $\mathcal{V}_1, \mathcal{V}_2$ such that each segment is a monotone path of T, and each connected component of $G \setminus (\mathcal{V}_1 \cup \mathcal{V}_2)$ is $(h+1,p,g-1)$-embedded. Moreover, $\mathcal{V}_1, \mathcal{V}_2$ and the embeddings can be computed in polynomial time.*

Let us sketch the proof of this key lemma. The idea is to select a suitable nontree edge $\{u,v\}$ of T, and to consider the two vortex-paths of the paths $T(u,r)$ and $T(v,r)$, where r is the root of T. Intuitively, in classical graphs on surface theory, we want that the fundamental cycle induced by $\{u,v\}$ in T reduces the Euler genus of Σ. Unfortunately, this fundamental cycle is not wholly embedded on Σ (due to vortices), and the use of the vortex-paths instead of the paths is required. Note that the vortex-path of $T(u,r)$ is not necessarily a path, it does not include all the vertices of $T(u,r)$. However, its projection is a path of Σ, and each segment is a monotone path of T. In fact, it is possible to reorganize the segments of the two vortex-paths such that the two projections plus the nontree edge form a closed curve C of Σ. (Indeed, the projections of the vortex-paths might intersect in some common vortex, but this can be avoid via exchanging segments between the two vortex-paths). Now, two cases occur. Case 1: If the removal of C reduces the Euler genus of Σ, then one can replace each side of C with a disc (it may happen that C has only one side if Σ is nonorientable). This disc forms the face of a new vortex by grouping all the vortices C intersects. This new embedding decreases by one g and increases the number of vortices h by one in the worst-case (this occurs, for instance, if C intersects one vortex and has two sides). Case 2: If the removal of C does not reduce the Euler genus of Σ, then one region of $\Sigma \setminus C$, say R, must be planar. We can then construct a new cycle C' thanks to a new vortex-path that either reduces the Euler genus of Σ (Case 1), or provides a planar region strictly including R (by at least one edge). By finiteness of G, Case 1 always occurs, and so an $(h+1,p,g-1)$-embedding of G can always be constructed by removing two vortex-paths.

We will now carefully apply iteratively Lemma 2 to the graph $\tilde{C}_1$ in order to get a $(h+g)$-nearly planar graph. For that we assume the following:

1. G_s is connected;
2. $\mathcal{T}_s$ is tree-decomposition of G_s;
3. C_s is a bag of $\mathcal{T}_s$; and
4. $\tilde{C}_s$ is $(h-1+s, h, g+1-s)$-embedded (the torso is w.r.t. $\mathcal{T}_s$).

For $s = 1$ we check all the conditions hold. Here is the main loop:

Main loop:

1. If $C_s \cap C_{\Sigma_0} = \varnothing$ or $\tilde{C}_s$ is $(h-1+s, h, 0)$-embedded, we jump directly to Step 3. The former case may occur if, for instance, C_s is a subgraph of C_0's vortices.

2. Choose a root $r_s \in C_s \cap C_{\Sigma_0}$, and perform from r_s a minimum cost path tree A_s in G_s spanning the vertices of C_s.

3. Let T_s be the graph $A_s \cap C_s$ augmented with the edges between any two vertices u, v satisfying the two conditions:
   (a) u and v are both in the same joint set of C_s; and
   (b) $A_s(u,v)$ is monotone and $A_s(u,v) \cap C_s = \{u,v\}$.

   LEMMA 3. *T_s is a spanning tree of $\tilde{C}_s$.*

   Note that by construction, for any pair $u, v \in C_s$, $V(T_s(u,v)) \subseteq V(A_s(u,v))$, and that if $T_s(u,v)$ is monotone, then so is $A_s(u,v)$.

4. We now apply Lemma 2 to the graph $\tilde{C}_s$ with the spanning tree T_s, and we obtain the vortex-paths $\mathcal{V}_j = Q_0^j \cup X_1^j \cup Y_1^j \cup \cdots \cup Q_{t_j}^j$ for $j \in \{1,2\}$. Assume that each segment Q_i^j is a path going from u_i^j to v_i^j. By Lemma 2, $Q_i^j = T_s(u_i^j, v_i^j)$ is monotone.

5. Finally, we update $S = S \cup P_s$ where:
$$P_s = \bigcup_{j\in\{1,2\}} \left(A_s(u_0^j, v_0^j) \cup \bigcup_{i=1}^{t_j} \left(X_i^j \cup Y_i^j \cup A_s(u_i^j, v_i^j) \right) \right)$$

6. It is clear that $\mathcal{V}_1 \cup \mathcal{V}_2 \subseteq P_s$, since $A_s(u_i^j, v_i^j)$ includes the segment $T_s(u_i^j, v_i^j) = Q_i$ for all i,j. Therefore, according Lemma 2, each connected components of $G_s \setminus P_s$ is $(h-1+s+1, h, g+1-s-1)$-embedded, i.e., $(h-1+(s+1), h, g+1-(s+1))$-embedded.

7. Observe that P_s is a set of minimum cost paths in G_s. Indeed, $A_s(u_i^j, v_i^j)$ is monotone in A_s (by monotonicity of $T_s(u_i^j, v_i^j)$), and thus is a minimum cost path. The number of paths in P_s is at most $2 + (t_1 + t_2) \cdot (2p+1)$ because each vortex bag has at most $p+1$ vertices (we count only p vertices since the perimeter vertex is a part of an entering or leaving segment). In a vortex-path, the number of segments cannot exceed the total number of vortices, so by assumption $t_j \leqslant h-1+s$. We have also that $p \leqslant h$, therefore $k_s \leqslant 2 + 2(h-1+s)\cdot(2h+1)$, where k_i denotes the number of paths used to defined P_i.

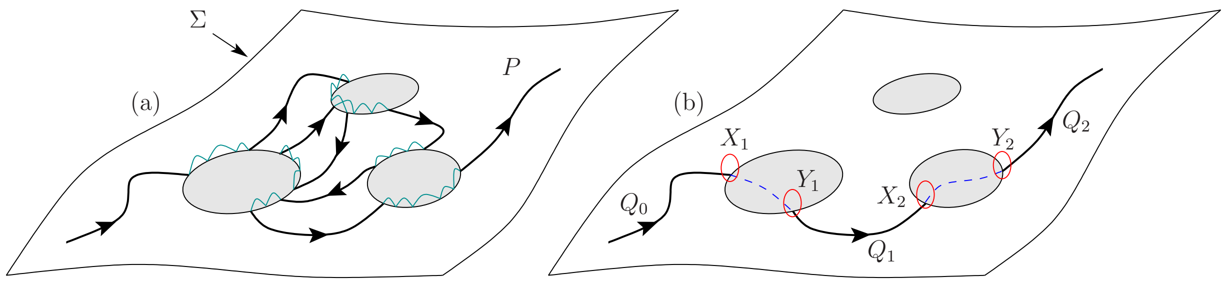

Figure 1: (a) A path P intersecting three vortices; (b) its vortex-path $\mathcal{V} = Q_0 \cup X_1 \cup Y_1 \cup \cdots \cup Q_2$, and its projection.

8. Let G_{s+1} be the largest connected component of $G_s \setminus P_s$. Let $\mathcal{T}_{s+1} = \mathcal{T}_s \cap G_{s+1}$, and $C_{s+1} = G_{s+1} \cap C_s$. $\mathcal{T}_{s+1}$ is a tree-decomposition for G_{s+1}, and C_{s+1} is a bag of $\mathcal{T}_{s+1}$. An important observation is that the joint sets of $\tilde{C}_{s+1}$ are included in those of $\tilde{C}_s$ since cliques of $\tilde{C}_s$ were not split by the removal of P_s. Therefore $\tilde{C}_{s+1}$ corresponds to a component of the graph computed from $\tilde{C}_s$ by Lemma 2. Therefore, $\tilde{C}_{s+1}$ is $(h-1+(s+1), h, g+1-(s+1))$-embedded as required.

   All the pre-conditions of the main loop are satisfied, and we continue with $s = s+1$.

Step 3: Split the nearly planar center

At this stage $s \leqslant g+1$, and $S = P_0 \cup \cdots \cup P_{s-1}$. Currently the total number of minimum cost paths in S is bounded by

$$h + \sum_{i=1}^{s-1} k_i \;\leqslant\; h + \sum_{s=1}^{g} (2 + 2(h-1+s) \cdot (2h+1)) \;=\; O(hg(h+g)) \,.$$

Moreover, either $C_s \cap C_{\Sigma_0} = \varnothing$, or $\tilde{C}_s$ is $(h+g, h, 0)$-embeddable, i.e., $\tilde{C}_s$ is $(h+g)$-nearly planar. We eliminate the former case.

Lemma 4. *If $C_s \cap C_{\Sigma_0} = \varnothing$, then there was a vortex bag X of C_0 such that the connected components of $G_s \setminus X$ have at most $n/2$ vertices.*

According to Lemma 4, if $C_s \cap C_{\Sigma_0} = \varnothing$, we can find in polynomial time the vortex bag X, and set $P_s = X \cap G_s$. As $X \cap G_s$ is a proper subset of the vortex bag X — indeed the perimeter vertex of X was belonging to C_{Σ_0} and thus is missing from C_s — we get $k_s = |X \cap G_s| \leqslant h$. Thus we have proved that $S = P_0 \cup \cdots \cup P_s$ is a k-path separator for G for $k = O(hg(h+g))$ which is bounded by some function depending only on H since we have seen that $g < |E(H)|$.

Our goal now is to find a separator in C_s that also splits G_s in small components. For that we need some definitions and the use of two lemmas.

A *clique-weight* for a graph G is a pair $(\mathcal{K}, \omega)$ where $\mathcal{K}$ is a set of cliques of G, and ω a function associating with each clique $K \in \mathcal{K}$ a real $\omega(K) \geqslant 0$. For every subgraph A of G having a clique-weight $(\mathcal{K}, \omega)$, we define the *weight* of A as:

$$f(A) \;=\; \sum_{\substack{K \cap A \neq \varnothing \\ K \in \mathcal{K}}} \omega(K) \,.$$

A clique-weight for G is a non-trivial generalization of a vertex-weight function. This can be seen by setting $\mathcal{K}$ as the collection of cliques each one composed of one single vertex of G. So that $\omega(\{u\})$ is the weight of vertex u, and $f(A)$ equals the sum over the vertices of A of the function ω. Unlike vertex-weight functions, observe that if $A \cup B \subseteq G$ and $A \cap B \neq \varnothing$, then $f(A) + f(B) \nleqslant f(G)$ for general clique-weights (for instance if there is a clique intersecting A and B). However, $f(A) + f(B) \leqslant f(G)$ whenever A and B are connected components of $G \setminus S$ for any S.

A *half-size separator* for G with clique-weight $(\mathcal{K}, \omega)$ is a set of vertices whose removal leaves connected components of weight at most half the weight of G, i.e., of weight at most $f(G)/2$.

The motivation for this generalization is the following simple result, that states that every half-separator, for a suitable clique-weight, is actually the wanted separator for the whole graph.

Lemma 5. *Let C be a center of tree-decomposition of an n-vertex graph G. Then, in polynomial time, one can construct a clique-weight $(\mathcal{K}, \omega)$ for the torso $\tilde{C}$ such that, given any half-size separator S of $\tilde{C}$, the connected components of $G \setminus S$ have at most $n/2$ vertices.*

So virtually, we can forget for a while about the small connected components joining C and concentrate our attention on half-size separators of $\tilde{C}$.

The following result is non-trivial variant of the Lipton-Tarjan planar separator [33], and is a generalization of the three-leaves tree separator of Thorup [44]. It applies to a more general graph (a planar graph is just a 0-nearly planar graph) and with a more general vertex-weight function (a vertex-weight function is just a particular clique-weighting).

Lemma 6. *Let G be a nearly planar graph with a clique-weight and a spanning tree T rooted in the planar part of G. In polynomial time, one can construct a half-size separator for G composed of at most three vortex-paths whose segments consist of monotone paths of T.*

Final construction:

At this step $C_s \cap C_{\Sigma_0} \neq \varnothing$. So, we went out from Step 2 because $\tilde{C}_s$ is $(h+g)$-nearly planar. Let A_s and T_s be the trees constructed as in Point 2 and Point 3 of the main loop of Step 2. A_s is a minimum cost tree rooted at some vertex of the planar part of G_s, and by Lemma 3, T_s is a spanning tree of $\tilde{C}_s$.

We now construct the clique-weight $(\mathcal{K}, \omega)$ for $\tilde{C}_s$ in G_s (with tree-decomposition $\mathcal{T}_s$) as done in Lemma 5, and we apply Lemma 6 to the $(h+g)$-nearly planar graph $\tilde{C}_s$ and its spanning tree T_s.

Let $\mathcal{V}_1$, $\mathcal{V}_2$, $\mathcal{V}_3$ be the vortex-paths constructed in Lemma 6, where $\mathcal{V}_j = Q_0^j \cup X_1^j \cup Y_1^j \cup \cdots \cup Q_{t_j}$ with $Q_i^j = T_s(u_i^j, v_i^j)$ for all i, j. Similarly to Point 5 of Step 2, we update the separator $S = S \cup P_s$ with:

$$P_s = \bigcup_{j\in\{1,2,3\}} \left(A_s(u_0^j, v_0^j) \cup \bigcup_{i=1}^{t_j} \left(X_i^j \cup Y_i^j \cup A_s(u_i^j, v_i^j) \right) \right)$$

Note that $\mathcal{V}_1 \cup \mathcal{V}_2 \cup \mathcal{V}_3 \subseteq P_s$, since $A_s(u_i^j, v_i^j)$ includes the segment $T_s(u_i^j, v_i^j) = Q_i$ for all i, j. Therefore, according to Lemma 6, P_s is a half-size separator of $\tilde{C}_s$. By Lemma 5, all the connected components of $G_s \setminus P_s$ have at most $|V(G_s)|/2 \leqslant n/2$ vertices.

We complete the proof by observing that P_s is a set of minimum cost paths in G_s, as $A_s(u_i^j, v_i^j)$ is monotone and thus is a minimum cost path in G_s for all i, j. The number of paths in P_s is at most $3 + (t_1 + t_2 + t_3) \cdot (2(h+g)+1) \leqslant 3 + 3(h+g) \cdot (2(h+g)+1) = O(hg(h+g))$.

In all the cases we have proved that $S = P_0 \cup \cdots \cup P_s$ is a k-path separator for G_0 for $k = O(hg(h+g)) = O(h|E(H)| \cdot (h + |E(H)|))$, which is bounded by some function of H as claimed. This completes the proof of Theorem 1.

Note. The above proof of Theorem 1 can be strengthened to construct a k-path vertex-weighted separator, that is a separator S that splits G (having edge and vertex-weights) in components of vertex-weight at most half of the total vertex-weight of G, S still composed of minimum cost paths as defined by property (P1). For that, lemmas 1 and 5 can be easily adapted.

4. SMALL-WORLDS

Definition 3 (Augmentation distribution). *An* augmentation distribution *of a graph $G = (V, E)$ is a function $\mathcal{D} : V \times V \to \mathbb{R}^+$ such that for any $v \in V$ the function $\mathcal{D}(v, \cdot)$ is a distribution function on V, i.e., $\forall v \in V$, $\sum_{u\in V} \mathcal{D}(v, u) = 1$.*

Definition 4 (Augmenting a graph). *Given a weighted graph $G = (V, E, \omega)$ and an augmentation distribution $\mathcal{D}$, let $\langle G, \mathcal{D}\rangle$ be a distribution on graphs with vertex set V formed by adding to E, for each $v \in V$ one directed edge $e = (v, u)$ where u is independently chosen via $\mathcal{D}(v, \cdot)$ and setting the weight of the edge to be $\omega(e) = d_G(v, u)$.*

Let G be a weighted undirected n-vertex k-path separable graph with aspect ratio Δ. Distance is extended naturally to sets, given sets $U, V \subseteq V(G)$, let $d(U, V) = \min\{d(u, v) \mid u \in U, v \in V\}$. For any subgraph H of G, let $S(H)$ be the k-path separator of H that separates H in to small components, i.e., of size at most $\frac{1}{2}|V(H)|$.

We build a rooted tree $\mathcal{T}$ called the *decomposition tree* of G as follows. The vertices of $\mathcal{T}$ are subgraphs of G. The root of $\mathcal{T}$ is G. For any $H \in V(\mathcal{T})$, its children $J_1, \ldots, J_t$ are the connected components induced by $H \setminus S(H)$. Observe that the depth of $\mathcal{T}$ is at most $\log n$.

The augmentation distribution. Fix a vertex $v \in G$, let $H_1(v), H_2(v), \ldots, H_r(v)$ be the path in $\mathcal{T}$ starting from the root G containing all the vertices $H \in V(\mathcal{T})$ such that $v \in H$. Let $\tau \in \{1, \ldots, r\}$ be a uniformly random variable.

Given τ let Q be a uniformly random path out of the paths in $S(H_\tau(v))$. Given the path Q of $P_i \in S(H_\tau(v))$ and the graph $J = H \setminus \bigcup_{j<i} P_j$ we define the following landmarks $L = L(Q) \subseteq Q$.

Let $Q = x_1, \ldots, x_t$ be the vertices on the path. Let $x_c \in Q$ be a vertex such that $d_J(v, x_c) = d_J(v, Q)$, and let $d = d_J(v, x_c)$. Now we consider separately the paths $Q_{(-1)} = x_1, \ldots, x_c$ and $Q_{(+1)} = x_c, \ldots, x_t$. For each $j \in \{-1, 1\}$: (1) For each $i \in \{0, 1, 2, 3, \ldots, 10\}$ add to L the first vertex in $Q_{(j)}$ such that $d_{Q_{(j)}}(x_c, x) \geqslant (i/2) \cdot d$. (2) For each $i \in \{0, 1, 2, 3, \ldots, \lceil\log\Delta\rceil\}$ add to L the first vertex in $Q_{(j)}$ such that $d_{Q_{(j)}}(x_c, x) \geqslant 2^i \cdot d$.

Observe that $|L| = O(\min\{t, \log\Delta\})$. After vertex v chooses a component index τ and chooses a path Q, it chooses uniformly at random a landmark ℓ in $L = L(Q)$. This completes the distribution for v.

Claim 1. *Let L be the set of landmarks chosen by v in graph J in respect to path Q. For any $x \in Q$, there exists $\ell \in L$ such that $d_Q(\ell, x) \leqslant \frac{3}{4} d_J(v, x)$.*

Proof. Let $d = d_J(v, x_c)$. There are two cases to consider. If $d_J(v, x) \leqslant 4d$ then $d_Q(x_c, x) \leqslant 5d$ since Q is a shortest path in J. Therefore there exists $0 \leqslant i \leqslant 10$ such that $\ell \in L$ is the first vertex on Q with $d_Q(x_c, \ell) > id/2$ and $d_Q(\ell, x) \leqslant d/2 \leqslant d_J(v, x)/2$ since $d_J(v, x) \geqslant d$.

Otherwise if $d_J(v, x) > 4d$ then $\frac{3}{4} d_J(v, x) \leqslant d_J(v, x) - d \leqslant d_Q(x_c, x) \leqslant d_J(v, x) + d \leqslant \frac{5}{4} d_J(v, x)$. Let $i \in \{0, 1, 2, 3, \ldots, \lceil\log\Delta\rceil\}$ be the index such that $d_J(v, x)/2 < 2^i d \leqslant d_J(v, x)$ then for the corresponding $\ell \in L$ with $d_Q(x_c, \ell) \geqslant 2^i d$ we have $d_Q(\ell, x) \leqslant \frac{3}{4} d_J(v, x)$. □

Theorem 3. *For every connected n-vertex k-path separable weighted graph G with aspect ratio Δ there exists a distribution of long-range edges, computable in polynomial time, such that the greedy routing performs in $O(k^2 \log^2 n \log^2 \Delta)$ expected number of hops.*

Proof. Consider any target $t \in V(G)$. Let $\mathcal{T}$ be the decomposition tree of G. Let $H_1, H_2, \ldots, H_r$ be the path in $\mathcal{T}$ starting from the root G containing all the vertices $H \in V(\mathcal{T})$ such that $t \in H$. For each H_i, let $Q_1^i, \ldots, Q_k^i$ be the paths (possibly empty or one vertex paths) of $S(H_i)$. Let $\mathcal{P} = \left\{Q_j^i \mid 1 \leqslant i \leqslant r, 1 \leqslant j \leqslant k\right\}$, observe that $|\mathcal{P}| = O(k \log n)$ and $\mathcal{P} \neq \varnothing$.

For any $Q \in \mathcal{P}$, let $x(Q) \in Q$ be a vertex such that $d(x(Q), t) = \min_{y\in Q} d(y, t)$. Let $\Phi : V \times \mathcal{P} \to \mathbb{R}$ be the function:

$$\Phi(u, Q) = \max\{d(u, t) - d(x(Q), t), 0\}.$$

The following is immediate from the definition of Φ.

Claim 2.
(1) If $d(v, t) \leqslant d(u, t)$ then $\Phi(v, Q) \leqslant \Phi(u, Q)$ for all $Q \in \mathcal{P}$.
(2) If $\forall Q \in \mathcal{P} : \Phi(v, Q) = 0$ then $p = t$.

Consider a current vertex u that wants to reach $t \neq u$. Let R be a shortest path on G from u to t, let $R' = R \setminus \{u\}$. Let $Q \in \mathcal{P}$ be the first path that intersects R', formally $Q = \operatorname{argmin}_{Q\in\mathcal{P}} \{d_{R'}(u, Q)\}$. Observe that $\Phi(u, Q) > 0$. We will now show that with probability $O(k \log n \log\Delta)^{-1}$ we reach a vertex v such that $\Phi(v, Q) \leqslant \frac{3}{4}\Phi(u, Q)$.

Let i, τ be the indexes such that $Q \in P_i \in S(H_\tau(u))$. Let $J = H_\tau(u) \setminus \bigcup_{j<i} P_j$. Since Q is the first path to intersect R then $d(u, x(Q)) = d_J(u, x(Q))$.

Hence with probability $O(k \log n \log \Delta)^{-1}$ vertex u chooses $\tau \in O(\log n)$ and $Q \in O(k)$ and the landmark $\ell \in O(\log \Delta)$ from Claim 1 that covers $x(Q)$. If this happens then u has an edge to some vertex ℓ such that $d(\ell, x(Q)) \leqslant d_Q(\ell, x(Q)) \leqslant \frac{3}{4} d_J(u, x(Q)) = \frac{3}{4} d(u, x(Q))$.

Hence, given such an ℓ, any greedy step that u performs will lead to a vertex v such that $\Phi(v, Q) \leqslant \Phi(\ell, Q) \leqslant \frac{3}{4}\Phi(u, Q)$.

Since such an event can happen at most $O(\log \Delta)$ times for each of the $O(k \log n)$ paths in $\mathcal{P}$ then by linearity the expected greedy diameter is at most $O(k^2 \log^2 n \log^2 \Delta)$. □

Note 1: when all separator paths are actually one vertex (as in the case of bounded treewidth graphs) then the greedy diameter can be reduced to $O(k^2 \log^2 n)$ by setting $L(Q)$ to simply be the single vertex of $Q = \{q\}$.

Note 2: if all separator paths form a graph with diameter δ and additionally the graph is unweighted then the greedy diameter reduces to $O(\log^2 n + \delta \log n)$. After u chooses τ, instead of randomly choosing one of the vertices in $S(H_\tau(u))$ it chooses the closest vertex from $S(H_\tau(u))$. Now consider a target t and a current vertex u and let H be the minimal component of $\mathcal{T}$ that contains both. If $d_H(u, S(H)) \leqslant 2\delta$ then u will cross $S(H)$ after $O(\delta)$ steps. Otherwise, with probability $O(\log n)^{-1}$, vertex u will have a long range contact to x the closest vertex in $S(H)$. Since $d_H(u, S(H)) > 2\delta$ then $d_H(x, t) \leqslant \delta + d_H(S(H), t)$. Hence if this event occurs then this level of the hierarchy will be crossed in $O(\delta)$ steps. Hence the expected number of hops to cross a separator is $O(\log n + \delta)$.

Note 3: if G is not k-path separable but is (k, α)-doubling separable (see Section 5) then instead of choosing $O(\log \Delta)$ landmarks we choose $O(2^{O(\alpha)} \log \Delta)$ landmarks using the construction of Slivkins [41]. The expected greedy diameter becomes $O(2^{O(\alpha)} k^2 \log^2 n \log^2 \Delta)$.

Corollary 1. *For any n-vertex graph G there exists a long range link augmentation $\mathcal{D}$ such that the expected hop count of greedy routing on $\langle G, \mathcal{D} \rangle$ is*

1. $O(k^2 \log^2 n)$, *if G is a weighted treewidth-k graph;*
2. $O(2^{O(\alpha)} k^2 \log^2 n \log^2 \Delta)$, *if G is a weighted (k, α)-doubling separable graph with aspect ratio Δ.*

5. ABOUT K-PATH SEPARABILITY

5.1 Lower bounds for sparse graphs

Although our results applies to a large family of sparse graphs, graphs excluding K_r have $O(r\sqrt{\log r} \cdot n)$ edges [43], no techniques can achieve similar performances for general sparse graphs. More precisely:

Theorem 5. *For every fixed $0 \leqslant \varepsilon < 2$, there are unweighted graphs with at most n vertices and with $O(n)$ edges on which every stretch-$(1+\varepsilon)$ distance labeling scheme with L-bit labels, or on which every stretch-$(1+\varepsilon)$ labeled routing scheme with L-bit labels and routing tables implies $L = \Omega(\sqrt{n})$. In particular such sparse graphs are not k-path separable for $k = o(\sqrt{n}/\log^2 n)$.*

Proof. Plug the classical lower bounds into a graph of $N = \lceil \sqrt{n} \rceil$ vertices and so with less than $N^2 = O(n)$ edges. It is known that, for every stretch $s < 3$, there are unweighted N-vertex graphs for which every stretch-s routing scheme requires $\Omega(N)$-bit labels [24], or for which every stretch-s distance labeling scheme requires $\Omega(N)$-bit labels [25, 45].

Let us fix $\varepsilon = 1$ (or any value $0 < \varepsilon < 2$). From Theorem 2, every k-separable graph supports a 2-approximate distance labeling scheme with L-bit labels where $L = O(k \log^2 n)$. So for the above sparse graph, $L = \Omega(\sqrt{n})$, and thus it cannot be k-separable for $k = o(\sqrt{n}/\log^2 n)$. □

5.2 Strongly separable graphs

Let us first observe that the graph composed of a path of $n/2$ vertices and of a stable set of $n/2$ vertices with all edges between the path and stable vertices, contains $K_{n/2,n/2}$ as minor and is a 1-path separable (the whole path) if the path edges have weight 1 and the other edges weight $n/2$. Thus the family of $O(1)$-path separable graphs does not reduce to $K_{O(1)}$-minor-free graphs.

A k-path separator S for G is said *strong* if $S = P_0$, i.e., S reduces to the union of at most k minimum cost paths in G. Although [44] proved that planar graphs are strongly 3-path separable, we stress that for a fixed k this natural definition actually captures much less graphs. For instance, we construct some K_6-minor-free graphs that are not strongly $\sqrt{n}$-path separable, in contrast Theorem 1 implies that they are $O(1)$-path separable.

Unless explicitly expressed, graphs are weighted connected graphs.

Theorem 6.

1. *Every planar graph is strongly 3-path separable [44].*
2. *Every H-minor-free graph with n vertices is strongly $O(\sqrt{n})$-separable.*
3. *There are unweighted K_6-minor-free graphs with n vertices for which every strong k-path separator requires $k = \Omega(\sqrt{n})$.*

Proof. It is known that K_r-minor-free graphs with n vertices have treewidth $O(r^{3/2}\sqrt{n})$ [6]. By Theorem 7, we conclude that the graphs excluding a fixed minor of at most r vertices are strongly $O(\sqrt{n})$-separable.

Let G be the graph composed of a $t \times t$ mesh augmented with a universal vertex, i.e., a vertex that neighbors all the vertices of the mesh. Since the mesh is K_5-minor-free, G is K_6-minor free. Let $n = t^2 + 1$ be the number of vertices of G.

Let S be any strong k-path separator for G. We remark that, since G has diameter two, any union of k shortest paths in G contains at most $3k$ vertices. In other words, $|V(S)| \leqslant 3k$.

Assume that $k < t/3$. Then, $|V(S)| < t$, and thus S cannot intersect all the rows or all the columns of the underlying $t \times t$ mesh of G. It follows that their must exists a connected component in $G \setminus S$ with at least $\sum_{i=1}^{t} i = t(t+1)/2$ vertices. Indeed, we can check that, for every $c \leqslant t$, the largest connected component in a $t \times t$ mesh in which c vertices have been removed, has at least $\sum_{i=t-c}^{t} \max\{i, t\}$ vertices, obtained by removing the diagonal vertices from coordinates $(c+1, 1)$ to (t, c).

Note that $t(t+1)/2 > (t^2+1)/2 = n/2$ for $t > 1$. This is a contradiction with the fact that all connected components of $G \setminus S$ must have at most $n/2$ vertices. Therefore, $k \geqslant t/3 = \Omega(\sqrt{n})$. □

THEOREM 7. *Every treewidth-r graph is K_{r+2}-minor-free and is strongly $(r+1)$-path separable. Moreover, there are unweighted treewidth-r graphs for which every k-path separator requires $k \geqslant r/2$.*

PROOF. Using Lemma 1, every connected graph having a tree-decomposition $\mathcal{T}$ of width w has a strong $(w+1)$-path separator. So, it suffices to consider a tree-decomposition of optimal width r for treewidth-r graphs.

Consider now the graph $G = K_{r,n-r}$ with $n \geqslant 2r$. Its treewidth is r (cf. [31]), and every set of $r-1$ vertices does not disconnect G. Moreover, every shortest path in G includes at most two vertices of each stable. Thus, if $n \geqslant 2r$, in order to obtain connected components of size at most $n/2$ we need to disconnect G. And, to disconnect G we need to delete at least $r/2$ shortest paths. □

5.3 More general separators

Observe that a 3D mesh has no k-path separators for bounded k, whereas there exists a "2D mesh separator". And this graph appears to be an easy graph for solving object location problem. This leads to the following natural extension of the k-path separability definition.

Recall that a subgraph H of G is *isometric* if $d_H(x,y) = d_G(x,y)$ for all $x,y \in H$. H is of doubling dimension α if for every radius r and vertex $x \in H$, the radius-$2r$ ball in H centered at x can be covered by at most 2^α radius-r balls of H.

A weighted graph G with n vertices is (k,α)-*doubling separable* is there exists a subgraph S, called (k,α)-doubling separator, satisfying conditions P2 and P3 of the previous definition and where condition P1 is replaced by:

(P1') $S = P_0 \cup P_1 \cup \cdots$, where each subgraph P_i is the union of k_i isometric subgraphs of doubling dimension at most α in $G \setminus \bigcup_{j<i} P_j$.

So a k-path separator is nothing else than a $(k,1)$-doubling separator. Based on $O((\alpha/\varepsilon)^{O(\alpha)} \log \Delta)$-bit $(1+\varepsilon)$-approximate labels [42], we can actually show the following:

THEOREM 8. *For every n-vertex (k,α)-doubling separable weighted graph G, and for every $\varepsilon > 0$, there exists a data-structure of $O(\tau \cdot n \log n)$ space supporting $(1+\varepsilon)$-approximate distance queries in $O(\tau \log n)$ time, where $\tau \leqslant k(\alpha/\varepsilon)^{O(\alpha)}$. Moreover, the data-structure is polynomial time constructible, and can be distributed as a $(1+\varepsilon)$-approximate distance labeling with $O(\tau \cdot \log n)$ space labels.*

6. CONCLUSION

We showed that weighted graphs excluding a fixed minor have a k-path separator for some constant k, i.e., a recursive pruning decomposition in k minimum cost paths. Our scheme is polynomial time constructible, and the associated data-structures have optimal size up to a $\log n$ factor, and allow us to solve several object location problems: 1) approximate compact distance oracle, 2) distance labeling, 3) compact routing, and 4) small-worldization. The approximation factor is $1+\varepsilon$ for any fixed $\varepsilon > 0$.

We generalized our decomposition to separator isometric subgraphs of low doubling dimension in order to capture more sparse graphs. Unfortunately, we have showed that efficient distance oracles cannot be achieved on the class of all sparse n-vertex graphs (even unweighted graphs).

We leave open several questions:

1) Proves or disprove that K_r-minor-free graphs are k-path separable for $k = r^{O(1)}$. By Theorem 7, $k \geqslant r/2$. A stronger lower bound would be interesting. Currently, an upper bound on k (actually an upper bound on the treewidth) is known only for graphs excluding a planar minor H: $k \leqslant 20^{2(2|V(H|+4|E(H)|)^5}$. Note that this is potentially greater than $20^{10^6 r^5}$ for a maximal planar graph H on r vertices! If H is an $r \times r$ grid, then $k \leqslant 20^{2r^5}$. This latter value is conjectured to be only $O(r^2 \log r)$ [40].

2) Is there is a constant $c > 0$ such that every weighted bounded degree graph is $(\log^c n, c \log\log n/\log\log\log n)$-doubling separable? By Theorem 8, this would imply $(1+\varepsilon)$-approximate solutions with distributed data-structures (labels) of poly-log space for these graphs.

7. REFERENCES

[1] I. ABRAHAM AND C. GAVOILLE, *Object location using path separators*, Research Report RR-1394-06, LaBRI, University of Bordeaux 1, France, Mar. 2006.

[2] I. ABRAHAM, C. GAVOILLE, A. V. GOLDBERG, AND D. MALKHI, *Routing in networks with low doubling dimension*, in 26th International Conference on Distributed Computing Systems (ICDCS), IEEE Computer Society Press, July 2006. To appear.

[3] I. ABRAHAM, C. GAVOILLE, AND D. MALKHI, *Compact routing for graphs excluding a fixed minor*, in 19th International Symposium on Distributed Computing (DISC), vol. 3724 of LNCS, Springer, Sept. 2005, pp. 442–456.

[4] ———, *On space-stretch trade-offs: Lower bounds*, in 18^{th} Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA), ACM Press, July 2006.

[5] J. ALBER, H. FERNAU, AND R. NIEDERMEIER, *Graph separators: A parameterized view*, Journal of Computer and System Sciences, 67 (2003), pp. 808–832.

[6] N. ALON, P. D. SEYMOUR, AND R. THOMAS, *A separator theorem for graphs with an excluded minor and its applications*, in 22nd Annual ACM Symposium on Theory of Computing (STOC), ACM Press, May 1990, pp. 293–299.

[7] ———, *Planar separators*, SIAM Journal on Discrete Mathematics, 7 (1994), pp. 184–193.

[8] B. AWERBUCH AND D. PELEG, *Sparse partitions*, in 31st Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society Press, Oct. 1990, pp. 503–513.

[9] ———, *Routing with polynomial communication-space trade-off*, SIAM Journal on Discrete Mathematics, 5 (1992), pp. 151–162.

[10] H. T.-H. CHAN, A. GUPTA, B. M. MAGGS, AND S. ZHOU, *On hierarchical routing in doubling metrics*, in 16th Symposium on Discrete Algorithms (SODA), ACM-SIAM, Jan. 2005.

[11] E. D. DEMAINE, F. V. FOMIN, M. T. HAJIAGHAYI, AND D. M. THILIKOS, *Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs*, Journal of the ACM, 52 (2005), pp. 866–893.

[12] E. D. DEMAINE AND M. HAJIAGHAYI, *Graphs excluding a fixed minor have grids as large as*

treewidth, with combinatorial and algorithmic applications through bidimensionality, in 15th Symposium on Discrete Algorithms (SODA), ACM-SIAM, Jan. 2005, pp. 682–689.

[13] R. Diestel and R. Thomas, *Excluding a countable clique*, Journal of Combinatorial Theory, Series B, 76 (1999), pp. 41–67.

[14] H. N. Djidjev and S. M. Venkatesan, *Reduced constants for simple cycle graph separation*, Acta Informatica, 34 (1997), pp. 231–243.

[15] P. Duchon, N. Hanusse, E. Lebhar, and N. Schabanel, *Could any graph be turned into a small-world?*, Research Report 2004-62, LIP, École Normale Supérieure de Lyon, Dec. 2004.

[16] ———, *Could any graph be turned into a small-world?*, Theoretical Computer Science, 355 (2006), pp. 96–103.

[17] U. Feige, M. Hajiaghayi, and J. R. Lee, *Improved approximation algorithms for minimum-weight vertex separators*, in 37th Annual ACM Symposium on Theory of Computing (STOC), ACM Press, May 2005.

[18] P. Flocchini and F. L. Luccio, *A new compact routing technique on series parallel networks*, Theory of Computing Systems, 36 (2003), pp. 137–157.

[19] P. Fraigniaud, *A new perspective on the small-world phenomenon: Greedy routing in tree-decomposed graphs*, in 13th Annual European Symposium on Algorithms (ESA), vol. 3669 of LNCS, Springer, Oct. 2005, pp. 791–802.

[20] P. Fraigniaud and C. Gavoille, *Routing in trees*, in 28th International Colloquium on Automata, Languages and Programming (ICALP), vol. 2076 of LNCS, Springer, July 2001, pp. 757–772.

[21] G. N. Frederickson and R. Janardan, *Designing networks with compact routing tables*, Algorithmica, 3 (1988), pp. 171–190.

[22] ———, *Efficient message routing in planar networks*, SIAM Journal on Computing, 18 (1989), pp. 843–857.

[23] B. Fu, *Theory and application of width bounded geometric separator*, in 23rd Annual Symposium on Theoretical Aspects of Computer Science (STACS), vol. 3884 of LNCS, Springer, Feb. 2006, pp. 277–288.

[24] C. Gavoille and M. Gengler, *Space-efficiency of routing schemes of stretch factor three*, Journal of Parallel and Distributed Computing, 61 (2001), pp. 679–687.

[25] C. Gavoille, D. Peleg, S. Pérennès, and R. Raz, *Distance labeling in graphs*, Journal of Algorithms, 53 (2004), pp. 85–112.

[26] J. R. Gilbert, J. P. Hutchinson, and R. E. Tarjan, *A separation theorem for graphs of bounded genus*, Journal of Algorithms, 5 (1984), pp. 391–407.

[27] M. Grohe, *Local tree-width, excluded minors, and approximation algorithms*, Combinatorica, 23 (2003), pp. 613–632.

[28] A. Gupta, A. Kumar, and R. Rastogi, *Traveling with a pez dispenser (or, routing issues in MPLS)*, SIAM Journal on Computing, 34 (2005), pp. 453–474.

[29] J. Kleinberg, *The small-world phenomenon: An algorithmic perspective*, in 32nd Annual ACM Symposium on Theory of Computing (STOC), ACM Press, May 2000, pp. 163–170.

[30] ———, *Complex networks and decentralized search algorithms*, in International Congress of Mathematicians (ICM), ACM Press, Aug. 2006. To appear.

[31] T. Kloks, *Treewidth: Computations and Approximations*, vol. 842 of LNCS, Springer-Verlag, June 1994.

[32] K. A. Laing, *Brief announcement: Name-independent compact routing in trees*, in 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, July 2004, pp. 382–382.

[33] R. J. Lipton and R. E. Tarjan, *A separator theorem for planar graphs*, SIAM Journal on Applied Mathematics, 36 (1979), pp. 177–189.

[34] B. Mohar and C. Thomassen, *Graphs on Surfaces*, The Johns Hopkins university Press, 2001.

[35] S. A. Plotkin, S. B. Rao, and W. D. Smith, *Shallow excluded minors and improved graph decomposition*, in 5th Symposium on Discrete Algorithms (SODA), ACM-SIAM, 1994, pp. 462–470.

[36] B. Reed and D. R. Wood, *A linear time algorithm to find a separation in a graph with an excluded minor*, 2005. Full version of EuroComb '05.

[37] N. Robertson and P. D. Seymour, *Graph minors. V. Excluding a planar graph*, Journal of Combinatorial Theory, Series B, 41 (1986), pp. 92–114.

[38] ———, *Graph minors. XIII. The disjoint paths problem*, Journal of Combinatorial Theory, Series B, 63 (1995), pp. 65–110.

[39] ———, *Graph minors. XVI. Excluding a non-planar graph*, Journal of Combinatorial Theory, Series B, 89 (2003), pp. 43–76.

[40] N. Robertson, P. D. Seymour, and R. Thomas, *Quickly excluding a planar graph*, Journal of Combinatorial Theory, Series B, 62 (1994), pp. 323–348.

[41] A. Slivkins, *Distance estimation and object location via rings of neighbors*, in 24th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, July 2005, pp. 41–50. Appears earlier as Cornell CIS technical report TR2005-1977.

[42] K. Talwar, *Bypassing the embedding: Algorithms for low dimensional metrics*, in 36th Annual ACM Symposium on Theory of Computing (STOC), ACM Press, June 2004, pp. 281–290.

[43] A. Thomason, *The extremal function for complete minors*, Journal of Combinatorial Theory, Series B, 81 (2001), pp. 318–338.

[44] M. Thorup, *Compact oracles for reachability and approximate distances in planar digraphs*, Journal of the ACM, 51 (2004), pp. 993–1024.

[45] M. Thorup and U. Zwick, *Approximate distance oracles*, Journal of the ACM, 52 (2005), pp. 1–24.