$O(1)$ Steiner Point Removal in Series-Parallel Graphs

D. Ellis Hershkowitz
Carnegie Mellon University, Pittsburgh, PA, USA

Jason Li
Berkeley, CA, USA

Abstract

We study how to vertex-sparsify a graph while preserving both the graph's metric and structure. Specifically, we study the Steiner point removal (SPR) problem where we are given a weighted graph $G = (V, E, w)$ and terminal set $V' \subseteq V$ and must compute a weighted minor $G' = (V', E', w')$ of G which approximates G's metric on V'. A major open question in the area of metric embeddings is the existence of $O(1)$ multiplicative distortion SPR solutions for every (non-trivial) minor-closed family of graphs. To this end prior work has studied SPR on trees, cactus and outerplanar graphs and showed that in these graphs such a minor exists with $O(1)$ distortion.

We give $O(1)$ distortion SPR solutions for series-parallel graphs, extending the frontier of this line of work. The main engine of our approach is a new metric decomposition for series-parallel graphs which we call a hammock decomposition. Roughly, a hammock decomposition is a forest-like structure that preserves certain critical parts of the metric induced by a series-parallel graph.

2012 ACM Subject Classification Theory of computation → Sparsification and spanners; Theory of computation → Shortest paths; Theory of computation → Approximation algorithms analysis

Keywords and phrases Metric embeddings, graph algorithms, vertex sparsification

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.66

Related Version *Full Version*: `https://arxiv.org/pdf/2104.00750.pdf`

Funding Supported in part by NSF grants CCF-1527110, CCF-1618280, CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808, a Sloan Research Fellowship, funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC grant agreement 949272) and Swiss National Foundation (project grant 200021-184735).
D. Ellis Hershkowitz: Supported by the Air Force Office of Scientific Research under award number FA9550-20-1-0080.

1 Introduction

Graph sparsification and metric embeddings aim to produce compact representations of graphs that approximately preserve desirable properties of the input graph. For instance, a great deal of work has focused on how, given some input graph G, we can produce a simpler graph G' whose metric is a good proxy for G's metric; see, for example, work on tree embeddings [4, 16], distance oracles [36, 35] and graph spanners [2, 1] among many other lines of work. Simple representations of graph metrics enable faster and more space efficient algorithms, especially when the input graph is very large. For this reason these techniques are the foundation of many modern algorithms for massive graphs.

Compact representations of graphs are particularly interesting when we assume that G is a member of a minor-closed graph family such as tree, cactus, series-parallel or planar graphs.[1] As many algorithmic problems are significantly easier on such families – see e.g. [26, 3, 37] – it is desirable that G' is not only a simple approximation of G's metric but that it also belongs to the same family as G.

Steiner point removal (SPR) formalizes the problem of producing a simple G' in the same graph family as G that preserves G's metric. In SPR we are given a weighted graph $G = (V, E, w)$ and a terminal set $V' \subseteq V$ where $V \setminus V'$ are called "Steiner points." We must return a weighted graph $G' = (V', E', w')$ where:

1. G' is a minor of G;
2. $d_G(u, v) \leq d_{G'}(u, v) \leq \alpha \cdot d_G(u, v)$ for every $u, v \in V'$;

and our aim is to minimize the multiplicative distortion α. We refer to a G' with distortion α as an α-SPR solution. In the above d_G and $d_{G'}$ give the distances in G and G' respectively.

If we only required that G' satisfies the second condition then we could always achieve $\alpha = 1$ by letting G' be the complete graph on V' where $w'(\{u, v\}) = d_G(u, v)$ for every $u, v \in V'$. However, such a G' forfeits any nice structure that G may have exhibited. Thus, the first condition ensures that if G belongs to a minor-closed family then so does G'. The second condition ensures that G''s metric is a good proxy for G's metric. G' is simpler than G since it is a graph only on V' while G' is a proxy for G's metric by approximately preserving distances on V'.

As [22] observed, even for the simple case of trees we must have $\alpha > 1$. For example, consider the star graph with unit weight edges where V' consists of the leaves of the star. Any tree $G' = (V', E', w')$ has at least two vertices u and v whose connecting path consists of at least two edges. On the other hand, the length of any edge in G' is at least 2 and so $d_{G'}(u, v) \geq 4$. Since $d_G(u, v) = 2$ it follows that $\alpha \geq 2$. While this simple example rules out the possibility of 1-SPR solutions on trees, it leaves open the possibility of small distortion solutions for minor-closed families.

In this vein several works have posed the existence of $O(1)$-SPR solutions for minor-closed families as an open question: see, for example, [5, 19, 8, 30, 11] among other works. A line of work (summarized in Figure 1) has been steadily making progress on this question for the past two decades. [22] showed that trees (i.e. connected K_3-minor-free graphs) admit 8-SPR solutions.[20] recently gave a simpler proof of this result. [8] proved this was tight by showing that $\alpha \geq 8$ for trees which remains the best known lower bound for K_h-minor-free graphs. In an exciting recent work, [19] reduced $O(1)$-SPR in K_h-minor-free graphs to computing "$O(1)$ scattering partitions" and showed how to compute these partitions for several graph classes, including cactus graphs (i.e. all connected F-minor-free graphs where F is K_4 missing one edge). Lastly, a work of [5] generalizes these results by showing that outerplanar graphs (i.e. graphs which are both K_4 and $K_{2,3}$-minor-free) have $\alpha = O(1)$ solutions.

[1] A graph G' is a minor of a graph G if G' can be attained (up to isomorphism) from G by edge contractions as well as vertex and edge deletions. A graph is F-minor-free if it does not have F as a minor. A family of graphs $\mathcal{G}$ is said to be minor-closed if for any $G \in \mathcal{G}$ if G' is a minor of G then $G' \in \mathcal{G}$. A seminal work of Robertson and Seymour [34] demonstrated that every minor-closed family of graphs is fully characterized by a finite collection of "forbidden" minors. In particular, if $\mathcal{G}$ is a minor-closed family then there exists a finite collection of graphs $\mathcal{M}$ where $G \in \mathcal{G}$ iff G does not have any graph in $\mathcal{M}$ as a minor. Here and throughout this work we will use "minor-closed" to refer to all non-trivial minor-closed families of graphs; in particular, we exclude the family of all graphs which is minor-closed but trivially so.

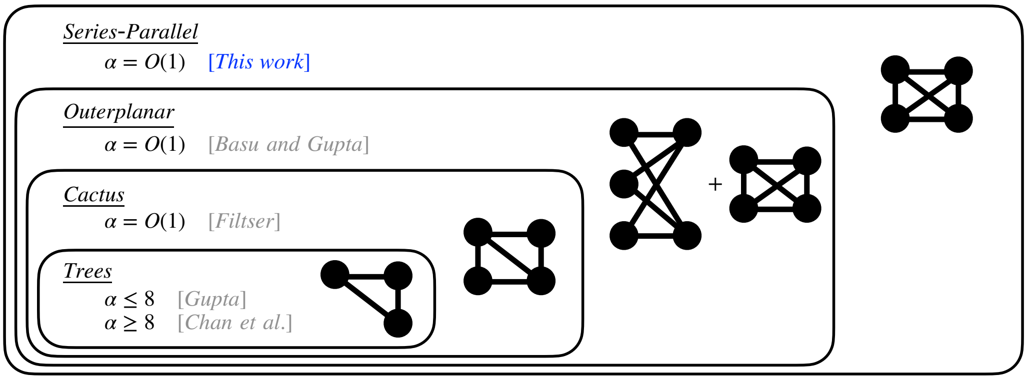

Figure 1 A summary of the SPR distortion for (connected) K_h-minor-free graphs achieved in prior work and our own. Graph classes illustrated according to containment. We also give the forbidden minors for each graph family.

1.1 Our Contributions

In this work, we advance the state-of-the-art for Steiner point removal in minor-closed graph families. We show that series-parallel graphs (i.e. graphs which are K_4-minor-free) have $O(1)$-SPR solutions. The following theorem summarizes the main result of our paper.

▶ **Theorem 1.** *Every series-parallel graph $G = (V, E, w)$ with terminal set $V' \subseteq V$ has a weighted minor $G' = (V', E', w')$ such that for any $u, v \in V'$ we have*

$$d_G(u, v) \leq d_{G'}(u, v) \leq O(1) \cdot d_G(u, v).$$

Moreover, G' is poly-time computable by a deterministic algorithm.

Series-parallel graphs are a strict superset of all of the aforementioned graph classes for which $O(1)$-SPR solutions were previously known; again, see Figure 1. Series-parallel graphs are one of the most well-studied graph classes in metric embeddings and serve as a frequent test bed for making progress on long-standing open questions. For example, series-parallel graphs are one of the few graph classes for which the well-known GNRS conjecture in metric embeddings has been successfully proven [25]. For further examples see, among many other works, those of [23, 7] and [13].

Relation to Prior Results

From a metric-embeddings perspective, series-parallel graphs are significantly more complex than outerplanar graphs (the largest minor-free graph class for which $O(1)$-distortion Steiner point removal was known prior to our work). For example, [24] showed that outerplanar graphs can be embedded into “dominating tree metrics” with constant distortion but that such an embedding for series-parallel graphs incurs $\Omega(\log n)$ distortion. Likewise, outerplanar graphs embed isometrically into l_1 which is known to not be possible for series-parallel graphs; see [33] and [9] for details. Thus, the metrics induced by series-parallel graphs often behave very differently and in less well-structured ways than those induced by outerplanar graphs.

Furthermore, the techniques on which we rely are quite different than those of [5] for the outerplanar case. At least two aspects of these techniques may be of independent interest. We defer a more thorough overview of our techniques to Section 4 but briefly highlight these two points now.

A New Approach for Steiner Point Removal

First, much of our approach generalizes to any K_h-minor-free graph so our approach seems like a promising avenue for future work on $O(1)$-SPR in minor-closed families. Specifically, we prove our result by beginning with the "chops" used by [29] to build low diameter decompositions for K_h-minor-free graphs. For input $\Delta > 0$ and root vertex r, these chops consist of deleting any edge which for some $i \in \mathbb{Z}$ has endpoints at distance $i\Delta$ and $i\Delta + 1$ from r; removing such edges partitions the input graph into width Δ "annuli." We begin with these chops but then slightly perturb them to respect the shortest path structure of the graph, resulting in what we call $O(1)$-scattering chops. We argue that the result of repeating such scattering chops is a scattering partition which by the results of [19] can be used to construct an $O(1)$-SPR solution.

The key to this strategy is arguing that series-parallel graphs admit a certain structure – which we call a hammock decomposition – that enables one to perform these perturbations in a principled way. If one could demonstrate a similar structure for K_h-minor-free graphs or otherwise demonstrate the existence of $O(1)$-scattering chops for such graphs, then the techniques laid out in our work would immediately give $O(1)$-SPR solutions for all K_h-minor-free graphs.

New Geometric Structure for Series-Parallel Graphs

Second, our hammock decompositions are a new metric decomposition for series-parallel graphs which may be interesting in their own right. We give significantly more detail in the full version of our work but briefly summarize our decomposition for now. We show that for any fixed BFS tree T_{BFS} there is a forest-like subgraph which contains all shortest paths between cross edges of T_{BFS}.[2] Specifically, the "nodes" of this forest are not vertices but highly structured subgraphs of the input series-parallel graph which we call hammock graphs. A hammock graph consists of two subtrees of the BFS tree and the cross edges between them.

Our hammock decompositions stand in contrast to the fact that the usual way in which one embeds a graph into a tree – by way of dominating tree metrics – are known to incur distortion $\Omega(\log n)$ in series-parallel graphs [25]. Furthermore, our decomposition can be seen as a metric-strengthening of the classic nested ear decompositions for series-parallel graphs of [28] and [15]. In general, a nested ear decomposition need not reflect the input metric. However, not only can one almost immediately recover a nested ear decomposition from a hammock decomposition, but the output nested ear decomposition interacts with the graph's metric in a highly structured way (see the full version of our work).

Open Questions Resolved

Lastly, we note that, in addition to making progress on the existence of $O(1)$-SPR solutions for every minor-closed family, our work also settles several open questions. The existence of $O(1)$-SPR solutions for series-parallel graphs was stated as an open question by both [5] and [8]; our result answers this question in the affirmative. Furthermore, [20] posed the existence of $O(1)$ scattering partitions for outerplanar and series-parallel graphs as an open question; we prove our main result by showing that series-parallel graphs admit $O(1)$ scattering partitions, settling both of these questions.

---

[2] Here and throughout this work a cross edge is an edge that is in the input graph but not in T_{BFS}.

2 Related Work

We briefly review additional related work.

Since the introduction of SPR by [22], a variety of works have studied the bounds achievable for well-behaved families of graphs for several very similar problems. [30] studied a problem like SPR but where distances in G must be exactly preserved by G' and the number of Steiner vertices – that is, vertices not in V' – must be made as small as possible; this work showed that while $O(k^4)$ Steiner vertices suffice (where $k = |V'|$) for general graphs, better bounds are possible for well-behaved families of graphs. More generally, [11] studied how to trade off between the number of terminals and distortion of G', notably showing $(1+\epsilon)$ distortion is possible in planar graphs with $\tilde{O}(k^2/\epsilon^2)$ Steiner vertices. [14] showed that in minor-closed graphs distances can be preserved up to $O(1)$ multiplicative distortion in expectation by a distribution over minors as opposed to preserving distances deterministically with a single minor as in SPR.

A variety of recent works have also studied how to find minors which preserve properties of G other than G's metric. [14] studied a flow/cut version of SPR where the goal is for G' to be a minor of G just on the specified terminals while preserving the congestion of multicommodity flows between terminals: this work showed that a convex combination of planar graphs can preserve congestion on V' up to a constant while for general graphs a convex combination of trees preserves congestion up to an $O(\log k)$. Similarly, [31] studied how to find minimum-size planar graphs which preserve terminal cuts. [21] studied how to find a minor of a directed graph with as few Steiner vertices and which preserves the reachability relationships between all k terminals, showing that $O(k^3)$ vertices suffices for general graphs but $O(\log k \cdot k^2)$ vertices suffices for planar graphs.

There has been considerable effort in the past few years on developing good SPR solutions for general graphs. [27] gave $O(\log^5 k)$-SPR solutions for general graphs. This was improved by [10] who gave $O(\log^2 k)$-SPR solutions which was, in turn, improved by [18] and [17] who gave $O(\log k)$-SPR solutions for general graphs. We also note that [19] also achieved similar results by way of scattering partitions, albeit with a worse poly-log factor as well as the first $O(1)$-SPR solutions for bounded pathwidth graphs.

3 Preliminaries

Before giving an overview of our approach we summarize the characterization of series-parallel graphs we use throughout this work as well as the scattering partition framework of [19] on which we build.

3.1 Characterizations of Series-Parallel Graphs

There are some minor inconsistencies in the literature regarding what is considered a series-parallel graph and so we clarify which notion of series-parallel we use throughout this paper. Some works – e.g. [15] – take series parallel graphs to be those which can be computed by iterating parallel and series compositions of graphs. Call these series-parallel A graphs.[3]

[3] The following is a definition of series-parallel A graphs due to [15]. A graph is two-terminal if it has a distinct source s and sink t. Let G and H be two two-terminal graphs with sources s and s' and sinks t and t'. Then the series composition of G and H is the graph resulting from identifying t and s' as the same vertex. The parallel composition of G and H is the graph resulting from identifying s and s' as the same vertex and t and t' as the same vertex. A two-terminal series-parallel graph is a two-terminal graph which is either a single edge or the graph resulting from the series or parallel composition of two two-terminal series-parallel graphs. A graph is series-parallel A if it has some pair of vertices with respect to which it is two-terminal series-parallel.

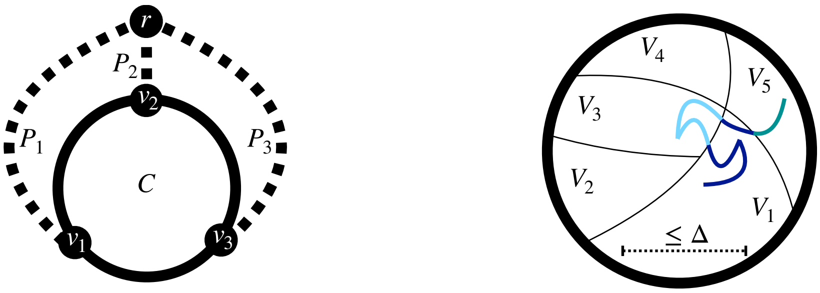

(a) A clawed cycle.

(b) A scattering partition.

Figure 2 In (a) we illustrate a clawed cycle where the cycle C is given in solid black and each path is given in dotted black. In (b) we illustrate a scattering partition with $\tau = 3$ and how one path P of length at most Δ is incident to at most three parts where we color the subpaths of P according to the incident part.

Strictly speaking, series-parallel A graphs are not even minor-closed as they are not closed under edge or vertex deletion. Other works – e.g. [19] – take series-parallel graphs to be graphs whose biconnected[4] components are each series-parallel A graphs; call these series-parallel B graphs. Series-parallel B graphs clearly contain series-parallel A graphs and, moreover, are minor-closed. For the rest of this work we will use the more expansive series-parallel B notion; henceforth we use "series-parallel" to mean series-parallel B.

It is well-known that a graph is K_4-minor-free iff it is series-parallel [6]. Similarly a graph has treewidth at most 2 iff it is series-parallel [6]. In this work we will use an alternate definition in terms of "clawed cycles" which we illustrate in Figure 2a.[5]

▶ **Definition 2** (Clawed Cycle). *A clawed cycle is a graph consisting of a root r, a cycle C and three paths P_1 P_2 and P_3 from r to vertices $v_1, v_2, v_3 \in C$ where $v_1 \neq v_2 \neq v_3$*

The fact that series-parallel graphs are exactly those that do not have any clawed cycles as a subgraph was proven by [12]; we give a proof for completeness.

▶ **Lemma 3** ([12]). *A graph G is series-parallel iff it does not contain a clawed cycle as a subgraph.*

Proof. K_4 is itself a clawed cycle and so a graph with no clawed cycle subgraphs is K_4-minor-free and therefore series-parallel. If a graph contains a clawed cycle then we can construct a K_4 minor by arbitrarily contracting the graph into v_1, v_2, v_3 and r, as defined in Definition 2. ◀

3.2 Scattering Partitions

Our result will be based on a new graph partition introduced by [19], the scattering partition. Roughly speaking, a scattering partition of a graph is a low-diameter partition which respects the shortest path structure of the graph; see Figure 2b.[6]

[4] A connected component C is biconnected if C remains connected even after the deletion of any one vertex in C.
[5] We note that clawed cycles are also called "embedded Wheatstone bridge."
[6] We drop one of the parameters of the definition of [19] as it will not be necessary for our purposes.

▶ **Definition 4** (Scattering Partition). *Given a weighted graph $G = (V, E, w)$, a partition $\mathcal{P} = \{V_i\}_i$ of V is a (τ, Δ) scattering partition if:*

1. ***Connected:*** *Each $V_i \in \mathcal{P}$ is connected;*
2. ***Low Weak Diameter:*** *For each $V_i \in \mathcal{P}$ and $u, v \in V_i$ we have $d_G(u, v) \leq \Delta$;*
3. ***Scattering:*** *Every shortest path P in G of length at most Δ satisfies $|\{V_i : V_i \cap P \neq \emptyset\}| \leq \tau$.*

[19] extended these partitions to the notion of a scatterable graph.

▶ **Definition 5** (Scatterable Graph). *A weighted graph $G = (V, E, w)$ is τ-scatterable if it has a (τ, Δ)-scattering partition for every $\Delta \geq 0$.*

We will say that G is deterministic poly-time τ-scatterable if for every $\Delta \geq 0$ a (τ, Δ)-scattering partition is computable in deterministic poly-time.

As a concrete example of a τ-scatterable graph and as observed by [19] notice that all trees are $O(1)$-scatterable. In particular, suppose we are given a tree and a $\Delta > 0$. If we fix a root vertex r and then delete any edge which for some $i \in \mathbb{Z}$ has endpoints at distance $\frac{i \cdot \Delta}{2}$ and $\frac{i \cdot \Delta}{2} + 1$ from r this breaks the input tree into connected components. Each component has diameter at most Δ by construction. Furthermore, it is easy to see that any path of length at most Δ is incident to a constant number of these components and so these components indeed form a scattering partition with $\tau = O(1)$. This construction is essentially a single chop of the aforementioned KPR strategy. However, while a KPR chop can be used to construct scattering partitions on trees, as we will see shortly, KPR chops on series-parallel graphs do not, in general, result in scattering partitions.

Lastly, the main result of [19] is that solving SPR reduces to showing that every induced subgraph is scatterable. In the following $G[A]$ is the subgraph of G induced by the vertex set A.

▶ **Theorem 6** ([19]). *A weighted graph $G = (V, E, w)$ with terminal set $V' \subseteq V$ has an $O(\tau^3)$-SPR solution if $G[A]$ is τ-scatterable for every $A \subseteq V$. Furthermore, if $G[A]$ is deterministic poly-time scatterable for every $A \subseteq V$ then the $O(\tau^3)$-SPR solution is computable in deterministic poly-time.*

4 Intuition and Overview of Techniques

We now give intuition and a high-level overview of our techniques. As discussed in the previous section, solving SPR with $O(1)$ distortion for any fixed graph reduces to showing that the subgraph induced by every subset of vertices is $O(1)$-scatterable. Moreover, since every subgraph of a K_h-minor-free graph is itself a K_h-minor-free graph, it follows that in order to solve SPR on any fixed K_h-minor-free graph, it suffices to argue that every K_h-minor-free graph is $O(1)$-scatterable.

Thus, the fact that we dedicate the rest of this document to showing is as follows.

▶ **Theorem 7.** *Every series-parallel graph G is deterministic, poly-time $O(1)$-scatterable.*

Combining this with Theorem 6 immediately implies Theorem 1.

4.1 General Approach

Given a series-parallel graph G and some $\Delta \geq 1$, our goal is to compute an $(O(1), \Delta)$-scattering partition for G. Such a partition has two non-trivial properties to satisfy: (1) each constituent part must have weak diameter at most Δ and (2) each shortest path of length at most Δ must be in at most $O(1)$ parts (a property we will call "scattering").

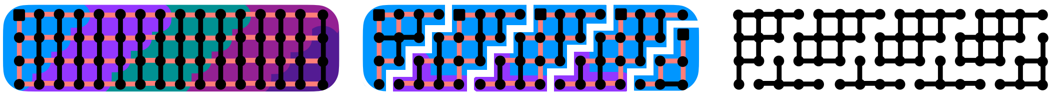

(a) First Δ-chop. **(b)** Second Δ-chop. **(c)** Resulting connected components.

Figure 3 Two levels of Δ-chops on the grid graph for $\Delta = 3$. We give the edges of the BFS trees we use in pink; roots of these trees are given as squares. Background colors give the annuli of nodes.

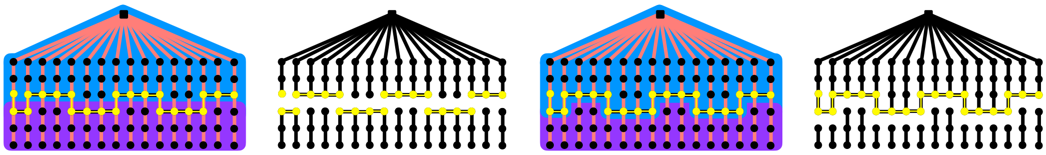

(a) A Δ-chop. **(b)** Components not scattering. **(c)** A perturbed Δ-chop. **(d)** Components scattering.

Figure 4 An example (of an outerplanar graph) where a Δ-chop does not produce a scattering partition but how perturbing said chop does. Here, we imagine that the root is at the top of the graph and each edge incident to the root has length $\Delta - 3$. We highlight the path P that either ends up in many or one connected component depending on whether we perturb our Δ-chop in yellow.

A well-known technique of [29] – henceforth "KPR" – has proven useful in finding so-called low diameter decompositions for K_h-minor-free graphs and so one might reasonably expect these techniques to prove useful for finding scattering partitions. Specifically, KPR shows that computing low diameter decompositions in a K_h-minor-free graph can be accomplished by $O(h)$ levels of recursive "Δ-chops". Fix a root r and a BFS tree T_{BFS} rooted at r. Then, a Δ-chop consists of the deletion of every edge with one vertex at depth $i \cdot \Delta$ and another vertex at depth $i \cdot \Delta + 1$ for every $i \in \mathbb{Z}_{\geq 1}$; that is, it consists of cutting edges between each pair of adjacent Δ-width annuli. KPR proved that if one performs a Δ-chop and then recurses on each of the resulting connected component then after $O(h)$ levels of recursive depth in a K_h-minor free graph the resulting components all have diameter at most $O(\Delta)$. We illustrate KPR on the grid graph in Figure 4.

Thus, we could simply apply Δ-chops $O(h)$ times to satisfy our diameter constraints (up to constants) and hope that the resulting partition is also scattering. Unfortunately, it is quite easy to see that (even after just one Δ-chop!) a path of length at most Δ can end up in arbitrarily many parts of the resulting partition. For example, the highlighted shortest path in Figures 4a and 4b repeatedly moves back and forth between two annuli and ends up in arbitrarily many parts after a single Δ-chop. Nonetheless, this example is suggestive of the basic approach of our work. In particular, if we merely perturbed our first Δ-chop to cut "around" said path as in Figures 4c and 4d then we could ensure that this path ends up in a small number of partitions.

More generally, the approach we take in this work is to start with the KPR chops but then slightly perturb these chops so that they do not cut *any* shortest path of length at most Δ more than $O(1)$ times. That is, all but $O(1)$ edges of any such path will have both vertices in the same (perturbed) annulus. We then repeat this recursively on each of the resulting connected components to a constant recursion depth. Since each subpath of a shortest path of length at most Δ is itself a shortest path with length at most Δ, we know that each such shortest path is broken into a constantly-many-more shortest paths at each level of recursion. Moreover, since we recurse a constant number of times, each path ends up in a constant number of components.

Implementing this strategy requires meeting two challenges. First, it is not clear that the components resulting from KPR still have low diameter if we allow ourselves to perturb our chops. Second, it is not clear how to perturb a chop so that it works *simultaneously* for every shortest path. Solving the first challenge will be somewhat straightforward while solving the second will be significantly more involved. In particular, what makes the second challenge difficult is that we cannot, in general, perturb a chop on the basis of one violated shortest path as in the previous example; doing so might cause other paths to be cut too many times which will then require additional, possibly conflicting, perturbations and so on. Rather, we must somehow perturb our chops in a way that takes every shortest path into account all at once.

4.2 Scattering Chops

The easier issue to solve will be how to ensure that our components have low diameter even if we perturb our chops. Here, by closely tracking various constants through a known analysis of KPR we show that the components resulting from KPR with (boundedly) perturbed cuts are still low diameter.

We summarize this fact and the above discussion with the idea of a scattering chop. A (τ, Δ)-scattering chop consists of cutting all edges at *about* every Δ levels in the BFS tree in such a way that no shortest path of length at most Δ is cut more than τ times. Our analysis shows that if all K_h-minor-free graphs admit $(O(1), \Delta)$-scattering chops for every Δ then they are also $O(1)$-scatterable and therefore also admit $O(1)$-SPR solutions; this holds even for $h > 4$.

4.3 Hammock Decompositions and How to Use Them

The more challenging issue we must overcome is how to perturb our chops so that every shortest path of length at most Δ is only cut $O(1)$ times. Moreover, we must do so in a way that does not perturb our boundaries by too much so as to meet the requirements of a scattering chop. We solve this issue with our new metric decomposition for series-parallel graphs, the hammock decomposition.

Consider a shortest path P of length at most Δ. Such a path can be partitioned into a (possibly empty) prefix consisting of only edges in T_{BFS}, a middle portion whose first and last edges are cross edges of T_{BFS} and a (possibly empty) suffix which also only has edges in T_{BFS}. Thus, if we want to compute a scattering chop, it suffices to guarantee that any shortest path of length at most Δ which is either fully contained in T_{BFS} or which is between two cross edges of T_{BFS} is only cut $O(1)$ times by our chop; call the former a BFS path and the latter a cross edge path.

Next, notice that all BFS paths are only cut $O(1)$ times by our initial KPR chops. Specifically, each BFS path can be partitioned into a subpath which goes "up" in the BFS tree and a subpath which goes "down" in the BFS tree. As our initial KPR chops are Δ apart and each such subpath is of length at most Δ, each such subpath is cut at most $O(1)$ times. Thus provided our perturbations do not interfere *too much* with the initial structure of our KPR chops we should expect that our BFS paths will only be cut $O(1)$ times.

Thus, our goal will be to perturb our KPR chops to not cut any cross edge path more than $O(1)$ times while mostly preserving the initial structure of our KPR chops. Our hammock decompositions will allow us to do exactly this. They will have two key components.

The first part is a "forest of hammocks." Suppose for a moment that our input graph had a forest subgraph F that contained all cross edge paths of our graph which were also shortest paths. Then, it is not too hard to see how to use F to perturb our chops to be scattering for all cross edge paths. Specifically, for each tree T in our forest F we fix an arbitrary root and then process edges in a BFS order. Edges which we process will be marked or unmarked where initially all edges are unmarked. To process an edge $e = \{u, v\}$ we do the following. If e is marked or u and v both belong to the same annulus then we do nothing. Otherwise, e is unmarked and u is in some annulus A but v is in some other annulus A' (before any perturbation). We then propagate A an additional $\Theta(\Delta)$ deeper into T; that is if we imagine that v is the child of u in F then we move all descendants of u in F within $\Theta(\Delta)$ of u into A. We then mark all edges in T whose endpoints are descendants of u and within $\Theta(\Delta)$ of u. A simple amortized analysis shows that after performing these perturbations every cross edge path is cut $O(1)$ times: if we think of following a cross edge path from one endpoint to the other, then each time this path is cut there must be at least $\Omega(\Delta)$ many edges we get to traverse until the next time it is cut again.

Unfortunately, it is relatively easy to see that such an F may not exist in a series-parallel graph. The forest of hammocks component of our decompositions is a subgraph which will be "close enough" to such an F, thereby allowing us to perturb our chops similarly to the above strategy. As mentioned in the introduction, a hammock graph consists of two subtrees of a BFS tree and the cross edges between them. A forest of hammocks is a graph partitioned into hammocks where every cycle is fully contained in one of the constituent hammocks. While the above perturbation will guarantee that our cross edge paths are not cut too often, it is not clear that such a perturbation does not change the structure of our initial chops in a way that causes our BFS paths to be cut too many times.

The second part of our hammock decompositions is what we use to guarantee that our BFS paths are not cut too many times by preserving the structure of our initial KPR chops. Specifically, the forest structure of our hammocks will reflect the structure of T_{BFS}. In particular, we can naturally associate each hammock H_i with a single vertex, namely the LCA of any u and v where u is in one tree of H_i and v is in the other. Then, our forest of hammocks will satisfy the property that if hammock H_i is a "parent" of hammock H_j in our forest of hammocks then the LCA corresponding to H_i is an ancestor of the LCA corresponding to H_j in T_{BFS}; even stronger, the LCA of H_j will be contained in H_i. Roughly, the fact that our forest of hammocks mimics the structure of T_{BFS} in this way will allow us to argue that the above perturbation does not alter the initial structure of our KPR chops too much, thereby ensuring that BFS paths are not cut too many times.

The computation of our hammock decompositions constitutes the bulk of our technical work but is somewhat involved. The basic idea is as follows. We will partition all cross edges into equivalence classes where each cross edge in an equivalence class shares an LCA in T_{BFS} (though there may be multiple, distinct equivalence classes with the same LCA). Each such equivalence class will eventually correspond to one hammock in our forest of hammocks. To compute our forest of hammocks we first connect up all cross edges in the same equivalence class. Next we connect our equivalence classes to one another by cross edge paths which run between them. We then extend our hammocks along paths towards their LCAs to ensure the above-mentioned LCA properties. Finally, we add so far unassigned subtrees of T_{BFS} to our hammocks. We will argue that when this process fails it shows the existence of a K_4-minor and, in particular, a clawed cycle.

5 Notation and Conventions

Before proceeding to our formal results we specify the notation we use throughout this work as well as some of the assumptions we make on our input series-parallel graph without loss of generality (WLOG).

Graphs. For a weighted graph $G = (V, E, w)$, we let $V(G) = V$ and $E(G) = E$ give the vertex set and edge sets of G respectively. We will sometimes abuse notation and use G to stand for V or E when it is clear from context if we mean G's vertex or edge set. Our weight function on edges is $w : E \to \mathbb{Z}_{\geq 1}$. Given graphs G and H, we will use the notation $H \subseteq G$ to indicate that H is a subgraph of G. The weak diameter of a subgraph H is $\max_{u,v \in V(H)} d_G(u, v)$.

Assumption of Unique Shortest Paths and Unit Weights. We will assume throughout this work that in our input series-parallel graph for any vertices u and v the shortest path between u and v is unique and that $w(e) = 1$ for every e. It is easy to see that our algorithms extend to non-unique shortest paths and the non-unit weight edge cases by standard techniques. In particular, one can randomly perturb the initial weights of the input graph so as to guarantee the uniqueness of shortest paths. Similarly, one can expand each edge of weight $w(e)$ into a path of $w(e)$ edges while preserving series-parallelness and the metric on the nodes from the original graph which suffices for our purposes.

Induced Graphs and Edges. Given an edge set E and disjoint vertex sets V_1 and V_2, we let $E(V_1, V_2) := \{e = \{v_1, v_2\} \in E : v_1 \in V_1, v_2 \in V_2\}$ be all edges between V_1 and V_2. Given graph $G = (V, E)$ and a vertex set $U \subseteq V$, we let $G[U] = (U, E_U)$ be the "induced subgraph" of G where $\{u', v'\} \in E_U$ iff $\{u', v'\} \in E$. Given a collection of subgraphs $\mathcal{H} = \{H_i\}_i$ of a graph we call $G[\mathcal{H}] := (\bigcup_i V(H_i), \bigcup_i E(H_i))$ the induced subgraph of $\mathcal{H}$. Similarly, we will let $E(\mathcal{H}) := \bigcup_i E(H_i)$ give the edges of $\mathcal{H}$. We emphasize that it is not necessarily the case that $G[\mathcal{H}] = G[V(\mathcal{H})]$.

Paths. Given a path $P = (v_0, v_1, \ldots, v_k, v_{k+1})$ we will use $\text{internal}(P) := \{v_1, \ldots, v_k\}$ to refer to the internal vertices of P. We will say that a path P is between two vertex sets U and W if its first and last vertices are in U and W respectively and $\text{internal}(P) \cap U = \emptyset$ and $\text{internal}(P) \cap W = \emptyset$. We will sometimes abuse notation and use P and $E(P)$ interchangeably. We will also sometimes say such a path is "from" U to W interchangeably with a path is "between" U and W. We will use $P \oplus P'$ to refer to the concatenation of two paths which share an endpoint throughout this paper. For a tree T, we will let $T(u, v)$ stand for the unique path between u and v in T for $u, v \in V(T)$. We will sometimes assume that a path from a vertex set to another vertex set is directed in the natural way.

BFS Tree Notation. For much of this work we will fix a series-parallel graph $G = (V, E)$ along with a fixed but arbitrary root $r \in V$ and a fixed but arbitrary BFS tree T_{BFS} with respect to r. When we do so we will let $E_c := E \setminus E(T_{\text{BFS}})$ be all cross edges of T_{BFS}. For $u, v \in V$, if $u \in T_{\text{BFS}}(r, v) \setminus \{v\}$ then we say that u is an ancestor of v. In this case, we also say that v is a descendant of u. If u is an ancestor of v or v is an ancestor of u then we say that u and v are related; otherwise, we say that u and v are unrelated. For two vertices $u, v \in V$ we will use the notation $u \prec v$ to indicate that v is an ancestor of u and we will use the notation $u \preceq v$ to indicate that v is an ancestor of or equal to u. It is easy to verify that $\preceq$ induces a partial order. We let $T_{\text{BFS}}(v) := T_{\text{BFS}}[\{v\} \cup \{u \in V : u \text{ is a descendant of } v\}]$ be the subtree of T_{BFS} rooted at v. Given a connected subgraph $T \subseteq T_{\text{BFS}}$, we will let $\text{high}(T)$ be the vertex in $V(T)$ which is an ancestor of all vertices in $V(T)$. Given a path $P \subseteq T_{\text{BFS}}$ we will say that P is monotone if $\text{high}(P)$ is an ancestor of all vertices in P and there is some vertex $\text{low}(P)$ which is a descendant of all vertices in P. We let $h(v)$ give the height of a vertex in T_{BFS} (where we imagine that the nodes furthest from r are at height 0). We let $\text{LCA}(e)$ be the least common ancestor of u and v in T_{BFS} for each $e = \{u, v\} \in E$.

Miscellaneous. We will use $\sqcup$ for disjoint set union throughout this paper. That is $A \sqcup B$ is equal to $A \cup B$ but indicates that $A \cap B = \emptyset$.

6 Perturbing KPR and Scattering Chops

In this section we show that KPR still gives low diameter components even if its boundaries are perturbed and therefore somewhat "fuzzy." We then observe that this fact shows that "$O(1)$-scattering chops" imply the existence of $O(1)$-scattering partitions for K_h-minor-free graphs and therefore $O(1)$-SPR solutions.

6.1 Perturbing KPR

We will repeatedly take the connected components of annuli with "fuzzy" boundaries. We formalize this with the idea of a c-fuzzy Δ-chop; see Figure 5a for an illustration.

▶ **Definition 8** (c-Fuzzy Δ-Chop). *Let $G = (V, E, w)$ be a weighted graph with root r and fix $0 \le c < 1$ and $\Delta \ge 1$. Then a c-fuzzy Δ-chop is a partition of V into "fuzzy annuli" $\mathcal{A} = \{A_1, A_2, \ldots\}$ where for every i and $v \in A_i$ we have*

$$(i-1)\Delta - \frac{c\Delta}{2} \le d(r, v) < i \cdot \Delta + \frac{c\Delta}{2}.$$

As each fuzzy annulus in a fuzzy chop may contain many connected components we must be careful when specifying how recursive application of these chops break a graph into connected components; hence the following definitions. Given fuzzy annulus A_i, we let $\mathcal{C}_i$ be the connected components of A_i.

▶ **Definition 9** (Components Resulting from a c-Fuzzy Δ-Chop). *Let $G = (V, E, w)$ be a weighted graph and let $\mathcal{C}$ be a partition of V into connected components. Then we say that $\mathcal{C}$ results from one level of c-fuzzy Δ-chops if there is a c-fuzzy Δ-chop $\mathcal{A}$ with respect to some root $r \in V$ satisfying $\mathcal{C} = \bigcup_{i:A_i \in \mathcal{A}} \mathcal{C}_i$. Similarly, for $h \ge 2$ we say that $\mathcal{C}$ results from h-levels of c-fuzzy Δ-chops if there is some $\mathcal{C}'$ which results from one level of c-fuzzy Δ-chops and $\mathcal{C}$ is the union of the result of $h-1$ levels of c-fuzzy Δ-chops on each $C' \in \mathcal{C}'$.*

We will now claim that taking $h-1$ levels of fuzzy chops in a K_h-minor-free graph will result in a connected, low weak diameter partition. In particular, we show the following lemma, the main result of this section.

▶ **Lemma 10.** *Let Δ and h satisfy $2 \le h$, $\Delta \ge 1$ and fix constant $0 \le c < 1$. Suppose $\mathcal{C}$ is the result of $h-1$ levels of c-fuzzy Δ-chops in a K_h-minor-free weighted graph G. Then, the weak diameter of every $C \in \mathcal{C}$ is at most $O(h \cdot \Delta)$.*

For the rest of this section we identify the nodes of a minor of graph G with "supernodes." In particular, we will think of each of the vertices of the minor as corresponding to a disjoint, connected subset of vertices in G (a supernode) where the minor can be formed from G (up to isomorphism) by contracting the constituent nodes of each such supernodes.

Our proof will closely track a known analysis of KPR [32]. The sketch of this strategy is as follows. We will argue that if we fail to produce parts with low diameter then we have found K_h as a minor. Our proof will be by induction on the number of levels of fuzzy chops. Suppose $\mathcal{C}$ is produced by $h-1$ levels of fuzzy chops; in particular, suppose $\mathcal{C}$ is produced by taking some fuzzy chop to get $\mathcal{C}'$ and then taking $h-2$ levels of fuzzy chops on each $C' \in \mathcal{C}'$. Also assume that there is some $C \in \mathcal{C}$ which has *large* diameter. Then, C must result from taking $h-2$ levels of fuzzy chops on some $C' \in \mathcal{C}$ where C' lies in some fuzzy annulus A_i of G. By our inductive hypothesis it follows that C contains K_{h-1} as a minor. Our goal is to add one more supernode to this minor to get a K_h minor. We will do so by finding disjoint paths of length $O(\Delta)$ in the annulus above A_i from each of the K_{h-1} supernodes all of which converge on a single connected component. By adding these paths to their respective supernodes in the K_{h-1} minor and adding the connected component on which these paths converge to our collection of supernodes, we will end up with a K_h minor.

The main challenge in this strategy is to show how to find paths as above which are disjoint. We will do so by choosing these paths from a "representative" from each supernode where initially the representatives are $\Omega(h\Delta)$ far-apart and grow at most $O(\Delta)$ closer at each level of chops; since we do at most $O(h)$ levels of chops, the paths we choose will never intersect.

To formalize this strategy we must state a few definitions which will aid in arguing that these representatives are far apart.

▶ **Definition 11** (Δ-Dense). *Given sets $S, U \subseteq V$ we say S is Δ-dense in U if $d(u, S) \leq \Delta$ for every $u \in U$.*

▶ **Definition 12** (R-Represented). *A K_h minor is R-represented by set $S \subseteq V$ if each supernode $V_i \subseteq V$ of the minor in G contains a representative $v_i \in S \cap V_i$ and these representatives are pairwise at least R apart in G.*

Since V is clearly $(1+c)\Delta$-dense in V, we can set S to V and j to $h-1$ in the following lemma to get Lemma 10.

▶ **Lemma 13.** *Fix $0 \leq c < 1$ and $h > j \geq 0$. Let S be any set which is $(1+c)\Delta$-dense in V and suppose $\mathcal{C}$ is the result of j levels of c-fuzzy Δ-chops and some $C \in \mathcal{C}$ has weak diameter more than $22h\Delta$. Then there exists a K_{j+1} minor which is $8(h-j)\Delta$-represented by S.*

Proof. Our proof is by induction on j. The base case of $j = 0$ is trivial as K_1 is a minor of any graph with a supernode $+\infty$-represented by any single vertex in V.

Now consider the inductive step on graph $G = (V, E)$. Fix some set S which is $(1+c)\Delta$-dense in V and let $\mathcal{C}$ be the result of j levels of c-fuzzy chops using root r with some $C' \in \mathcal{C}$ of diameter more than $22h\Delta$. Suppose C' is in fuzzy annulus A_k and suppose that C' is the result of applying $j-1$ levels of c-fuzzy chops to some C which resulted from 1 level of c-fuzzy chops in G; note that C is a connected component of A_k and that C' is contained in C.

As an inductive hypothesis we suppose that any $j-1$ levels of c-fuzzy Δ-chops on any graph H which results in a cluster of weak diameter more than $22h\Delta$ demonstrates the existence of a K_j minor in H which is $8(h-j+1)\Delta$-represented by any set S' which is $(1+c)\Delta$-dense in $V(H)$. Here weak diameter is with respect to the distances induced by the original input graph.

Thus, by our inductive hypothesis we therefore know that C contains a K_j minor which is $8(h-j+1)\Delta$-represented by any $S' \subseteq V(C)$ which is $(1+c)\Delta$-dense in $V(C)$. In particular, we may let S' be the "upper boundary" of C; that is, we let S' be all vertices v in C such that the shortest path from v to r does not contain any vertices in C. Clearly the shortest path from any vertex in C to r intersects a node in S'; moreover, when restricted to C this shortest path has length at most $\Delta + c\Delta$ (since C is contained in A_k) which is to say that S' is $(1+c)\Delta$-dense in C. Thus, by our induction hypothesis there is a K_j minor in C which is $8(h-j+1)\Delta$-represented by S'. Let $V_1, \ldots, V_j$ be the nodes in the supernodes of the K_j minor.

We now describe how to extend the K_j minor to a K_{j+1} minor which is $8(h-j)\Delta$-represented by S. We may assume that $k \geq 9h+1$; otherwise the distance from every node in A_k to r would be at most $(9h+1)\Delta + \frac{c\Delta}{2} \leq (9h+\frac{3}{2})\Delta$ and so the weak diameter of C' would be at most $(18h+3)\Delta \leq 21h\Delta$, contradicting our assumption on C''s diameter. It follows that for every $v \in A_k$ we have

$$d(v,r) \geq (k-1)\Delta - \frac{c\Delta}{2} \geq 9h\Delta - \frac{c\Delta}{2} \geq 8h\Delta. \tag{1}$$

We first describe how we grow each supernode V_i from the K_j minor to a new supernode V_i'. Let v_i be the representative in S' for V_i. Consider the path which consists of following the shortest path from v_i to r for distance 2Δ and then continuing on to the nearest node in S; let v_i' be this nearest node; this path from v_i to v_i' has length at most $(3+c)\Delta$ since S is $(1+c)\Delta$-dense in $V(G)$. Let V_i' be the union of V_i with the vertices in this path. Since each of these paths is of length at most $(3+c)\Delta \leq 4\Delta$, it follows that each of these paths for each i must be disjoint since each v_i is at least $8(h-j+1)\Delta > 8\Delta$ apart. Further, every v_i' must also, therefore, be at least $8(h-j)\Delta$ apart. Therefore, we let these v_i' form the representatives in S for each of the V_i'.

We now describe how we construct the additional supernode, V_0, which we add to our minor to get a K_{j+1} minor. V_0 will "grow" from the root to S and each of the V_i'. In particular, let $u_i \in V_i'$ be the node in V_i' which is closest to r and let P_i be the shortest path from r to u_i, excluding u_i. Similarly, let v_0' be the node in S closest to r and let P_0 be the shortest path from r to v_0', including v_0'. Then, we let V_0 be $P_0 \cup P_1 \cup \ldots \cup P_j$ and we let v_0' be the representative for V_0 in S. We claim that for every $i \geq 1$ we have

$$d(P_0, V_i') \geq 8(h-j)\Delta. \tag{2}$$

In particular, notice that since S is $(1+c)\Delta$-dense in $V(G)$ we know that $d(r, v_0') \leq (1+c)\Delta \leq 2\Delta$ and since $d(v,r) \geq 8h\Delta$ for every $v \in A_k$ by Equation (1) and $d(v_i', A_k) \leq (3+c)\Delta \leq 4\Delta$, it follows that $d(P_0, V_i') \geq (8h-6)\Delta \geq 8(h-j)\Delta$. Consequently, $d(v_0', v_i') \geq 8(h-j)\Delta$ for every $i \geq 1$. Thus, our representatives of each supernode are appropriately far apart.

It remains to show that our supernodes indeed form a K_{j+1} minor; clearly by construction they are all pair-wise adjacent and so it remains only to show that they are all disjoint from one another. We already argued above that for $i, j \geq 1$ any V_i' and V_j' are disjoint so we need only argue that V_0' is disjoint from each V_i' for $i \geq 1$. P_0 must be disjoint from each V_i' for $i \geq 1$ by Equation (2) and so we need only verify that P_i is disjoint from V_j' for $i, j \geq 1$;

By construction if $i = j$ we know that P_i is disjoint from V_j' so we assume $i \neq j$ and that P_i intersects V_j' for the sake of contradiction. Notice that each P_i has length at most $k\Delta + \frac{c\Delta}{2} - 2\Delta = k\Delta + c\Delta - 2\Delta - \frac{c\Delta}{2} < (k-1)\Delta - \frac{c\Delta}{2}$ by how we construct V_i'. Thus, P_i must be disjoint from A_k. It follows that if P_i intersects V_j' then it must intersect $V_j' \setminus V_j$. However, since $d(v_i, v_j) \geq 8(h-j+1)\Delta \geq 16\Delta$ and the length of paths $V_i' \setminus V_i$ and $V_j' \setminus V_j$ are at most 4Δ we know that $d(V_i' \setminus V_i, V_j' \setminus V_j) \geq 8(h-j)\Delta \geq 8\Delta$. Thus, after intersecting $V_j' \setminus V_j$ and then continuing on to a vertex adjacent to $V_i' \setminus V_i$, we know P_i must travel at least 8Δ; since the vertices of P_i are monotonically further and further from r, and the vertex in $V_j' \setminus V_j$ that P_i intersects must be distance at least $(k-1)\Delta - \frac{c\Delta}{2} - 4\Delta \geq (k-5)\Delta$ from r, then the last vertex of P_i must be distance at least $(k+3)\Delta$ from r, meaning P_i must intersect annulus A_k, a contradiction. ◀

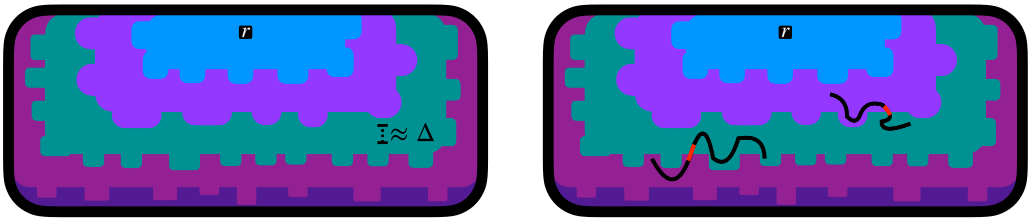

(a) A c-fuzzy Δ-chop. **(b)** Visualizing some paths.

Figure 5 A c-fuzzy Δ-chop that is 1-scattering. We draw each fuzzy annulus in a distinct color. In (b) we visualize some shortest paths of length at most Δ and highlight cut edges in red.

6.2 Scattering Chops

Using Lemma 10 we can reduce computing a good scattering partition and therefore computing a good SPR solution to finding what we call a scattering chop. The following definitions are somewhat analogous to Definition 4 and Definition 5. However, notice that the second definition is for a family of graphs (as opposed to a single graph as in Definition 5). We illustrate a τ-scattering chop in Figure 5.

▶ **Definition 14** (τ-Scattering Chop)**.** *Given a weighted graph $G = (V, E, w)$, let $\mathcal{A}$ be a c-fuzzy Δ-chop with respect to some root $r \in V$. $\mathcal{A}$ is a τ-scattering chop if each shortest path of length at most Δ has at most τ edges cut by $\mathcal{A}$ where we say that an edge is cut by $\mathcal{A}$ if it has endpoints in different fuzzy annuli of $\mathcal{A}$.*

▶ **Definition 15** (τ-Scatter-Choppable Graphs)**.** *A family of graphs $\mathcal{G}$ is τ-scatter-choppable if there exists a constant $0 \leq c < 1$ such that for any $G \in \mathcal{G}$ and $\Delta \geq 1$ there is some τ-scattering and c-fuzzy Δ-chop $\mathcal{A}$ with respect to some root.*

We will say that $\mathcal{G}$ is deterministic poly-time τ-scatter-choppable if the above chop $\mathcal{A}$ for each $G \in \mathcal{G}$ can be computed in deterministic poly-time.

Lastly, we conclude that to give an $O(1)$-scattering partition – and therefore to give an $O(1)$-SPR solution – for a K_h-minor-free graph family it suffices to show that such a family is $O(1)$-scatter choppable.

▶ **Lemma 16.** *Fix a constant $h \geq 2$ and let $\mathcal{G}_h$ be all K_h-minor-free graphs. Then, if $\mathcal{G}_h$ is τ-scatter-choppable then every $G \in \mathcal{G}_h$ is $O(\tau^{h-1})$-scatterable.*

Proof. The claim is almost immediate from Lemma 10 and the fact that all subpaths of a shortest path are themselves shortest paths.

In particular, first fix a sufficiently small constant c' to be chosen later. Then, consider a $G \in \mathcal{G}_h$ and fix a Δ. By assumption we know that G is τ-scatter-choppable and since each subgraph of G is in $\mathcal{G}_h$ so too is each subgraph of G. Thus, we may let $\mathcal{C}$ be the connected components resulting from $h-1$ levels of c-fuzzy and $(c'\Delta)$-chops which are τ-scattering.

We claim that for sufficiently small c' we have that $\mathcal{C}$ is a $\left(\frac{\tau^{h-1}}{c'}, \Delta\right)$-scattering partition. By Lemma 10 the diameter of each part in $\mathcal{C}$ is at most $O(c' \cdot h \cdot \Delta) \leq \Delta$ for sufficiently small c'. Next, consider a shortest path P of length at most Δ. We can partition the edges of P into at most $\frac{1}{c'}$ shortest paths $P_1, P_2, \ldots$, each of length at most $c' \cdot \Delta$. Thus, it suffices to show that each P_i satisfies $|\{C \in \mathcal{C} : P_i \cap C \neq \emptyset\}| \leq \tau^{h-1}$.

We argue by induction on the number of levels of chops that after $h' < h$ chops we have $|\{C \in \mathcal{C} : P_i \cap C \neq \emptyset\}| \leq \tau^{h'}$. Suppose we perform just one chop; i.e. $h' = 1$. Then, since our chops are τ-scattering we know that P will be cut at most τ times and so be incident to at most τ components of $\mathcal{C}$ as required. Next, suppose we perform $h' > 1$ levels of chops. Then our top-level chop will partition the vertices of P_i into at most τ components. By induction and the fact that each subpath of P_i is itself a shortest path of length at most $c'\Delta$, we know that the vertices of P_i in each such component are broken into at most $\tau^{h'-1}$ components and so P_i will be incident to at most $\tau^{h'}$ components as required. As we perform $h-1$ levels of chops, it follows that $\mathcal{C}$ is indeed a $\left(\frac{\tau^{h-1}}{c'}, \Delta\right)$-scattering partition. ◀

References

1 Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020.

2 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.

3 Sanjeev Arora, Michelangelo Grigni, David R Karger, Philip N Klein, and Andrzej Woloszyn. A polynomial-time approximation scheme for weighted planar graph tsp. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, volume 98, pages 33–41, 1998.

4 Yair Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Symposium on Foundations of Computer Science (FOCS)*, pages 184–193. IEEE, 1996.

5 Amitabh Basu and Anupam Gupta. Steiner point removal in graph metrics, 2008. Unpublished Manuscript, available from `https://www.ams.jhu.edu/~abasu9/papers/SPR.pdf`.

6 Hans L Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical computer science*, 209(1-2):1–45, 1998.

7 Bo Brinkman, Adriana Karagiozova, and James R Lee. Vertex cuts, random walks, and dimension reduction in series-parallel graphs. In *Annual ACM Symposium on Theory of Computing (STOC)*, pages 621–630, 2007.

8 T-H Hubert Chan, Donglin Xia, Goran Konjevod, and Andrea Richa. A tight lower bound for the steiner point removal problem on trees. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 70–81. Springer, 2006.

9 Chandra Chekuri, Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Embedding k-outerplanar graphs into l1. *SIAM Journal on Discrete Mathematics*, 20(1):119–136, 2006.

10 Yun Kuen Cheung. Steiner point removal—distant terminals don't (really) bother. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1353–1360. SIAM, 2018.

11 Yun Kuen Cheung, Gramoz Goranci, and Monika Henzinger. Graph minors for preserving terminal distances approximately-lower and upper bounds. In *International Colloquium on Automata, Languages and Programming (ICALP)*, 2016.

12 Richard J Duffin. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications*, 10(2):303–318, 1965.

13 Yuval Emek and David Peleg. A tight upper bound on the probabilistic embedding of series-parallel graphs. *SIAM Journal on Discrete Mathematics*, 23(4):1827–1841, 2010.

14 Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Racke, Inbal Talgam-Cohen, and Kunal Talwar. Vertex sparsifiers: New results from old techniques. *SIAM Journal on Computing*, 43(4):1239–1262, 2014.

15 David Eppstein. Parallel recognition of series-parallel graphs. *Information and Computation*, 98(1):41–55, 1992.

16 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004.

17 Arnold Filtser. Steiner point removal with distortion o (log k). In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1361–1373. SIAM, 2018.

18 Arnold Filtser. Steiner point removal with distortion o(log k), using the noisy-voronoi algorithm. *arXiv preprint*, 2018. `arXiv:1808.02800`.

19 Arnold Filtser. Scattering and sparse partitions, and their applications. In *International Colloquium on Automata, Languages and Programming (ICALP)*, 2020.

20 Arnold Filtser, Robert Krauthgamer, and Ohad Trabelsi. Relaxed voronoi: A simple framework for terminal-clustering problems. In *SIAM Symposium on Simplicity in Algorithms (SOSA)*, 2018.

21 Gramoz Goranci, Monika Henzinger, and Pan Peng. Improved guarantees for vertex sparsification in planar graphs. *SIAM Journal on Discrete Mathematics*, 34(1):130–162, 2020.

22 Anupam Gupta. Steiner points in tree metrics don't (really) help. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 682–690, 2001.

23 Anupam Gupta, Robert Krauthgamer, and James R Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Symposium on Foundations of Computer Science (FOCS)*, pages 534–543. IEEE, 2003.

24 Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and l/sub 1/-embeddings of graphs. In *Symposium on Foundations of Computer Science (FOCS)*, pages 399–408. IEEE, 1999.

25 Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and l1-embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.

26 Frank Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4(3):221–225, 1975.

27 Lior Kamma, Robert Krauthgamer, and Huy L Nguyen. Cutting corners cheaply, or how to remove steiner points. *SIAM Journal on Computing*, 44(4):975–995, 2015.

28 Samir Khuller. Ear decompositions. *SigAct News*, 20(1):128, 1989.

29 Philip Klein, Serge A Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Annual ACM Symposium on Theory of Computing (STOC)*, pages 682–690, 1993.

30 Robert Krauthgamer, Huy L Nguyen, and Tamar Zondiner. Preserving terminal distances using minors. *SIAM Journal on Discrete Mathematics*, 28(1):127–141, 2014.

31 Robert Krauthgamer and Havana Inbal Rika. Refined vertex sparsifiers of planar graphs. *SIAM Journal on Discrete Mathematics*, 34(1):101–129, 2020.

32 James R. Lee and Cyrus Rashtchian. A simpler proof of the kpr theorem. `https://tcsmath.wordpress.com/tag/klein-plotkin-rao/`, January 2012.

33 Haruko Okamura and Paul D Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B*, 31(1):75–81, 1981.

34 Neil Robertson and Paul D Seymour. Graph minors. xx. wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.

35 Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 261–272. Springer, 2005.

36 Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.

37 Thomas Victor Wimer. Linear algorithms on k-terminal graphs. *PhD Thesis*, 1987. AAI8803914.