REFINED VERTEX SPARSIFIERS OF PLANAR GRAPHS[*]

ROBERT KRAUTHGAMER[†] AND HAVANA (INBAL) RIKA[†]

Abstract. We study the following version of cut sparsification. Given a large edge-weighted network G with k terminal vertices, compress it into a smaller network H with the same terminals, such that every minimum terminal cut in H approximates the corresponding one in G, up to a factor $q \geq 1$ that is called the quality. (The case $q = 1$ is known also as a mimicking network.) We provide new insights about the structure of minimum terminal cuts, leading to new results for cut sparsifiers of *planar* graphs. Our first contribution identifies a subset of the minimum terminal cuts, which we call elementary, that generates all the others. Consequently, H is a cut sparsifier if and only if it preserves all the elementary terminal cuts (up to this factor q). Our second and main contribution is to refine the known bounds in terms of $\gamma = \gamma(G)$, which is defined as the minimum number of faces that are incident to all the terminals in a planar graph G. We prove that the number of elementary terminal cuts is $O((2k/\gamma)^{2\gamma})$ (compared to $O(2^k)$ terminal cuts) and furthermore obtain a mimicking network of size $O(\gamma 2^{2\gamma} k^4)$, which is near-optimal as a function of γ. Our third contribution is a duality between cut sparsification and distance sparsification for certain planar graphs, when the sparsifier H is required to be a minor of G. This duality connects problems that were previously studied separately, implying new results, new proofs of known results, and equivalences between open gaps.

Key words. algorithms, graph theory, combinatorics

AMS subject classifications. 68R10, 68R05, 05C85

DOI. 10.1137/17M1151225

1. Introduction. A very powerful paradigm when manipulating a huge graph G is to *compress* it, in the sense of transforming it into a small graph H (or alternatively, into a succinct data structure) that maintains certain features (quantities) of G, like distances, cuts, or flows. The basic idea is that once the compressed graph H is computed in a preprocessing step, further processing can be performed on H instead of on G, using less resources like running time and memory, or achieving better accuracy when the solution is approximate. This paradigm has lead to remarkable successes, such as faster running time for fundamental problems, and the introduction of important concepts, from spanners [36] to cut and spectral sparsifiers [3, 38]. In these examples, H is a subgraph of G with the *same* vertex set but *sparse* and is sometimes called an edge sparsifier. In contrast, we aim to reduce the number of vertices in G, using so-called vertex sparsifiers.

In the vertex-sparsification scenario, G has k designated vertices called *terminals*, and the goal is to construct a small graph H that contains these terminals and maintains some of their features inside G, like distances or cuts. Throughout, a *k-terminal network*, denoted $G = (V, E, T, c)$, is an undirected graph (V, E) with edge weights $c : E \to \mathbb{R}_+$ and terminals set $T \subset V$ of size $|T| = k$. As usual, a *cut* is a partition of the vertices, and its *cutset* is the set of edges that connect between different parts.

Interpreting the edge weights as capacities, the *cost* of a cut $(W, V \setminus W)$ is the total weight of the edges in the respective cutset.

We say that a cut $(W, V \setminus W)$ *separates* a terminals subset $S \subset T$ from $\bar{S} := T \setminus S$ (or in short that it is *S-separating*) if all of S is on one side of the cut and $\bar{S}$ on the other side, i.e., $W \cap T$ equals either S or $\bar{S}$. We denote by $\text{mincut}_G(S)$ the minimum cost of an S-separating cut in G, where by a consistent tie-breaking mechanism, such as edge-weights perturbation, we assume throughout that the minimum is attained by only one cut, which we call the *minimum terminal cut* (of S).

DEFINITION 1.1. *A network $H = (V_H, E_H, T, c_H)$ is a* cut sparsifier *of $G = (V, E, T, c)$ with quality $q \geq 1$ and size $s \geq k$ (or in short, a (q, s)-cut-sparsifier) if its size is $|V_H| \leq s$ and*

$$(1) \qquad \forall S \subset T, \qquad \text{mincut}_G(S) \leq \text{mincut}_H(S) \leq q \cdot \text{mincut}_G(S).$$

In words, (1) requires that every minimum terminal cut in H approximates the corresponding one in G. Throughout, we consider only $S \neq \emptyset, T$, although for brevity we will not write it explicitly.

Two special cases are particularly important for us. One is quality $q = 1$, or a $(1, s)$-cut-sparsifier, which is known in the literature as a *mimicking network* and was introduced by [23]. The second case is a cut sparsifier H that is furthermore a minor of G, and then we call it a *minor cut sparsifier*, and similarly for a *minor mimicking network*. In all our results, the sparsifier H is actually a minor of G, which can be important in some applications; for instance, if G is planar, then H admits planar-graph algorithms.

In known constructions of mimicking networks ($q = 1$), the sparsifier's size s highly depends on the number of constraints in (1) that are really needed. Naively, there are at most 2^k constraints, one for every minimum terminal cut (this can be slightly optimized, e.g., by symmetry of S and $\bar{S}$). This naive bound was used to design, for an arbitrary network G, a mimicking network whose size s is exponential in the number of constraints, namely $s \leq 2^{2^k}$ [23]. A slight improvement, that is still doubly exponential in k, was obtained by using the submodularity of cuts to reduce the number of constraints [26]. For a planar network G, the mimicking network size was improved to a polynomial in the number of constraints, namely $s \leq k^2 2^{2k}$ [29], and this bound is actually near-optimal, due to a very recent work showing that some planar graphs require $s = 2^{\Omega(k)}$ [25]. In this paper we explore the structure of minimum terminal cuts more deeply, by introducing technical ideas that are new and different from previous work like [29].

Our approach. We take a closer look at the mimicking network size s of planar graphs, aiming at bounds that are more sensitive to the given network G. For example, we would like to "interpolate" between the very special case of an outerplanar G, which admits a mimicking network of size $s = O(k)$ [8], and an arbitrary planar G for which $s \leq 2^{O(k)}$ is known and optimal [29, 25]. Our results employ a graph parameter $\gamma(G)$, defined next.

DEFINITION 1.2 (terminal face cover). *The* terminal face cover $\gamma = \gamma(G)$ *of a planar k-terminal network G with a given drawing[1] is the minimum number of faces that are incident to all the k terminals, and thus $1 \leq \gamma \leq k$.*

This graph parameter $\gamma(G)$ is well known to be important algorithmically. For example, it can be used to control the running time of algorithms for shortest-path

---

[1]We can let γ refer to the best drawing of G, and then our results might be nonalgorithmic.

TABLE 1
Known and new bounds for mimicking networks.

Graphs	Size	Minor	Reference
General	$2^{2^k} \approx 2^{\binom{(k-1)}{(k-1)/2}}$	no	[23, 26]
Planar	$O(k^2 2^{2k})$	yes	[29]
Planar	$O(k 2^{2k})$	yes	Corollary 3.2
Planar $\gamma = \gamma(G)$	$O(\gamma 2^{2\gamma} k^4)$	yes	Corollary 4.7
Planar $\gamma(G) = 1$	$O(k^4)$	yes	Corollary 4.3
Planar $\gamma(G) = 1$	$O(k^2)$	no	[20]
General	$2^{\Omega(k)}$	no	[29, 26] lower bound
Planar	$2^{\Omega(k)}$	no	[25] lower bound

problems [19, 11], for cut problems [10, 4], and for multicommodity flow problems [33]. For the complexity of computing an optimal/approximate face cover $\gamma(G)$, see [5, 19].

When $\gamma = 1$, all the terminals lie on the boundary of the same face, which we may assume to be the outerface. This special case was famously shown by Okamura and Seymour [35] to have a flow-cut gap of 1 (for multicommodity flows). Later work showed that for general gamma, the flow-cut gap is at most 3γ [30, 9], and after the current results were announced this bound was further improved to $O(\sqrt{\log \gamma})$ [27, 17].

1.1. Main results and techniques. We provide new bounds for mimicking networks of planar graphs. In particular, our main result refines the previous bound so that it depends exponentially on $\gamma(G)$ rather than on k. This yields much smaller mimicking networks in important cases; for instance, when $\gamma = O(1)$ we achieve size $s = \text{poly}(k)$. See Table 1 for a summary of known and new bounds. Technically, we develop two methods to decompose the minimum terminal cuts into "more basic" subsets of edges, and then represent the constraints in (1) using these subsets. This is equivalent to reducing the number of constraints and leads (as we hinted above) to a smaller sparsifier size s. A key difference between the methods is that the first one in effect restricts attention to a subset of the constraints in (1), while the second method uses alternative constraints.

Decomposition into elementary cutsets. Our first decomposition method identifies (in every graph G, even nonplanar) a subset of minimum terminal cuts that "generates" all the other ones, as follows. First, we call a cutset *elementary* if removing its edges disconnects the graph into exactly two connected components (Definition 2.1). We then show that every minimum terminal cut in G can be decomposed into a disjoint union of elementary ones (Theorem 2.5), and we use this to conclude that if all the elementary cutsets in G are well approximated by those in H, then H is a cut sparsifier of G (Corollary 2.6).

Combining this framework with prior work on planar sparsifier [29], we devise the following bound that depends on $\mathcal{T}_e(G)$, the set of elementary cutsets in G.

- Generic bound: Every planar graph G has a mimicking network of size $s = O(k) \cdot |\mathcal{T}_e(G)|^2$; see Theorem 3.1.

Trivially $|\mathcal{T}_e(G)| \leq 2^k$, and we immediately achieve $s = O(k 2^{2k})$ for all planar graphs (Corollary 3.2). This improves over the known bound [29] slightly (by factor k) and stems directly from the restriction to elementary cutsets (which are simple cycles in the planar-dual graph).

Using the same generic bound, we further obtain mimicking networks whose size is *polynomial* in k (but inevitably exponential in γ), starting with the base case $\gamma = 1$ and then building on it, as follows.

- Base case: If $\gamma(G) = 1$, then $|\mathcal{T}_e(G)| \leq O(k^2)$ and thus G has a mimicking network of size $s = O(k^4)$;[2] see Theorem 4.1 and Corollary 4.3.
- General case, first bound: If $\gamma(G) \geq 1$, then $|\mathcal{T}_e(G)| \leq (2k/\gamma)^{2\gamma}$ and thus G has a mimicking network of size $s = O(k(2k/\gamma)^{4\gamma})$; see Theorem 4.4 and Corollary 4.5.

The last bound on $|\mathcal{T}_e(G)|$ is clearly wasteful (for $\gamma = k$, it is roughly quadratically worse than the trivial bound). To avoid overcounting of edges that belong to multiple elementary cutsets, we devise a better decomposition.

Further decomposition of elementary cutsets. Our second method decomposes each elementary cutset even further, in a special way such that we can count the underlying fragments (special subsets of edges) without repetitions, and this yields our main result.

- General case, second bound. When $\gamma(G) \geq 1$, there are $O(2^\gamma k^2)$ subsets of edges, such that every elementary cutset in G can be decomposed into a disjoint union of some of these subsets. Thus, G admits a mimicking network of size $O(\gamma 2^{2\gamma} k^4)$; see Theorem 4.6 and Corollary 4.7.

Additional results. First, all our cut sparsifiers are also approximate flow sparsifiers, by straightforward application of the known bounds on the flow-cut gap; see section 4.4. Second, our decompositions easily yield a succinct data structure that stores all the minimum terminal cuts of a planar graph G. Its storage requirement depends on $|\mathcal{T}_e(G)|$, which is bounded as above; see section 5 for details.

Finally, we show a duality between cut and distance sparsifiers (for certain graphs) and derive new relations between their bounds, as explained next.

1.2. Cuts versus distances. Although in several known scenarios cuts and distances are closely related, the following notion of distance sparsification was studied separately, with no formal connections to cut sparsifiers [22, 6, 2, 28, 24, 21, 13, 12, 16, 18].

DEFINITION 1.3. *A network $H = (V_H, E_H, T, c_H)$ is called a (q, s)-distance-approximating minor (DAM) of $G = (V, E, T, c)$ if it is a minor of G, its size is $|V_H| \leq s$, and*

$$(2) \qquad \forall t, t' \in T, \qquad d_G(t, t') \leq d_H(t, t') \leq q \cdot d_G(t, t'),$$

where $d_G(\cdot, \cdot)$ is the shortest-path metric in G with respect to $c(\cdot)$ as edge lengths.

We emphasize that the well-known planar duality between cuts and cycles does *not* directly imply a duality between cut and distance sparsifiers. We nevertheless do use this planar-duality approach, but we need to break "shortest cycles" into "shortest paths," which we achieve by adding new terminals (ideally not too many).

- Fix $k, q, s \geq 1$. Then all planar k-terminal networks with $\gamma = 1$ admit a minor (q, s)-cut sparsifier if and only if all these networks admit a $(q, O(s))$-DAM; see Theorems 6.1 and 6.2.

This result yields new cut-sparsifier bounds in the special case $\gamma = 1$ (see section 6.3). Notice that in this case of $\gamma = 1$ the flow-cut gap is 1 [35], hence the three problems of minor sparsification (of distances, of cuts, and of flows) all have the same asymptotic bounds and gaps.

---

[2] The generic bound implies $s = O(k^5)$, but we can slightly improve it in this case.

This duality can be extended to general $\gamma \geq 1$ (including $\gamma = k$), essentially at the cost of increasing the number of terminals, as follows. If for some functions $q(\cdot)$ and $s(\cdot)$, all planar k-terminal networks with given γ admit a $(q(k), s(k))$-DAM, then all networks in this class admit also a minor $(q(\gamma 2^\gamma k^2), s(\gamma 2^\gamma k^2))$-cut sparsifier. For $\gamma = k$, we can add only $k2^k$ new terminals instead of $k^3 2^k$. We omit the proof of this extension, as applying it to the known bounds for DAM yields alternative proofs for known/our cut-sparsifier bounds, but no new results. For example, using the reduction together with the known upper bound of $(1, k^4)$-DAM, we get that every planar k-terminal network with $\gamma(G) = k$ admits a minor mimicking network of size $O((k2^k)^4)$.

Comparison with previous techniques. Probably the closest notion to duality between cut sparsification and distance sparsification is Räcke's powerful method [37], adapted to vertex sparsification as in [7, 15, 32]. However, in his method the cut sparsifier H is inherently randomized; this is acceptable if H contains only the terminals, because we can take its "expectation" $\bar{H}$ (a complete graph with expected edge weights), but it is calamitous when H contains nonterminals, and then each randomized outcome has different vertices. Another related work, by Chen and Wu [10], reduces multiway-cut in a planar network with $\gamma(G) = 1$ to a minimum Steiner tree problem in a related graph G'. Their graph transformation is similar to one of our two reductions, although they show a reduction that goes in one direction rather than an equivalence between two problems.

1.3. Related work. Cut and distance sparsifiers were studied extensively in recent years, in an effort to optimize their two parameters, quality q and size s. The foregoing discussion is arranged by the quality parameter, starting with $q = 1$, then $q = O(1)$, and finally quality that grows with k.

Cut sparsification. Let us start with $q = 1$. Apart from the already mentioned work on a general graph G [23, 26, 29], there are also bounds for specific graph families, like bounded-treewidth or planar graphs [8, 29, 25]. For planar G with $\gamma(G) = 1$, there is a recent tight upper bound $s = O(k^2)$ [20] (independent of our work), where the sparsifier is planar but is *not* a minor of the original graph.

We proceed to a constant quality q. Chuzhoy [14] designed an $(O(1), s)$-cut sparsifier, where s is polynomial in the total capacity incident to the terminals in the original graph, and certain graph families (e.g., bipartite) admit sparsifiers with $q = 1 + \epsilon$ and $s = \text{poly}(k/\epsilon)$ [1].

Finally, we discuss the best quality known when $s = k$, i.e., the sparsifier has only the terminals as vertices. In this case, it is known that $q = O(\log k / \log \log k)$ [34, 31, 7, 15, 32], and there is a lower bound $q = \Omega(\sqrt{\log k})$ [32]. For networks that exclude a fixed minor (e.g., planar) it is known that $q = O(1)$ [15], and for trees $q = 2$ [21] (where the sparsifier is *not* a minor of the original tree).

Distance sparsification. A separate line of work studied the trade-off between the quality q and the size s of a DAM. For $q = 1$, every graph admits DAM of size $s = O(k^4)$ [28], and there is a lower bound of $s = \Omega(k^2)$ even for planar graphs [28]. Independently of our work, Goranci, Henzinger, and Peng [20] recently constructed, for planar graphs with $\gamma(G) = 1$, a $(1, O(k^2))$-distance sparsifier that is planar but *not* a minor of the original graph. Proceeding to quality $q = O(1)$, planar graphs admit a DAM with $q = 1 + \epsilon$ and $s = O(k \log k/\epsilon)^2$ [13], and certain graph families, such as trees and outerplanar graphs, admit a DAM with $q = O(1)$ and $s = O(k)$ [22, 2, 6, 28]. When $s = k$ (the sparsifier has only the terminals as vertices), then known quality is $q = O(\log k)$ for every graph [24, 12, 16]. Additional trade-offs and lower bounds can be found in [6, 28, 13].

1.4. Preliminaries. Let $G = (V, E, T, c)$ be a k-terminal network, and denote its k terminals by $T = \{t_1, \ldots, t_k\}$. We assume without loss of generality that G is connected, as otherwise we can construct a sparsifier for each connected component separately. For every $S \subset T$, let $\operatorname{argmincut}_G(S)$ denote the *argument* of the minimizer in $\operatorname{mincut}_G(S)$, i.e., the minimum-cost cutset that separates S from $\bar{S} = T \setminus S$ in G. We assume that the minimum is unique by a perturbation of the edge weights. Throughout, when G is clear from the context, we use the shorthand

$$(3) \qquad\qquad E_S := \operatorname{argmincut}_G(S).$$

Similarly, $CC(E_S)$ is shorthand for the set of connected components of the graph $G \setminus E_S$. Define the *boundary* of $W \subseteq V$, denoted $\delta(W)$, as the set of edges with exactly one endpoint in W, and observe that for every connected component $C \in CC(E_S)$ we have $\delta(C) \subseteq E_S$. By symmetry, $E_S = E_{\bar{S}}$. And since G is connected and $S \neq \emptyset, T$, we have $E_S \neq \emptyset$ and $|CC(E_S)| \geq 2$. In addition, by the minimality of E_S, every connected component $C \in CC(E_S)$ contains at least one terminal.

LEMMA 1.4 (see [29, Lemma 2.2]). *For every two subsets of terminals $S, S' \subset T$ and their corresponding minimum cuts $E_S, E_{S'}$, every connected component $C \in CC(E_S \cup E_{S'})$ contains at least one terminal.*

2. Elementary cutsets in general graphs. In this section we define a special set of cutsets called elementary cutsets (Definition 2.1) and prove that these elementary cutsets generate all other relevant cutsets, namely, the minimum terminal cutsets in the graph (Theorem 2.5). Therefore, to produce a cut sparsifier, it is enough to preserve only these elementary cutsets (Corollary 2.6). In the following discussion, we fix a network $G = (V, E, T, c)$ and employ the notation E_S, $CC(E_S)$ and $\delta(W)$ set up in section 1.4.

DEFINITION 2.1 (elementary cutset). *Fix $S \subset T$. Its minimum cutset E_S is called an* elementary cutset *if $|CC(E_S)| = 2$.*

DEFINITION 2.2 (elementary component). *A subset $C \subseteq V$ is called an* elementary component *if $\delta(C)$ is an elementary cutset for some $S \subset T$, i.e., $\delta(C) = E_{C \cap T}$ and $|CC(\delta(C))| = 2$.*

Although the following two lemmas are quite straightforward, they play a central role in the proof of Theorem 2.5.

LEMMA 2.3. *Fix a subset $S \subset T$ and its minimum cutset E_S. The boundary of every $C \in CC(E_S)$ is itself the minimum cutset separating the terminals $T \cap C$ from $T \setminus C$ in G, i.e., $\delta(C) = E_{T \cap C}$.*

Proof. Assume toward contradiction that $\delta(C) \neq E_{T \cap C}$. Since both sets of edges separate between the terminals $T \cap C$ and $T \setminus C$, then $c(E_{T \cap C}) < c(\delta(C))$. Let us replace the edges $\delta(C)$ by the edges $E_{T \cap C}$ in the cutset of E_S and call this new set of edges E'_S, i.e., $E'_S = (E_S \setminus \delta(C)) \cup E_{T \cap C}$. It is clear that $c(E'_S) < c(E_S)$. We will prove that E'_S is also a cutset that separates between S and $\bar{S}$ in the graph G, contradicting the minimality of E_S.

Assume without loss of generality that $T \cap C \subseteq S$, and consider $E_S \setminus \delta(C)$. By the minimality of E_S all the neighbors of C contain terminals of $\bar{S}$, therefore the cutset $E_S \setminus \delta(C)$ separates the terminals $S \setminus (T \cap C)$ from $\bar{S} \cup (T \cap C)$ in G. Now consider $E'_S = (E_S \setminus \delta(C)) \cup E_{T \cap C}$ and note that the connected component $C_{C + N_S(C)}$ contains all the terminals $T \cap C$ and some terminals of $\bar{S}$. This cutset E'_S clearly separates

$T \cap C$ from all other terminals and also separates $S \setminus (T \cap C)$ from $\bar{S} \cup (T \cap C)$. Altogether this cutset separates between S and $\bar{S}$ in G, and the lemma follows. $\square$

LEMMA 2.4. *For every $S \subset T$, at least one component in $CC(E_S)$ is elementary.*

Proof. Fix $S \subset T$. Lemma 2.3 yields that $\delta(C) = E_{C \cap T}$ for every $C \in CC(E_S)$; thus it is left to prove that there exists $C' \in CC(E_S)$ such that $|CC(\delta(C'))| = 2$. For simplicity, we shall represent our graph G as a bipartite graph $\mathcal{G}_S$ whose vertices and edges are $CC(E_S)$ and E_S, respectively, i.e., we get $\mathcal{G}_S$ by contracting every $C \in CC(E_S)$ in G into a vertex v_C. Let $V_1(\mathcal{G}_S) = \{v_C : C \cap T \subseteq S\}$ and $V_2(\mathcal{G}_S) = \{v_C : C \cap T \subseteq \bar{S}\}$ be the partition of $V(\mathcal{G}_S)$ into two sets. By the minimality of E_S the graph $\mathcal{G}_S$ is connected, and each of $V_1(\mathcal{G}_S)$ and $V_2(\mathcal{G}_S)$ is an independent set.

For every connected component $C \in CC(E_S)$, it is easy to see that $|CC(\delta(C))| = 2$ if and only if $\mathcal{G}_S \setminus \{v_C\}$ is connected. Since $\mathcal{G}_S$ is connected, it has a spanning tree and thus $\mathcal{G}_S \setminus \{v_{C'}\}$ is connected for every leaf $v_{C'}$ of that spanning tree, and the lemma follows. $\square$

THEOREM 2.5 (decomposition into elementary cutsets). *For every $S \subset T$, the minimum cutset E_S can be decomposed into a disjoint union of elementary cutsets.*

The idea of the proof is to iteratively decrease the number of connected components in $CC(E_S)$ by uniting an elementary connected component with all its neighbors (while recording the cutset between them), until we are left with only one connected component—all of V.

Proof of Theorem 2.5. We will need the following definition. Given $S \subset T$ and its minimum cutset E_S, we say that two connected components $C, C' \in CC(E_S)$ are *neighbors with respect to E_S* if E_S has an edge from C to C'. We denote by $N_S(C) \subseteq CC(E_S)$ the set of neighbors of C with respect to E_S. Observe that removing $\delta(C)$ from the cutset E_S is equivalent to uniting the connected component C with all its neighbors $N_S(C)$. Denoting this new connected component by $C_{C+N_S(C)}$ we get that $CC(E_S \setminus \delta(C)) = \left(CC(E_S) \setminus (\{C\} \cup N_S(C)) \right) \cup \{C_{C+N_S(C)}\}$.

Let E_S be a minimum cutset that separates S from $\bar{S}$. By Lemma 2.4, there exists a component $C \in CC(E_S)$ that is elementary, and by Lemma 2.3, $\delta(C) = E_{T \cap C}$. Assume without loss of generality that $T \cap C \subseteq S$ (rather than $\bar{S}$), and unite C with all its neighbors $N_S(C)$. Now, we would like to show that this step is equivalent to "moving" the terminals in C from S to $\bar{S}$. Clearly, the new cutset $E_S \setminus \delta(C)$ separates the terminals $S' = S \setminus (T \cap C)$ from $T \setminus S' = \bar{S} \cup (T \cap C)$, but to prove that

$$(4) \qquad\qquad E_{S'} = E_S \setminus \delta(C),$$

we need to argue that this new cutset has *minimum cost* among those separating S' from $\bar{S}'$. To this end, assume to the contrary; then $E_{S'}$ must have a strictly smaller cost than $E_S \setminus \delta(C)$, because both cutsets separate S' from $T \setminus S'$. Now similarly to the proof of Lemma 2.3, it follows that $E_{S'} \cup \delta(C)$ separates S from $\bar{S}$ and has a strictly smaller cost than E_S, which contradicts the minimality of E_S.

Using (4), we can write $E_S = \delta(C) \cup E_{S'}$ and continue iteratively with $E_{S'}$ while it is nonempty (i.e., $|CC(E_{S'})| > 1$). Formally, the theorem follows by induction on $|CC(E_S)|$. $\square$

To easily examine all the elementary cutsets in a graph G, we define

$$\mathcal{T}_e(G) := \{S \subset T \; : \; |CC(E_S)| = 2\}.$$

Using Theorem 2.5, the cost of every minimum terminal cut can be recovered, in a

certain manner, from the costs of the elementary cutsets of G, and this yields the following corollary.

COROLLARY 2.6. *Let H be a k-terminal network with same terminals as G. If $\mathcal{T}_e(G) = \mathcal{T}_e(H)$ and*

$$(5) \qquad \forall S \in \mathcal{T}_e(G), \qquad \mathrm{mincut}_G(S) \leq \mathrm{mincut}_H(S) \leq q \cdot \mathrm{mincut}_G(S),$$

then H is a cut-sparsifier of G of quality q.

Proof. Given G and H as above, we only need to prove (1). To this end, fix $S \subset T$. Observe that for every $\varphi \subseteq \mathcal{T}_e(G) = \mathcal{T}_e(H)$, the set $\cup_{S' \in \varphi} \mathrm{argmincut}_G(S')$ is S-separating in G if and only if for every $t \in S$ and $t' \in \bar{S}$ there exists $S' \in \varphi$ such that without loss of generality $t \in S'$ and $t' \notin S'$. Thus, if φ is a partition of S, i.e., $S = \cup_{S' \in \varphi} S'$, then $\cup_{S' \in \varphi} \mathrm{argmincut}_G(S')$ is S-separating in G. Since the same arguments hold also for H, we get the following:

$$\underset{S' \in \varphi}{\cup} \mathrm{argmincut}_G(S') \text{ is } S\text{-separating in } G \Leftrightarrow \underset{S' \in \varphi}{\cup} \mathrm{argmincut}_H(S') \text{ is } S\text{-separating in } H.$$

By Theorem 2.5, there exists $\varphi_G \subseteq \mathcal{T}_e(G)$ such that $S = \cup_{S' \in \varphi_G} S'$ and $\mathrm{argmincut}_G(S) = \cup_{S' \in \varphi_G} \mathrm{argmincut}_G(S')$, thus

$$\mathrm{mincut}_H(S) \leq \sum_{S' \in \varphi_G} \mathrm{mincut}_H(S') \leq q \cdot \sum_{S' \in \varphi_G} \mathrm{mincut}_G(S') = q \cdot \mathrm{mincut}_G(S).$$

Applying Theorem 2.5 to H together with an analogous argument yields that $\mathrm{mincut}_G(S) \leq \mathrm{mincut}_H(S)$, which proves (1) and the corollary follows. □

3. Mimicking networks for planar graphs. We now present an application of our results in section 2. We begin with a bound on the mimicking network size for a planar graph G as a function of the number of elementary cutsets (Theorem 3.1). We then obtain an upper bound of $O(k2^{2k})$ for every planar network (Corollary 3.2), which improves the previous work [29, Theorem 1.1] by a factor of k, thanks to the use of elementary cuts. The underlying reason is that the previous analysis in [29] considers all the 2^k possible terminal cutsets, and each of them is a collection of at most k simple cycles in the dual graph G^*. We can consider only the elementary cutsets by Corollary 2.6, and each of them is a simple cycle in G^* by Definition 2.1. Thus, we consider a total of 2^k simple cycles, saving a factor of k over the earlier naive bound of $k2^k$ simple cycles.

THEOREM 3.1. *Every planar network G, in which $|CC(E_S \cup E_{S'})| \leq \alpha \; \forall \; S, S' \in \mathcal{T}_e(G)$, admits a minor mimicking network H of size $O(\alpha \cdot |\mathcal{T}_e(G)|^2)$.*

The proof of this theorem appears in section 3.1. It is based on applying the machinery of [29], but restricting the analysis to elementary cutsets.

COROLLARY 3.2. *Every planar network G admits a minor mimicking network of size $O(k2^{2k})$.*

Proof. Apply Theorem 3.1, using an easy bound $\alpha = O(k)$ from Lemma 1.4, and a trivial bound on the number of elementary cutsets $|\mathcal{T}_e(G)| \leq 2^k$. □

3.1. Proof of Theorem 3.1. Given a k-terminal network G and $\alpha > 0$ such that $|CC(E_S \cup E_{S'})| \leq \alpha$ for every $S, S' \in \mathcal{T}_e(G)$, we prove that it admits a minor mimicking network H of size $O(\alpha |\mathcal{T}_e(G)|^2)$. Let $\hat{E} = \cup_{S \in \mathcal{T}_e(G)} E_S$, and construct

H by contracting every connected component of $G \setminus \hat{E}$ into a single vertex. Notice that edge contractions can only increase the cost of any minimum terminal cut, and that in our construction edges of an elementary cutset of G are never contracted. Thus, the resulting H is a minor of G that maintains all the elementary cutsets of G, and by Corollary 2.6 H maintains all the terminal mincuts of G. We proceed to bound the number of connected components in $G \setminus \hat{E}$, as this will clearly be the size of our mimicking network H. The crucial step here is to use the planarity of G by employing the dual graph of G denoted by G^* (for basic notions of planar duality see Appendix A).

Loosely speaking, the elementary cutsets in G correspond to cycles in the dual graph G^*, and thus we consider the dual edges of $\hat{E}$, which may be viewed as a subgraph of G^* comprising (many) cycles. We then use Euler's formula and the special structure of this subgraph of cycles; more specifically, we count its meeting vertices, which turns out to require the aforementioned bound of α for two sets of terminals S, S'. This gives us a bound on the number of faces in this subgraph, which in turn is exactly the number of connected components in the primal graph (Lemma 3.4). Observe that removing edges from a graph G can disconnect it into (one or more) connected components. The next lemma characterizes this behavior in terms of the dual graph G^*. Let $V_m(G)$ be all the vertices in the graph G with degree ≥ 3, and call them *meeting vertices* of G. The following lemma bounds the number of meeting vertices in two elementary cuts by $O(\alpha)$.

LEMMA 3.3. *For every two subsets of terminals $S, S' \in \mathcal{T}_e(G)$, the dual graph $G^*[E_S^* \cup E_{S'}^*]$ has at most 2α meeting vertices.*

Proof sketch. For simplicity denote by $G_{SS'}^*$ the graph $G^*[E_S^* \cup E_{S'}^*]$. By our assumption, the graph $G \setminus (E_S \cup E_{S'})$ has at most α connected components. By Lemma A.2 every connected component in $G \setminus (E_S \cup E_{S'})$ corresponds to a face in $G_{SS'}^*$. Therefore, $G_{SS'}^*$ has at most α faces. Let $V_{SS'}, E_{SS'}$, and $F_{SS'}$ be the vertices, edges, and faces of the graph $G_{SS'}^*$. Note that the degree of every vertex in that graph is at least 2. Thus, by the degree-sum formula (the total degree of all vertices equals twice the number of edges), $2|E_{SS'}| \geq 2|V_{SS'} \setminus V_m(G_{SS'}^*)| + 3|V_m(G_{SS'}^*)|$ and so $|E_{SS'}| \geq |V_{SS'}| + \frac{1}{2}V_m(G_{SS'}^*)$. Together with the Euler formula we get that $\alpha \geq |F_{SS'}| \geq |E_{SS'}| - |V_{SS'}| \geq \frac{1}{2}|V_m(G_{SS'}^*)|$, and the lemma follows. □

LEMMA 3.4. *The dual graph $G^*[\hat{E}^*]$ has at most $O(\alpha|\mathcal{T}_e(G)|^2)$ faces. Thus, $G \setminus \hat{E}$ has at most $O(\alpha|\mathcal{T}_e(G)|^2)$ connected components.*

Proof sketch. For simplicity denote by $\hat{G}^*$ the graph $G^*[\hat{E}^*]$, and let $E_m(\hat{G}^*)$ be all the edges in $\hat{G}^*$ that are incident to meeting vertices. Fix an elementary subset of terminals $S \in \mathcal{T}_e(G)$. By Lemma 3.3 there are at most 2α meeting vertices in $G^*[E_S^* \cup E_{S'}^*]$ for every $S' \in \mathcal{T}_e(G)$. Summing over all the different S' in $\mathcal{T}_e(G)$ we get that there are at most $2\alpha|\mathcal{T}_e(G)|$ meeting vertices on the cycle E_S^* in the graph $\hat{G}^*$. Since the degree of every vertex in $G^*(E_S^*)$ is 2, we get that

$$|E_S^* \cap E_m(\hat{G}^*)| \leq 2|V(G^*(E_S^*)) \cap V_m(\hat{G}^*)| \leq 4\alpha|\mathcal{T}_e(G)|.$$

Again summing over at most $|\mathcal{T}_e(G)|$ different elementary subsets S we get that $|E_m(\hat{G}^*)| \leq 4\alpha|\mathcal{T}_e(G)|^2$. Plugging it into the Euler formula for the graph $G^*[\hat{E}^*]$, together with the inequality $|E(\hat{G}^*) \setminus E_m(\hat{G}^*)| \leq |V(\hat{G}^*) \setminus V_m(\hat{G}^*)|$ by the fact that the two sides represent the edges and vertices of a graph consisting of vertex-disjoint paths (because its maximum degree is at most 2), we get the following:

$$|F(\hat{G}^*)| = |E(\hat{G}^*)| - |V(\hat{G}^*)| + 1 + |CC(\hat{G}^*)|$$
$$\leq |E_m(\hat{G}^*)| - |V_m(\hat{G}^*)| + 1 + |CC(\hat{G}^*)|$$
$$\leq 4\alpha|\mathcal{T}_e(G)|^2 + 1 + |CC(\hat{G}^*)|.$$

Since $|\mathcal{T}_e(G)| \geq k$ it is left to bound $|CC(\hat{G}^*)|$ by k. Assume toward contradiction that $|CC(\hat{G}^*)| \geq k + 1$; thus there exists a connected component W in $\hat{G}^*$ that does not contains a terminal face of G^*. By the construction of $\hat{E}^*$, W contains at least one elementary shortest cycle that separates between terminal faces of G^* in contradiction. Finally, Lemma A.2 with $M = \hat{E}$ yields that $|CC(G \setminus \hat{E})| = |F(G^*[\hat{E}^*])| = O(\alpha|\mathcal{T}_e(G)|^2)$ and the lemma follows.   □

Recall that we construct our mimicking network H by contracting every connected component of $G \setminus \hat{E}$ into a single vertex. By Lemma 3.4 we get that H is a minor of G of size $O(\alpha|\mathcal{T}_e(G)|^2)$ and Theorem 3.1 follows.

4. Mimicking networks for planar graphs with bounded $\gamma(G)$. In this section, the setup is that $G = (V, E, T, c)$ is a planar k-terminal network with terminal face cover $\gamma = \gamma(G)$. Let $f_1, \ldots, f_\gamma$ be faces that are incident to all the terminals, and let k_i denote the number of terminals incident to face f_i. We can in effect assume that $\sum_{i=1}^{\gamma} k_i = k$, because we can count each terminal as incident to only one face and "ignore" its incidence to the other $\gamma - 1$ faces (if any).

Our goal is to construct for G a mimicking network H and bound its size as a function of k and $\gamma(G)$. Our construction of H is the same as in Theorem 3.1, and the challenge is to bound its size. The implications to flow sparsifiers are discussed in section 4.4.

All terminals are on one face. We start with the basic case $\gamma = 1$, i.e., all the terminals are on the same face, which we can assume to be the outerface. The idea is to apply Theorem 3.1. The first step is to characterize all the elementary cutsets, which yields immediately an upper bound on their number. The second step is to analyze the interaction between any two elementary cutsets.

THEOREM 4.1. *In every planar k-terminal network G with $\gamma(G) = 1$, the number of elementary cutsets is $|\mathcal{T}_e(G)| \leq \binom{k}{2}$.*

The proof, appearing in section 4.1, is based on two observations that view the outerface as a cycle of vertices: (1) every elementary cutset disconnects the outerface's cycle into two paths, which we call intervals (see Definition 4.8), and (2) every such interval can be identified by the terminals it contains. It then follows that every elementary cutset E_S is uniquely determined by two terminals, leading to the required bound.

The next lemma bounds the interaction between any two elementary cutsets. Its proof appears at the end of section 4.1.

LEMMA 4.2. *For every planar k-terminal network G with $\gamma(G) = 1$, and for every $S, S' \in \mathcal{T}_e(G)$, there are at most four connected components in $G \setminus (E_S \cup E_{S'})$.*

COROLLARY 4.3. *Every planar k-terminal network G with $\gamma(G) = 1$ admits a minor mimicking network of size $s = O(k^4)$.*

Proof. Apply Theorem 3.1 with $\alpha = 4$ from Lemma 4.2 and $|\mathcal{T}_e(G)| \leq k^2$ from Theorem 4.1.   □

All terminals are on γ faces, first bound. Our first (and weaker) bound for the general case $\gamma \geq 1$ follows by applying Theorem 3.1. To this end, we bound the number

of elementary cutsets by $(2k/\gamma)^{2\gamma}$ in Theorem 4.4, whose proof is in section 4.2, and then conclude a mimicking network size of $O(k(2k/\gamma)^{4\gamma})$ in Corollary 4.5.

THEOREM 4.4. *In every planar k-terminal network G with $\gamma = \gamma(G)$, the number of different elementary cutsets is $|\mathcal{T}_e(G)| \leq 2^{2\gamma} \left(\Pi_{i=1}^{\gamma} k_i^2 \right) \leq (2k/\gamma)^{2\gamma}$.*

COROLLARY 4.5. *Every planar k-terminal network G with $\gamma = \gamma(G)$ admits a mimicking network of size $s = O\left(k(2k/\gamma)^{4\gamma} \right)$.*

Proof. Apply Theorem 3.1, using $\alpha = O(k)$ from Lemma 1.4 and $|\mathcal{T}_e(G)| \leq (2k/\gamma)^{2\gamma}$ from Theorem 4.4. □

All terminals are on γ faces, second bound. Our second (and improved) result for the general case $\gamma \geq 1$ follows by a refined analysis of the elementary cutsets. While our bound of $(2k)^{2\gamma}$ on the number of elementary cutsets is tight, it leads to a wasteful mimicking network size (for example, plugging the worst-case $\gamma = k$ into Corollary 4.5 is inferior to the bound in Corollary 3.2). The reason is that this approach overcounts edges of the mimicking network, and we therefore devise a new proof strategy that decomposes each elementary cutset even further, in a special way that lets us to count the underlying fragments (special subsets of edges) without repetitions. We remark that the actual proof works in the dual graph G^* and decomposes a simple cycle into (special) paths.

THEOREM 4.6 (further decomposition of elementary cutsets). *Every planar k-terminal network G with $k_1, \ldots, k_\gamma$ as above has $p = 2^\gamma \left(1 + \sum_{i,j=1}^{\gamma} k_i k_j \right) \leq O(2^\gamma k^2)$ subsets of edges $E_1, \ldots, E_p \subset E$, such that every elementary cutset in G can be decomposed into a disjoint union of some of these E_i's, and each of E_i contains exactly two edges from the boundaries of the faces $f_1, \ldots, f_\gamma$.*

We prove this theorem in section 4.3. The main difficulty is to define subsets of edges that are contained in elementary cutsets and are also easy to identify. We implement this identification by attaching to every such subset a three-part label. We prove that each label is unique, and we count the number of different possible labels, which obviously bounds the number of such "special" subsets of edges.

COROLLARY 4.7. *Every planar k-terminal network G with $\gamma = \gamma(G)$ admits a minor mimicking network of size $s = O(\gamma 2^{2\gamma} k^4)$.*

A slightly weaker bound of $O(2^{2\gamma} k^5)$ on the mimicking network size follows easily from Theorem 3.1 by replacing elementary cutsets with our "special" subsets of edges. To this end, it is easy to verify that all arguments about elementary cutsets hold also for the "special" subsets of edges. This includes the bound $\alpha = O(k)$, because if every two elementary cutsets intersect at most $O(k)$ times, then certainly every two "special" subsets (which are subsets of elementary cutsets) intersect at most $O(k)$ times. We can thus apply Theorem 3.1 with $\alpha = O(k)$ and replace $|\mathcal{T}_e(G)|$ with $O(2^\gamma k^2)$ that we have by Theorem 4.6. The stronger bound in Corollary 4.7 follows by showing that $\alpha = O(\gamma)$ for "special" subsets of edges.

Proof of Corollary 4.7. Given a planar k-terminal network G with $\gamma(G) = \gamma$, use Theorem 4.6 to decompose the elementary terminal cutsets of G into $p = O(2^\gamma k^2)$ subsets of edges $E_1, \ldots, E_p$ as stated above. Since each of E_i has exactly two edges from the boundaries of the faces $f_1, \ldots, f_\gamma$, then for every E_i and E_j there are at most $O(\gamma)$ connected components in $G \setminus (E_i \cup E_j)$ that contain terminals. Let E_S and $E_{S'}$ be elementary cutsets of G such that $E_i \subseteq E_S$ and $E_j \subseteq E_{S'}$. By Lemma 1.4, each connected component in $G \setminus (E_S \cup E_{S'})$ must contain at least one terminal. Thus, each connected component in $G \setminus (E_i \cup E_j)$ must contain at least one terminal, which

bounds the number of its connected components by $O(\gamma)$. Apply Theorem 3.1 with $|\mathcal{T}_e(G)| = O(2^\gamma k^2)$ and $\alpha = O(\gamma)$ and the corollary follows. $\qquad\square$

4.1. Proof of Theorem 4.1 and Lemma 4.2. In this section we prove Theorem 4.1, which bounds the number of elementary cutsets when $\gamma = 1$. We start with a few definitions and lemmas. Let $G = (V, E, T, c)$ be a connected planar k-terminal network, such that the terminals $t_1, \ldots, t_k$ are all on the same face in that order. Assume without loss of generality that this special face is the outerface f_∞. We refer to this outerface as a clockwise-ordered cycle $\langle v_1^\infty, v_2^\infty, \ldots, v_l^\infty \rangle$, such that for every two terminals t_i, t_j if $v_x^\infty = t_i$ and $v_y^\infty = t_j$, then $i < j$ if and only if $x < y$.

DEFINITION 4.8. *An* interval *of f_∞ is a subpath $\hat{I} = \langle v_i^\infty, v_{i+1}^\infty, \ldots, v_j^\infty \rangle$ if $i \leq j$ and in the case where $i > j$ $\hat{I} = \langle v_i^\infty, \ldots, v_l^\infty, v_1^\infty, \ldots, v_j^\infty \rangle$. Denote its vertices by $V(\hat{I})$ or, slightly abusing notation, simply by $\hat{I}$.*

Two trivial cases are a single vertex $\langle v_i^\infty \rangle$ if $i = j$ and the entire outerface cycle f_∞ if $i = j - 1$.

DEFINITION 4.9. *Given $W \subseteq V$, an interval $\hat{I} = \langle v_i^\infty, v_{i+1}^\infty, \ldots, v_j^\infty \rangle$ is called* maximal *with respect to W if $V(\hat{I}) \subseteq W$ and no interval in W strictly contains $\hat{I}$, i.e., $v_{i-1}^\infty, v_{j+1}^\infty \notin W$. Let $\mathcal{I}(W)$ be the set of all maximal intervals with respect to W, and let the* order *of $W \subseteq V$ be $|\mathcal{I}(W)|$.*

Observe that $\mathcal{I}(W)$ is a unique partition of $W \cap V(f_\infty)$, hence the order of W is well defined. Later on, we apply Definition 4.9 to connected components $C \in CC(E_S)$, instead of arbitrary subsets $W \subseteq V$. For example, in Figure 1, $\mathcal{I}(C_3) = \{I_3, I_5, I_9\}$, and the order of C_3 is $|\mathcal{I}(C_3)| = 3$.

LEMMA 4.10. *For every subset $S \subset T$, $\cup_{C \in CC(E_S)} \mathcal{I}(C)$ is a partition of $V(f_\infty)$.*

Proof. Fix $S \subset T$ and its minimum cutset E_S. Then $CC(E_S)$ is a partition of the vertices V into connected components. It induces a partition also of $V(f_\infty)$, i.e., $V(f_\infty) = \cup_{C \in CC(E_S)} \big(C \cap V(f_\infty)\big)$. By Definition 4.9, each $C \cap V(f_\infty)$ can be further partitioned into maximal intervals, given by $\mathcal{I}(C)$. Combining all these partitions, the lemma follows. See Figure 1 for an illustration. $\qquad\square$
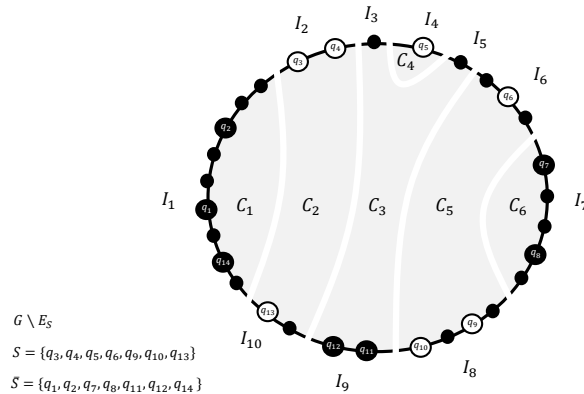

FIG. 1. *A minimum cutset E_S that partitions the graph G into 6 connected components $CC(E_S) = \{C_1, \ldots, C_6\}$ and partitions the outerface vertices into 10 intervals $\mathcal{I}(CC(E_S)) = \{I_1, \ldots, I_{10}\}$. Large nodes represent terminals, colored according to whether they lie in S or $\bar{S}$.*

LEMMA 4.11. *For every $S \subset T$ and its minimum cutset E_S, if $C \in CC(E_S)$ is an elementary connected component in G, then $|I(C)| = 1$.*

Proof. Since C is elementary, there are exactly two connected components C and C' in $CC(\delta(C))$. By Lemma 1.4 each of C and C' contains at least one terminal. Since all the terminals are on the outerface, each of C and C' contains at least one interval. Assume toward contradiction that C contains at least two maximal intervals I_1 and I_3; then there must be at least two intervals I_2 and I_4 in C' that appear on the outerface in an alternating order, i.e., I_1, I_2, I_3, I_4. Let v_i be a vertex in the interval I_i, and denote by P_{13} and P_{24} a path that connects between v_1, v_3 and between v_2, v_4 correspondingly. Note that P_{13} is contained in C and P_{24} is contained in C'. Moreover note that these two paths must intersect each other, giving a contradiction, and the lemma follows. □

We are ready to prove Theorem 4.1. Recall that for every $S \subset T$, if E_S is an elementary cutset, then by Lemma 4.11 each of S and $\bar{S}$ must be a single interval. Hence they must be of the form $\{t_i, t_{i+1}, \ldots, t_j\}$ and $\{t_{j+1}, \ldots, t_{i-1}\}$ allowing wrap-around. Thus, we can characterize S and $\bar{S}$ by the pairs (i, j) and $(j + 1, i - 1)$, respectively. There are at most $k(k-1)$ such different pairs, since $S \neq T$ and thus $j \neq i - 1$. By the symmetry between S and $\bar{S}$, we should divide that number by 2 and Theorem 4.1 follows.

Proof of Lemma 4.2. Let E_S and $E_{S'}$ be two elementary minimum cutsets, and let C_S and $C_{\bar{S}}$ be the two elementary connected components in $CC(E_S)$. By Lemma 4.11, each of C_S and $C_{\bar{S}}$ contains exactly one maximal interval denoted by I_S and $I_{\bar{S}}$, respectively, and similarly denote $C_{S'}, C_{\bar{S}'}, I_{S'}$, and $I_{\bar{S}'}$ for $E_{S'}$. Since each of the cutsets E_S and $E_{S'}$ intersects the cycle of the outerface in exactly two edges, the cutset $E_S \cup E_{S'}$ intersects the cycle of the outerface in at most four edges. Therefore the graph $G \setminus (E_S \cup E_{S'})$ has at most four maximal intervals. By Lemma 1.4, every connected component in $CC(E_S \cup E_{S'})$ must contain at least one terminal. Since all the terminals lie on the outerface, any connected component that contains a terminal must also contain an interval. Every interval is contained in exactly one connected component. Thus, there are at most four connected components in $CC(E_S \cup E_{S'})$, and the lemma follows. See Figure 2 for an illustration. □

4.2. Proof of Theorem 4.4. In this section we prove Theorem 4.4, which bounds the number of elementary cutsets when $\gamma > 1$. Since we assume (by perturbation) that there is a one-to-one correspondence between $S \subseteq T$ and E_S, it suffices
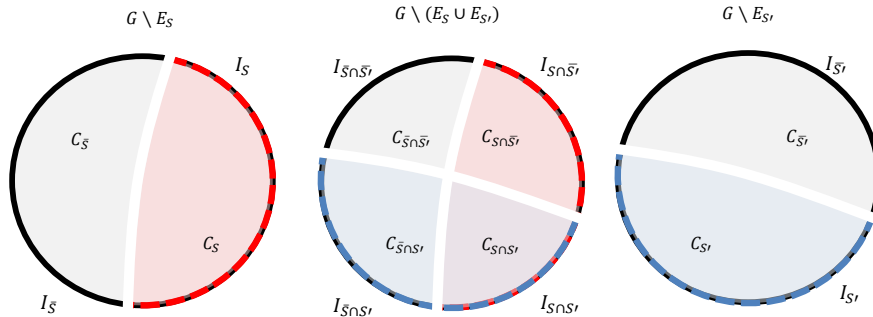

FIG. 2. *The union of two elementary cutsets, $E_S \cup E_{S'}$, disconnects G into (at most) four connected components and the outerface into four intervals.*

to bound the number of different ways that an elementary cutset can partition the
terminals into S and $\bar{S}$. We achieve the latter by two observations, which are ex-
tensions of the ideas in Theorem 4.1. First, an elementary cutset can break each of
the γ faces into at most two paths, which overall splits the terminals into at most
2γ subsets. As each subset (path) can lie either in S or in $\bar{S}$, there are at most $2^{2\gamma}$
different ways to partition T into S and $\bar{S}$ (this bound includes cases where two paths
from the same face lie both in S or both in $\bar{S}$, which is equivalent to not breaking
the face into two paths). Second, there are $\binom{k_i}{2}$ ways that the face f_i can be broken
into two paths by elementary cutsets, which gives overall $\Pi_{i=1}^{\gamma}\binom{k_i}{2}$ ways to break all
the γ faces simultaneously. Combining these two observations leads to the required
bound.

We start with a few definitions. Let $G = (V, E, F, T, c)$ be a k-terminal network
with γ faces $f_1, \ldots, f_\gamma$, where each f_i contains the k_i terminals T_i (breaking ties
arbitrarily), where $T = \cup_{i=1}^{\gamma} T_i$ and thus $k = \sum_{i=1}^{\gamma} k_i$. Denote the terminals in T_i
by $t_1^i, \ldots, t_{k_i}^i$, where the order is by a clockwise order around the boundary of f_i,
starting with an arbitrary terminal; for simplicity, we shall write t_j instead of t_j^i when
the face f_i is clear from the context. Let G^* be the dual graph of G. The graph G^*
has k terminal faces $\{f_{t_j^i}\}$ that are dual to the terminals $\{t_j^i\}$ of G and has γ special
vertices $W = \{w_1, \ldots, w_\gamma\}$ that are dual to the faces $f_1, \ldots, f_\gamma$ of G (see Appendix A
for basic notions of planar duality).

We label each $S \in \mathcal{T}_e(G)$ (and its elementary cycle E_S^*) by two vectors $\bar{x}, \bar{y}$, as
follows. Since E_S^* is a simple cycle, it visits every vertex $w_i \in W$ at most once.
If it does visit w_i, then exactly two cycle edges are incident to w_i, and these two
edges naturally partition the faces around w_i into two subsets. Moreover, each subset
appears as a contiguous subsequence if the faces around w_i are scanned in a clockwise
order. In particular, the terminal faces $f_{t_1}, \ldots, f_{t_{k_i}}$ are partitioned into two subsets,
whose indices can be written as $\{t_{x_i}, \ldots, t_{y_i-1}\}$ and $\{t_{y_i}, \ldots, t_{x_i-1}\}$, for some $x_i, y_i \in
[k_i]$, under two conventions: (i) we allow wraparound, i.e., $t_{k_i+1} = t_1$ and so forth; (ii)
if $x_i = y_i$, then we have a trivial partition of T_i, where one subset is T_i and the other
is $\emptyset$. Observe that one of these subsets is contained in S and the other in $\bar{S}$; thus we
can assume that $\{t_{x_i}, \ldots, t_{y_i-1}\} \subseteq S$ and $\{t_{y_i}, \ldots, t_{x_i-1}\} \subset \bar{S}$. If the cycle E_S^* does
not visit w_i, then we simply define $x_i = y_i = 1$, which represents a trivial partitioning
of T_i. The labels are now defined as $\bar{x} = (x_1, \ldots, x_\gamma)$ and $\bar{y} = (y_1, \ldots, y_\gamma)$.

We now claim that G^* has at most $2^{2\gamma}$ elementary cycles with the same label
$(\bar{x}, \bar{y})$. To see this, fix $\bar{x}, \bar{y} \in [k_1] \times \cdots \times [k_\gamma]$ and modify G^* into a plane graph $G_{\bar{x},\bar{y}}^*$
with at most 2γ terminal faces, as follows. For every w_i, create a single terminal face
$f_{x_i}^i$ by "merging" faces around w_i, starting from $f_{t_{x_i}}$ and going in a clockwise order
until $f_{t_{y_i-1}}$ (inclusive). Then merge similarly the faces from $f_{t_{y_i}}$ and until $f_{t_{x_i-1}}$
into a single terminal face $f_{y_i}^i$. If $x_i = y_i$, then the two merging operations above
are identical and thus (as an exception) create only one terminal face denoted $f_{x_i}^i$.
Formally, a merge of two faces is implemented by removing the edge incident to w_i
that goes between the relevant faces. Observe that removing these edges in G^* can be
described in G as contracting the path around the boundary of the face f_i from the
terminal t_{x_i} to t_{y_i-1} and similarly from the terminal t_{y_i} to t_{x_i-1}; see Figure 3. It is
easy to verify that the modified graph $G_{\bar{x},\bar{y}}^*$ is planar and that every elementary cycle
E_S^* in G^* with this label $(\bar{x}, \bar{y})$ is also an elementary cycle in $G_{\bar{x},\bar{y}}^*$ that separates the
new terminal faces in a certain way. Usually, the new terminal faces are separated
into $\{f_{x_i}^i\}_{i=1}^{\gamma}$ and $\{f_{y_i}^i\}_{i=1}^{\gamma}$, except that when $x_i = y_i$, we have only one new terminal
face $f_{x_i}^i$, which should possibly be included with the y_i's instead of with the x_i'. Since
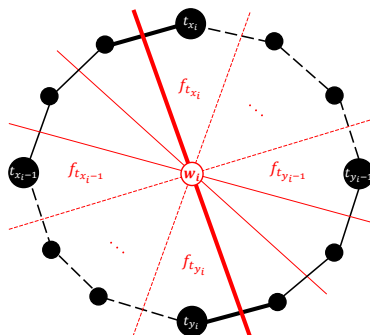
FIG. 3. *A simple cycle E_S^* separates the faces around w_i into two subsets. The primal graph G is shown in black and its dual G^* in red. Thicker lines are used for edges of the elementary cutset E_S and of the cycle E_S^*. Dashed lines represent dual edges that are removed by merging faces and primal edges that are contracted.*

$G_{\bar{x},\bar{y}}^*$ has at most 2γ terminal faces, it can have at most $2^{2\gamma}$ elementary cycles (one for each subset). This shows that for every label $(\bar{x}, \bar{y})$, there are at most $2^{2\gamma}$ different elementary cycles in G^*, as claimed.

Finally, the number of distinct labels $(\bar{x}, \bar{y})$ is clearly bounded by $\Pi_{i=1}^{\gamma} k_i^2$ and the above claim applies to each of them. By the inequality of arithmetic and geometric means $\Pi_{i=1}^{\gamma} k_i^2 \leq (k/\gamma)^{2\gamma}$. Therefore, the total number of different elementary cycles in G^* is at most $(2k/\gamma)^{2\gamma}$, and Theorem 4.4 follows.

4.3. Proof of Theorem 4.6. In this section we prove Theorem 4.6, which actually decomposes the elementary cutsets in a "bounded manner" when $\gamma > 1$. The idea is to consider the dual graph, which has γ special vertices, and elementary cycles. Since every elementary cycle is a simple cycle, it visits each of the γ vertices at most once, and thus we can decompose the elementary cycles into paths, such that the two endpoints of every path belong to the γ vertices. The challenging part is to count how many distinct paths are there.

We shall use the notation introduced in the beginning of section 4.2. In particular, the graph G has terminals $T = \cup_{i=1}^{\gamma} T_i$, where $T_i = \{t_1^i, \dots, t_{k_i}^i\}$ are the terminals on the boundary of special face f_i, and for simplicity we omit i when it is clear from the context. The dual graph, denoted G^*, has terminal faces $\{f_{t_j^i}\}$ and special vertices $W = \{w_1, \dots, w_\gamma\}$. Let $v_\infty \in V$ be the vertex whose dual face f_{v_∞} is the outerface of G^*.

Informally, the next definition determines whether f_v, the face dual to a vertex $v \in V$, lies "inside" or "outside" a circuit M^* in G^*. It works by counting how many times a path from v to v_∞ "crosses" M^* and evaluating it modulo 2 (i.e., its parity). The formal definition is more technical because it involves fixing a path, but the ensuing claim shows the value is actually independent of the path. Moreover, we need to properly define a "crossing" between a path Φ in G and a circuit in G^*; to this end, we view the path Φ as a sequence of faces in G^* that goes from f_v to f_{v_∞} and at each step "crosses" an edge of G^*.

DEFINITION 4.12 (parity of a dual face). *Let f_v be the dual face to a vertex $v \in V$, and fix a simple path in G between v and v_∞, denoted Φ. Let M^* be a circuit in G^*,*

and observe that its edges $E(M^*)$ form a multiset. Define the parity *of f_v with respect to M^* to be*

$$\mathrm{Par}(f_v, M^*) := \left(\sum_{e \in E(\Phi)} \mathrm{Count}(e^*, E(M^*)) \right) \bmod 2,$$

where $\mathrm{Count}(a, A)$ is the number of times an element a appears in a multiset A.

The next claim justifies the omission of the path Φ in the notation $\mathrm{Par}(f_v, M^*)$.

CLAIM 4.13. *Fix $v \in V$, and let Φ and Φ' be two paths in G between v and v_∞. Then for every circuit M^* in G^*,*

$$\sum_{e \in E(\Phi)} \mathrm{Count}(e^*, E(M^*)) = \sum_{e \in E(\Phi')} \mathrm{Count}(e^*, E(M^*)) \pmod 2.$$

Proof. Fix a vertex $v \in V$ and its dual face f_v. Fix also a circuit M^* and a decomposition of it into simple cycles. We say that a simple cycle in G^* (like one from the decomposition of M^*) *contains* the face f_v if that cycle separates f_v from the outerface f_{v_∞}. Let Φ be a path between v and v_∞. By the Jordan curve theorem, the path's dual edges $\{e^* : e \in E(\Phi)\}$ intersect a simple cycle in G^* an odd number of times if and only if that simple cycle contains the dual face f_v. By summing this quantity over the simple cycles in the decomposition of M^*, we get that

$$\sum_{e \in E(\Phi)} \mathrm{Count}(e^*, E(M^*)) = 1 \pmod 2$$

if and only if f_v is contained in an odd number of these simple cycles. The latter is clearly independent of the path Φ, which proves the claim. $\square$

Given a circuit M^* in G^*, we use the above definition to partition the terminals T into two sets according to their parity, namely,

$$T_{odd}(M^*) := \{t \in T : \mathrm{Par}(f_t, M^*) = 1\},$$
$$T_{even}(M^*) := \{t \in T : \mathrm{Par}(f_t, M^*) = 0\}.$$

Given $S \in \mathcal{T}_e(G)$, recall that E_S^* is the shortest cycle which is S-separating in G^* (i.e., it separates between the terminal faces S^* and $\bar{S}^*$). Since E_S^* is an elementary cycle, it separates the plane into exactly two regions, which implies, without loss of generality, $T_{odd}(E_S^*) = S$ and $T_{even}(E_S^*) = \bar{S}$. Moreover, E_S^* is a simple cycle and thus goes through every vertex of W at most once. We decompose E_S^* into $|W \cap V(E_S^*)|$ paths in the obvious way, where the two endpoints of each path, and only they, are in W, and we let Π_S denote this collection of paths in G^*. There are two exceptional cases here. First, if $|W \cap V(E_S^*)| = 1$, then we let Π_S contain one path whose two endpoints are the same vertex (so actually a simple cycle). Second, if $|W \cap V(E_S^*)| = 0$, then we let $\Pi_S = \emptyset$ (we will deal with this case separately later). Now define the set

$$\Pi := \bigcup_{S \in \mathcal{T}_e(G)} \Pi_S$$

as the collection of all the paths that are obtained in this way over all possible $S \in \mathcal{T}_e(G)$. Notice that if the same path is contained in multiple sets Π_S, then it is included in the set Π only once (in fact, this "overlap" is what we are trying to leverage).
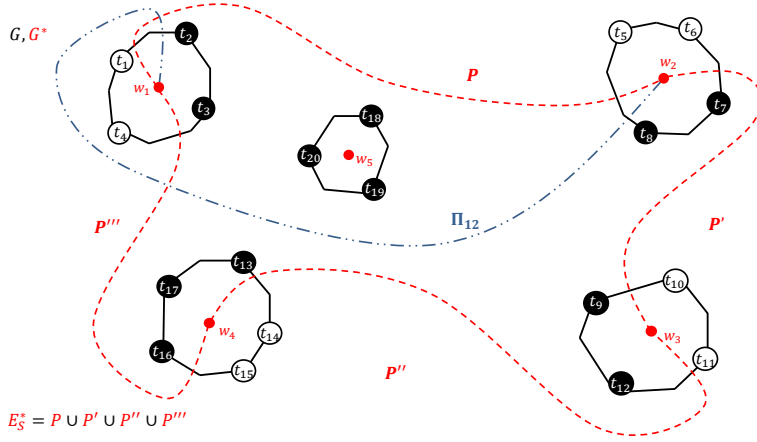
FIG. 4. *A planar 20-terminal network with $\gamma = 5$. Let $S \subset T$ be all the black terminals; then E_S^* (red dashed line) is split into 4 paths P, P', P'', P'''. The label of P, for example, is (1) w_1 and w_2, (2) t_1, t_2 and t_5, t_8, (3) $T_{odd}(P \cup \Pi_{12}) = \{t_1, t_2, t_3, t_4, t_{18}, t_{19}, t_{20}\}$, and is computed using Π_{12} (blue dashed line).*

Now give to each path $P \in \Pi$ a label that consists of three parts: (1) the two endpoints of P, say, $w_i, w_j \in W$; (2) the two successive terminals on each of the faces f_i and f_j, which describe where the path P enters vertices w_i and w_j, say, between t_{x-1}^i, t_x^i and between t_y^j, t_{y+1}^j; and (3) the set $T_{odd}(P \cup \Pi_{ij})$, where Π_{ij} is the shortest path (or any other fixed path) that agrees with parts (1) and (2) of the label and does not go through W, i.e., the shortest path between w_i and w_j that enters them between t_{x-1}^i, t_x^i and t_y^j, t_{y+1}^j and does not go through any other vertex in W. This includes the exceptional case $i = j$, in which P is actually a simple cycle; see Figure 4.

We proceed to show that each label is given to at most one path in Π (which will be used to bound $|\Pi|$). Assume toward contradiction that two different paths $P, P' \in \Pi$ get the same label, and suppose $c(P') < c(P)$. Suppose P is the path between w_i to w_j in E_S^* for $S \in T_e(G)$ and that P' is the path between the same endpoints (because of the same label) in $E_{S'}^*$, for another $S' \in T_e(G)$. By construction, the paths P and P' are simple, because E_S^* and $E_{S'}^*$ are elementary cycles, and only their endpoint vertices are from W.

The key to arriving at a contradiction is the next lemma. In these proofs, a path P is viewed as a multiset of edges $E(P)$, and the union and subtraction operations are applied to multisets. In particular, the union of two paths with the same endpoints gives a circuit.

LEMMA 4.14. *The circuit $(E_S^* \setminus P) \cup P'$ is S-separating.*

To prove this lemma, we will need the following two claims.

CLAIM 4.15. *Let A, B and C be (the edge sets of) simple paths in G^* between the same $w_i, w_j \in W$. Then*

$$\forall t \in T, \qquad \text{Par}(f_t, A \cup C) = \text{Par}(f_t, A \cup B) + \text{Par}(f_t, B \cup C) \pmod{2}.$$

Proof. Fix $t \in T$ and a path Φ between t and v_∞. Since A, B, and C are simple paths,

$$\sum_{e \in E(\Phi)} \mathrm{Count}(e^*, A \cup C) = |E^*(\Phi) \cap A| + |E^*(\Phi) \cap C|,$$

$$\sum_{e \in E(\Phi)} \mathrm{Count}(e^*, A \cup B) = |E^*(\Phi) \cap A| + |E^*(\Phi) \cap B|,$$

$$\sum_{e \in E(\Phi)} \mathrm{Count}(e^*, B \cup C) = |E^*(\Phi) \cap B| + |E^*(\Phi) \cap C|.$$

Summing the three equations above modulo 2 yields

$$\mathrm{Par}(f_t, A \cup C) + \mathrm{Par}(f_t, A \cup B) + \mathrm{Par}(f_t, B \cup C) = 0 \pmod 2,$$

which proves the claim. □

CLAIM 4.16. *Let $\triangle$ be the symmetric difference between two sets. For every three simple paths A, B, and C between $w_i, w_j \in W$,*

$$T_{odd}(A \cup C) = T_{odd}(A \cup B) \triangle T_{odd}(B \cup C).$$

Proof. Observe that $T_{odd}(A \cup B) \triangle T_{odd}(B \cup C)$ contains all $t \in T$ for which exactly one of $\mathrm{Par}(f_t, A \cup B)$ and $\mathrm{Par}(f_t, B \cup C)$ is equal to 1, which by Claim 4.15 is equivalent to having $\mathrm{Par}(f_t, A \cup C) = 1$. □

Proof of Lemma 4.14. To set up some notation, let $Q := E_S^* \setminus P$ be a simple path between w_i and w_j. Since E_S^* is a simple cycle that contains P, we can write $E_S^* = Q \cup P$.

The idea is to swap the path P in E_S^* with the other path P', which for the sake of analysis is implemented in two steps. The first step replace P (in E_S^*) with Π_{ij}, which gives the circuit $(E_S^* \setminus P) \cup \Pi_{ij} = Q \cup \Pi_{ij}$. The second step replaces Π_{ij} with P', which results in the circuit $Q \cup P' = (E_S^* \setminus P) \cup P'$. Now apply Claim 4.16 twice, once to the simple paths $A = Q$, $B = P$, and $C = \Pi_{ij}$, and once to the simple paths $A = Q$, $B = \Pi_{ij}$, and $C = P'$; we get that

$$T_{odd}(Q \cup \Pi_{ij}) = T_{odd}(E_S^*) \triangle T_{odd}(P \cup \Pi_{ij}),$$
$$T_{odd}(Q \cup P') = T_{odd}(Q \cup \Pi_{ij}) \triangle T_{odd}(\Pi_{ij} \cup P').$$

By plugging the first equality above into the second one, and observing that $T_{odd}(\Pi_{ij} \cup P) = T_{odd}(\Pi_{ij} \cup P')$ because P and P' have the same label, we obtain that

(6) $$T_{odd}(Q \cup P') = T_{odd}(E_S^*).$$

Finally, it is easy to verify that the circuit $Q \cup P'$ must separate between $T_{odd}(Q \cup P')$ and $T_{even}(Q \cup P')$. Using (6) and the fact that E_S^* is an elementary cycle, we know that $T_{odd}(Q \cup P') = T_{odd}(E_S^*) = S$, and thus $T_{even}(Q \cup P') = T \setminus S$. It follows that $Q \cup P'$ is S-separating, as required. □

Lemma 4.14 shows that the circuit $(E_S^* \setminus P) \cup P'$ is S-separating while also having lower cost than E_S^*. This contradicts the minimality of E_S^* and shows that the paths in Π have distinct labels. Thus, $|\Pi|$ is at most the number of distinct labels, and we will bound the latter using the following claim.

CLAIM 4.17. *Let $P \in \Pi$ be a path between w_i and w_j, and let $r \in [\gamma]$. Then*

$$\forall t, t' \in T_r, \qquad \mathrm{Par}(f_t, P \cup \Pi_{ij}) = \mathrm{Par}(f_{t'}, P \cup \Pi_{ij}),$$

where Π_{ij} is the shortest path with the same parts (1) and (2) of the label as P and does not go through any other vertices of W.

Proof. Since $t, t' \in T_r$, their dual faces f_t and $f_{t'}$ share w_r on their boundary. P and Π_{ij} are simple paths in G^* with the same endpoints, and thus $P \cup \Pi_{ij}$ is a circuit in G^*, which by construction does not go through any vertex w_r with $r \neq i, j$. Fix a path Φ in G between t and v_∞. We can extend it into a path Φ' between t' and v_∞, by taking a path $A_{t't}$ in G that goes around the face f_r between t' and t (both are on the face f_r, because $t, t' \in T_r$), and letting $\Phi' := A_{t't} \cup \Phi$.

Since P and Π_{ij} agree on the same parts (1) and (2) of the label, then $P \cup \Pi_{ij}$ have exactly two edges between some two successive terminals on each of the faces f_i and f_j. Thus, if $r \neq i, j$, then $|A_{t't} \cap (P \cup \Pi_{ij})| = 0$. If $r = i$ or $r = j$ but $i \neq j$, then $|A_{t't} \cap (P \cup \Pi_{ij})|$ is either 0 or 2. And if $r = i = j$, then $|A_{t't} \cap (P \cup \Pi_{ij})|$ is either 0, 2, or 4. Therefore, if we examine the parities of f_t and $f_{t'}$ with respect to $P \cup \Pi_{ij}$ using the paths Φ and $\Phi' = A_{t't} \cup \Phi$, respectively, we conclude that these parities are equal, as required.   □

We can now bound the number of possible labels of a path $P \in \Pi$. There are γ^2 possibilities for part 1 of the label, i.e., the endpoints $w_i, w_j \in W$ of P (note that we may have $i = j$). Given this data, there are $k_i k_j$ possibilities for part 2, i.e., between which two terminals the path P exits w_i and enters w_j. Furthermore, the number of possibilities for part 3 is the number of different subsets $T_{odd}(P \cup \Pi_{ij})$. By Claim 4.17 for every $r \in [\gamma]$ either $T_r \subseteq T_{odd}(P \cup \Pi_{ij})$ or $T_r \cap T_{odd}(P \cup \Pi_{ij}) = \emptyset$. Thus, the number of different subsets $T_{odd}(P \cup \Pi_{ij})$ is the number of different subsets of $\{T_1, \ldots, T_\gamma\}$, which is at most 2^γ. Altogether we get that there are at most $2^\gamma \sum_{i,j=1}^\gamma k_i \cdot k_j$ different labels.

Finally, there are also cycles E_S^* for $S \in \mathcal{T}_e(G)$ that do not go through any vertices of W, i.e., $W \cap V(E_S^*) = \emptyset$. Thus, they are not included in Π, so we count them now separately. Recall that without loss of generality $T_{odd}(E_S^*) = S$, i.e., every such cycle E_S^* is identified uniquely by a different subset $T_{odd}(\cdot)$. Since by Claim 4.17 there are at most 2^γ such subsets, we get that there are at most 2^γ such cycles. Adding them to our calculation, Theorem 4.6 follows.

4.4. Flow sparsifiers. Okamura and Seymour [35] proved that in every planar network with $\gamma(G) = 1$, the flow-cut gap is 1 (as usual, flow refers here to multicommodity flow between terminals). It follows immediately (see, e.g., [1]) that for such a graph G, every (q, s)-cut-sparsifier is itself also a (q, s)-flow-sparsifier of G. Thus, Corollary 4.3 implies the following.

COROLLARY 4.18. *Every planar k-terminal network G with $\gamma = 1$ admits a minor $(1, O(k^4))$-flow-sparsifier.*

Chekuri, Shepherd, and Weibel [9, Theorem 4.13] proved that in every planar network G, the flow-cut gap is at most $3\gamma(G)$, and thus Corollary 4.7 implies the following.

COROLLARY 4.19. *Every planar k-terminal network G with $\gamma = \gamma(G)$ admits a minor $(3\gamma, O(\gamma 2^{2\gamma} k^4))$-flow-sparsifier.*

5. Terminal-cuts scheme. In this section we present applications of our results in section 2 to data structures that store all the minimum terminal cuts in a graph G. As our focus is on the data structure's memory requirement, we do not discuss its query time. We start with a formal definition of such a data structure, and then provide our bounds of $\tilde{O}(|\mathcal{T}_e|)$ bits for general graphs (Theorem 5.2) and $\tilde{O}(2^\gamma k^2)$ bits

for planar graphs (Corollaries 5.3 and 5.4). In comparison, a trivial data structure for general graphs uses $\tilde{O}(2^k)$ bits, by storing the cost of all the terminal mincuts explicitly.

DEFINITION 5.1. *A terminal-cuts scheme (TC-scheme)* is a data structure that uses a storage (memory) M to support the following two operations on a k-terminal network $G = (V, E, T, c)$, where $n = |V|$ and $c : E \to \{1, \dots, n^{O(1)}\}$:
1. *Preprocessing, denoted P, which gets as input the network G and builds M.*
2. *Query, denoted R, which gets as input a subset of terminals S and uses M (without access to G) to output the cost of the minimum cutset E_S.*

We usually assume a machine word size of $O(\log n)$ bits, because even if G has only unit-weight edges, the cost of a cut might be $O(n^2)$, which is not bounded in terms of k.

THEOREM 5.2. *Every k-terminal network $G = (V, E, T, c)$ admits a TC-scheme with storage size of $O(|\mathcal{T}_e(G)|(k + \log n))$ bits, where $\mathcal{T}_e(G)$ is the set of elementary cutsets in G.*

Proof. We construct a TC-scheme as follows. In the preprocessing stage, given G, the TC-scheme stores $\langle S, c(E_S) \rangle$ for every $S \in \mathcal{T}_e(G)$, where S is written using k bits. The cost of every cutset is at most $|E| \cdot n^{O(1)} = \text{poly}(n)$, and thus the storage size of the TC-scheme is $O(|\mathcal{T}_e(G)|(k + \log n))$ bits, as required. Now given a subset $S \subset T$, the query operation $R(S; P(G))$ outputs

(7)
$$\min \Big\{ \sum_{S' \in \varphi} \text{mincut}_G(S') : \ \varphi \subseteq \mathcal{T}_e(G) \text{ s.t. } \bigcup_{S' \in \varphi} \text{argmincut}_G(S') \text{ is } S\text{-separating in } G \Big\}.$$

Since for every $\varphi \subseteq 2^T$, the cutset $\cup_{S' \in \varphi} \text{argmincut}_G(S')$ is S-separating in G if and only if $|\mathbb{1}_S(t_i) - \mathbb{1}_S(t_j)| \le \sum_{S' \in \varphi} |\mathbb{1}_{S'}(t_i) - \mathbb{1}_{S'}(t_j)| \ \forall \ i, j \in [k]$, the calculation in (7) can be done with no access to G. Clearly, $\text{mincut}_G(S) \le R(S; P(G))$. By Theorem 2.5, there is $\varphi \subseteq \mathcal{T}_e(G)$ such that $\text{argmincut}_G(S) = \cup_{S' \in \varphi} \text{argmincut}_G(S')$ and $\text{mincut}_G(S) = \sum_{S' \in \varphi} \text{mincut}_G(S')$. Thus, $R(S'; P(G)) = \text{mincut}_G(S')$.   □

COROLLARY 5.3. *Every planar k-terminal network G with $\gamma = 1$ admits a TC-scheme with storage size of $O(k^2 \log n)$ bits, i.e., $O(k^2)$ words.*

Proof. If G is a planar k-terminal network with $\gamma = 1$, then by Theorem 4.1 every $S \in \mathcal{T}_e(G)$ is equal to $\{t_i, t_{i+1}, \dots, t_j\}$ for some $i, j \in [k]$ and $|\mathcal{T}_e(G)| = \binom{k}{2}$ (recall that all the terminals $t_1, \dots, t_k$ are on the outerfaces of G in order). Thus, we can specify S via these two indices i and j, using only $O(\log k) \le O(\log n)$ bits (instead of k). The storage bound follows.   □

THEOREM 5.4. *Every planar k-terminal network G with $\gamma = \gamma(G)$ admits a TC-scheme with storage size of $O\big(2^\gamma \big(1 + \sum_{i,j=1}^\gamma k_i \cdot k_j\big)\big(\gamma + \log n\big)\big) \le O(2^\gamma k^2(\gamma + \log n))$ bits.*

Proof sketch. If G is a planar k-terminal network with bounded γ, then Theorem 4.6 characterizes $2^\gamma k^2$ special subsets of edges together with some small addition information for each such subset denoted by *label*. It further proves that all the elementary cuts can be restored using only the special subsets and their labels. As each label can be stored by at most $O(\gamma)$ bits, the storage bound follows.   □

6. Cut-sparsifier versus DAM in planar networks. In this section we prove the duality between cuts and distances in planar graphs with all terminals on the

outerface. Although the duality between shortest cycles and minimum cuts in planar graphs is known, the main difficulty is to transform all the shortest cycles into shortest paths without blowing up the number of terminals in the graph. We prove this duality using the following two theorems, and applications of them can be found in section 6.3.

THEOREM 6.1. *Let $G = (V, E, T, c)$ be a planar k-terminal network with all its terminals T on the outerface. One can construct a planar k-terminal network $G' = (V', E', T', c')$ with all its terminals T' on the outerface such that if G' admits a (q, s)-DAM, then G admits a minor $(q, O(s))$-cut-sparsifier.*

THEOREM 6.2. *Let $G = (V, E, T, c)$ be a planar k-terminal network with all its terminals T on the outerface. One can construct a planar k-terminal network $G' = (V', E', T', c')$ with all its terminals T' on the outerface such that if G' admits minor (q, s)-cut-sparsifier, then G admits a $(q, O(s))$-DAM.*

6.1. Proof of Theorem 6.1.

Construction of the reduction. The idea is to first use the duality of planar graphs in order to convert every minimum terminal cut into a shortest cycle, and then "open" every shortest cycle into a shortest path between two terminals, which in turn are preserved by a (q, s)-DAM. More formally, given a plane k-terminal network $G = (V, E, F, T, c)$ with all its terminals $T = \{t_1, \ldots, t_k\}$ on the outerface in a clockwise order, we first construct its dual graph G_1 where the boundaries of all its k terminal faces $T(G_1) = \{f_{t_1}, \ldots, f_{t_k}\}$ share the same vertex v_{f_∞}, and second we construct G_2 by the graph G_1 where the vertex v_{f_∞} is split into k different vertices $v_{f_\infty}^{i,i+1}$, and every edge (v_{f_∞}, v^*) that is embedded between (or on) the two terminal faces f_{t_i} and $f_{t_{i+1}}$ in G_1 corresponds to a new edge $(v_{f_\infty}^{i,i+1}, v^*)$ in G_2 with the same length. See Figure 5 from left to right for an illustration, and see Appendix A for basic notions of planar duality. In the following, $f + f'$ denotes a new face that is the union of two faces f and f':

$$V(G_2) := \big(V(G_1) \setminus \{v_{f_\infty}\}\big) \cup \{v_{f_\infty}^{1,2}, \ldots, v_{f_\infty}^{k-1,k}, v_{f_\infty}^{k,1}\},$$
$$E(G_2) := \big(E(G_1) \setminus \{(v_{f_\infty}, v^*) \; : \; v^* \in V(G_1)\}\big)$$
$$\cup \{(v_{f_\infty}^{i,i+1}, v^*) \; : \; i \in [k], (v_{f_\infty}, v^*) \in E(G_1), v^* \text{ between } f_{t_i}, f_{t_{i+1}}\}^3,$$
$$F(G_2) := \big(F(G_1) \setminus \{f_\infty, f_{t_1}, \ldots, f_{t_k}\}\big) \cup \{f_\infty + f_{t_1} + \cdots + f_{t_k}\},$$
$$T(G_2) := \{v_{f_\infty}^{1,2}, \ldots, v_{f_\infty}^{k-1,k}, v_{f_\infty}^{k,1}\}.$$

Let H_2 be a (q, s)-DAM of G_2. Since it is a minor of G_2, both are planar k-terminals network such that all their terminals are on their outerface in the same clockwise order. Hence, we can use H_2 and the same reduction above, but in reverse operations, in order to construct a $(q, O(s))$-cut-sparsifier H for G. First, we "close" all the shortest paths in H_2 into cycles by merging its k terminals $v_{f_\infty}^{i-1,i}$ into one vertex called v_{f_∞}, and we denote this new graph by H_1. Note that H_1 has k new faces $f_{t_1}, \ldots, f_{t_k}$, where each face f_{t_i} was created by uniting the two terminals $v_{f_\infty}^{i-1,i}, v_{f_\infty}^{i,i+1}$ of H_2. These k new faces of H_1 will be its k terminal faces. Second, we argue that the dual graph of H_1 is our requested cut-sparsifier of G, which we denote by H. See Figure 5 from right to left for an illustration.

Analysis of the reduction. The key element of the reduction's proof is the duality between every shortest cycle in G_1 to a shortest path in G_2, which we formally stated
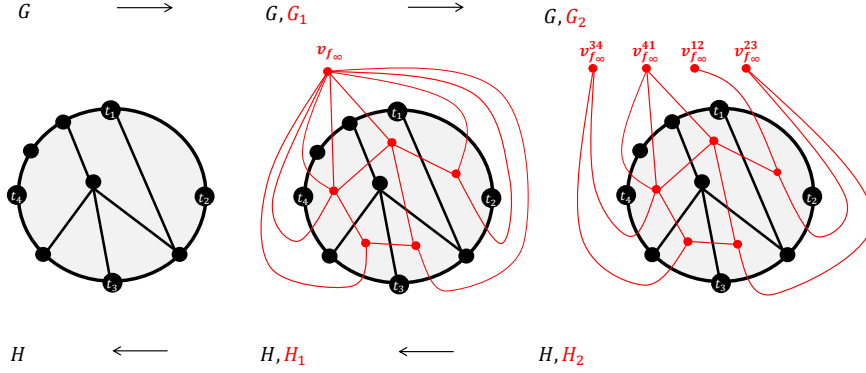
---

[3]We allow wraparound, i.e., $v^{k,k+1} = v^{k,1}$.

FIG. 5. *The first graph (in black) is the original graph G. The second is its dual graph G_1, in red. The third graph H and its new terminals are in red. We get that graph by "splitting" the outerface vertex v_{f_∞} to k new vertices, which are the new terminals.*

in the following lemma. Given G and its dual graph $G_1 = G^*$ as stated above, for every subsets of terminals $S \subset T(G)$ we denote by $S^* \subset T(G^*)$ the corresponding set of terminal faces, i.e., $S^* = \{f_{t_i} : \forall i \in [k] \text{ s.t. } t_i \in S\}$.

LEMMA 6.3. *Every shortest circuit that separates between the terminal faces $S^*_{i,j}$ and $\bar{S}^*_{i,j}$ in G_1 corresponds to a shortest path between the two terminals $v^{i-1,i}_{f_\infty}$ and $v^{j,j+1}_{f_\infty}$ in G_2, and vice versa.*

$$V(H_1) := \big(V(H_2) \setminus \{v^{1,2}_{f_\infty}, \ldots, v^{k-1,k}_{f_\infty}, v^{k,1}_{f_\infty}\}\big) \cup \{v_{f_\infty}\},$$
$$E(H_1) := \big(E(H_2) \setminus \{(v^{i,i+1}_{f_\infty}, v^*) : i \in [k]\}\big) \cup \{(v_{f_\infty}, v) : (v^{i,i+1}_{f_\infty}, v^*) \in E(H_2)\},$$
$$F(H_1) := \big(F(H_2) \setminus \{f_\infty + f_{t_1} + \cdots + f_{t_k}\}\big) \cup \{f_\infty, f_{t_1}, \ldots, f_{t_k}\},$$
$$T(H_1) := \{f_{t_1}, \ldots, f_{t_k}\}.$$

Proof. First direction—circuits to distances. Let $\mathcal{C}$ be a minimum circuit that separates between the terminal faces S^*_{ij} and $\bar{S}^*_{ij}$ in G_1 (assume without loss of generality $i \leq j$). By Theorem 2.5 that circuit is a union of l shortest cycles disjoint of each other for some $l \geq 1$. We prove that this circuit corresponds to a simple path in G_2 between the terminals $v^{i-1,i}_{f_\infty}$ and $v^{j,j+1}_{f_\infty}$ with the same weight using in an induction on l.

Induction base: $l = 1$. The circuit $\mathcal{C}$ contains exactly one simple cycle C that separates between the terminal faces S^*_{ij} and $\bar{S}^*_{ij} = S^*_{(j+1)(i-1)}$ in G_1. So the vertex v_{f_∞} appears in C exactly once, i.e., $C = \langle v_{f_\infty}, v_1, v_2, \ldots, v_x, v_{f_\infty} \rangle$. According to our construction, the graph G_2 contains the same vertices and edges as G_1, except for the vertex v_{f_∞} and all the edges incident to it. Therefore, $\langle v_1, v_2, \ldots, v_x \rangle$ is a simple path in G_2. Moreover, since without loss of generality the vertex v_1 is embedded between the terminal faces $f_{t_{i-1}}, f_{t_i}$ and the vertex v_x is embedded between the terminal faces $f_{t_j}, f_{t_{j+1}}$, we get that $(v^{i-1,i}_{f_\infty}, v_1)$, $(v_x, v^{j,j+1}_{f_\infty}) \in E(G_2)$. Thus, $\langle v^{i-1,i}_{f_\infty}, v_1, v_2, \ldots, v_x, v^{j,j+1}_{f_\infty} \rangle$ is a simple path in G_2 with the same weight as C.

Induction step. Assume that if $\mathcal{C}$ has $l' < l$ cycles, then it corresponds to a simple path in G_2 between the terminals $v^{i-1,i}_{f_\infty}$ and $v^{j,j+1}_{f_\infty}$ with the same weight, and prove it for l. There are two cases:

- If neither of the cycles in the circuit is nested, then without loss of generality all the cycles $C \in \mathcal{C}$ bound terminal faces of S^*_{ij}. Let $C \in \mathcal{C}$ be the cycle that

bounds the terminal faces $f_{t_i}, \ldots, f_{t_x}$ where $i < x < j$. Thus $\{C\}$ is a simple circuit that separates between the terminal faces S_{ix}^* to $\bar{S}_{ix}^* = S_{(x+1)(i-1)}^*$, and $\mathcal{C} \setminus \{C\}$ is a simple circuit that separates between the terminal faces $S_{(x+1)j}^*$ to $\bar{S}_{(x+1)j}^* = S_{(j+1)x}^*$ in G_1. By the inductive assumption these two circuits correspond to two simple paths in G_2 with the same weights. The first path is between the two terminals $v_{f_\infty}^{i-1,i}$ and $v_{f_\infty}^{x,x+1}$, and the second is between the two terminals $v_{f_\infty}^{x,x+1}$ and $v_{f_\infty}^{j,j+1}$, which form a simple path from $v_{f_\infty}^{i-1,i}$ to $v_{f_\infty}^{j,j+1}$ in G_2 with the same weight as $\mathcal{C}$.

- There are nested cycles in the circuit. Let $C \in \mathcal{C}$ be a simple cycle that separates between S_{xy}^* and $\bar{S}_{xy}^*$ in G_2 and contains at least one cycle of $\mathcal{C} \setminus \{C\}$. If $i < x \le y < j$ or $x < i \le j < y$, then $\mathcal{C} \setminus \{C\}$ separates between S_{ij}^* to $\bar{S}_{ij}^*$ in contradiction to the minimality of $\mathcal{C}$. Therefore either $i = x \le j < y$ or $i < x \le j = y$. Assume without loss of generality that the first case holds, i.e., C is a minimum circuit that separates between S_{iy}^* and $\bar{S}_{iy}^*$, and $\mathcal{C} \setminus \{C\}$ is a minimum circuit that separates between the terminal faces $S_{(y+1)j}^*$ to $\bar{S}_{(y+1)j}^*$ in G_1. By the inductive assumption these two circuits correspond to two simple paths in G_2 with the same weights. The first simple path is between the two terminals $v_{f_\infty}^{i-1,i}$ and $v_{f_\infty}^{y,y+1}$, and the second simple path is between the two terminals $v_{f_\infty}^{y,y+1}$ and $v_{f_\infty}^{j,j+1}$. Uniting these two paths forms a simple path between $v_{f_\infty}^{i-1,i}$ to $v_{f_\infty}^{j,j+1}$ in G_2 with the same weight as $\mathcal{C}$ as we required.

Second direction—distances to cuts. Let P be a shortest path between the terminals $v_{f_\infty}^{i-1,i}$ and $v_{f_\infty}^{j,j+1}$ in G_1 (assume $i \le j$), and let l be the number of terminals in that path (including the two terminals in its endpoints). It is easy to verify that replacing each terminal $v_{f_\infty}^{x,x+1}$ in P with the vertex v_{f_∞} transforms it into a circuit in G_1 with $l-1$ disjoint simple cycles and with the same weight of P. We prove that this circuit separates between the terminal faces $S_{ij}^* = \{f_{t_i}, \ldots, f_{t_j}\}$ and $\bar{S}_{ij}^*$ in G_1 by an induction on l.

Induction base: $l = 2$. That is, the only terminals on the path P are those on the endpoints. Thus, all the inner vertices on that path are nonterminal vertices, i.e., $P = \langle v_{f_\infty}^{i-1,i}, v_1, v_2, \ldots, v_x, v_{f_\infty}^{j,j+1}\rangle$. Substitute the terminals $v_{f_\infty}^{i-1,i}$ and $v_{f_\infty}^{j,j+1}$ of G_2 with the vertex v_{f_∞} of G_1 and get $C = \langle v_{f_\infty}, v_1, v_2, \ldots, v_x, v_{f_\infty}\rangle$. According to our construction, $\langle v_1, v_2, \ldots, v_x\rangle$ is a simple path in G_1, and $(v_{f_\infty}^{i-1,i}, v_1), (v_l, v_{f_\infty}^{j,j+1}) \in E(G_2)$ if and only if $(v_{f_\infty}, v_1), (v_x, v_{f_\infty}) \in E(G_1)$. Therefore, C is a simple cycle in G_1 and the two edges that incident to the vertex v_{f_∞} are embedded between the terminal faces $f_{t_{i-1}}$ to f_{t_i} and f_{t_j} to $f_{t_{j+1}}$ in G_1. Thus, C separates between S_{ij}^* to $\bar{S}_{ij}^*$ and has the same weight as P.

Induction step. Assume that if P has $l' < l$ inner terminals, then it corresponds to a simple circuit with l' cycles that separates between the terminal faces S_{ij}^* and $\bar{S}_{ij}^*$ in G_1, and prove it for $l' = l$. Let $v_{f_\infty}^{x,x+1}$ be some inner terminal in the path P that breaks it into two simple subpaths P_1 and P_2, i.e., P_1 is a simple path between $v_{f_\infty}^{i-1,i}$ to $v_{f_\infty}^{x,x+1}$ and P_2 is a simple path between $v_{f_\infty}^{x,x+1}$ to $v_{f_\infty}^{j,j+1}$ in G_2. Since both of these paths have less than l terminals we can use the inductive assumption and get that P_1 corresponds to a circuit $\mathcal{C}_1$ in G_1 with the same weight that separates between the terminals S_{ix}^* and $\bar{S}_{ix}^*$ and $\bar{S}_{ix}^*$, and P_2 corresponds to a circuit $\mathcal{C}_2$ in G_1 with the same weight that separates between the terminals $S_{(x+1)j}^*$ and $\bar{S}_{(x+1)j}^*$. If $i \le x \le j$, then $S_{ij}^* = S_{ix}^* \cup S_{(x+1)j}^*$. And if $i \le j < x$ (symmetric to the case were $x < i \le j$), then

$S^*_{(x+1)j} = S^*_{(j+1)x}$ and so $S^*_{ij} = S^*_{ix} \setminus S^*_{(j+1)x}$. In both cases we get that $\mathcal{C}_1 \cup \mathcal{C}_2$ is a simple circuit in G_1 with the same weight as P that separates between the terminal faces S^*_{ij} and $\bar{S}^*_{ij}$ in G_1, and the lemma follows. □

LEMMA 6.4. *The elementary cuts* $\mathcal{T}_e(G)$ *and* $\mathcal{T}_e(H)$ *are equal, and* $\mathrm{mincut}_G(S) \leq \mathrm{mincut}_H(S) \leq q \cdot \mathrm{mincut}_G(S)$ *for every* $S \in \mathcal{T}_e(G)$.

Proof. Let us call a shortest path between two terminals *elementary* if all the internal vertices on the path are Steiner, and denote by D_e all the terminal pairs that the shortest path between them is elementary. Moreover, recall that every elementary subset $S \in \mathcal{T}_e(G)$ is of the form $\{t_i, t_{i+1}, \ldots, t_j\}$, and denote it S_{ij} and $\bar{S}_{ij} = S_{(j+1)(i-1)}$ for simplicity.

By Lemma 6.3 a shortest circuit that separates between S^*_{ij} to $\bar{S}^*_{ij}$ in G_1 contains l elementary cycles if and only if a shortest path between the terminals $v^{i-1,i}_{f_\infty}$ and $v^{j,j+1}_{f_\infty}$ in G_2 contains $l+1$ terminals (including the endpoints). Notice that Lemma 6.3 holds also in the graphs H_2 and H_1, therefore $\mathcal{T}_e(G_1) = D_e(G_2)$ and $D_e(H_2) = \mathcal{T}_e(H_1)$. In addition, the equalities $\mathcal{T}_e(G) = \mathcal{T}_e(G_1)$ and $\mathcal{T}_e(H_1) = \mathcal{T}_e(H)$ hold by the duality between cuts and circuits, and $D_e(G_2) = D_e(H_2)$ because of the triangle inequality in the distance metric. Altogether we get that $\mathcal{T}_e(G) = \mathcal{T}_e(H)$.

Again by the duality between cuts and circuits and by Lemma 6.3 on the two pairs of graphs G, G_2 and H_2, H_1 we get that $\mathrm{mincut}_G(S_{ij}) = d_{G_2}(v^{i-1,i}_{f_\infty}, v^{j,j+1}_{f_\infty})$ and $\mathrm{mincut}_H(S_{ij}) = d_{H_2}(v^{i-1,i}_{f_\infty}, v^{j,j+1}_{f_\infty})$. Since H_2 is a (q,s)-DAM of G_2 we get that $\mathrm{mincut}_G(S_{ij}) \leq \mathrm{mincut}_H(S_{ij}) \leq q \cdot \mathrm{mincut}_G(S_{ij})$ and the lemma follows. □

LEMMA 6.5. *The size of* H *is* $O(s)$.

Proof. Given that H_2 is a (q,s)-DAM, i.e., $|V(H_2)| = s$, we need to prove that $|V(H)| = O(|V(H_2)|)$. Note that by the reduction construction $|V(H)| = |F(H_1)| = |F(H_2)| + k - 1$. Moreover, we can assume that H_2 is a simple planar graph (if it has parallel edges, we can keep the shortest one). Thus, $|E(H_2)| \leq 3|V(H_2)| + 6$. Plug it into Euler's formula to get $|F(H_2)| \leq 2|V(H_2)| + 8$. Since $s \geq k$ we derive that $|V(H)| \leq 2s + 8 + k - 1 = O(s)$ and the lemma follows. □

Proof of Theorem 6.1. Assume H_2 is a (q,s)-DAM of G_2 and let H be the graph that was constructed from H_2. By Lemmas 6.5 and 6.4 the graph H is a $(q, O(s))$-cut-sparsifier of G. Since H_2 is a minor of G_2, and a minor operation is closed under planar duality, H is furthermore a minor of G and the theorem follows. □

6.2. Proof of Theorem 6.2.

Construction of the reduction. The idea is to first "close" the shortest paths between every two terminals into shortest cycles that separate between terminal faces, and then use the planar duality between cuts and cycles to get that every shortest cycle corresponds to a minimum terminal cut that in turn is preserved by a (q,s)-cut-sparsifier. More formally, assume a plane k-terminal network $G = (V, E, F, T, c)$ with all its terminals $T = \{t_1, \ldots, t_k\}$ on the outerface in a clockwise order. First, construct a graph G_1 by adding to G a new vertex v_{f_∞} and connect it to all its k terminals t_i using edges with 0 capacity. Note that G_1 has k new faces $f_{1,2}, \ldots, f_{k-1,k}, f_{k,1}$, where each $f_{i,i+1}$ was created by adding the two new edges (v_{f_∞}, t_i) and (v_{f_∞}, t_{i+1}). These k new faces will be the terminals of G_1.

$$V(G_1) := V \cup \{v_{f_\infty}\},$$
$$E(G_1) := E \cup \{(v_{f_\infty}, t_i) \ : \ t_i \in T\},$$

$$F(G_1) := F \cup \{f_{1,2}, \ldots, f_{k-1,k}, f_{k,1}\},$$
$$T(G_1) := \{f_{1,2}, \ldots, f_{k-1,k}, f_{k,1}\}.$$

Second, we denote by G_2 the dual graph of G_1, where its k terminals are $T(G_2) = \{v_{i,i+1} : f_{i,i+1} \in T(G_1)\}$. Moreover, the new vertex v_{f_∞} in G_1 corresponds to the outerface f_∞ of G_2, the k new edges (v_{f_∞}, t_i) we added to G_1 are the edges that lie on the outerface of G_2, and the vertices on the outerface of G_2 are the k terminals $v_{i,i+1}$ in a clockwise order. See Figure 5 from left to right for an illustration, and see Appendix A for basic notions of planar duality.

Let H_2 be a (q, s)-cut-sparsifier and a minor of G_2. Since H_2 is a minor of G_2, then both are plane graphs with all their terminals on the outerface in the same clockwise order, and there is an edge with capacity 0 on the outerface that connects between every two adjacent terminals. Hence, we can use H_2 and the same reduction above (but in opposite order of operations) in order to construct a $(q, O(s))$-DAM H of G as follows. First, let H_1 be the dual graph of H_2, where every minimum terminal cut in H_2 is equivalent to a shortest cycle that separates terminal faces. Notice that again each terminal face $f_{i,i+1}$ in H_1 contains the two edges (v_{f_∞}, t_i) and (v_{f_∞}, t_{i+1}) with capacity 0 on their boundary. Second, we "open" each shortest cycle in H'^* into a shortest path between terminals by removing the vertex v_{f_∞} and all its incidence edges and denote this new graph by H. The terminals of H are all the vertices $v \in V(H_1)$ such that (v_{f_∞}, v) is an edge in H_1, which are equal to the original terminals of G. See Figure 6 from right to left for an illustration.

$$V(H) := V(H_1) \setminus \{v_{f_\infty}\},$$
$$E(H) := E(H_1) \setminus \{(v_{f_\infty}, t_i) : t_i \in T\},$$
$$F(H) := F(H_1) \setminus \{f_{1,2}, \ldots, f_{k-1,k}, f_{k,1}\},$$
$$T(H) := T(G).$$

Analysis of the reduction.
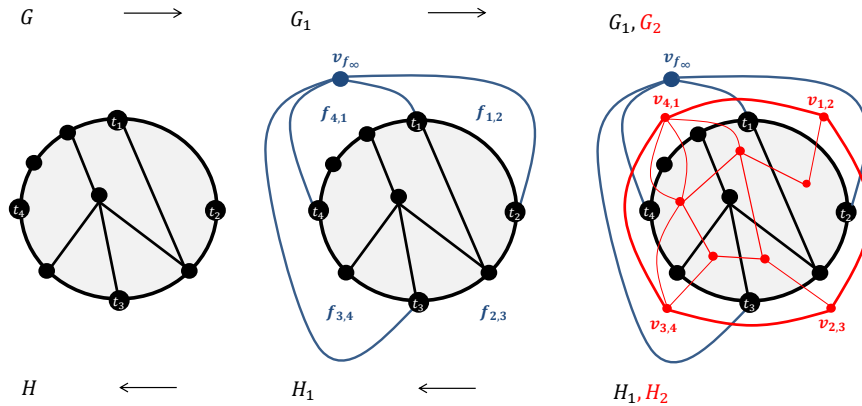
LEMMA 6.6. *The size of H is $O(s)$.*

FIG. 6. *The first graph (in black) is the original graph G. The second is the graph G_1, where the additional vertex and edges and terminal faces are in blue. The third graph G_2 and its terminals are in red. The bold red edges are the dual of the blue edges, and both of them have capacity 0.*

Proof. Given that H_2 is a (q, s)-cut-sparsifier, i.e., $|V(H_2)| = s$, we will prove that $|V(H)| = O(|V(H_2)|)$. We can assume that H_2 is a simple planar graph (if not, we can replace all the parallel edges between every two vertices by one edge where its capacity is the sum over all the capacities of these parallel edges), thus $|E(H_2)| \geq \frac{3}{2}|F(H_2)|$. Plug it into Euler's formula to get $|F(H_2)| \leq 2|V(H_2)| - 4 = 2s - 4$. By the reduction construction $|V(H)| + 1 = |V(H_1)| = |F(H_2)| = O(s)$, and the lemma follows. □

LEMMA 6.7. *The graph H is a minor of G.*

Proof. Given that H_2 is a minor of G_2, and that minor is close under deletion and contraction of edges, we get that H_1 is a minor of G_1. Now by deleting the same vertex v_{f_∞} together with all its incidence edges from both G_1 and H_1, we get the graphs G and H correspondingly. Therefore H is a minor of G, and the lemma follows. □

LEMMA 6.8. *The graph H preserves all the distances between every two terminals by factor q, i.e., $d_G(t_i, t_j) \leq d_H(t_i, t_j) \leq q \cdot d_G(t_i, t_j)$ for every $t_i, t_j \in T$.*

Proof. Notice that connecting all the terminals to a new vertex using edges with capacity 0 is equivalent to uniting all the terminals into one vertex, and also splitting the vertex v_{f_∞} to k new terminals is equivalent to disconnecting all the terminals by deleting that vertex. Thus our reduction is equivalent to the reduction of Theorem 6.1. In particular, Lemma 6.3 holds on the graphs G, G_1 and on the graphs H, H_1 correspondingly, i.e., every shortest path between two terminals t_i and t_j in G (or H) corresponds to a minimum circuit in G_1 (or H_1) that separates between the terminal faces $\{f_{i,i+1}, \ldots, f_{j-1,j}\}$ to $\{f_{j,j+1}, \ldots, f_{i-1,i}\}$ and vice versa.

Let $S_{(i,i+1),(j-1,j)} = \{v_{i,i+1}, \ldots, v_{j-1,j}\}$ be a set of terminals in G_2, where every terminal $v_{l,l+1}$ corresponds to the terminal face $f_{l,l+1}$ in G_1. By the duality between cuts and circuits we get that $d_G(t_i, t_j) = \text{mincut}_{G_2}(S_{(i,i+1),(j-1,j)})$ and $d_H(t_i, t_j) = \text{mincut}_{H_2}(S_{(i,i+1),(j-1,j)})$. Since H_2 is a (q, s)-cut-sparsifier of G_2 we derive the inequalities $d_G(t_i, t_j) \leq d_H(t_i, t_j) \leq q \cdot d_G(t_i, t_j)$ and the lemma follows. □

Proof of Theorem 6.2. By Lemmas 6.6, 6.7, and 6.8 the k-terminal network H is a $(q, O(s))$-DAM of G and the theorem follows. □

6.3. Duality applications. By Theorems 6.1 and 6.2, every k-terminal network G with $\gamma = 1$ admits a (q, s)-DAM if and only if it admits a minor $(q, O(s))$-cut-sparsifier. Hence, every new upper or lower bound result, especially for $q > 1$, on DAM also holds for the minor cut-sparsifier problem and vice versa. For example, the upper bound of $(1 + \epsilon, (k/\epsilon)^2)$-DAM for planar networks [13] yields the following new theorem.

THEOREM 6.9. *Every planar network G with $\gamma = 1$ admits a minor $(1 + \epsilon, \tilde{O}((k/\epsilon)^2)$-cut-sparsifier for every $\epsilon > 0$.*

As already mentioned, by recent independent work [20] these networks also admit a $(1, O(k^2))$-sparsifier that is planar but *not* a minor of G.

In addition, we can apply known upper and lower bounds for $(1, s)$-DAM to the minor mimicking network problem (i.e., a cut-sparsifier of quality 1). In particular, the known $(1, k^4)$-DAM [28] yields an alternative proof for Corollary 4.3, and the known lower bound of $(1, \Omega(k^2))$-DAM (which is shown on grid graphs) [28] yields an alternative proof for a lower bound shown in [29].

Appendix A. Planar duality.   Using planar duality we bound the size of mimicking networks for planar graphs (Theorem 3.1), and we further use it to prove

the duality between cuts in distances (Theorems 6.1 and 6.2) Recall that every planar graph G has a dual graph G^*, whose vertices correspond to the faces of G and whose faces correspond to the vertices of G, i.e., $V(G^*) = \{v_f^* : f \in F(G)\}$ and $F(G^*) = \{f_v^* : v \in V(G)\}$. Thus the terminals $T = \{t_1, \ldots, t_k\}$ of G correspond to the terminal faces $T(G_1) = \{f_{t_1}, \ldots, f_{t_k}\}$ in G^*, which for the sake of simplicity we may refer to as terminals as well. Every edge $e = (v, u) \in E(G)$ with capacity $c(e)$ that lies on the boundary of two faces $f_1, f_2 \in F(G)$ has a dual edge $e^* = (v_{f_1}^*, v_{f_2}^*) \in E(G^*)$ with the same capacity $c(e^*) = c(e)$ that lies on the boundary of the faces f_v^* and f_u^*. For every subset of edges $M \subset E(G)$, let $M^* := \{e^* : e \in M\}$ denote the subset of the corresponding dual edges in G^*.

The following theorem describes the duality between two different kinds of edge sets—minimum cuts and minimum circuits—in a plane multigraph. It is a straight-forward generalization of the case of st-cuts (whose dual are cycles) to three or more terminals.

A *circuit* is a collection of cycles (not necessarily disjoint) $\mathcal{C} = \{C_1, \ldots, C_l\}$. Let $\mathcal{E}(\mathcal{C}) = \cup_{i=1}^l C_i$ be the set of edges that participate in one or more cycles in the collection (note it is not a multiset, so we discard multiplicities). The capacity of a circuit $\mathcal{C}$ is defined as $\sum_{e \in \mathcal{E}(\mathcal{C})} c(e)$.

THEOREM A.1 (duality between cutsets and circuits). *Let G be a connected plane multigraph, let G^* be its dual graph, and fix a subset of the vertices $W \subseteq V(G)$. Then, $M \subset E(G)$ is a cutset in G that has minimum capacity among those separating W from $V(G) \backslash W$ if and only if the dual set of edges $M^* \subseteq E(G^*)$ is actually $\mathcal{E}(\mathcal{C})$ for a circuit $\mathcal{C}$ in G^* that has minimum capacity among those separating the corresponding faces $\{f_v^* : v \in W\}$ from $\{f_v^* : v \in V(G) \setminus W\}$.*

LEMMA A.2 (the dual of a connected component). *Let G be a connected plane multigraph, let G^* be its dual, and fix a subset of edges $M \subset E(G)$. Then $W \subseteq V$ is a connected component in $G \setminus M$ if and only if its dual set of faces $\{f_v^* : v \in W\}$ is a face of $G^*[M^*]$.*

Fix $S \subset T$. We call E_S^* elementary circuit if E_S is an elementary cutset in G. Note that by Lemma A.2 E_S^* is an elementary circuit if and only if the graph $G^* \setminus E_S^*$ has exactly two faces. Thus $\mathcal{T}_e(G) = \mathcal{T}_e(G^*)$, and the circuit E_S^* has exactly one minimum cycle in G^*. For the sake of simplicity we later on use the term cycle instead of circuit when we refer to an elementary minimum circuit.

REFERENCES

[1] A. ANDONI, A. GUPTA, AND R. KRAUTHGAMER, *Towards $(1 + \epsilon)$-approximate flow sparsifiers*, in Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms, 2014, pp. 279–293, https://doi.org/10.1137/1.9781611973402.20.

[2] A. BASU AND A. GUPTA, *Steiner Point Removal in Graph Metrics*, manuscript, 2008, http://www.math.ucdavis.edu/~abasu/papers/SPR.pdf.

[3] A. A. BENCZÚR AND D. R. KARGER, *Randomized approximation schemes for cuts and flows in capacitated graphs*, SIAM J. Comput., 44 (2015), pp. 290–319, https://doi.org/10.1137/070705970.

[4] C. BENTZ, *A simple algorithm for multicuts in planar graphs with outer terminals*, Discrete Appl. Math., 157 (2009), pp. 1959–1964, https://doi.org/10.1016/j.dam.2008.11.010.

[5] D. BIENSTOCK AND C. L. MONMA, *On the complexity of covering vertices by faces in a planar graph*, SIAM J. Comput., 17 (1988), pp. 53–76, https://doi.org/10.1137/0217004.

[6] H. T. Chan, D. Xia, G. Konjevod, and A. W. Richa, *A tight lower bound for the Steiner point removal problem on trees*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, Lecture Notes in Comput. Sci. 4110, Springer, New York, 2006, pp. 70–81, https://doi.org/10.1007/11830924_9.

[7] M. Charikar, T. Leighton, S. Li, and A. Moitra, *Vertex sparsifiers and abstract rounding algorithms*, in 51st Annual Symposium on Foundations of Computer Science, 2010, pp. 265–274, https://doi.org/10.1109/FOCS.2010.32.

[8] S. Chaudhuri, K. V. Subrahmanyam, F. Wagner, and C. D. Zaroliagis, *Computing mimicking networks*, Algorithmica, 26 (2000), pp. 31–49, https://doi.org/10.1007/s004539910003.

[9] C. Chekuri, F. B. Shepherd, and C. Weibel, *Flow-cut gaps for integer and fractional multiflows*, J. Combin. Theory Ser. B, 103 (2013), pp. 248–273, https://doi.org/10.1016/j.jctb.2012.11.002.

[10] D. Z. Chen and X. Wu, *Efficient algorithms for k-terminal cuts on planar graphs*, Algorithmica, 38 (2004), pp. 299–316, https://doi.org/10.1007/s00453-003-1061-2.

[11] D. Z. Chen and J. Xu, *Shortest path queries in planar graphs*, in 32nd Annual ACM Symposium on Theory of Computing, 2000, pp. 469–478, https://doi.org/10.1145/335305.335359.

[12] Y. K. Cheung, *Steiner point removal—distant terminals don't (really) bother*, in Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, 2018, pp. 1353–1360, https://doi.org/10.1137/1.9781611975031.89.

[13] Y. K. Cheung, G. Goranci, and M. Henzinger, *Graph minors for preserving terminal distances approximately—lower and upper bounds*, in 43rd International Colloquium on Automata, Languages, and Programming, 2016, pp. 131:1–131:14, https://doi.org/10.4230/LIPIcs.ICALP.2016.131.

[14] J. Chuzhoy, *On vertex sparsifiers with Steiner nodes*, in Proceedings of the 44th Symposium on Theory of Computing Conference, 2012, pp. 673–688, https://doi.org/10.1145/2213977.2214039.

[15] M. Englert, A. Gupta, R. Krauthgamer, H. Räcke, I. Talgam-Cohen, and K. Talwar, *Vertex sparsifiers: New results from old techniques*, SIAM J. Comput., 43 (2014), pp. 1239–1262, https://doi.org/10.1137/130908440.

[16] A. Filtser, *Steiner point removal with distortion* O(log k), in Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, 2018, pp. 1361–1373, https://doi.org/10.1137/1.9781611975031.90.

[17] A. Filtser, *A Face Cover Perspective to ℓ_1 Embeddings of Planar Graphs*, https://arxiv.org/abs/1903.02758, 2019.

[18] A. Filtser, R. Krauthgamer, and O. Trabelsi, *Relaxed Voronoi: A simple framework for terminal-clustering problems*, in Proceedings of the 2nd Symposium on Simplicity in Algorithms, 2019, pp. 10:1–10:14, https://doi.org/10.4230/OASIcs.SOSA.2019.10.

[19] G. N. Frederickson, *Planar graph decomposition and all pairs shortest paths*, J. ACM, 38 (1991), pp. 162–204, https://doi.org/10.1145/102782.102788.

[20] G. Goranci, M. Henzinger, and P. Peng, *Improved guarantees for vertex sparsification in planar graphs*, in Proceedings of the 25th Annual European Symposium on Algorithms, LIPIcs. Leibniz Int. Proc. Inform. 87, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2017, pp. 44:1–44:14, https://doi.org/10.4230/LIPIcs.ESA.2017.44.

[21] G. Goranci and H. Räcke, *Vertex sparsification in trees*, in Approximation and Online Algorithms—14th International Workshop, 2016, pp. 103–115, https://doi.org/10.1007/978-3-319-51741-4_9.

[22] A. Gupta, *Steiner points in tree metrics don't (really) help*, in Proceedings of the 12th Annual Symposium on Discrete Algorithms, 2001, pp. 220–227, http://dl.acm.org/citation.cfm?id=365411.365448.

[23] T. Hagerup, J. Katajainen, N. Nishimura, and P. Ragde, *Characterizing multiterminal flow networks and computing flows in networks of small treewidth.*, J. Comput. System Sci., 57 (1998), pp. 366–375, https://doi.org/10.1006/jcss.1998.1592.

[24] L. Kamma, R. Krauthgamer, and H. L. Nguyen, *Cutting corners cheaply, or how to remove Steiner points*, SIAM J. Comput., 44 (2015), pp. 975–995, https://doi.org/10.1137/140951382.

[25] N. Karpov, M. Pilipczuk, and A. Zych-Pawlewicz, *An exponential lower bound for cut sparsifiers in planar graphs*, in Proceedings of the 12th International Symposium on Parameterized and Exact Computation, 2017, pp. 24:1–24:11, https://doi.org/10.4230/LIPIcs.IPEC.2017.24.

[26] A. Khan and P. Raghavendra, *On mimicking networks representing minimum terminal cuts*, Inform. Process. Lett., 114 (2014), pp. 365–371, https://doi.org/10.1016/j.ipl.2014.02.011.

[27] R. Krauthgamer, J. R. Lee, and H. Rika, *Flow-cut gaps and face covers in planar graphs*, in Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, 2019, pp. 525–534, https://doi.org/10.1137/1.9781611975482.33.

[28] R. Krauthgamer, H. L. Nguyen, and T. Zondiner, *Preserving terminal distances using minors*, SIAM J. Discrete Math., 28 (2014), pp. 127–141, https://doi.org/10.1137/120888843.

[29] R. Krauthgamer and I. Rika, *Mimicking networks and succinct representations of terminal cuts*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 1789–1799, https://doi.org/10.1137/1.9781611973105.128.

[30] J. R. Lee and A. Sidiropoulos, *On the geometry of graphs with a forbidden minor*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing, 2009, pp. 245–254, https://doi.org/10.1145/1536414.1536450.

[31] F. T. Leighton and A. Moitra, *Extensions and limits to vertex sparsification*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, 2010, pp. 47–56, https://doi.org/10.1145/1806689.1806698.

[32] K. Makarychev and Y. Makarychev, *Metric extension operators, vertex sparsifiers and Lipschitz extendability*, Israel J. Math., 212 (2016), pp. 913–959, https://doi.org/10.1007/s11856-016-1315-8.

[33] K. Matsumoto, T. Nishizeki, and N. Saito, *An efficient algorithm for finding multicommodity flows in planar networks*, SIAM J. Comput., 14 (1985), pp. 289–302, https://doi.org/10.1137/0214023.

[34] A. Moitra, *Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size*, in Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, 2009, pp. 3–12, https://doi.org/10.1109/FOCS.2009.28.

[35] H. Okamura and P. Seymour, *Multicommodity flows in planar graphs*, J. Combin. Theory Ser. B, 31 (1981), pp. 75–81, https://doi.org/10.1016/S0095-8956(81)80012-3.

[36] D. Peleg and J. D. Ullman, *An optimal synchronizer for the hypercube*, SIAM J. Comput., 18 (1989), pp. 740–747, https://doi.org/10.1137/0218050.

[37] H. Räcke, *Optimal hierarchical decompositions for congestion minimization in networks*, in Proceedings of the 40th Annual ACM Symposium on Theory of Computing, 2008, pp. 255–264, https://doi.org/10.1145/1374376.1374415.

[38] D. A. Spielman and S.-H. Teng, *Spectral sparsification of graphs*, SIAM J. Comput., 40 (2011), pp. 981–1025, https://doi.org/10.1137/08074489X.