SIAM J. DISCRETE MATH.
Vol. 33, No. 3, pp. 1482–1502

ON MERGINGS IN ACYCLIC DIRECTED GRAPHS*

GUANGYUE HAN†

Abstract. Consider an acyclic directed graph G with sources $s_1, s_2, \ldots, s_n$ and sinks $r_1, r_2, \ldots, r_n$. For $i = 1, 2, \ldots, n$, let c_i denote the size of the minimum edge cut between s_i and r_i, which, by Menger's theorem, implies that there exists a group of c_i edge-disjoint paths from s_i to r_i. Although they are edge disjoint within the same group, the above-mentioned edge-disjoint paths from different groups may merge with each other (or, roughly speaking, share a common subpath). In this paper we show that by choosing these paths appropriately, the number of mergings among all these edge-disjoint paths is always bounded by a finite function $\mathcal{M}(c_1, c_2, \ldots, c_n)$, which is independent of the size of G. Moreover, we prove some elementary properties of $\mathcal{M}(c_1, c_2, \ldots, c_n)$, derive exact values of $\mathcal{M}(1, c)$ and $\mathcal{M}(2, c)$, and establish a scaling law of $\mathcal{M}(c_1, c_2)$ when one of the parameters is fixed.

Key words. merging, network flow theory, graph structure

AMS subject classification. 05C21

DOI. 10.1137/17M1150633

1. Introduction.

Let $G = (E, V)$ be an acyclic directed graph with the edge set E and vertex set V. In this paper, an edge $e \in E$ linking a vertex $v_1 \in V$ to another vertex $v_2 \in V$ will be represented by (v_1, v_2) and, more generally, a directed path β consisting of vertices $v_1, v_2, \ldots, v_{\ell+1}$, ordered according to the direction of the path, will be represented by $(v_1, v_2, \ldots, v_{\ell+1})$, where each v_j will be referred to as the *direct predecessor* of v_{j+1} on β (see Figure 1(a) for a quick example). We say $m \geq 2$ directed paths $\beta_1, \beta_2, \ldots, \beta_m$ *merge* at an edge $e = (u, v) \in E$ if (1) e belongs to all β_j and (2) there exist $j_1 \neq j_2$ such that the direct predecessors of u on β_{j_1} and β_{j_2} are distinct; see Figure 1 for some illustrative examples.

We are primarily interested in the case that G has n distinct sources $s_1, s_2, \ldots, s_n$, n distinct sinks $r_1, r_2, \ldots, r_n$, and for each i, the size of the minimum edge cut between s_i and r_i is c_i, which, by Menger's theorem [9], implies the existence of $\alpha_i = \{\alpha_{i,1}, \alpha_{i,2}, \ldots, \alpha_{i,c_i}\}$, a set of edge-disjoint paths from s_i to r_i. Throughout the paper, for each feasible i, an element in α_i will be referred to as an α_i*-path*, and an edge on some α_i-path will be referred to as an α_i*-edge*, and the set of all α_i-edges will be denoted by $E(\alpha_i)$. For any edge e, let $\alpha_i(e)$ denote the α_i-path passing through e (note that $\alpha_i(e)$ is well-defined if and only if e is an α_i-edge), and let $\alpha(e) = \{\alpha_1(e), \alpha_2(e), \ldots, \alpha_n(e)\}$. Let $\mathcal{G}(c_1, c_2, \ldots, c_n)$ denote the set of all such G.

Apparently, an α_{i_1}-path merges with an α_{i_2}-path only if $i_1 \neq i_2$. An edge $e \in E$ is said to be a *merging* with respect to $\alpha_1, \alpha_2, \ldots, \alpha_n$ if there exist $i_1 \neq i_2$ such that some α_{i_1}-path and some α_{i_2}-path merge at e. Let $M(G; \alpha_1, \alpha_2, \ldots, \alpha_n)$ denote the number of mergings in G with respect to $\alpha_1, \alpha_2, \ldots, \alpha_n$.

Noting that the choices of each α_i may not be unique, we let $\Lambda_i(G)$ denote the set of all possible α_i. The main result of this paper is that one can choose $\alpha_i \in \Lambda_i(G)$ for

*Received by the editors October 4, 2017; accepted for publication (in revised form) May 14, 2019; published electronically August 22, 2019. An abstract version [8] of this work has been presented at the 2009 IEEE Information Theory Workshop, Volos, Greece, IEEE, Piscataway, NJ, 2009, pp. 271–275.

https://doi.org/10.1137/17M1150633

Funding: This work is supported by the Research Grants Council of the Hong Kong Special Administrative Region, China, under Project 17301017 and by the National Natural Science Foundation of China, under Project 61871343.

†The University of Hong Kong, Pokfulam, Hong Kong (ghan@hku.hk).

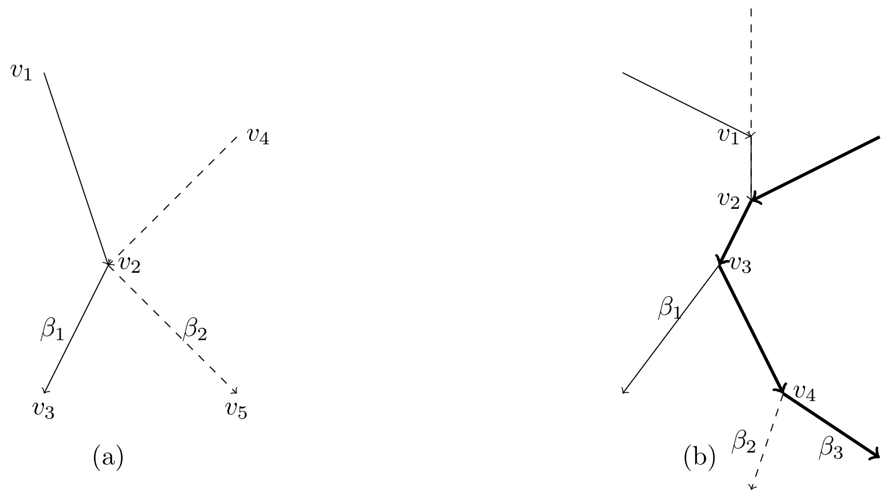

FIG. 1. *In* (a)*, the path* $\beta_1 = (v_1, v_2, v_3)$ *consists of two concatenated edges:* $e_1 = (v_1, v_2)$ *and* $e_2 = (v_2, v_3)$*. Note that though paths* β_1 *and* β_2 *share the vertex* v_2*, they do not share any edges, so* β_1 *and* β_2 *do not merge. In* (b)*,* β_1 *and* β_2 *merge at the edge* (v_1, v_2)*, however, not at* (v_2, v_3)*;* β_1*, and* β_3 *merge at the edge* (v_2, v_3)*;* β_2*, and* β_3 *merge at the edge* (v_2, v_3)*;* β_1*,* β_2*, and* β_3 *merge at the edge* (v_2, v_3)*.*

all feasible i such that $M(G; \alpha_1, \alpha_2, \dots, \alpha_n)$ is upper bounded by a finite function of $\alpha_1, \alpha_2, \dots, \alpha_n$, which is independent of the size of the graph G. Here we remark that for any fixed i, the the Edmonds–Karp algorithm [5] (an efficient implementation of the classical Ford–Fulkerson method [6]) can find a minimum edge cut and a set of edge-disjoint paths between s_i and r_i in polynomial time. On the other hand though, the fact that the *link disjoint problem*, which essentially asks if there are two edge-disjoint paths from s_i, s_j to r_i, r_j for any $i \neq j$, is NP-complete [7] suggests the intricacy of the scenarios where multiple pairs of sources and sinks are involved.

The following definition introduces some fundamental notions to be examined in this paper.

DEFINITION 1.1. *For any* $G \in \mathcal{G}(c_1, c_2, \dots, c_n)$*, we define*

$$M(G) \triangleq \min_{\alpha_i \in \Lambda_i(G):\ i=1,2,\dots,n} M(G; \alpha_1, \alpha_2, \dots, \alpha_n)$$

and, furthermore, for any $c_1, c_2, \dots, c_n$*, we define*

$$\mathcal{M}(c_1, c_2, \dots, c_n) = \sup_{G \in \mathcal{G}(c_1, c_2, \dots, c_n)} M(G).$$

Colloquially, for the purpose of minimizing the number of mergings by appropriately choosing all α_i-paths, $M(G)$ is the result of our best decision for a given graph G, and $\mathcal{M}(c_1, c_2, \dots, c_n)$ is the worst performance of our best decisions among all possible graphs $G \in \mathcal{G}(c_1, c_2, \dots, c_n)$.

Example 1.2. It can be easily verified that the graph G in Figure 2 belongs to $\mathcal{G}(2, 2)$. Let

$$\alpha_1 = \{\alpha_{1,1}, \alpha_{1,2}\} = \{(s_1, v_1, v_2, v_3, v_4, r_1), (s_1, v_5, v_6, v_7, v_8, r_1)\},$$

$$\alpha_2 = \{\alpha_{2,1}, \alpha_{2,2}\} = \{(s_2, v_1, v_2, v_7, v_8, r_2), (s_2, v_5, v_6, v_3, v_4, r_2)\}.$$

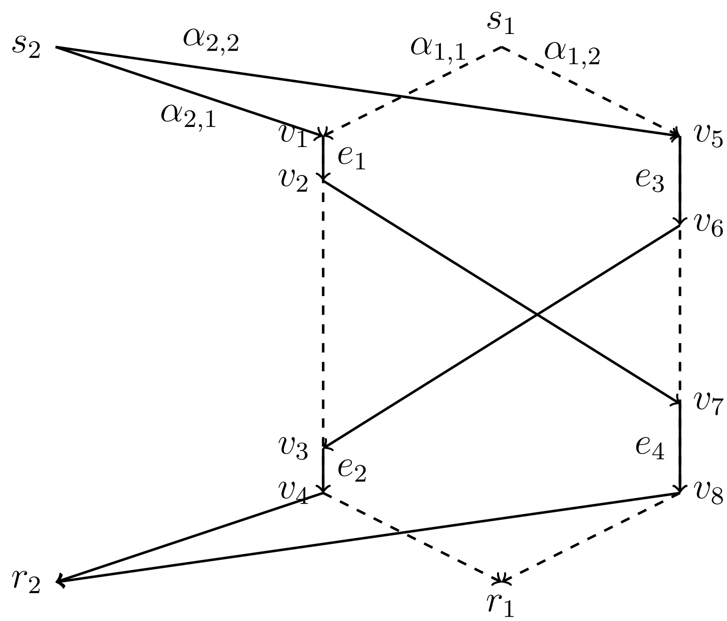

FIG. 2. *An illustrative example.*

Apparently, e_1, e_2, e_3, e_4 are the only 4 mergings with respect to α_1, α_2 and so

$$M(G; \alpha_1, \alpha_2) = 4.$$

Now, with the following alternative set of two edge-disjoint paths from s_1 to r_1,

$$\alpha_1' = \{\alpha_{1,1}', \alpha_{1,2}'\} = \{(s_1, v_1, v_2, v_7, v_8, r_1), (s_1, v_5, v_6, v_3, v_4, r_1)\},$$

one verifies that e_1, e_3 are the only 2 mergings with respect to α_1', α_2 and, thereby,

$$M(G; \alpha_1', \alpha_2) = 2$$

and, furthermore, the choice of α_1' and α_2 achieves $M(G) = 2$. Note that it will be established in Theorem 4.4 that $\mathcal{M}(2,2) = 5$, that is to say, for any graph in $\mathcal{G}(2,2)$, there is a way to choose α_1, α_2 such that the number of mergings in the graph with respect to α_1, α_2 is at most 5, which can be achieved by some graph in $\mathcal{G}(2,2)$.

At first glance, $\mathcal{M}(c_1, c_2, \ldots, c_n)$ may be infinite. However, the following theorem, which is the main result in this paper, asserts its finiteness.

THEOREM 1.3. *For any $c_1, c_2, \ldots, c_n$, we have*

$$\mathcal{M}(c_1, c_2, \ldots, c_n) \ < \ \infty.$$

Theorem 1.3 follows from Proposition 3.4, which establishes the finiteness of $\mathcal{M}$ with two parameters via an explicit upper bound, and Proposition 3.5, which upper bounds $\mathcal{M}$ with multiple parameters using a sum of $\mathcal{M}$ with two parameters.

Remark 1.4. Theorem 1.3 does not hold for cyclic directed graphs: For the cyclic directed graph G in Figure 3, $\alpha_{2,1}$ merges with $\alpha_{1,2}$ at e_1, e_2, e_3, e_4, e_5 and it can be verified that $M(G) = 5$. Moreover, in a similar fashion, one can construct a graph G' such that $\alpha_{2,1}$ merges with $\alpha_{1,2}$ at $e_1, e_2, \ldots, e_\ell$ with $M(G') = \ell$ for an arbitrary ℓ, which means that $\mathcal{M}(2,2)$, if defined on cyclic directed graphs, is in fact infinity.

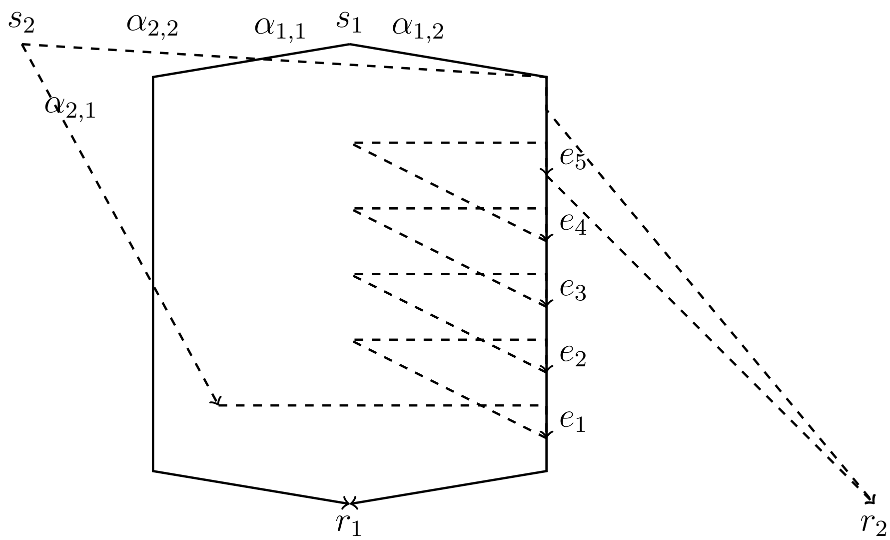

FIG. 3. *An illustrative example.*

Theorem 1.3 has a number of variants. For one example, the definition of $\mathcal{M}$ carries over almost verbatim to the cases that some of the sources and/or some of the sinks are identical, for which Theorem 1.3 still holds true. For another example, we note that (I) each c_i can be redefined as the size of a mininum vertex cut between s_i and r_i; (II) each α_i can be redefined as a set of c_i vertex-disjoint paths; and (III) the notion of merging can be redefined as follows: we say $m \geq 2$ directed paths $\beta_1, \beta_2, \ldots, \beta_m$ *merge* at $v \in V$ if (1) v belong to all β_j and (2) there exist $j_1 \neq j_2$ such that the direct predecessors of v on β_{j_1} and β_{j_2} are distinct. It can be easily verified that with the above-mentioned new definitions in place, the definition of $\mathcal{M}$ carries over in a straightforward fashion and Theorem 1.3 still holds true.

As elaborated below, Theorem 1.3 and its variants can be of use in certain practical situations where a network features multiple sources and/or multiple sinks.

In particular, the case that all r_i are identical is of relevance to transportation networks. More specifically, consider the traffic in a monocentric city during the morning rush hour (see, e.g., [11, 1, 2]), where a very large number of commuters travel from home (presumably in the suburban areas) to the workplace (presumably in the downtown area). Apparently, the bulk of the suburban traffic (i.e., traffic *outside* the downtown area) in such a situation is "largely" loopless: before reaching the downtown area, it makes no sense for any traveler to first go substantially further away from the workplace in order to eventually get to work, or take a circuitous detour in any part of his/her journey. As a result, we can assume, through contracting the downtown area to a single destination and orientating the occupied roads, that the suburban morning traffic is acyclic with multiple origins $s_1, s_2, \ldots, s_n$ and one single destination r. Assume the travel demand from s_i to r is c_i, which in turn demands traffic assignments on a set of c_i edge-disjoint paths of unit capacity from s_i to r. Then, Theorem 1.3 implies that under the optimal routing strategy, the minimum number of traffic mergings is always upper bounded by $\mathcal{M}(c_1, c_2, \ldots, c_n)$, which is independent of the size of the underlying transportation network. Generally speaking, noting that traffic mergings naturally give rise to congestions in transportation networks [3, 10], we expect that Theorem 1.3 can lead to an enhanced understanding on the level of traffic congestion in transportation networks under system-optimum route choices.

In constrast, the case that all s_i are identical is of relevance to communication networks. Below, we briefly mention the connection between our work and the theory of network coding [12]: Let G be an acyclic directed communication network with one sender s and n receivers $r_1, r_2, \ldots, r_n$. Again, let c_i denote the size of the minimum edge cut between s and r_i. It has been shown that if network coding is employed at some intermediate nodes, then information can be simultaneously transmitted from s to r_i at full rate c_i for all i. The so-called *network encoding complexity* refers to the least number of encoding nodes required for a feasible network coding scheme (see [8] and references therein). It turns out for the above-mentioned network, its network encoding complexity is always bounded by $\mathcal{M}(c_1, c_2, \ldots, c_n)$.

Here we remark that an independent work [4] has conducted an in-depth analysis on routings on vertex-disjoint paths for graphs of a given treewidth, which has implications to relevant aspects of graph theory and computational complexity. It has been observed that a slightly modified proof of the main results in [4] can be used to establish $\mathcal{M}(c, c) = O(c^4)$, whereas, by comparison, our Proposition 3.4 implies that $\mathcal{M}(c, c) = O(c^3)$. It however remains to be seen if our arguments can be modified to yield a sharper upper bound for the setting in [4] since the graphs considered therein are more general and in particular may not necessarily be acyclic.

The remainder of the paper is organized as follows. First of all, notations and terminology that will be used in our proofs will be introduced in section 2. And in section 3, we will prove Theorem 1.3, the main result in this paper. Some properties, exact values, and a scaling law of $\mathcal{M}$ will be established in sections 4; more specifically, we will show that for any positive integer c, $\mathcal{M}(1, c) = c$ (Theorem 4.3) $\mathcal{M}(2, c) = 3c - 1$ (Theorem 4.4), and a scaling law of $\mathcal{M}(c_1, c_2)$ with one of the parameters fixed (Theorem 4.6).

2. Notations and terminologies. Let $G \in \mathcal{G}(c_1, c_2, \ldots, c_n)$. For an edge e in G, we will use $h(e)$ and $t(e)$ to denote its *head* and *tail*, respectively. And we say a vertex v is *reachable* from another vertex u if there is a directed path from u to v. For a directed path β containing two vertices u, v with v reachable from u, let $\beta[u, v]$ denote the *segment* of β starting from u and ending at v. For two distinct vertices u, v, we say u is *smaller* than v (or, equivalently, v is *larger* than u) if v is reachable from u. Similarly, for two distinct edges e, f in G, we say e is smaller than f (or, equivalently, f is larger than e) if $t(f)$ is reachable from $h(e)$; if, in addition, e, f, and the connecting path from $h(e)$ to $t(f)$ all belong to a path β, we say e is smaller than f on β. For a set of vertices $v_1, v_2, \ldots, v_k$ in G, define $G|v_1, \ldots, v_k)$ to be the part of G that is "upstream" of $v_1, v_2, \ldots, v_k$ or, more rigorously, $G|v_1, \ldots, v_k)$ is the subgraph of G induced on the set of vertices, each of which is smaller than or equal to some v_i, $i = 1, 2, \ldots, k$.

Now, choose $\alpha = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ such that each $\alpha_i \in \Lambda_i(G)$, $i = 1, 2, \ldots, n$. We say α_i is *reroutable* if there exists a different set α_i' of c_i edge-disjoint paths from s_i to r_i; and, more specifically, we say α_i can be *rerouted* to α_i' through removing $E(\alpha_i) - E(\alpha_i')$ from the α_i-paths and then adding $E(\alpha_i') - E(\alpha_i)$ to form the α_i'-paths. Here, we remark that it is possible that $E(\alpha_i) = E(\alpha_i')$ even if α_i and α_i' are different, in which case, we say α_i is *equivalent to* α_i', denoted by $\alpha_i \sim \alpha_i'$ (see Figure 4 for an example). And we say G is reroutable with respect to $\alpha_1, \alpha_2, \ldots, \alpha_n$ if some α_i is reroutable. Note that for a nonreroutable G, the choice of α_i is unique, so $M(G)$ is simply $M(G; \alpha_1, \alpha_2, \ldots, \alpha_n)$. For any fixed i, reverse the directions of those

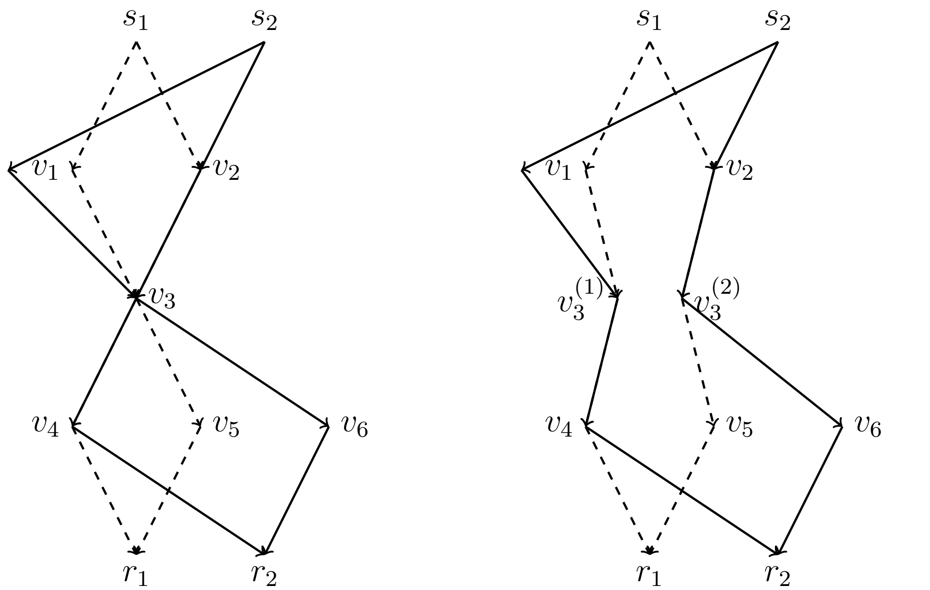

FIG. 4. *For the graph on the left, let* $\alpha_{1,1} = (s_1, v_1, v_3, v_5, r_1)$, $\alpha_{1,2} = (s_1, v_2, v_3, v_4, r_1)$, $\alpha_{2,1} = (s_2, v_1, v_3, v_4, r_2)$, $\alpha_{2,2} = (s_2, v_2, v_3, v_6, r_2)$, $\alpha'_{1,1} = (s_1, v_1, v_3, v_4, r_1)$, *and* $\alpha'_{1,2} = (s_1, v_2, v_3, v_5, r_1)$. *As illustrated by the graph on the right, during the corresponding type (D) operation, we first reset* $\alpha_{1,1}$ *to be* $\alpha'_{1,1}$ *and reset* $\alpha_{1,2}$ *to be* $\alpha'_{1,2}$; *and then split* v_3 *into* $v_3^{(1)}, v_3^{(2)}$; *and finally re-route* $\alpha_{1,1}$, $\alpha_{2,1}$ *through* $v_3^{(1)}$ *and* $\alpha_{1,2}$, $\alpha_{2,2}$ *through* $v_3^{(2)}$.

edges which do not belong to any α_i-path to obtain a new (possibly cyclic) graph $\tilde{G}$. For any two vertices v, v' in G, if there is a directed path $(v, u_1, u_2, \ldots, u_\ell, v')$ in $\tilde{G}$, we say v' is *semireachable* from v along α_i, and more specifically, we may also say v *semireaches* v' via $u_1, u_2, \ldots, u_\ell, v'$.

Example 2.1. Again, consider the graph in Figure 2. Note that e_1 is smaller than both e_2 and e_4, and so is e_3. $G|s_1, s_2)$ only consists of two isolated vertices s_1, s_2; $G|v_1, v_5)$ is the subgraph of G induced on the set of vertices $\{s_1, s_2, v_1, v_5\}$; $G|v_4, v_8)$ is the subgraph of G induced on the set of vertices $\{s_1, s_2, v_1, v_5, v_2, v_6, v_3, v_7, v_4, v_8\}$; and $G|r_1, r_2)$ is just G itself.

Let $\alpha_1, \alpha_2, \alpha'_1$ be chosen as in Example 1.2. Note that α_1 is reroutable, since there exists another group α'_1 of two edge-disjoint paths from s_1 to r_1 and, hence, G is reroutable with respect to α_1, α_2; similarly, one can verify that α_2 is also reroutable. It is also easy to check, by definition, that v_3 semireaches v_7 along α_2 via v_2, v_7; v_2 semireaches v_7 along α_2 via v_7; and v_7 semireaches itself along α_1 via v_2, v_3, v_6, v_7; and v_7 also semireaches itself along α_2 via v_6, v_3, v_2, v_7.

In the remainder of this section, we focus on the case that $n = 2$ and consider the following 4 types of operations on G with $\alpha = \{\alpha_1, \alpha_2\}$ chosen as above:

(A) If a vertex v is isolated, then remove v.
(B) If an edge e does not belong to any α_j-path, $j = 1, 2, \ldots, n$, then remove e.
(C) If, for a vertex v, there is only one incoming edge (v', v) and only one outgoing edge (v, v''), we then contract the edge (v, v'') into one single vertex, which, for any α_j-path passing through (v, v''), naturally yields a corresponding new α_j-path.

(D) A vertex v in G is said to be *splittable* if the set of its incoming links can be partitioned into $\ell \geq 2$ subsets, $\{e_{i,1}, e_{i,2}, \ldots, e_{i,j_i}\}$, and the set of its outgoing links can be partitioned into ℓ subsets, $\{f_{i,1}, f_{i,2}, \ldots, f_{i,k_i}\}$, $i = 1, 2, \ldots, \ell$, such that for any feasible i, $\cup_{l=1}^{j_i} \alpha'(e_{i,l}) = \cup_{l=1}^{k_i} \alpha'(f_{i,l})$ for some $\alpha' = \{\alpha'_1, \alpha'_2\}$ with $\alpha'_i \sim \alpha_i$, $i = 1, 2$. Note that it can be easily verified that

$$M(G; \alpha'_1, \alpha'_2) \leq M(G; \alpha_1, \alpha_2). \tag{2.1}$$

For any such v, reset α to be α', split v into ℓ copies, $v^{(1)}, v^{(2)}, \ldots, v^{(\ell)}$, and for each feasible i, reroute all paths in $\alpha(e_{i,l})$ through $v^{(i)}$ instead of v to form the corresponding new α-paths (see Figure 4 for an illustrative example).

A repeated application of the above 4 types of operations to G until there are no feasible operations left will yield $\hat{G} \in \mathcal{G}(c_1, c_2)$, a *reduced* version of G, with $\hat{\alpha} = \{\hat{\alpha}_1, \hat{\alpha}_2\}$, where $\hat{\alpha}_j \in \Lambda_j(G)$ corresponds to α_j, $j = 1, 2$. Noting that a type (D) operation will not increase the number of mergings and all other types of operations will keep the number of mergings, we conclude that

$$M(G; \alpha_1, \alpha_2) \leq M(\hat{G}; \hat{\alpha}_1, \hat{\alpha}_2), \tag{2.2}$$

where the strict inequality holds if (2.1) strictly holds for at least one type (D) operation during the procedure. Here, we remark that a rerouting of $\hat{\alpha}$ that strictly reduces the number of mergings naturally corresponds to a rerouting of α that strictly reduces the number of mergings, which will be implicitly used throughout this paper. We say G is reduced with respect to α_1, α_2 if all the 4 types of operations are unfeasible for G. If, in addition, G is nonreroutable, we simply say G is reduced and rewrite $M(G; \alpha_1, \alpha_2)$ as $|G|_{\mathcal{M}}$ without referencing α_1, α_2 since each α_j is the only element in $\Lambda_i(G)$.

3. Proof of Theorem 1.3.

We first need the following lemma.

LEMMA 3.1. *Let $G \in \mathcal{G}(c_1, c_2)$ be reduced with respect to α_1, α_2, where $\alpha_i \in \Lambda_i(G)$, $i = 1, 2$. Then, a rerouting of α_1 to α'_1 will strictly decrease the number of mergings in G, that is,*

$$M(G; \alpha_1, \alpha_2) < M(G; \alpha'_1, \alpha_2). \tag{3.1}$$

Proof. Let V_0 denote the set of vertices in G whose in-degrees are at least 2. Since G is reduced with respect to α_1, α_2, we have $M(G; \alpha_1, \alpha_2) = |V_0|$. Let G' denote the subgraph of G induced on all the α'_1-paths and all the α_2-paths, and let V'_0 denote the set of vertices in G' whose in-degrees are at least 2. Similarly as above, it holds that $M(G; \alpha'_1, \alpha_2) = |V'_0|$. Obviously, $V'_0 \subseteq V_0$, which means that $|V'_0| \leq |V_0|$. Now, one verifies that each nonterminal vertex in G either (has exactly two incoming edges and one outgoing edge) or (has exactly one incoming edge and two outgoing edges), which then implies that $E(\alpha_1) - E(\alpha'_1)$ is nonempty and so there exists at least one vertex $v \in V_0$ such that $v \notin V'_0$ and, thereby, $V'_0 \subset V_0$, which further implies (3.1), as desired. ∎

Remark 3.2. It follows from Lemma 3.1 that to compute $\mathcal{M}(c_1, c_2)$, it is enough to take the supremum in Definition 1.1 over all graph G such that G is nonreroutable and reduced.

We also need the following key lemma, which gives the necessary and sufficient conditions for a reduced graph in $\mathcal{G}(c_1, c_2)$ to be reroutable.

LEMMA 3.3. *Let $G \in \mathcal{G}(c_1, c_2)$ be reduced with respect to α_1, α_2, where $\alpha_i \in \Lambda_i(G)$. The following statements are equivalent:*

(a) *α_1 (resp., α_2) is reroutable.*

(b) *There exists a merging e such that $h(e)$ semireaches itself along α_1 (resp., α_2).*

(c) *There exists a merging e such that $t(e)$ semireaches itself along α_1 (resp., α_2).*

Proof. We first establish the equivalence between (b) and (c). If $h(e)$ semireaches itself via $u_1, u_2, \ldots, u_\ell, h(e)$, then u_1 must be a tail of certain merging $\hat{e}$, which semireaches itself via $u_2, u_3, \ldots, h(e), u_1$, establishing (b) $\Rightarrow$ (c). A similar argument can be used to establish (c) $\Rightarrow$ (b). So, in the following, we only prove the equivalence between (a) and (b).

(a) $\Rightarrow$ (b). Suppose that α_1 can be rerouted to α_1' and that $E(\alpha_1) - E(\alpha_1')$ consists of the following α_1-edges: $e_1, e_2, \ldots, e_\ell$. Without loss of generality, assume that $\{\alpha_{1,1}, \alpha_{1,2}, \ldots, \alpha_{1,k}\}$ is the set of α_1-paths which contain at least one e_j, that is,

$$\bigcup_{j=1}^{\ell} \alpha_1(e_j) = \{\alpha_{1,1}, \alpha_{1,2}, \ldots, \alpha_{1,k}\};$$

and assume that $e_1, e_2, \ldots, e_k$ are the smallest edges from $E(\alpha_1) - E(\alpha_1')$ on $\alpha_{i,1}, \alpha_{i,2}, \ldots, \alpha_{i,k}$, respectively. Since G is reduced with respect to α_1, α_2, $E(\alpha_1') - E(\alpha_1)$ only consists of α_2-edges and, furthermore, there are some α_2-edges $f_1, f_2, \ldots, f_k$ such that for each $i = 1, 2, \ldots k$, there exists $1 \leq k_i \leq \ell$ such that $t(f_i) = t(e_i)$, $h(f_i) = h(e_{k_i})$. It then follows that one can find a subset $\{\hat{k}_1, \hat{k}_2, \ldots, \hat{k}_s\}$ of $\{1, 2, \ldots, k\}$ such that $h(f_{\hat{k}_1}) \in \alpha_{1,\hat{k}_2}$, $h(f_{\hat{k}_2}) \in \alpha_{1,\hat{k}_3}, \ldots, h(f_{\hat{k}_s}) \in \alpha_{1,\hat{k}_1}$, which implies that there is a merging whose head semireaches itself along α_1 (see Figure 5 for an illustrative example).

(b) $\Rightarrow$ (a). Assume that $h(e)$ semireaches itself along α_1 via $v_1, v_2, \ldots, v_\ell$. Let

$$C = \{(h(e), v_1), (v_1, v_2), (v_2, v_3), \ldots, (v_{\ell-1}, v_\ell), (v_\ell, h(e))\}.$$

Let $A = \{(u, w) \in C : (w, u) \text{ is a reversed } \alpha_2\text{-edge}\}$, and let $B = C - A$. Then, it can be readily verified that the subgraph of G induced on $(E(\alpha_i) \cup A) - B$ consists of c_i connected components, each of which is a path from s_i to r_i (see Figure 5 for an illustrative example). ∎

Before the proof of Theorem 1.3, we will first prove the following proposition, which will establish the finiteness of $\mathcal{M}(c_1, c_2)$.

PROPOSITION 3.4. *For any c_1, c_2,*

$$\mathcal{M}(c_1, c_2) \leq c_1 c_2 (c_1 + c_2)/2.$$

Proof. It suffices to prove that for any $G \in \mathcal{G}(c_1, c_2)$ with $\alpha = \{\alpha_1, \alpha_2\}$, where $\alpha_i \in \Lambda_i(G)$, $i = 1, 2$, if

$$M(G; \alpha_1, \alpha_2) \geq c_1 c_2 (c_1 + c_2)/2 + 1, \tag{3.2}$$

then α can be routed, say, to $\alpha' = \{\alpha_1', \alpha_2'\}$, such that

$$M(G; \alpha_1', \alpha_2') < M(G; \alpha_1, \alpha_2). \tag{3.3}$$

First of all, we repeatedly apply the 4 types of operations to G as in section 2 to obtain its reduced version $\hat{G}$. If the strict inequality in (2.2) holds, then, as argued in

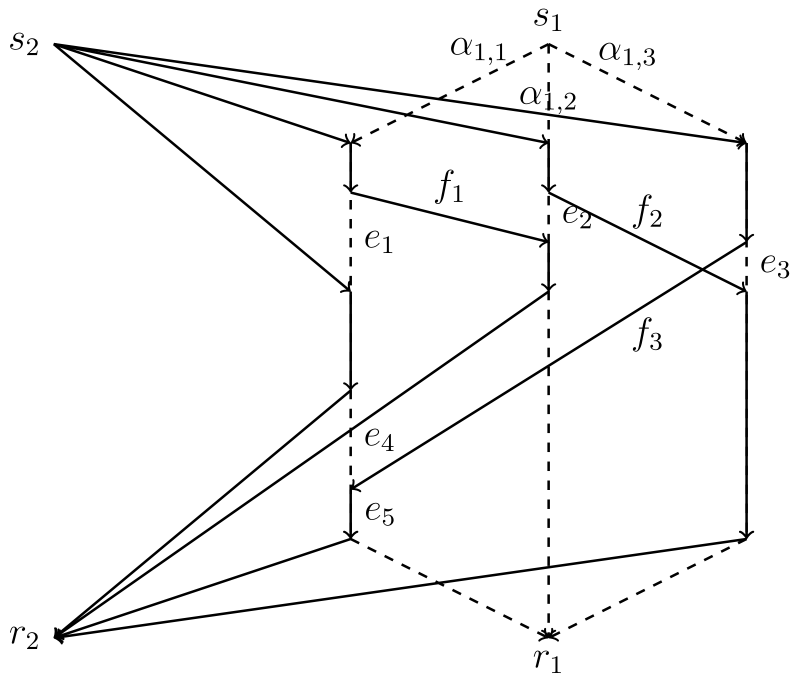

FIG. 5. *For the merging e_5, $t(e_5)$ semireaches itself along α_1 via $t(f_3)$, $h(e_3)$, $t(f_2)$, $h(e_2)$, $t(f_1)$, $h(e_1)$, $t(e_4)$, $t(e_5)$. One verifies that α_1 can be rerouted to skip e_1, e_2, e_3, e_4 and pass through f_1, f_2, f_3 instead.*

section 2, the existence of α' then immediately follows. So, in the following, we only need to prove (3.3) for the case that G is reduced with respect to α; more specifically, we will show that if G is reduced with respect to α and (3.2) holds, then either α_1 or α_2 is reroutable, which, by Lemma 3.1, implies (3.3).

Consider the following operations on G: we first delete all the edges that are both α_1-edges and α_2-edges, which are necessarily mergings due to the assumption that G is reduced, then we reverse the directions of the remaining α_2-edges. Note that after the above operations, for any directed path in $\hat{G}$, each edge is either an α_1-edge or a reversed α_2-edge. Suppose that there is a directed cycle $(v_1, v_2, \ldots, v_\ell)$ in $\hat{G}$, where $v_\ell = v_1$ and $e_i \triangleq (v_i, v_{i+1})$ is a reversed α_2-edge for any odd i and an α_1-edge for any even i. It can be verified that all v_j belong to $V_{\mathcal{M}}$, where $V_{\mathcal{M}}$ denotes the set of all the tails and heads of all the mergings. It then follows that v_1 semireaches itself along α_1 via $v_2, v_3, \ldots, v_{\ell-1}, v_\ell, v_1$, which implies α_1 is reroutable.

So, in the following, we assume that $\hat{G}$ is acyclic. Note that in $\hat{G}$, s_1, r_2 have outdegree c_1, c_2, respectively, s_2, r_1 has in-degree c_1, c_2, respectively, and any vertex in $V_{\mathcal{M}}$ has in-degree 1 and out-degree 1. It then immediately follows that $\hat{G}$ consists of c_1+c_2 pairwise vertex-disjoint paths, each of which starts from either s_1 or r_2, ends at either s_2 or r_1, and consists of a sequence of concatenated edges that alternates between an α_1-edge and a reversed α_2-edge. It then follows from

$$|V_{\mathcal{M}}| = 2M(G; \alpha_1, \alpha_2) \geq c_1c_2(c_1+c_2)+1,$$

that out of the c_1+c_2 edge-disjoint paths, there must be at least one path, say, γ, that contains more than c_1c_2 vertices in $V_{\mathcal{M}}$. It then follows that there are two distinct vertices $u, v \in V_{\mathcal{M}}$ on γ and i_0, j_0 such that u corresponds to a merging by α_{1,i_0} and α_{2,j_0}, and so does v. Note that if u is smaller (resp., larger) than v on α_{1,i_0}, then u will also be smaller (resp., larger) than v on α_{2,j_0}, otherwise we would have a cycle formed by concatenating $\alpha_{1,i_0}[u, v]$ and $\alpha_{2,j_0}[v, u]$ in G, which contradicts the assumption that G is acyclic.

In what follows, we assume that $\gamma[u, v] = (u, w_1, w_2, \ldots, w_\ell, v)$ and we consider the following conditions:

- u is smaller (larger) than v on α_{1,i_0};
- u is the tail (head) of the corresponding merging in G, v is the tail (head) of the corresponding merging in G.

Ignoring the parenthesized words for the moment, one verifies that v semireaches itself along α_2 via $w_\ell, w_{\ell-1}, \ldots, w_1$ and vertices on $\alpha_{1,i_0}[u, v]$ (ordered by the direction of the path α_{1,i_0}), implying α_2 is reroutable. A similar argument can be applied to other cases when any parenthesized words replace the words before them, which completes the proof. ∎

We are now ready for the proof of Theorem 1.3. With Proposition 3.4 established, it suffices to prove the following proposition.

Proposition 3.5. *For any $c_1, c_2, \ldots, c_n$, we have*

$$\mathcal{M}(c_1, c_2, \ldots, c_n) \leq \sum_{i<j} \mathcal{M}(c_i, c_j).$$

Proof. It suffices to prove that for $n \geq 3$

$$\mathcal{M}(c_1, c_2, \ldots, c_n) \leq \mathcal{M}(c_1, c_2, \ldots, c_{n-1}) + \sum_{i<n} \mathcal{M}(c_i, c_n). \tag{3.4}$$

Consider any graph $G \in \mathcal{G}(c_1, c_2, \ldots, c_n)$ with $\alpha_i \in \Lambda_i(G)$, $i = 1, 2, \ldots, n$. Let $\bar{G}$ denote the subgraph of G induced on all α_j-paths, $j = 1, 2, \ldots, n-1$. Note that, through rerouting if necessary, we can assume that $\alpha_1, \alpha_2, \ldots, \alpha_{n-1}$ are chosen such that

$$M(\bar{G}; \alpha_1, \alpha_2, \ldots, \alpha_{n-1}) \leq \mathcal{M}(c_1, c_2, \ldots, c_{n-1}).$$

A merging e is said to be *new* if e is an α_n-edge, and $\alpha_n(e)$ merges with each (well-defined) $\alpha_j(e)$, $j = 1, 2, \ldots, n-1$, at e. We will prove that if the number of new mergings between α_n and $\alpha_1, \alpha_2, \ldots, \alpha_{n-1}$ is larger than or equal to

$$l \triangleq \mathcal{M}(c_1, c_n) + \mathcal{M}(c_2, c_n) + \cdots + \mathcal{M}(c_{n-1}, c_n) + 1,$$

certain reroutings can be done to strictly reduce the number of new mergings. Apparently, this is sufficient to imply (3.4) and then the proposition. For ease of presentation, in the remainder of this proof, we assume that all mergings in G are new.

Now, label all the new mergings as $e_1, e_2, \ldots, e_l$. Then, by the pigeonhole principle, there exists some $i = 1, 2, \ldots, n-1$ such that α_i merges with α_n for more than $\mathcal{M}(c_i, c_n)$ times. As a consequence, the subgraph of G induced on $\{\alpha_i, \alpha_n\}$ is reroutable; in other words, either α_i or α_n can be rerouted in this induced subgraph of G. If such a rerouting is in fact a rerouting of α_n, then the number of new mergings between α_n and $\alpha_1, \alpha_2, \ldots, \alpha_{n-1}$ will be strictly decreased after such a rerouting. So in the following we assume that the rerouting between every α_i and α_n, if it exists, is a rerouting of α_i. Then, after the rerouting of α_i, there are at least

$$\mathcal{M}(c_1, c_{k+1}) + \cdots + \mathcal{M}(c_{i-1}, c_{k+1}) + \mathcal{M}(c_{i+1}, c_{k+1}) + \cdots + \mathcal{M}(c_k, c_{k+1}) + 1$$

of the e_j's, at which the new α_i does not merge. This implies that there exists at least one e_j such that none of α_i's, $i = 1, 2, \ldots, n-1$, merge with α_n at e_j. So the number of new mergings between $\alpha_1, \alpha_2, \ldots, \alpha_{n-1}$ and α_n strictly decreases after the reroutings of all α_i's. ∎

4. Properties, exact values, and a scaling law. With the finiteness of $\mathcal{M}$ established in Theorem 1.3, one natural question is to compute the values of $\mathcal{M}$, which seems to be fairly difficult, even for small parameters. In this section, among other results, we will derive the values of $\mathcal{M}$ for certain special parameters. Some propositions on properties of $\mathcal{M}$ as a function of its parameters will be established as well; these propositions, besides helping to derive the values of $\mathcal{M}$, are of interest in their own right.

The following proposition shows that $\mathcal{M}$ is "sup-linear" in all its parameters.

PROPOSITION 4.1. *For any $c_{1,0}, c_{1,1}, c_2, \ldots, c_n$, we have*

$$\mathcal{M}(c_{1,0} + c_{1,1}, c_2, \ldots, c_n) \geq \mathcal{M}(c_{1,0}, c_2, \ldots, c_n) + \mathcal{M}(c_{1,1}, c_2, \ldots, c_n).$$

Proof. We only prove the proposition for the case $n = 2$, the case with a generic n being similar.

For any $c_{1,0}, c_{1,1}$, and c_2, consider the following acyclic directed graph G (see Figure 6 for an illustrative example) with 2 sources s_1, s_2 and 2 sinks r_1, r_2 such that

1. there is a set α_1 of $c_{1,0} + c_{1,1}$ edge-disjoint paths from s_1 to r_1, here $\alpha_1 = \alpha_1^{(0)} \cup \alpha_1^{(1)}$, where $\alpha_1^{(0)}$ and $\alpha_1^{(1)}$ are mutually exclusive, consisting of $c_{1,0}$, $c_{1,1}$ edge-disjoint paths, respectively, and there is a set α_2 of c_2 edge-disjoint paths from s_2 to r_2;
2. mergings by $\alpha_1^{(0)}, \alpha_2$ and mergings by $\alpha_1^{(1)}, \alpha_2$ are "sequentially isolated" on α_2 in the sense that on each α_2-path, the smallest $\alpha_1^{(1)}$-merging is larger than the largest $\alpha_1^{(0)}$-merging;
3. the number of mergings in the subgraph of G induced on $\alpha_1^{(0)}$ and α_2 achieves $\mathcal{M}(c_{1,0}, c_2)$, and the number of mergings in the subgraph of G induced on $\alpha_1^{(1)}$ and α_2 achieves $\mathcal{M}(c_{1,1}, c_2)$.

It can be verified that for such G, the size of the minimum edge cut between s_1 and r_1 is $c_{1,0} + c_{1,1}$, and the size of the minimum edge cut between s_2 and r_2 is c_2, and

$$M(G; \alpha_1, \alpha_2) = \mathcal{M}(c_{1,0}, c_2) + \mathcal{M}(c_{1,1}, c_2),$$

which implies that

$$\mathcal{M}(c_{1,0} + c_{1,1}, c_2) \geq \mathcal{M}(c_{1,0}, c_2) + \mathcal{M}(c_{1,1}, c_2). \qquad ∎$$

The following proposition gives a lower bound on $\mathcal{M}$ with multiple parameters using $\mathcal{M}$ with two parameters.

PROPOSITION 4.2. *For any $c_1, c_2, \ldots, c_n$ and any fixed k with $1 \leq k \leq n$, we have*

$$\mathcal{M}(c_1, c_2, \ldots, c_n) \geq \sum_{i \leq k, j \geq k+1} \mathcal{M}(c_i, c_j).$$

Proof. For any $c_1, c_2, \ldots, c_n$, consider the following directed graph G (see Figure 7 for an illustrative example) with n sources $s_1, s_2, \ldots, s_n$ and n sinks $r_1, r_2, \ldots, r_n$ such that for any fixed k with $1 \leq k \leq n$,

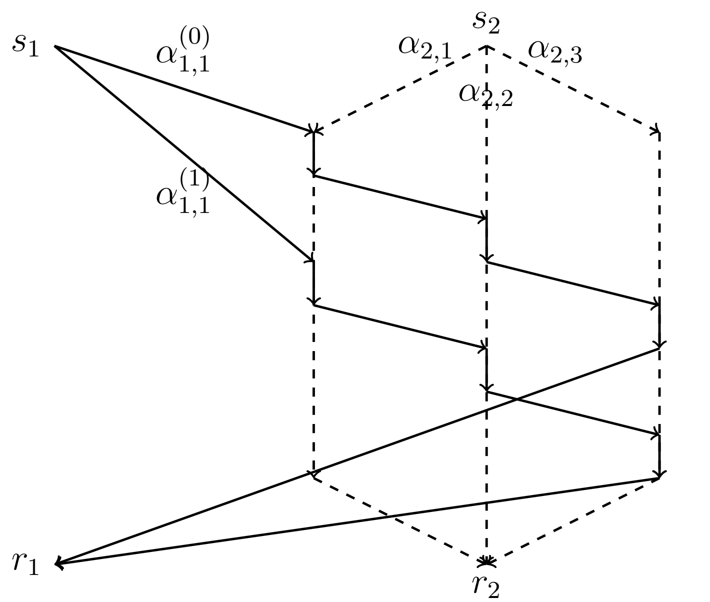

FIG. 6. *Each merging by $\alpha_{1,1}^{(0)}$ and an α_2-path is smaller than than that by $\alpha_{1,2}^{(0)}$ and the same α_2-path.*

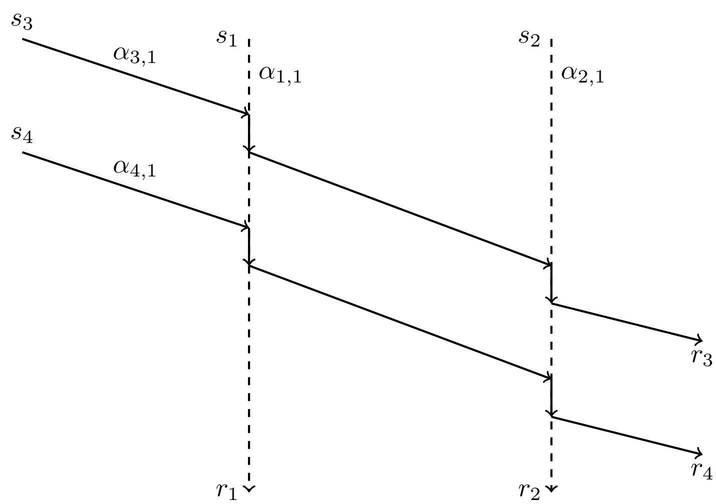

FIG. 7. *On the path $\alpha_{i,1}$, $i = 1, 2$, the merging with $\alpha_{3,1}$ is smaller than that with $\alpha_{4,1}$, and on the path $\alpha_{i,1}$, $i = 3, 4$, the merging with $\alpha_{1,1}$ is smaller than that with $\alpha_{2,1}$.*

1. there is a set α_i of c_i edge-disjoint paths from s_i to r_i for each i;
2. all α_i's, $i \leq k$, do not merge with each other, and all α_j's, $j \geq k+1$, do not merge with each other;
3. for any i with $i \leq k$, mergings by α_i and all α_j's, $j \geq k+1$, are sequentially isolated on α_i in the sense that on each α_i-path, for any $j_1 < j_2$ with $j_1, j_2 \geq k+1$, the smallest α_{j_2}-merging is larger than the largest α_{j_1}-merging.

Similarly for any j with $j \geq k+1$, mergings by α_j and all α_i's, $i \leq k$, are sequentially isolated on α_j;

4. the number of mergings in the subgraph of G induced on any α_i, $i \leq k$, and any α_j, $j \geq k+1$, achieves $\mathcal{M}(c_i, c_j)$.

One checks that for such a graph G, the size of the minimum edge cut between s_i and r_i is c_i, and

$$M(G) = \sum_{i \leq k, j \geq k+1} \mathcal{M}(c_i, c_j),$$

which implies that

$$\mathcal{M}(c_1, c_2, \ldots, c_n) \geq \sum_{i \leq k, j \geq k+1} \mathcal{M}(c_i, c_j). \qquad \square$$

The following theorem is straightforward. We give a proof for completeness.

THEOREM 4.3. *For any c,*

$$\mathcal{M}(1, c) = c.$$

Proof. Consider a nonreroutable and reduced graph $G \in \mathcal{G}(1, c)$ with $\alpha_i \in \Lambda_i(G)$, $i = 1, 2$. If $\alpha_{1,1}$ merges with some α_2-path, say, $\alpha_{2,j}$, at mergings e and $\hat{e}$, then we can reroute $\alpha_{1,1}$ by replacing $\alpha_{1,1}[t(e), t(\hat{e})]$, the subpath of $\alpha_{1,1}$ starting from $t(e)$ to $t(\hat{e})$, by $\alpha_{2,j}[t(e), t(\hat{e})]$, the subpath of $\alpha_{2,j}$ starting from $t(e)$ to $t(\hat{e})$. This contradicts the assumption that G is nonreroutable, which implies that $\alpha_{1,1}$ can be chosen to merge with each α_2-path for at most once, which further implies that

$$\mathcal{M}(1, c) \leq c.$$

For the other direction, by Proposition 4.1, we have

$$\mathcal{M}(1, c) \geq \sum_{i=1}^{c} \mathcal{M}(1, 1) = c,$$

the last equality follows from the simple fact that $\mathcal{M}(1, 1) = 1$. $\square$

THEOREM 4.4. *For any c,*

$$\mathcal{M}(2, c) = 3c - 1.$$

Proof. The upper bound direction. We first show that

$$\mathcal{M}(2, c) \leq 3c - 1.$$

Consider a nonreroutable and reduced $G \in \mathcal{G}(2, c)$ with $\alpha_i \in \Lambda_i(G)$, $i = 1, 2$. In this proof, for notational convenience, we rewrite α_1, α_2 as ψ, ϕ, respectively. Assume that out of c ϕ-paths, there are k ϕ-paths, say $\phi_1, \phi_2, \ldots, \phi_k$, each of which merges for at least 3 times. Notice that when $k = 0$, the total number of mergings in G is upper bounded by $2c$, which trivially implies the upper bound direction. So, in this proof, we only consider the case when $k \geq 1$. For $i = 1, 2, \ldots, k$, assume that ϕ_i sequentially merges at edges $e_{i,1}, e_{i,2}, \ldots, e_{i,m_i}$. Let $\ell(i, j)$ denote the index of the ψ-path which $e_{i,j}$ belongs to. Since each ϕ-path has to merge with ψ_1, ψ_2 alternately, we have $\ell(i, j) = \ell(i, k)$ if $j = k \mod 2$.

Note that for each pair of mergings $e_{i,j}, e_{i,j+2}$, there must exist at least one merging, say $f_{i,j}$, which is in-between $e_{i,j}$ and $e_{i,j+2}$ on $\psi_{\ell(i,j)}$. It can be readily verified (see Figure 8 for an illustrative example) that

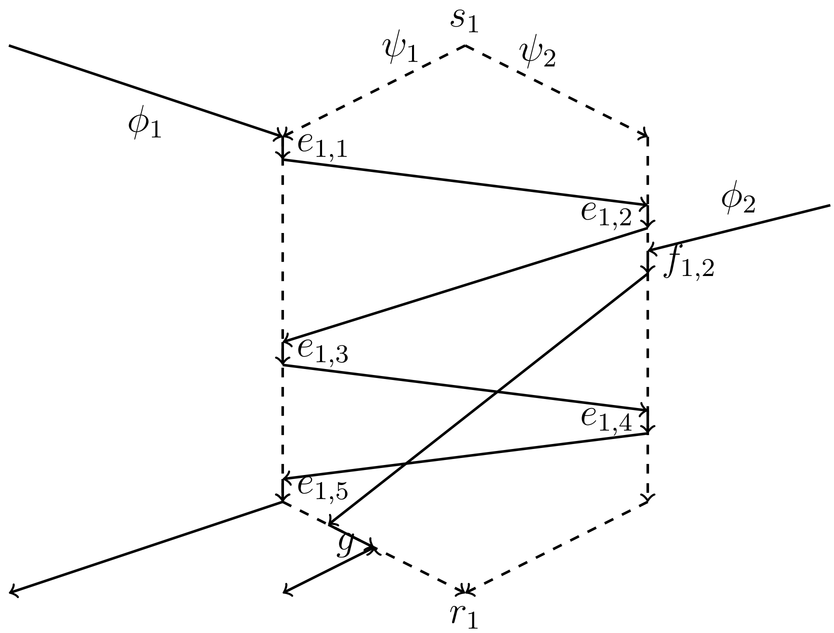

FIG. 8. *On ϕ_2, the next merging after $f_{1,2}$, say, g, will have to be in-between $e_{1,1}$ and $e_{1,3}$, and as a result, the nonreroutability of G forbids the third merging. Indeed, g cannot be in-between s_1 and $e_{1,1}$ since the acyclicity assumption would be violated. And g cannot be (in-between $e_{1,3}$ and $e_{1,5}$) or (in-between $e_{1,5}$ and r_1) since it would mean ψ is reroutable; for example, if g is in-between $e_{1,5}$ and r_1, then $t(g)$ can semireach itself along ψ via $h(f_{1,2}), t(e_{1,4}), h(e_{1,3}), t(e_{1,5}), h(e_{1,5}), t(g)$.*

- for any i, j, $\phi(f_{i,j})$ merges with ψ-paths at most twice; and
- if $\phi(f_{i,j}) = \phi(f_{i,k})$ for any $j < k$, then necessarily $k = j + 1$.

Now, for the associated ϕ-paths of all $f_{i,j}$, we claim the following.

CLAIM 4.5. For any fixed i, one can choose all $f_{i,j}$ such that for $j \neq k$, $\phi(f_{i,j}) \neq \phi(f_{i,k})$.

The claim can be shown via an inductive argument on the length of path ϕ_i. The case when $m_i = 3$ is trivial. Now suppose the claim is established for $m_i = l$ and assume that $f_{i,2}, f_{i,3}, \ldots, f_{i,l+1}$ all belong to different ϕ-paths. We next show that the claim is also true for $m_i = l + 1$. We will consider each of the following cases:

Case 1: $\phi(f_{i,1}) \neq \phi(f_{i,2})$. For this case, by induction assumptions, the claim is trivially true.

Case 2: $\phi(f_{i,1}) = \phi(f_{i,2})$. For this case, either (in-between $e_{i,1}$ and $f_{i,1}$ on $\psi_{\ell(i,1)}$) or (in-between $f_{i,2}$ and $e_{i,4}$ on $\psi_{\ell(i,2)}$), there must be a merging, whose associated ϕ-path is different from that of $f_{i,1}$ and $f_{i,2}$. Otherwise, $t(e_{i,4})$ would semireach itself along ϕ via $h(f_{i,2}), t(f_{i,1}), h(e_{i,4}), t(e_{i,4})$, which implies ϕ is reroutable, a contradiction.

Case 2.1: In-between $e_{i,1}$ and $f_{i,1}$ on $\psi_{\ell(i,1)}$, there is a merging $f'_{i,1}$ such that $\phi(f'_{i,1}) \neq \phi(f_{i,1})$. For this case, one can simply reset $f_{i,1}$ to be $f'_{i,1}$, then the claim immediately follows.

Case 2.2: In-between $f_{i,2}$ and $e_{i,4}$ on $\psi_{\ell(i,2)}$, there is a merging $f'_{i,2}$ such that $\phi(f'_{i,2}) \neq \phi(f_{i,1})$. For this case, we have the following subcases.

Case 2.2.1: $\phi(f'_{i,2}) \neq \phi(f_{i,3})$. For this case, we can simply reset $f_{i,2}$ to be $f'_{i,2}$ to establish the claim.

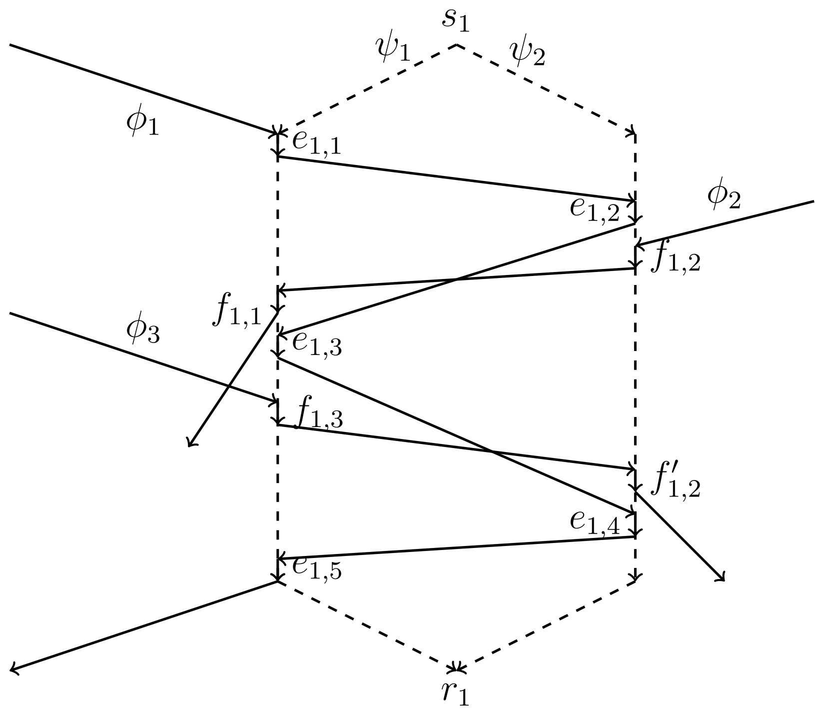

FIG. 9. *If there are no mergings between $f_{1,3}$ and $e_{1,5}$, then $t(e_{1,5})$ can semireach itself along ϕ via $h(f_{1,3}), t(f'_{1,2}), h(f_{1,2}), t(f_{1,1}), h(e_{1,1}), t(e_{1,2}), h(e_{1,2}), t(e_{1,3}), h(e_{1,3}), t(e_{1,4}), h(e_{1,4}), t(e_{1,5})$.*

Case 2.2.2: $\phi(f'_{i,2}) = \phi(f_{i,3})$. For $j = 2, 3, \ldots, m_i - 3$, we say $f_{i,j}$ is of *type* I if (there exists exactly one merging $f'_{i,j}$ in-between $f_{i,j}$ and $e_{i,j+2}$) and ($f'_{i,j}$, $f_{i,j+1}$ belong to the same ϕ-path).

Case 2.2.2.1: All $f_{i,j}$, $j = 2, 3, \ldots, m_i - 3$, are of type I. For this case, consider f_{i,m_i-2}. One checks that there must exist at least a merging, say, f'_{i,m_i-2}, in-between f_{i,m_i-2} and e_{i,m_i} on $\psi_{\ell(i,m_i-2)}$, since otherwise $t(e_{i,m_i})$ would semireach itself along ϕ via

$$h(f_{i,m_i-2}), t(f'_{i,m_i-3}), h(f_{i,m_i-3}), \ldots, t(f_{i,1}) \text{ and vertices on } \phi_i[h(e_{i,1}), t(e_{i,m_i})],$$

which implies ϕ is reroutable, a contradiction (see Figure 9 for an illustrative example). Then we can reset f_{i,m_i} to be f'_{i,m_i}, f_{i,m_i-1} to be $f'_{i,m_i-1}, \ldots, f_{i,2}$ to be $f'_{i,2}$. One checks that each of the newly defined $f_{i,j}$ belongs to a different ϕ-path.

Case 2.2.2.2: For some $2 \le k \le m_i - 3$, $f_{i,k}$ is not of type I. Let $2 \le k \le m_i - 3$ be the smallest index such that $f_{i,k}$ is not of type I, meaning either

- there is no merging in-between $f_{i,k}$ and $e_{i,k+2}$ on $\psi_{\ell(i,k+2)}$; or
- there is a merging, say $f'_{i,k}$, in-between $f_{i,k}$ and $e_{i,k+2}$ on $\psi_{\ell(i,k+2)}$; however, $\phi(f'_{i,k}) \ne \phi(f_{i,k+1})$.

The first case implies that $t(e_{i,k+2})$ semireaches itself along ϕ by itself via

$$h(f_{i,k}), t(f'_{i,k-1}), h(f_{i,k-1}), \ldots, t(f_{i,1}) \text{ and vertices on } \phi_i[h(e_{i,1}), t(e_{i,k+2})],$$

a contradiction to the fact that G is nonreroutable; while for the second case, one can reset $f_{i,k}$ to be $f'_{i,k}$, $f_{i,k-1}$ to be $f'_{i,k-1}, \ldots,$ and $f_{i,2}$ to be $f'_{i,2}$ to establish the claim.

One also verifies that for any $i, j = 1, 2, \ldots, k, \phi_i$ and ϕ_j are "well-separated"; more precisely, one of the pair, say ϕ_i, must be "smaller" than the other one, ϕ_j, in the sense that the mergings by ϕ_i on ψ_1, ψ_2 must be smaller than the mergings by ϕ_j on ψ_1, ψ_2, respectively. Through renumbering, if necessary, we assume that for any $1 \leq i < j \leq k$, ϕ_i is always smaller than ϕ_j. Then with this, one checks that for any $1 \leq i_1 < i_2 \leq k$, f_{i_1,j_1} and f_{i_2,j_2} share the same ϕ-path if and only if $i_2 = i_1 + 1$ and $j_1 = m_{i_1} - 2, j_2 = 1$. Thus, by Claim 4.5, there must exist at least $(m_1 - 2 + m_2 - 2 + \cdots + m_k - 2) - (k - 1)$ ϕ-paths, each of which contains some $f_{i,j}$, and again each of these ϕ-paths can merge at most twice.

Now, we conclude that (below, $|\cdot|_{\mathcal{M}}$ is a shorthand notation for the number of mergings; see the end of section 2)

$$|G|_{\mathcal{M}} \leq m_1 + m_2 + \cdots + m_k + 2(c - k),$$

where

$$(m_1 - 2) + (m_2 - 2) + \cdots + (m_k - 2) - (k - 1) \leq |\{\phi(f_{i,j})\}| \leq c - k,$$

which further implies that

$$|G|_{\mathcal{M}} \leq 3c - 1.$$

So, we have established the upper bound direction.

The lower bound direction. To show

$$\mathcal{M}(2, c) \geq 3c - 1,$$

it suffices to construct a nonreroutable graph G with $M(G) = 3c - 1$. For instance, we can first choose ϕ_1 to alternately merge with ψ_1, ψ_2 for $c + 1$ times at $e_1, e_2, \ldots, e_{c+1}$. Next we choose each ϕ_i, $i = 2, 3, \ldots, c$, to merge exactly twice, while ensuring that, for all $i < j$, ϕ is smaller than ϕ_j in the sense that the merged subpaths by ϕ_i on ψ_1, ψ_2 are smaller than the merged subpaths by ϕ_j on ψ_1, ψ_2, respectively. Moreover we also require that ϕ_{2i} first merges with ψ_1 in-between e_{2i-1} and e_{2i+1}, and then merge with ψ_2 in-between e_{2i-2} and e_{2i}, and that ϕ_{2i+1} first merges with ψ_2 in-between e_{2i} and e_{2i+2}, and then merges with ψ_1 in-between e_{2i-1} and e_{2i+1} (see an example graph in Figure 10 for the case $c = 3$). It can be checked that such a graph is nonreroutable and the number of mergings is $3c - 1$. ∎

We next prove that when fixing one parameter, $\mathcal{M}(c_1, c_2)$ grows at most linearly with respect to the other parameter.

THEOREM 4.6. *For any fixed c_1, there exists a positive constant C_{c_1} such that for all c_2,*

$$\mathcal{M}(c_1, c_2) \leq C_{c_1} c_2.$$

Proof. For notational simplicity, in this proof, we rewrite c_1, c_2 as k, l, respectively, that is, we will prove that for any fixed k, there exists a positive constant C_k such that for all l,

$$\mathcal{M}(k, l) \leq C_k l.$$

We proceed by induction on k. It follows from $\mathcal{M}(1, l) = l$ (see Theorem 4.3) that for the case when $k = 1$, the theorem is true with $C_1 = 1$. Now for any $k \geq 2$, assume that for any $i = 1, 2, \ldots, k - 1$, there exists a positive constant C_i such that for all l,

$$\mathcal{M}(i, l) \leq C_i l;$$

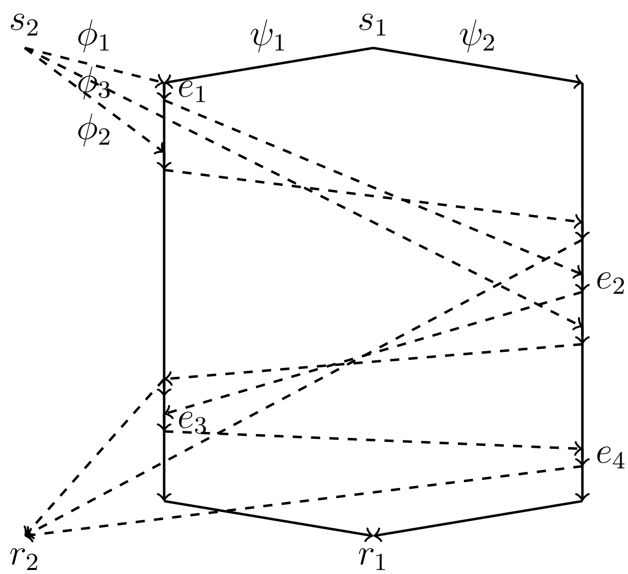

FIG. 10. *An example graph in $\mathcal{G}(2,3)$ achieving $\mathcal{M}(2,3)$.*

we next show that there exists a positive constant C_k such that for all l,

$$\mathcal{M}(k,l) \leq C_k l.$$

Consider $G \in \mathcal{G}(k,l)$ and assume G is nonreroutable and reduced. Rewrite

$$\alpha_1 = \{\alpha_{1,1}, \alpha_{1,2}, \ldots, \alpha_{1,k}\}, \quad \alpha_2 = \{\alpha_{2,1}, \alpha_{2,2}, \ldots, \alpha_{2,l}\}$$

as

$$\psi = \{\psi_1, \psi_2, \ldots, \psi_k\}, \quad \phi = \{\phi_1, \phi_2, \ldots, \phi_l\},$$

respectively. We only need to prove that there exists C_k such that

$$|G|_{\mathcal{M}} \leq C_k l.$$

Consider the following iterative procedure on G, where, for notational simplicity, we treat a graph as the union of its vertex set and edge set.

Initialization. Set

$$\mathbb{S}^{(0)} = \emptyset, \qquad \mathbb{R}^{(0)} = G.$$

Finding a normal block. Now for an arbitrary yet fixed $K > 0$ (we will choose K large enough later) and each $j = 1, 2, \ldots, k$, pick merging $e_{0,j}$ such that $e_{0,j}$ belongs to path ψ_j and (roughly speaking, $\mathbb{R}^{(0)}|h(e_{0,1}), h(e_{0,2}), \ldots, h(e_{0,k}))$ as below means the part of $\mathbb{R}^{(0)}$ that is upstream of $e_{0,1}, e_{0,2}, \ldots, e_{0,k}$; see the first paragraph of section 2 for the precise definition)

$$|\mathbb{R}^{(0)}|h(e_{0,1}), h(e_{0,2}), \ldots, h(e_{0,k}))|_{\mathcal{M}} = K.$$

Without loss of generality, we can assume that, within $\mathbb{R}^{(0)}|h(e_{0,1}), h(e_{0,2}), \ldots, h(e_{0,k}))$, $e_{0,j}$ is the largest merging on ψ_j (one can set $h(e_{0,j})$ to be s_1 if such merging does not exist on ψ_j). Now, set

$$\mathbb{L}^{(1)} = \mathbb{R}^{(0)}|h(e_{0,1}), h(e_{0,2}), \ldots, h(e_{0,k}))$$

and, subsequently,

$$\mathbb{S}^{(1)} = \mathbb{S}^{(0)} \cup \mathbb{L}^{(1)}, \qquad \mathbb{R}^{(1)} = \mathbb{R}^{(0)} - \mathbb{L}^{(1)}.$$

If a merging is the smallest (resp., the largest) on a ϕ-path, we say it is an *x-terminal* (resp., *y-terminal*) merging on the ϕ-path, or simply a *ϕ-terminal* merging. Suppose that we have already obtained

$$\mathbb{L}^{(i)} = \mathbb{R}^{(i-1)}|h(e_{i-1,1}), h(e_{i-1,2}), \ldots, h(e_{i-1,k}))$$

and

$$\mathbb{S}^{(i)} = \mathbb{S}^{(i-1)} \cup \mathbb{L}^{(i)}, \qquad \mathbb{R}^{(i)} = \mathbb{R}^{(i-1)} - \mathbb{L}^{(i)},$$

where $\mathbb{L}^{(i)}$ contains exactly K mergings and at least one ϕ-terminal merging. We then try to pick within $\mathbb{R}^{(i)}$ a merging $e_{i,j}$ on each ψ_j such that

$$|\mathbb{R}^{(i)}|h(e_{i,1}), h(e_{i,2}), \ldots, h(e_{i,k})|_{\mathcal{M}} = K,$$

where each $e_{i-1,j}$, $j = 1, 2, \ldots, k$, is chosen to be largest merging on ψ_j, and there is at least one ϕ-terminal merging in $\mathbb{R}^{(i)}|h(e_{i,1}), h(e_{i,2}), \ldots, h(e_{i,k}))$. If such $e_{i,j}$ exist or $|\mathbb{R}^{(i)}|_{\mathcal{M}} < K$, we then set

$$\mathbb{L}^{(i+1)} = \mathbb{R}^{(i)}|h(e_{i,1}), h(e_{i,2}), \ldots, h(e_{i,k}))$$

and, subsequently,

$$\mathbb{S}^{(i+1)} = \mathbb{S}^{(i)} \cup \mathbb{L}^{(i+1)}, \qquad \mathbb{R}^{(i+1)} = \mathbb{R}^{(i)} - \mathbb{L}^{(i+1)};$$

furthermore, for the case $|\mathbb{R}^{(i)}|_{\mathcal{M}} < K$, we will terminate the procedure. So far, for any obtained "block" $\mathbb{L}^{(i+1)}$, either we have $(|\mathbb{L}^{(i+1)}|_{\mathcal{M}} < K)$ or $(|\mathbb{L}^{(i+1)}|_{\mathcal{M}} = K$ and there are at least one ψ-terminal mergings in $\mathbb{L}^{(i+1)})$; such a block $\mathbb{L}^{(i+1)}$ is said to be *normal*. If $|\mathbb{R}^{(i)}| \geq K$, however, we cannot find a normal block, we continue the procedure and define a *singular* $\mathbb{L}^{(i+1)}$ in the following.

Finding a singular block. A merging within $\mathbb{S}^{(i)}$ is said to be *critical* with respect to $\mathbb{S}^{(i)}$ if its associated ϕ-path, after the said merging, does not merge anymore within $\mathbb{S}^{(i)}$. Now, let $\{\beta_j^{(i)}\}$ denote the set of all critical mergings with respect to $\mathbb{S}^{(i)}$, and let $\bar{\mathbb{T}}^{(i)}$ denote the set of all the mergings whose heads or tails are semireachable by the head of some $\beta_j^{(i)}$ along ϕ. One verifies at least one ψ-path in $\{\psi(\beta_j^{(i)})\}$ does not contain any mergings within $\bar{\mathbb{T}}^{(i)}$ (since otherwise ψ can be proven to be reroutable, a contradiction).

Assume that $f_{i,1}, f_{i,2}, \ldots, f_{i,m_i}, 1 \leq m_i \leq k-1$, are the largest mergings within $\bar{T}^{(i)}$, and they belong to paths $\psi_{j_{i,1}}, \psi_{j_{i,2}}, \ldots, \psi_{j_{i,m_i}}$, respectively. Now, we set

$$\mathbb{L}^{(i+1)} = \mathbb{R}^{(i)}|h(f_{i,1}), h(f_{i,2}), \ldots, h(f_{i,m_i}))$$

and define

$$\mathbb{T}^{(i)} = \bigcup_{j=1}^{m_i} \psi_{j_{i,j}}[h(e_{i-1,j_{i,j}}), h(f_{i,j})].$$

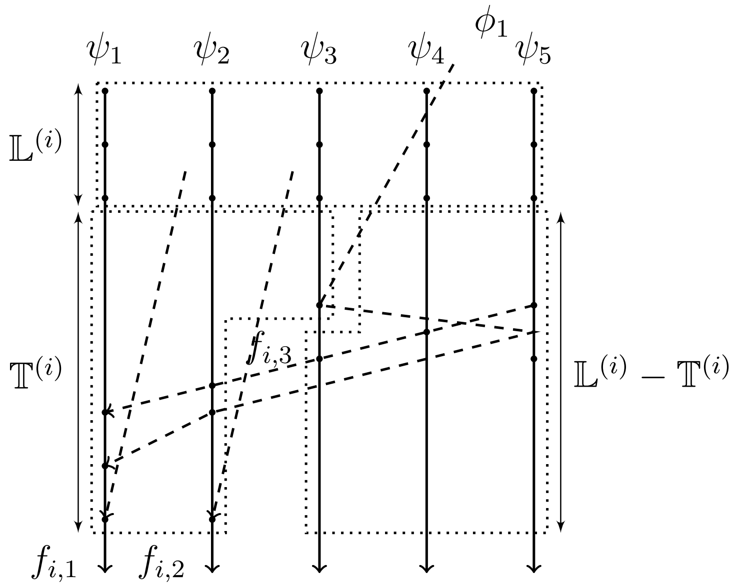

FIG. 11. *This diagram illustrates how to find a singular block, where mergings are represented by solid dots. Note that ϕ_1 is an excursive path with respect to $f_{i,3}$.*

Here, let us note that $\psi_{j_{i,j}}[h(e_{i-1,j_{i,j}}), h(f_{i,j})]$ is the segment of $\psi_{j_{i,j}}$ that is within $\mathbb{R}^{(i)}$ and before $h(f_{i,j})$ or, more formally,

$$\psi_{j_{i,j}}[h(e_{i-1,j_{i,j}}), h(f_{i,j})] = \psi_{j_{i,j}}[s_1, h(f_{i,j})] \cap \mathbb{R}^{(i)}.$$

Let x_i and y_i denote the number of x-terminal and y-terminal mergings in the ϕ-paths in $\mathbb{L}^{(i)}$, respectively. Note that for any $f_{i,j}$, $j = 1, 2, \ldots, m_i$, the associated ϕ-path, from $f_{i,j}$, may merge outside $\mathbb{T}^{(i)}$ the next time; if this ϕ-path merges within $\mathbb{T}^{(i)}$ again after a number of mergings outside $\mathbb{T}^{(i)}$, we call it an *excursive* ϕ-path (with respect to $f_{i,j}$; see Figure 11 for an illustrative example). One checks that there are at most $m_i - 1$ excursive ϕ-paths (since, otherwise, we can find a cycle in G, which is a contradiction). Note that for any merging within $\mathbb{T}^{(i)}$ other than $f_{i,j}$, $j = 1, 2, \ldots, m_i$, say g, the associated ϕ-path, from g, can only merge within $\mathbb{T}^{(i)}$. Now, consider all ϕ-paths that contains at least one merging within $\mathbb{L}^{(i+1)}$, the number of connected components of such ϕ-paths is upper bounded by $y_{i+1} + m_i$ (restricted to $\mathbb{T}^{(i)}$, an excursive path can be split into multiple connected components). Then, by the induction assumptions,

$$|\mathbb{L}^{(i+1)} \cap \mathbb{T}^{(i)}|_{\mathcal{M}} \le C_{m_i}(y_{i+1} + m_i) \le C_{k-1} y_{i+1} + C_{k-1}(k-1).$$

One also checks that there exists at least one $\psi_{j_{i,j}}$, $j = 1, 2, \ldots, m_i$, which does not merge with any ϕ-paths within $\mathbb{L}^{(i+1)} - \mathbb{T}^{(i)}$ (since otherwise we can find a cycle in G, which is a contradiction). Also, it is clear that all nonexcursive ϕ-paths that contain at least one merging within $\mathbb{L}^{(i+1)} - \mathbb{T}^{(i)}$ must have x-terminal mergings in $\mathbb{L}^{(i+1)}$, and the number of involved connected components of ϕ-paths is at most $x_{i+1} + m_i - 1$. Thus, by the induction assumptions,

$$|\mathbb{L}^{(i+1)} - \mathbb{T}^{(i)}|_{\mathcal{M}} \le C_{k-1}(x_{i+1} + m_i - 1) \le C_{k-1} x_{i+1} + C_{k-1}(k-2).$$

It then immediately follows that

$$|\mathbb{L}^{(i+1)}|_{\mathcal{M}} \leq C_{k-1}(x_{i+1} + y_{i+1}) + C_{k-1}(2k-3).$$

Now, we claim that if K is chosen such that $K \geq \mathcal{M}(k-1,k)+1$, then necessarily $x_{i+1} + y_{i+1} \geq 1$. Indeed, let $z_i = \sum_{j=1}^{i}(x_i - y_i)$, that is, z_i is the number of ϕ-paths that will continue to merge within $\mathbb{R}^{(i)}$. Then if $z_i \geq k$, one verifies that at least one ϕ-path merges within $\mathbb{L}^{(i+1)}$, however, not within $\mathbb{R}^{(i+1)}$, which means $y_{i+1} \geq 1$; if $z_i \leq k-1$, then an x-terminal merging must exist within $\mathbb{L}^{(i+1)}$, which implies that $x_{i+1} \geq 1$.

Recursive application. We continue these operations in an iterative fashion to further obtain normal blocks and singular blocks until there are no mergings left in G.

Suppose that upon the termination of the above procedure, l_1 singular blocks $\mathbb{L}_{j_1}, \mathbb{L}_{j_2}, \ldots, \mathbb{L}_{j_{l_1}}$ and l_2 normal blocks are found. Note that except for the last normal one, each block has at least one ϕ-terminal merging, which implies that

$$l_1 + l_2 \leq 2l + 1.$$

In any case, we will have for some $C_k > 0$,

$$\begin{aligned} |G|_{\mathcal{M}} &\leq Kl_2 + \sum_{i=1}^{l_1}[C_{k-1}(x_{j_i} + y_{j_i}) + C_{k-1}(2k-3)] \\ &\leq 2C_{k-1}l + (K - C_{k-1})l_2 + C_{k-1}(2k-3)l_1 \\ &\leq C_k l. \end{aligned}$$

∎

Remark 4.7. Theorem 4.6 partially confirms the following conjecture.

CONJECTURE 4.8. *There exists a positive constant C such that for any c_1, c_2, we have*

$$\mathcal{M}(c_1, c_2) \leq Cc_1c_2.$$

Note that this conjecture, together with the easily verifiable fact that $\mathcal{M}(c,c) \geq c^2$, implies that $\mathcal{M}(c,c)$ is exactly of order c^2.

Acknowledgment. We thank the anonymous reviewers for their valuable comments and suggestions, which have greatly improved the presentation of the paper.

REFERENCES

[1] R. Arnott, R. A. de Palma, and R. Lindsey, *Economics of a bottleneck*, J. Urban Econ., 27 (1990), pp. 111–130.
[2] R. Arnott, A. de Palma, and R. Lindsey, *Recent developments in the bottleneck model*, in Road Pricing, Traffic Congestion and the Environment, Elgar Economics, Cheltenham, UK, 1998, pp. 161–179.
[3] M. Beckmann, C. McGuire, and C. Winsten, *Studies in the Economics of Transportation*, Yale University Press, New Haven, CT, 1956.
[4] C. Chekuri and J. Chuzhoy, *Degree-*3 *Treewidth Sparsifiers*, in Proceedings of the 2015 ACM-SIAM Symposium on Discrete Algorithms, San Diego, CA, SIAM, Philadelphia, 2015, pp. 242–255.
[5] J. Edmonds and R. M. Karp, *Theoretical improvements in algorithmic efficiency for network flow problems*, J. ACM, 19 (1972), pp. 248–264.
[6] L. R. Ford and D. R. Fulkerson, *Maximal flow through a network*, Canad. J. Math., 8 (1956), pp. 399–404.

[7] S. Fortune, J. Hopcroft, and J. Wyllie, *The directed subgraph homeomorphism problem*, Theoret. Comput. Sci., 10 (1980), pp. 111–121.
[8] G. Han, *Menger's Paths with Minimum Mergings*, Proceedings of the 2009 IEEE Information Theory Workshop, Volos, Greece, IEEE, Piscataway, NJ, 2009, pp. 271–275.
[9] K. Menger, *Zur allgemeinen Kurventhoerie*, Fund. Math., 10 (1927), pp. 96–115.
[10] Y. Sheffi, *Urban Transportation Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1985.
[11] W. Vickrey, *Congestion theory and transport investment*, Amer. Econ. Rev., 59 (1969), pp. 251–261.
[12] R. W. Yeung, S.-Y. R. Li, N. Cai, and Z. Zhang, *Network Coding Theory*, Now Publishers, Hanover, MA, 2005.