

Minimum Number of Hubs in Networks

Easton Li Xu Guangyue Han
The University of Hong Kong The University of Hong Kong
email: xulimc@gmail.com email: ghan@hku.hk

August 23, 2013

Abstract

In this paper, a hub refers to a non-terminal vertex of degree at least three. We study the minimum number of hubs needed to guarantee certain flow demand constraints imposed between multiple pairs of sources and sinks. We prove that under the constraints, regardless of the size or the topology of the network, such minimum number is always upper bounded and we derive tight upper bounds for some special parameters. In particular, for two pairs of sources and sinks, we present a novel path-searching algorithm, the analysis of which is instrumental for the derivations of the tight upper bounds.

1 Introduction

Consider a network $G = (V, E)$, where V denotes the set of vertices in G , and E denotes the set of edges in G . A vertex in G is said to be a *source* if it is only incident with outgoing edges, and a *sink* if it is only incident with incoming edges. Often, a source or sink is referred to as a *terminal* vertex. A non-terminal vertex is said to be a *hub* if its degree is greater than or equal to 3. In this paper, we are primarily concerned with the minimum number of hubs needed when certain constraints on the flow demand between multiple pairs of sources and sinks are imposed. The flow demand constraints considered in this paper will be in terms of the vertex-cuts between pairs of sources and sinks. This can be justified by a vertex version of the max-flow min-cut theorem [1], which states that for a network with infinite edge-capacity and unit vertex-capacity, the maximum flow between one source and one sink is equal to the minimum vertex-cut between them. Here, we remark that with appropriately modified setup, our results can be stated in terms of edge-cuts as well.

More precisely, for given $C_1, C_2, \dots, C_n \in \mathbb{N}$, let $\mathcal{N}(C_1, C_2, \dots, C_n)$ denote the set of all finite networks G (see Figure 1 for an example) such that

- there are n sources S_1, S_2, \dots, S_n and n sinks R_1, R_2, \dots, R_n in G ;
- all edges in G , except those incident with a source or sink, are undirected (alternatively, bi-directional);
- for each feasible i , the minimum vertex-cut from S_i to R_i is C_i .

Now, define

$$\mathcal{H}(C_1, C_2, \dots, C_n) \triangleq \sup_{G \in \mathcal{N}(C_1, C_2, \dots, C_n)} \min_{\substack{\widehat{G} \subset G \\ \widehat{G} \in \mathcal{N}(C_1, C_2, \dots, C_n)}} \mathcal{H}(\widehat{G}),$$




Figure 1: An illustrative graph in $\mathcal{N}(2, 2)$ with 8 hubs.

where $\mathcal{H}(\widehat{G})$ denotes the number of hubs in \widehat{G} . The above definition can be roughly interpreted as follows: for a given G , we try to find a subgraph \widehat{G} that contains the minimum number of hubs required to satisfy the vertex-cut constraints, and $\mathcal{H}(C_1, C_2, \dots, C_n)$ gives us the minimum number corresponding to the worst-case scenarios among all possible G .

At first glance, $\mathcal{H}(C_1, C_2, \dots, C_n)$ can be infinite. One of the main (and somewhat surprising) results in this paper, Theorem 6.1, however, states that for any given C_1, C_2, \dots, C_n , $\mathcal{H}(C_1, C_2, \dots, C_n)$ is in fact finite. With finiteness confirmed, we are primarily interested in computing the value of $\mathcal{H}(C_1, C_2, \dots, C_n)$. We say a graph G is *minimal* in $\mathcal{N}(C_1, C_2, \dots, C_n)$, if $G \in \mathcal{N}(C_1, C_2, \dots, C_n)$, but for any $e \in E$, $G \setminus \{e\} \notin \mathcal{N}(C_1, C_2, \dots, C_n)$. It then follows from the fact that every graph in $\mathcal{N}(C_1, C_2, \dots, C_n)$ has at least one minimal subgraph that

$$\mathcal{H}(C_1, C_2, \dots, C_n) = \sup_{\substack{G \in \mathcal{N}(C_1, C_2, \dots, C_n) \\ G \text{ is minimal}}} \mathcal{H}(G).$$

The vertex-connectivity version of the classical Menger's theorem [4] states that for a network with one pair of source S and sink R with the minimum vertex-cut between them being C , there exist C vertex-disjoint paths connecting S and R , which immediately implies that $\mathcal{H}(C) = 0$ for any given C . Theorem 6.1 states that for any given $G \in \mathcal{N}(C_1, C_2, \dots, C_n)$, one can always find a subgraph \widehat{G} of G such that $\mathcal{H}(\widehat{G})$ is upper bounded by a constant, which is independent of the choice of G . In some sense, Theorem 6.1 can be viewed as a generalization of the vertex-connectivity version of Menger's theorem.

Mathematically, the proposed problem of computing $\mathcal{H}(C_1, C_2, \dots, C_n)$ is a natural combinatorial optimization problem. On a more practical side, hubs in networks naturally correspond to more costly vertices. For instance, in a transportation network, as opposed to "relaying" vertices with degree 2, hubs may have to be better equipped for traffic scheduling; for this reason, when designing the route map, an airline may need to avoid running too many airline hubs to reduce the cost. So, as might be expected, $\mathcal{H}(C_1, C_2, \dots, C_n)$ is of significance to cost-minimizing resource allocation in transportation networks.

To the best of our knowledge, the proposed problem of computing or estimating \mathcal{H} has not yet been examined previously and there is little related work in the vast literature of graph theory. On the other hand, to a great extent, this work is motivated by the study of network encoding complexity (see [3] and references therein), where the number of encoding vertices in directed networks is of primary concern. Moreover, our approaches to tackle the problem are influenced by those in network encoding complexity theory, particularly, to a greater extent, those in [2, 5].

The remainder of the paper is organized as follows. In Section 2, we give necessary and sufficient conditions (Theorem 2.1) for a graph being minimal in $\mathcal{N}(C_1, C_2)$. In Section 3, we introduce the notion of a representation of a graph in $\mathcal{N}(C_1, C_2)$ and we present the structural decomposition theorem (Theorem 3.4) for representations of minimal graphs in $\mathcal{N}(C_1, C_2)$. We will introduce in Section 4 a novel path-searching algorithm, the analysis of which will aptly produce an upper bound on $\mathcal{H}(C_1, C_2)$ for any given C_1, C_2 . In Section 5, we derive the value of $\mathcal{H}(C_1, C_2)$ (Theorem 5.1), which is a main result in this paper. Another main result is Theorem 6.1, which establishes the finiteness of $\mathcal{H}(C_1, C_2, \dots, C_n)$, $n \geq 3$, through a recursive bounding argument. The remaining part of Section 6 will be devoted to the derivations of the values of \mathcal{H} for some special parameters (Theorem 6.2, 6.3).

2 Minimal Graphs in $\mathcal{N}(C_1, C_2)$

We first introduce some notation and terminologies that will be used throughout the paper.

A sequence of edges in G , $e_i = (u_i, v_i)$, $i = 1, 2, \dots, d$, with $v_i = u_{i+1}$, $i = 1, 2, \dots, d-1$, can be linked to form a *path* p , denoted by

$$p = e_1 \circ e_2 \circ \dots \circ e_d; \quad (1)$$

furthermore, p is called a *cycle* if $v_d = u_1$. A path p of the form (1) is said to be *directed* each e_i is oriented such that all of them concatenate.

For a path β and two vertices u, v of β , let $\beta[u, v]$ denote the subpath of β between u and v . For a directed path $\beta = e_1 \circ e_2 \circ \dots \circ e_n$, let $h(\beta)$, $t(\beta)$ denote the head, tail of β , respectively; and we say e_i is *smaller* than e_j on β , denoted by $e_i < e_j$, if $h(e_j)$ is the head of the directed path $\beta[e_i, e_j]$ (or alternatively, e_j is *larger* than e_i , denoted by $e_j > e_i$).

For a graph $G \in \mathcal{N}(C_1, C_2, \dots, C_n)$, by the vertex-connectivity version of Menger's theorem, for each i , one can find a set α_i of C_i vertex-disjoint ϕ -paths from S_i to R_i . If the choice of α_i is unique, α_i is said to be *non-reroutable*, otherwise it is said to be *reroutable*. G is said to be *non-reroutable* if all α_i 's are non-reroutable, *reroutable* otherwise. For any index set $\{j_1, j_2, \dots, j_k\} \subset \{1, 2, \dots, n\}$, let $G_{\alpha_{j_1} \alpha_{j_2} \dots \alpha_{j_k}}$ denote the subgraph of G induced on the edges of $\alpha_{j_1}, \alpha_{j_2}, \dots, \alpha_{j_k}$ -paths. G is said to be a (C_1, C_2, \dots, C_n) -*graph* if $G = \bigcup_{i=1}^n G_{\alpha_i}$, that is, each edge in G belongs to some α_i -path. In order to compute $\mathcal{H}(C_1, C_2, \dots, C_n)$, it is enough to only consider all (C_1, C_2, \dots, C_n) -graphs, since $\bigcup_{i=1}^n G_{\alpha_i}$ is a subgraph of G and also in $\mathcal{N}(C_1, C_2, \dots, C_n)$.

Sections 2 to 5 will be devoted to derive the value of $\mathcal{H}(C_1, C_2)$. For notational convenience only, we often rewrite α_1, α_2 as ϕ, ψ , respectively.

For a (C_1, C_2) -graph G , an edge in G is said to be *public* if it is shared by a ϕ -path and a ψ -path, *private* otherwise. Evidently, for each i , from S_i to R_i , each ϕ or ψ -path in G induces a natural orientation to all its edges. We note that a public edge in G may have opposite ϕ -direction and ψ -direction (such "inconsistency" will be dealt with in Section 3).

We say G can be *naturally orientable* if each public edge in G has the same natural ϕ and ψ -direction. A cycle $e_1 \circ e_2 \circ \dots \circ e_d$ in G , where $e_i = (u_i, v_i)$, $v_i = u_{i+1}$ for $i = 1, 2, \dots, d-1$ and $v_d = u_1$, is said to be a *ϕ -consistent cycle*, if it satisfies the following property: for any $1 \leq i \leq d$, if e_i belongs to a ϕ -path, then its natural ϕ -direction is from u_i to v_i . And we similarly define a *ψ -consistent cycle*.

The following theorem gives necessary and sufficient conditions for a (C_1, C_2) -graph being minimal.

Theorem 2.1. *The following three statements are equivalent for a (C_1, C_2) -graph G :*

- (i) G is minimal;

(ii) G is non-reroutable;

(iii) G has no ϕ or ψ -consistent cycle.

Proof. **1. (ii) \Rightarrow (i):** Any edge e in G must belong to some ϕ or ψ -path. After deleting e from G , we no longer find C_i vertex-disjoint paths from S_i to R_i for some $i \in \{1, 2\}$. So $G \setminus \{e\} \notin \mathcal{N}(C_1, C_2)$, and therefore G is minimal.

2. (ii) \Rightarrow (iii): Suppose, by way of contradiction, that G has a ϕ -consistent cycle O , which can be written in the following form

$$p_1 \circ e_1 \circ p_2 \circ e_2 \circ \cdots \circ p_d \circ e_d,$$

where each p_i is a subpath of some ϕ -path, and each e_i is a private ψ -edge. Furthermore, we assume O has the smallest d (the number of private ψ -edges) among all ϕ -consistent cycles. Then, each p_i belongs to a different ϕ -path (since otherwise we can always find a ϕ -consistent cycle with fewer private ψ -edges), which further implies that $d \leq C_1$. Suppose $p_i = \phi_i[u_i, v_i]$ for $1 \leq i \leq d$, $e_i = (v_i, u_{i+1})$ for $1 \leq i \leq d-1$ and $e_d = (v_d, u_1)$. Then one can find another group of C_1 vertex-disjoint paths $\hat{\phi} = \{\hat{\phi}_1, \hat{\phi}_2, \dots, \hat{\phi}_{C_1}\}$ from S_1 to R_1 in G , where

$$\hat{\phi}_i = \begin{cases} \phi_1[S_1, u_1] \circ (u_1, v_d) \circ \phi_d[u_d, R_1] & \text{for } i = 1; \\ \phi_i[S_1, u_i] \circ (u_i, v_{i-1}) \circ \phi_{i-1}[v_{i-1}, R_1] & \text{for } 2 \leq i \leq d; \\ \phi_i & \text{for } d+1 \leq i \leq C_1, \end{cases}$$

which contradicts the assumption that G is non-reroutable. With a parallel argument, we conclude that ψ -consistent cycles do not exist either.

3. (i) or (iii) \Rightarrow (ii): Suppose, by contradiction, that G is reroutable, and furthermore, by symmetry, that there exists another group of C_1 vertex-disjoint paths $\hat{\phi} = \{\hat{\phi}_1, \hat{\phi}_2, \dots, \hat{\phi}_{C_1}\}$ from S_1 to R_1 in G with $\hat{\phi}_i$ sharing the same outgoing edge from S_1 as ϕ_i for every i . Pick a $\hat{\phi}$ -path, say, $\hat{\phi}_{i_1}$, such that $\hat{\phi}_{i_1} \neq \phi_{i_1}$, and let v_{i_1} denote the smallest vertex on ϕ_{i_1} (under the natural ϕ -direction) where they leave each other. Assume that, after v_{i_1} , $\hat{\phi}_{i_1}$ first meets some ϕ -path, say, ϕ_{i_2} at the vertex u_{i_1} . Denote by v_{i_2} the smallest vertex where $\hat{\phi}_{i_2}$ and ϕ_{i_2} leave each other. Assume that, after v_{i_2} , $\hat{\phi}_{i_2}$ first meets some ϕ -path, say, ϕ_{i_3} at the vertex u_{i_2} . Continue the procedure in a similar manner to obtain an index sequence $i_1, i_2, \dots, i_t, \dots$, and similarly define v_{i_t} 's and u_{i_t} 's. Pick the smallest k such that $i_k = i_j$ for some $j < k$. Notice that $v_{i_{t+1}}$ is smaller than u_{i_t} on $\phi_{i_{t+1}}$, which easily follows from three facts:

- 1) $\hat{\phi}_{i_{t+1}}$ first takes apart from $\phi_{i_{t+1}}$ at $v_{i_{t+1}}$;
- 2) $\hat{\phi}_{i_t}$ meets $\phi_{i_{t+1}}$ at u_{i_t} ;
- 3) $\hat{\phi}_{i_t}$ and $\hat{\phi}_{i_{t+1}}$ are vertex-disjoint.

Hence, we conclude that

$$\begin{aligned} & \hat{\phi}_{i_j}[v_{i_j}, u_{i_j}] \circ \phi_{i_{j+1}}[u_{i_j}, v_{i_{j+1}}] \circ \hat{\phi}_{i_{j+1}}[v_{i_{j+1}}, u_{i_{j+1}}] \circ \phi_{i_{j+2}}[u_{i_{j+1}}, v_{i_{j+2}}] \circ \\ & \cdots \circ \hat{\phi}_{i_{k-2}}[v_{i_{k-2}}, u_{i_{k-2}}] \circ \phi_{i_{k-1}}[u_{i_{k-2}}, v_{i_{k-1}}] \circ \hat{\phi}_{i_{k-1}}[v_{i_{k-1}}, u_{i_{k-1}}] \circ \phi_{i_j}[u_{i_{k-1}}, v_{i_j}] \end{aligned}$$

is a ϕ -consistent cycle in G .

Since all ψ -paths are vertex-disjoint and the above cycle does not contain any terminal vertex, at least one edge e in this cycle does not belong to any ψ -path (that is, e is a private ϕ -edge). Notice that each edge of $\hat{\phi}_{i_t}[v_{i_t}, u_{i_t}]$, $j \leq t \leq k-1$, is a private ψ -edge, and each edge of $\phi_{i_{t+1}}[u_{i_t}, v_{i_{t+1}}]$,

$j \leq t \leq k - 1$, does not belong to any $\hat{\phi}$ -path. So, e must belong to $\phi_{i_{t+1}}[u_{it}, v_{i_{t+1}}]$ for some t with $j \leq t \leq k - 1$, and hence e does not belong to any $\hat{\phi}$ -path. In the graph $G \setminus \{e\}$, we can find a set $\hat{\phi}$ of C_1 vertex-disjoint paths from S_1 to R_1 , and a set ψ of C_2 vertex-disjoint paths from S_2 to R_2 . Thus $G \setminus \{e\} \in \mathcal{N}(C_1, C_2)$ and G is not minimal, a contradiction. \square

Remark 2.2. *The third part of the proof of the theorem has actually proved that for a (C_1, C_2) -graph G , if ϕ (resp. ψ) is reroutable, then there exists a private ϕ (resp. ψ)-edge e such that $G \setminus \{e\} \in \mathcal{N}(C_1, C_2)$. This fact will be used later in the paper.*

3 Representations and Structural Decomposition

In this section, we will transform a minimal (C_1, C_2) -graph G into G° , its representation, through the following steps:

Step 1 [Remove Relays]: In this step, we remove all non-terminal vertices in G with degree 2. In more detail, for any non-terminal vertex v with $d(v) = 2$ and $e_1 = (v, u_1)$, $e_2 = (v, u_2)$ being the two edges incident with v , where $u_1 \neq u_2$, delete edges e_1 , e_2 and vertex v , and then add a new edge (u_1, u_2) . Let G_1 denote the resulting graph.

Step 2 [Stretch Crossings]: We say a ϕ -path and a ψ -path *cross* at vertex v if they share v , but not any edges incident with v . In this step, we convert each crossing into a pair of degree 3 vertices. In more detail, for any vertex v in G_1 with $d(v) = 4$ and $e_1 = (v, u_1)$, $e_2 = (v, u_2)$ being the two ϕ -edges incident with v , $e_3 = (v, u_3)$, $e_4 = (v, u_4)$ being the two ψ -edges incident with v , where all u_i are all distinct, delete edges e_1 , e_2 , e_3 , e_4 and vertex v , and then add two vertices v_1 , v_2 and edges (v_1, v_2) , (u_1, v_1) , (u_3, v_1) , (u_2, v_2) and (u_4, v_2) . Let G_2 denote the resulting graph.

Step 3 [Match Directions]: For any public edge e in G_2 with inconsistent ϕ and ψ -direction, we will perform the following operations to obtain consistency: Assume edge $e = (u, v)$ belongs to both path $\phi_i = \phi_i[S_1, w_1] \circ (w_1, u) \circ (u, v) \circ (v, w_2) \circ \phi_i[w_2, R_1]$ and path $\psi_j = \psi_j[S_2, w_3] \circ (w_3, v) \circ (v, u) \circ (u, w_4) \circ \psi_j[w_4, R_2]$. We delete edges (w_3, v) , (u, w_4) , add edges (w_3, u) , (v, w_4) , and then obtain a new ψ -path

$$\psi_j[S_2, w_3] \circ (w_3, u) \circ (u, v) \circ (v, w_4) \circ \psi_j[w_4, R_2].$$

Let G° denote the resulting graph, which, evidently, is naturally orientable; let \vec{G}° denote the directed version of G° , equipped with the consistent natural orientation induced on all its ϕ and ψ -paths. Apparently, $G^\circ \in \mathcal{N}(C_1, C_2)$ with its ϕ and ψ -paths determined by the original ones, and all the non-terminal vertices in G° are hubs.

The obtained G° after the above three steps is said to be a *representation* of G . The following lemma states that G° must be a minimal (C_1, C_2) -graph as well.

Lemma 3.1. *The representation G° of a minimal (C_1, C_2) -graph G is also minimal.*

Proof. First, since G is minimal, by Theorem 2.1, G is non-reroutable. Evidently, G_1 is also non-reroutable, and thus minimal. By way of contradiction, suppose that G° is not minimal. Again, by Theorem 2.1, G° is reroutable. Notice that in Step 2 and 3, both crossings or inconsistently oriented public edges are converted into consistently oriented public edges. Then, by Remark 2.2, one can find a private edge e such that $G^\circ \setminus \{e\} \in \mathcal{N}(C_1, C_2)$, implying $G_1 \setminus \{e\}$ also belongs to $\mathcal{N}(C_1, C_2)$, which contradicts the minimality of G_1 . \square

Let $\mathcal{N}^\circ(C_1, C_2)$ denote the subset of $\mathcal{N}(C_1, C_2)$ consisting of all networks G such that G is minimal, naturally orientable and all the non-terminal vertices in G are of degree 3. Apparently, $\mathcal{N}^\circ(C_1, C_2)$ is the set of the representations of all minimal (C_1, C_2) -graphs. The following theorem says that in order to compute $\mathcal{H}(C_1, C_2)$, it is enough to only consider all graphs in $\mathcal{N}^\circ(C_1, C_2)$.

Theorem 3.2.

$$\mathcal{H}(C_1, C_2) = \sup_{G \in \mathcal{N}^\circ(C_1, C_2)} \mathcal{H}(G).$$

Proof. The “ \leq ” direction follows from the observation that

$$\mathcal{H}(G) = \mathcal{H}(G_1) \leq \mathcal{H}(G_2) = \mathcal{H}(G^\circ),$$

and the “ \geq ” direction immediately follows from Lemma 3.1. \square

A path in $G^\circ \in \mathcal{N}^\circ(C_1, C_2)$ is said to be an *alternating* path if all its edges are private edges and the terminal pair of this path is one of the following: (S_1, S_2) , (S_1, R_1) , (R_2, S_2) or (R_2, R_1) .

Lemma 3.3. *An alternating path has the following properties:*

1. *Each of its ϕ -edges is only adjacent to ψ -edges, and each of its ψ -edges is only adjacent to ϕ -edges.*
2. *Each of its ϕ (resp. ψ)-edge belongs to a different ϕ (resp. ψ)-path.*

Proof. 1. This follows from the fact that after Step 1, vertices with degree 2 have been removed and thus no two private ϕ (or ψ)-edges are adjacent.

2. We show that in any alternating path, each ϕ -edge belongs to a different ϕ -path. Suppose, by contradiction, that for an alternating path $e_1 \circ e_2 \circ \dots \circ e_d$, where $e_i = (u_i, v_i)$, two edges e_k, e_l , $k < l$, both belong to ϕ_t . If e_l is smaller than e_k on path ϕ_t in \vec{G}° , then

$$\phi_t[u_l, v_k] \circ e_{k+1} \circ e_{k+2} \circ \dots \circ e_{l-1}$$

is a ϕ -consistent cycle in G° , which, by Theorem 2.1, gives us a contradiction. If e_k is smaller than e_l on path ϕ_t in \vec{G}° , then the subpath $\phi_t[v_k, u_l]$ in fact can be expressed as $f_1 \circ f_2 \circ \dots \circ f_{2p-1}$, where $f_1, f_3, \dots, f_{2p-1}$ are public and $f_2, f_4, \dots, f_{2p-2}$ are private. Then

$$(e_{k+1} \circ f_1) \circ f_2 \circ \dots \circ f_{2p-2} \circ (f_{2p-1} \circ e_{l-1}) \circ e_{l-2} \circ \dots \circ e_{k+2}$$

is a ψ -consistent cycle in G° , which, by Theorem 2.1, gives us a contradiction. With a parallel argument, we conclude that each private ψ -edge belongs to a different ψ -path. \square

We next present the structural decomposition theorem of a representation $G^\circ \in \mathcal{N}^\circ(C_1, C_2)$, which, roughly speaking, states that after deleting public edges in G° , each connected component in the resulting graph is an alternating path. More precisely, letting G_p° denote the subgraph of G° induced on all private edges in G° , we have

Theorem 3.4. *G_p° consists of $(C_1 + C_2)$ alternating paths.*

Proof. Consider \vec{G}_p° , the naturally oriented version of G_p° . Obviously, the degree of any non-terminal vertex in \vec{G}_p° is 2. Now, starting from S_1 , traverse along an outgoing ϕ -edge, say, e_1 , and then traverse against the ψ -edge adjacent to e_1 , say, e_2 , and then along a ϕ -edge adjacent to e_2 , say, e_3 , and then against a ψ -edge adjacent to e_3 Continue the procedure in this fashion, we will always reach R_1 or S_2 , since otherwise the set of edges that we have traversed will contain a cycle, which is both ϕ -consistent and ψ -consistent. Evidently, a similar argument can be applied to the case when we start from an incoming edge incident with R_2 . It then follows that one can find a set of $(C_1 + C_2)$ vertex-disjoint paths in G_p° from $\{S_1, R_2\}$ to $\{S_2, R_1\}$; moreover, it can easily be checked that each edge in G_p° belongs to one of the above-mentioned vertex-disjoint paths. \square

Remark 3.5. Let v be a non-terminal vertex of an alternating path in G_p° and let e, e' be the two edges incident with v . It then follows from the proof of Theorem 3.4 that if v is the head (resp. tail) of e in \vec{G}_p° , then it is also the head (resp. tail) of e' .

An alternating path in G_p° is said to be an $S_1S_2, S_1R_1, R_2S_2, R_2R_1$ -alternating path if its terminal pair is $(S_1, S_2), (S_1, R_1), (R_2, S_2)$ or (R_2, R_1) , respectively. A path is said to be an S_1 -alternating path if it is either an S_1S_2 -alternating path or S_1R_1 -alternating path, similarly, a path is also referred to as an R_2 -alternating path if it is either an R_2S_2 -alternating path or R_2R_1 -alternating path. Apparently, there are C_1 S_1 -alternating paths and C_2 R_2 -alternating paths. For two vertices u, v of an S_1 (resp. R_2)-alternating path L , we say u is on the right of v in L if v is “nearer” to S_1 (resp. R_2) than u in L (more precisely, the number of edges between u and S_1 (resp. R_2) in L is more than that between v and S_1 (resp. R_2)) (the word “right” arises since we will “position” the vertices of a path in a plane in Section 4 for an easy illustration). For two edges e, e' in L , we say e is on the right of e' if one of two vertices incident with e is on the right of the two ones incident with e' in L .

Example 3.6. Figure 1 shows a naturally oriented minimal $(2, 2)$ -graph with 12 hubs. From S_1 to R_1 , there are two vertex-disjoint ϕ -paths $e_1 \circ e_3 \circ e_7 \circ e_9 \circ e_{13}$ and $e_4 \circ e_8 \circ e_{10} \circ e_{14} \circ e_{16}$; and from S_2 to R_2 , there are two vertex-disjoint ψ -paths $e_2 \circ e_3 \circ e_6 \circ e_8 \circ e_{12}$ and $e_5 \circ e_9 \circ e_{11} \circ e_{14} \circ e_{15}$. The edges e_3, e_8, e_9, e_{14} are public since each of them is shared by some ϕ -path and some ψ -path. The others are all private. Then we can find that G_p° consists of four alternating paths: two S_1 -alternating paths

$$e_1 \circ e_2, e_4 \circ e_6 \circ e_7 \circ e_5,$$

and two R_2 -alternating paths

$$e_{15} \circ e_{16}, e_{12} \circ e_{10} \circ e_{11} \circ e_{13}.$$

Moreover, $e_4 \circ e_6 \circ e_7 \circ e_5$ is an S_1S_2 -alternating path, where edge e_7 is on the right of edge e_6 , vertex B is on the right of vertex C .

4 A Path-Searching Algorithm

In this section, we introduce an algorithm, the analysis of which will be instrumental for deriving the value of $\mathcal{H}(C_1, C_2)$. Before rigorously describing the algorithm, we roughly illustrate its idea.

Consider a representation G° of a minimal (C_1, C_2) -graph G . Imagine on a two-dimensional plane, each alternating path $e_1 \circ e_2 \circ \dots \circ e_d$ in G_p° is “positioned” into a “double deck” (see Figure 2 for an example), where e_{i+1} is on the right of e_i , and all the tails (in \vec{G}_p°) are on the upper deck and heads on the lower deck (see Remark 3.5). A vertex on the upper deck is referred to as an *upper vertex* if it is not a source, and a vertex on the lower deck is referred to as a *lower vertex* if it is not a sink. Notice that there is one more lower vertices than upper vertices in an S_1S_2 -alternating path, while one less in an R_2R_1 -alternating path, and there are equally many of them in an S_1R_1 -alternating path or an R_2S_2 -alternating path.

In each S_1S_2 -alternating (resp. R_2R_1 -alternating) path, one of the lower (resp. upper) vertices will be labeled as a *choke vertex* (which could be a “bottleneck” for the path extending procedure in the algorithm), which is initially the rightmost lower (resp. rightmost upper) vertex of the path and will be updated during the execution of the algorithm.

As summarized below, the algorithm will iteratively find a set of the so-called *interconnecting* paths in G° , which will “link” the double decks.




Figure 2: In this R_2R_1 -alternating path, x_1, x_2, x_3 are upper vertices and y_1, y_2 are lower vertices.

- In the beginning of each iteration, we initialize the interconnecting path as a public edge, whose tail is an unoccupied lower vertex. See STEP 2.
- We first traverse along edges in \vec{G}° to extend the interconnecting paths to link the double decks. When we reach a choke vertex, we check if it is possible to switch paths (see Figure 3) and find new ways to extend the interconnecting paths. See STEP 3(A), 3(B).
- Then we traverse against edges in \vec{G}° to extend the interconnecting paths in a parallel manner. See STEP 5(A), 5(B).
- At the end of k -th iteration, a set of k interconnecting paths is found, whose vertices are all labeled as occupied. See STEP 6.

As proven later, the algorithm will produce all interconnecting paths and upon termination of the algorithm, all hubs in G° are occupied.

Let Δ be the number of S_1S_2 -alternating paths, then $\Delta \leq \min\{C_1, C_2\}$, and moreover, the numbers of R_2R_1 -alternating paths, S_1R_1 -alternating paths and R_2S_2 -alternating paths are $\Delta, C_1 - \Delta, C_2 - \Delta$, respectively. The following algorithm will find Δ interconnecting paths in G° .

Algorithm 4.1. *Input:* G° ; *Output:* a set \mathcal{I} of Δ interconnecting paths.

STEP 1: Initialize the algorithm.

Orient each edge in the direction of ϕ and ψ -paths;
 Label all hubs as “unoccupied”;
FOR each S_1S_2 -alternating path **DO**
 Its choke vertex := its rightmost lower vertex;
FOR each R_2R_1 -alternating path **DO**
 Its choke vertex := its rightmost upper vertex;
 $\mathcal{I} := \emptyset$;
 $n := 1$.

STEP 2: Initialize the forward extension.

$v :=$ an arbitrarily picked unoccupied lower vertex and label v as “occupied”;
 $f :=$ the public edge whose tail is v ;
 $u := h(f)$ and label u as “occupied”;
 $P := f$.

STEP 3(A): Prepend a private edge.

$L :=$ the alternating path to which u belongs;

IF L is an R_2R_1 -alternating path **AND** u is the choke vertex of L **THEN BEGIN**

IF there are no unoccupied upper vertices in L **THEN**

Go to **STEP 4**;

ELSE BEGIN

$x_0 :=$ the rightmost unoccupied upper vertex of L ;

$(x_0, y_0) \circ (y_0, x_1) \circ \dots \circ (x_d, y_d) \circ (y_d, u) :=$ the subpath of L between x_0 and u ;

IF $d = 0$ **THEN BEGIN**

The choke vertex of $L := x_0$;

Label y_0 as “occupied”;

$e := (u, y_0)$;

$P := P \circ e$;

END

ELSE (*viz.* $d > 0$) **BEGIN**

FOR $i := 1$ **TO** d **DO**

$P_i :=$ the interconnecting path in \mathcal{I} containing (x_i, y_i) ;

$\mathcal{I} := \mathcal{I} \setminus \{P_1, P_2, \dots, P_d\}$;

The choke vertex of $L := x_0$;

Label y_0 as “occupied”;

$e := (x_1, y_0)$;

$\hat{P} := P$;

$P := P_1[t(P_1), x_1] \circ e$;

FOR $i := 1$ **TO** $d - 1$ **DO**

$P_i := P_{i+1}[t(P_{i+1}), x_{i+1}] \circ (x_{i+1}, y_i) \circ P_i[y_i, h(P_i)]$;

$P_d := \hat{P}[t(\hat{P}), u] \circ (u, y_d) \circ P_d[y_d, h(P_d)]$;

$\mathcal{I} := \mathcal{I} \cup \{P_1, P_2, \dots, P_d\}$;

END

END

END

ELSE BEGIN

IF $u \in$ an S_1 -alternating path **THEN**

$e :=$ the private ψ -edge whose tail is u ;

ELSE (*viz.* $u \in$ an R_2 -alternating path)

$e :=$ the private ϕ -edge whose tail is u ;

Label $h(e)$ as “occupied”;

$P := P \circ e$.

END

STEP 3(B): Prepend a public edge.

$f :=$ public edge whose tail is $h(e)$;
 $u := h(f)$ and label u as “occupied”;
 $P := P \circ f$;
 Go to **STEP 3(A)**.

STEP 4: Initialize the backward extension.

$\mathcal{I} := \mathcal{I} \cup \{P\}$;
 $P :=$ the interconnecting path in \mathcal{I} whose tail is v ;
 $\mathcal{I} := \mathcal{I} \setminus \{P\}$;
 $w := v$.

STEP 5(A): Append a private edge.

$L :=$ the alternating path to which w belongs;

IF L is an S_1S_2 -alternating path **AND** w is the choke vertex of L **THEN BEGIN**

IF there is no unoccupied lower vertex in L **THEN**

Go to **STEP 6**;

ELSE BEGIN

$y_0 :=$ the rightmost unoccupied lower vertex of L ;

$(y_0, x_0) \circ (x_0, y_1) \circ \cdots \circ (y_d, x_d) \circ (x_d, w) :=$ the subpath of L between y_0 and w ;

IF $d = 0$ **THEN BEGIN**

The choke vertex of $L := y_0$;

Label x_0 as “occupied”;

$e := (x_0, w)$;

$P := e \circ P$;

END

ELSE (*viz.* $d > 0$) **BEGIN**

FOR $i := 1$ **TO** d **DO**

$P_i :=$ the interconnecting path in \mathcal{I} containing (x_i, y_i) ;

$\mathcal{I} := \mathcal{I} \setminus \{P_1, P_2, \dots, P_d\}$;

The choke vertex of $L := y_0$;

Label x_0 as “occupied”;

$e := (x_0, y_1)$;

$\hat{P} := P$;

$P := e \circ P_1[y_1, h(P_1)]$;

FOR $i := 1$ **TO** $d - 1$ **DO**

$P_i := P_i[t(P_i), x_i] \circ (x_i, y_{i+1}) \circ P_{i+1}[y_{i+1}, h(P_{i+1})]$;

$P_d := P_d[t(P_d), x_d] \circ (x_d, w) \circ \hat{P}[w, h(\hat{P})]$;

$\mathcal{I} := \mathcal{I} \cup \{P_1, P_2, \dots, P_d\};$

END

END

END

ELSE BEGIN

IF $w \in$ an R_2 -alternating path **THEN**

$e :=$ the private ϕ -edge whose head is w ;

ELSE (*viz.* $w \in$ an S_1 -alternating path)

$e :=$ the private ψ -edge whose head is w ;

Label $t(e)$ as “occupied”;

$P := e \circ P$.

END

STEP 5(B): Append a public edge.

$f :=$ public edge whose head is $t(e)$;

$w := t(f)$ and label w as “occupied”;

$P := f \circ P$;

Go to **STEP 5(A)**.

STEP 6: Terminate the iteration.

IF $n = \Delta$ **THEN**

Terminate the algorithm;

$\mathcal{I} := \mathcal{I} \cup \{P\}$;

$n := n + 1$;

Go to **STEP 2**.

Remark 4.2. 1. For Step 2 in each iteration, an unoccupied lower vertex always exists since there is at least one S_1S_2 -alternating path with its chock vertex left unoccupied.

2. Each iteration of Algorithm 4.1 consists of forward and backward extensions, and each interconnecting path starts from an S_1S_2 -alternating path and ends at an R_2R_1 -alternating path. The algorithm cannot be simplified into a “one-direction extension” version: for each iteration, the extending procedure will terminate at an S_1S_2 -alternating path with only one unoccupied lower vertex, but, in the beginning of each iteration, such an alternating path may not exist (to see this, notice that more than two unoccupied vertices of an alternating path can become “occupied”, if we switch paths during one iteration).

We will need the following two lemmas in the next section.

Lemma 4.3. *For a representation G° of a minimal (C_1, C_2) -graph, upon the termination of Algorithm 4.1, each private edge of any interconnecting path belongs to a different alternating path.*




Figure 3: Switch interconnecting paths when a choke vertex is met.

Proof. By contradiction, suppose that for an interconnecting path $P = e_0 \circ e_1 \circ e_2 \circ \dots \circ e_{2n}$, two private edges e_{2k} and e_{2l} , $k < l$, belong to the same alternating path L . Let O be the cycle

$$P[h(e_{2k}), t(e_{2l})] \circ L[t(e_{2l}), h(e_{2k})].$$

If e_{2l} is on the right of e_{2k} in L , then O is a ψ -consistent cycle in G° . If e_{2k} is on the right of e_{2l} in L , then O is a ϕ -consistent cycle in G° . By Theorem 2.1, both cases imply that G° is not minimal, which, by Lemma 3.1, further implies that G is not minimal, a contradiction. \square

Lemma 4.4. *For a representation G° of a minimal (C_1, C_2) -graph G , upon the termination of Algorithm 4.1, \mathcal{I} consists of at most $\min\{C_1, C_2\}$ interconnecting paths, and each hub in G° is in exactly one of the interconnecting paths in \mathcal{I} .*

Proof. Each interconnecting path in \mathcal{I} starts from a lower vertex of a distinct S_1S_2 -alternating path and ends to an upper vertex of a distinct R_2R_1 -alternating path. So the number of interconnecting paths is Δ with $\Delta \leq \min\{C_1, C_2\}$. Now upon the termination of Algorithm 4.1, pick an unoccupied lower vertex, and then execute the algorithm from STEP 3(B); or pick an unoccupied upper vertex, and then execute the algorithm from STEP 3(A). Since all choke vertices have been unoccupied, the algorithm will fail to terminate, violating the fact that G° is finite. \square

5 $\mathcal{H}(C_1, C_2)$

Theorem 5.1.

$$\mathcal{H}(C_1, C_2) = 2C_1C_2.$$

Proof. The “ \leq ” direction: For a representation G° of a minimal (C_1, C_2) -graph G , apply Algorithm 4.1 to obtain a set \mathcal{I} of Δ interconnecting paths. In the t -th iteration of the algorithm, the forward extension stops at the choke vertex of some R_2R_1 -alternating path. Let L_t denote this alternating path. Note that after t iterations, 1) all hubs of L_t become occupied; 2) by Lemma 4.3, each of the t obtained interconnecting paths contains at most two hubs of L_t ; 3) one obtained interconnecting path contains exactly one hub of L_t . Hence, we deduce that the number of hubs of L_t is at most $2t - 1$, and therefore the total number of hubs of all R_2R_1 -alternating paths is at most $\sum_{t=1}^{\Delta} (2t - 1) = \Delta^2$. Similarly, the total number of hubs of all S_1S_2 -alternating paths is at most Δ^2 as well. Again, by Lemma 4.3, the total number of hubs of all S_1R_1 and R_2S_2 -alternating paths is at most $2\Delta(C_1 + C_2 - 2\Delta)$. By Lemma 4.4, each hub in G° belongs to some interconnecting path in \mathcal{I} . Therefore,

$$|\mathcal{H}(G^\circ)| \leq 2\Delta^2 + 2\Delta(C_1 + C_2 - 2\Delta) = 2\Delta(C_1 + C_2 - \Delta) \leq 2C_1C_2,$$

where the last inequality follows from $\Delta \leq \min\{C_1, C_2\}$.

The “ \geq ” direction: We only need to construct a minimal (C_1, C_2) -graph G with $2C_1C_2$ hubs (see Figure 4 for an example). The graph G can be described as follows: $G \in \mathcal{N}^\circ(C_1, C_2)$ is naturally oriented, and there is a set of C_1 vertex-disjoint paths $\phi = \{\phi_1, \phi_2, \dots, \phi_{C_1}\}$ from S_1 to R_1 and a set of C_2 vertex-disjoint paths $\psi = \{\psi_1, \psi_2, \dots, \psi_{C_2}\}$ from S_2 to R_2 . Furthermore, in \vec{G} , the directed version of G , paths ϕ_i and ψ_j meet at vertex $\lambda_{i,j}$ and depart at vertex $\mu_{i,j}$, for $1 \leq i \leq C_1, 1 \leq j \leq C_2$, and

- on path ϕ_i , $\lambda_{i,1} < \mu_{i,1} < \lambda_{i,2} < \mu_{i,2} < \dots < \lambda_{i,C_2} < \mu_{i,C_2}$;
- on path ψ_j , $\lambda_{1,j} < \mu_{1,j} < \lambda_{2,j} < \mu_{2,j} < \dots < \lambda_{C_1,j} < \mu_{C_1,j}$.

□




Figure 4: A minimal $(3,3)$ -graph with 18 hubs.

6 $\mathcal{H}(C_1, C_2, \dots, C_n)$

In this section, we are concerned with \mathcal{H} with more than two parameters, which turns out to be much more difficult to compute than \mathcal{H} with two parameters.

The following theorem establishes the finiteness of \mathcal{H} with more than two parameters.

Theorem 6.1. *For any given $C_1, C_2, \dots, C_n \in \mathbb{N}$,*

$$\mathcal{H}(C_1, C_2, \dots, C_n) < \infty.$$

Proof. We will use an inductive argument on n . Notice that the case when $n = 2$ has been established in Theorem 5.1. By way of induction, we assume that the theorem is true for $n \leq k-1$ and proceed to prove it for $n = k$.

Consider any minimal (C_1, C_2, \dots, C_k) -graph G . Let α_i denote the set of C_i vertex-disjoint paths from S_i to R_i for $1 \leq i \leq k$. After necessary rerouting of $\alpha_1, \alpha_2, \dots, \alpha_{k-1}$ within $G_{\alpha_1 \alpha_2 \dots \alpha_{k-1}}$, we can assume that $G_{\alpha_1 \alpha_2 \dots \alpha_{k-1}}$ is minimal and thus

$$N_1 \triangleq \mathcal{H}(G_{\alpha_1 \alpha_2 \dots \alpha_{k-1}}) \leq \mathcal{H}(C_1, C_2, \dots, C_{k-1}) < \infty.$$

Let \widehat{G} denote the subgraph of $G_{\alpha_1 \alpha_2 \dots \alpha_{k-1}}$ induced on all the edges, each of which is simultaneously an $\alpha_1, \alpha_2, \dots, \alpha_{k-1}$ -edge, and let ω denote the number of connected components in \widehat{G} . Obviously, each connected component of \widehat{G} is in fact a path, and therefore we have $\omega \leq N_1/2$.

In this proof, we say a hub in G is *new* if it is a hub in G , however, not one in $G_{\alpha_1 \alpha_2 \dots \alpha_{k-1}}$. And we say a new hub is *global*, if this hub belongs to \widehat{G} , *local*, if this hub is in $G \setminus \widehat{G}$.

Then, by the induction hypothesis, for any distinct $j_1, j_2, \dots, j_{k-2} \in \{1, 2, \dots, k-1\}$, we deduce that the number of new hubs in $G_{\alpha_{j_1} \alpha_{j_2} \dots \alpha_{j_{k-2}}} \setminus \widehat{G}$ is upper bounded by $\mathcal{H}(C_{j_1}, C_{j_2}, \dots, C_{j_{k-2}}, C_k)$. So the number of local new hubs is at most

$$N_2 \triangleq \sum_{i=1}^{k-1} \mathcal{H}(C_1, \dots, C_{i-1}, C_{i+1}, \dots, C_{k-1}, C_k) < \infty.$$

For each α_k -path, we “cut” it at each of its local new hubs and then “divide” the path into “segments”, each of which, evidently, is a subpath of the original α_k -path. Let $\check{\alpha}_k$ denote the set of all obtained subpaths. Then we have $N_3 \triangleq |\check{\alpha}_k| \leq C_k + N_2$.

From the subgraph $\widehat{G} \cup \check{\alpha}_k$, we construct an $(\underbrace{\omega, \omega, \dots, \omega}_{k-1}, N_3)$ -graph G' through the following procedure:

1. Add sources S'_1, S'_2, \dots, S'_k , sinks R'_1, R'_2, \dots, R'_k .
2. For any $i \in \{1, 2, \dots, k-1\}$ and for any connected component of \widehat{G} , whose natural α_i -direction is from one end vertex, say v_1 , to the other end vertex, say v_2 , add two directed edges (S'_i, v_1) and (v_2, R'_i) (notice that a connected component may have opposite α_{i_1} -direction and α_{i_2} -direction for different i_1, i_2). Then, for any $1 \leq i \leq k-1$, we obtain a group α'_i of ω vertex-disjoint paths from S'_i to R'_i .
3. For any $\check{\alpha}_k$ -path, whose natural α_k -direction is from one end vertex, say v_1 , to the other end vertex, say v_2 , add two directed edges (S'_k, v_1) and (v_2, R'_k) . Then, we obtain a group α'_k of N_3 vertex-disjoint paths from S'_k to R'_k .

Obviously, the number of global new hubs in G is just $\mathcal{H}(G')$. It follows from the minimality of G and the observation that for any $1 \leq i \leq k$, any α'_i -consistent cycle in G' naturally corresponds an α_i -consistent cycle in G that G' is a minimal $(\underbrace{\omega, \omega, \dots, \omega}_{k-1}, N_3)$ -graph. We then proceed to deduce

that α'_k is non-reroutable in G' , since otherwise there exists an edge e which is exclusively owned by α'_k -paths (this follows from a parallel argument leading to Remark 2.2) such that $G' \setminus \{e\} \in \mathcal{N}(\underbrace{\omega, \omega, \dots, \omega}_{k-1}, N_3)$, violating the fact that G' is minimal.

Now, after necessary rerouting of α'_i within G' , we assume that G'_i , the subgraph of G' consisting of all α'_i -paths and α'_k -paths, is non-reroutable and thus, by Theorem 2.1, minimal. Notice that any hub v in G' must be a hub of some G'_i , since otherwise G' contains an edge incident with v that does not belong to any G'_i and therefore is not minimal. Hence, by Theorem 5.1, we obtain that

$$\mathcal{H}(G') \leq \sum_{i=1}^{k-1} \mathcal{H}(G'_i) \leq (k-1)\mathcal{H}(\omega, N_3) = 2(k-1)\omega N_3 \leq (k-1)N_1(C_k + N_2).$$

Therefore,

$$\mathcal{H}(G) \leq N_1 + N_2 + (k-1)N_1(C_k + N_2) < \infty.$$

The proof is then complete. \square

Theorem 6.2. *For any $n \geq 0$, $C_1, C_2 \geq 1$, we have*

$$\mathcal{H}(C_1, C_2, \underbrace{1, 1, \dots, 1}_n) = 2(C_1 C_2 + n).$$

Proof. The case when $n = 0$ is nothing but Theorem 5.1. So we only have to prove the theorem when $n \geq 1$.

The “ \leq ” direction: We will establish this direction using an inductive argument on n . Suppose the inequality holds when $n < t$. Consider a minimal $(C_1, C_2, \underbrace{1, 1, \dots, 1}_t)$ -graph G with C_1 vertex-disjoint paths $\phi_1, \phi_2, \dots, \phi_{C_1}$ from S_1 to R_1 , C_2 vertex-disjoint paths $\psi_1, \psi_2, \dots, \psi_{C_2}$ from S_2 to R_2 , and a path β_i from S_i to R_i for $3 \leq i \leq t+2$. Let G_1 be the subgraph of G induced on

$$\phi_1, \phi_2, \dots, \phi_{C_1}, \psi_1, \psi_2, \dots, \psi_{C_2}, \beta_3, \beta_4, \dots, \beta_{t+1}.$$

Now we split S_1 into C_1 copies $S_1^{(1)}, S_1^{(2)}, \dots, S_1^{(C_1)}$; R_1 into C_1 copies $R_1^{(1)}, R_1^{(2)}, \dots, R_1^{(C_1)}$; S_2 into C_2 copies $S_2^{(1)}, S_2^{(2)}, \dots, S_2^{(C_2)}$; R_2 into C_2 copies $R_2^{(1)}, R_2^{(2)}, \dots, R_2^{(C_2)}$, such that ϕ_i has starting point $S_1^{(i)}$ and ending point $R_1^{(i)}$ for $1 \leq i \leq C_1$; ψ_i has starting point $S_2^{(i)}$ and ending point $R_2^{(i)}$ for $1 \leq i \leq C_2$. Let G_2 denote the resulting graph and ω be the number of weakly connected components (which means connected components when disregarding the orientation) in G_2 . Now for any $1 \leq i < j \leq C_1$, we identify $S_1^{(i)}$ and $S_1^{(j)}$ if they belong to the same component and we perform similar operations on $R_1^{(i)}$'s, $S_2^{(i)}$'s and $R_2^{(i)}$'s. Let \widehat{G} denote the resulting graph, which consists of ω connected components. Note that the i -th component \widehat{G}_i is in fact a minimal $(C_{1,i}, C_{2,i}, \underbrace{1, 1, \dots, 1}_{t_i})$ -

graph, where $\sum_{i=1}^{\omega} C_{1,i} = C_1$, $\sum_{i=1}^{\omega} C_{2,i} = C_2$ and $\sum_{i=1}^{\omega} t_i = t - 1$. Notice that for any i , at least one of $C_{1,i}$, $C_{2,i}$ and t_i is nonzero, but some $C_{1,i}$'s, $C_{2,i}$'s, t_i 's can be zero, for which case, a $(C_{1,i}, C_{2,i}, \underbrace{1, 1, \dots, 1}_{t_i})$ -graph can be interpreted as if all zero-valued parameters are simply dropped.

For each component \widehat{G}_i , by the induction hypothesis, we have

$$\mathcal{H}(\widehat{G}_i) \leq \begin{cases} 2(C_{1,i}C_{2,i} + t_i) & \text{if } C_{1,i}, C_{2,i} \geq 1, \\ 2(C_{2,i} + t_i - 1) & \text{if } C_{1,i} = 0, C_{2,i} \geq 1, \\ 2(C_{1,i} + t_i - 1) & \text{if } C_{1,i} \geq 1, C_{2,i} = 0, \\ 2(t_i - 1) & \text{if } C_{1,i}, C_{2,i} = 0. \end{cases}$$

Notice that the right hand side of the above inequality can be unified as

$$2[C_{1,i} + C_{2,i} + (C_{1,i} - 1)^+ \cdot (C_{2,i} - 1)^+ + t_i - 1],$$

for any $C_{1,i}, C_{2,i} \geq 0$; here, $(\cdot)^+ = \max\{\cdot, 0\}$.

A hub in G is said to be *new* if it is not a hub in \widehat{G} . Notice that the path β_t “meets” each component at most once, yielding at most two new hubs (more precisely, it meets one edge in \widehat{G}_i , then departs from one edge in \widehat{G}_i , and thereafter, it will never meet any edge in \widehat{G}_i again), otherwise G is not minimal. Hence,

$$\begin{aligned} \mathcal{H}(G) &\leq 2\omega + \sum_{i=1}^{\omega} \mathcal{H}(\widehat{G}_i) \\ &\leq 2\omega + \sum_{i=1}^{\omega} 2[C_{1,i} + C_{2,i} + (C_{1,i} - 1)^+ \cdot (C_{2,i} - 1)^+ + t_i - 1] \\ &\leq 2\omega + 2(C_1 + C_2 + t - 1 - \omega) + 2\left[\sum_{i=1}^{\omega} (C_{1,i} - 1)^+\right] \cdot \left[\sum_{i=1}^{\omega} (C_{2,i} - 1)^+\right] \\ &\leq 2(C_1 + C_2 + t - 1) + 2(C_1 - 1)(C_2 - 1) \\ &= 2(C_1C_2 + t). \end{aligned}$$

The “ \geq ” direction: It suffices to construct a minimal $(C_1, C_2, \underbrace{1, 1, \dots, 1}_n)$ -graph with $2(C_1C_2 + n)$ hubs; see Figure 5 for an example. It turns out the graph G , described below in




Figure 5: A minimal $(2, 2, 1, 1)$ -graph with 12 hubs.

detail, is such a graph: G is naturally orientable; there is C_1 vertex-disjoint paths $\phi_1, \phi_2, \dots, \phi_{C_1}$ from S_1 to R_1 , C_2 vertex-disjoint paths $\psi_1, \psi_2, \dots, \psi_{C_2}$ from S_2 to R_2 , and a path β_i from S_i to

R_i for $3 \leq i \leq n+2$; paths ϕ_i and ψ_j meet at vertex $\lambda_{i,j}$ and depart at vertex $\mu_{i,j}$ for $1 \leq i \leq C_1$ and $1 \leq j \leq C_2$; paths β_i and ϕ_1 meet at vertex γ_i and depart at vertex δ_i for $3 \leq i \leq n+2$. Furthermore, in \vec{G} , we have

- for $3 \leq i \leq n+2$, on path β_i ,

$$\gamma_i < \delta_i;$$

- on path ϕ_1 ,

$$\gamma_3 < \delta_3 < \gamma_4 < \delta_4 < \cdots < \gamma_{n+2} < \delta_{n+2} < \lambda_{1,1} < \mu_{1,1} < \lambda_{1,2} < \mu_{1,2} < \cdots < \lambda_{1,C_2} < \mu_{1,C_2};$$

- for $2 \leq i \leq C_1$, on path ϕ_i ,

$$\lambda_{i,1} < \mu_{i,1} < \lambda_{i,2} < \mu_{i,2} < \cdots < \lambda_{i,C_2} < \mu_{i,C_2};$$

- for $1 \leq j \leq C_2$, on path ψ_j ,

$$\lambda_{1,j} < \mu_{1,j} < \lambda_{2,j} < \mu_{2,j} < \cdots < \lambda_{C_1,j} < \mu_{C_1,j}.$$

It can be easily checked that this graph is minimal. The proof is then complete. \square

Theorem 6.3.

$$\mathcal{H}(2, 2, 2) = 12.$$

Proof. First, it can be verified that the graph in Figure 6 is a minimal $(2, 2, 2)$ -graph with 12 hubs, which implies that

$$\mathcal{H}(2, 2, 2) \geq 12.$$




Figure 6: a minimal $(2, 2, 2)$ -graph with 12 hubs

So we only need to prove the other direction. By contradiction, suppose that a minimal $(2, 2, 2)$ -graph G has at least 13 hubs, and there exists a set of two vertex-disjoint paths $\phi = \{\phi_1, \phi_2\}$ from S_1 to R_1 , a set of two vertex-disjoint paths $\psi = \{\psi_1, \psi_2\}$ from S_2 to R_2 , and a set of two vertex-disjoint paths $\xi = \{\xi_1, \xi_2\}$ from S_3 to R_3 .




Figure 7: A minimal but reroutable $(2, 2, 2)$ -graph: the graph is minimal, but we can choose either $\{e_1 \circ e_3 \circ e_7, e_2 \circ e_6 \circ e_8\}$ or $\{e_1 \circ e_5 \circ e_8, e_2 \circ e_4 \circ e_7\}$ to be a set two vertex-disjoint paths from S_3 to R_3 .

Here, we note that when there are three pairs of sources and sinks, the equivalence statements as in Theorem 2.1 do not hold any more: it turns out that a minimal graph can be reroutable (see Figure 7).

We first consider the case G is minimal and non-reroutable. Let $G_{\phi\psi}$ be the subgraph of G induced on the edges of ϕ -paths and ψ -paths; similarly, we define $G_{\phi\xi}$ and $G_{\psi\xi}$. By Theorem 2.1, these three subgraphs are all minimal, since they are all non-reroutable. Suppose $G_{\phi\psi}$ has the most hubs among them. Every hub in G belongs to at least one of these three subgraphs, so we have

$$13 \leq \mathcal{H}(G) \leq \mathcal{H}(G_{\phi\psi}) + \mathcal{H}(G_{\phi\xi}) + \mathcal{H}(G_{\psi\xi}) \leq 3\mathcal{H}(G_{\phi\psi}),$$

and hence $\mathcal{H}(G_{\phi\psi}) \geq 5$.

Now, we transform $G_{\phi\psi}$ into a corresponding graph $G_{\phi\psi}^\bullet$ by shrinking each public edge into a vertex (this can be regarded as the “inverse” operation of Step 2 in Section 3). In more detail, for a public edge (v_1, v_2) , say $e_1 = (v_1, u_1), e_2 = (v_1, u_3)$ are two private edges incident with v_1 , and $e_3 = (v_2, u_2), e_4 = (v_2, u_4)$ are two private edges incident with v_2 . Then we delete edges e_1, e_2, e_3, e_4 and vertices v_1, v_2 , and then add a new vertex v and edges $(v, u_1), (v, u_2), (v, u_3), (v, u_4)$. If $G_{\phi\psi}$ has at least five hubs, then as shown in Figure 8, up to isomorphism, $G_{\phi\psi}^\bullet$ has three possibilities (note that the first and second are different).

Next, we examine in the ways one can add ξ -paths into $G_{\phi\psi}$ to form G such that G is minimal. In the following, a hub in G is said to be *new* if it is not a hub in $G_{\phi\psi}$, and an edge is said to be *new* if it does not belong to $G_{\phi\psi}$ and not incident with S_3 or R_3 .

Case 1: This case is shown in Figure 8(a), where each edge is labeled. Since $G_{\phi\psi}^\bullet$ has four hubs, $G_{\phi\psi}$ has at most eight hubs. And since $\mathcal{H}(G) \geq 13$, ξ_1, ξ_2 have to be added to generate at least 5 new hubs. Without loss of generality, say, ξ_1 contains at least three new hubs. Observe that each of these new hubs is incident with exactly one ξ_1 -edge that does not belong to any ϕ or ψ -path, and S_3 and R_3 are also incident with one such ξ_1 -edge. So at least $\lceil (3+2)/2 \rceil = 3$ edges of ξ_1 are exclusively owned by ξ_1 , and thereby at least one of them, say (v_1, v_2) , is a new edge (not incident with S_3 or R_3). We next discuss the possible locations of (v_1, v_2) ; here, notice that it is possible that either v_1 or v_2 is not a new hub.

Suppose $v_1 \in \phi_2^a$ (this means v_1 is one of vertices in ϕ_2^a , including its two end vertices). Then, one verifies that if $v_2 \in \phi_1^b, \phi_1^c, \phi_2^a, \phi_2^b, \phi_2^c, \psi_1^a, \psi_1^b, \psi_1^c, \psi_2^a$ or ψ_2^b , then ϕ is reroutable; if $v_2 \in \phi_1^a$ or ψ_2^c ,




Figure 8: Three possibilities of $G_{\phi\psi}^{\bullet}$

then ψ is reroutable. For example, if $v_2 \in \psi_2^c$, then we can find another two vertex-disjoint ψ -paths

$$\begin{aligned}\psi'_1 &= \psi_1, \\ \psi'_2 &= \psi_2[S_2, A] \circ \phi_2[A, v_1] \circ (v_1, v_2) \circ \psi_2[v_2, R_2]\end{aligned}$$

from S_2 to R_2 , and thus $G_{\phi\psi}$ is reroutable. And if $v_2 \in \phi_1^c$, then we can find another two vertex-disjoint paths

$$\begin{aligned}\phi'_1 &= \phi_1[S_1, B] \circ \psi_2[B, A] \circ \phi_2[A, R_1], \\ \phi'_2 &= \phi_2[S_1, v_1] \circ (v_1, v_2) \circ \phi_1[v_2, R_1]\end{aligned}$$

from S_1 to R_1 , and thus $G_{\phi\psi}$ is reroutable.

Suppose $v_1 \in \phi_2^b$. Then, one verifies that if $v_2 \in \phi_1^a, \phi_1^c, \phi_2^a, \phi_2^b, \phi_2^c, \psi_1^a, \psi_1^b, \psi_2^a$ or ψ_2^b , then ϕ is reroutable; and if $v_2 \in \phi_1^b, \psi_1^c$ or ψ_2^c , ψ is reroutable.

By symmetry, if $v_1 \in \phi_1^a, \phi_1^b, \phi_1^c, \phi_2^a, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b$ or ψ_2^c , $G_{\phi\psi}$ is reroutable.

Hence, we only need to check the last possible case: all new edges (including (v_1, v_2)) are incident with public edges in $G_{\phi\psi}$. By symmetry, say, $v_1 \in A$. Then, if $v_2 \in A$ or B , ψ is reroutable; if $v_2 \in C$, ϕ is reroutable; if $v_2 \in D$, there must be a new edge (v'_1, v'_2) of ξ_2 with $(v'_1 \in B$ and $v'_2 \in C)$ or $(v'_1 \in D$ and $v'_2 \in A)$ (see Figure 9(a)(b)), since otherwise $G \setminus \{(v_1, v_2)\}$ is still in $\mathcal{N}(2, 2, 2)$, which contradicts the fact that G is minimal. However, in both cases, we can still find another two vertex-disjoint paths

$$\begin{aligned}\phi'_1 &= \phi_1[S_1, v'_1] \circ (v'_1, v'_2) \circ \phi_2[v'_2, R_1], \\ \phi'_2 &= \phi_2[S_1, v_1] \circ (v_1, v_2) \circ \phi_1[v_2, R_1]\end{aligned}$$

from S_1 to R_1 , and thus $G_{\phi\psi}$ is reroutable.

Case 2: This case is shown in Figure 8(b). Similarly as Case 1, we can find a new edge (v_1, v_2) of ξ_1 .

Suppose $v_1 \in \phi_1^a$. Then, one verifies that if $v_2 \in \phi_1^a, \phi_1^b, \phi_2^b, \phi_2^c, \phi_2^d, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, \psi_2^c$ or ψ_2^d , then ϕ is reroutable; if $v_2 \in \phi_2^a$, then ψ is reroutable.

Suppose $v_1 \in \phi_2^a$. Then, one verifies that if $v_2 \in \phi_1^b, \phi_2^a, \phi_2^b, \phi_2^c, \phi_2^d, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, \psi_2^c$ or ψ_2^d , then ϕ is reroutable; if $v_2 \in \phi_1^a$, then ψ is reroutable.




Figure 9: The possible locations of the new edge (v'_1, v'_2) .

Suppose $v_1 \in \phi_2^b$. Then, one verifies that if $v_2 \in \phi_1^a, \phi_1^b, \phi_2^a, \phi_2^b, \phi_2^c, \phi_2^d, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, \psi_2^c$ or ψ_2^d , then $G_{\phi\psi}$ is reroutable.

By symmetry, if $v_1 \in \phi_1^b, \phi_2^c, \phi_2^d, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, \psi_2^c$ or ψ_2^d , then $G_{\phi\psi}$ is reroutable.

Hence, we only need to check the last possible case: all new edges (including (v'_1, v'_2)) are incident with public edges in $G_{\phi\psi}$. By symmetry, say $v_1 \in A$ or B . If $v_1 \in A$ and $v_2 \in A, C$ or D , then ϕ is reroutable; if $v_1 \in A$ and $v_2 \in B$, then ψ is reroutable; If $v_1 \in B$ and $v_2 \in A, B$ or D , ψ is reroutable; If $v_1 \in B$ and $v_2 \in C$, there must be a new edge (v'_1, v'_2) of ξ_2 with $v'_1 \in B$ and $v'_2 \in C$ (see Figure 9(c)), since otherwise $G \setminus \{(v_1, v_2)\}$ is still in $\mathcal{N}(2, 2, 2)$, which contradicts the fact that G is minimal. However, we can still find another two vertex-disjoint paths

$$\begin{aligned}\phi'_1 &= \phi_1[S_1, v_1] \circ (v_1, v_2) \circ \phi_2[v_2, R_1], \\ \phi'_2 &= \phi_2[S_1, v'_2] \circ (v'_2, v'_1) \circ \phi_1[v'_1, R_1]\end{aligned}$$

from S_1 to R_1 , and thus $G_{\phi\psi}$ is reroutable.

Case 3: This case is shown in Figure 8(c), where each edge is also labeled. Since $G_{\phi\psi}^\bullet$ has three hubs, $G_{\phi\psi}$ has at most six hubs. And since $\mathcal{H}(G) \geq 13$, ξ_1, ξ_2 have to be added to generate at least 7 new hubs. So at least $\lceil (7 + 2 \times 2)/2 \rceil = 6$ edges of ξ_1 and ξ_2 do not belong to any ϕ or ψ -path, and thereby at least two of them, say, (v_1, v_2) and (v'_1, v'_2) are two new edges. We next discuss their possible locations.

Suppose $v_1 \in \phi_1^a$. Then, one verifies that if $v_2 \in \phi_1^a, \phi_1^b, \phi_1^c, \phi_2^b, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, A$ or C , then ϕ is reroutable; if $v_2 \in \phi_2^a, \psi_1^c$ or B , then ψ is reroutable.

Suppose $v_1 \in \phi_1^b$. Then, one verifies that if $v_2 \in \phi_1^a, \phi_1^b, \phi_1^c, \phi_2^a, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, A$ or C , then ϕ is reroutable; if $v_2 \in \phi_2^b, \psi_1^c$ or B , then ψ is reroutable.

Suppose $v_1 \in \phi_2^a$. Then, one verifies that if $v_2 \in \phi_1^b, \phi_1^c, \phi_2^a, \phi_2^b, \psi_1^b, \psi_1^c, \psi_2^a, \psi_2^b, B$ or C , then ϕ is reroutable; if $v_2 \in \phi_1^a, \psi_1^c$ or A , then ψ is reroutable.

Suppose $v_1 \in A$. Then, one verifies that if $v_2 \in \phi_1^a, \phi_1^b, \phi_1^c, \psi_1^a, \psi_1^b, \psi_2^a, \psi_2^b, A$ or C , then ϕ is reroutable; if $v_2 \in \phi_2^a, \phi_2^b, \psi_1^c$ or B , then ψ is reroutable.

By symmetry, if $v_1 \in \psi_1^a, \psi_1^b$ or ψ_2^a , then $G_{\phi\psi}$ is reroutable.

Let G_B be the subgraph of G induced on the edges of $\psi_1[B, R_2]$ and $\phi_2[B, R_1]$ and G_C be the subgraph induced on the edges of $\phi_1[C, R_1]$ and $\psi_2[C, R_2]$. If $v_1, v_2 \in G_B$, either ϕ or ψ is reroutable, and hence $v_1 \notin G_B$ or $v_2 \notin G_B$. Similarly, we deduce that $v_1 \notin G_C$ or $v_2 \notin G_C$. Then we just need to check the last possibility: $v_1, v'_1 \in G_B$ and $v_2, v'_2 \in G_C$, for which we can also easily rule out the subcases when (1) $v_1 \in \phi_1$ and $v'_1 \in \psi_2$ (2) $v_2 \in \phi_2$ and $v'_2 \in \psi_1$ (3) $v_1, v'_1 \in B$ (4)

$v_2, v'_2 \in C$ (5) $v_1, v'_1 \in \phi_2$ and $v_2, v'_2 \in \phi_1$ (6) $v_1, v'_1 \in \psi_1$ and $v_2, v'_2 \in \psi_2$. So, in the following, we examine the remaining two subcases:

(1) $v_1, v'_1 \in \phi_2$ and $v_2, v'_2 \in \psi_2$. By symmetry, we assume that $v_1 < v'_1$ on ϕ_2 in $\vec{G}_{\phi\psi}$. If $v_2 < v'_2$, then ξ is reroutable; If $v_2 > v'_2$, then ϕ or ξ is reroutable.

(2) $v_1, v'_1 \in \psi_1$ and $v_2, v'_2 \in \phi_1$. By symmetry, we assume $v_1 < v'_1$ on ψ_1 in $\vec{G}_{\phi\psi}$. If $v_2 < v'_2$, then ξ is reroutable; If $v_2 > v'_2$, then ψ is reroutable. Now, we are ready to conclude that any minimal and non-reroutable graph in $\mathcal{N}(2, 2, 2)$ has at most 12 hubs.

We next consider the case when G is a minimal but reroutable $(2, 2, 2)$ -graph. Suppose ξ is reroutable in the subgraph $G_{\psi\xi}$. By Theorem 2.1, there exists a ξ -consistent cycle

$$p_1 \circ e_1 \circ p_2 \circ e_2 \circ \cdots \circ p_d \circ e_d,$$

where each p_i is a subpath of some ξ -path and each e_i is a private ψ -edge (in $G_{\psi\xi}$). We further assume that d (the number of private ψ -edges) is the smallest among all possible ξ -consistent cycles in $G_{\psi\xi}$. Then, each e_i belongs to a different ψ -path, which implies $d \leq 2$. Since G is minimal, each ξ -private edge (in $G_{\psi\xi}$) in p_1, p_2, \dots, p_d must belong to ϕ -paths as well. So, each edge of the ξ -consistent cycle in G belongs to either some ϕ -path or ψ -path. Note that for any edge e of this cycle, $G_{\phi\psi} \setminus \{e\}$ does not belong to $\mathcal{N}(C_1, C_2)$, since otherwise, together with the observation that ξ can be rerouted using edges in $G \setminus \{e\}$, we deduce that $G \setminus \{e\} \in \mathcal{N}(C_1, C_2, C_3)$, which contradicts the minimality of G . Then we can find a minimal subgraph $\widehat{G}_{\phi\psi}$ of $G_{\phi\psi}$, which is of either Case 1 (when $d = 2$) or Case 2 (when $d = 1$). Let G_ξ be the subgraph induced on the edges of ξ -paths. Then $G = \widehat{G}_{\phi\psi} \cup G_\xi$, by the minimality of G . Notice that all contradictions in Case 1 and Case 2 arise from the rerouting of ϕ or ψ -paths. So we can follow similar arguments and conclude $\mathcal{H}(G) \leq 12$. The proof is then complete. \square

Acknowledgements: We are indebted to Weiping Shang for insightful discussions during the early stage of this work.

References

- [1] B. Bollobás, *Modern Graph Theory*, Springer, 1998.
- [2] G. Han, *Menger's Paths with Minimum Mergings*, Proc. 2009 IEEE Inform. Theory Workshop on Netw. and Inform. Theory, 2009, pp. 271-275.
- [3] M. Langberg, A. Sprintson and J. Bruck, *The Encoding Complexity of Network Coding*, IEEE Trans. Inform. Theory, vol. 52, no. 6, Jun. 2006, pp. 2386-2397.
- [4] K. Menger, *Zur allgemeinen Kurventheorie*, Fundam. Math., vol. 10, 1927, pp. 96-115.
- [5] E. L. Xu, W. Shang and G. Han, *A Graph Theoretical Approach to Network Encoding Complexity*, Proc. 2012 Int. Symp. on Inform. Theory and its Appl., 2012, pp. 396-400.