博士論文

Combinatorial Optimization on Group-Labeled Graphs
(群ラベル付きグラフにおける組合せ最適化)

Yutaro Yamaguchi (山口 勇太郎)

Combinatorial Optimization on Group-Labeled Graphs

Yutaro Yamaguchi

Preface

The maximum matching and disjoint paths problems have been studied as central topics in combinatorial optimization since its beginning at the middle of the 20th century. Through the research in these topics, various useful concepts and techniques have been developed such as good characterizations by combinatorial duality and efficient algorithms using augmenting paths, which are utilized even today.

One of the highlights on these two topics is Mader's theorem (1978) for openly disjoint A-paths, where an A-path is a simple path between two vertices in a prescribed vertex subset A. He gave a good characterization by a min-max duality for the maximum number of openly disjoint A-paths in an undirected graph, which commonly generalizes the Tutte–Berge formula for maximum matchings in non-bipartite graphs and Menger's theorem for the maximum number of disjoint s–t paths.

To the problem of finding a maximum number of openly disjoint A-paths, Lovász (1980) first provided a solution by reducing it to the matroid matching problem. The matroid matching problem unifies two tractable generalizations of the maximum matching problem in bipartite graphs (the maximum matching problem in non-bipartite graphs and the matroid intersection problem), but itself is intractable in general. It was however impressive that Mader's problem, which also unifies two generalizations of bipartite matching (non-bipartite matching and disjoint s–t paths), can be solved via matroid matching.

As a further extension of Mader's problem, Chudnovsky, Geelen, Gerards, Goddyn, Lohman, and Seymour (2006) introduced the problem of finding disjoint "non-zero" A-paths in group-labeled graphs, which also includes some interesting problems in topological graph theory. They showed a min-max duality extending Mader's theorem, and later Chudnovsky, Cunningham, and Geelen (2008) proposed a polynomial-time combinatorial algorithm for this problem. Pap (2006–2008) introduced a slightly more general model, and also suggested a border between disjoint A-paths problems that enjoy nice structure (e.g., good characterizations and efficient algorithms) and those who do not.

One of the main topics of this thesis is to analyze these extended problems via matroid matching, and to clarify the border between tractable and intractable problems. We show that the most general setting among those suspected as tractable reduces to the matroid matching problem, and clarify when the reduced problem has nice structure such as tractability, good characterizations, and linear representations of reduced matroids.

Other than openly disjoint paths and topological conditions, various settings on graphs can be simply formulated using group-labeled graphs. For instance, a variety of NP-hard problems such as the Hamiltonian path problem and the k-disjoint paths problem can be formulated as finding a path of a designated label in a group-labeled graph. This fact implies that combinatorial optimization on group-labeled graphs is pretty challenging even if we focus on just one path, e.g., to determine whether a given group-labeled graph contains an s–t path whose label is in a prescribed subset of the underlying group.

As the first nontrivial step in this direction, since the situation forbidding only one label is quite easy, we investigate the problem of finding an s–t path with two labels forbidden. This problem in fact includes the 2-disjoint paths problem in undirected graphs, which was well-studied in 1980s. Inspired by and with the aid of Seymour's characterization (1980) for 2-disjoint paths, we give the first nontrivial characterization of group-labeled graphs in terms of the possible labels of s–t paths. Moreover, based on our characterization, we provide an efficient algorithm for finding an s–t path with two labels forbidden.

Acknowledgments

Let me express my thanks and gratitude to those who supported me during my research.

First of all, I would like to show my greatest appreciation to Professor Satoru Iwata, my supervisor. During my master's course and Ph.D. course, he allowed me to pursue the research as I wanted, gave a variety of useful comments and advices when I needed, and was always willing to discuss together. I am really happy to have studied under him for this five years.

I am deeply grateful to the collaborators. Through a joint work with Professor Shin-ichi Tanigawa, I learned the basics of research such as how to do research with others and how to write and submit a paper. Professor Yasushi Kawase told me interesting problems, one of which was the starting point of our joint work. He also taught me a variety of tips on writing and drawing by computers, which helped me a lot while writing this thesis. The key idea in Chapter 7 was given by Professor Yusuke Kobayashi. I greatly appreciate countless illuminating discussions with him, which helped me to write up a long paper.

I would like to express my gratitude to Professor Hiroshi Imai, Professor Kunihiko Sadakane, Professor Ken-ichi Kawarabayashi, and Professor Hiroshi Hirai, who read this thesis and provided a number of instructive and constructive comments.

I am thankful to all current and former members of Mathematical Informatics Laboratories #2 and #7 in University of Tokyo. In particular, I thank Tasuku Soma, Yu Yokoi, and Naoki Ito, with whom I shared a lot of time in the university. I have studied with Tasuku as a same-grade colleague since our master's course at Kyoto University, and learned from him a variety of things not only about our research themes. Yu and Naoki, the roommates in our laboratory, were always cooperative to keep our room comfortable, which supported my happy life in the laboratory. I am also indebted to Professor Kazuo Murota, Professor Kunihiko Sadakane, Professor Akiko Takeda, Professor Hiroshi Hirai, Professor Yusuke Kobayashi, and Professor Yuji Nakatsukasa for their helpful comments in numerous seminars.

I would like to offer my special thanks to Dr. Ryohei Fujimaki and all members of NEC Laboratories who supported my research internship in Japan and in Silicon Valley, California, U.S. That exciting experience made me grow drastically. I also thank Professor Jim Geelen and Professor Gyula Pap, who gave excellent talks on various research topics related to mine and discussed those topics with me.

I also appreciate the financial support by JSPS Fellowship for Young Scientists, by JST ERATO Kawarabayashi Large Graph Project, and by JST CREST.

Finally, I would like to thank my parents and friends for their warm support.

Contents

Chapter 1

Introduction

Combinatorial optimization on graphs is, for a given graph (with some additional data) and requirement, to make an optimal decision, where the optimality is defined by some objective function on the set of feasible decisions. Such a situation appears in the real world in various ways, e.g., when we want to know an efficient routing for a trip, a costless design of network structure (transportation networks, VLSI circuits, and so on), or a vulnerable point of networks by some security reason (e.g., [95]).

From the theoretical aspect, a number of various combinatorial optimization problems on graphs have been studied, which are roughly classified into two types: tractable problems and intractable problems. In general, an optimization problem is said to be *tractable* if it can be solved efficiently, i.e., there exists an algorithm whose running time is polynomially bounded with respect to the input size. One of the main aims in this field is to clarify what kinds of problems are tractable, and moreover to understand the reason behind the tractability, e.g., what structure of tractable problems makes them tractable and what approaches are effective in each nice structure.

In this thesis, we focus on combinatorial optimization on group-labeled graphs. The concept of *group-labeled graphs* is a marriage of simple but deep combinatorial and algebraic structures, graphs and groups, which is expected to enjoy nice properties. In short, group-labeled graphs are directed graphs with each edge labeled by an element of a fixed group. They are also known as *gain graphs* or *voltage graphs* (examples of *biased graphs* introduced by Zaslavsky [96]), and appear in a variety of fields such as network flow theory, scheduling theory, rigidity theory, and so on (see, e.g., [98]).

Depending on the choice of groups Γ for labeling, group-labeled graphs can formulate various situations in simple manners. For instance, $\mathbb{Z}_2$-labeled graphs are equivalent to *signed graphs*, in which each edge has a sign, i.e., $+$ or $-$. Signed graphs were first introduced by Harary [30] in a context of social psychology. They have also been handled in and applied to Ising model and correlation clustering.

Another example is a *periodic graph*, which is introduced by Collatz [10] in a context of spectral graph theory. Roughly speaking, infinite graphs with periodic structure representable as $\mathbb{Z}^d$-labeled graphs are called periodic graphs, where d is a positive integer. Literally, periodic graphs appear in modeling various periodic structures in the real world such as crystals and VLSI circuits. They are utilized also in theoretical situations, e.g., as *dependence graphs* of uniform recurrence equations [44].

As seen above, by introducing groups (algebraic structure) into graphs (combinatorial structure), a variety of situations can be simply formulated. From the viewpoint of combinatorial optimization, most of such formulations naturally generate combinatorial optimization problems on group-labeled graphs, which sometimes unify several apparently unrelated classical problems or bring about a new idea for modeling other situations. In exchange for expressive power, problems on group-labeled graphs tend to be pretty hard and often intractable. We aim to reveal what types of problems on group-labeled graphs are tractable and why so are they, based on classical approaches in combinatorial optimization and graph theory. Furthermore, through such research in high-level models, we expect to acquire novel knowledge and interpretation in these classical theories.

In what follows, we present background and our contributions. In Section 1.1, we describe historical background of combinatorial optimization, featuring matchings and disjoint paths, which are closely related to our work. Next, we show significant results on these two topics with precise statements separately in Sections 1.2 and 1.3. In Section 1.4, we review problems on paths in group-labeled graphs, which are dealt with as the main subject in this thesis. Finally, in Section 1.5, we sum up our contributions to this field and the organization of this thesis.

1.1 Historical Background

The theory of combinatorial optimization started in the middle of the 20th century with unifying independent works on combinatorial structure that had existed before. On graphs in particular, as mentioned at the beginning, a number of various problems have been studied, e.g., the shortest path problem, the minimum-cost spanning tree problem, and the maximum flow problem as well as its dual, the minimum cut problem. Here we focus on the following two types of problems, which are closely related to the problems on group-labeled graphs dealt with in this thesis: the maximum matching problem and the disjoint paths problem. Both of these two problems have been studied as central topics in combinatorial optimization as well as in graph theory.

The first explicit appearance of matchings in graphs is in problems on matrices in the early 20th century. Frobenius [19] characterized mixed matrices containing no nonzero constant entry whose determinants are factorizable, and later Kőnig [51] gave a simpler proof by regarding it as a problem on matchings in bipartite graphs. Kőnig [52, 53] also proved the so-called Birkhoff–von Neumann Theorem (every nonnegative doubly stochastic matrix can be represented as a convex combination of permutation matrices, which was rediscovered later independently by Birkhoff [2] and von Neumann [67]) by showing that any k-regular bipartite graph can be decomposed into k disjoint perfect matchings.

From the viewpoint of combinatorial optimization on graphs, the first milestone of the study of matchings is Kőnig's min-max theorem [54], which claims that the maximum cardinality of a matching in any bipartite graph is equal to the minimum cardinality of a vertex cover in it. This theorem is a good characterization by a combinatorial duality, i.e., one can easily certify that a given bipartite graph has a matching of size k or not by finding such a matching or a vertex cover of size less than k, respectively. Independently, Frobenius [20] proved a special case of Kőnig's theorem when a bipartite graph has a perfect matching, and Hall [29] gave another characterization for the existence of single-side perfect matchings, which is equivalent to Kőnig's theorem. The proof of Kőnig's theorem is based on the so-called augmenting-path technique, which leads to a simple algorithm for finding a maximum matching in a bipartite graph.

After decades, Tutte [85] gave a necessary and sufficient condition for general graphs (i.e., not only for bipartite graphs) to admit a perfect matching, which generalizes Frobenius' theorem for bipartite graphs. Later, Berge [1] observed that Tutte's characterization leads to a min-max duality, which extends Kőnig's theorem. The min-max theorem is well-known today as the Tutte–Berge formula (Theorem 1.1), whose proof is not constructive as Kőnig's one.

The first efficient algorithm for non-bipartite matching was devised by Edmonds [13]. His algorithm is also based on the augmenting-path idea, and contains a novel technique, shrinking a blossom, which is the reason why it is called the blossom algorithm. Originated by Edmonds' blossom algorithm, a variety of efficient algorithms have been developed, e.g., combinatorial algorithms due to Micali and Vazirani [65] and Goldberg and Karzanov [28] and algebraic algorithms due to Mucha and Sankowski [66] and Harvey [31] (cf. Theorems 1.2–1.4). It is intriguing that a combinatorial problem appearing from matrix theory is solved efficiently by algebraic techniques based on matrix theory.

The research in the disjoint paths problem has its origin in a min-max duality for the maximum number of disjoint s–t paths due to Menger [64]. The concept of disjoint s–t paths generalizes matchings in bipartite graphs, and Menger's theorem in fact implies Kőnig's theorem as a special case. On the other hand, it was pointed out by Tutte [86] that disjoint s–t paths can be formulated as matroid intersection, which is another generalization of bipartite matching. A simple but the first efficient algorithm for finding a maximum number of disjoint s–t paths was given as a special case of an augmenting-path algorithm for the maximum flow problem due to Ford and Fulkerson [16].

As a noteworthy result making a connection between matchings and disjoint paths, Mader [63] showed a min-max duality for openly disjoint A-paths, which commonly generalizes the Tutte–Berge formula and Menger's theorem, where an A-path is a path between distinct vertices in a prescribed vertex subset A. The problem of finding a maximum number of openly disjoint A-paths was first solved by Lovász [59, 61] via a reduction to the matroid matching problem, which unifies two tractable generalizations of bipartite matching (non-bipartite matching and matroid intersection) but is intractable in general. It is intriguing again that a deep combinatorial optimization problem on graphs unifying two generalizations of bipartite matching was first solved via another deep unification.

Other than numerous studies for such maximization-type disjoint paths problems, the following setting is also well-investigated: for given multiple pairs (s_i, t_i) of two vertices, to find disjoint s_i–t_i paths if exist. This kind of disjoint paths problem tends to be intractable, but some cases are known to be tractable. The deepest result is in graph minor theory due to Robertson and Seymour as follows [76]: when a positive integer k is fixed (i.e., considered as a constant), for given k pairs (s_i, t_i) of two vertices in an undirected graph, one can find disjoint s_i–t_i paths in polynomial time if exist. When $k = 2$ in the same setting, there are earlier independent results for its tractability due to Seymour [80], Shiloach [81], and Thomassen [84], which may suggest that there exists some essential gap between the case of $k = 2$ and $k \geq 3$.

1.2 Maximum Matching Problem

For an undirected graph $G = (V, E)$, an edge set $M \subseteq E$ is called a *matching* in G if $|V(M)| = 2|M|$, i.e., all $2|M|$ end vertices of $|M|$ edges in M are distinct. The *maximum matching problem* is to find a matching of the maximum cardinality in a given undirected graph.

Maximum Matching Problem

Input: An undirected graph $G = (V, E)$.

Goal: Find a matching $M \subseteq E$ in G such that $|M|$ is maximum.

As a milestone of the study of this problem, Tutte [85] first characterized undirected graphs that admit a perfect matching (a matching M covering all vertices, i.e., $V(M) = V$). Later, Berge [1] observed that Tutte's characterization leads to a min-max duality, which is known today as the *Tutte–Berge formula* as follows.

Theorem 1.1 (Tutte [85], Berge [1]). *Let $G = (V, E)$ be an undirected graph. Then, the maximum cardinality of a matching in G is equal to the minimum value of*

$$\frac{1}{2}\left(|V| + |U| - \mathrm{odd}(G - U)\right), \tag{1.1}$$

taken over all vertex subsets $U \subseteq V$, where $\mathrm{odd}(G - U)$ *denotes the number of connected components of $G - U$ that consist of odd number of vertices.*

Originated by the *blossom algorithm* of Edmonds [13], a variety of efficient algorithms for this problem have been developed. We just remark the fastest ones currently known. Let ω denote the matrix multiplication exponent, which is at most 2.373 due to Le Gall [57].

Theorem 1.2 (Mucha–Sankowski [66], Harvey [31]). *One can solve the maximum matching problem in* $\mathrm{O}(|V|^{\omega})$ *time by a randomized algorithm.*

Theorem 1.3 (Goldberg–Karzanov [28]). *One can solve the maximum matching problem in* $\mathrm{O}(\sqrt{|V|} \cdot |E| \cdot \log(|V|^2/|E|)/\log |V|)$ *time by a deterministic algorithm.*

Theorem 1.4 (Micali–Vazirani [65], and see also [87, 88]). *One can solve the maximum matching problem in* $\mathrm{O}(\sqrt{|V|} \cdot |E|)$ *time by a deterministic algorithm.*

Related to maximum matchings in undirected graphs, Gallai [25,27] and Edmonds [13] independently showed a structure of undirected graphs, which decomposes them into three parts as follows. This result is currently known as the *Edmonds–Gallai decomposition* or the *Edmonds–Gallai structure theorem.*

Theorem 1.5 (Gallai [25,27], Edmonds [13]). *For an undirected graph* $G = (V, E)$*, define*

$$\begin{aligned} D &:= \{\, v \in V \mid \exists M \subseteq E\colon \text{a maximum matching in } G \text{ with } v \notin V(M)\,\},\\ A &:= N_G(D) \;=\; \{\, v \in V \setminus D \mid \exists u \in D,\ \exists e = uv \in E\,\},\\ C &:= V \setminus (A \cup D). \end{aligned}$$

Then, the vertex set A *attains the minimum of* (1.1) *as* U *in Theorem* 1.1*, and* $G[D]$ *and* $G[C]$ *consist of odd and even connected components of* $G - A$*, respectively.*

1.3 Disjoint Paths Problem

1.3.1 s–t paths

We refer to the following problem as *Menger's disjoint paths problem*: for a given graph $G = (V, E)$ and distinct vertices $s, t \in V$, to find a maximum number of disjoint s–t paths in G. Here, the term "disjoint" has two meanings: "edge-disjoint" and "openly disjoint." Two paths are said to be *edge-disjoint* if they do not share any edge, and *openly disjoint* if they do not share any inner vertex (end vertices are sharable). More than two paths are said to be disjoint if they are pairwise disjoint, i.e., every two among them are disjoint. Here, we focus on the case of undirected graphs as follows.

Edge-Disjoint s–t Paths Problem

Input: An undirected graph $G = (V, E)$ and distinct vertices $s, t \in V$.

Goal: Find a family $\mathcal{P}$ of edge-disjoint s–t paths in G such that $|\mathcal{P}|$ is maximum.

Openly Disjoint s–t Paths Problem

Input: An undirected graph $G = (V, E)$ and distinct vertices $s, t \in V$.

Goal: Find a family $\mathcal{P}$ of openly disjoint s–t paths in G such that $|\mathcal{P}|$ is maximum.

By using the concept of line graphs, the former problem can be reduced to the latter problem. For an undirected graph $G = (V, E)$, the *line graph* $\hat{G} = (\hat{V}, \hat{E})$ of G is defined as follows: $\hat{V} := E$ and $\hat{E} := \{\, e_1e_2 \mid e_1, e_2 \in E\colon \text{adjacent in } G \,\}$, i.e., each edge in G is a vertex in $\hat{G}$ and vice versa, and two vertices are adjacent (connected by an edge) in $\hat{G}$ if and only if the corresponding edges are adjacent (share one of their end vertices) in G.

For an instance of the edge-disjoint s–t paths problem, let $G' = (V \cup \{s', t'\},\ E \cup \{e_s, e_t\})$ be the undirected graph obtained from G by adding two new vertices s' and t' with edges $e_s = s's$ and $e_t = tt'$. Then, openly disjoint e_s–e_t paths in the line graph of G' is essentially equivalent to edge-disjoint s–t paths in G (e.g., we can obtain the latter from the former by transforming its vertex set to an edge set in G and scraping redundant edges if necessary), and this equivalence completes a reduction of the edge-disjoint setting to the openly disjoint setting.

In [64], Menger provided the first good characterization for this kind of disjoint paths problems as follows: in short, the maximum number of disjoint s–t paths is equal to the minimum size of a *hitting set*, whose removal disconnects the terminals s and t. The original theorem is for the openly disjoint setting, which also leads to an analogous characterization for the edge-disjoint setting through the above reduction.

Theorem 1.6 (Menger [64])**.** *Let $G = (V, E)$ be an undirected graph, $s, t \in V$ distinct and nonadjacent vertices, and $k \in \mathbb{Z}_{>0}$ a positive integer. Then, there exist k openly disjoint s–t paths in G if and only if, for every vertex set $X \subseteq V \setminus \{s, t\}$ with $|X| \leq k - 1$, the terminals s and t are contained in the same connected component of $G - X$.*

Theorem 1.7 (Menger [64])**.** *Let $G = (V, E)$ be an undirected graph, $s, t \in V$ distinct vertices, and $k \in \mathbb{Z}_{>0}$ a positive integer. Then, there exist k edge-disjoint s–t paths in G if and only if, for every edge set $F \subseteq E$ with $|F| \leq k - 1$, the terminals s and t are contained in the same connected component of $G - F$.*

In [18], it was remarked that Gallai [23] had given a constructive proof for Menger's theorem and pointed out its extendability to the case of directed graphs. Ford and Fulkerson [16] first provided an algorithm for the maximum flow problem, which generalizes the edge-disjoint s–t paths problem in directed graphs. Based on this algorithm, one can obtain a simple polynomial-time algorithm for Menger's problem in every setting (edge-disjoint or openly disjoint, and undirected or directed).

1.3.2 A-paths

Let $G = (V, E)$ be an undirected graph. For a vertex set $A \subseteq V$, an *A-path* is a path between distinct vertices in A that does not intersect A in between. In this situation, each vertex in A is called a *terminal*. Note that, when $A = \{s, t\}$ for distinct vertices $s, t \in V$, an A-path is equivalent to an s–t path. Besides, when $A = V$, an A-path is equivalent to an edge in E.

The *disjoint A-paths problem* is, for a given undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$, to find a maximum number of disjoint A-paths in G. Also here, the term "disjoint" has several meanings: "vertex-disjoint," "edge-disjoint," and "openly disjoint." Two paths are said to be *vertex-disjoint* if they do not share any vertex (not only inner vertex but also end vertex), and recall that more than two paths are said to be disjoint if they are pairwise disjoint. In this thesis, we also use the term "packing" to indicate the vertex-disjoint setting.

Vertex-Disjoint (Packing) A-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint A-paths in G such that $|\mathcal{P}|$ is maximum.

Edge-Disjoint A-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of edge-disjoint A-paths in G such that $|\mathcal{P}|$ is maximum.

Openly Disjoint A-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of openly disjoint A-paths in G such that $|\mathcal{P}|$ is maximum.

As mentioned in Section 1.1, the openly disjoint A-paths problem commonly generalizes the maximum matching problem and Menger's disjoint paths problem. Specifically, the maximum matching problem is a special case of the vertex-disjoint setting with $A = V$, Menger's problem in undirected graphs is a special case with $A = \{s, t\}$ (in both of the edge-disjoint and openly disjoint settings), and the openly disjoint setting is the most general among these as follows.

To formulate the vertex-disjoint setting, add a copy v' of each terminal $v \in A$ with an edge $e_v = vv'$, and let $A' := \{\, v' \mid v \in A \,\}$ be a new terminal set. For the edge-disjoint setting, after the above procedure, consider the line graph and let $A' := \{\, e_v \mid v \in A \,\}$ be a new terminal set (cf. the reduction for Menger's problems in Section 1.3.1).

Since the packing A-paths problem is one of the main topics in this thesis, we review this problem in detail in Chapter 3 (see also Fig. 1.1). Here we just give an overview.

Gallai [24] solved the vertex-disjoint A-paths problem by reducing it to the maximum matching problem (hence, these two problems are essentially equivalent), and gave a min-max formula that extends the Tutte–Berge formula (Theorem 1.1). He also mentioned the openly disjoint setting, for which Mader [63] provided a min-max formula later. As described before, Mader's theorem for openly disjoint A-paths commonly generalizes Gallai's theorem (as well as the Tutte–Berge formula) and Menger's theorem (see Section 3.2 for the detail).

From the viewpoint of tractability, the first efficient algorithm for the openly disjoint A-paths problem was proposed by Lovász [59,61] via a reduction to the matroid matching problem. Later, Schrijver [78, Section 73.1a] pointed out that his reduction admits a linear representation, which leads to much faster algorithms via the linear matroid parity problem. These reductions provide us an interpretation of the tractability from a wide framework of combinatorial optimization that is not conventional on graphs. Since they are closely related to our results, we describe them more specifically in Sections 4.1 and 5.1, respectively.

1.3.3 k-disjoint paths

Let k be a positive integer. We refer to the following problem as the *k-disjoint paths problem*: for a given graph $G = (V, E)$ and $2k$ distinct vertices s_i, t_i ($i = 1, 2, \ldots, k$), to find vertex-disjoint s_i–t_i paths in G if exist. Recall that the term "vertex-disjoint" means that any vertex cannot be shared by any two distinct paths.

This problem is known to be NP-hard when k is a part of the input [43]. Besides, in the directed case (i.e., when G is a directed graph and any directed edge cannot be traversed in the backward direction), this problem is NP-hard even if $k = 2$ [17]. In contrast, when G is an undirected graph and k is fixed, this problem can be solved in polynomial time with the aid of graph minor theory [76]. In this thesis, we focus only on the last setting as follows.

k-disjoint Paths Problem

Input: An undirected graph $G = (V, E)$ and $2k$ distinct vertices $s_i, t_i \in V$ ($i = 1, 2, \ldots, k$).

Goal: Find vertex-disjoint s_i–t_i paths P_i ($i = 1, 2, \ldots, k$) in G, or certify the nonexistence.

For the case of $k = 2$ in particular, Seymour [80], Shiloach [81], and Thomassen [84] independently devised elementary polynomial-time algorithms that do not rely on graph minor theory. The key of these algorithms is the following characterization of undirected graphs in terms of the existence of 2-disjoint paths, which can be tested in polynomial time.

Theorem 1.8 (Seymour [80]). *Let $G = (V, E)$ be an undirected graph and $s_1, t_1, s_2, t_2 \in V$ distinct vertices. Then, there exist two vertex-disjoint paths P_i connecting s_i and t_i ($i = 1, 2$) if and only if there is no family of disjoint vertex sets $X_1, X_2, \ldots, X_k \subseteq V \setminus \{s_1, t_1, s_2, t_2\}$ such that*

1. *$N_G(X_i) \cap X_j = \emptyset$ for distinct $i, j \in \{1, 2, \ldots, k\}$,*
2. *$|N_G(X_i)| \leq 3$ for $i = 1, 2, \ldots, k$, and*
3. *if G' is the graph obtained from G by deleting X_i and adding a new edge joining each pair of distinct vertices in $N_G(X_i)$ for each $i \in \{1, 2, \ldots, k\}$, then G' can be embedded on a plane so that s_1, s_2, t_1, t_2 are on the outer boundary in this order.*

1.4 Paths in Group-Labeled Graphs

Recall that a group-labeled graph is a directed graph with each edge labeled by an element of a fixed group. In group-labeled graphs, the label of an undirected path (allowed to traverse a directed edge in backward direction) is naturally defined by the group operation along it, where the label of each backward edge is inversed (see Section 2.2.1 for more details). Using the labels of paths, we can formulate various problems as those on group-labeled graphs. In what follows, we briefly describe significant related works on paths in group-labeled graphs, where a path is said to be *zero* if its label is the identity element of the underlying group, and *non-zero* otherwise.

1.4.1 Non-zero paths

Chudnovsky, Geelen, Gerards, Goddyn, Lohman, and Seymour [7] introduced the *packing non-zero A-paths problem*, which generalizes Mader's openly disjoint A-paths problem and has several interesting applications in topological graph theory, e.g., to find disjoint non-contractible (or non-separating) cycles in a graph embedded on a surface (see [7, Section 2] for more details). The objective is to maximize the number of vertex-disjoint non-zero A-paths in a given group-labeled graph with a terminal set A. They showed a min-max duality that extends Mader's theorem, and Chudnovsky, Cunningham, and Geelen [6] proposed a polynomial-time combinatorial algorithm for this problem with an Edmonds–Gallai-type structure theorem (see Section 3.3 for the details).

Later, Pap [72, 73] introduced a more general problem, called *packing non-returning A-paths*, and provided a simple proof for a generalized min-max formula and another combinatorial approach to his problem (see Section 3.4 for the details). Moreover, he suspected in his thesis [71] that the tractability of these packing A-paths problems is derived from a combinatorial property enjoyed by non-zero and non-returning models, which he called the *triple exchange property* (see Section 3.5 for the details).

There are several works on packing non-zero cycles in group-labeled graphs. We just refer to related papers [48, 90].

1.4.2 Zero paths

While it is rather easy to handle non-zero paths, zero paths are much difficult. In what follows, we focus only on paths between fixed distinct terminals s and t. Note that finding a zero s–t path is equivalent to finding an s–t path of a designated label α, since the latter reduces to the former by adding a new vertex s' as a new source (i.e., the objective is to find a zero s'–t path) with a new directed edge from s' to s with label α^{-1}.

As a simple example, the Hamiltonian path problem can be formulated as finding a zero path in $\mathbb{Z}$-labeled graphs. The *Hamiltonian path problem* is a popular NP-hard problem [42] (see also [26, Section 3.1.4]), and here we consider the following setting: for a given directed graph $G = (V, E)$ and distinct vertices $s, t \in V$, to find a directed s–t path intersecting all vertices in V exactly once (called a *directed s–t Hamiltonian path*) if exists. For an instance of this problem, assign the label $1 \in \mathbb{Z}$ to all directed edges. Then, a directed s–t Hamiltonian path in the original graph is an s–t path of label $|V| - 1$ in the resulting $\mathbb{Z}$-labeled graph and vice versa, which implies that finding a zero s–t path in $\mathbb{Z}$-labeled graphs is NP-hard.

Another special case is the k-disjoint paths problem[1]. For an undirected graph $G = (V, E)$ and $2k$ distinct vertices $s_i, t_i \in V$ $(i = 1, 2, \ldots, k)$, we construct the following A_{2k-1}-labeled graph, where A_{2k-1} denotes the alternating group of degree $2k - 1$, i.e., the group of even permutations on $\{1, 2, \ldots, 2k - 1\}$. Orient each edge arbitrarily (i.e., replace each edge between two vertices u and v with a directed edge from u to v), assign the label $\mathrm{id} \in A_{2k-1}$ to each directed edge (note that the orientation is not essential since $\mathrm{id}^{-1} = \mathrm{id}$), and add a directed edge from t_i to s_{i+1} with label $(2i-1\ \ 2i+1\ \ 2i) \in A_{2k-1}$ for each $i = 1, 2, \ldots, k - 1$. Then, each k-disjoint paths in G corresponds to an s_1–t_k path of label

$$\sigma := (2k-3\ \ 2k-1\ \ 2k-2) \cdots (3\ 5\ 4)(1\ 3\ 2)$$

in the resulting A_{2k-1}-labeled graph and vice versa. Note that this σ is a unique permutation that maps 1 to $2k - 1$ among those which can be constructed in this A_{2k-1}-labeled graph. This correspondence implies that finding a zero s–t path in A_{2k-1}-labeled graphs is not easier than the k-disjoint paths problem. It should be noted that, when $k = 2$, the alternating group $A_{2k-1} = A_3$ is isomorphic to $\mathbb{Z}_3 = \mathbb{Z}/3\mathbb{Z}$, which is abelian (and when $k \geq 3$, it is non-abelian).

As an extension of the solution to the k-disjoint paths problem [76], Huynh [37] gave a polynomial-time algorithm for finding k-disjoint zero paths in a Γ-labeled graph for any fixed finite abelian group Γ and any fixed positive integer k, which also relies on sophisticated results in graph minor theory. No other result for finding a zero path in group-labeled graphs is known, and the case of non-abelian groups is still open even if they are finite.

1.4.3 Paths in periodic graphs

Periodic graphs are infinite graphs with periodic structure representable as $\mathbb{Z}^d$-labeled graphs, where d is a positive integer. There are several algorithmic studies on paths in periodic graphs. They may have some relation to our work but only a little is currently recognized, because of the difference between paths in periodic graphs and those in group-labeled graphs. Specifically, a path in a periodic graph corresponds to a walk (which is allowed to traverse an edge multiple times) in the corresponding $\mathbb{Z}^d$-labeled graph (called the *static graph*).

[1]This observation was described in [37, p. 11]. However, their reduction is inadequate, which cannot distinguish two pairs of disjoint s_i–t_i path and s_j–t_j path and disjoint s_i–s_j path and t_i–t_j path for any distinct i, j.

Orlin [68] first solved the problem of finding a path between two vertices in 1-dimensional periodic graphs, Iwano [38] gave an algorithm for the 2-dimensional case, and Cohen and Megiddo [9] did it for the general dimensions. This problem is equivalent to finding a zero s–t walk in $\mathbb{Z}^d$-labeled graphs. Finding a zero s–t walk in group-labeled graphs is translated into testing whether the subgroup generated by given elements contains a designated element or not (cf. Proposition 2.11), which can be solved with the aid of the extended Euclidean algorithm when the underlying group is $\mathbb{Z}^d$.

In contrast to the tractability of undirected paths, directed paths in periodic graphs are pretty hard to deal with. Wanke [89] analyzed the complexity of finding a directed path in finitely restricted areas of periodic graphs and showed that it is NP-hard, e.g., even when the static graph consists of only one vertex and $d = 1$. Höfting and Wanke [34] formulated the shortest path problem in periodic graphs (which is also NP-hard) as integer program and provided a pseudopolynomial-time algorithm for the fixed dimension case. As a recent breakthrough, Fu [21] showed the tractability of the shortest path problem in 2-dimensional planar periodic graphs. She pointed out that the algorithm of Höfting and Wanke [34] runs in weakly polynomial time for 2-dimensional planar periodic graphs, using the result of Iwano and Steiglitz [39] on the planarity of periodic graphs. In addition, she introduced the concept of *coherence* of periodic graphs by capturing a nice combinatorial property, and devised a strongly polynomial-time algorithm for the shortest path problem in 2-dimensional coherent planar periodic graphs.

1.5 Our Contributions

Throughout this thesis, we deal with problems on paths in group-labeled graphs. The main results are divided into two parts: characterizations and algorithms for packing A-paths in group-labeled graphs (Fig. 1.1 shows our contributions to this topic), and those for finding an s–t path in group-labeled graphs. In both parts, we aim to reveal the boundary between tractable and intractable problems on group-labeled graphs as well as on graphs.

The first main result is a reduction of the packing non-zero A-paths problem to the matroid matching problem, which extends Lovász' reduction of Mader's openly disjoint A-paths problem. With the aid of a generalized frame matroid introduced by Tanigawa [82], we reduce packing non-zero A-paths to a tractable case of matroid matching, and give alternative proofs for the min-max formula due to Chudnovsky et al. [7] and the polynomial-time solvability (cf. [6]) by applying Lovász' theory on matroid matching [59, 61]. In addition, we show a possible extension of our reduction to a further generalized model of packing A-path (with triple exchange property), which does not necessarily lead to a good characterization or a polynomial-time algorithm.

The second main result is a characterization of groups for which a generalized problem of packing A-paths in group-labeled graphs, called the *subgroup-forbidden model*, which is equivalent to Pap's problem [72, 73] (see Section 3.4 for the details), can be solved much faster via the linear matroid parity problem. We provide a necessary and sufficient condition for the groups in question to admit a reasonable reduction to linear matroid parity, which extends the reduction of Mader's problem due to Schrijver [78, Section 73.1a]. In the case of packing non-zero A-paths in particular, a large class admits our reduction. This fact leads to an $\mathrm{O}(|V|^\omega)$-time algorithm (recall that $\omega < 2.373$ denotes the matrix multiplication exponent [57]) with the aid of a linear matroid parity algorithm due to Cheung, Lau, and Leung [5], which is much faster than the algorithm of Chudnovsky et al. [6] requiring $\mathrm{O}(|V|^5)$ time.

The above two reductions are also applicable to a weighted setting, in which we are given length of each edge and required to minimize the total length of a designated number of disjoint paths. Such a setting may be reasonable in several situations, e.g., cost minimization in VLSI design. By an ingenious transformation of the input graphs, we can handle this setting via matroid matching. Furthermore, with the aid of weighted linear matroid parity algorithms thanks to Iwata [40] and Pap [75], our reduction leads to the first polynomial-time algorithm for the weighted version of Mader's openly disjoint A-paths problem.

The third main result is a characterization for group-labeled graphs with two distinct terminals s and t which have exactly two possible labels of s–t paths. It is easy to characterize group-labeled graphs with two distinct terminals s and t such that all s–t paths have the same label, by using the concept of the balancedness of group-labeled graphs (see Section 2.2.3). In contrast, our case (having exactly two possible labels of s–t paths) is much more difficult. As an evidence, it includes the nonexistence of 2-disjoint paths in undirected graphs (see Chapter 7), and our characterization is deeply inspired by Seymour's theorem (Theorem 1.8). Furthermore, using our characterization, we provide a polynomial-time algorithm for finding an s–t path in group-labeled graphs with two labels forbidden.

The rest of this thesis is organized as follows. In Chapter 2, we define necessary concepts and notations, and show elementary key properties. Chapter 3 is devoted to reviewing the packing A-paths problem.

We present the first main result in Chapter 4: a reduction of packing non-zero A-paths to matroid matching and its applications. In Chapter 5, we clarify a necessary and sufficient condition for a reduction of the subgroup-forbidden model to linear matroid parity as the second main result. We discuss their extension to a weighted version in Chapter 6.

As the third main result, in Chapter 7, we provide a solution to finding an s–t path with two labels forbidden using a characterization of group-labeled graphs with exactly two possible labels of s–t paths.

Finally, Chapter 8 concludes this thesis.

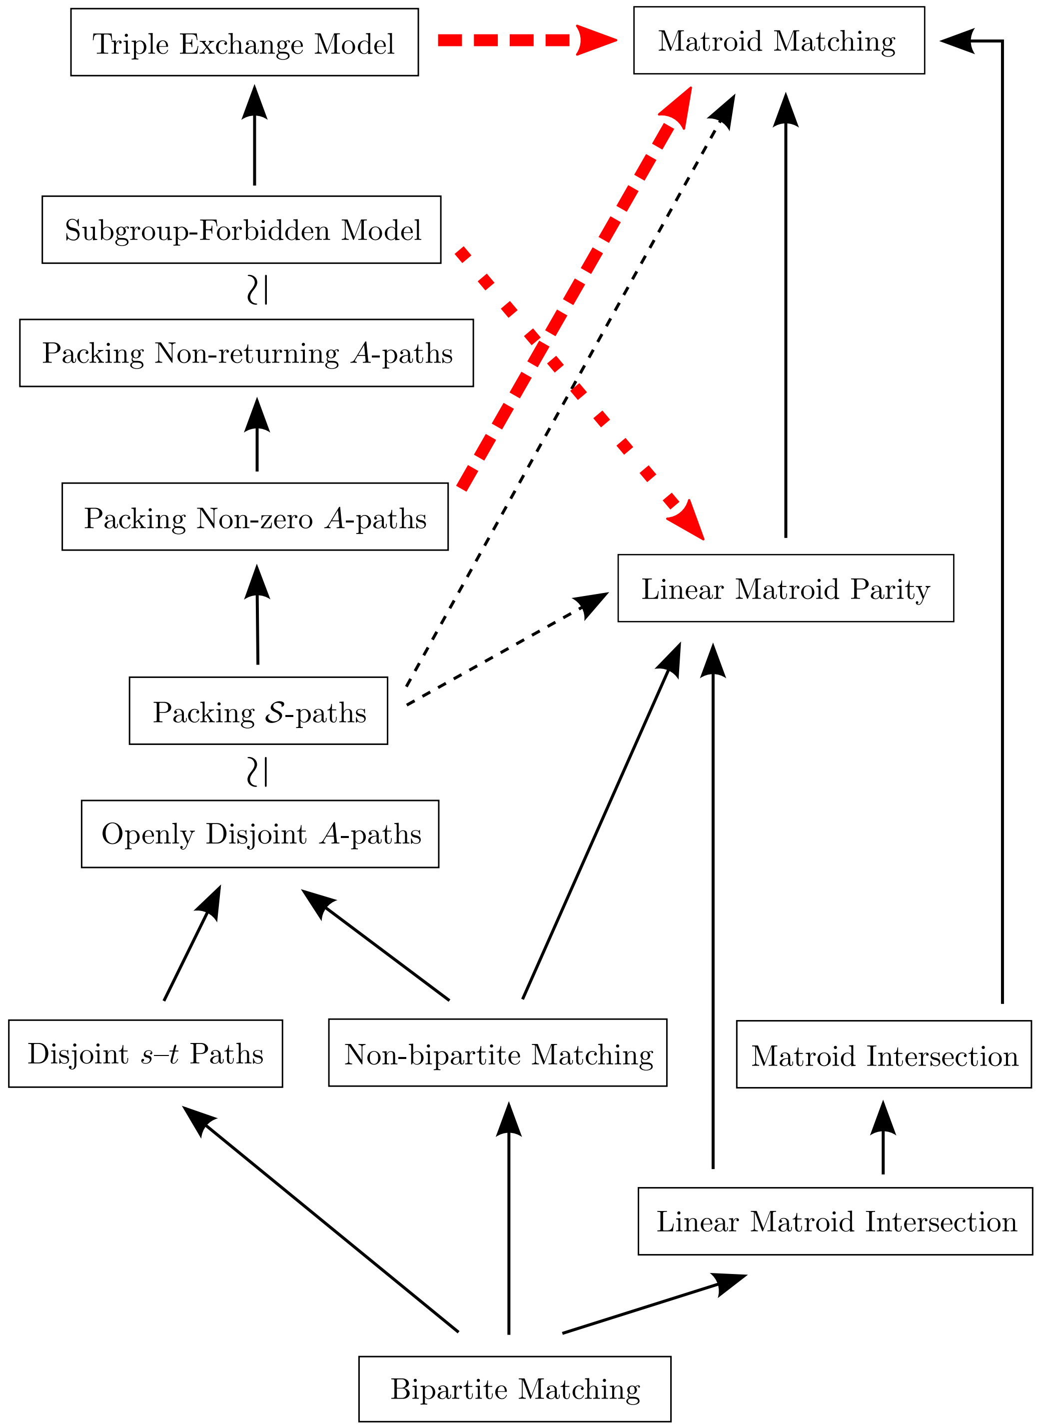

Figure 1.1: Problems related to packing A-paths (see Chapter 3 for the details). Each filled, dashed, or dotted arrow means a generalization, a complete reduction, or a conditional reduction (with a necessary and sufficient condition), respectively. Bold and red ones represent our results.

Chapter 2

Preliminaries

First of all, we mention basic notations throughout this thesis. Let $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$ be the sets of integers, rationals, and reals, respectively. They are sometimes with a constraint subscript, e.g., $\mathbb{R}_{>0}$ denotes the set of positive reals. For $n \in \mathbb{Z}_{\geq 0}$, define $[n] := \{1, 2, \ldots, n\} \subsetneq \mathbb{Z}_{>0}$. We often identify a singleton $\{x\}$ with its element x when it makes no confusion, e.g., for an operator f defined on a power set 2^X, we denote $f(\{x\})$ simply by $f(x)$ for each $x \in X$. For a set X and an element e, define $X + e := X \cup \{e\}$ when $e \notin X$, and $X - e := X \setminus \{e\}$ when $e \in X$.

2.1 Graphs

In this section, we define terms and notations for graphs used throughout this thesis, where the term "graph" indicates both of an undirected graph and a directed graph. For a graph G, we denote by $V(G)$ and $E(G)$ the vertex set and edge set of G, respectively. We often denote a graph G by a pair of its vertex set and edge set, i.e., $G = (V(G), E(G))$. In what follows, let $G = (V, E)$ be a graph.

2.1.1 Basic notations

For distinct vertices $u, v \in V$, we denote by uv an edge connecting u and v. When G is an undirected graph, the representation uv is equivalent to vu. When G is directed, the edge uv is directed from u to v, which is not equivalent to vu. A *self-loop* on a vertex $v \in V$ (an edge whose end vertices are both v) is also denoted by vv. Throughout this thesis, an edge in directed graphs is called an *arc* when its direction is important.

Let $X \subseteq V$ be a vertex set and $F \subseteq E$ an edge set. We denote by $E(X)$ the set of edges *induced by* X and by $V(F)$ the set of end vertices of edges in F, i.e., $E(X) := \{\, xy \in E \mid \{x, y\} \subseteq X \,\}$ and $V(F) := \{\, x \in V \mid \exists y \in V \colon xy \in F \text{ or } yx \in F \,\}$. Let $G[X] := (X, E(X))$ denote the subgraph of G *induced by* X, and $G[F] := (V, F)$ the subgraph of G *restricted to* F. Define $G - X := G[V \setminus X]$ and $G - F := G[E \setminus F]$.

For a vertex set $X \subseteq V$, we denote by $\delta_G(X)$ the set of edges *incident to* X in G and by $N_G(X)$ the set of vertices *adjacent to* X (called *neighbors* of X) in G. Besides, when G is directed, $\delta_G^{\mathrm{in}}(X)$ and $\delta_G^{\mathrm{out}}(X)$ denote the set of arcs that *enter* and *leave* X in G, respectively, and $N_G^{\mathrm{in}}(X)$ and $N_G^{\mathrm{out}}(X)$ the set of *in-neighbors* and *out-neighbors* of X in G, respectively. To sum up,

$$\begin{aligned}
\delta_G(X) &:= \{\, e = xy \in E \mid x \neq y \text{ and } |\{x, y\} \cap X| = 1 \,\},\\
\delta_G^{\mathrm{in}}(X) &:= \{\, e = yx \in E \mid x \in X \text{ and } y \in V \setminus X \,\},\\
\delta_G^{\mathrm{out}}(X) &:= \{\, e = xy \in E \mid x \in X \text{ and } y \in V \setminus X \,\},\\
N_G(X) &:= \{\, y \in V \setminus X \mid \delta_G(X) \cap \delta_G(y) \neq \emptyset \,\},\\
N_G^{\mathrm{in}}(X) &:= \{\, y \in V \setminus X \mid \delta_G^{\mathrm{in}}(X) \cap \delta_G^{\mathrm{out}}(y) \neq \emptyset \,\},\\
N_G^{\mathrm{out}}(X) &:= \{\, y \in V \setminus X \mid \delta_G^{\mathrm{out}}(X) \cap \delta_G^{\mathrm{in}}(y) \neq \emptyset \,\}.
\end{aligned}$$

2.1.2 Walks and paths

For vertices $v_0, v_1, \ldots, v_l \in V$ and edges $e_1, e_2, \ldots, e_l \in E$ with $e_i = v_{i-1}v_i$ or $e_i = v_i v_{i-1}$ for each $i \in [l]$, a sequence $W = (v_0, e_1, v_1, e_2, v_2, \ldots, e_l, v_l)$ is called a *walk* (in particular, a v_0–v_l *walk*) in G. The walk W is said to be *directed* if G is directed and every arc e_i is directed from v_{i-1} to v_i, i.e., $e_i = v_{i-1}v_i$ for every $i \in [l]$. We call v_0 and v_l the *end vertices* of W, and each v_i ($i \in [l-1]$) an *inner vertex* on W. Note that some vertex can be an end vertex and an inner vertex. For a pair (i, j) with $0 \leq i < j \leq l$, let $W[v_i, v_j]$ denote the subwalk $(v_i, e_{i+1}, v_{i+1}, \ldots, e_j, v_j)$ of W when it is uniquely determined. Let $\bar{W}$ denote the reversed walk of W, i.e., $\bar{W} = (v_l, e_l, \ldots, v_1, e_1, v_0)$. For a u–v walk W_1 and a v–w walk W_2, we denote by $W_1 * W_2$ the u–w walk obtained by concatenating W_1 and W_2.

The walk W is called a *trail* if $e_1, e_2, \ldots, e_l$ are distinct, and a *path* if $v_0, v_1, \ldots, v_l$ are distinct. Note that each trail in G can be regarded as a subgraph of G. We say that W is *closed* (in particular, a *closed* v_0*-walk*) if $v_0 = v_l$, and W is called a *cycle* if it is closed and $v_0, v_1, \ldots, v_{l-1}$ are distinct. A walk is said to be *trivial* if it is of length 0 (i.e., $l = 0$), and *nontrivial* otherwise. Note that any trivial walk is a trail, a path, and a cycle.

Let $A \subseteq V$ be a vertex set, and $\mathcal{S} = \{A_1, A_2, \ldots, A_k\}$ a partition of A, i.e., $\bigcup_{i=1}^{k} A_i = A$ and $A_i \cap A_j = \emptyset$ if $i \neq j$. A nontrivial walk is called an A*-walk* if its end vertices are both in A, and an A-walk is called an A*-path* if it is a path and its inner vertices are disjoint from A. An $\mathcal{S}$*-path* is an A-path whose end vertices are in distinct classes of $\mathcal{S}$, i.e., one is in A_i and the other is in A_j for distinct $i, j \in [k]$.

2.1.3 Connectivity

For a positive integer $k \in \mathbb{Z}_{>0}$, a vertex set $X \subsetneq V$ with $|X| = k$ is called a k*-cut* in G if $G - X$ is not connected. A graph is called k*-connected* if it contains more than k

vertices and no k'-cut for every $k' < k$. A *k-connected component* of G is a maximal k-connected induced subgraph $G[X]$ ($X \subseteq V$ with $|X| > k$). Similarly, an edge set $F \subseteq E$ with $|F| = k$ is called a *k-edge-cut* in G if $G - F$ is not connected. A graph is called *k-edge-connected* if it contains no k-edge-cut. A *k-edge-connected component* of G is a maximal k-edge-connected induced subgraph of G.

For nonempty vertex sets $X, Y, Z \subseteq V$, we say that X *separates* Y *and* Z *in* G if every two vertices $y \in Y \setminus X$ and $z \in Z \setminus X$ are contained in different connected components of $G - X$. In particular, each of Y and Z always separates Y and Z since $Y \setminus Y = Z \setminus Z = \emptyset$. Besides, if X separates Y and Z in G and $Y \setminus X \neq \emptyset \neq Z \setminus X$, then X is an $|X|$-cut in G. Similarly, for an edge set $F \subseteq E$ and nonempty vertex sets $Y, Z \subseteq V$, we say that F *separates* Y *and* Z *in* G if every two vertices $y \in Y$ and $z \in Z$ are contained in different connected components of $G - F$.

Using Menger's theorems (Theorems 1.6 and 1.7), we can characterize the existence of redundant vertices in considering A-paths (including s–t paths) or s–t trails by the connectivity of graphs as follows.

Proposition 2.1. *Let $A \subseteq V$ be a terminal set with $|A| \geq 2$, and suppose that G is connected. Then, every vertex $v \in V$ is contained in some A-path in G if and only if the graph G' obtained from G by adding a new vertex r so that r is adjacent to all terminals in A is 2-connected.*

Proof. First of all, we note that each A-path in G corresponds to a cycle in G' that contains the new vertex r (equivalent to two openly disjoint r–v paths for some vertex $v \in V$) and vice versa, since the two vertices adjacent to r in such a cycle are distinct terminals.

["If" part] Suppose that G' is 2-connected, i.e., $G' - u$ is connected for any $u \in V + r$. Fix an arbitrary vertex $v \in V$. Since G' is 2-connected, by Theorem 1.6, there exist two openly disjoint r–v paths in G', which form a cycle C containing r and v.

["Only if" part] Suppose that G' is not 2-connected, i.e., there exists a vertex $u \in V + r$ such that $G' - u$ is disconnected. Since $G = G' - r$ is connected, we have $u \in V$. Then, there exists a connected component of $G' - u$ that contains no terminal in A, since r is adjacent to all terminals in A. Fix an arbitrary vertex $v \in V \setminus A$ in such a connected component. Then, by Theorem 1.6, we cannot take two openly disjoint r–v paths in G', which means that there is no cycle in G that contains r and v. $\square$

Proposition 2.2. *Let $s, t \in V$ be distinct terminals. Then, every vertex $v \in V$ is contained in some s–t trail in G if and only if the graph G' obtained from G by adding a new vertex r so that r is adjacent to both s and t is 2-edge-connected.*

Proof. The proof is almost the same as the previous one with $A = \{s, t\}$. The only difference is that each s–t trail in G corresponds to a closed trail in G' that contains the new vertex r (equivalent to two edge-disjoint r–v paths for some vertex $v \in V$), and we use Theorem 1.7. $\square$

It should be remarked that 2-connected components and 2-edge-connected components of a given graph can be computed in polynomial time [35], which implies that one can find the set of redundant vertices in polynomial time.

2.1.4 Associated matrices

A graph is said to be *simple* if it contains neither a self-loop nor parallel edges (multiple edges connecting the same pair of distinct vertices). For a simple undirected graph $G = (V, E)$ with a total order $\prec$ on the vertex set V, the *Tutte matrix* T_G of G is a $V \times V$ matrix defined as follows: for two vertices $u, v \in V$ (possibly $u = v$), the (u, v)-entry t_{uv} of T_G is

$$t_{uv} = \begin{cases} x_e & (e = uv \in E \text{ and } u \prec v), \\ -x_e & (e = uv \in E \text{ and } v \prec u), \\ 0 & (\text{otherwise}), \end{cases}$$

where x_e $(e \in E)$ are algebraically independent indeterminates.

The *incidence matrix* B_G of a directed graph $G = (V, E)$ is a $V \times E$ matrix defined as follows: for each vertex $w \in V$ and each arc $e = uv \in E$, the (w, e)-entry $b_{w,e}$ of B_G is

$$b_{w,e} = \begin{cases} 1 & (w = u \neq v), \\ -1 & (w = v \neq u), \\ 0 & (\text{otherwise}). \end{cases}$$

The *generic incidence matrix* $\hat{B}_G$ is also a $V \times E$ matrix, whose (w, e)-entry $\hat{b}_{w,e}$ $(w \in V,\ e = uv \in E)$ is

$$\hat{b}_{w,e} = \begin{cases} x_e^+ + x_e^- & (w = u = v), \\ x_e^+ & (w = u \neq v), \\ x_e^- & (w = v \neq u), \\ 0 & (\text{otherwise}), \end{cases}$$

where x_e^+, x_e^- $(e \in E)$ are algebraically independent indeterminates. Note that, by substituting ± 1 for $x_e^\pm$ (respectively) for each arc $e \in E$, we obtain the incidence matrix B_G from the generic incidence matrix $\hat{B}_G$.

2.2 Group-Labeled Graphs

2.2.1 Definitions and notations

Throughout this thesis, let Γ be a group, for which we usually use multiplicative notation with denoting the identity element by 1_Γ. We sometimes use additive notation with denoting the identity element by 0, e.g., when $\Gamma \simeq \mathbb{Z}$. When we focus on the computational complexity, we assume that the following procedures can be done in constant time for any $\alpha, \beta \in \Gamma$: getting the inverse element $\alpha^{-1} \in \Gamma$, computing the product $\alpha\beta \in \Gamma$, and testing the identification $\alpha = \beta$. A Γ-*labeled graph* is a directed graph G with a mapping $\psi_G \colon E(G) \to \Gamma$ called a *label function*.

Let $G = (V, E)$ be a Γ-labeled graph and $W = (v_0, e_1, v_1, e_2, v_2, \dots, e_l, v_l)$ a walk in G. The *label* $\psi_G(W)$ of W in G is defined as the product $\psi_G(e_l) \cdots \psi_G(e_2) \cdot \psi_G(e_1)$ if W is directed (i.e., $e_i = v_{i-1}v_i$ for every $i \in [l]$), and otherwise by replacing $\psi_G(e_i)$ with $\psi_G(e_i)^{-1}$ in the product for each $i \in [l]$ with $e_i = v_i v_{i-1}$. Note that, for the reversed walk $\bar{W}$ of W, we have $\psi_G(\bar{W}) = \psi_G(W)^{-1}$. In particular, when we consider undirected walks, since an arc from a vertex $u \in V$ to a vertex $v \in V$ with label $\alpha \in \Gamma$ is essentially equivalent to the reversed arc vu with label α^{-1}, we identify such two arcs.

We say that W is *balanced* (or a *zero walk*) if $\psi_G(W) = 1_\Gamma$, and *unbalanced* (or a *non-zero walk*) otherwise. A Γ-labeled graph G' (or its edge set $E(G')$) is also said to be *balanced* if G' contains no unbalanced cycle[1], and *unbalanced* otherwise. In particular, if $\psi_{G'}(e) = 1_\Gamma$ for every edge $e \in E(G')$, then G' is obviously balanced, and we say such G' to be *trivially labeled*. Moreover, for a vertex set $A \subseteq V(G')$, a balanced Γ-labeled graph G' (or $E(G')$) is said to be A-*balanced* if it contains no non-zero A-path.

2.2.2 Shifting operation

The following operation is often used in dealing with group-labeled graphs (see, e.g., [6, 7, 49, 83]), which is sometimes referred to as *switching*.

Definition 2.3 (Shifting)**.** Let $G = (V, E)$ be a Γ-labeled graph. For a vertex $v \in V$ and an element $\alpha \in \Gamma$, *shifting* (*the label function* ψ_G) *by* α *at* v means the following operations: update ψ_G to $\psi'_G \colon E \to \Gamma$ defined by, for each $e \in E$,

$$\psi'_G(e) := \begin{cases} \alpha \cdot \psi_G(e) \cdot \alpha^{-1} & (e = vv, \text{ i.e., } e \text{ is a self-loop on } v), \\ \alpha \cdot \psi_G(e) & (e \in \delta_G^{\text{in}}(v), \text{ i.e., } e \text{ enters } v), \\ \psi_G(e) \cdot \alpha^{-1} & (e \in \delta_G^{\text{out}}(v), \text{ i.e., } e \text{ leaves } v), \\ \psi_G(e) & (\text{otherwise}). \end{cases}$$

By the definition, shifting at $v \in V$ does not change the label of any walk whose end vertices are not v, and neither that of any closed v-walk C up to conjugate, i.e., $\psi'_G(C) = \alpha \cdot \psi_G(C) \cdot \alpha^{-1}$. Note that shiftings at distinct vertices $u, v \in V$ do not interfere with each other. This implies that the order of applications of shifting does not make any effect on the resulting label function. Besides, a sequence of shiftings by $\alpha, \beta \in \Gamma$ in this order at a vertex $v \in V$ is equivalent to shifting by $\beta\alpha \in \Gamma$ at v.

[1] The balancedness of a cycle does not depend on the choices of the direction and the end vertex, since $\psi_G(\bar{C}) = \psi_G(C)^{-1}$ and $\psi_G(C') = \psi_G(e_1) \cdot \psi_G(C) \cdot \psi_G(e_1)^{-1}$, where $C = (v_0, e_1, v_1, e_2, v_2, \dots, e_l, v_l = v_0)$ and $C' = (v_1, e_2, v_2, \dots, e_l, v_l = v_0, e_1, v_1)$ are cycles in a Γ-labeled graph G.

Let G and H be Γ-labeled graphs. We say that G is *equivalent to* H if G is obtained from H by a sequence of shiftings, where we identify two arcs from a vertex u to a vertex w with label $\alpha \in \Gamma$ and from w to u with label α^{-1}. Note that, if G is obtained from H by shifting by $\alpha_v \in \Gamma$ at each vertex $v \in V$, then H is obtained from G by shifting by α_v^{-1} at each v. Moreover, for a vertex set $A \subseteq V(G)$, we say that G is *A-equivalent to* H if G is equivalent to H and $\alpha_v = 1_\Gamma$ for every vertex $v \in A$, i.e., shifting at any vertex in A is not necessary to obtain G from H. It should be noted that, if G is A-equivalent to H, then we have $\psi_G(W) = \psi_H(W)$ for every A-walk W in G (as well as in H), since only shifting at an end vertex of a walk changes its label.

We here describe a useful procedure using shifting. Recall that Γ denotes a group for which several elementary procedures can be done in constant time.

Proposition 2.4. *Let $G = (V, E)$ be a Γ-labeled graph. For any edge set $F \subseteq E$ such that $G[F]$ is a forest (i.e., contains no cycle), there exists a Γ-labeled graph H equivalent to G such that $H[F]$ is trivially labeled. Moreover, one can find such H in $\mathrm{O}(|V| + |E|)$ time.*

Proof. For each connected component T in $G[F]$ (i.e., T is a maximal tree in $G[F]$), perform the following procedure. Choose a root $r \in V(T)$ arbitrarily, and let $X := \{r\}$. Here, we may assume that all arcs in $E(T)$ are directed toward r (by reversing some edges with inversing their labels if necessary), i.e., for each vertex $v \in V(T)$, there uniquely exists a directed v–r path in T. While $X \neq V(T)$, take an in-neighbor $v \in N_T^{\mathrm{in}}(X)$, apply shifting the current label function ψ by $\psi(e)$ at v for a unique arc $e \in \delta_T^{\mathrm{out}}(v)$ from v to X in T (so that $\psi(e) = 1_\Gamma$ after the shifting), and update $X := X + v$.

The above procedure takes $\mathrm{O}(|V| + |E|)$ time in total, since it just performs breadth first search once for each T and shifting at most $|V| - 1$ times in total (note that the label of each arc changes at most twice). After that, the label function ψ satisfies that $\psi(e) = 1_\Gamma$ for every edge $e \in F$, and hence the resulting graph is desired H. $\square$

Using this procedure, we can prove the following characterization for the balancedness.

Proposition 2.5. *For a Γ-labeled graph $G = (V, E)$, the following conditions are equivalent:*

(1) *G is balanced,*

(2) *G is equivalent to a trivially-labeled Γ-labeled graph,*

(3) *G is r-equivalent to a trivially-labeled Γ-labeled graph for some vertex $r \in V$, and*

(4) *G is r-equivalent to a trivially-labeled Γ-labeled graph for any vertex $r \in V$.*

Moreover, each condition can be tested in $\mathrm{O}(|V| + |E|)$ time.

Proof. [(2) $\Longrightarrow$ (1)] A trivially-labeled Γ-labeled graph is obviously balanced. Besides, any shifting at any vertex does not change the label of any cycle up to conjugate, and hence it does not violate the balancedness.

[(4) $\Longrightarrow$ (3) and (3) $\Longrightarrow$ (2)] Obvious.

[(1) $\Longrightarrow$ (4)] For each connected component of G, take a spanning tree, and let $F \subseteq E$ be the set of edges which are contained in one of the spanning trees. Then, $G[F]$ is a forest. Apply Proposition 2.4 to this F, we can obtain a Γ-labeled graph H equivalent to G such that $H[F]$ is trivially labeled, and H itself is in fact trivially labeled.

Suppose to the contrary that $\psi_H(e) \neq 1_\Gamma$ for some $e \in E(H) = E$. Since the edge set F forms spanning trees of all connected components of H as well as of G, there uniquely exists a cycle C in $H[F + e]$, which traverses e. The label $\psi_G(C)$ of C is conjugate to $\psi_H(C) = \psi_H(e) \neq 1_\Gamma$ (since $\psi_H(e) = 1_\Gamma$ for every edge $e \in F$), which contradicts that G is balanced.

Recall that, in the procedure to obtain H from G (cf. the proof of Proposition 2.4), we can choose a root of each tree arbitrarily, and we do not perform shifting at any root. Thus, we complete this proof by choosing r as a root of the spanning tree of the connected component of G that contains r. □

Proposition 2.5 can be extended as follows. This extended version plays an important role in considering the labels of A-paths.

Proposition 2.6. *For a vertex set $A \subseteq V$, a Γ-labeled graph $G = (V, E)$ is A-balanced if and only if G is A-equivalent to a trivially-labeled Γ-labeled graph. Moreover, one can test whether G is A-balanced or not in* $\mathrm{O}(|V| + |E|)$ *time.*

Proof. "If" part is easy to see similarly to the first part of the proof of Proposition 2.5. A trivially-labeled Γ-labeled graph whose vertex set is V is A-balanced for any subset $A \subseteq V$, and any shifting at a vertex not in A does not change the label of any A-path.

To see the converse direction, consider a similar procedure to that in the last part of the proof of Proposition 2.5, i.e., for each connected component of G, take a spanning tree and apply shifting to make it trivially labeled. The only difference is that, for each connected component of G that contains a vertex in A, we must choose its root r (which can be chosen arbitrarily in the previous proof) so that $r \in A$.

After the procedure, the resulting graph is trivially labeled by the previous proof. Suppose that we applied shifting by $\alpha_v \in \Gamma$ at each vertex $v \in V$ in the procedure, where we define $\alpha_r := 1_\Gamma$ for each root $r \in V$. Then, it suffices to show that $\alpha_v = 1_\Gamma$ for every vertex $v \in A$. Suppose to the contrary that $\alpha_v \neq 1_\Gamma$ for some vertex $v \in A$. Clearly, v was not chosen as a root. Let T_v be the spanning tree taken in the procedure that contains v, and $r_v \in V(T_v) \cap A$ its root. Then, the original label of the v–r_v path P in T_v is $\alpha_v \neq 1_\Gamma$. If P is an A-path, then it contradicts that G is A-balanced. Otherwise, P intersects some other vertices in A, and at least one subpath of P is a non-zero A-path (since otherwise $\psi_G(P) = 1_\Gamma$), a contradiction. □

2.2.3 Labels of walks

The balancedness of group-labeled graphs just takes care of the labels of cycles. However, it characterizes the labels of all walks as follows.

Proposition 2.7. *A Γ-labeled graph $G = (V, E)$ is balanced if and only if every two s–t walks in G have the same label for every fixed pair of two vertices $s, t \in V$ (possibly $s = t$).*

Proof. By Proposition 2.5, G is balanced if and only if G is equivalent to a trivially-labeled Γ-labeled graph, in which the labels of all walks are obviously 1_Γ. Fix two vertices $s, t \in V$ (possibly $s = t$). Recall that any shifting at a vertex $v \in V \setminus \{s, t\}$ does not change the label of any s–t walk. Moreover, when a label function $\psi' \colon E \to \Gamma$ is obtained from a label function $\psi \colon E \to \Gamma$ by shifting by $\alpha \in \Gamma$ at s, for every s–t walk W in G, we have $\psi'(W) = \psi(W) \cdot \alpha^{-1}$ if $s \neq t$, and $\psi'(W) = \alpha \cdot \psi(W) \cdot \alpha^{-1}$ otherwise. Similarly, when ψ' is obtained from ψ by shifting by $\alpha \in \Gamma$ at t, for every s–t walk W in G, we have $\psi'(W) = \alpha \cdot \psi(W)$ if $s \neq t$, and $\psi'(W) = \alpha \cdot \psi(W) \cdot \alpha^{-1}$ otherwise. This means that any shifting makes no effect on whether every two s–t walks in G have the same label or not, and hence we have done. $\square$

With the aid of Menger's theorems (Theorems 1.6 and 1.7), analogous characterizations work for s–t paths and s–t trails in nonredundant graphs. Recall that there are simple characterizations for the existence of redundant vertices (Propositions 2.1 and 2.2).

Proposition 2.8. *Let $G = (V, E)$ be a Γ-labeled graph with distinct vertices $s, t \in V$ in which every vertex $v \in V$ is contained in some s–t path, and suppose that G contains no self-loop. Then, G is balanced if and only if every two s–t paths in G have the same label.*

Proof. "Only if" part is obvious from Proposition 2.7, and we show "if" part using Proposition 2.1 as follows. Suppose that G is unbalanced, i.e., G contains an unbalanced cycle C of length at least 2 since G contains no self-loop. Let G'_C be the graph obtained from G by adding two new vertices r and v_C so that r is adjacent to s and t, and v_C is adjacent to all vertices in $V(C)$. By Proposition 2.1 and Theorem 1.6, there exist two openly disjoint r–v_C paths in G'_C (note that adding v_C does not violate the 2-connectivity of G' in Proposition 2.1 since $|V(C)| \geq 2$).

Take two openly disjoint r–v_C paths P_s and P_t in G'_C so that $|E(P_s) \cup E(P_t)|$ is minimum and P_x intersects $x \in \{s, t\}$. Then, each of P_s and P_t contains exactly one vertex in C, say x and y, respectively. Without loss of generality, assume that x is the end vertex of C (recall that the choice of the end vertex makes no effect on the balancedness of the cycle). By concatenating $P_s[s, x]$, either $C[x, y]$ or $\bar{C}[x, y]$, and $P_t[y, t]$, we can construct two s–t paths in G of distinct labels, since $\psi_G(C[x, y]) \neq \psi_G(\bar{C}[x, y])$ (otherwise, $\psi_G(C) = \psi_G(C[y, x]) \cdot \psi_G(C[x, y]) = \psi_G(\bar{C}[x, y])^{-1} \cdot \psi_G(C[x, y]) = 1_\Gamma$, a contradiction). $\square$

Proposition 2.9. *Let $G = (V, E)$ be a Γ-labeled graph with distinct vertices $s, t \in V$ in which every vertex $v \in V$ is contained in some s–t trail. Then, G is balanced if and only if every two s–t trails in G have the same label.*

Proof. The proof is almost the same as the previous one. It suffices to remark that we use Theorem 1.7 and Proposition 2.2 instead of Theorem 1.6 and Proposition 2.1. $\square$

In another direction, Proposition 2.7 can be extended to the possible labels of walks in an arbitrary Γ-labeled graph as follows.

Proposition 2.10. *Let $G = (V, E)$ be a connected Γ-labeled graph with a vertex $s \in V$. The set of possible labels of closed s-walks in G coincides with the subgroup Γ_s of Γ that is generated by $\{\, \psi_H(e) \mid e \in E \,\}$, where H is an arbitrary Γ-labeled graph s-equivalent to G that contains a trivially-labeled spanning tree.*

Proof. Fix an arbitrary Γ-labeled graph H s-equivalent to G that contains a trivially-labeled spanning tree T. Note that such H exists by Propositions 2.4 and 2.5. Since H is s-equivalent to G, we have $\psi_H(C) = \psi_G(C)$ for every closed s-walk C in G. Hence, it suffices to show that the set L of possible labels of closed s-walks in H coincides with the subgroup Γ_s of Γ that is generated by $S = \{\, \psi_H(e) \mid e \in E \,\}$.

It is obvious that $L \subseteq \Gamma_s$, since the label of any walk in H is a product of some elements in S and their inverses. To the contrary, for any $\alpha \in \Gamma_s$, one can construct a closed s-walk of label α as follows. Suppose that $\alpha = \alpha_k \cdots \alpha_2\alpha_1$, where $\alpha_i \in S$ or $\alpha_i^{-1} \in S$ for each $i \in [k]$. For the sake of simplicity, we assume that H contains the reversed arc $\bar{e}$ with label $\psi_H(e)^{-1}$ for each arc $e \in E$ (note that the two arcs e and $\bar{e}$ are essentially equivalent and identified usually). Then, there exists an arc $e_i = u_iv_i \in E$ with $\psi_H(e_i) = \alpha_i$ for each $i \in [k]$. By concatenating a unique s–u_1 path in the trivially-labeled spanning tree T, unique v_i–u_{i+1} paths in T for all $i \in [k-1]$, and a unique v_k–s path in T using the arcs $e_i = u_iv_i$ ($i \in [k]$), we obtain a desired closed s-walk of label α. $\square$

Proposition 2.11. *Let $G = (V, E)$ be a connected Γ-labeled graph with distinct vertices $s, t \in V$. The set of possible labels of s–t walks in G coincides with the left coset $\alpha_t^{-1}\Gamma_s$, where Γ_s is the subgroup of Γ in Proposition 2.10 and shifting by $\alpha_t \in \Gamma$ at t is applied to obtain from G an s-equivalent Γ-labeled graph H that contains a trivially-labeled spanning tree.*

Proof. Fix an arbitrary Γ-labeled graph H s-equivalent to G that contains a trivially-labeled spanning tree T, and suppose that shifting by $\alpha_t \in \Gamma$ at t is applied to obtain H from G. Let G' be the Γ-labeled graph obtained from G by shifting by α_t at t, which changes the label $\gamma \in \Gamma$ of each s–t walk in G into $\alpha_t\gamma$. Hence, it suffices to show that the set of possible labels of s–t walks in G' coincides with the subgroup Γ_s. Moreover, since H is $\{s, t\}$-equivalent to G', it suffices to prove it for H.

The proof is almost the same as Proposition 2.10. The only difference is that, when we construct an s–t walk of label $\alpha \in \Gamma_s$, we have to replace the last v_k–s path with a unique v_k–t path in the trivially-labeled spanning tree T. □

Corollary 2.12. *Let $G = (V, E)$ be a connected Γ-labeled graph with distinct vertices $s, t \in V$. The set of possible labels of t–s walks in G coincides with the right coset $\Gamma_s \alpha_t$, where the subgroup $\Gamma_s \leq \Gamma$ and the element $\alpha_t \in \Gamma$ are the same as those in Proposition* 2.11.

2.3 Matroids and Delta-Matroids

2.3.1 Definitions and notations

A matroid is defined by a pair of a finite set and a family of its subsets, which abstracts combinatorial structure of the linear independence of sets of vectors. We here describe only necessary concepts and properties utilized in this thesis, and see, e.g,, [70] for more details.

A pair $(E, \mathcal{I})$ of a finite set E and a family $\mathcal{I} \subseteq 2^E$ is called a *matroid* if the following conditions hold:

(I0) $\emptyset \in \mathcal{I}$,

(I1) $X \subseteq Y \in \mathcal{I} \implies X \in \mathcal{I}$, and

(I2) $X, Y \in \mathcal{I}$ and $|X| < |Y| \implies \exists e \in E \colon X + e \in \mathcal{I}$.

Let $\mathbf{M} = (E, \mathcal{I})$ be a matroid. The set E is called the *ground set* of $\mathbf{M}$, and each subset $X \subseteq E$ is said to be *independent* in $\mathbf{M}$ if $X \in \mathcal{I}$ and *dependent* otherwise. A maximal independent set is called a *base*, and a minimal dependent set is called a *circuit*. Note that every two bases are of the same size by Condition (I2).

The *rank function* $r_{\mathbf{M}} \colon E \to \mathbb{Z}_{\geq 0}$ of $\mathbf{M}$ is defined by

$$r_{\mathbf{M}}(X) := \max\{\, |Y| \mid Y \subseteq X,\ Y \in \mathcal{I} \,\} \qquad (X \subseteq E).$$

The rank function $r \colon E \to \mathbb{Z}_{\geq 0}$ of any matroid on a finite set E satisfies the following conditions:

(R0) $0 \leq r(X) \leq |X|$ for every subset $X \subseteq E$,

(R1) $X \subseteq Y \subseteq E \implies r(X) \leq r(Y)$, and

(R2) $r(X) + r(Y) \geq r(X \cup Y) + r(X \cap Y)$ for every subsets $X, Y \subseteq E$.

To the contrary, such a set function r in fact coincides with the rank function of some matroid on E, and hence we can define a matroid by a pair of a finite set E and a set function $r \colon 2^E \to \mathbb{Z}_{\geq 0}$ satisfying Conditions (R0)–(R2).

Let $X \subseteq E$ be a subset. The *restriction of* $\mathbf{M}$ *to* X is a matroid on X such that a subset $Y \subseteq X$ is independent if and only if Y is independent in $\mathbf{M}$. We denote by $\mathbf{M}|X$ the restriction of $\mathbf{M}$ to X. The *contraction of* $\mathbf{M}$ *by* X is a matroid on $E \setminus X$ such that a subset $Y \subseteq E \setminus X$ is independent if and only if $Y \cup B_X$ is independent in $\mathbf{M}$ for any base $B_X \subseteq X$ in $\mathbf{M}|X$. We denote by $\mathbf{M}/X$ the contraction of $\mathbf{M}$ by X. From the viewpoint of the rank functions, the restriction of a matroid is literally the restriction of the domain. The rank function $r_{\mathbf{M}/X} \colon E \setminus X \to \mathbb{Z}_{\geq 0}$ of the contraction of $\mathbf{M}$ by X is written by

$$r_{\mathbf{M}/X}(Y) = r_{\mathbf{M}}(Y \cup X) - r_{\mathbf{M}}(X) \qquad (Y \subseteq E \setminus X). \tag{2.1}$$

A *delta-matroid* is a pair of a finite set E and a nonempty family $\mathcal{F} \subseteq 2^E$ of subsets of E with the following condition:

(DM) $X, Y \in \mathcal{F}$ and $e \in X \triangle Y \implies \exists f \in X \triangle Y \colon X \triangle \{e, f\} \in \mathcal{F}$,

where $X \triangle Y := (X \setminus Y) \cup (Y \setminus X)$ denotes the *symmetric difference* between X and Y. For a delta-matroid $\mathbf{D} = (E, \mathcal{F})$, the set E is called the *ground set* of $\mathbf{D}$, and each subset $X \subseteq E$ is said to be *feasible* in $\mathbf{D}$ if $X \in \mathcal{F}$ and *infeasible* otherwise.

It should be noted that f may coincide with e in Condition (DM). When $f \neq e$ for every $X, Y \in \mathcal{F}$ and every $e \in X \triangle Y$, every two feasible sets have the same parity (i.e., either all feasible sets are odd-size, or all are even-size). Such a delta-matroid is called an *even delta-matroid*. The following property is well-known.

Proposition 2.13. *Let $\mathbf{D} = (E, \mathcal{F})$ be an even delta-matroid, and $X \in \mathcal{F}$ a feasible set not maximal (i.e., there exists a feasible set $Y \in \mathcal{F}$ with $X \subsetneq Y$). Then, there exist distinct elements $e, f \in E \setminus X$ such that $X \cup \{e, f\} \in \mathcal{F}$.*

Proof. Suppose that $X \in \mathcal{F}$ is not maximal. Take a feasible set $Y \in \mathcal{F}$ with $X \subsetneq Y$ so that $|X \triangle Y| = |Y \setminus X|$ is minimized. Note that $|X \triangle Y|$ is even since $\mathbf{D}$ is an even delta-matroid. If $|X \triangle Y| = 2$, then $\{e, f\} = Y \setminus X$ must hold in Condition (DM).

Suppose that $|X \triangle Y| \geq 4$. Then, by Condition (DM), for each $e \in Y \setminus X$ we have $Y' := Y \setminus \{e, f\} \in \mathcal{F}$ for some $f \in (Y \setminus X) - e$ (since $\mathbf{D}$ is an even delta-matroid). It is obvious that $|X \triangle Y'| = |Y' \setminus X| = |Y \setminus X| - 2 < |X \triangle Y|$, a contradiction. □

Any matroid is a delta-matroid, and the base family of any matroid forms an even delta-matroid. In particular, when $\mathcal{F}$ is the base family of a matroid, $|X \cap \{e, f\}| = |Y \cap \{e, f\}| = 1$ hold for every $X, Y \in \mathcal{F}$ and every $e \in X \triangle Y$ in Condition (DM), since every two bases in a fixed matroid are of the same size.

2.3.2 Examples

We here show several well-known examples of matroids that are related to our work.

Linear (Vector, Matric) Matroids. Let E be a finite set, $\mathbb{F}$ a field, and $d \in \mathbb{Z}_{>0}$ a positive integer. For a multiset $V = \{ v_e \in \mathbb{F}^d \mid e \in E \}$ of vectors indexed by E, the *vector matroid* is defined on E so that a subset $X \subseteq E$ is independent if and only if the multiset $V(X) := \{ v_e \mid e \in X \}$ of vectors corresponding to X is linearly independent. In other words, for a matrix $Z = (v_e)_{e \in E} \in \mathbb{F}^{d \times E}$, a subset $X \subseteq E$ is independent if and only if the submatrix $Z(X) := (v_e)_{e \in X} \in \mathbb{F}^{d \times X}$ corresponding to X is column-full-rank. This matroid is also called the *matric matroid* of Z. The rank function $r \colon 2^E \to \mathbb{Z}_{\geq 0}$ is written by $r(X) = \dim \operatorname{span} V(X) = \operatorname{rank} Z(X)$ for each $X \subseteq E$.

A matroid on a finite set E is said to be *linear* (or *linearly representable over* $\mathbb{F}$) if it coincides with the matric matroid of some matrix over a field $\mathbb{F}$ whose column set is indexed by E. Moreover, it is said to be *linear represented* if we are given such a matrix.

Graphic (Cycle) Matroids. Let $G = (V, E)$ be an undirected graph. The *cycle matroid* of G is defined on the edge set E so that a subset $F \subseteq E$ is independent if and only if $G[F]$ is a forest (i.e., contains no cycle). The rank function $r_G \colon 2^E \to \mathbb{Z}_{\geq 0}$ is written by

$$r_G(F) = \sum_{F' \in \operatorname{comp}(F)} \bigl(|V(F')| - 1\bigr) \qquad (F \subseteq E),$$

where $\operatorname{comp}(F)$ denotes the partition of F into the edge sets of the connected components of $G[F]$, i.e., $\{ G[F'] \mid F' \in \operatorname{comp}(F) \}$ is the set of connected components of $G[F]$.

A matroid on a finite set E is said to be *graphic* if it coincides with the cycle matroid of some undirected graph whose edge set is E. It should be noted that any graphic matroid is linear. Indeed, for any undirected graph $G = (V, E)$, the matric matroid of its incidence matrix B_G (over any field) coincides with the cycle matroid of G.

Bicircular Matroids. Let $G = (V, E)$ be an undirected graph. The *bicircular matroid* of G is defined on the edge set E so that a subset $F \subseteq E$ is independent if and only if each connected component of $G[F]$ contains at most one cycle. The rank function $\hat{r}_G \colon 2^E \to \mathbb{Z}_{\geq 0}$ is written by

$$\hat{r}_G(F) = \sum_{F' \in \operatorname{comp}(F)} \bigl(|V(F')| - 1 + \hat{\rho}_G(F')\bigr) \qquad (F \subseteq E),$$

where $\hat{\rho}_G \colon 2^E \to \mathbb{Z}$ is defined as

$$\hat{\rho}_G(F') := \begin{cases} 1 & (G[F'] \text{ contains a cycle}), \\ 0 & (\text{otherwise}). \end{cases}$$

It is worth remarking that any bicircular matroid is also linear. Indeed, for any undirected graph, its generic incidence matrix represents this matroid.

Frame Matroids. Let $G = (V, E)$ be a Γ-labeled graph. The *frame matroid* (or the *bias matroid*) of G is defined on the edge set E so that a subset $F \subseteq E$ is independent if and only if each connected component of $G[F]$ contains at most one cycle, which is unbalanced [97]. The rank function $r_G \colon 2^E \to \mathbb{Z}_{\geq 0}$ is written by

$$r_G(F) = \sum_{F' \in \mathrm{comp}(F)} \bigl(|V(F')| - 1 + \rho_G(F')\bigr) \qquad (F \subseteq E),$$

where $\rho_G \colon 2^E \to \mathbb{Z}$ is defined as

$$\rho_G(F') := \begin{cases} 1 & (G[F'] \text{ is unbalanced}), \\ 0 & (\text{otherwise}). \end{cases}$$

It should be remarked that the frame matroids of group-labeled graphs commonly generalize the cycle matroids and the bicircular matroids of undirected graphs as follows. The frame matroid of a trivially-labeled group-labeled graph coincides with the cycle matroid of the underlying undirected graph. For an undirected graph $G = (V, E)$, let $\Gamma := \mathbb{Z}_2^E$, and assign the label $\mathbf{1}_e \in \mathbb{Z}_2^E$ (which denotes the element such that all entries but the e-th entry is $0 \in \mathbb{Z}_2$ and the e-th entry is $1 \in \mathbb{Z}_2$) to each edge $e \in E$ (note that the direction can be ignored since $-\mathbf{1}_e = \mathbf{1}_e$). Then, the resulting Γ-labeled graph contains no balanced cycle, and hence its frame matroid coincides with the bicircular matroid of G.

Frame matroids play an important role also in the matroid representation theory (see, e.g., [70]). In particular, the following theorem is closely related to our work (cf. Theorem 5.6).

Theorem 2.14 (Dowling [12], and see [70, Theorem 6.10.10])**.** *Let Γ be a group, and $\mathbb{F}$ a field. The frame matroids of all Γ-labeled graphs are linearly representable over $\mathbb{F}$ if and only if Γ is isomorphic to a subgroup of its multiplicative group $\mathbb{F}^\times$.*

Matching (Delta-)Matroids. For an undirected graph $G = (V, E)$, the *matching matroid* is defined on the vertex set V so that a subset $X \subseteq V$ is independent if and only if X is covered by some matching in G, i.e., the family of independent set is written by

$$\{\, X \subseteq V \mid \exists M \subseteq E \colon |M| = 2|V(M)| \text{ and } X \subseteq V(M) \,\}.$$

Moreover, let $\mathcal{F}$ be the family of vertex subsets each of which is exactly covered by some matching in G, i.e.,

$$\mathcal{F} := \{\, X \subseteq V \mid \exists M \subseteq E \colon |M| = 2|V(M)| \text{ and } X = V(M) \,\}.$$

Then, the pair $(V, \mathcal{F})$ is a delta-matroid (in particular, an even delta-matroid), which is called the *matching delta-matroid* of G.

The matching matroid of any undirected graph G is linear, and the Tutte matrix of G in fact represents this matroid.

2.4 Matroid Matching

2.4.1 Formulations

The *matroid matching problem* was introduced by Lawler [56] as a common generalization of the maximum matching problem and the matroid intersection problem, which are both fundamental problems in combinatorial optimization. This problem is literally to find a maximum matching under a matroidal constraint (for a given undirected graph $G = (V, E)$ and matroid $\mathbf{M} = (V, \mathcal{I})$ on the vertex set V, to find a maximum matching $M \subseteq E$ in G with $V(M) \in \mathcal{I}$), and it has several equivalent formulations such as the *matroid parity problem* and the *matchoid problem*. We adopt one of them which is also called the *polymatroid matching problem* (see [62, Section 11.1] for the details).

Definition 2.15. For a positive integer $d \in \mathbb{Z}_{>0}$, a *d-polymatroid* is a pair of a finite set E and an integer-valued set function $f\colon 2^E \to \mathbb{Z}_{\geq 0}$ with the following conditions:

(P0) $0 \leq f(X) \leq d|X|$ for every subset $X \subseteq E$,

(P1) $X \subseteq Y \subseteq E \implies f(X) \leq f(Y)$, and

(P2) $f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y)$ for every subsets $X, Y \subseteq E$.

Note that 1-polymatroids are equivalent to matroids (cf. Section 2.3.1).

A subset $M \subseteq E$ is called a *matching* in a 2-polymatroid (E, f) if $f(M) = 2|M|$. The *matroid matching problem* is formulated as finding a matching in a given 2-polymatroid (E, f) of the maximum cardinality, whose value is denoted by $\nu(E, f)$.

Matroid Matching Problem

Input: A 2-polymatroid (E, f).

Goal: Find a matching $M \subseteq E$ in (E, f) such that $|M|$ is maximum.

It is easy to reformulate the literal formulation in this form as follows. For a given undirected graph $G = (V, E)$ and matroid $\mathbf{M} = (V, \mathcal{I})$, define an integer-valued set function $f\colon 2^E \to \mathbb{Z}_{\geq 0}$ by $f(X) := r_{\mathbf{M}}(V(X))$ for each subset $X \subseteq E$, where $r_{\mathbf{M}}\colon 2^V \to \mathbb{Z}_{\geq 0}$ denotes the rank function of $\mathbf{M}$. It follows from the properties of the rank function of a matroid that f satisfies Conditions (P0)–(P2) with $d = 2$, and it is also easily checked that $f(X) = 2|X|$ if and only if X is a matching in G with $r_{\mathbf{M}}(V(X)) = |V(X)|$.

The *matroid intersection problem* is a fundamental problem in combinatorial optimization, which admits a good characterization [14] and polynomial-time algorithms originated by [14, 55]. In this problem, we are given two matroids $\mathbf{M}_1, \mathbf{M}_2$ on the same ground set E, and required to find a maximum common independent set. This problem can be formulated as the matroid matching problem by defining an integer-valued set function $f\colon 2^E \to \mathbb{Z}_{\geq 0}$ as $f := r_{\mathbf{M}_1} + r_{\mathbf{M}_2}$.

Let $d \in \mathbb{Z}_{>0}$ be a positive integer. As an extension of the linear matroids (cf. Section 2.3.2), a d-polymatroid is said to be *linear* if it can be represented by some matrix. Specifically, for a field $\mathbb{F}$, a d-polymatroid (E, f) is *linearly representable over* $\mathbb{F}$ if there exists a matrix $Z = (Z_e)_{e \in E} \in \mathbb{F}^{r \times dE}$ obtained by concatenating $r \times d$ matrices $Z_e \in \mathbb{F}^{r \times d}$ ($e \in E$) such that $f(X) = \operatorname{rank} Z(X)$ for every subset $X \subseteq E$, where $r \in \mathbb{Z}_{>0}$ is a positive integer and $Z(X) = (Z_e)_{e \in X}$ denotes the submatrix of Z obtained by extracting the corresponding columns.

The matroid matching problem is called the *linear matroid parity problem* if the input 2-polymatroid is linearly represented. We call a subset $M \subseteq E$ a *matching* for a matrix $Z \in \mathbb{F}^{r \times 2E}$ if $\operatorname{rank} Z(M) = 2|M|$, and let $\nu(Z)$ denotes the maximum cardinality of a matching for Z (recall that $\nu(E, f)$ denotes the maximum cardinality of a matching in a 2-polymatroid (E, f)).

Linear Matroid Parity Problem

Input: A finite set E and a matrix $Z \in \mathbb{F}^{r \times 2E}$ over a field $\mathbb{F}$, where $r \in \mathbb{Z}_{>0}$.

Goal: Find a matching $M \subseteq E$ for Z such that $|M|$ is maximum.

2.4.2 Key theorem

In this section, we describe key theorems for matroid matching due to Lovász [59, 61]. Before we state the theorem, we define necessary concepts and notations.

Let (E, f) be a 2-polymatroid, and $X \subseteq E$ a subset. The *contraction of* (E, f) *by* X is a 2-polymatroid $(E \setminus X, f_X)$ such that $f_X(Y) = f(Y \cup X) - f(X)$ for every subset $Y \subseteq E \setminus X$ (cf. (2.1)). Besides, the *span* of X is defined as $\operatorname{span}_f(X) := \{\, e \in E \mid f(X \cup \{e\}) = f(X) \,\}$.

A subset $C \subseteq E$ is called a *circuit* in (E, f) if $f(C) = 2|C| - 1$ and $f(C - e) = 2|C| - 2$ (i.e., $C - e$ is a matching) for every element $e \in C$. A subset $D \subseteq E$ is called a *double circuit* in (E, f) if $f(D) = 2|D| - 2$ and $f(D - e) = 2|D| - 3$ for every element $e \in D$. It is known that every double circuit $D \subseteq E$ has a unique partition $\{D_1, D_2, \ldots, D_r\}$ such that $\{\, D \setminus D_i \mid i = 1, 2, \ldots, r \,\}$ coincides with the set of all circuits included in D (cf. [60]). A double circuit is said to be *nontrivial* if it is partitioned into at least three parts, i.e., $r \geq 3$. The *kernel* of a double circuit D is defined to be $\bigcap_{1 \leq i \leq r} \operatorname{span}_f(D \setminus D_i)$, i.e., the set of all elements spanned by all circuits in D.

Let k be a nonnegative integer. A subset $F \subseteq E$ is called a *k-flower* if $f(F) = 2|F| - 1$ and $|F| = k + 1$, and a *k-double-flower* if $f(F) = 2|F| - 2$, $|F| = k + 2$, and F includes no matching of size $k + 1$. It is easy to see that every k-flower $F \subseteq E$ has a unique partition $\{M_F, C_F\}$ such that M_F is a matching and C_F is a circuit (hence, $f(F) = f(M_F) + f(C_F)$), and every k-double-flower $F \subseteq E$ has a unique partition $\{M_F, D_F\}$ such that M_F is a matching and D_F is a double circuit (hence, $f(F) = f(M_F) + f(D_F)$) (cf. [59]).

Theorem 2.16 (Lovász [59], and see also [62, Theorem 11.2.7])**.** *Let (E, f) be a 2-polymatroid and define $\nu := \nu(E, f)$. Then, at least one of the following conditions holds:*

(i) $f(E) = 2\nu + 1$;

(ii) *E has a partition $\{E_1, E_2\}$ into nonempty subsets with $\nu = \nu(E_1, f) + \nu(E_2, f)$;*

(iii) *E has an element e which is contained in the span of every maximum matching*;

(iv) *(E, f) contains a nontrivial ν-double-flower.*

This theorem suggests an algorithmic approach to the matroid matching problem. If we encounter (i) or (ii), then we can identify a solution or reduce the problem to smaller ones. Similarly, if we encounter (iii), we can consider the contraction of (E, f) by e, which also reduces the problem size. The only difficult situation is Case (iv). However, as shown in [59], if the rank of the kernel is nonzero for every nontrivial double circuit in any contraction of (E, f), then the problem can be reduced to a smaller one by contracting an element in the kernel. The proof of this theorem is constructive, which means that it leads to a polynomial-time algorithm for the matroid matching problem if we can appropriately reduce the problem whenever we encounter (iv) (see [61, 62] for the complete description).

2.4.3 Algorithms

In general, the matroid matching problem cannot be solved in polynomial time. Specifically, there exists a 2-polymatroid for which an exponentially large number of evaluations of a function value are required to find a maximum matching. We show such an example due to Lovász [61] (also by Jensen and Korte [41]) as follows.

Let E be a finite set, and (E, f_1) and (E, f_2) 2-polymatroids defined as follows: for a positive integer $k \leq |E|$ and a subset $Y \subseteq E$ with $|Y| = k$,

$$f_1(X) := \begin{cases} 2|X| & (|X| < k), \\ 2k - 1 & (|X| = k), \\ 2k & (|X| > k), \end{cases}$$

$$f_2(X) := \begin{cases} 2|X| & (|X| < k), \\ 2k - 1 & (|X| = k \text{ and } X \neq Y), \\ 2k & (|X| > k \text{ or } X = Y). \end{cases}$$

It is easy to see these two are indeed 2-polymatroids. The 2-polymatroid (S, f_1) is uniform, and each subset $X \subseteq E$ with $|X| = k - 1$ is a maximum matching in (S, f_1). On the other hand, the subset Y is a unique maximum matching in (S, f_2).

To find maximum matchings in these two 2-polymatroids, we have to distinguish them, which requires at least $\binom{|E|}{k}$ evaluations of function values, since $f_1(X) = f_2(X)$ for every subset $X \subseteq E$ with $|X| = k$ and $X \neq Y$. When $k = \lfloor |E|/2 \rfloor$ for example, the number is exponentially large with respect to $|E|$, and this implies that any algorithm needs an exponentially large number of evaluations of function values for finding a maximum matching in (S, f_1) or (S, f_2).

While the general setting has such difficulty, Lovász [61] proposed a polynomial-time algorithm for a large class of this problem (cf. Theorem 2.16). Originated by his algorithm, several efficient algorithms for the linear matroid parity problem (which always admits Lovász' algorithm) have been developed.

Gabow and Stallmann [22] provided a combinatorial algorithm inspired by Edmonds' blossom algorithm for the maximum matching problem. Their algorithm runs in $\mathrm{O}(r^3|E|)$ time (or $\mathrm{O}(r^\omega|E|)$ time if fast matrix multiplication is used), which is known to be the fastest among the deterministic algorithms even today. Orlin [69] gave a simpler one with the same running time bound under no use of fast matrix multiplication, which solves the problem as a sequence of the matroid intersection problem. Cheung, Lau, and Leung [5] proposed an algebraic algorithm for this problem inspired by Harvey's work [31], which runs in $\mathrm{O}(r^2|E|)$ time (or $\mathrm{O}(r^{\omega-1}|E|)$ time if fast matrix multiplication is used).

Theorem 2.17 (Gabow–Stallmann [22])**.** *One can solve the linear matroid parity problem in* $\mathrm{O}(r^\omega|E|)$ *time by a deterministic algorithm.*

Theorem 2.18 (Cheung–Lau–Leung [5])**.** *One can solve the linear matroid parity problem in* $\mathrm{O}(r^{\omega-1}|E|)$ *time by a randomized algorithm.*

2.4.4 Weighted problems

A weighted version of the matroid matching problem is as follows. For a 2-polymatroid (E, f), a subset $B \subseteq E$ is called a *parity base* in (E, f) if $f(B) = 2|B| = f(E)$. The *weighted matroid matching problem* is, for a given 2-polymatroid (E, f) and weight $w \in \mathbb{R}^E$, to find a parity base $B \subseteq E$ in (E, f) such that the total weight $w(M) := \sum_{e \in M} w_e$ is minimum.

Weighted Matroid Matching Problem

Input: A 2-polymatroid (E, f) and a weight $w \in \mathbb{R}^E$.

Goal: Find a parity base $B \subseteq E$ in (E, f) such that $w(B)$ is minimum.

Also for the linearly represented case, we call a subset $B \subseteq E$ a *parity base* for a matrix $Z \in \mathbb{F}^{r \times 2E}$ if $\operatorname{rank} Z(B) = 2|B| = \operatorname{rank} Z$. The *weighted linear matroid parity problem* is analogously defined as follows.

Weighted Linear Matroid Parity Problem

Input: A finite set E, a matrix $Z \in \mathbb{F}^{r \times 2E}$ over a field $\mathbb{F}$, and a weight $w \in \mathbb{R}^E$, where $r \in \mathbb{Z}_{>0}$.

Goal: Find a parity base $B \subseteq E$ for Z such that $w(B)$ is minimum.

Camerini, Galbiati, and Maffioli [4] first showed that the weighted linear matroid parity problem can be solved in pseudopolynomial time (where the weight is assumed to be an integer vector), and later Cheung et al. [5] devised a faster pseudopolynomial-time algorithm. It was recently announced by Iwata [40] that this problem is solved in strongly polynomial time (estimated as $\mathrm{O}(r^3|E|)$ time). Independently, Pap [75] also announced a strongly polynomial-time algorithm for an equivalent weighted problem, whose running time bound is nontrivial.

Chapter 3

Packing A-paths

In this chapter, we review the packing A-paths problem, which is one of the main problems in this thesis. The basic objective is to find a maximum number of "vertex-disjoint" A-paths in a given undirected graph. Furthermore, we consider several constrained versions of this problem, in which some paths are forbidden to be used.

This chapter is organized as follows. In Section 3.1, we focus on the unconstrained setting, i.e., every A-path can be used. Next in Section 3.2, we consider a constrained version in which a partition $\mathcal{S}$ of the terminal set A is given and only $\mathcal{S}$-paths (A-paths connecting distinct classes of $\mathcal{S}$) can be used. This setting is called *Mader's $\mathcal{S}$-paths problem*, since this problem is in fact equivalent to the openly disjoint A-paths problem, for which Mader [63] showed a min-max duality. Sections 3.3 and 3.4 are devoted to the packing A-paths problems in group-labeled graphs in which only the identity element and all elements in a fixed proper subgroup are forbidden, respectively. Finally in Section 3.5, we mention a further generalization using a family of admissible A-paths with a certain structure, which was first pointed out by Pap [71].

3.1 Unconstrained Problem

For the sake of convenience, we describe the problem formulation again (the first appearance is in Section 1.3.2). The *packing A-paths problem* is, for a given undirected graph $G = (V, E)$ and vertex set $A \subseteq V$ (called a *terminal set*), to find a maximum number of vertex-disjoint A-paths in G. Recall that the term "vertex-disjoint" means that any vertex cannot be shared by any two distinct paths.

Packing A-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint A-paths in G such that $|\mathcal{P}|$ is maximum.

As mentioned in Introduction, Gallai [24] first solved this problem by reducing it to the maximum matching problem. To the contrary, the maximum matching problem is a special case of this problem with $A = V$, and hence these two problems are essentially equivalent. Here, we describe Gallai's reduction and min-max theorem, which extends the Tutte–Berge formula (Theorem 1.1) for maximum matching.

Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set. For each non-terminal $v \in V \setminus A$, add a copy v' so that v' is adjacent to v itself and its neighbors, i.e., all vertices in $N_G(v) + v$. Let $G' = (V', E')$ be the resulting graph. Then, G' has a trivial matching $M = \{\, vv' \in E' \mid v \in V \setminus A \,\}$ of size $|V| - |A|$.

Let M^* be a maximum matching in G'. Then, the symmetric difference $M \triangle M^*$ forms vertex-disjoint A-paths and cycles in G', since all non-terminals in $V' \setminus A$ are covered by M. Note that each of such A-paths in G' corresponds to an A-path in G obtained by identifying each non-terminal and its copy, and vice versa (i.e., an A-path in G intersecting $v_0, v_1, v_2, \ldots, v_{l-1}, v_l \in V$ in this order corresponds to an A-path in G' intersecting $v_0, v_1, v_1', v_2, v_2', \ldots, v_{l-1}, v_{l-1}', v_l \in V'$ in this order). Hence, finding a maximum matching in G' is equivalent to finding a maximum number of vertex-disjoint A-paths in G. In addition, the maximum number of vertex-disjoint A-paths is equal to $|M^*| - |M| = |M^*| - |V| + |A|$.

Theorem 3.1 (Gallai [24]). *Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set. Then, the maximum number of vertex-disjoint A-paths in G is equal to the minimum value of*

$$|U| + \sum_{H \in \mathrm{comp}(G-U)} \left\lfloor \frac{|V(H) \cap A|}{2} \right\rfloor, \tag{3.1}$$

taken over all vertex subsets $U \subseteq V$, where $\mathrm{comp}(G - U)$ *denotes the set of connected components of $G - U$.*

This theorem may not seem to generalize the Tutte–Berge formula (Theorem 1.1), but it does indeed as follows. Recall that the maximum matching problem is a special case with $A = V$. Then, "$\cap A$" in the second term of (3.1) is not necessary. If $|V(H)|$ are even for all $H \in \mathrm{comp}(G - U)$, then the sum is equal to $(|V| - |U|)/2$, and hence the value of (3.1) coincides with (1.1). Otherwise, the value of the sum decreases by $1/2$ per odd component of $G - U$. Thus, we have the equality between (3.1) and (1.1), which means that Theorem 3.1 generalizes Theorem 1.1.

The author gave an alternative proof for Gallai's theorem in [92]. The proof also leads to an Edmonds–Gallai-type structure theorem for this problem, which extends the Edmonds–Gallai structure theorem (Theorem 1.5) and is a special case of such structure for Mader's $\mathcal{S}$-paths shown by Sebő and Szegő [79] (see Section 3.2).

3.2 Mader's $\mathcal{S}$-paths

Mader's $\mathcal{S}$-paths problem is, for a given undirected graph $G = (V, E)$ and terminal set $A \subseteq V$ with its partition $\mathcal{S}$, to find a maximum number of vertex-disjoint $\mathcal{S}$-paths in G.

Mader's $\mathcal{S}$-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$ with its partition $\mathcal{S}$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint $\mathcal{S}$-paths in G such that $|\mathcal{P}|$ is maximum.

The original formulation of Mader's problem is the openly disjoint A-paths problem as follows, which is in fact equivalent to the above problem. Recall that the term "openly disjoint" means that any vertex cannot be shared by two paths as an inner vertex, i.e., in this situation, any non-terminal in $V \setminus A$ is not sharable and each terminal in A can be shared by an arbitrary number of A-paths.

Openly Disjoint A-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of openly disjoint A-paths in G such that $|\mathcal{P}|$ is maximum.

Suppose that we are given an instance of the former problem, i.e., an undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$ partitioned as $\mathcal{S} = \{A_1, A_2, \dots, A_k\}$. For each $i \in [k]$, we add a new vertex a_i so that a_i is adjacent to all terminals in A_i, and let $A' := \{\, a_i \mid i \in [k] \,\}$ be a new terminal set. Then, vertex-disjoint $\mathcal{S}$-paths in G correspond to openly disjoint A'-paths in the resulting graph, and vice versa.

To reduce the latter to the former to the contrary, it suffices to subdivide each edge incident to each terminal $a_i \in A$ (i.e., add a new vertex v_e for each edge $e = ua_i \in \delta_G(a_i)$, remove the edge e, and add two edges uv_e and a_iv_e), let A_i be the set of new vertices adjacent to a_i and $\mathcal{S} := \{\, A_i \mid a_i \in A \,\}$, and remove all original terminals in A. Here, we may assume that there is no edge connecting two terminals in A in the original graph G, since such an edge is always used in a maximum number of openly disjoint A-paths. This operation can be regarded as an inverse operation of the previous one, and hence vertex-disjoint $\mathcal{S}$-paths in the resulting graph correspond to openly disjoint A-paths in G and vice versa.

Mader [63] originally showed a min-max duality for openly disjoint A-paths, and here we describe the $\mathcal{S}$-paths version as follows. Recall that the openly disjoint A-paths problem is the most general among the three settings in Section 1.3.2, and commonly generalizes the maximum matching problem and Menger's disjoint problem as a result. Hence, Mader's theorem commonly extends the Tutte–Berge formula (as well as Gallai's theorem) and Menger's theorems.

Theorem 3.2 (Mader [63]). *Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set with its partition $\mathcal{S}$. Then, the maximum number of vertex-disjoint $\mathcal{S}$-paths in G is equal to the minimum value of*

$$|U| + \sum_{W \in \mathcal{W}} \left\lfloor \frac{|W \cap A_F|}{2} \right\rfloor, \tag{3.2}$$

taken over all vertex subsets $U \subseteq V$ and all partitions $\mathcal{W}$ of $V \setminus U$ such that $G' := G - U - \bigcup_{W \in \mathcal{W}} E(W)$ contains no $\mathcal{S}$-path, where $F := E(G')$ and $A_F := A \cup V(F)$.

Gallai's theorem (Theorem 3.1) and Menger's theorem (Theorem 1.6) follows from Mader's theorem as follows. In both cases, the weak duality (i.e., max $\leq$ min) is clear, and we focus on their minimizers attaining the equality.

The packing A-paths problem can be formulated as Mader's $\mathcal{S}$-paths problem with $\mathcal{S} = \{\, \{a\} \mid a \in A \,\}$ (every A-path is also an $\mathcal{S}$-path). In this case, for each vertex set $U \subseteq V$, a partition $\{\, V(H) \mid H \in \mathrm{comp}(G - U) \,\}$ of $V \setminus U$ satisfies the condition in Theorem 3.2, and makes (3.2) = (3.1), as follows. Since $E(G') = F = \emptyset$, clearly G' in Theorem 3.2 contains no A-path. Also, (3.2) = (3.1) is obvious from $A_F = A$.

Menger's openly disjoint paths problem can be formulated as Mader's $\mathcal{S}$-paths problem by subdividing each edge incident to the terminal s or t, letting A_s and A_t be the set of new vertices adjacent to s and t, respectively, define $A := A_s \cup A_t$ and $\mathcal{S} := \{A_s, A_t\}$, and removing s and t. In this case, for each vertex set $X \subseteq V \setminus \{s, t\}$ separating s and t in the original graph G, the vertex set $U := X$ and a partition $\{\, \{v\} \mid v \in (V \setminus \{s, t\}) \cup A \,\}$ satisfy the condition in Theorem 3.2, attain the minimum of (3.2) under $U = X$, and make (3.2) = $|X|$. Note that each terminal in A has exactly one incident edge, and hence it suffices to consider the case of $U \cap A = \emptyset$ for the minimum value of (3.2).

Sebő and Szegő [79] showed an Edmonds–Gallai-type structure theorem for Mader's $\mathcal{S}$-paths. Preliminary to their structure theorem, we restate Mader's theorem using the concept of $\mathcal{S}$-subpartitions. Let $G = (V, E)$ be an undirected graph, and $A \subseteq V$ a terminal set partitioned as $\mathcal{S} = \{A_1, A_2, \ldots, A_k\}$. A family $\mathcal{X} = (X_0; X_1, X_2, \ldots, X_k)$ of disjoint vertex sets is called an *$\mathcal{S}$-subpartition* if $A_i \subseteq X_0 \cup X_i$ for every $i \in [k]$, and we denote the subgraph $G - X_0 - \bigcup_{i=1}^{k} E(X_i)$ simply by $G - \mathcal{X}$.

Theorem 3.3 (Restatement of Theorem 3.2). *Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set with its partition $\mathcal{S} = \{A_1, A_2, \ldots, A_k\}$. Then, the maximum number of vertex-disjoint $\mathcal{S}$-paths in G is equal to the minimum value of*

$$|X_0| + \sum_{H \in \mathrm{comp}(G - \mathcal{X})} \left\lfloor \frac{V(H) \cap X}{2} \right\rfloor, \tag{3.3}$$

taken over all $\mathcal{S}$-subpartitions $\mathcal{X} = (X_0; X_1, X_2, \ldots, X_k)$, where $X := \bigcup_{i=1}^{k} X_i$

For $i \in [k]$, a vertex $v \in V$ is said to be *i-reachable* if there exists a path from a terminal in $A_i \in \mathcal{S}$ to v in G that is vertex-disjoint from some maximum family of vertex-disjoint $\mathcal{S}$-paths in G, and *i-touched* if v is not i-reachable and some neighbor of v (i.e., some vertex in $N_G(v)$) is i-reachable. If such $i \in [k]$ is unique, then v is said to be *uniquely i-reachable* and *uniquely i-touched*, respectively. Besides, if there exist distinct such $i, j \in [k]$, then *multiply reachable* and *multiply touched*, respectively. We simply say that v is *reachable* or *touched* if it is i-reachable or i-touched, respectively, for some $i \in [k]$.

Theorem 3.4 (Sebő–Szegő [79]). *Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set with its partition $\mathcal{S} = \{A_1, A_2, \dots, A_k\}$, and define*

$$\begin{aligned}
D^* &:= \{\, v \in V \mid v \text{ is multiply reachable} \,\}, \\
X_i^D &:= \{\, v \in V \mid v \text{ is uniquely } i\text{-reachable and } j\text{-touched for some } j \in [k] - i \,\} \quad (i \in [k]), \\
X_i^* &:= \{\, v \in V \mid v \text{ is uniquely } i\text{-reachable and not touched} \,\} \quad (i \in [k]), \\
X_0 &:= \{\, v \in V \mid v \text{ is not reachable and multiply touched} \,\}, \\
X_j^C &:= \{\, v \in V \mid v \text{ is not reachable and uniquely } j\text{-touched for some } j \in [k] \,\}, \\
C^* &:= \{\, v \in V \mid v \text{ is neither reachable nor touched} \,\},
\end{aligned}$$

$$\begin{aligned}
X_i &:= X_i^* \cup X_i^C \cup X_i^D \qquad (i \in [k]), \\
D &:= D^* \cup \bigcup_{i=1}^{k} X_i^D, \qquad C := C^* \cup \bigcup_{i=1}^{k} X_i^C.
\end{aligned}$$

Then, the $\mathcal{S}$-partition $\mathcal{X} := (X_0; X_1, X_2, \dots, X_k)$ attains the minimum of (3.3) *in Theorem* 3.3*, and $G[D]$ and $G[C]$ consists of odd and even connected components of $G - \mathcal{X}$, respectively, where a connected component $H \in \mathrm{comp}(G - \mathcal{X})$ is said to be* odd *or* even *if $|V(H) \cap X|$ is odd or even, respectively.*

Table 3.1: Edmonds–Gallai-type decomposition due to Sebő–Szegő [79].

<table>
<tr><th>touched \ reachable</th><th>not</th><th>uniquely for i</th><th>multiply</th></tr>
<tr><td>not</td><td>C^*</td><td>X_i^*</td><td rowspan="3">D^*</td></tr>
<tr><td>uniquely for j</td><td>X_j^C</td><td rowspan="2">X_i^D</td></tr>
<tr><td>multiply</td><td>X_0</td></tr>
</table>

Lovász [59] showed a reduction of Mader's $\mathcal{S}$-paths problem to the matroid matching problem, and gave an alternative proof for Mader's theorem and the first polynomial-time algorithm through his matroid matching algorithm [61]. Later, Schrijver [78] pointed out that Lovász' reduction admits a linear representation, which leads to much faster algorithms via linear matroid parity algorithms. Their reductions are closely related to our results, and shown in Sections 4.1 and 5.1, respectively.

3.3 Non-zero Model

The *packing non-zero A-paths problem* (we often refer to this problem as the *non-zero model*) was introduced by Chudnovsky, Geelen, Gerards, Goddyn, Lohman, and Seymour [7]. In this problem, we intend to maximize the number of vertex-disjoint non-zero A-paths in a given group-labeled graph, i.e., the identity element of the underlying group is the forbidden label. In what follows, let us fix the underlying group Γ.

Packing Non-zero A-paths Problem

Input: A Γ-labeled graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint non-zero A-paths in G such that $|\mathcal{P}|$ is maximum.

This problem commonly generalizes Mader's $\mathcal{S}$-paths problem and the packing *odd-length* A-paths problem as follows.

To formulate Mader's problem, we use $\mathbb{Z}$-labeled graphs. Suppose that we are given an undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$ partitioned as $\mathcal{S} = \{A_1, A_2, \ldots, A_k\}$. We may assume that there is no edge between two terminals in the same class of $\mathcal{S}$, since such an edge cannot be used in vertex-disjoint $\mathcal{S}$-paths. First, orient each edge between two terminals in A arbitrarily and assign an arbitrary nonzero label in $\mathbb{Z}$ to the resulting arc. Next, orient each edge between a terminal in $A_i \in \mathcal{S}$ and a non-terminal in $V \setminus A$ so that the resulting arc leaves A_i, and assign the label $i \in \mathbb{Z}$ to it. Finally, orient each edge between two non-terminals arbitrarily and assign the label $0 \in \mathbb{Z}$ to the resulting arc. Then, an $\mathcal{S}$-path in G is a non-zero A-path in this $\mathbb{Z}$-labeled graph and vice versa.

To handle the parity constraint, we use $\mathbb{Z}_2$-labeled graphs, where $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z} = (\{0, 1\}, + \pmod 2)$. For a given undirected graph $G = (V, E)$ with a terminal set $A \subseteq V$, orient each edge arbitrarily and assign the label $1 \in \mathbb{Z}_2$ to the resulting arc. Since we have $-1 = 1$ in $\mathbb{Z}_2$, and hence this orientation is not essential. Then, the label of a path in the resulting $\mathbb{Z}_2$-labeled graph is the parity of its length, and hence "non-zero" in such a $\mathbb{Z}_2$-labeled graph is equivalent to "odd-length" in the underlying undirected graph.

For vertex-disjoint non-zero A-paths, Chudnovsky et al. [7] showed the following min-max duality, which generalizes Mader's theorem (Theorem 3.2).

Theorem 3.5 (Chudnovsky–Geelen–Gerards–Goddyn–Lohman–Seymour [7])**.** *Let $G = (V, E)$ be a Γ-labeled graph and $A \subseteq V$ a terminal set. Then, the maximum number of vertex-disjoint non-zero A-paths in G is equal to the minimum value of*

$$|X| + \sum_{H \in \mathrm{comp}(G' - X - E_0)} \left\lfloor \frac{|V(H) \cap A'|}{2} \right\rfloor, \tag{3.4}$$

taken over all A-equivalent Γ-labeled graphs G', all vertex subsets $X \subseteq V$, and all A' with $A \subseteq A' \subseteq V$, where $E_0 := \{\, e = uv \in E \mid u, v \in A',\ \psi_{G'}(e) = 1_\Gamma \,\}$.

Chudnovsky et al. [6] proposed a combinatorial algorithm for packing non-zero A-paths, which runs in polynomial time. They also showed an Edmonds–Gallai-type structure at the same time, and utilize it to determine when their algorithm should halt.

Theorem 3.6 (Chudnovsky–Cunningham–Geelen [6])**.** *One can solve the packing non-zero A-paths problem in* $\mathrm{O}(|V|^5)$ *time by a deterministic algorithm.*

To state their structure theorem, we define several necessary concepts, which are closely related to those introduced for the structure theorem of Sebő and Szegő (Theorem 3.4). For an element $\alpha \in \Gamma$, a vertex $v \in V$ is said to be *α-reachable* if there exists a path of label α from a terminal in A to v in G that is vertex-disjoint from some maximum family of vertex-disjoint non-zero A-paths in G. If such $\alpha \in \Gamma$ is unique, then v is said to be *uniquely α-reachable*, and *multiply reachable* otherwise (i.e., if there exist distinct such $\alpha, \beta \in \Gamma$). Besides, v is said to be *α-touched* if v is not reachable and there exists an arc $e = uv \in E$ and an element $\beta \in \Gamma$ such that $u \in V$ is β-reachable and $\alpha = \psi_G(e) \cdot \beta$. We simply say that v is *reachable* or *touched* if v is α-reachable or α-touched, respectively, for some $\alpha \in \Gamma$.

Theorem 3.7 (Chudnovsky–Cunningham–Geelen [6])**.** *Let $G = (V, E)$ be a Γ-labeled graph and $A \subseteq V$ a terminal set. Suppose that every uniquely reachable vertex in G is 1_Γ-reachable, and every touched non-terminal in G is 1_Γ-touched* (*since there exists an A-equivalent Γ-labeled graph, this assumption is not essential*)*. Define*

$$\begin{aligned} D &:= \{\, v \in V \mid v \text{ is reachable} \,\}, \\ D_1 &:= \{\, v \in V \mid v \text{ is uniquely reachable} \,\}, \\ A' &:= A \cup N_G(D) \cup D_1, \\ E_0 &:= \{\, e = uv \in E \mid u, v \in A',\ \psi_G(e) = 1_\Gamma \,\}, \\ X &:= N_{G - E_0}(D). \end{aligned}$$

Then, the vertex subsets X and A' with $G' := G$ attains the minimum of (3.4) *in Theorem* 3.5*, and $G[D] - E_0$ coincides with the union of odd connected components of $G - X - E_0$ with respect to the terminal set A'.*

As a natural question, one may be concerned with "openly disjoint" non-zero A-paths in Γ-labeled graphs. The following problem is the simplest special case of this setting with $\Gamma \simeq \mathbb{Z}_2$ and $|A| = 2$, even which is NP-hard.

Openly Disjoint Odd s–t Paths Problem

Input: An undirected graph $G = (V, E)$ and distinct vertices $s, t \in V$.

Goal: Find a family $\mathcal{P}$ of openly disjoint odd s–t paths such that $|\mathcal{P}|$ is maximum.

For a variant problem in which we intend to maximize the number of edge-disjoint odd s–t trails, Churchley, Mohar, and Wu [8] showed a weak duality and an approximation algorithm. They also mentioned that, for disjoint odd s–t paths, the minimum size of a hitting set can be arbitrarily large (cf. [8, Fig. 1]). We show that this problem is NP-hard as suspected in [7].

Theorem 3.8 (Personal communication with Kawarabayashi and Kobayashi [47])**.** *The openly disjoint odd s–t paths problem is NP-hard.*

Proof. The following problem was shown to be NP-hard by Even, Itai, and Shamir [15].

Simple 2-commodity Integral Flow in Undirected Graphs

Input: An undirected graph $G = (V, E)$, four vertices $s_1, t_1, s_2, t_2 \in V$ (possibly $s_1 = s_2$ or $t_1 = t_2$), and positive integers R_1 and R_2.

Goal: Find a family $\mathcal{P}$ of edge-disjoint R_i s_i–t_i paths ($i = 1, 2$).

First, we convert this problem to an openly disjoint version (cf. the reduction of Menger's edge-disjoint paths problem to the openly disjoint setting shown in Section 1.3.1). Suppose that we are given an instance of the simple 2-commodity integral flow problem. We add four distinct vertices s'_1, t'_1, s'_2, t'_2 as copies of the original terminals s_1, t_1, s_2, t_2, and connect each copy s'_i or t'_i with each original s_i or t_i, respectively, by R_i parallel edges. In addition, we add extra four distinct vertices $s''_1, t''_1, s''_2, t''_2$, and connect them with the copies s'_1, t'_1, s'_2, t'_2 by single edges e_1, f_1, e_2, f_2, respectively.

Let $\hat{G}$ be the line graph of the resulting graph. Then, G contains edge-disjoint R_i s_i–t_i paths ($i = 1, 2$) if and only if $\hat{G}$ contains openly disjoint R_i e_i–f_i paths ($i = 1, 2$). Moreover, it is easy to construct the former from the latter by transforming its vertex set to an edge set in G and scraping redundant edges if necessary. Since there are exactly R_i edges around e_i and f_i, respectively, the following problem generalizes the current problem (finding openly disjoint R_i e_i–f_i paths in $\hat{G}$), and hence it is also NP-hard.

Complete 2-disjoint Paths in Undirected Graphs

Input: An undirected graph $G = (V, E)$ and four distinct vertices $s_i, t_i \in V$ ($i = 1, 2$) with $|\delta_G(s_i)| = |\delta_G(t_i)|$.

Goal: Find a family $\mathcal{P}$ of openly disjoint s_i–t_i paths ($i = 1, 2$) with $|\mathcal{P}| = \sum_{i=1}^{2} |\delta_G(s_i)|$.

We reduce this problem to the openly disjoint odd s–t paths problem as follows. Suppose that we are given an instance of the complete 2-disjoint paths problem. First, we subdivide each edge $e = uw \in E$, i.e., remove the edge e and add a new vertex v_e with two new edges uv_e and wv_e. Next, we create $|\delta_G(s_2)|$ copies of s_2 and $|\delta_G(t_1)|$ copies of t_1 so that each copy is adjacent to exactly one of the original neighbors and vice versa, and remove original s_2 and t_1. Finally, we connect $s := s_1$ with each copy of s_2 and $t := t_2$

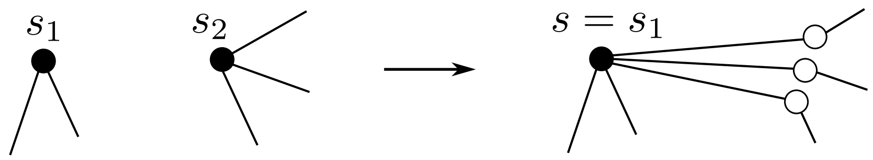

Figure 3.1: Splitting s_2 and connecting them with $s = s_1$.

with each copy of t_1 by a single edge, respectively (see Fig. 3.1). Then, each s_i–t_i paths ($i \in \{1, 2\}$) in G corresponds to an odd s–t path in the resulting graph G' (and vice versa), which contains at most $|\delta_{G'}(s)| = |\delta_G(s_1)| + |\delta_G(s_2)|$ openly disjoint s–t paths. □

Note that, because of the parity constraint, finding a maximum number of edge-disjoint odd s–t paths in undirected graphs is not reduced straightforwardly to the openly disjoint setting (scraping redundant edges may change the parity of length in reconstructing s–t paths). It is however derived by the same technique as Fig. 3.1 directly from the simple 2-commodity integral flow problem that the edge-disjoint version is also NP-hard.

3.4 Subgroup-Forbidden Model

The *packing non-returning A-paths problem* (we often refer to this problem as the *non-returning model*) was introduced by Pap [71, 72]. In this problem, the group Γ is the symmetric group S_d of degree d for some integer $d \geq 2$, i.e., each element of Γ is a permutation on the set $[d]$. An A-path P from $u \in A$ to $v \in A$ is said to be *non-returning* with respect to a map $\omega\colon A \to [d]$ if $\omega(v) \neq \psi_G(P)(\omega(u))$. The objective is to find a maximum number of vertex-disjoint non-returning A-paths in a given S_d-labeled graph with respect to a given map $\omega\colon A \to [d]$.

Packing Non-returning A-paths Problem

Input: An S_d-labeled graph $G = (V, E)$, a terminal set $A \subseteq V$, and a map $\omega\colon A \to [d]$, where $d \in \mathbb{Z}_{\geq 2}$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint non-returning A-paths in G with $|\mathcal{P}|$ maximum.

Pap [71] showed the equivalence between the non-returning model and the following model, which we refer to as the *subgroup-forbidden model*. In this model, we are given a proper subgroup Γ' of Γ and we can use only A-paths whose labels are NOT in Γ'. We say that an A-path is Γ'-*admissible* if its label is not in Γ'.

Subgroup-Forbidden Model

Input: A Γ-labeled graph G, a terminal set $A \subseteq V(G)$, and a proper subgroup Γ' of Γ.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint Γ'-admissible A-paths in G with $|\mathcal{P}|$ maximum.

It is easy to formulate the non-returning model as the subgroup-forbidden model as follows. Without loss of generality, we may assume that a map $\omega\colon A \to [d]$ satisfies $\omega(a) = d$ for every terminal $a \in A$ (otherwise, for each terminal $a \in A$ with $\omega(a) \neq d$, add a new vertex a' and an arc aa' with label $(\omega(a)\ \ d) \in S_d$, and update $A := A - a + a'$ and $\omega(a') := d$). Then, an A-path is non-returning if and only if its label permutation does not fix d, and hence the set of forbidden labels is the proper subgroup S_{d-1} of S_d.

The converse direction is not so easy. Pap [71] proved it via another model shown in the next section (Theorem 3.13).

Pap [72] gave a min-max duality for vertex-disjoint non-returning A-paths, which generalizes Theorem 3.5, and proved it simply inspired by Schrijver's proof for Mader's theorem [77]. Here, we state his theorem in the form for the subgroup-forbidden model.

Theorem 3.9 (Pap [72])**.** *Let $G = (V, E)$ be a Γ-labeled graph, $A \subseteq V$ a terminal set, and Γ' a proper subgroup of Γ. Then, the maximum number of vertex-disjoint Γ'-admissible A-paths in G is equal to the minimum value of the maximum number of vertex-disjoint A_F-paths in $G - F$ taken over all edge subsets $F \subseteq E$ such that $G[F]$ contains no Γ'-admissible A-path, where $A_F := A \cup V(F)$.*

The weak duality (max $\leq$ min) is easy to see as follows. Fix an arbitrary edge subset $F \subseteq E$ such that $G[F]$ is balanced and contains no Γ'-admissible A-path. Then, for each Γ'-admissible A-path P in G, there exists an A_F-path P' in $G - F$ with $E(P') \subseteq E(P)$, since F cannot include $E(P)$ itself. Hence, for each family of vertex-disjoint Γ'-admissible A-paths in G, there exists a same-size family of vertex-disjoint A_F-paths in $G - F$.

By using Gallai's theorem for vertex-disjoint A-paths (Theorem 3.1), the maximum number of vertex-disjoint A_F-paths in $G - F$ is equal to the minimum value of

$$|X| + \sum_{H \in \mathrm{comp}(G-X-F)} \left\lfloor \frac{|V(H) \cap A_F|}{2} \right\rfloor, \tag{3.5}$$

taken over all vertex subsets $X \subseteq V$. This form implies the min-max duality for vertex-disjoint non-zero A-paths due to Chudnovsky et al. (Theorem 3.5) with the aid of Proposition 2.6 as follows.

For each A-balanced edge set $F \subseteq E$ (i.e., $G[F]$ is balanced and contains no non-zero A-path), $G[F]$ is A-equivalent to a trivially-labeled Γ-labeled graph. Then, by shifting G to be G' so that $G'[F]$ is trivially labeled and taking $A' := A_F$, we have (3.4) $\leq$ (3.5). To the contrary, for each A' with $A \subseteq A' \subseteq V$, the edge set $E_0 = \{\, e = uv \in E \mid u, v \in A',\ \psi_{G'}(e) = 1_\Gamma \,\}$ forms no non-zero A-path. Hence, by taking $F := E_0$, we have (3.4) $\geq$ (3.5) (note that $A_F \subseteq A'$).

In [94], the author extended the results of Chudnovsky et al. [6] to the subgroup-forbidden model: a polynomial-time algorithm and an Edmonds–Gallai-type structure theorem. The extension itself is not surprising but reasonable. However, it involves a slightly change of the computational time bound because of the following reason.

Suppose that there exist parallel arcs e_1, e_2 with distinct labels between the same pair of two vertices. Then, for every A-path P_1 traversing e_1, at least one of P_1 itself and the A-path P_2 obtained from P_1 by replacing e_1 with e_2 is non-zero. Hence, when we consider non-zero A-paths in a group-labeled graph, we may assume that there are at most two arcs between the same pair of two vertices. This assumption bounds the number of arcs in the input graph by $\mathrm{O}(|V|^2)$.

In the subgroup-forbidden model, however, $\Omega(|\Gamma'|)$ parallel arcs between a pair of two vertices can be necessary. This implies that we cannot bound the computational time only by the number of vertices, and the author showed the following bound.

Theorem 3.10 (Yamaguchi [94, Section 4.5])**.** *One can solve the subgroup-forbidden model in* $\mathrm{O}(|V|^5 + |E| \cdot |V|)$ *time by a deterministic algorithm.*

3.5 Axiomatic Model

Here we introduce the most general setting of packing A-paths problem as the *axiomatic model.* In this model, we are given a family of A-paths as the family of *admissible* A-paths as follows. We say that a family $\mathcal{F}$ of A-paths is *symmetric* if $\bar{P} \in \mathcal{F}$ for every $P \in \mathcal{F}$.

Axiomatic Model

Input: An undirected graph G, a terminal set $A \subseteq V(G)$, and a symmetric family $\mathcal{F}$ of A-paths in G.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint A-paths in $\mathcal{F}$ such that $|\mathcal{P}|$ is maximum.

For example, Mader's $\mathcal{S}$-paths problem is a special case such that $\mathcal{F}$ is the set of $\mathcal{S}$-paths for a given partition $\mathcal{S}$ of A. In the subgroup-forbidden model, $\mathcal{F}$ is the set of Γ'-admissible A-paths for a given proper subgroup Γ' of Γ. Pap [71] introduced the concept of *triple exchange* by extracting a property enjoyed by these tractable families $\mathcal{F}$.

Definition 3.11 (Weak Triple Exchange)**.** Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set. A symmetric family $\mathcal{F}$ of A-paths in G is *weakly triple exchangeable* if it satisfies the following condition: for every A-path $P \in \mathcal{F}$, inner vertex $v \in V(P) \setminus A$, terminal $a \in A \setminus V(P)$, and a–v path Q in G openly disjoint from P, at least one of the two A-paths obtained by extending Q along P is in $\mathcal{F}$.

Definition 3.12 (Strong Triple Exchange)**.** Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set. A symmetric family of A-paths in G is *strongly triple exchangeable* if it coincides with the set of A-paths contained in some symmetric family $\mathcal{F}$ of A-walks with the following condition: for every A-walk $P \in \mathcal{F}$, vertex $v \in V(P)$, terminal $a \in A$, and a–v walk Q in G, at least one of the two A-walks obtained by extending Q along P is in $\mathcal{F}$.

By the above definitions, a strongly triple exchangeable family is weakly triple exchangeable, since P, v, a, Q in Definition 3.11 can be taken as those in Definition 3.12, respectively. The difference between the two property is rather large, since, while the weak triple exchange property is defined in a completely finite way (the number of A-paths in a finite graph is finite), the strong triple exchange property requires an infinite combinatorial structure (the number of A-walks can be infinite even in a finite graph).

Pap [71] pointed out the following two properties.

Theorem 3.13 (Pap [71, Claims 3.21–3.23])**.** *The non-returning model, the subgroup-forbidden model, and the axiomatic model with the strong triple exchange property are equivalence.*

Theorem 3.14 (Pap [71, Theorem 3.24] (cf. matching delta-matroids in Section 2.3.2))**.** *In the axiomatic model with the weak triple exchange property, the family of terminal subsets each of which is exactly covered by some vertex-disjoint admissible A-paths is the feasible family of a delta-matroid on the terminal set A.*

Theorems 3.9 and 3.13 imply that the axiomatic model with the strong triple exchange property admits the following good characterization.

Corollary 3.15. *Let $G = (V, E)$ be an undirected graph, $A \subseteq V$ a terminal set, and $\mathcal{F}$ a symmetric family of A-paths in G with the strong triple exchange property. Then, the maximum number of vertex-disjoint A-paths in $\mathcal{F}$ is equal to the minimum value of the maximum number of vertex-disjoint A_F-paths in $G - F$ taken over all edge subsets $F \subseteq E$ such that $G[F]$ contains no A-path in $\mathcal{F}$, where $A_F := A \cup V(F)$.*

Corollary 3.15 raises a natural question: does the axiomatic model with the "weak" triple exchange property admits the same good characterization? It is still open now, and we have two progresses related to this question, which imply the border of the tractability of packing A-paths.

One is a reduction to the matroid matching problem, which is given in Section 4.5. While a reduction of the non-zero model leads to an alternative proof for the min-max duality (Theorem 3.5) as shown in Section 4.4, the reduction of the axiomatic model with weak triple exchange does not necessarily give a good characterization.

The other is a simple counterexample indicating that the axiomatic model without weak triple exchange does not admit the same good characterization. Define an undirected graph $G = (V, E)$, a terminal set $A \subseteq V$, and a symmetric family $\mathcal{F}$ of A-paths in G as follows (see Fig. 3.2):

$$\begin{aligned} V &:= \{ a_1, a_2, a_3, a_4, v \}, \\ E &:= \{ e_{12} = a_1a_2,\ e_{13} = a_1a_3,\ e_2 = va_2,\ e_3 = va_3,\ e_4 = va_4 \}, \\ A &:= V - v \ = \ \{ a_1, a_2, a_3, a_4 \}, \end{aligned}$$

and let $\mathcal{F}$ be the set of A-paths disjoint from a_4. Then, every A-path in $\mathcal{F}$ is $(A-a_4)$-path, and hence the maximum number of vertex-disjoint A-paths in $\mathcal{F}$ is 1. On the other hand, since the dual variable $F \subseteq E$ cannot contain any A-path in $\mathcal{F}$ (i.e., any $(A-a_4)$-path in G), it is one of the followings: $\emptyset$, $\{e_2\}$, $\{e_3\}$, $\{e_4\}$, $\{e_2, e_4\}$, and $\{e_3, e_4\}$. It is easily checked that, whenever F is one of the above six edge subsets, $G - F$ contains two vertex-disjoint A_F-paths, where $A_F := A \cup V(F)$.

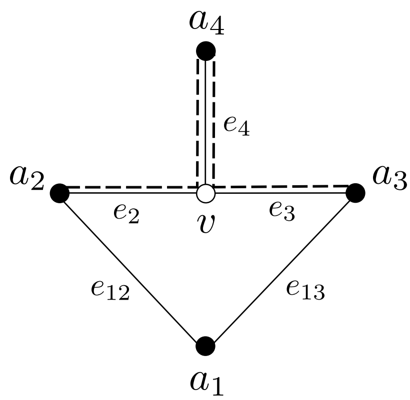

Figure 3.2: Counterexample for the duality, where the dashed A-paths are not in $\mathcal{F}$.

Chapter 4

Packing A-paths via Matroid Matching

In this chapter, we discuss reductions of the packing A-paths problems to the matroid matching problem. Lovász [59] showed a reduction of Mader's $\mathcal{S}$-paths problem to the matroid matching problem, which leads to an alternative proof for Mader's min-max theorem (Theorem 3.2) and the first polynomial-time algorithm. We extend his work to packing A-paths in group-labeled graphs, and discuss further extendability to the axiomatic model with the weak triple exchange property.

This chapter is based on [83] and organized as follows. In Section 4.1, we briefly review Lovász' reduction [59, Section 3] of Mader's problem. In Section 4.2, we introduce an extension of the frame matroids of group-labeled graphs, which is a key concept in our reduction. Section 4.3 is devoted to showing a reduction of the packing non-zero A-paths problem to the matroid matching problem, which extends Lovász' reduction of Mader's problem. In Section 4.4, based on our reduction, we give alternative proofs for the min-max formula due to Chudnovsky et al. [7] and the polynomial-time solvability (cf. [6]), with the aid of Lovász' theory on matroid matching. Finally, in Section 4.5, we mention a possible extension of our reduction to the axiomatic model with the weak triple exchange property, which does not necessarily lead to a good chacracterization or a polynomial-time algorithm.

4.1 Lovász' Reduction of Mader's $\mathcal{S}$-paths

For the sake of convenience, we restate the two problems here (the first appearances are in Sections 3.2 and 2.4.1).

Mader's $\mathcal{S}$-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$ with its partition $\mathcal{S}$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint $\mathcal{S}$-paths in G such that $|\mathcal{P}|$ is maximum.

Matroid Matching Problem

Input: A 2-polymatroid (E, f).

Goal: Find a matching $M \subseteq E$ in (E, f) such that $|M|$ is maximum.

For an instance of Mader's $\mathcal{S}$-paths problem, define a set function $f_{G,\mathcal{S}} \colon 2^E \to \mathbb{Z}_{\geq 0}$ by

$$f_{G,\mathcal{S}}(F) := \sum_{F' \in \mathrm{comp}(F)} \big(2|V(F')| - 2 + \rho_{G,\mathcal{S}}(F') - |V(F') \cap A|\big) \qquad (F \subseteq E), \tag{4.1}$$

where recall that $\mathrm{comp}(F)$ denotes the partition of F according to the connected components, and $\rho_{G,\mathcal{S}} \colon 2^E \to \mathbb{Z}_{\geq 0}$ is defined as

$$\rho_{G,\mathcal{S}}(F') := \begin{cases} 0 & (V(F') \cap A = \emptyset), \\ 1 & (V(F') \cap A_i \neq \emptyset \text{ for exactly one } A_i \in \mathcal{S}), \\ 2 & (\text{otherwise}). \end{cases} \tag{4.2}$$

Then, $(E, f_{G,\mathcal{S}})$ is a 2-polymatroid.

The following theorem characterizes the matchings in $(E, f_{G,\mathcal{S}})$, which also leads to a connection between those and vertex-disjoint $\mathcal{S}$-paths in G.

Theorem 4.1 (Lovász [59, Lemma 3.3])**.** *An edge subset $F \subseteq E$ is a matching in $(E, f_{G,\mathcal{S}})$ if and only if $G[F]$ is a forest such that each connected component contains at most one A-path, which is an $\mathcal{S}$-path.*

Lovász [59] enumerated all nontrivial double circuits, and coped with each case separately. Based on the analysis, he gave an alternative proof for Mader's theorem (Theorem 3.2) and showed the polynomial-time solvability of Mader's $\mathcal{S}$-paths problem using Theorem 2.16 and his matroid matching algorithm [61]. The details are shown in the following sections throughout our extended reduction.

4.2 Extension of Frame Matroids

Let $G = (V, E)$ be a Γ-labeled graph. Recall that the rank function $r_G \colon 2^E \to \mathbb{Z}_{\geq 0}$ of the frame matroid of G (cf. Section 2.3.2) is written by

$$r_G(F) = \sum_{F' \in \mathrm{comp}(F)} \big(|V(F')| - 1 + \rho_G(F')\big) \qquad (F \subseteq E),$$

where $\rho_G \colon 2^E \to \mathbb{Z}$ is defined as

$$\rho_G(F') := \begin{cases} 1 & (G[F'] \text{ is unbalanced}), \\ 0 & (\text{otherwise}). \end{cases}$$

Tanigawa [82] extended the construction of the union of frame matroids by using structures of the underlying group Γ. The idea is to replace the term ρ_G with a function taking fractional values. For a nonempty subset $X \subseteq \Gamma$, we denote by $\langle X \rangle$ the subgroup of Γ generated by X.

Definition 4.2. A set function $\rho\colon 2^\Gamma \to \mathbb{R}_{\geq 0}$ is called a *symmetric polymatroidal function over* Γ if ρ satisfies the following conditions:

(SP0) $\rho(\emptyset) = 0$,

(SP1) $X \subseteq Y \subseteq \Gamma \implies \rho(X) \leq \rho(Y)$,

(SP2) $\rho(X) + \rho(Y) \geq \rho(X \cup Y) + \rho(X \cap Y)$ for every $X, Y \subseteq \Gamma$,

(SP3) $\rho(\langle X \rangle) = \rho(X)$ for every nonempty subset $X \subseteq \Gamma$, and

(SP4) $\rho(\alpha X \alpha^{-1}) = \rho(X)$ for every nonempty subset $X \subseteq \Gamma$ and every element $\alpha \in \Gamma$.

Let $\rho\colon 2^\Gamma \to \mathbb{R}_{\geq 0}$ be a symmetric polymatroidal function, $F \subseteq E$ a connected edge set (i.e., $\mathrm{comp}(F) = \{F\}$), and $r \in V(F)$ a vertex. We denote by $\langle F \rangle_r$ the set of possible labels of closed r-walk in $G[F]$, which is a subgroup of Γ by Proposition 2.10. Note that, for any $r' \in V(F)$, there exists an element $\alpha \in \Gamma$ such that $\langle F \rangle_{r'} = \alpha \cdot \langle F \rangle_r \cdot \alpha^{-1}$, since we can construct a closed r'-walk in $G[F]$ by extending any closed r-walk in $G[F]$ using an r'–r path in $G[F]$ and vice versa. Hence, the value of $\rho(\langle F \rangle_r)$ does not depend on the choice of the vertex $r \in V(F)$ by Condition (SP4), and we simply denote by $\rho\langle F \rangle$ the value. We define a set function $g_\rho\colon 2^E \to \mathbb{R}_{\geq 0}$ by

$$g_\rho(F) := \sum_{F' \in \mathrm{comp}(F)} \bigl(|V(F')| - 1 + \rho\langle F' \rangle\bigr) \qquad (F \subseteq E). \tag{4.3}$$

Theorem 4.3 (Tanigawa [82])**.** *Let $\rho\colon 2^\Gamma \to [0, 1]$ be a symmetric polymatroidal function over a group Γ, and $G = (V, E)$ a Γ-labeled graph. Then, the set function $g_\rho\colon 2^E \to \mathbb{R}_{\geq 0}$ defined as* (4.3) *is monotone and submodular.*

Suppose that ρ takes fractional values, i.e., $\rho\colon 2^\Gamma \to \{0, \frac{1}{d}, \frac{2}{d}, \dots, \frac{d-1}{d}, 1\}$ for some positive integer $d \in \mathbb{Z}_{>0}$. Then, if we define a set function $f_\rho\colon 2^E \to \mathbb{Z}_{\geq 0}$ by

$$f_\rho(F) := d \cdot g_\rho(F) \qquad (F \subseteq E),$$

(E, f_ρ) is a d-polymatroid (cf. Definition 2.15).

Now we construct a symmetric polymatroidal function of rank 2. For nontrivial groups Γ_1 and Γ_2, the free product $\Gamma_1 * \Gamma_2$ is the group consisting of all words $\gamma_1\gamma_2\cdots\gamma_m$ of arbitrary finite length $m \in \mathbb{Z}_{\geq 0}$, where each letter γ_i is a nonidentity element of Γ_1 or Γ_2 and adjacent letters γ_i and γ_{i+1} belong to different groups. The identity element of $\Gamma_1 * \Gamma_2$ is defined to be the empty word. See, e.g., [32] for more details on the free product.

Lemma 4.4. *Let Γ_1 and Γ_2 be disjoint nontrivial groups, and Γ the free product of Γ_1 and Γ_2. Suppose that a set function $\rho\colon 2^\Gamma \to \mathbb{Z}$ is defined by*

$$\rho(X) := \begin{cases} 0 & (X \text{ is trivial, i.e., } X = \emptyset \text{ or } X = \{1_\Gamma\}), \\ 1 & (X \text{ is nontrivial and } X \subseteq \gamma\Gamma_i\gamma^{-1} \text{ for some } i \in \{1,2\} \text{ and some } \gamma \in \Gamma), \\ 2 & (\text{otherwise}). \end{cases}$$

Then, ρ is symmetric polymatroidal over Γ.

Proof. Clearly, ρ satisfies Conditions (SP0) and (SP1) in Definition 4.2. Also, it satisfies (SP3) and (SP4), since $\gamma\Gamma_i\gamma^{-1}$ is a subgroup of Γ.

Let $\mathcal{G} := \{\, \gamma\Gamma_i\gamma^{-1} \mid i \in \{1,2\},\ \gamma \in \Gamma \,\}$, and $X, Y \in \mathcal{G}$. We then have

$$X \neq Y \implies X \cap Y = \{1_\Gamma\}. \tag{4.4}$$

To see this, let $X = \gamma_X\Gamma_i\gamma_X^{-1}$ and $Y = \gamma_Y\Gamma_j\gamma_Y^{-1}$. The case of $i \neq j$ is obvious, since the middle element of any nonempty word in $X \cap Y$ must be in $\Gamma_1 \cap \Gamma_2 = \emptyset$. Otherwise, $i = j$ and $\gamma_X \neq \gamma_Y$. This case is also clear.

In what follows, we show the submodularity (SP2), i.e., $\rho(X) + \rho(Y) \geq \rho(X \cap Y) + \rho(X \cup Y)$ for every $X, Y \subseteq \Gamma$. If $\rho(Y) = 0$ (or symmetrically $\rho(X) = 0$), then we have $\rho(X \cap Y) = \rho(Y)$ and $\rho(X) = \rho(X \cup Y)$. Thus, (SP2) holds with the equality.

If $\rho(Y) = 2$, then we have $\rho(Y) = \rho(X \cup Y)$ and $\rho(X) \geq \rho(X \cap Y)$, since ρ is monotone and the range of the value of ρ is $\{0,1,2\}$. This leads to (SP2).

Finally, suppose that $\rho(X) = 1$ and $\rho(Y) = 1$. Note that there are unique $\Gamma_X \in \mathcal{G}$ with $X \subseteq \Gamma_X$ and unique $\Gamma_Y \in \mathcal{G}$ with $Y \subseteq \Gamma_Y$. By (4.4), if $\Gamma_X \neq \Gamma_Y$, then $\rho(X \cup Y) = 2$ and $\rho(X \cap Y) = 0$; otherwise, $\rho(X \cup Y) = 1$ and $\rho(X \cap Y) \leq 1$. Thus we have done. $\square$

4.3 Reduction of Packing Non-zero A-paths

With the aid of the extension of frame matroids, we show a reduction of packing non-zero A-paths to matroid matching, which generalizes Lovász' reduction of Mader's problem reviewed in Section 4.1. We first restate the problem (the first appearance is in Section 3.3).

Packing Non-zero A-paths Problem

Input: A Γ-labeled graph $G = (V, E)$ and a terminal set $A \subseteq V$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint non-zero A-paths in G such that $|\mathcal{P}|$ is maximum.

Let Γ' be a group consisting of two elements $1_{\Gamma'}$ and $\bullet$ (then, Γ' is isomorphic to $\mathbb{Z}/2\mathbb{Z}$, i.e., $\bullet \neq 1_{\Gamma'}$ and $\bullet^2 = 1_{\Gamma'}$). We denote by $\Gamma^\bullet$ the free product of Γ and Γ', and define a function $\rho\colon 2^{\Gamma^\bullet} \to \mathbb{Z}_{\geq 0}$ by

$$\rho(X) := \begin{cases} 0 & (X \text{ is trivial}),\\ 1 & (X \text{ is nontrivial, and } X \subseteq \gamma\Gamma\gamma^{-1} \text{ or } X = \{1_{\Gamma^\bullet}, \gamma \bullet \gamma^{-1}\} \text{ for some } \gamma \in \Gamma^\bullet),\\ 2 & (\text{otherwise}),\end{cases}$$

for each $X \subseteq \Gamma^\bullet$. By Lemma 4.4, this ρ is a symmetric polymatroidal function over $\Gamma^\bullet$.

Let $\tilde{G} = (V, \tilde{E})$ be a $\Gamma^\bullet$-labeled graph obtained from G by attaching a new self-loop $\ell_{v,\gamma}$ at each vertex $v \in V$ with label $\psi_{\tilde{G}}(\ell_{v,\gamma}) = \gamma \bullet \gamma^{-1}$ for each element $\gamma \in \Gamma$. Let L be the set of new self-loops and, for each vertex set $U \subseteq V$, define $L_U := \{\, \ell_{v,1_\Gamma} \in L \mid v \in U \,\}$. By Theorem 4.3, the function $f_G\colon 2^{\tilde{E}} \to \mathbb{Z}_{\geq 0}$ defined by

$$f_G(F) := \sum_{F' \in \mathrm{comp}(F)} \bigl(2|V(F')| - 2 + \rho\langle F'\rangle\bigr) \qquad (F \subseteq \tilde{E})$$

is monotone and submodular, and thus $(\tilde{E}, f_G)$ is a 2-polymatroid.

We consider the contraction $(\tilde{E} \setminus L_A,\ f_{G,A})$ of $(\tilde{E}, f_G)$ by L_A, and its restriction to E. The rank function is characterized as follows. Note that Lemma 4.5 implies that our 2-polymatroids extend Lovász' 2-polymatroids defined by (4.1) and (4.2), by considering the formulation of Mader's $\mathcal{S}$-paths problem as packing non-zero A-paths shown in Section 3.3.

Lemma 4.5. *For a subset $F \subseteq E$,*

$$f_{G,A}(F) = \sum_{F' \in \mathrm{comp}(F)} \bigl(2|V(F')| - 2 + \rho_{G,A}(F') - |V(F') \cap A|\bigr),$$

where $\rho_{G,A}\colon 2^E \to \mathbb{Z}$ is written by

$$\rho_{G,A}(F') = \begin{cases} 2 & (F' \text{ is not } A\text{-balanced, and } |V(F') \cap A| \geq 1),\\ 1 & \begin{array}{l}(F' \text{ is } A\text{-balanced and } |V(F') \cap A| \geq 1, \text{ or}\\ \ F' \text{ is unbalanced and } |V(F') \cap A| = 0),\end{array}\\ 0 & (\text{otherwise}).\end{cases}$$

Proof. For an edge set $F \subseteq E$, define $L_A(F) := L_{V(F) \cap A}$. Let us first check

$$\rho_{G,A}(F) = \rho\langle F \cup L_A(F)\rangle$$

for any connected edge set $F \subseteq E$.

Suppose that $V(F) \cap A = \emptyset$. We then have $\langle F \cup L_A(F)\rangle_v = \langle F\rangle_v$ for every vertex $v \in V(F)$, since $L_A(F) = \emptyset$. If F is balanced, then $\langle F\rangle_v$ is trivial, which implies $\rho\langle F \cup L_A(F)\rangle = 0 = \rho_{G,A}(F)$. Otherwise (i.e., if F is unbalanced), $\langle F\rangle$ is a nontrivial subgroup of Γ, and hence $\rho\langle F \cup L_A(F)\rangle = 1 = \rho_{G,A}(F)$.

Suppose next that $|V(F) \cap A| \geq 1$. If F is A-balanced, then we can see $\rho\langle F \cup L_A(F)\rangle =$

$1 = \rho_{G,A}(F)$ as follows. Take a terminal $a \in V(F) \cap A$. Then, the label of every path in $G[F]$ from a to any other terminal in $V(F) \cap A$ is the identity as F is A-balanced, and hence $\langle F \cup L_A(F) \rangle_a = \{1_{\Gamma^\bullet}, \bullet\}$.

Otherwise, $G[F]$ contains a non-zero A-path or an unbalanced cycle. In the former case, take an end vertex a of a non-zero A-path in $G[F]$. Then, $\langle F \cup L_A(F) \rangle_a$ contains $\bullet$ and $\gamma^{-1} \bullet \gamma$, where γ is the label of the non-zero A-path. Since γ is not the identity, $\rho\langle F \cup L_A(F) \rangle = 2 = \rho_{G,A}(F)$. In the latter case (i.e., when F is unbalanced with $|V(F) \cap A| \geq 1$), we can similarly see $\rho\langle F \cup L_A(F) \rangle = 2 = \rho_{G,A}(F)$ as follows. Take a terminal $a \in V(F) \cap A$. Then, $\langle F \cup L_A(F) \rangle_a$ contains $\bullet$ and γ, where γ is the nonzero label of a closed a-walk in $G[F]$. This implies $\rho\langle F \cup L_A(F) \rangle = 2$ by the definition of ρ.

Note that, for any vertex set $U \subseteq V$, we have $f(L_U) = |L_U|$. Since $f_{G,A}$ is defined as the contraction of f_G by L_A, we complete the proof as follows: for any edge set $F \subseteq E$,

$$\begin{aligned} f_{G,A}(F) &= f_G(F \cup L_A) - f_G(L_A) \\ &= f_G(F \cup L_A(F)) - f_G(L_A(F)) \\ &= \sum_{F' \in \mathrm{comp}(F)} \big(2|V(F')| - 2 + \rho\langle F' \cup L_A(F') \rangle - f_G(L_A(F'))\big) \\ &= \sum_{F' \in \mathrm{comp}(F)} \big(2|V(F')| - 2 + \rho_{G,A}(F') - |V(F') \cap A|\big) . \end{aligned}$$ □

Lemma 4.5 implies that $f_{G,A}$ is invariant up to the A-equivalence of G. Analogously to Theorem 4.1, the following lemma characterizes the matchings in $(E, f_{G,A})$.

Lemma 4.6. *A subset $F \subseteq E$ is a matching in $(E, f_{G,A})$ if and only if*

- *$G[F]$ contains no cycle, and*
- *for each $F' \in \mathrm{comp}(F)$, we have $|V(F') \cap A| \leq 2$ and the A-path between the two terminals is non-zero if $|V(F') \cap A| = 2$.*

Proof. By Lemma 4.5, it suffices to check the statement for each connected edge set F.

If $G[F]$ contains a cycle, then $f_{G,A}(F) \leq 2|V(F)| - 1 < 2|F|$ by Lemma 4.5, since $\rho_{G,A}(F) - |V(F) \cap A| \leq 1$ follows from the fact that $|V(F) \cap A| \geq 1$ whenever $\rho_{G,A}(F) = 2$. Hence, F is not a matching.

Suppose that $G[F]$ contains no cycle. If $|V(F) \cap A| \leq 1$, then $\rho_{G,A}(F) = |V(F) \cap A|$ and hence $f_{G,A}(F) = 2|V(F)| - 2 = 2|F|$, i.e., F is a matching. If $|V(F) \cap A| = 2$, then $\rho_{G,A}(F) = |V(F) \cap A|$ if and only if the A-path between the two terminals is non-zero, and hence F is a matching if and only if the A-path is non-zero. If $|V(F) \cap A| \geq 3$, then $\rho_{G,A}(F) < |V(F) \cap A|$ and hence $f_{G,A}(F) \leq 2|V(F)| - 1$. Thus, F is not a matching. □

We have the following relation between the maximum objective values of the two problems. Recall that $\nu(E, f_{G,A})$ denotes the maximum size of a matching in $(E, f_{G,A})$, and let $\mu(G, A)$ denote the maximum number of vertex-disjoint non-zero A-paths in G.

Theorem 4.7. *If G is connected and $A \neq \emptyset$, then $\nu(E, f_{G,A}) = |V| - |A| + \mu(G, A)$.*

Proof. Let us simply denote $\nu := \nu(E, f_{G,A})$ and $\mu := \mu(G, A)$. Let $F \subseteq E$ be a maximum matching in $(E, f_{G,A})$. We denote by c_i the number of connected components of $G[F]$ containing exactly $i \in \{0, 1, 2\}$ terminals, where each isolated terminal contributes c_1. By Lemma 4.6 and the maximality of F, we may assume $c_2 \leq \mu$ and that each connected component of $G[F]$ contains one or two terminals. Therefore, $\nu = |F| = |V| - (c_1 + c_2) = |V| - |A| + c_2 \leq |V| - |A| + \mu$.

The converse direction can be easily seen as follows. Let F be the edge set of a maximum number of vertex-disjoint non-zero A-paths. By Lemma 4.6, F is a matching extendable to be a maximal matching $\hat{F}$ so that each connected component of $G[\hat{F}]$ contains one or two terminals. Hence, $\nu \geq |\hat{F}| = |V| - |A| + \mu$. □

The proof implies that, for any maximum matching $F \subseteq E$ in $(E, f_{G,A})$, the corresponding subgraph $G[F]$ contains a maximum number of vertex-disjoint non-zero A-paths in G, which can be extracted by the depth first search from each terminal in A by Lemma 4.6. This observation completes our reduction.

4.4 Applications

4.4.1 Enumeration of nontrivial double circuits

The key observation for applying Lovász' theory on matroid matching is to enumerate all nontrivial double circuits of $(E, f_{G,A})$. The following lemma is an extension of [59, Lemma 3.4] and shows that $(E, f_{G,A})$ is a relatively simple 2-polymatroid.

Lemma 4.8. *For a Γ-labeled graph $G = (V, E)$ and a terminal set $A \subseteq V$, every nontrivial double circuit $D \subseteq E$ in $(E, f_{G,A})$ is a tree whose leaves are all terminals, and is in one of the following forms (see Fig. 4.1).*

- **D1** *$V(D) \cap A = \{a_1, a_2, a_3\}$, and all three A-paths in $G[D]$ are zero paths and intersect at a non-terminal $v \in V(D) \setminus A$.*
- **D2** *$V(D) \cap A = \{a_1, a_2, a_3, a_4\}$, and $G[D]$ contains six A-paths. At most one of the six A-paths is a zero path, and all A-paths intersect at a non-terminal $v \in V(D) \setminus A$.*
- **D3** *$V(D) \cap A = \{a_1, a_2, a_3, a_4\}$, and $G[D]$ contains vertex-disjoint a_1–a_2 path and a_3–a_4 path, which are connected by a path between two non-terminals $u, v \in V(D) \setminus A$. The a_1–a_2 path is a zero path and intersects v, and the others are non-zero.*
- **D4** *$V(D) \cap A = \{a_1, a_2, a_3, a_4\}$, and $G[D]$ contains only four A-paths. Three of the four connecting a_1, a_2, a_3 intersect at a non-terminal $v \in V(D) \setminus A$, and the rest connects a_3 and a_4. The A-path between a_1 and a_2 is a zero path, and the others are non-zero.*
- **D5** *$V(D) \cap A = \{a_1, a_2, a_3, a_4\}$, and $G[D]$ contains only three A-paths, which are non-zero and intersect at a_1.*

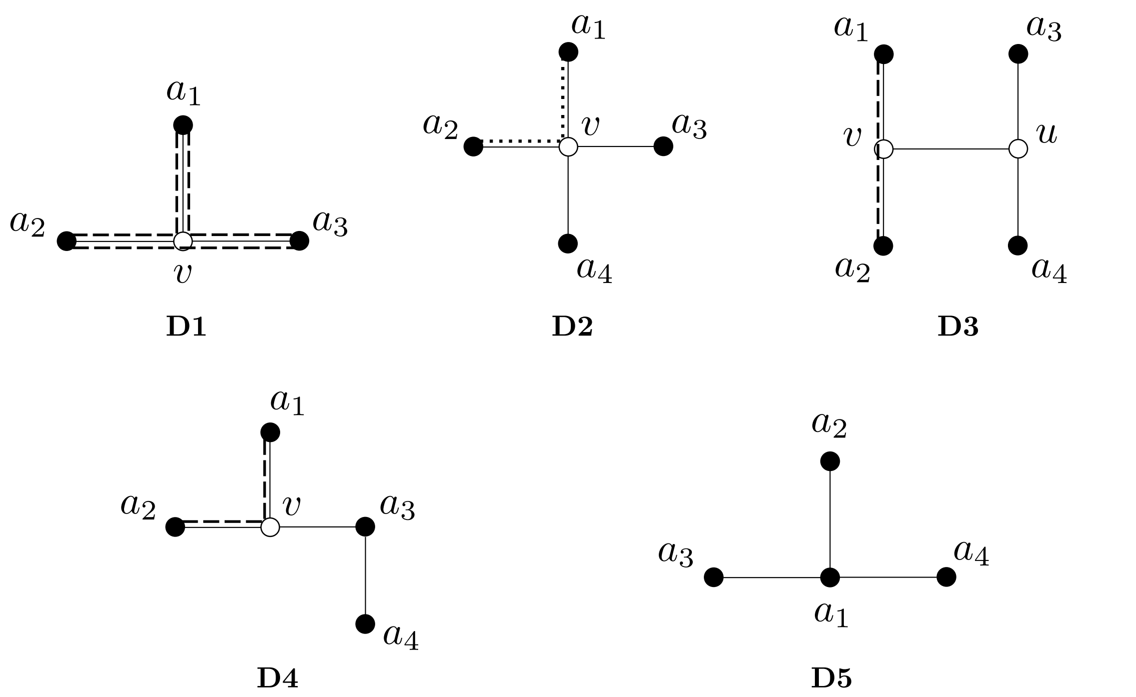

Figure 4.1: Nontrivial double circuits, where the dashed A-paths are zero paths, the dotted A-path has an arbitrary label, and the others are non-zero.

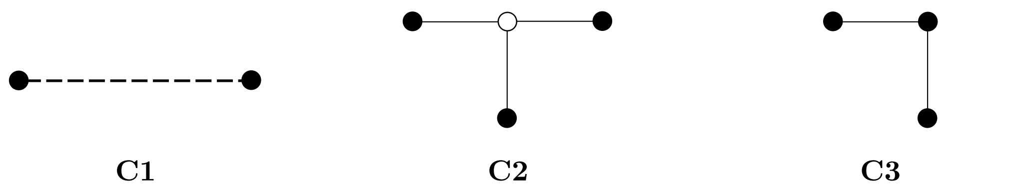

Figure 4.2: Circuits with no cycle, where only the dashed A-path is not non-zero.

Proof. We first enumerate all the patterns of circuits in $(E, f_{G,A})$ that contain no cycle. Let $C \subseteq E$ be such a circuit. By Lemma 4.5, one can easily check that C is connected. Hence, we have $|V(C)| = |C| + 1$, and $\rho_{G,A}(C) = |V(C) \cap A| - 1$ by Lemma 4.5. This means that $(\rho_{G,A}(C), |V(C) \cap A|) = (0, 1)$, $(1, 2)$, or $(2, 3)$.

The first case cannot occur, since $\rho_{G,A}(C) \geq 1$ whenever $|V(C) \cap A| \geq 1$. If C forms a zero A-path, then C is a circuit (Type **C1** in Fig. 4.2). Otherwise, we have $|V(C) \cap A| = 3$. Since each leaf of $G[C]$ must be a terminal (otherwise, the deletion of the edge incident to a non-terminal leaf decreases the value of $f_{G,A}$ by 2), C is of Type **C2** or **C3** in Fig. 4.2.

Let $D \subseteq E$ be a nontrivial double circuit in $(E, f_{G,A})$. We shall prove that $G[D]$ is connected and contains no cycle. If D is not connected, then D consists of two connected components each of which is a circuit in $(E, f_{G,A})$, and hence D is trivial.

Thus D is connected, i.e., we have the following equality by Lemma 4.5:

$$2|D| - 2 \;=\; f_{G,A}(D) \;=\; 2|V(D)| - 2 + \rho_{G,A}(D) - |V(D) \cap A|.$$

Therefore, by the definition of $\rho_{G,A}$, we have $2(|D| - |V(D)|) = \rho_{G,A}(D) - |V(D) \cap A| \leq 1$, which implies $|D| \leq |V(D)|$. This means that $G[D]$ contains at most one cycle.

Suppose that there exists exactly one cycle in $G[D]$, and let $C \subseteq D$ be its edge set. Then, we have $\rho_{G,A}(D) = |V(D) \cap A|$, i.e., $(\rho_{G,A}(D), |V(D) \cap A|) = (0, 0)$, $(1, 1)$, or $(2, 2)$. In the first or second case, $D - e$ is a matching for every edge $e \in C$, a contradiction. Also when $G[D]$ contains two A-paths in the third case, $D - e$ is a matching for some edge $e \in C$ since $G[D]$ is unbalanced (implying that at least one of the two A-paths is non-zero) by $\rho_{G,A}(D) = 2$. Suppose that $(\rho_{G,A}(D), |V(D) \cap A|) = (2, 2)$ and $G[D]$ contains only one A-path, which is a zero path. Then, $G[D]$ consists of the zero A-path and an unbalanced cycle joined at exactly one vertex. This means that D is decomposed into one circuit of Type **C1** and one forming an unbalanced cycle, and hence D is trivial, a contradiction.

Thus, D is a tree, and hence we have $|V(D)| = |D| + 1$ and $\rho_{G,A}(D) - |V(D) \cap A| = -2$. Since $|V(D) \cap A| \geq 1$ implies $\rho_{G,A}(D) \geq 1$, it suffices to consider two cases: $|V(D) \cap A| = 3$ and $|V(D) \cap A| = 4$. Similarly to the case of circuits, each leaf of $G[D]$ is a terminal.

Suppose that $|V(D) \cap A| = 3$. We then have $\rho_{G,A}(D) = 1$, and hence $G[D]$ contains no non-zero A-path. If $G[D]$ contains only two A-paths, then D is trivial. Therefore, $G[D]$ contains three A-paths, and hence it is of Type **D1**.

Suppose that $|V(D) \cap A| = 4$. We then have $\rho_{G,A}(D) = 2$, and hence $G[D]$ contains at least one non-zero A-path. Every tree with at most four leaves is one of the following forms: a path, three paths joined at one vertex (Fig. 4.1 **D4** and **D5**), four paths joined at one vertex (Fig. 4.1 **D2**), and two vertex-disjoint paths connected by a path between their internal vertices (Fig. 4.1 **D3**). If some A-paths in $G[D]$ have labels violating the conditions of **D2**–**D5** in the statement, then one can easily check from Lemma 4.6 and the list of circuits given in Fig. 4.2 that D is trivial, there is an edge $e \in D$ such that $D - e$ is a matching, or $f_{G,A}(D) \neq 2|D| - 2$. $\square$

As we reviewed in Section 2.4.2, it is preferable in applying Theorem 2.16 that the kernel of every double circuit has the rank at least one. By contracting an element in the kernel, one can reduce the problem to a smaller one. In our situation, we observe that the kernel of a double circuit of Type **D1**, **D2**, **D3**, or **D4** in Lemma 4.8 contains some loop $\ell_{v,\gamma} \in L$ in $(\tilde{E}, f)$, and hence, by contracting it, we can reduce the problem size appropriately. For the completeness, we shall formalize this fact in terms of packing non-zero A-paths as follows. In the following two lemmas, let $\nu := \nu(E, f_{G,A})$.

Lemma 4.9. *Let $F \subseteq E$ be a ν-double-flower containing a double circuit of Type* **D1**, **D2**, **D3**, *or* **D4**. *Then, there exists an A-equivalent Γ-labeled graph G' such that $\nu(E, f_{G',A+v}) < \nu$, where v is the vertex specified in Lemma* 4.8.

Proof. Let D be the double circuit in F. By Proposition 2.6, there exists an A-equivalent Γ-labeled graph G' such that all edges along each zero A-path in $G[D]$ have the identity label. Then, observe that every circuit in D spans the loop $\ell_{v,1_\Gamma} \in L$ on v with $\psi_{\tilde{G}'}(\ell_{v,1_\Gamma}) = \bullet$, i.e., the kernel of D contains $\ell_{v,1_\Gamma}$, and hence $f_{G',A+v}(C) = f_{G,A}(C) - 1$ for every circuit C in D. In the rest of this proof, we simply denote $f_{G,A}, f_{G',A+v}$ by f_1, f_2, respectively.

Suppose to the contrary that (E, f_2) has a matching M of size ν. Let us choose such M so that $|M \cap F|$ is maximized. We then have $f_2(F) = f_1(F) - 1 = 2\nu + 1 > 2\nu = f_2(M)$. Thus, there exists an edge $e \in D \setminus \mathrm{span}_{f_2}(M)$. Since $M + e$ cannot be a matching in (E, f_1), we have $f_1(M + e) \leq 2\nu + 1$. Furthermore, since $f_2(M + e) \leq f_1(M + e)$ and $e \notin \mathrm{span}_{f_2}(M)$, we obtain $f_1(M + e) = f_2(M + e) = 2\nu + 1$. This implies that $M + e$ is a ν-flower both in (E, f_1) and in (E, f_2).

Let C be the circuit in $M + e$ with respect to (E, f_2). Then, C is also the circuit of $M + e$ in (E, f_1), since otherwise $M + e$ becomes a matching. If $C \not\subseteq F$, then, for any edge $e' \in C \setminus F$, $M' := M + e - e'$ is a matching in (E, f_2) of size ν with $|M' \cap F| > |M \cap F|$. This contradicts the choice of M. If $C \subseteq F$, then $f_2(C) = f_1(C) - 1$, which contradicts the fact that C is a circuit both in (E, f_1) and in (E, f_2). $\square$

In Mader's $\mathcal{S}$-paths case, Lovász introduced a notion of a *regular set* to solve the case when we encounter a double circuit of Type **D5**. He claimed that the set of edges incident to a_1 forms a regular set in [59, Lemma 3.5 (b)]. This claim turns out to be false, but at least we can apply the proof idea of [59, Lemma 1.6] to accomplish our purpose as follows.

Lemma 4.10. *Let F be a ν-double-flower containing a double circuit of Type* **D5**. *Then, $\nu(E(G - a_1), f_{G-a_1,A-a_1}) < \nu$, where a_1 is the vertex specified in Lemma* 4.8.

Proof. We simply denote $f_{G,A}$ and $f_{G-a_1,A-a_1}$ by f_1 and f_2, respectively, and let $E' := E(G - a_1)$. Suppose to the contrary that (E', f_2) has a matching M of size ν. Let us choose such M so that $|M \cap F|$ is maximized, and let $\tilde{M} := \mathrm{span}_{f_1}(M)$. Let D be the double circuit in F, and let E_1 be the set of edges in D incident to a_1. Observe that $E_1 \cap C \neq \emptyset$ for every circuit C in D since D is of Type **D5**.

Suppose that $F \not\subseteq \tilde{M} \cup E_1$. Take an edge $e \in F \setminus (\tilde{M} \cup E_1)$. Then, $M + e$ is a ν-flower in (E', f_2) and the circuit C in $M + e$ is not included in F since every circuit in D intersects E_1. Therefore, for an edge $e' \in C \setminus F$, $M' := M + e - e'$ is a matching in (E', f_2) of size ν with $|M' \cap F| > |M \cap F|$, which contradicts the choice of M.

Suppose that $F \subseteq \tilde{M} \cup E_1$. Recall that $E_1 \subseteq D$ and every A-path in $G[D]$ is non-zero. Therefore, if there exists an edge $e \in E_1$ incident to a connected component $M' \in \mathrm{comp}(\tilde{M})$ that does not contain a non-zero A-path, then $\rho_{G,A}(M' + e) = \rho_{G,A}(M') + 1$ holds, and hence $f_1(\tilde{M} + e) = f_1(\tilde{M}) + 2$. This however implies $f_1(M + e) = f_1(M) + 2$, and hence $M + e$ is a matching of size $\nu + 1$ in (E, f_1), a contradiction. Otherwise (i.e., if each connected component of $G[\tilde{M}]$ around a_1 contains a non-zero A-path), we have $f_1(\tilde{M} \cup E_1) = f_1(\tilde{M}) + 1 = f_1(M) + 1 = 2\nu + 1$. However, since $F \subseteq \tilde{M} \cup E_1$, we also have $f_1(\tilde{M} \cup E_1) \geq f_1(F) = 2\nu + 2$, a contradiction. $\square$

4.4.2 Alternative proof for min-max duality

First, we restate the min-max formula of Chudnovsky et al. [7] (Theorem 3.5). For a Γ-labeled graph $G = (V, E)$ and vertex subsets $A, X \subseteq V$, we define

$$t(G, A; X) := |X| + \sum_{H \in \mathrm{comp}(G-X-E_0)} \left\lfloor \frac{|V(H) \cap A|}{2} \right\rfloor,$$

where $E_0 := \{\, e = uv \in E \mid u, v \in A,\ \psi_G(e) = 1_\Gamma \,\}$.

Theorem 4.11 (Restatement of Theorem 3.5)**.** *Let $G = (V, E)$ be a Γ-labeled graph and $A \subseteq V$ a terminal set. Then,*

$$\mu(G, A) = \min_{G', A', X} t(G', A'; X),$$

where the minimum is taken over all A-equivalent Γ-labeled graphs G' and all vertex subsets $A', X \subseteq V$ with $A \subseteq A'$.

Proof. The direction of $\leq$ is easy to see as follows. For any A-equivalent Γ-labeled graphs G', any vertex subset $X \subseteq V$, and any A' with $A \subseteq A' \subseteq V$, we have

$$\mu(G, A) = \mu(G', A) \leq \mu(G', A' \cup X) \leq t(G', A'; X).$$

The first equality holds since shifting at any non-terminal in $V \setminus A$ does not change the label of any A-path. The next inequality holds since each non-zero A-path contains at least one non-zero $(A' \cup X)$-path as its subpath. The final inequality holds, since at most $|X|$ paths in a family of vertex-disjoint $(A' \cup X)$-paths can intersect X and each connected component of $G - X$ contains a family of vertex-disjoint A'-paths of size at most half of the number of terminals in it.

We now show the converse direction by using Lovász' theorem (Theorem 2.16) for the associated 2-polymatroid $(E, f_{G,A})$. The proof is done by induction on $|V \setminus A| + |E|$. We may assume that G is connected, A is nonempty, and there is no A-path consisting of a single edge with the identity label. By Theorem 2.16, we split the proof into four cases.

<u>**Case 1.** When (i) in Theorem 2.16 holds.</u>

We then have $f_{G,A}(E) = 2\nu(E, f_{G,A}) + 1$, and by Theorem 4.7 and Lemma 4.5,

$$\begin{aligned}
\mu(G, A) &= \nu(E, f_{G,A}) - |V| + |A| \\
&= \frac{f_{G,A}(E) - 1}{2} - |V| + |A| \\
&= \frac{2|V| - 3 + \rho_{G,A}(E) - |A|}{2} - |V| + |A| \\
&= \frac{\rho_{G,A}(E) + |A| - 3}{2}.
\end{aligned} \tag{4.5}$$

If $\rho_{G,A}(E) = 1$, then $|A| \geq 2$ and E is A-balanced, and hence G is A-equivalent to a trivially-balanced Γ-labeled graph G'. This implies that $t(G', V; \emptyset) = 0$, which is no more than $\mu(G, A)$ by (4.5). If $\rho_{G,A}(E(G)) = 2$, then $|A|$ is odd, and hence by (4.5), we have

$$\mu(G, A) = \frac{|A| - 1}{2} = \left\lfloor \frac{|A|}{2} \right\rfloor = t(G, A; \emptyset).$$

Case 2. When (ii) in Theorem 2.16 holds.

Then, there is a partition $\{E_1, E_2\}$ of E such that $\nu(E, f_{G,A}) = \nu(E_1, f_{G,A}) + \nu(E_2, f_{G,A})$. By Theorem 4.7,

$$\begin{aligned} \mu(G, A) &= \nu(E, f_{G,A}) - |V| + |A| \\ &= \nu(E_1, f_{G,A}) + \nu(E_2, f_{G,A}) - |V \setminus A| \\ &= \sum_{i=1}^{2} (\nu(E_i, f_{G,A}) - |V(E_i) \setminus A|) + |(V(E_1) \cap V(E_2)) \setminus A| \\ &= \sum_{i=1}^{2} \mu(G[E_i], A) + |(V(E_1) \cap V(E_2)) \setminus A|. \end{aligned} \tag{4.6}$$

By the induction hypothesis, for each $i \in \{1, 2\}$, there exist G_i, A_i, X_i such that G_i is A-equivalent to $G[E_i]$, $A \subseteq A_i \subseteq V$, $X_i \subseteq V$, and $\mu(G[E_i], A) = t(G_i, A_i; X_i)$. Define $X := X_1 \cup X_2 \cup (V(E_1) \cap V(E_2) \setminus A)$ and $A' := A_1 \cup A_2$. For each $i \in \{1, 2\}$, let E_i' be the set of edges in E_i that are contained in $G - X$. Since $V(E_1') \cap V(E_2') \subseteq A$ by $(V(E_1) \cap V(E_2)) \setminus A \subseteq X$, there exists an A-equivalent Γ-labeled graph G' such that $G_i[E_i'] = G'[E_i']$ for each $i \in \{1, 2\}$. We then have $\mu(G, A) \geq t(G', A'; X)$ by (4.6).

Case 3. When (iii) in Theorem 2.16 holds.

Then, there exists an edge $e = uv \in E$ contained in the span of every maximum matching in $(E, f_{G,A})$.

We first show that, for any maximum matching M in $(E, f_{G,A})$ with $e \notin M$,

if e connects distinct connected subsets of M, say $M_1, M_2 \in \mathrm{comp}(M)$,
then $G[M_i]$ contains a non-zero A-path for each $i \in \{1, 2\}$. (4.7)

To see this, observe first

$$\sum_{i=1,2} (\rho_{G,A}(M_i) - |V(M_i) \cap A|) = 2 + \rho_{G,A}(M_1 \cup M_2 + e) - |V(M_1 \cup M_2 + e) \cap A|,$$

by $f_{G,A}(M) = f_{G,A}(M + e)$ and Lemma 4.5. Moreover, since F is a matching, we have $\rho_{G,A}(M_i) = |V(M_i) \cap A|$, which means $2 + \rho_{G,A}(M_1 \cup M_2 + e) = |V(M_1 \cup M_2 + e) \cap A|$. Therefore, $(\rho_{G,A}(M_1 \cup M_2 + e), |V(M_1 \cup M_2 + e) \cap A|) = (0, 2)$, $(1, 3)$, or $(2, 4)$. However, if $|V(M_1 \cup M_2 + e) \cap A| \geq 1$ then $\rho_{G,A}(M_1 \cup M_2 + e) \geq 1$, and if $|V(M_1 \cup M_2 + e) \cap A| \geq 3$ then $\rho_{G,A}(M_1 \cup M_2 + e) \geq 2$ since $G[M_1]$ or $G[M_2]$ contains a non-zero A-path in this case. It thus follows that $|V(M_1 \cup M_2 + e)| = 4$ holds and $G[M_i]$ contains a non-zero A-path for each $i \in \{1, 2\}$.

Suppose that $u, v \in A$. We may assume that $\psi_G(e) \neq 1_\Gamma$ since otherwise we can delete the edge $e = uv \in E$ and use induction. Let $\mathcal{P}$ be a maximum family of vertex-disjoint non-zero A-paths in G, and $V_\mathcal{P}$ and $E_\mathcal{P}$ its vertex set and edge set, respectively. Suppose that $u \notin V_\mathcal{P}$. Then, we must have $v \in V_\mathcal{P}$. If we extend $E_\mathcal{P}$ to a maximum matching M, then u and v belong to distinct connected components of M and moreover the component that contains u does not contain a non-zero A-path. This however contradicts (4.7). Thus, every maximum family of vertex-disjoint non-zero A-paths intersects the terminal u, and hence we can delete u (by adding u to X) and use induction to complete the proof.

Suppose that $u \notin A$. If the addition of u to A does not increase the value of μ, then we can use induction since $|V \setminus A|$ decreases. Otherwise, there are $\mu + 1$ vertex-disjoint non-zero $(A + u)$-paths $P_0, P_1, \ldots, P_\mu$ in G such that u is an end vertex of P_0. Let $a \in A$ be the other end vertex of P_0. If G contains no A-path traversing the edge $e = uv$, then we can delete e and use induction to complete the proof. Hence, we assume that there exists an A-path Q traversing e in G.

Let $A_\mathcal{P} := A \cup (\bigcup_{i=1}^{\mu} V(P_i))$, and Q' be the subpath of Q that is an $A_\mathcal{P}$-path traversing the edge e. We walk along P_0 from a until we hit Q' first, and then continue walking along Q' so that we traverse e until the end of Q'. The resulting path, denoted by P_0', is an $A_\mathcal{P}$-path which starts at a and traverses e. Then, an edge set $(E(P_0') - e) \cup (\bigcup_{i=1}^{\mu} E(P_i))$ is a matching extendable to be a maximum matching M with $e \notin M$. Furthermore, e connects distinct connected components of $G[M]$, at least one of which does not contain a non-zero A-path. This again contradicts (4.7), and we complete the proof.

Case 4. When (iv) in Theorem 2.16 holds.

Then, there exists a nontrivial ν-double-flower $F \subseteq E$. By Lemma 4.8, F contains a double circuit D of one of the five types. If D is of Type **D1**, **D2**, **D3**, or **D4**, then, by Theorem 4.7, Lemma 4.9, and the induction hypothesis, we have

$$\begin{aligned} \mu(G, A) &= \nu(E, f_{G,A}) - |V| + |A| \\ &\geq \nu(E, f_{G',A+v}) + 1 - |V| + |A + v| - 1 \\ &= \mu(G', A + v) \\ &= t(G'', A'; X), \end{aligned} \tag{4.8}$$

for some A-equivalent Γ-labeled graph G', $(A + v)$-equivalent G'', and vertex subsets $A', X \subseteq V$ with $A + v \subseteq A'$. Note that G'' is A-equivalent to G and $A \subseteq A' \subseteq V$. Thus, we complete the proof by (4.8).

If D is of Type **D5**, then, by Lemma 4.10 instead of Lemma 4.9, we have

$$\begin{aligned}\mu(G, A) &= \nu(E, f_{G,A}) - |V| + |A| \\ &\geq \nu(E(G - a_1), f_{G-a_1,A-a_1}) + 1 - |V - a_1| + |A - a_1| \\ &= \mu(G - a_1, A - a_1) + 1 \\ &= t(G', A'; X) + 1,\end{aligned}$$

for some Γ-labeled graph G' A-equivalent to $G - a_1$, and vertex subsets $A', X \subseteq V - a_1$ with $A - a_1 \subseteq A'$. Define $X' := X + a_1$. Then, we have $t(G', A'; X) + 1 = t(G'', A'; X')$, where G'' is A-equivalent to G (and $G'' - a_1 = G'$), $A \subseteq A' \subseteq V$, and $X' \subseteq V$. This completes the proof. $\square$

4.4.3 Algorithm via matroid matching

The high-level idea of Lovász' matroid matching algorithm [61] is to maintain a nonempty family of same-size matchings in a 2-polymatroid (E, f) with "improving" it repeatedly. Most parts of his algorithm work for general 2-polymatroids, and the crucial step using linear representations is as follows (cf. [61, Algorithm 3.2]): for a positive integer $k \in \mathbb{Z}_{>0}$, a nontrivial k-double-flower F, and a matching M of size k such that $\mathrm{span}(M)$ does not include the kernel of the double circuit in F, to find a matching of size $k + 1$. Based on the enumeration of nontrivial double circuits (Lemma 4.8), we show that this step can be done for our 2-polymatroid.

Claim 4.12. *Let $(E, f_{G,A})$ be the 2-polymatroid defined in Section 4.3, $k \in \mathbb{Z}_{>0}$ a positive integer, $F \subseteq E$ a nontrivial k-double-flower in $(E, f_{G,A})$, and $M \subseteq E$ a matching of size k such that $\mathrm{span}_{f_{G,A}}(M)$ does not include the kernel of the double circuit in F. Then, one can find a matching of size $k + 1$ in $(E, f_{G,A})$ in polynomial time.*

Proof. Let D be the nontrivial double circuit in F. We prove the following two cases separately: when D is of one of Types **D1**–**D4** in Lemma 4.8, and when it is of Type **D5**.

Case 1. When D is of one of Types **D1**–**D4**.

As seen in the beginning of the proof of Lemma 4.9, there exists an A-equivalent Γ-labeled graph in which all edges along each zero A-paths in $G[D]$ have the identity label. Without loss of generality (by shifting at v in advance if necessary), we assume that G itself satisfies this condition. Then, the kernel of D is $\{\ell_{v,1_\Gamma}\}$ as observed in the proof of Lemma 4.9. Hence, the assumption implies $\ell_{v,1_\Gamma} \notin \mathrm{span}_{f_{G,A}}(M)$. This means

$$f_{G,A+v}(M) = f_{G,A}(M + \ell_{v,1_\Gamma}) - f_{G,A}(\ell_{v,1_\Gamma}) = (f_{G,A}(M) + 1) - 1 = f_{G,A}(M) = 2|M|,$$

i.e., M is a matching of size k also in $(E, f_{G,A+v})$. The proof of Lemma 4.9 claims that, if $k = \nu(E, f_{G,A})$, then this $(E, f_{G,A+v})$ contains no matching of size k. In other words, if $(E, f_{G,A+v})$ contains a matching of size k, then $(E, f_{G,A})$ contains a larger matching, which is indeed constructed as follows.

We simply denote $f_{G,A}$ and $f_{G,A+v}$ by f_1 and f_2, respectively. Since $f_2(F) = f_1(F) - 1 = 2k+1 > 2k = f_2(M)$, there exists an edge $e \in D \setminus \mathrm{span}_{f_2}(M)$. If $f_1(M+e) = 2k+2$, then we have done. Suppose that $f_1(M+e) \leq 2k+1$. Since $f_1(M+e) \geq f_2(M+e) \geq 2k+1$, we have $f_1(M+e) = f_2(M+e) = 2k+1$, and hence $M+e$ is a k-flower both in (E, f_1) and in (E, f_2).

Let C be the circuit in $M+e$ with respect to f_2. Then, so is it with respect to f_1, since otherwise $M+e$ is a matching in (E, f_1). If $C \not\subseteq F$, then, for any edge $e' \in C \setminus F$, $M' := M + e - e'$ is a matching in (E, f_2) with $|M' \cap F| > |M \cap F|$. Since $f_1(M+e) = f_2(M+e)$, the kernel $\{\ell_{v,1_\Gamma}\}$ of D is not spanned by M' as well as $M+e$. Therefore, by replacing M with M', after at most k iterations, we can find a matching of size $k+1$ in (E, f_1).

Case 2. When D is of Type **D5**.

It is easy to check that the kernel of D is $\{\, \ell_{a_1,\gamma} \in L \mid \gamma \in \Gamma \setminus \{1_\Gamma\} \,\}$. If the connected component of $G[M]$ containing a_1 contains another terminal in A, then the A-path in it is non-zero by Lemma 4.6, which implies that the kernel is spanned by M. Hence, the assumption means that the connected component of $G[M]$ containing a_1 contains no other terminal in A.

Without loss of generality, we may assume that G is connected and contains no redundant vertex, i.e., every vertex is contained in some A-path. Then, $G - a_1$ is connected by Proposition 2.1. This implies that every connected component of $G[M] - a_1$ can be connected to another connected component by an edge not in M, and hence one can take a matching M' of size k in $(E(G-a_1), f_{G-a_1,A-a_1})$. Lemma 4.10 claims that if $k = \nu(E, f_{G,A})$, then this $(E(G-a_1), f_{G-a_1,A-a_1})$ contains no matching of size k. In other words, if $(E(G-a_1), f_{G-a_1,A-a_1})$ contains a matching of size k, then $(E, f_{G,A})$ contains a larger matching, which is indeed constructed as follows.

We simply denote $f_{G,A}$ and $f_{G-a_1,A-a_1}$ by f_1 and f_2, respectively, and let $E' := E(G-a_1)$, $E_1 := \delta_G(a_1) \cap D$, and $\tilde{M} := \mathrm{span}_{f_1}(M')$.

Suppose that $F \not\subseteq \tilde{M} \cup E_1$. Take an edge $e \in F \setminus (\tilde{M} \cup E_1)$. If $f_1(M'+e) = 2k+2$, then we have done. Otherwise, we have $f_1(M'+e) = f_2(M'+e) = 2k+1$, i.e., $M'+e$ is a k-flower both in (E, f_1) and in (E', f_2). The circuit C in $M'+e$ is not included in F (since each circuit in F must contain some edge in E_1), and hence there exists an edge $e' \in C \setminus F$ such that $M'' := M' + e - e'$ is a matching with $|M'' \cap F| > |M' \cap F|$. Since we have $M'' \cap \delta_G(a_1) = \emptyset$ as well as $M'' \cap \delta_G(a_1) = \emptyset$, by replacing M' with M'', after at most k iterations, we can find a matching of size $k+1$.

Otherwise, $F \subseteq \tilde{M} \cup E_1$. If there exists an edge $e \in E_1$ incident to a connected component $M'' \in \mathrm{comp}(\tilde{M})$ that does not contain a non-zero A-path, then $\rho_{G,A}(M''+e) = \rho_{G,A}(M'') + 1$, and hence $f_1(\tilde{M}+e) = f_1(\tilde{M}) + 2$. This implies $f_1(M'+e) =$

$f_1(M') + 2 = 2k + 2$, i.e., $M' + e$ is a matching of size $k + 1$ in (E, f_1). Otherwise, each connected component of $G[\tilde{M}]$ around a_1 contains a non-zero A-path. Then, we have $f_1(\tilde{M} \cup E_1) = f_1(\tilde{M}) + 1 = 2k + 1$, which contradicts that $f_1(\tilde{M} \cup E_1) \geq f_1(F) = 2k + 2$ (recall that $F \subseteq \tilde{M} \cup E_1$). $\square$

4.5 Extension to Axiomatic Model

For the sake of convenience, we restate the problem. Recall that a family $\mathcal{F}$ of A-paths is said to be *symmetric* if $\bar{P} \in \mathcal{F}$ for every $P \in \mathcal{F}$.

Axiomatic Model

Input: An undirected graph G, a terminal set $A \subseteq V(G)$, and a symmetric family $\mathcal{F}$ of A-paths in G.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint A-paths in $\mathcal{F}$ such that $|\mathcal{P}|$ is maximum.

Definition 4.13 (Weak Triple Exchange)**.** Let $G = (V, E)$ be an undirected graph and $A \subseteq V$ a terminal set. A symmetric family $\mathcal{F}$ of A-paths in G is *weakly triple exchangeable* if it satisfies the following condition: for every A-path $P \in \mathcal{F}$, inner vertex $v \in V(P) \setminus A$, terminal $a \in A \setminus V(P)$, and a–v path Q in G openly disjoint from P, at least one of the two A-paths obtained by extending Q along P is in $\mathcal{F}$.

Theorem 4.14. *The axiomatic model with the weak triple exchange property reduces to the matroid matching problem.*

Proof. Let $G = (V, E)$ be an undirected graph, $A \subseteq V$ a terminal set, and $\mathcal{F}$ be a symmetric family of A-paths in G. Define a set function $f_{\mathcal{F}} \colon 2^E \to \mathbb{Z}_{\geq 0}$ by

$$f_{\mathcal{F}}(F) := \sum_{F' \in \mathrm{comp}(F)} \bigl(2|V(F')| - 2 + \rho_{\mathcal{F}}(F') - |V(F') \cap A|\bigr), \tag{4.9}$$

for each edge set $F \subseteq E$, where $\rho_{\mathcal{F}} \colon 2^E \to \mathbb{Z}$ is defined as

$$\rho_{\mathcal{F}}(F') := \begin{cases} 2 & (G[F'] \text{ contains an } A\text{-path in } \mathcal{F} \text{ or a cycle and } |V(F') \cap A| \geq 1), \\ 1 & \begin{array}{l}(G[F'] \text{ contains no } A\text{-path in } \mathcal{F} \text{ and no cycle and } |V(F') \cap A| \geq 1, \text{ or} \\ \ G[F'] \text{ contains a cycle and } |V(F') \cap A| = 0),\end{array} \\ 0 & (\text{otherwise}). \end{cases}$$

Claim 4.15. *If $\mathcal{F}$ is weakly triple exchangeable, then $(E, f_{\mathcal{F}})$ is a 2-polymatroid.*

Proof. It is obvious that we have $f_{\mathcal{F}}(\emptyset) = 0$. Let us fix arbitrary edge sets $X, Y \subseteq E$ with $X \subseteq Y$. For each $Y_i \in \mathrm{comp}(Y)$, define $\mathcal{X}_i := \{\, X_j \in \mathrm{comp}(X) \mid X_j \subseteq Y_i \,\}$.

To see the monotonicity, it suffices to show

$$2|V(Y_i)|-2+\rho_{\mathcal{F}}(Y_i)-|V(Y_i)\cap A| \geq \sum_{X_j \in \mathcal{X}_i} \left(2|V(X_j)| - 2 + \rho_{\mathcal{F}}(X_j) - |V(X_j) \cap A|\right), \quad (4.10)$$

for each $Y_i \in \mathrm{comp}(Y)$. Since the left-hand side of (4.10) is nonnegative by $|V(Y_i)| \geq 2$, we may assume $\mathcal{X}_i \neq \emptyset$. Since $Y_i \supseteq \bigcup_{X_j \in \mathcal{X}_i} X_j$ and each $X_j \in \mathcal{X}_i$ is connected, we have

$$\begin{aligned} 2|V(Y_i)| - |V(Y_i) \cap A| &= |V(Y_i)| + |V(Y_i) \setminus A| \\ &\geq \sum_{X_j \in \mathcal{X}_i} \left(|V(X_j)| + |V(X_j) \setminus A|\right) \\ &= \sum_{X_j \in \mathcal{X}_i} \left(2|V(X_j)| - |V(X_j) \cap A|\right). \end{aligned}$$

Let us take $k \in \arg\max_j \{\, \rho_{\mathcal{F}}(X_j) \mid X_j \in \mathcal{X}_i \,\}$. Then, $\rho_{\mathcal{F}}(Y_i) \geq \rho_{\mathcal{F}}(X_k)$ by $Y_i \supseteq X_k$. Since $\rho_{\mathcal{F}}(F) - 2 \leq 0$ for every connected edge set $F \subseteq E$, we have shown the inequality (4.10).

To see the submodularity, it suffices to show

$$f_{\mathcal{F}}(X+e) - f_{\mathcal{F}}(X) \geq f_{\mathcal{F}}(Y+e) - f_{\mathcal{F}}(Y)$$

for every edge $e = uv \in E \setminus Y$. Note that, by simple observation, the addition of one edge increases the value of $f_{\mathcal{F}}$ by at most 2.

Case 1. When $u, v \in V(X_i)$ for some $X_i \in \mathrm{comp}(X)$.

Let us take $Y_i \in \mathrm{comp}(Y)$ with $X_i \subseteq Y_i$. Suppose that $f_{\mathcal{F}}(X+e) = f_{\mathcal{F}}(X)$. Then, $\rho_{\mathcal{F}}(X_i+e) = \rho_{\mathcal{F}}(X_i) =: \rho$. Since $G[X_i+e]$ contains a cycle traversing e, we have $\rho = 1$ or 2. By the monotonicity of $\rho_{\mathcal{F}}$, it suffices to check the case of $\rho = 1$ and $\rho_{\mathcal{F}}(Y_i) = 1$. In this case, $G[X_i]$ contains a cycle and no terminal, and so does $G[Y_i]$, which implies $\rho_{\mathcal{F}}(Y_i+e) = 1$.

The addition of one edge connecting the same connected component increases the value of $f_{\mathcal{F}}$ by at most 1, since it adds no new vertex to the component. Then, the case of $f_{\mathcal{F}}(X+e) = f_{\mathcal{F}}(X)+1$ is obvious.

Case 2. When $u \in V(X_i)$ and $v \in V(X_j)$ for distinct $X_i, X_j \in \mathrm{comp}(X)$.

Let us take $Y_i, Y_j \in \mathrm{comp}(Y)$ with $X_i \subseteq Y_i$ and $X_j \subseteq Y_j$. Note that we may have $Y_i = Y_j$. If $\rho_{\mathcal{F}}(X_i) = \rho_{\mathcal{F}}(X_j) = 2$, then $f_{\mathcal{F}}(X+e) = f_{\mathcal{F}}(X)$. In this case, we have also $\rho_{\mathcal{F}}(Y_i) = \rho_{\mathcal{F}}(Y_j) = 2$, and hence $f_{\mathcal{F}}(Y+e) = f_{\mathcal{F}}(Y)$ whether $Y_i = Y_j$ or not. If $\rho_{\mathcal{F}}(X_i) = 2$ and $\rho_{\mathcal{F}}(X_j) = 1$, then $f_{\mathcal{F}}(X+e) = f_{\mathcal{F}}(X)+1$. In this case, we have $\rho_{\mathcal{F}}(Y_i) = 2$ and $\rho_{\mathcal{F}}(Y_j) \geq 1$, and hence $f_{\mathcal{F}}(Y+e) \leq f_{\mathcal{F}}(Y)+1$ whether $Y_i = Y_j$ or not.

Suppose $\rho_{\mathcal{F}}(X_i) = \rho_{\mathcal{F}}(X_j) = 1$ and $\rho_{\mathcal{F}}(X_i \cup X_j + e) = 1$. In this case, we have $f_{\mathcal{F}}(X+e) = f_{\mathcal{F}}(X)+1$, and hence it suffices to show $f_{\mathcal{F}}(Y+e) \leq f_{\mathcal{F}}(Y)+1$. Then, we assume $Y_i \neq Y_j$. We observe either that $G[X_i]$ and $G[X_j]$ contain cycles and no terminal, or that $G[X_i]$ and $G[X_j]$ contain neither an A-path in $\mathcal{F}$ nor a cycle, $|V(X_i) \cap A| \geq 1$, and $|V(X_j) \cap A| \geq 1$.

In the former case, if $\rho_{\mathcal{F}}(Y_i \cup Y_j + e) = 2$, then at least one of $G[Y_i]$ and $G[Y_j]$ contains a terminal, and hence $\rho_{\mathcal{F}}(Y_i) = 2$ or $\rho_{\mathcal{F}}(Y_j) = 2$, since each of $G[Y_i]$ and $G[Y_j]$ contains a cycle. This means $f_{\mathcal{F}}(Y + e) \leq f_{\mathcal{F}}(Y) + 1$ since $\rho_{\mathcal{F}}(Y_i) \geq \rho_{\mathcal{F}}(X_i) = 1$ and $\rho_{\mathcal{F}}(Y_j) \geq \rho_{\mathcal{F}}(X_j) = 1$.

In the latter case, suppose that $\rho_{\mathcal{F}}(Y_i) = \rho_{\mathcal{F}}(Y_j) = 1$. Then, $G[Y_i \cup Y_j]$ contain neither an A-path in $\mathcal{F}$ nor a cycle. Suppose that $G[Y_i \cup Y_j + e]$ contains an A-path $P \in \mathcal{F}$. Let $Q \notin \mathcal{F}$ be an A-path in $X_i \cup X_j + e$ from $a \in V(X_i) \cap A$ to $b \in V(X_j) \cap A$. Since both P and Q traverse the edge e and $Y_i \cup Y_j + e$ is a tree, we can partition P and Q into P_1, P_2, P_3 and Q_1, Q_2, Q_3, respectively, so that $E(P_1) \subseteq Y_i \setminus E(Q)$, $E(P_3) \subseteq Y_j \setminus E(Q)$, $E(Q_1) \subseteq Y_i \setminus E(P)$, $E(Q_3) \subseteq Y_j \setminus E(P)$, and $P_2 = Q_2$.

Since $P \neq Q$, we may assume that P_1 and Q_1 are both nonempty. Let $R_1 := P_1 * P_2 * P_3 = P$, $R_2 := Q_1 * P_2 * P_3$, and $R_3 := P_1 * \bar{Q}_1$. By the weak triple exchange property, $R_2 \in \mathcal{F}$ since R_3 is contained in $G[Y_j]$. If $P_3 = Q_3 = \emptyset$, then $R_2 = Q \notin \mathcal{F}$, a contradiction. If P_3 and Q_3 are both nonempty, then, by weak triple exchange, $Q_1 * P_2 * Q_3$ is also in $\mathcal{F}$, but this A-path coincides with Q, a contradiction again. □

Analogously to Lemma 4.6 and Theorem 4.7, we can show the following claims by almost the same arguments using Claim 4.15 instead of Lemma 4.5. These claims complete this proof since, for any maximum matching M in $(E, f_{\mathcal{F}})$, $G[M]$ contains a maximum family of vertex-disjoint A-paths in $\mathcal{F}$, which can be extracted by breadth first search.

Here, suppose that $\mathcal{F}$ satisfies the weak triple exchange property.

Claim 4.16. *A subset $F \subseteq E$ is a matching in $(E, f_{\mathcal{F}})$ if and only if*

- *$G[F]$ contains no cycle, and*
- *for each $F' \in \mathrm{comp}(F)$, we have $|V(F') \cap A| \leq 2$ and the A-path between the two terminals is in $\mathcal{F}$ if $|V(F') \cap A| = 2$.*

Claim 4.17. *If G is connected and $A \neq \emptyset$, then the maximum number of vertex-disjoint A-paths in $\mathcal{F}$ is equal to $\nu(E, f_{\mathcal{F}}) - |V| + |A|$.* □

Theorem 4.14 does not necessarily lead to a good characterization or an efficient algorithm, since we do not mention its nontrivial double circuits, i.e., how to reduce the problem size when we encounter (iv) in Theorem 2.16. In particular, the 2-polymatroid $(E, f_{\mathcal{F}})$ defined by (4.9) has another type of nontrivial double circuits, each of which is obtained by gluing two A-paths not in $\mathcal{F}$ as Fig. 4.3.

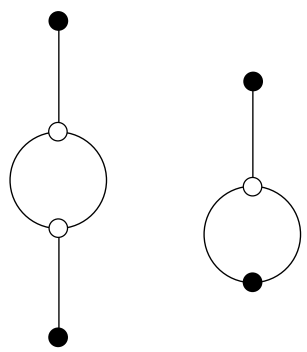

Figure 4.3: New nontrivial double circuits in $(E, f_{\mathcal{F}})$, where any A-path is not in $\mathcal{F}$.

Chapter 5

Packing A-paths via Linear Matroid Parity

In this chapter, we investigate the reducibility of the problem of packing A-paths to the linear matroid parity problem, which is a special case of the matroid matching problem with the input 2-polymatroids linearly represented. In particular, for the subgroup-forbidden model, we clarify when the problem admits a reasonable reduction.

This chapter is based on [94] and organized as follows. In Section 5.1, we review a reduction of Mader's $\mathcal{S}$-paths problem due to Schrijver [78, Section 73.1a], which in fact gives a linear representation of Lovász' 2-polymatroid in Section 4.1. Section 5.2 is devoted to presenting our result for the subgroup-forbidden model: a necessary and sufficient condition for the groups in question to admit a reasonable reduction extending Schrijver's one. The proof is divided into the sufficiency and necessity parts, which are shown in Sections 5.3 and 5.4, respectively. In Section 5.5, we apply our result to various cases and show that a large class admits our reduction, which leads to fast algorithms.

5.1 Schrijver's Reduction of Mader's $\mathcal{S}$-paths

For the sake of convenience, we restate the two problems again (the first appearances are in Sections 3.2 and 2.4.1).

Mader's $\mathcal{S}$-paths Problem

Input: An undirected graph $G = (V, E)$ and a terminal set $A \subseteq V$ with its partition $\mathcal{S}$.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint $\mathcal{S}$-paths in G such that $|\mathcal{P}|$ is maximum.

Linear Matroid Parity Problem

Input: A finite set E and a matrix $Z \in \mathbb{F}^{r \times 2E}$ over a field $\mathbb{F}$, where $r \in \mathbb{Z}_{>0}$.

Goal: Find a matching $M \subseteq E$ for Z such that $|M|$ is maximum.

For an instance of Mader's problem, we shall construct an associated matrix $Z \in \mathbb{Q}^{2V \times 2E}$ as follows.

Associate each edge $e = uw \in E$ with a 2-dimensional linear subspace

$$L_e := \{\, x \in (\mathbb{Q}^2)^V \mid x(u) + x(w) = \mathbf{0},\ x(v) = \mathbf{0} \ \ (v \in V \setminus \{u, w\}) \,\}$$

of $(\mathbb{Q}^2)^V$. For each terminal $a \in A_i$ $(i \in [k])$, define a 1-dimensional linear subspace

$$Q_a := \{\, x \in (\mathbb{Q}^2)^V \mid x(a) \in \left\langle \binom{1}{i} \right\rangle,\ x(v) = \mathbf{0} \ \ (v \in V - a) \,\}$$

of $(\mathbb{Q}^2)^V$, where $\langle x \rangle := \{\, kx \mid k \in \mathbb{F} \,\}$ for a vector $x \in \mathbb{F}^r$ over a field $\mathbb{F}$.

Let $Q := \sum_{a \in A} Q_a$ and $\mathcal{E} := \{\, L_e/Q \mid e \in E \,\}$. Note that $\dim(L_e/Q) = 2$ for every edge $e \in E$, since we may assume that no edge connects two terminals in the same class of $\mathcal{S}$. Let us construct a matrix $Z = (Z_e)_{e \in E} \in \mathbb{Q}^{2V \times 2E}$ associated with $\mathcal{E}$ by enumerating the bases of L_e/Q for all edges $e \in E$, i.e., $Z_e := (b_e, c_e) \in \mathbb{Q}^{2V \times 2}$ for each edge $e \in E$, where $\{b_e, c_e\}$ is an arbitrary fixed basis of L_e/Q. Then, each edge set $F \subseteq E$ is a matching for Z if and only if $\dim(L_F/Q) = 2|F|$, where $L_F := \sum_{e \in F} L_e$.

The above matrix Z is also obtained as follows: starting with the Kronecker product $B_G \otimes I_2 \in \mathbb{Q}^{2V \times 2E}$ of the incidence matrix $B_G \in \mathbb{Q}^{V \times E}$ of G (where each edge in G is assumed to be arbitrarily oriented) and the 2×2 identity matrix $I_2 \in \mathbb{Q}^{2 \times 2}$, appropriately eliminate it by using a basis $x \in (\mathbb{Q}^2)^V$ of Q_a for each terminal $a \in A_i$ $(i \in [k])$, which is defined by $x(a) := \binom{1}{i}$ and $x(v) := \mathbf{0}$ $(v \in V - a)$.

The following theorem shows a connection between the maximum number of vertex-disjoint $\mathcal{S}$-paths in G and the maximum cardinality of a matching for Z, which completes the reduction. A proof and how to construct vertex-disjoint paths in G from a matching for Z are described in Section 5.3 for a generalized version of this reduction (cf. the proof of Theorem 4.7).

Theorem 5.1 (Schrijver [78, (73.20)]). *Suppose that G is connected and $A \neq \emptyset$. Then, the maximum number of vertex-disjoint $\mathcal{S}$-paths in G is equal to $\nu(Z) - |V| + |A|$.*

By applying Theorems 2.17 and 2.18, we obtain deterministic and randomized algorithms for Mader's $\mathcal{S}$-paths problem which runs in $\mathrm{O}(|V|^\omega \cdot |E|)$ time and $\mathrm{O}(|V|^{\omega-1} \cdot |E|)$ time, respectively. Since we may assume that the input graph is simple (and hence $|E| = \mathrm{O}(|V|^2)$), these bounds are already better than $\mathrm{O}(|V|^5)$, which is obtained by applying the algorithm for packing non-zero A-paths due to Chudnovsky et al. [6] (see Theorem 3.6). In addition, Cheung et al. [5, Section 5.1.3] gave a further improvement as follows by utilizing a nice property of an associated matrix Z.

Theorem 5.2 (Cheung–Lau–Leung [5]). *One can solve Mader's $\mathcal{S}$-paths problem in $\mathrm{O}(|V|^\omega)$ time by a randomized algorithm.*

Since we utilize their algorithm and speeding-up argument, here we give a brief description of their ideas. It is based on the following matrix formulation.

Theorem 5.3 (Lovász [58])**.** *Let E be a finite set and $Z = (Z_e)_{e \in E} \in \mathbb{F}^{r \times 2E}$ a matrix over a field $\mathbb{F}$, where $r \in \mathbb{Z}_{>0}$. Suppose that $Z_e = (b_e, c_e) \in \mathbb{F}^{r \times 2}$ for each $e \in E$, and define*

$$Y := \sum_{e \in E} x_e \cdot \left(b_e^\top c_e - c_e^\top b_e \right), \tag{5.1}$$

where x_e $(e \in E)$ are algebraically independent indeterminates. Then, $2\nu(Z) = \operatorname{rank} Y$.

By Theorem 5.3, one can compute the maximum cardinality of a matching for $Z \in \mathbb{F}^{r \times 2E}$ by computing the rank of $Y \in \tilde{\mathbb{F}}^{r \times r}$, where $\tilde{\mathbb{F}}$ denotes the field obtained by adding the indeterminates x_e $(e \in E)$ to $\mathbb{F}$. Moreover, one can find a maximum matching by finding a minimal subset $F \subseteq E$ such that $\operatorname{rank} Y_F = \operatorname{rank} Y$, where

$$Y_F := \sum_{e \in F} x_e \cdot \left(b_e^\top c_e - c_e^\top b_e \right).$$

Their algorithm first assigns random values for the indeterminates x_e $(e \in E)$ in Y in Theorem 5.3. Let $Y' \in \mathbb{F}^{r \times r}$ be the resulting matrix. For each element $e \in E$, the algorithm checks whether $\tilde{Y}' := Y' - x'_e \cdot \left(b_e^\top c_e - c_e^\top b_e \right)$ has the same rank as Y' or not, where $x'_e \in \mathbb{F}$ denotes the value assigned to the indeterminate x_e. If $\operatorname{rank} \tilde{Y}' = \operatorname{rank} Y'$, it updates $Y' := \tilde{Y}'$. The update can be done quickly with the aid of a small area update formula due to Harvey [31] based on the Sherman–Morrison–Woodbury formula [91].

It is in fact needed to reformulate the problem (from finding a maximum matching to testing the existence of a parity base), but we omit the details since it is not essential here (see [5, Sections 4 and 6.5]). The significant fact is that their algorithm is based on Lovász' matrix formulation in Theorem 5.3 and small area update.

Since an associated matrix $Z \in \mathbb{Q}^{2V \times 2E}$ in Schrijver's reduction is sparse based on the incidence matrix B_G of G, each indeterminate x_e appears at most 16 entries of Y in (5.1). This property makes each small area update performable in constant time, which leads to the computational time bound in Theorem 5.2. The same sparsity is assumed also in our reduction (cf. Property 5.4).

5.2 Reducibility of Subgroup-Forbidden Model

We first restate the problem again (the first appearance is in Section 3.4). Let Γ be a group, and remember that an A-path is Γ'-*admissible* if its label is not in a proper subgroup Γ' of Γ.

Subgroup-Forbidden Model

Input: A Γ-labeled graph G, a terminal set $A \subseteq V(G)$, and a proper subgroup Γ' of Γ.

Goal: Find a family $\mathcal{P}$ of vertex-disjoint Γ'-admissible A-paths in G with $|\mathcal{P}|$ maximum.

5.2.1 Coherent representation

We introduce two natural and significant properties which are satisfied by Schrijver's reduction. Let $\mathbb{F}$ be a field. For a Γ-labeled graph $G = (V, E)$ and a terminal set $A \subseteq V$, we try to construct an associated matrix $Z = (Z_{v,e}) \in \mathbb{F}^{2V \times 2E}$, where $Z_{v,e} \in \mathbb{F}^{2\times 2}$ denotes the submatrix of Z corresponding to a vertex $v \in V$ and an edge $e \in E$.

The first property guarantees the sparsity of Z. In order to admit the improving argument due to Cheung et al. [5, Section 5.1.3], which leads to a speeding up of their algorithm (Theorem 5.2), the associated matrix Z is desired to be sparse based on the incidence matrix B_G of G as follows.

Property 5.4. *$Z_{v,e} = O$ for each edge $e = uw \in E$ and each vertex $v \in V \setminus \{u, w\}$.*

The second property is regarding a connection between the two problems. It is natural that a feasible or infeasible solution to the original problem is also feasible or infeasible, respectively, in the reduced problem.

Property 5.5. *For each A-path P in G, its edge set $E(P)$ is a matching for Z if and only if P is Γ'-admissible.*

For an instance of the subgroup-forbidden model, we call a matrix $Z \in \mathbb{F}^{2V\times 2E}$ a *coherent representation* if Z satisfies Properties 5.4 and 5.5.

5.2.2 Necessary and sufficient condition

For a positive integer $n \in \mathbb{Z}_{>0}$ and a field $\mathbb{F}$, we define $\mathrm{PGL}(n, \mathbb{F}) := \mathrm{GL}(n, \mathbb{F})/\{\, kI_n \mid k \in \mathbb{F}\setminus\{\mathbf{0}\}\,\}$, where $\mathrm{GL}(n, \mathbb{F})$ denotes the general linear group of degree n over $\mathbb{F}$ (i.e., the set of all nonsingular $n \times n$ matrices over $\mathbb{F}$ with the ordinary multiplication) and $I_n \in \mathrm{GL}(n, \mathbb{F})$ the $n \times n$ identity matrix. Each element of PGL is denoted by its representative in GL. For a fixed group Γ and its proper subgroup Γ', we refer to the subgroup-forbidden model as the Γ'-*forbidden model* if its input proper subgroup Γ' is fixed. We are now ready to state the main theorem.

Theorem 5.6. *Let Γ be a group, Γ' its proper subgroup, and $\mathbb{F}$ a field. Then, the following two statements are equivalent.*

(i) *The Γ'-forbidden model can be reduced to the linear matroid parity problem with a coherent representation over $\mathbb{F}$. That is, there exists a coherent representation $Z \in \mathbb{F}^{2V\times 2E}$ for any Γ-labeled graph $G = (V, E)$ and any terminal set $A \subseteq V$, under the fixed proper subgroup Γ'.*

(ii) *There is a homomorphism $\rho\colon \Gamma \to \mathrm{PGL}(2, \mathbb{F})$ with $\Gamma' = \left\{\, \alpha \in \Gamma \mid \rho(\alpha)\binom{1}{0} \in \left\langle \binom{1}{0} \right\rangle \,\right\}$.*

By our reduction, a number of special cases of the subgroup-forbidden model can be solved via linear matroid parity. Naive applications of Theorems 2.17 and 2.18 lead to deterministic $\mathrm{O}(|E|\cdot|V|^{\omega})$-time and randomized $\mathrm{O}(|E|\cdot|V|^{\omega-1})$-time algorithms, respectively.

Moreover, because of Property 5.4, we can improve the latter bound to $\mathrm{O}(|V|^\omega + |E|)$ by the same argument as [5, Section 5.1.3]. Note that, when Γ' is finite, since $|\Gamma'| + 1$ edges are enough between each pair of two vertices (cf. Section 3.4), we may assume $|E| = \mathrm{O}(|\Gamma'| \cdot |V|^2)$.

On the other hand, as shown in Theorem 3.10, the extension of the algorithm of Chudnovsky et al. [6] for the non-zero model requires $\mathrm{O}(|V|^5 + |E| \cdot |V|)$ time. Compared to this running time bound, a coherent representation, if exists, leads to a much faster algorithm. When $|E| = \mathrm{O}(|V|^\omega)$ in particular, it reduces the computational time bound to less than the square root (from $\mathrm{O}(|V|^5)$ to $\mathrm{O}(|V|^\omega)$).

5.3 Reduction Procedure (Proof of Sufficiency Part)

In this section, we prove that Condition (ii) is sufficient for (i) in Theorem 5.6 by showing a procedure to construct a coherent representation. Let $\rho\colon \Gamma \to \mathrm{PGL}(2, \mathbb{F})$ be a homomorphism with $\Gamma' = \{\, \alpha \in \Gamma \mid \rho(\alpha)Y = Y \,\}$, where $Y := \left\langle \binom{1}{0} \right\rangle$ is a 1-dimensional linear subspace of $\mathbb{F}^2$ spanned by the vector $\binom{1}{0} \in \mathbb{F}^2$. Fix an arbitrary Γ-labeled graph $G = (V, E)$ and an arbitrary terminal set $A \subseteq V$. The construction procedure is shown in an analogous way to Schrijver's reduction of Mader's $\mathcal{S}$-paths problem (Section 5.1).

Associate each arc $e = uw \in E$ to a 2-dimensional linear subspace

$$L_e := \{\, x \in (\mathbb{F}^2)^V \mid \rho(\psi_G(e))x(u) + x(w) = \mathbf{0},\ x(v) = \mathbf{0} \ \ (v \in V \setminus \{u, w\}) \,\} \tag{5.2}$$

of $(\mathbb{F}^2)^V$. For each terminal $a \in A$, define a 1-dimensional linear subspace

$$Q_a := \{\, x \in (\mathbb{F}^2)^V \mid x(a) \in Y,\ x(v) = \mathbf{0} \ \ (v \in V - a)\}$$

of $(\mathbb{F}^2)^V$. Let $Q_{A'} := \sum_{a \in A'} Q_a$ for each subset $A' \subseteq A$, and $Q := Q_A$.

Let $\mathcal{E} := \{\, L_e/Q \mid e \in E \,\}$. Note that $\dim(L_e/Q) = 2$ for every edge $e \in E$, since we may assume that no edge with label in Γ' connects two terminals. Let us construct a matrix $Z = (Z_e)_{e \in E} \in \mathbb{F}^{2V \times 2E}$ associated with $\mathcal{E}$ by enumerating the bases of L_e/Q for all edges $e \in E$, i.e., $Z_e := (b_e, c_e) \in \mathbb{F}^{2V \times 2}$ for each edge $e \in E$, where $\{b_e, c_e\}$ is an arbitrary fixed basis of L_e/Q in $(\mathbb{F}^2)^V/Q$. Then, each edge set $F \subseteq E$ is a matching for Z if and only if $\dim(L_F/Q) = 2|F|$, where $L_F := \sum_{e \in F} L_e$.

The above matrix Z is also obtained as follows: starting with the Kronecker product $Z' := B_G \otimes I_2 \in \mathbb{F}^{2V \times 2E}$ of the incidence matrix $B_G \in \mathbb{F}^{V \times E}$ of G and the 2×2 incidence matrix $I_2 \in \mathbb{F}^{2 \times 2}$, replace $Z'_{w,e} = -I_2$ with $-\rho(\psi_G(e))$ for each arc $e = uw \in E$, and eliminate it by using a basis $x \in (\mathbb{F}^2)^V$ of Q_a for each terminal $a \in A$, which is defined by $x(a) := \binom{1}{0}$ and $x(v) := \mathbf{0}$ for each $v \in V - a$.

Analogously to Theorem 5.1 (cf. Theorem 4.7), the following claim holds. Here, recall that $\nu(Z)$ denotes the maximum cardinality of a matching for Z, and let $\mu(G, A)$ denote the maximum number of vertex-disjoint Γ'-admissible A-paths in G.

Claim 5.7. *Suppose that G is connected and $A \neq \emptyset$. Then, $\mu(G, A) = \nu(Z) - |V| + |A|$.*

In order to prove Claim 5.7, we characterize the matchings for Z analogously to Lemma 4.6. Recall that $\mathrm{comp}(F)$ denotes the partition of an edge set $F \subseteq E$ according to the connected components.

Claim 5.8. *If an edge set $F \subseteq E$ is a matching for Z, then each $F' \in \mathrm{comp}(F)$ satisfies one of the following conditions*:

- $|V(F') \cap A| = 0$ *and $G[F']$ contains at most one cycle*;
- $|V(F') \cap A| \leq 2$, *$G[F']$ contains no cycle, and the A-path between the two terminals is Γ'-admissible if $|V(F') \cap A| = 2$.*

Proof. For distinct $F_1, F_2 \in \mathrm{comp}(F)$, the intersection of the corresponding subspaces is trivial, i.e., $L_{F_1} \cap L_{F_2} = \{\mathbf{0}\}$. Hence, it suffices to prove for a connected edge set $F \subseteq E$, for which we have $|F| \geq |V(F)| - 1$.

Define a linear subspace $X_F := L_F + Q_{A(F)}$ of $(\mathbb{F}^2)^V$, where $A(F) := V(F) \cap A$. Since every $x \in X_F$ has at most $2|V(F)|$ nonzero entries, we have $\dim(X_F) \leq 2|V(F)| \leq 2|F|+2$. By $L_F/Q = L_F/Q_{A(F)}$ and $\dim(Q_{A(F)}) = |A(F)| = |V(F) \cap A|$, we have

$$\begin{aligned}
\dim(L_F/Q) &= \dim(L_F) - \dim(L_F \cap Q_{A(F)}) \\
&= \dim(X_F) - \dim(Q_{A(F)}) \\
&\leq 2|F| + 2 - |V(F) \cap A|.
\end{aligned}$$

Since $\dim(L_F/Q) = 2|F|$ by the assumption, this inequality implies $|V(F) \cap A| \leq 2$.

Suppose that $G[F]$ contains a cycle. Then, by $|V(F)| \leq |F|$, we have $\dim(X_F) \leq 2|V(F)| \leq 2|F|$, which implies $|V(F) \cap A| \leq 0$. Therefore, $G[F]$ contains exactly one cycle C and $|V(F) \cap A| = 0$.

Suppose that $G[F]$ contains no cycle and $|V(F) \cap A| = 2$. Let $P = (v_0, e_1, v_1, \ldots, e_k, v_k)$ be the unique A-path in $G[F]$. Take vectors $x_i \in L_{e_i}$ $(i = 1, 2, \ldots, k)$ so that $x_1(v_0) = \binom{1}{0}$ and $x_{i+1}(v_i) + x_i(v_i) = \mathbf{0}$ for every $i \in [k-1]$. By (5.2), each $x_i \in L_{e_i}$ satisfies $x_i(v_i) = -\rho(\psi_G(e_i))x_i(v_{i-1})$, and hence we have $x_k(v_k) = -\rho(\psi(P))x_1(v_0)$. If $\psi(P) \in \Gamma'$, then $\rho(\psi(P))\binom{1}{0} \in Y$. This implies that $\{\, x_i \in L_{e_i} \mid i \in [k] \,\}$ is linearly dependent in $(\mathbb{F}^2)^V/Q$ as follows, which contradicts that F is a matching for Z:

$$\sum_{i=1}^{k} x_i(v_0) \in Y \setminus \{\mathbf{0}\}, \qquad \sum_{i=1}^{k} x_i(v_k) \in Y, \qquad \sum_{i=1}^{k} x_i(v_j) = \mathbf{0} \quad (j \in [k-1]). \tag{5.3}$$

□

Claim 5.9. *An edge set $F \subseteq E$ is a matching for Z if each $F' \in \mathrm{comp}(F)$ satisfies the second condition in Claim 5.8, i.e., $|V(F') \cap A| \leq 2$, $G[F']$ contains no cycle, and the A-path between the two terminals is Γ'-admissible if $|V(F') \cap A| = 2$.*

Proof. Similarly to the proof of Claim 5.8, we may assume that F is connected. Suppose that there exists an edge set $F \subseteq E$ that satisfies the assumption but is not a matching for Z. Let us take such an edge set F so that $|F|$ is minimized.

Suppose that $G[F]$ has a non-terminal leaf $v \in V \setminus A$. Let $e = vw \in F$ be the incident edge. Since L_e has two degrees of freedom at v-th entry, we have $\dim(L_F/Q) = \dim(L_{F-e}/Q) + 2$, which contradicts the choice of F.

Thus, every leaf is a terminal. Recall that $|V(F) \cap A| \leq 2$. Since any tree has at least two leaves, we have $|V(F) \cap A| = 2$. Hence, F forms a Γ'-admissible A-path $P = (v_0, e_1, v_1, \dots, e_k, v_k)$. Since F is not a matching for Z, there exists a set of vectors $x_i \in L_{e_i}$ $(i = 1, 2, \dots, k)$ such that $\{\, x_i \in L_{e_i} \mid i \in [k] \,\}$ is linearly dependent in $(\mathbb{F}^2)^V/Q$ (satisfying (5.3) in particular). By (5.2) and (5.3), we have $\mathbf{0} \neq x_1(v_0) \in Y$ and $x_k(v_k) = -\rho(\psi_G(P))x_1(v_0) \in Y$, which imply $\psi_G(P) \in \Gamma'$, a contradiction. □

We are now ready to prove Claim 5.7. The proof is analogous to that for Theorem 4.7.

Proof of Claim 5.7. Let us simply denote $\nu := \nu(Z)$ and $\mu := \mu(G, A)$. Let $F \subseteq E$ be a maximum matching for Z. We denote by c_i the number of connected components of $G[F]$ containing exactly $i \in \{0, 1, 2\}$ terminals, where each isolated terminal contributes c_1. By Claim 5.8 and the maximality of F, we may assume $c_2 \leq \mu$ and that each connected component of $G[F]$ contains one or two terminals. Note that, if there exists an edge set $F' \in \mathrm{comp}(F)$ with no terminal in $G[F']$, and hence including a cycle C by the maximality of F, then F remains a matching after replacing an edge $e \in E(C)$ with an edge connecting $G[F' - e]$ and another connected component of $G[F]$. Therefore, $\nu = |F| = |V| - (c_1 + c_2) = |V| - |A| + c_2 \leq |V| - |A| + \mu$.

The converse direction can be easily seen as follows. Let $F \subseteq E$ be the edge set of a maximum number of vertex-disjoint Γ'-admissible A-paths. By Claim 5.9, F is a matching extendable to be a maximal matching $\hat{F}$ so that each connected component of $G[\hat{F}]$ contains one or two terminals, and hence $|V| - |A| + \mu \leq |\hat{F}| \leq \nu$. □

By the same observation as that just after the proof of Theorem 4.7, one can construct a maximum number of vertex-disjoint Γ'-admissible A-paths in G from a maximum matching F for Z by the depth first search from each terminal, which can be done in $\mathrm{O}(|V|)$ time since $G[F]$ is a forest. Thus, we conclude the proof of the sufficiency.

5.4 Extreme Case (Proof of Necessity Part)

To prove the necessity of Condition (ii) for (i) in Theorem 5.6, we construct an extreme example. Let Γ/Γ' denote the left cosets $\{\, \alpha\Gamma' \mid \alpha \in \Gamma \,\}$, for which we denote each element by its representative.

For each $i \in \{1,2,3\}$, let $G_i = (V_i, E_i)$ be a star defined as follows:

$$\begin{aligned} V_i &:= \{v_i\} \cup \{\, u_{ij}^\alpha \mid j \in \{1,2\},\ \alpha \in \Gamma/\Gamma' \,\}, \\ E_i &:= \{\, u_{ij}^\alpha v_i \mid j \in \{1,2\},\ \alpha \in \Gamma/\Gamma' \,\}. \end{aligned}$$

Define a Γ-labeled graph $G = (V, E)$ as follows:

$$\begin{aligned} V &:= V_1 \cup V_2 \cup V_3, \\ E &:= E_1 \cup E_2 \cup E_3 \cup \{\, e_{ij}^\alpha = v_i v_j \mid 1 \le i < j \le 3,\ \alpha \in \Gamma \,\}, \\ \psi_G(u_{ij}^\alpha v_i) &:= \alpha \qquad (i \in \{1,2,3\},\ j \in \{1,2\},\ \alpha \in \Gamma/\Gamma'), \\ \psi_G(e_{ij}^\alpha) &:= \alpha \qquad (1 \le i < j \le 3,\ \alpha \in \Gamma). \end{aligned}$$

Suppose that, for the Γ-labeled graph G and a terminal set $A := V \setminus \{v_1, v_2, v_3\}$, there exists a coherent representation $Z =\in \mathbb{F}^{2V \times 2E}$, i.e., satisfying Properties 5.4 and 5.5. We denote by $Z_{v,e} \in \mathbb{F}^{2\times 2}$ the submatrix of Z corresponding to a vertex $v \in V$ and an edge $e \in E$, and by $b_e, c_e \in (\mathbb{F}^2)^V$ the column vectors corresponding to an edge $e \in E$.

Claim 5.10. *Let $i \in \{1,2,3\}$. For each arc $e = uv_i \in E_i$, we may assume that $b_e(u) = \mathbf{0} \neq c_e(u)$. Moreover, for two distinct arcs $e_1, e_2 \in E_i$ to v_i, $\{b_{e_1}(v_i), b_{e_2}(v_i)\}$ is linearly dependent if and only if $\psi_G(e_1) = \psi_G(e_2)$.*

Proof. Without loss of generality, let us fix $i = 1$. Take three arcs $e_j = u_j v_1 \in E_1$ $(j = 1,2,3)$ with $\alpha := \psi_G(e_1) = \psi_G(e_2) \neq \psi_G(e_3) =: \beta$ (see Fig. 5.1). Then, the two A-paths $(u_j, e_j, v_1, e_3, u_3)$ $(j = 1,2)$ are Γ'-admissible, and $(u_1, e_1, v_1, e_2, u_2)$ is not. Hence, the two edge sets $\{e_j, e_3\}$ $(j = 1,2)$ are matchings for Z, and $\{e_1, e_2\}$ is not by Property 5.5.

The latter implies that the 2×2 submatrix $Z_j := Z_{u_j, e_j}$ is singular for each $j \in \{1,2\}$. Since there exists an arc $e_4 = u_4 v_1 \in E_1$ such that $e_4 \neq e_3$ and $\psi_G(e_4) = \beta$, this holds also for $j = 3$, i.e., $Z_3 := Z_{u_3, e_3}$ is also singular. The former implies that Z_j is not the zero matrix, since otherwise the set $\{\, b_{e_k}, c_{e_k} \mid k \in \{j, 3\} \,\}$ of corresponding vectors is linearly dependent, a contradiction. Thus, we have $\operatorname{rank} Z_j = 1$ $(j = 1,2,3)$, and hence we may assume $b_{e_j}(u_j) = \mathbf{0} \neq c_{e_j}(u_j)$.

By this assumption and Property 5.4, for fixed $i \in \{1,2,3\}$, each edge set $\{e, e'\} \subseteq E_i$ with $e \neq e'$ is a matching for Z if and only if $\{b_e(v_i), b_{e'}(v_i)\}$ is linearly independent. Besides, by Property 5.5, the former is equivalent to the Γ'-admissibility of the A-path consisting of e and e'. □

Claim 5.11. *For each arc $e = v_i v_j \in E$, both $Z_{v_i, e}$ and $Z_{v_j, e}$ are nonsingular.*

Proof. Without loss of generality, let us fix $i = 1$ and $j = 2$. Take four arcs $e_{i1}, e_{k2} \in E_k$ $(k = 1,2)$ to v_k and $e = v_1 v_2 \in E$ with $\psi_G(e) = \beta$ such that $\alpha_1 := \psi_G(e_{11}) \neq \psi_G(e_{12}) =: \alpha_2$, and $\psi_G(e_{21}) \sim \beta\alpha_1 \nsim \beta\alpha_2 \sim \psi_G(e_{22})$ (see Fig. 5.2), where $\sim$ denotes the

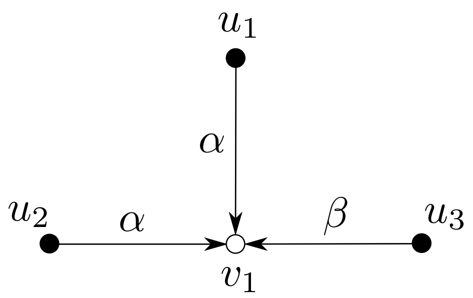

Figure 5.1: Observed part in Claim 5.10.

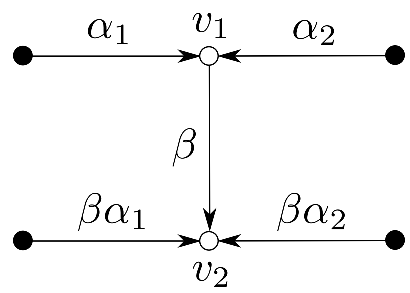

Figure 5.2: Observed part in Claim 5.11.

left equivalence with respect to Γ'. There exist four A-paths consisting of some of the above edges and traversing e. Two of these are Γ'-admissible, and the other two are not.

Since $\{e_{1l}, e\}$ and $\{e, e_{2l}\}$ are matchings for Z and $\{e_{1l}, e, e_{2l}\}$ is not for each $l \in \{1,2\}$, we have $b_{e_{k1}}(v_k), b_{e_{k2}}(v_k) \in \langle b_e(v_k)\rangle + \langle c_e(v_k)\rangle$ for each $k \in \{1,2\}$. By Claim 5.10, $\{b_{e_{k1}}(v_k), b_{e_{k2}}(v_k)\}$ is linearly independent for each $k \in \{1,2\}$. Hence, $\{b_e(v_k), c_e(v_k)\}$ is linearly independent, which implies that $Z_{v_k,e}$ is nonsingular. $\square$

By Claim 5.11, for every three arcs $e_1 = v_1v_2$, $e_2 = v_1v_3$, and $e_3 = v_3v_2$, we may assume that $Z_{v_1,e_1} = Z_{v_1,e_2} = Z_{v_3,e_3} = I_2$, and moreover $Z_{v_2,e_1} \sim Z_{v_3,e_2} \sim Z_{v_2,e_3}$ if $\psi_G(e_1) = \psi_G(e_2) = \psi_G(e_3)$ as follows. Here, $Z_1 \sim Z_2$ for nonsingular $Z_1, Z_2 \in \mathbb{F}^{2\times 2}$ means that $Z_1 = kZ_2$ for some $k \in \mathbb{F} \setminus \{0\}$.

Choose $e_1 = v_1v_2 \in E$ with $\psi_G(e_1) = 1_\Gamma$, and let $B := Z_{v_2,e_1}$. Redefine $b_e(v_2) := B^{-1}b_e(v_2)$ for each edge $e \in E_2$ and $Z_{v_2,e} := B^{-1}Z_{v_2,e}$ for each arc $e = v_iv_2 \in E$ ($i \in \{1,3\}$), and then we have $Z_{v_2,e_1} = I_2$. This redefinition leads to the linear dependence of $\{b_{f_1}(v_1), b_{f_2}(v_2)\}$ for each pair of edges $f_1 \in E_1$ and $f_2 \in E_2$ with $\psi_G(f_1) = \psi_G(f_2)$. Then, by the similar redefinition around v_3, each two-vector subset of $\{\, b_{f_i}(v_i) \mid i \in \{1,2,3\} \,\}$ is linearly dependent for $f_i \in E_i$ ($i = 1,2,3$) with $\psi_G(f_1) = \psi_G(f_2) = \psi_G(f_3)$.

The following claim concludes the proof of the necessity. It should be noted that the 1-dimensional subspace Y of $\mathbb{F}^2$ in Claim 5.12 can be replaced with $\left\langle \binom{1}{0} \right\rangle$ by an appropriate basis transformation, which leads to the same condition in Theorem 5.6.

Claim 5.12. *For each edge $e = v_1v_2 \in E$, let $\rho(\psi_G(e)) := Z_{v_2,e}$. Then, $\rho\colon \Gamma \to \mathrm{PGL}(2,\mathbb{F})$ is homomorphic. Moreover, for $f = uv_1 \in E_1$ with $\psi_G(f) = 1_\Gamma$, $Y := \langle b_f(v_1)\rangle$ satisfies $\Gamma' = \{\, \alpha \in \Gamma \mid \rho(\alpha)Y = Y \,\}$.*

Proof. Take seven arcs $e_{i1}, e_{i2} \in E_i$ ($i = 1,2$), $e_1 = v_1v_2$, $e_2 = v_1v_3$, and $e_3 = v_3v_2$ so that $\beta_1 = \beta_3\beta_2$, $\alpha_1 := \psi_G(e_{11}) \neq \psi_G(e_{12}) =: \alpha_2$, and $\psi_G(e_{21}) \sim \beta_1\alpha_1 \nsim \beta_1\alpha_2 \sim \psi_G(e_{22})$, where $\beta_i := \psi_G(e_i)$ ($i = 1,2,3$) (see Fig. 5.3).

For each $j \in \{1,2\}$, since A-paths formed by $\{e_{1j}, e_1, e_{2j}\}$ and $\{e_{1j}, e_2, e_3, e_{2j}\}$ are not Γ'-admissible, we have $b_{e_{2j}}(v_2) \in \langle \rho(\beta_1)b_{e_{1j}}(v_1)\rangle$ and $b_{e_{2j}}(v_2) \in \langle \rho(\beta_3)\rho(\beta_2)b_{e_{1j}}(v_1)\rangle$. Since $\langle b_{e_{i1}}(v_i)\rangle \neq \langle b_{e_{i2}}(v_i)\rangle$ ($i = 1,2$) by Claim 5.10, $\rho(\beta_1) \sim \rho(\beta_3)\rho(\beta_2)$ holds.

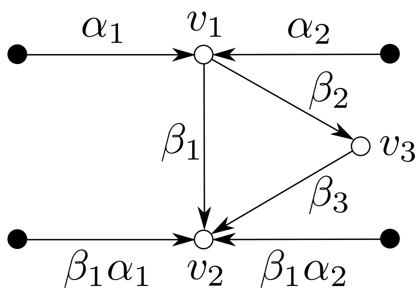

Figure 5.3: Observed part in Claim 5.12.

Suppose $\alpha_1 = 1_\Gamma$ and let $Y := \langle b_{e_{11}}(v_1)\rangle$. Then, $\beta_1 \in \Gamma' \iff \beta_1\alpha_1 \in \Gamma' \iff \psi_G(e_{21}) \sim 1_\Gamma \iff b_{e_{21}}(v_2) \in Y$ by Claim 5.10. Since $\{e_{11}, e_1, e_{21}\}$ is not a matching for Z, we have $b_{e_{21}}(v_2) \in \rho(\beta_1)Y$. Both Y and $\rho(\beta_1)Y$ are 1-dimensional subspaces and $b_{e_{21}}(v_2) \neq \mathbf{0}$, and hence $b_{e_{21}}(v_2) \in Y \iff Y = \rho(\beta_1)Y$. $\square$

5.5 Applications

In this section, we present a variety of important special cases of the subgroup-forbidden model, which admit coherent representations. Almost all following cases satisfies $|\Gamma'| = \mathrm{O}(1)$, and hence the running time of the linear matroid parity algorithm of Cheung et al. with a speeding up argument in [5, Section 5.1.3] is bounded by $\mathrm{O}(|V|^\omega)$, which is much better than $\mathrm{O}(|V|^5)$ derived from the algorithm of Chudnovsky et al. [6].

5.5.1 Infinite cyclic group $\mathbb{Z}$ (including Mader's $\mathcal{S}$-paths)

For the additive group on the set $\mathbb{Z}$ of integers and its trivial subgroup $\{0\}$, we have a desired homomorphism $\rho\colon \mathbb{Z} \to \mathrm{PGL}(2, \mathbb{Q})$ as follows, where $\mathbb{Q}$ denotes the rational field:

$$\rho(k) \;=\; \begin{pmatrix} 1 & 0 \\ i & 1 \end{pmatrix} \quad (i \in \mathbb{Z}). \tag{5.4}$$

Let $\mathcal{S} = \{A_1, A_2, \ldots, A_k\}$ be a partition of the terminal set A. Recall that Mader's $\mathcal{S}$-paths problem is a special case of the subgroup-forbidden model as follows: $\Gamma = (\mathbb{Z}, +)$, $\Gamma' = \{0\}$ and $\psi(e) = i - j$ for each arc $e = uv$ with $u \in A_i$ and $v \in A_j$, where $A_0 := V \setminus A$ and each edge is assumed to be oriented arbitrarily. In this setting, the homomorphism ρ defined by (5.4) leads to a coherent representation over $\mathbb{Q}$, which coincides with Schrijver's one by appropriate basis transformations.

5.5.2 Finite cyclic groups $\mathbb{Z}_n$ (including odd-length A-paths)

For the cyclic group $\mathbb{Z}_n := \mathbb{Z}/n\mathbb{Z}$ of degree $n \geq 2$ and its trivial subgroup $\{0\}$ (which is essentially equivalent to the setting in which $\Gamma = \mathbb{Z}$ and $\Gamma' = n\mathbb{Z}$), we have a desired homomorphism $\rho\colon \mathbb{Z}_n \to \mathrm{PGL}(2, \mathbb{R})$ as follows, where $\mathbb{R}$ denotes the real field:

$$\rho(i) \;=\; \begin{pmatrix} \cos\dfrac{k\pi}{n} & -\sin\dfrac{k\pi}{n} \\ \sin\dfrac{k\pi}{n} & \cos\dfrac{k\pi}{n} \end{pmatrix} \quad (i \in \mathbb{Z}_n)\,.$$

When $n = 2$ in particular, the following homomorphism $\rho'\colon \mathbb{Z}_2 \to \mathrm{PGL}(2, \mathbb{F})$ is also available, where $\mathbb{F}$ is an arbitrary field:

$$\rho'(0) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \rho'(1) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{5.5}$$

Recall that the packing odd-length A-paths problem is a special case of the subgroup-forbidden model as follows: $\Gamma = \mathbb{Z}_2$, $\Gamma' = \{0\}$, and $\psi(e) = 1$ for each edge e. This case is solvable via the linear matroid parity problem over an arbitrary field with the aid of the homomorphism ρ' defined as (5.5).

5.5.3 Dihedral groups D_n

Even when Γ is non-abelian, there exists a solvable case. Let D_n be the dihedral group of degree $n \geq 3$, i.e., $D_n = \langle\, r, R \mid r^n = R^2 = \mathrm{id},\; rR = Rr^{n-1} \,\rangle$, and Γ' its proper subgroup.

If $\Gamma' \subseteq \langle r \rangle$, then Γ' is generated only by r^m for some divisor m of n (possibly $m = n$). In this case, Γ' is normal and Γ/Γ' is isomorphic to the dihedral group D_m of degree m. Hence, we may assume that Γ' is trivial (i.e., $m = n$). Then, there exists a desired homomorphism $\rho\colon D_n \to \mathrm{PGL}(2, \mathbb{R})$ defined as follows:

$$\rho(r^i R^j) \;=\; \begin{pmatrix} \cos\dfrac{i\pi}{n} & -\sin\dfrac{i\pi}{n} \\ \sin\dfrac{i\pi}{n} & \cos\dfrac{i\pi}{n} \end{pmatrix} \begin{pmatrix} \cos\dfrac{\pi}{n+1} & \sin\dfrac{\pi}{n+1} \\ \sin\dfrac{\pi}{n+1} & -\cos\dfrac{\pi}{n+1} \end{pmatrix}^{j} \qquad \begin{pmatrix} 0 \leq i \leq n-1 \\ j \in \{0,1\} \end{pmatrix}.$$

Otherwise, we may assume that $R \in \Gamma'$ without loss of generality, and then Γ' is generated by R and r^m for some divisor m of n (possibly $m = n$, but $m \neq 1$). Let us take m minimum. Note that Γ' is not normal unless $m = 2$. In this case, the following homomorphism $\rho'\colon D_n \to \mathrm{PGL}(2, \mathbb{R})$ leads to a coherent representation:

$$\rho'(r^i R^j) \;:=\; \begin{pmatrix} \cos\dfrac{i\pi}{m} & -\sin\dfrac{i\pi}{m} \\ \sin\dfrac{i\pi}{m} & \cos\dfrac{i\pi}{m} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}^{j} \qquad \begin{pmatrix} 0 \leq i \leq n-1 \\ j \in \{0,1\} \end{pmatrix}.$$

Note that, in the latter case, the size of Γ' is not necessary to be constant and hence the running time bound of the algorithm of Cheung et al. depends on the number of edges, which is polynomially bounded by $|V|$ and $|\Gamma'| = 2n/m$.

5.5.4 Non-returning model

Recall that the non-returning model is an equivalent formulation of the subgroup-forbidden model (see Section 3.4), and in particular the former reduces to the latter as follows. Let Γ be the symmetric group S_n of degree $n \geq 2$, and $\Gamma' := \{\, \sigma \in S_n \mid \sigma(n) = n \,\} = S_{n-1}$. If $n \geq 3$, then $\Gamma' = S_{n-1}$ is not a normal subgroup of Γ. If $n = 2, 3$, then it reduces to previous examples by $S_2 \simeq \mathbb{Z}/2\mathbb{Z}$ and $S_3 \simeq D_3$. We shall clarify the case of $n \geq 4$.

Theorem 5.13. *The subgroup-forbidden model reduced from the non-returning model with the label set Ω admits a coherent representation if and only if $|\Omega| \leq 4$.*

Proof. Suppose that $n = |\Omega| \geq 4$ and there exists a desired homomorphism ρ. Then, ρ is faithful, i.e., it is isomorphic. To see this, suppose to the contrary that the kernel of ρ is not trivial, i.e., $\Gamma'' := \{\, \sigma \in \Gamma \mid \rho(\sigma) = I_2 \,\}$ contains a non-identity permutation. By the basic fact in group theory, the kernel Γ'' is a normal subgroup of Γ. On the other hand, S_{n-1} is not a normal subgroup of S_n if $n \geq 3$, since every permutation in S_{n-1} fixes n but some in $(k\ n)S_{n-1}(k\ n)$ does not for any $k \in [n-1]$.

Since $\Gamma' = \{\, \sigma \in \Gamma \mid \rho(\sigma)\binom{1}{0} = \binom{1}{0} \,\}$, for every $\sigma \in \Gamma'$, ρ has the following form:

$$\rho(\sigma) \;=\; \begin{pmatrix} 1 & a_\sigma \\ 0 & b_\sigma \end{pmatrix},$$

where $a_\sigma, b_\sigma \in \mathbb{F}$. Then, it is seen as follows that the characteristic of $\mathbb{F}$ is 3.

Let $p := a_{(1\ 2)},\ q := b_{(1\ 2)},\ r := a_{(1\ 2\ 3)},\ s := b_{(1\ 2\ 3)}$. Because of $(1\ 2)(1\ 2\ 3) = (2\ 3)$, $(1\ 2\ 3)(1\ 2) = (1\ 3)$, $(1\ 2\ 3)^2 = (1\ 3\ 2)$, and $(1\ 2)^2 = (2\ 3)^2 = (1\ 3)^2 = (1\ 2\ 3)^3 = (1\ 3\ 2)^3 =$ id, we have

$$p(q+1) = (r+ps)(1+qs) = (p+qr)(1+qs) = r(s^2+s+1) = 0,$$
$$q^2 = q^2 s^2 = s^3 = 1. \tag{5.6}$$

Hence, we have $s = 1$, and $r(1+1+1) = 0$. If $r = 0$, then $\rho\big((1\ 2\ 3)\big) = I_2$, contradicting the faithfulness of ρ. Thus we have $1+1+1 = 0$.

By (5.6), we have $q = \pm 1$. If $q = 1$, then $2p = 0$, which implies $p = 0$ and $\rho\big((1\ 2)\big) = I_2$, contradicting the faithfulness of ρ. Thus we have $q = -1 = 2$.

Now we have the following representation:

$$\rho\big((1\ 2)\big) = \begin{pmatrix} 1 & p \\ 0 & -1 \end{pmatrix}, \quad \rho(\mathrm{id}) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

$$\rho\big((1\ 3)\big) = \begin{pmatrix} 1 & p-r \\ 0 & -1 \end{pmatrix}, \quad \rho\big((1\ 2\ 3)\big) = \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix},$$

$$\rho\big((2\ 3)\big) = \begin{pmatrix} 1 & p+r \\ 0 & -1 \end{pmatrix}, \quad \rho\big((1\ 3\ 2)\big) = \begin{pmatrix} 1 & -r \\ 0 & 1 \end{pmatrix},$$

$$\rho\big((1\ 4)\big) = \begin{pmatrix} w_1 & x_1 \\ y_1 & z_1 \end{pmatrix}, \quad \rho\big((2\ 4)\big) = \begin{pmatrix} w_2 & x_2 \\ y_2 & z_2 \end{pmatrix}.$$

Since ρ is faithful, we have $p \neq p - r \neq p + r \neq p$, and hence at least two of $p_1 := p + r$, $p_2 := p - r$, and $p_3 := p$ are nonzero. By symmetry, without loss of generality, we may assume $p_1 \neq 0 \neq p_2$.

For each $i \in \{1, 2\}$, since $(j\ k)(i\ 4) = (i\ 4)(j\ k)$ ($\{j, k\} = \{1, 2, 3\} - i$) and $(i\ 4)^2 = \mathrm{id}$,

$$\begin{aligned} \begin{pmatrix} w_i + p_i y_i & x_i + p_i z_i \\ -y_i & -z_i \end{pmatrix} &= l_{i1} \begin{pmatrix} w_i & p_i w_i - x_i \\ y_i & p_i y_i - z_i \end{pmatrix}, \\ \begin{pmatrix} w_i^2 + x_i y_i & x_i(w_i + z_i) \\ y_i(w_i + z_i) & x_i y_i + z_i^2 \end{pmatrix} &= l_{i2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \end{aligned}$$

where $l_{i1}, l_{i2} \in \mathbb{F}$. If $n \geq 5$, then we have $y_i = 0 \neq w_i$ by $(i\ 4) \in \Gamma'$, and hence $l_{i1} = 1$. If $x_i = 0$, then $w_i = z_i$ by $p_i \neq 0$, i.e., $\rho\big((i\ 4)\big) = w_i I_2 \sim \rho(\mathrm{id})$, a contradiction. Otherwise, by $x_i(w_i + z_i) = 0$, we have $z_i = -w_i$, and hence $x_i - p_i w_i = p_i w_i - x_i$. This implies $x_i - p_i w_i = 0$, i.e., $\rho\big((j\ k)(i\ 4)\big) = w_i I_2 \sim \rho(\mathrm{id})$, a contradiction. Thus we have proved that the case of $n \geq 5$ does not admit a coherent representation.

If $n = 4$, then we have $y_i \neq 0$ by $(i\ 4) \notin \Gamma'$, and hence $l_{i1} = -1$ and $z_i = -w_i$. Therefore, we have the following equations:

$$w_i + p_i y_i = -w_i, \quad x_i - p_i w_i = -p_i w_i + x_i, \quad w_i = -p_i y_i - w_i.$$

The first and third equations imply $w_i = p_i y_i$, and the second holds obviously.

By patient but straightforward calculation, we get a desired projective representation ρ over $\mathbb{F} := \mathbb{F}_3 = \mathbb{Z}/3\mathbb{Z}$ as shown in Claim 5.14, which is an isomorphism to $\mathrm{PGL}(2, \mathbb{F}_3)$ but not to $\mathrm{GL}(2, \mathbb{F}_3)$. The correctness can be easily confirmed by checking $A^2 = B^2 = C^2 = I_2$, $AC = CA$, $ABA = BAB$, $BCB = CBC$, where $A := \rho(\sigma_{12})$, $B := \rho(\sigma_{23})$, $C := \rho(\sigma_{34})$, and $\sigma_{ij} := (i\ j)$, based on the basic fact (5.7) in group theory for generating S_4.

$$S_4 = \left\langle \sigma_{12},\, \sigma_{23},\, \sigma_{34} \;\middle|\; \begin{array}{l} \sigma_{12}^2 = \sigma_{23}^2 = \sigma_{34}^2 = \mathrm{id},\ \sigma_{12}\sigma_{34} = \sigma_{34}\sigma_{12}, \\ \sigma_{12}\sigma_{23}\sigma_{12} = \sigma_{23}\sigma_{12}\sigma_{23},\ \sigma_{23}\sigma_{34}\sigma_{23} = \sigma_{34}\sigma_{23}\sigma_{34} \end{array} \right\rangle. \tag{5.7}$$

Claim 5.14. *A mapping $\rho\colon S_4 \to \mathrm{PGL}(2, \mathbb{F}_3)$ defined as follows is a homomorphism with $S_3 = \left\{\, \sigma \in S_4 \mid \rho(\sigma)\binom{1}{0} = \binom{1}{0} \,\right\}$:*

$$\begin{aligned}
\rho(\mathrm{id}) &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, & \rho\big((1\ 2\ 3)\big) &= \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, & \rho\big((1\ 3\ 2)\big) &= \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}, \\
\rho\big((1\ 2)\big) &= \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, & \rho\big((2\ 3)\big) &= \begin{pmatrix} 1 & 1 \\ 0 & 2 \end{pmatrix}, & \rho\big((1\ 3)\big) &= \begin{pmatrix} 1 & 2 \\ 0 & 2 \end{pmatrix}, \\
\rho\big((1\ 4)\big) &= \begin{pmatrix} 1 & 0 \\ 1 & 2 \end{pmatrix}, & \rho\big((2\ 4)\big) &= \begin{pmatrix} 2 & 0 \\ 1 & 1 \end{pmatrix}, & \rho\big((3\ 4)\big) &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \\
\rho\big((1\ 2)(3\ 4)\big) &= \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}, & \rho\big((1\ 3)(2\ 4)\big) &= \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}, & \rho\big((1\ 4)(2\ 3)\big) &= \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}, \\
\rho\big((1\ 2\ 4)\big) &= \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, & \rho\big((1\ 4\ 2)\big) &= \begin{pmatrix} 2 & 0 \\ 1 & 2 \end{pmatrix}, & \rho\big((1\ 3\ 4)\big) &= \begin{pmatrix} 1 & 2 \\ 1 & 0 \end{pmatrix}, \\
\rho\big((1\ 4\ 3)\big) &= \begin{pmatrix} 0 & 2 \\ 1 & 2 \end{pmatrix}, & \rho\big((2\ 3\ 4)\big) &= \begin{pmatrix} 2 & 2 \\ 1 & 0 \end{pmatrix}, & \rho\big((2\ 4\ 3)\big) &= \begin{pmatrix} 0 & 2 \\ 1 & 1 \end{pmatrix}, \\
\rho\big((1\ 3\ 2\ 4)\big) &= \begin{pmatrix} 1 & 2 \\ 1 & 1 \end{pmatrix}, & \rho\big((1\ 4\ 2\ 3)\big) &= \begin{pmatrix} 2 & 2 \\ 1 & 2 \end{pmatrix}, & \rho\big((1\ 2\ 3\ 4)\big) &= \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}, \\
\rho\big((1\ 4\ 3\ 2)\big) &= \begin{pmatrix} 0 & 1 \\ 1 & 2 \end{pmatrix}, & \rho\big((1\ 3\ 4\ 2)\big) &= \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix}, & \rho\big((1\ 2\ 4\ 3)\big) &= \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}.
\end{aligned}$$

$\square$

Chapter 6

Shortest Disjoint A-paths via Weighted Matroid Matching

In this chapter, we extend the previous two reductions to a weighted setting in which we minimize the total length of a designated number of vertex-disjoint A-paths. We reduce this problem to weighted versions of matroid matching and linear matroid parity. In particular, a reduction to the weighted linear matroid parity problem leads to the efficient solvability by weighted linear matroid parity algorithms due to Iwata [40] and Pap [75].

This chapter is based on [93] and organized as follows. In Section 6.1, we give an overview of problems and results. Section 6.2 is devoted to presenting a key trick to extend our reductions, and the correctness of extended reductions are shown in Section 6.3. Finally, some remarks for applications of our reductions conclude this chapter in Section 6.4.

6.1 Overview

In this section, we focus on a weighted setting of packing non-zero A-paths as follows: for a given Γ-labeled graph $G = (V, E)$, terminal set $A \subseteq V$, nonnegative edge length $\ell \in \mathbb{R}^E_{\geq 0}$, and positive integer $k \in \mathbb{Z}_{>0}$, to find a family $\mathcal{P}$ of k vertex-disjoint non-zero A-paths in G with the total length $\ell(\mathcal{P}) := \sum_{P \in \mathcal{P}} \sum_{e \in E(P)} \ell_e$ minimum.

<u>**Shortest Disjoint Non-zero A-paths Problem**</u>

Input: A Γ-labeled graph $G = (V, E)$, a terminal set $A \subseteq V$, a nonnegative edge length $\ell \in \mathbb{R}^E_{\geq 0}$, and a positive integer $k \in \mathbb{Z}_{>0}$.

Goal: Find a family $\mathcal{P}$ of k vertex-disjoint non-zero A-paths in G with $\ell(\mathcal{P})$ minimum.

Even for such a weighted version of Mader's $\mathcal{S}$-paths problem (a special case of this problem), any polynomial-time algorithm was not known, while Karzanov [46] had shown one for a similar problem in the edge-disjoint A-paths setting (which is a further special case of Mader's setting), whose full proof had been left to an unpublished paper [45]. Karzanov's problem can be solved by finding shortest k vertex-disjoint $\mathcal{S}$-paths for possible k, where the number of iterations is at most $|A|/2$ and can be reduced to $\mathrm{O}(\log |A|)$ by binary search.

It should be remarked that Hirai and Pap [33] discussed a generalization of Karzanov's setting, in which each pair of two terminals has weight. Also, Pap [74] dealt with a weighted version of Mader's $\mathcal{S}$-paths problem in which weight is defined only on terminal pairs (no edge length or cost).

In Chapter 4, we show a reduction of the packing non-zero A-paths problem to the matroid matching problem. Furthermore, in Chapter 5, we clarify a necessary and sufficient condition for the groups in question to admit a reasonable reduction of the subgroup-forbidden model to the linear matroid parity problem. These results can be extended to the weighted settings, using a trick shown in the next section.

In particular, the following theorem can be derived from a reduction of packing non-zero A-paths to matroid matching shown in Section 4.3.

Theorem 6.1. *The shortest disjoint non-zero A-paths problem reduces to the weighted matroid matching problem.*

In the same way, the next theorem can be obtained from a more direct reduction to linear matroid parity shown in Section 5.3. Recall that $\mathrm{PGL}(n, \mathbb{F}) = \mathrm{GL}(n, \mathbb{F})/\{\, kI_n \mid k \in \mathbb{F} \setminus \{0\}\,\}$, where $\mathrm{GL}(n, \mathbb{F})$ denotes the general linear group of degree n over a field $\mathbb{F}$, $I_n \in \mathrm{GL}(n, \mathbb{F})$ the $n \times n$ identity matrix, and we denote by $\langle y \rangle$ the 1-dimensional subspace spanned by a vector $y \in \mathbb{F}^2$.

Theorem 6.2. *Let Γ be a group and $\mathbb{F}$ a field. If there exists a homomorphism $\rho\colon \Gamma \to \mathrm{PGL}(2, \mathbb{F})$ such that $\rho(\alpha)\binom{1}{0} \notin \left\langle \binom{1}{0} \right\rangle$ for every $\alpha \in \Gamma \setminus \{1_\Gamma\}$, then the shortest disjoint non-zero A-paths problem in Γ-labeled graphs reduces to the weighted linear matroid parity problem over $\mathbb{F}$.*

Remark. These theorems can be extended to the weak triple exchange model and the subgroup-forbidden model, respectively, by simple application of the original theorems (Theorems 4.14 and 5.6).

Recall that the weighted matroid matching problem is to find a minimum-weight parity base as follows, where a subset $B \subseteq E$ is called a parity base if B is a full-rank matching, i.e., $f(B) = 2|B| = f(E)$ or $\operatorname{rank} Z(B) = 2|B| = \operatorname{rank} Z$.

Weighted Matroid Matching Problem

Input: A 2-polymatroid (E, f) and a weight $w \in \mathbb{R}^E$.

Goal: Find a parity base $B \subseteq E$ in (E, f) such that $w(B)$ is minimum.

Weighted Linear Matroid Parity Problem

Input: A finite set E, a matrix $Z \in \mathbb{F}^{r \times 2E}$ over a field $\mathbb{F}$, and a weight $w \in \mathbb{R}^E$, where $r \in \mathbb{Z}_{>0}$.

Goal: Find a parity base $B \subseteq E$ for Z such that $w(B)$ is minimum.

Iwata [40] and Pap [75] announced that one can solve the weighted linear matroid parity problem in polynomial time. In particular, Iwata's algorithm is based on the linear matroid parity algorithm of Gabow and Stallmann [22], and has the same running time bound $\mathrm{O}(r^3|E|)$ (cf. Theorem 2.17).

6.2 Construction of Auxiliary Graph

First we construct a common auxiliary Γ-labeled graph from a given Γ-labeled graph $G = (V, E)$ and terminal set $A \subseteq V$. Without loss of generality, we assume that G is connected, and $|A| \geq 2k$ (since otherwise there cannot be a feasible solution).

The construction is summarized as follows. Add $|A| - 2k$ extra terminals so that each extra terminal is adjacent to every original terminal by an edge with an arbitrary non-zero label. Besides, add two other extra terminals b_1, b_2 so that b_1 and b_2 are adjacent by a non-zero edge and b_1 is adjacent to all original non-terminals.

Formally, for the vertex set, let a_i $(i = 1, 2, \ldots, |A| - 2k)$ and b_j $(j = 1, 2)$ be distinct vertices not in V, and define $A_1 := \{\, a_i \mid i = 1, 2, \ldots, |A| - 2k \,\}$, $A_2 := \{b_1, b_2\}$, $V' := V \cup A_1 \cup A_2$, and $A' := A \cup A_1 \cup A_2$. Next, for the arc set, let $E_1 := \{\, e_{it} = a_i t \mid a_i \in A_1,\ t \in A \,\}$, $E_2 := \{\, e_v = b_1 v \mid v \in V \setminus A \,\}$, and $E' := E \cup E_1 \cup E_2 \cup \{e' = b_1 b_2\}$. Finally, for the label function of $G' := (V', E')$, extend $\psi_G \colon E \to \Gamma$ to $\psi_{G'} \colon E' \to \Gamma$ as follows: for each arc $e \in E'$,

$$\psi_{G'}(e) := \begin{cases} \psi_G(e) & (e \in E), \\ \alpha & (e \in E' \setminus E), \end{cases}$$

where $1_\Gamma \neq \alpha \in \Gamma$.

For the resulting graph $G' = (V', E')$ with terminal set $A' \subseteq V'$, define a weight vector $w \in \mathbb{R}^{E'}$ as follows: for each $e \in E'$,

$$w_e := \begin{cases} \ell_e & (e \in E), \\ 0 & (e \in E' \setminus E). \end{cases} \tag{6.1}$$

Note that to minimize the total weight makes incentive to take edges in $E' \setminus E$ rather than in E, since $\ell_e \geq 0$ for every $e \in E$. Then, the constructions of an associated 2-polymatroid and a coherent representation shown in Sections 4.3 and 5.3, respectively, complete reductions of the shortest disjoint non-zero A-paths problem to the weighted matroid matching problem and the weighted linear matroid parity problem.

6.3 Correctness of Reductions

6.3.1 Weighted matroid matching case (proof of Theorem 6.1)

We denote the associated 2-polymatroid by (E', f). Then, the following property follows from Lemma 4.5.

Claim 6.3. $f(E') = 2|V'| - |A'| = 2(|V| - k + 1)$.

The following claim just rewrites Lemma 4.6. Recall that $\mathrm{comp}(F)$ denotes the partition of an edge set $F \subseteq E$ according to the connected components of $G[F]$.

Claim 6.4. *An edge set $F \subseteq E'$ is a matching in (E', f) (i.e., $f(F) = 2|F|$) if and only if*

- *$G[F]$ contains no cycle, and*
- *for each $F' \in \mathrm{comp}(F)$, we have $|V(F') \cap A'| \leq 2$ and the A'-path between the two terminals is non-zero if $|V(F') \cap A'| = 2$.*

Suppose that G contains k vertex-disjoint non-zero A-paths. By extending the edge set of such paths using edges in $E_1 \cup E_2 \cup \{e'\}$, we can obtain an edge set $F \subseteq E'$ such that $G'[F]$ is a spanning forest (i.e., $V(F) = V'$ and F contains no cycle), and for each $F' \in \mathrm{comp}(F)$, $|V(F') \cap A'| = 2$ and the A'-path between the two terminals is non-zero. This is because each unused terminal can be connected to an extra terminal in A_1 by an edge in E_1 (and vice versa), each unused non-terminal can be connected to the extra terminal $b_1 \in A_2$ by an edge in E_2, and $b_1, b_2 \in A_2$ are adjacent by a non-zero edge e'. Then, the number of connected components of $G'[F]$ is $k + (|A| - 2k) + 1 = |A| - k + 1$, and hence, by Claims 6.3 and 6.4, we have

$$f(F) = 2|F| = 2\left(|V'| - (|A| - k + 1)\right) = 2(|V| - k + 1) = f(E'),$$

which means that F is a parity base in (E', f). Therefore, for each family $\mathcal{P}$ of k vertex-disjoint non-zero A-paths in G, there exists a parity base $F_{\mathcal{P}}$ in (E', f) with $w(F_{\mathcal{P}}) = \ell(E(\mathcal{P}))$ (recall the definition (6.1) of the weight $w \in \mathbb{R}^{E'}$).

To the contrary, for each parity base F in (E', f), there exists a family $\mathcal{P}_F$ of k vertex-disjoint non-zero A-paths in G with $E(\mathcal{P}_F) \subseteq F$ (hence, we have $\ell(E(\mathcal{P}_F)) \leq w(F)$) as follows. Thus we have done, i.e., shortest k vertex-disjoint non-zero A-paths in G can be obtained by finding a minimum-weight parity base in (E', f).

Claim 6.5. *For a parity base F in (E', f), the subgraph $G'[F]$ is a spanning forest consisting of $|A| - k + 1$ connected components, each of which contains exactly one non-zero A'-path.*

Proof. The first condition in Claim 6.4 implies that F contains no cycle. Since $|A'| = |A| + (|A| - 2k) + 2 = 2(|A| - k + 1)$ and each $F' \in \mathrm{comp}(F)$ intersects at most two terminals in A' by the second condition in Claim 6.4, there are at least $|A| - k + 1$ connected components in $G'[F]$. Hence, we have

$$|F| \leq |V'| - (|A| - k + 1) = |V| - k + 1. \tag{6.2}$$

Recall that $2|F| = f(F) = f(E') = 2(|V| - k + 1)$ by Claim 6.3, which implies that the equality holds in (6.2). This means that $G'[F]$ has exactly $|A| - k + 1$ connected components, each of which contains exactly two terminals in A' and the A'-path between the two terminals is non-zero by the second condition in Claim 6.4. Then, F is obviously spanning. ☐

Claim 6.6. *For a parity base F in (E', f), the subgraph $G[F \cap E]$ contains exactly k vertex-disjoint non-zero A-paths in G.*

Proof. By Claim 6.5, there are $|A| - k + 1$ connected components in $G'[F]$ each of which contains exactly one non-zero A-paths. Since there is only one edge $e' = b_1 b_2 \in E'$ incident to the extra terminal $b_2 \in A_2$, the connected component containing b_2 must contain $b_1 \in A_2$. Besides, since each edge in E' incident to each extra terminal $a_i \in A_1$ ends an original terminal in A, the connected component containing a_i must contain some original terminal in A. The number of such connected components is $|A_1| = |A| - 2k$, and hence the number of the connected components containing non-zero A-paths in G is $|A| - k + 1 - (|A| - 2k + 1) = k$. ☐

6.3.2 Weighted linear matroid parity case (proof of Theorem 6.2)

We denote by $Z \in \mathbb{F}^{2V \times 2E}$ the coherent representation constructed in Section 5.3. Then, the following claim follows from Claims 5.8 and 5.9.

Claim 6.7. *An edge set $F \subseteq E'$ is a matching for Z (i.e., $\operatorname{rank} Z(F) = 2|F|$) if each $F' \in \operatorname{comp}(F)$ satisfies the following condition* (b) *or* (c)*, and only if* (a), (b)*, or* (c):

(a) $|V(F') \cap A'| = 0$ *and $G'[F']$ contains exactly one cycle*;

(b) $|V(F') \cap A'| \leq 1$ *and $G'[F']$ contains no cycle*;

(c) $|V(F') \cap A'| = 2$, *$G'[F']$ contains no cycle, and the A'-path between the two terminals is non-zero.*

This claim implies the following claim.

Claim 6.8. *An edge set $F \subseteq E'$ is a parity base for Z if F is spanning in G' with each $F' \in \operatorname{comp}(F)$ satisfying Condition* (c) *in Claim* 6.7*, and only if F is spanning in G' with* (a) *or* (c).

Proof. By the construction of Z, we have

$$\operatorname{rank} Z \leq 2|V'| - |A'| = 2(|V| - k + 1).$$

Suppose that G contains k vertex-disjoint non-zero A-paths. By extending the edge set of such paths using edges in $E_1 \cup E_2 \cup \{e'\}$, we can obtain a spanning forest $F \subseteq E'$ with each $F' \in \mathrm{comp}(F)$ satisfying Condition (c) (cf. the argument just after Claim 6.4). Since the number of connected components of $G'[F]$ is $k + (|A| - 2k) + 1 = |A| - k + 1$, by the "if" part of Claim 6.7, we have

$$\operatorname{rank} Z(F) = 2|F| = 2\left(|V'| - (|A| - k + 1)\right) = 2(|V| - k + 1) \geq \operatorname{rank} Z \geq \operatorname{rank} Z(F).$$

Hence, $\operatorname{rank} Z = 2(|V| - k + 1)$, and the "if" part follows from Claim 6.7.

The converse direction is also derived from Claim 6.7. Note that, for a parity base $F \subseteq E'$ for Z, there are at most $|A| - k + 1$ connected components in $G'[F]$ that contain no cycle because of the rank, and hence there cannot be a connected component of type (b). □

By the same discussion as that between Claims 6.4 and 6.5 in the previous section, shortest k vertex-disjoint non-zero A-paths in G can be obtained by finding a minimum-weight parity base for Z. Note again that, we have incentive to take extra edges in $E' \setminus E$ rather than original edges in E because of the definition (6.1) of the weight w, and hence the total length of any connected component of type (a) in Claim 6.7 in a minimum-wight parity base for Z is 0.

6.4 Remarks for Applications

As shown in Section 5.5, our reduction of the shortest disjoint non-zero A-paths problem (as well as its subgroup-forbidden extension) to the weighted linear matroid parity problem is applicable for various settings. In those cases, the problem can be solved in polynomial time with the aid of the weighted linear matroid parity algorithms thanks to Iwata [40] and Pap [75].

In particular, we should remark the case when Γ is a finitely generated abelian group and $k = 1$. By the fundamental theorem of finitely generated abelian groups, Γ is decomposed into (finite or infinite) cyclic groups, and suppose that we are given Γ as the direct product of p cyclic groups. In this case, one can find a shortest non-zero A-path by solving the weighted linear matroid parity problem repeatedly p times, since any cyclic group with its trivial subgroup admits a coherent representation (see Sections 5.5.1 and 5.5.2).

While Kobayashi and Toyooka [50] gave an algebraic algorithm for the case when $k = 1$, $|A| = 2$ (i.e., finding a shortest non-zero s–t path), and Γ is a finite abelian group inspired by the work of Björklund and Husfeldt [3] (which gave an algebraic algorithm for finding shortest 2-disjoint paths in an undirected graph), our result provides a combinatorial solution to a more general case. As shown in the beginning of Chapter 7, 2-disjoint paths in an undirected graph correspond to an s–t path of a designated label (or with two labels forbidden) in an $\mathbb{Z}_3$-labeled graph, and it is unknown whether one can find such a shortest s–t path in polynomial time or not.

Chapter 7

Finding an s–t Path with Two Labels Forbidden

In this chapter, we investigate the problem of finding an s–t path with two labels forbidden. This problem generalizes the 2-disjoint paths problem in undirected graphs as follows: as shown in Section 1.4.2, the k-disjoint paths problem can be formulated as finding a zero s–t path in A_{2k-1}-labeled graphs; when $k = 2$, the alternating group A_3 contains exactly three element (which is isomorphic to the cyclic group $\mathbb{Z}_3 = \mathbb{Z}/3\mathbb{Z}$), and hence finding a zero path is equivalent to finding a path with two labels forbidden. Our algorithm and characterization for this problem are strongly inspired by those for the 2-disjoint paths problem [80, 81, 84].

This chapter is based on [49] and organized as follows. Section 7.1 is devoted to presenting our results: the efficient solvability of the problem to find an s–t path with two labels forbidden, and a characterization of group-labeled graphs with exactly two possible labels of s–t paths. Their verifications are shown in Sections 7.2 and 7.3–7.4, which provide a concrete description of a polynomial-time algorithm for the problem with its correctness and a proof of our characterization, respectively.

7.1 Overview of Results

7.1.1 Preliminaries

We focus on s–t paths in this chapter, and hence a vertex that is not contained in any s–t path is redundant. To consider Γ-labeled graphs with no such redundant vertex, let us define the set $\mathcal{D}$ of all triplets (G, s, t) such that G is a Γ-labeled graph with $s, t \in V(G)$ in which every vertex is contained in some s–t path. In addition, let $l(G; s, t)$ denote the set of possible labels of s–t paths in a Γ-labeled graph G. Then, Proposition 2.8 is restated as follows.

Proposition 7.1. *For any triplet $(G, s, t) \in \mathcal{D}$, $|l(G; s, t)| = 1$ if and only if G is balanced.*

By Proposition 2.1, one can test whether a given triplet (G, s, t) is in $\mathcal{D}$ or not in polynomial time by decomposing $G + r + rs + rt$ into the 2-connected components (e.g., by [35]). Note that, if $(G, s, t) \notin \mathcal{D}$, then one obtains the maximal subgraph of G with no redundant vertex. Therefore, Propositions 2.5 and 7.1 lead to the following proposition.

Proposition 7.2. *Let G be a Γ-labeled graph with two distinct vertices $s, t \in V(G)$. Then, for any element $\alpha \in \Gamma$, one can test whether $l(G; s, t) \subseteq \{\alpha\}$ or not in polynomial time. Furthermore, if $l(G; s, t) \not\subseteq \{\alpha\}$, then one can find an s–t path P with $\psi_G(P) \neq \alpha$ in polynomial time.*

7.1.2 Algorithmic result

By Proposition 7.2, one can find a non-zero s–t path in a given Γ-labeled graph in polynomial time. The following theorem, one of our main results, is the first nontrivial extension of this property, which claims that not only one label but also another can be forbidden simultaneously.

Theorem 7.3. *Let G be a Γ-labeled graph with two distinct vertices $s, t \in V(G)$. Then, for any distinct elements $\alpha, \beta \in \Gamma$, one can test whether $l(G; s, t) \subseteq \{\alpha, \beta\}$ or not in polynomial time. Furthermore, if $l(G; s, t) \not\subseteq \{\alpha, \beta\}$, then one can find an s–t path P with $\psi_G(P) \notin \{\alpha, \beta\}$ in polynomial time.*

Such an algorithm is constructed based on characterizations of Γ-labeled graphs with exactly two possible labels of s–t paths, which are shown in Section 7.1.3. Our algorithm and a proof of this theorem are presented in Section 7.2. It should be mentioned that this theorem leads to a solution to the problem of finding an s–t path of a designated label in $\mathbb{Z}_3$-labeled graphs.

Corollary 7.4. *Let G be a $\mathbb{Z}_3$-labeled graph with two distinct vertices $s, t \in V(G)$. Then one can compute $l(G; s, t)$ in polynomial time. Furthermore, for each element $\alpha \in l(G; s, t)$, one can find an s–t path P with $\psi_G(P) = \alpha$ in polynomial time.*

7.1.3 Characterization

In this section, we provide a complete characterization of triplets $(G, s, t) \in \mathcal{D}$ with $l(G; s, t) = \{\alpha, \beta\}$ for some distinct $\alpha, \beta \in \Gamma$. We consider two cases separately: when $\alpha\beta^{-1} = \beta\alpha^{-1}$ and when $\alpha\beta^{-1} \neq \beta\alpha^{-1}$.

First, we give a characterization in the easier case: when $\alpha\beta^{-1} = \beta\alpha^{-1}$. Note that this case does not appear when $\Gamma \simeq \mathbb{Z}_3$. The following proposition holds analogously to Propositions 2.5 and 7.1, which characterize triplets $(G, s, t) \in \mathcal{D}$ with $|l(G; s, t)| = 1$ by the t-equivalence of G to a trivially-balanced Γ-labeled graph.

Proposition 7.5. *Let α and β be distinct elements in Γ with $\alpha\beta^{-1} = \beta\alpha^{-1}$. For any $(G, s, t) \in \mathcal{D}$, $l(G; s, t) = \{\alpha, \beta\}$ if and only if G is unbalanced and there exists a Γ-labeled graph G' which is $\{s, t\}$-equivalent to G such that*

$$\psi_{G'}(e) = \begin{cases} \alpha \text{ or } \beta & (e \in \delta^{\mathrm{out}}_{G'}(s), \text{ i.e., } e \text{ leaves } s), \\ \alpha^{-1} \text{ or } \beta^{-1} & (e \in \delta^{\mathrm{in}}_{G'}(s), \text{ i.e., } e \text{ enters } s), \\ 1_\Gamma \text{ or } \alpha\beta^{-1} & (\text{otherwise}), \end{cases} \tag{7.1}$$

for every arc $e \in E(G') = E(G)$. Moreover, one can find such G' in $\mathrm{O}(|V| + |E|)$ time if exists.

Proof. "If" part is easy to see as follows. Since G is not balanced, $|l(G; s, t)| \geq 2$ by Proposition 7.1. Furthermore, since $\alpha\beta^{-1} = \beta\alpha^{-1}$, the label of any s–t path in G' is α or β. Hence, the $\{s, t\}$-equivalence between G and G' leads to $l(G; s, t) = l(G'; s, t) = \{\alpha, \beta\}$.

The converse direction is rather difficult. Using Proposition 2.4, take an arbitrary spanning tree T of G and apply shifting at each $v \in V - t$ so that $\psi(e) = 1_\Gamma$ for every arc $e \in E(T)$, where ψ denotes the resulting label function. Since $l(G; s, t) = \{\alpha, \beta\}$ and $l(T; s, t) = 1_\Gamma$, we applied shifting by α or β at s. Hence, by shifting ψ by α^{-1} or β^{-1}, respectively, at s after the above procedure, we can obtain a Γ-labeled graph G' which is $\{s, t\}$-equivalent to G, and this G' is in fact desired one.

To see this, suppose to the contrary that some arc $e' \in E(G')$ does not satisfy (7.1), and let $E' \subsetneq E(G')$ be the set of arcs satisfying (7.1). Note that $E(T) \subseteq E'$, and hence $G'[E']$ is connected. Take an s–t path P in G' with $E(P) \setminus E' \neq \emptyset$ so that $|E(P) \setminus E'|$ is minimized.

If $|E(P) \setminus E'| = 1$, then $\psi_{G'}(P) \notin \{\alpha, \beta\}$, which contradicts $l(G'; s, t) = l(G; s, t) = \{\alpha, \beta\}$. Otherwise, we have $|E(P) \setminus E'| \geq 2$. Let $e_1, e_2 \in E(P) \setminus E'$ be the first two such arcs traversed in walking along P, and Q the subpath of P connecting e_1 and e_2 (hence, $E(Q) \subseteq E'$). Since $G'[E']$ is connected, there exists a path R from $u \in V(Q)$ to $w \in V(P) \setminus V(Q)$ in $G'[E']$. We can construct an s–t path P' from P by replacing $P[u, w]$ (or $P[w, u]$) with R (or $\bar{R}$) such that $\emptyset \neq E(P') \setminus E' \subsetneq E(P) \setminus E'$ (since $|E(P') \cap \{e_1, e_2\}| = 1$). This implies that $1 \leq |E(P') \setminus E'| \leq |E(P) \setminus E'| - 1$, which contradicts the choice of R. $\square$

We next discuss the main case, which is much more difficult: when $\alpha\beta^{-1} \neq \beta\alpha^{-1}$. The following theorem, one of our main results, completes a characterization of triplets $(G, s, t) \in \mathcal{D}$ with $l(G; s, t) = \{\alpha, \beta\}$ for some distinct $\alpha, \beta \in \Gamma$. The definition of the set $\mathcal{D}_{\alpha,\beta} \subseteq \mathcal{D}$, which appears in the theorem, is shown later through Definitions 7.7–7.11 in Section 7.1.4. In short, $(G, s, t) \in \mathcal{D}_{\alpha,\beta}$ if G is constructed by "gluing" together "nice" planar Γ-labeled graphs (and some trivial Γ-labeled graphs) and their derivations.

Theorem 7.6. *Let α and β be distinct elements in Γ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$. For any $(G, s, t) \in \mathcal{D}$, $l(G; s, t) = \{\alpha, \beta\}$ if and only if $(G, s, t) \in \mathcal{D}_{\alpha,\beta}$.*

Recall that $|l(G;s,t)| = 1$ if and only if G is balanced by Proposition 7.1, which can be easily tested by Proposition 2.5. Hence, these characterizations lead to the first nontrivial classification of Γ-labeled graphs in terms of the number of possible labels of s–t paths, and the classification is also complete when $\Gamma \simeq \mathbb{Z}_3$.

7.1.4 Trivial cases and new operations

Fix distinct elements $\alpha, \beta \in \Gamma$ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$. To characterize triplets $(G,s,t) \in \mathcal{D}$ with $l(G;s,t) = \{\alpha,\beta\}$, let us define several sets of triplets $(G,s,t) \in \mathcal{D}$ for which it is easy to see that $l(G;s,t) = \{\alpha,\beta\}$. Theorem 7.6 claims that any triplet $(G,s,t) \in \mathcal{D}$ with $l(G;s,t) = \{\alpha,\beta\}$ is in fact contained in one of them.

Suppose that a graph G is embedded on a plane. We call a unique unbounded face of G the *outer face* of G, and any other face an *inner face*. For a face F of G, let $\mathrm{bd}(F)$ denote the closed walk obtained by walking the boundary of F in an arbitrary direction from an arbitrary vertex on it.

Definition 7.7. For distinct elements $\alpha, \beta \in \Gamma$ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$, let $\mathcal{D}^0_{\alpha,\beta}$ be the set of all triplets $(G,s,t) \in \mathcal{D}$ satisfying one of the following conditions.

(A) There exists a Γ-labeled graph G' which is not balanced and is $\{s,t\}$-equivalent to G such that $\delta^{\mathrm{out}}_{G'}(s) = \delta_{G'}(s)$, $\delta^{\mathrm{in}}_{G'}(t) = \delta_{G'}(t)$, and either

- the label of every arc in $G' - s$ is 1_Γ and in $\delta^{\mathrm{out}}_{G'}(s)$ is α or β (see Fig. 7.1), or
- the label of every arc in $G' - t$ is 1_Γ and in $\delta^{\mathrm{in}}_{G'}(t)$ is α or β (see Fig. 7.2).

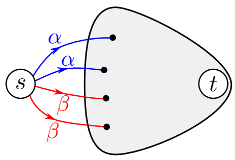

Figure 7.1: The former of Case (A).

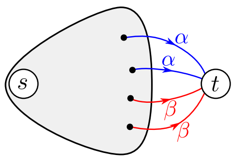

Figure 7.2: The latter of Case (A).

(B) G is $\{s,t\}$-equivalent to the Γ-labeled graph consisting of six vertices s, v_1, v_2, v_3, v_4, t, six arcs $sv_1, sv_2, v_1v_2, v_3v_4, v_3t, v_4t$ with label 1_Γ, and two pairs of two parallel arcs from v_i to v_{i+2} $(i = 1, 2)$ whose labels are both α and β (see Fig. 7.3).

(C) G can be embedded on a plane with the face set $\mathcal{F}$ (see Fig. 7.4) so that

- both s and t are on the boundary of the outer face $F_0 \in \mathcal{F}$,
- one s–t path along $\mathrm{bd}(F_0)$ is of label α and the other is of β, and
- there exists a unique inner face $F_1 \in \mathcal{F}$ whose boundary is unbalanced, i.e., $\psi_G(\mathrm{bd}(F_1)) \neq 1_\Gamma$ and $\psi_G(\mathrm{bd}(F)) = 1_\Gamma$ for any $F \in \mathcal{F} \setminus \{F_0, F_1\}$.

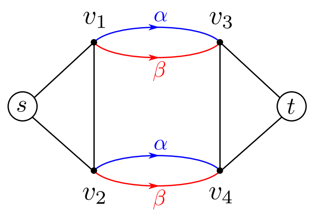

Figure 7.3: Case (B).

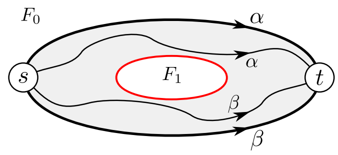

Figure 7.4: Case (C).

It is not difficult to see that $l(G; s, t) = \{\alpha, \beta\}$ for any triplet $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$.

Here, in order to define larger subsets of $\mathcal{D}$ than $\mathcal{D}^0_{\alpha,\beta}$, we introduce two new operations which do not make any effect on $l(G; s, t)$. Fix a triplet $(G = (V, E), s, t) \in \mathcal{D}$, and define $G[\![X]\!] := G[X \cup N_G(X)] - E(N_G(X))$ for a vertex set $X \subseteq V$.

Definition 7.8 (2-contraction)**.** For a vertex set $X \subseteq V \setminus \{s, t\}$ such that $N_G(X) = \{x, y\}$ for some distinct vertices $x, y \in V$ and $G[\![X]\!]$ is connected, the 2-*contraction* of X is the following operation (see Fig. 7.5):

- remove all vertices in X, and
- add a new arc from x to y with label α for each $\alpha \in l(G[\![X]\!]; x, y)$ if there is no such arc in G.

The resulting graph is denoted by $G/_2X$. A vertex set $X \subseteq V \setminus \{s, t\}$ is said to be 2-*contractible in* G if the 2-contraction of X can be performed in G and $G[\![X]\!] \neq G$.

Definition 7.9 (3-contraction)**.** For a vertex set $X \subseteq V \setminus \{s, t\}$ such that $|N_G(X)| = 3$, $G[X]$ is connected, and $G[\![X]\!]$ is balanced, the 3-*contraction* of X is the following operation (see Fig. 7.6):

- remove all vertices in X, and
- add a new arc from x to y with label $l(G[\![X]\!]; x, y)$ (which consists of a single element by Proposition 7.1) for each pair of $x, y \in N_G(X)$ if there is no such arc in G.

The resulting graph is denoted by $G/_3X$. A vertex set $X \subseteq V \setminus \{s, t\}$ is said to be 3-*contractible in* G if the 3-contraction of X can be performed in G.

The 2-contraction and the 3-contraction are analogous to the operation which is performed in Condition 3 in Theorem 1.8, and we use the same term "contraction" to refer to each of them. Any contraction does not change $l(G; s, t)$, since each s–t path cannot enter $G[\![X]\!]$ after leaving it once (i.e., cannot traverse arcs in $G[\![X]\!]$ intermittently). Moreover, we also have $(G', s, t) \in \mathcal{D}$ for the resulting graph G' after any contraction.

Using these two operations, we define two larger subsets of $\mathcal{D}$ than $\mathcal{D}^0_{\alpha,\beta}$.

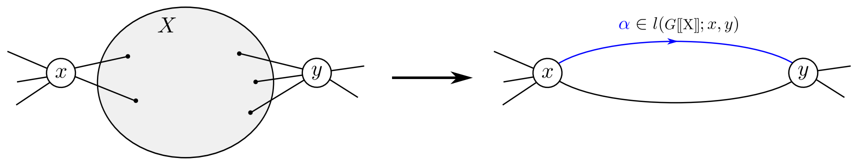

Figure 7.5: 2-contraction.

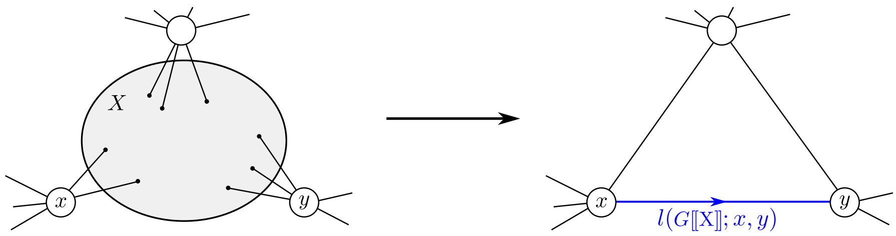

Figure 7.6: 3-contraction.

Definition 7.10. For distinct elements $\alpha, \beta \in \Gamma$ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$, we define $\mathcal{D}^1_{\alpha,\beta}$ as the minimal set of triplets $(G, s, t) \in \mathcal{D}$ with the following conditions:

- $\mathcal{D}^0_{\alpha,\beta} \subseteq \mathcal{D}^1_{\alpha,\beta}$, and
- if $(G/_3X, s, t) \in \mathcal{D}^1_{\alpha,\beta}$ for some 3-contractible $X \subseteq V \setminus \{s, t\}$, then $(G, s, t) \in \mathcal{D}^1_{\alpha,\beta}$.

Definition 7.11. For distinct elements $\alpha, \beta \in \Gamma$ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$, we define $\mathcal{D}_{\alpha,\beta}$ as the minimal set of triplets $(G, s, t) \in \mathcal{D}$ with the following conditions:

- $\mathcal{D}^1_{\alpha,\beta} \subseteq \mathcal{D}_{\alpha,\beta}$, and
- if $(G/_2X, s, t) \in \mathcal{D}_{\alpha,\beta}$ for some $X \subseteq V \setminus \{s, t\}$ such that either $G[\![X]\!]$ is balanced or $(G[\![X]\!], x, y) \in \mathcal{D}^1_{\alpha',\beta'}$, where $N_G(X) = \{x, y\}$ and $\alpha', \beta' \in \Gamma$ satisfy $\alpha'\beta'^{-1} \neq \beta'\alpha'^{-1}$, then $(G, s, t) \in \mathcal{D}_{\alpha,\beta}$.

Note that the first condition in Definition 7.11 can be replaced with $(G_0, s, t) \in \mathcal{D}_{\alpha,\beta}$, where G_0 consists of two parallel arcs from s to t whose labels are α and β.

It is easy to see that $l(G; s, t) = \{\alpha, \beta\}$ for any triplet $(G, s, t) \in \mathcal{D}_{\alpha,\beta}$ since any contraction does not change $l(G; s, t)$. A proof of the non-trivial direction ("only if" part of Theorem 7.6) is presented later in Section 7.4 (and sketched in Section 7.1.5).

7.1.5 Proof sketch

Since our proof of Theorem 7.6 shown in Section 7.4 is long, here we give its sketch.

To derive a contradiction, assume that there exist distinct elements $\alpha, \beta \in \Gamma$ and a triplet $(G, s, t) \in \mathcal{D}$ such that $\alpha\beta^{-1} \neq \beta\alpha^{-1}$, $l(G; s, t) = \{\alpha, \beta\}$, and $(G, s, t) \notin \mathcal{D}_{\alpha,\beta}$. We choose such $\alpha, \beta \in \Gamma$ and $(G, s, t) \in \mathcal{D}$ so that G is as small as possible.

Fix an arbitrary arc $e_0 \in \delta_G^{\mathrm{out}}(s)$ in G (where we assume $\delta_G^{\mathrm{out}}(s) = \delta_G(s)$), and define $G' := G - e_0$. By using the minimality of G, we can show that $(G', s, t) \in \mathcal{D}_{\alpha,\beta}$ (cf. Claims 7.19 and 7.20). We consider the following two cases separately: when $(G', s, t) \in \mathcal{D}^1_{\alpha,\beta}$ and when not (Sections 7.4.1 and 7.4.2, respectively).

In both cases, we can embed a graph $\tilde{G}$ obtained from G' by at most one 3-contraction on a plane so that the conditions of Case (C) in Definition 7.7 are satisfied (or derive a contradiction). By expanding a vertex set and adding e_0, we try to extend the planar embedding of $\tilde{G}$ to G. Then, we have one of the following cases.

- Such an extension is possible, i.e., G can be embedded on a plane with the conditions of Case (C) in Definition 7.7. This contradicts that $(G, s, t) \notin \mathcal{D}_{\alpha,\beta}$.
- G contains a contractible vertex set, which contradicts that G is a minimal counterexample (cf. Claims 7.17 and 7.18).
- We can construct an s–t path of label $\gamma \in \Gamma \setminus \{\alpha, \beta\}$ in G by using e_0 and some arcs in G', which contradicts that $l(G; s, t) = \{\alpha, \beta\}$.

In each case, we have a contradiction, which completes the proof. We remark that Theorem 1.8 plays an important role in this case analysis.

7.2 Algorithm

In this section, we give a proof of Theorem 7.3. That is, we present a polynomial-time algorithm to test whether $l(G; s, t) \subseteq \{\alpha, \beta\}$ or not for given distinct $\alpha, \beta \in \Gamma$ and to find an s–t path of label $\gamma \in \Gamma \setminus \{\alpha, \beta\}$ if $l(G; s, t) \not\subseteq \{\alpha, \beta\}$, in a given Γ-labeled graph $G = (V, E)$ with $s, t \in V$. It should be mentioned that, when $\Gamma \simeq \mathbb{Z}_3$, such an algorithm can compute $l(G; s, t)$ itself and find an s–t path of label α for each $\alpha \in l(G; s, t)$. Without loss of generality, we assume that G does not have parallel arcs with the same label.

7.2.1 Algorithm description

We separate our algorithm into two parts: to test whether $|l(G; s, t)| \leq 2$ or not and return at most two s–t paths which attain all labels in $l(G; s, t)$ when $|l(G; s, t)| \leq 2$, and to find three s–t paths whose labels are distinct when it has turned out that $|l(G; s, t)| \geq 3$.

We first present the former algorithm. Note again that this algorithm can compute $l(G; s, t)$ itself when $\Gamma \simeq \mathbb{Z}_3$. Throughout this algorithm, let $G' = (V', E')$ denote a temporary Γ-labeled graph currently considered.

TESTTWOLABELS(G, s, t)

Input A Γ-labeled graph $G = (V, E)$ and distinct vertices $s, t \in V$.

Output The set $l(G; s, t)$ of all possible labels of s–t paths in G with those which attain the labels if $|l(G; s, t)| \leq 2$, and "$|l(G; s, t)| \geq 3$" otherwise.

Step 0. Compute the maximal subgraph G' of G which contains no redundant vertex by the 2-connected component decomposition (cf. Proposition 2.1). Note that $(G', s, t) \in \mathcal{D}$ and $l(G'; s, t) = l(G; s, t)$.

Step 1. Test whether G' is balanced or not by Proposition 2.5. If G' is balanced, then halt with returning the label of an arbitrary s–t path in G with the path. Otherwise, by using an unbalanced cycle, obtain two s–t paths in G whose labels are distinct (cf. the proof of Proposition 2.8), say $\alpha, \beta \in \Gamma$. In the following steps, we check whether $l(G'; s, t) = \{\alpha, \beta\}$ or not.

Step 2. If $\alpha\beta^{-1} = \beta\alpha^{-1}$, then check the condition in Proposition 7.5. Return $\{\alpha, \beta\}$ with the two s–t paths in G obtained in Step 1 if it is satisfied, and "$|l(G; s, t)| \geq 3$" otherwise. Otherwise (i.e., if $\alpha\beta^{-1} \neq \beta\alpha^{-1}$), to make G' 2-connected, add to G' a new arc from s to t with label α (or β) if s and t are not adjacent in G'.

Step 3. While G' is not 3-connected and $|V'| \geq 4$, do the following procedure. Let $\{x, y\} \subsetneq V'$ be a 2-cut in G', and X the vertex set of a connected component of $G' - \{x, y\}$ with $X \cap \{s, t\} = \emptyset$ (such X exists, since s and t are adjacent in G'). Test whether $|l(G'[\![X]\!]; x, y)| \leq 2$ or not recursively by TESTTWOLABELS($G'[\![X]\!], x, y$). Update $G' \leftarrow G' /_2 X$ (2-contraction) if $|l(G'[\![X]\!]; x, y)| \leq 2$, and return "$|l(G; s, t)| \geq 3$" otherwise.

Step 4. While there exists a 3-contractible vertex set $X \subseteq V' \setminus \{s, t\}$, update $G' \leftarrow G' /_3 X$ (3-contraction).

Step 5. If $|V'| \leq 6$, then compute $l(G', s, t)$ by enumerating all s–t paths in G' and return the result. Otherwise, test whether $(G', s, t) \in \mathcal{D}^0_{\alpha,\beta}$ or not by Lemma 7.12. Return $\{\alpha, \beta\}$ with the s–t paths in G obtained in Step 1 if $(G', s, t) \in \mathcal{D}^0_{\alpha,\beta}$, and "$|l(G; s, t)| \geq 3$" otherwise.

Next, we show the latter algorithm, which finds three s–t paths whose labels are distinct when it has turned out that $|l(G; s, t)| \geq 3$. Also note again that this algorithm finds three s–t paths which attain all labels when $\Gamma \simeq \mathbb{Z}_3$.

FINDTHREEPATHS(G, s, t)

Input A Γ-labeled graph $G = (V, E)$ and distinct vertices $s, t \in V$ with $|l(G; s, t)| \geq 3$.

Output Three s–t paths in G whose labels are distinct.

Step 0. If $V = \{s, t\}$, then halt with returning three s–t paths each of which consists of a single arc from s to t in E.

Step 1. Test whether $|l(G-s; s', t)| \leq 2$ or not by TestTwoLabels$(G-s, s', t)$ for each neighbor $s' \in N_G(s) - t$.

Step 2. If $|l(G-s; s', t)| \leq 2$ for all $s' \in N_G(s) - t$, then we have already obtained s'–t paths which attain all labels in $l(G-s; s', t)$. Choose three s–t paths whose labels are distinct among the s–t paths obtained by extending such s'–t paths using an arc (possibly parallel arcs) $ss' \in E$ for each $s' \in N_G(s) - t$ and the s–t paths each of which consists of a single arc $st \in E$, and halt with returning them.

Step 3. Otherwise, we obtained $|l(G-s; \tilde{s}, t)| \geq 3$ for some neighbor $\tilde{s} \in N_G(s) - t$. Then, recursively by FindThreePaths$(G-s, \tilde{s}, t)$, find three $\tilde{s}$–t paths whose labels are distinct. Extend them using an arc $s\tilde{s} \in E$, and return the extended s–t paths.

7.2.2 Time complexity

Before starting the proof, we show the detailed procedure of Step 5 in TestTwoLabels.

Lemma 7.12. *Let $(G, s, t) \in \mathcal{D}$. Suppose that $G = (V, E)$ is 3-connected and contains no 3-contractible vertex set, $|V| > 6$, s and t are adjacent, and $\{\alpha, \beta\} \subseteq l(G; s, t)$ for some distinct $\alpha, \beta \in \Gamma$ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$. Then, one can test whether $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$ or not in polynomial time.*

Proof. Since $|V| > 6$, it is not necessary to consider Case (B) in Definition 7.7. Besides, Case (A) is easily checked by testing whether $G - s$ or $G - t$ is balanced or not. Hence, in what follows, we assume that (G, s, t) is not in Case (A) or (B) and focus on Case (C).

First, test the planarity of G. If G is not planar, then we can conclude $(G, s, t) \notin \mathcal{D}^0_{\alpha,\beta}$. Otherwise, compute an embedding of G on a plane in which both s and t are on the outer boundary (because of an arc $st \in E$, there exists a face on whose boundary both s and t are). It should be noted that such a planar embedding can be computed in polynomial time (e.g., by [36]). Since G is 3-connected, the face set is unique if there are no parallel arcs (see, e.g., [11, Chapter 4]). Although there may be parallel arcs in G, we can say that the number of parallel arcs is bounded as follows.

Claim. We may assume that there are no parallel arcs between s and t.

Suppose that there exist parallel arcs from s to t, which may be assumed to have distinct labels. Moreover, we may assume that there are exactly two such arcs $e_\alpha, e_\beta \in E$ with labels α, β, respectively, since otherwise, we have $|l(G; s, t)| \geq 3$ and hence we can conclude $(G, s, t) \notin \mathcal{D}^0_{\alpha,\beta}$. Since $|V| > 6$ and $(G, s, t) \in \mathcal{D}$, there exists an s–t path in $G - \{e_\alpha, e_\beta\}$, and let γ be its label. If $\alpha \neq \gamma \neq \beta$, then $|l(G; s, t)| \geq 3$. Otherwise, remove e_γ from G. Note that this removal does not violate the hypotheses of this lemma, and does not make any effect on whether $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$ or not.

Claim. We may assume that there exists at most one pair of parallel arcs.

Suppose that there exist parallel arcs from x to y with distinct labels, where $\{x, y\} \neq \{s, t\}$. Then, by the 3-connectivity of G, the parallel arcs form an inner face whose boundary is unbalanced (since otherwise $\{x, y\}$ is a 2-cut in G). Hence, there is a unique pair of such parallel arcs if $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$, since the existence of at least two pairs of parallel arcs immediately implies that there exist at least two inner faces whose boundaries are unbalanced.

Recall that we have to test whether there exists an embedding of G such that the outer boundary is unbalanced and there exists a unique inner face whose boundary is unbalanced. Since a pair of parallel arcs is unique if exists, there are at most two possible face sets of G. Furthermore, since there exists exactly one arc from s to t, both of the two faces whose boundaries share the arc $st \in E$ can be the outer face, i.e., there are two choices of the outer face. It can be done in polynomial time to check, in each of the at most four $(= 2 \times 2)$ cases, whether exactly one inner face has an unbalanced boundary or not, and hence one can do the whole procedure in polynomial time. □

We are now ready to prove Theorem 7.3.

Proof of Theorem 7.3. Recall that our goal is to test whether $|l(G; s, t)| \leq 2$ or not, and to find $\min\{3, |l(G; s, t)|\}$ s–t paths whose labels are distinct. These are achieved as follows. For the input triplet (G, s, t) (which may not be in $\mathcal{D}$), we first test whether $|l(G; s, t)| \leq 2$ or not by TESTTWOLABELS(G, s, t). If we obtain $|l(G; s, t)| \leq 2$, then we also obtain at most two s–t paths in G which attain all labels in $l(G; s, t)$. Otherwise, we can obtain three s–t paths whose labels are distinct by FINDTHREEPATHS(G, s, t). Hence, it suffices to show the correctness and polynomiality of these two algorithms.

The correctness of these two algorithms is almost obvious. It should be noted that we have $l(G'; s, t) = l(G; s, t)$ and $(G', s, t) \in \mathcal{D}$ at any step of TESTTWOLABELS(G, s, t). This follows from the fact that the 2-contractions in Step 3 and the 3-contractions in Step 4 do not change $l(G'; s, t)$ or violate $(G', s, t) \in \mathcal{D}$.

We finally confirm the polynomiality of the two algorithms. Let $T_{\text{labels}}(n)$ and $T_{\text{paths}}(n)$ denote the computational time of TESTTWOLABELS(G, s, t) and FINDTHREEPATHS(G, s, t), respectively, where n is the number of vertices in G. It is easy to see that TESTTWOLABELS runs in polynomial time, i.e., $T_{\text{labels}}(n)$ is polynomially bounded. Note that, in the recursion step (Step 3), we just divide the graph G' into two smaller graphs which have $|V'| - |X|$ and $|X| + 2$ vertices, and in the 3-contraction step (Step 4), it suffices to check all 3-cuts in G', whose number is $\mathrm{O}(n^3)$. For FINDTHREEPATHS, by a recurrence relation

$$T_{\text{paths}}(n) \leq n \cdot T_{\text{labels}}(n-1) + T_{\text{paths}}(n-1) + \text{poly}(n),$$

we have $T_{\text{paths}}(n) \leq n^2 \cdot T_{\text{labels}}(n) + \text{poly}(n)$. Hence, $T_{\text{paths}}(n)$ is polynomially bounded. □

7.3 Useful Lemmas

Before starting the proof of our characterization (Theorem 7.6), we show several lemmas. Fix distinct elements $\alpha, \beta \in \Gamma$ with $\alpha\beta^{-1} \neq \beta\alpha^{-1}$.

Lemma 7.13. *For any $(G = (V, E), s, t) \in \mathcal{D}_{\alpha,\beta}$, we have the following properties.*

(1) *Let G' be the graph obtained from G by shifting by $\gamma \in \Gamma$ at s. Then, $(G', s, t) \in \mathcal{D}_{\alpha',\beta'}$, where $\alpha' := \alpha\gamma^{-1}$ and $\beta' := \beta\gamma^{-1}$.*

(2) *Let G' be the graph obtained from G by adding a new vertex s' and a new arc $e' = s's$ with label $\gamma \in \Gamma$. Then, $(G', s', t) \in \mathcal{D}_{\alpha',\beta'}$, where $\alpha' := \alpha\gamma$ and $\beta' := \beta\gamma$.*

(3) *Let G' be a Γ-labeled graph such that $G = G' /_2 X$ for some $X \subseteq V(G') \setminus \{s, t\}$ with $(G'[\![X]\!], x, y) \in \mathcal{D}_{\alpha',\beta'}$, where $N_{G'}(X) = \{x, y\}$ and $\alpha', \beta' \in \Gamma$ satisfy $\alpha'\beta'^{-1} \neq \beta'\alpha'^{-1}$. Then, $(G', s, t) \in \mathcal{D}_{\alpha,\beta}$.*

Proof. (1) We first confirm that, if $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$, then $(G', s, t) \in \mathcal{D}^0_{\alpha',\beta'}$. The former of Case (A) and Case (C) are obvious (cf. Definition 7.7). In the latter of Case (A), apply shifting by γ at each $v \in V \setminus \{s, t\}$, and in Case (B), do so at v_1 and v_2.

We next show that, if $(G, s, t) \in \mathcal{D}^1_{\alpha,\beta}$, then $(G', s, t) \in \mathcal{D}^1_{\alpha',\beta'}$. Suppose that $(G, s, t) \in \mathcal{D}^1_{\alpha,\beta}$. Then, one can obtain a Γ-labeled graph $\tilde{G}$ such that $(\tilde{G}, s, t) \in \mathcal{D}^0_{\alpha,\beta}$ from G by applying 3-contractions. Since any shifting does not make effect on whether a Γ-labeled graph is balanced or not, the same 3-contractions can be applied to G', and we obtain a Γ-labeled graph $\tilde{G}'$ such that $(\tilde{G}', s, t) \in \mathcal{D}^0_{\alpha',\beta'}$ as a result. Thus we have done.

By the definition of $\mathcal{D}_{\alpha,\beta}$ (Definition 7.11), there exists a sequence $G_0, G_1, \ldots, G_r$ of Γ-labeled graphs satisfying the following conditions:

- $G_r = G$,
- G_0 consists of two vertices s and t and two parallel arcs e_α, e_β from s to t whose labels are α and β, respectively, and
- for each $i \in [r]$, $G_{i-1} = G_i /_2 X_i$ for some $X_i \subseteq V(G_i) \setminus \{s, t\}$ such that either $G_i[\![X_i]\!]$ is balanced or $(G_i[\![X_i]\!], x_i, y_i) \in \mathcal{D}^1_{\alpha_i,\beta_i}$, where $N_{G_i}(X_i) = \{x_i, y_i\}$ and $\alpha_i, \beta_i \in \Gamma$ satisfy $\alpha_i\beta_i^{-1} \neq \beta_i\alpha_i^{-1}$.

We prove that the same 2-contractions can be applied to G'.

Define $G'_r := G'$. Then, we can inductively construct a Γ-labeled graph $G'_{i-1} := G'_i /_2 X_i$, which coincides with the one obtained from G_{i-1} by shifting by γ at s. This means that we finally obtain a Γ-labeled graph G'_0 from G' by the 2-contractions of X_i ($i = r, r-1, \ldots, 1$), which satisfies $(G'_0, s, t) \in \mathcal{D}^0_{\alpha',\beta'}$ (in Cases (A) and (C)). Thus we have $(G', s, t) \in \mathcal{D}_{\alpha',\beta'}$, since either $G'_i[\![X_i]\!]$ is balanced or $(G'_i[\![X_i]\!], x_i, y_i) \in \mathcal{D}^1_{\alpha'_i,\beta'_i}$, where $\alpha'_i = \alpha_i$ and $\beta'_i = \beta_i$ if $s \notin \{x_i, y_i\}$, and $\alpha'_i = \alpha_i\gamma^{-1}$ and $\beta'_i = \beta_i\gamma^{-1}$ otherwise (assume $x_i = s$ without loss of generality by the symmetry of x_i and y_i).

(2) Similarly to the proof of (1), there exists a sequence $G_0, G_1, \ldots, G_r = G$ such that $G_0 = (\{s,t\}, \{e_\alpha, e_\beta\})$, and G_{i-1} is obtained from G_i by some appropriate 2-contraction. The same 2-contractions can be applied to G', and we obtain the Γ-labeled graph $G'_0 = (\{s', s, t\}, \{e', e_\alpha, e_\beta\})$, which satisfies $(G'_0, s', t) \in \mathcal{D}^0_{\alpha', \beta'}$ (in Cases (A) and (C)). Thus we have $(G', s, t) \in \mathcal{D}_{\alpha', \beta'}$.

(3) Similarly, there exists a sequence $H_0, H_1, \ldots, H_r = G'[\![X]\!]$ such that H_0 consists of two parallel arcs from x to y whose labels are α' and β', and H_{i-1} is obtained from H_i by some appropriate 2-contraction. The same 2-contractions can be applied to G', and we obtain G. This implies that $(G', s, t) \in \mathcal{D}_{\alpha,\beta}$. $\square$

By Lemma 7.13-(1), it suffices to consider the case when $\beta = 1_\Gamma$ and $\alpha^{-1} \neq \alpha$ (i.e., $\alpha^2 \neq 1_\Gamma$). The following lemma gives a useful characterization of $\mathcal{D}^0_{1_\Gamma, \alpha}$ in Case (C).

Lemma 7.14. *Suppose that $\alpha^{-1} \neq \alpha \in \Gamma$. For any triplet $(G = (V, E), s, t) \in \mathcal{D}^0_{1_\Gamma, \alpha}$ in Case* (C) *in Definition* 7.7*, there exists a Γ-labeled graph G' which is $\{s,t\}$-equivalent to G and embeddable with the following conditions* (*see Fig.* 7.7)*.*

1. *The arc set E is partitioned into E^0 and E^1 (i.e., $E^0 \cup E^1 = E$ and $E^0 \cap E^1 = \emptyset$), where $E^i := \{\, e \in E \mid \psi_{G'}(e) = \alpha^i \,\}$ $(i = 0, 1)$.*
2. *There exists an s–t path $P = (s = u_0, e_1, u_1, \ldots, e_l, u_l = t)$ along the outer boundary of $G' - E^1$ such that*
   - *every arc in E^1 is embedded on the outer face of $G' - E^1$ and is from $u_i \in V(P)$ to $u_j \in V(P)$ for some $i < j$, and*
   - *for any distinct arcs $e_1 = u_{i_1}u_{j_1}$, $e_2 = u_{i_2}u_{j_2} \in E^1$, one of two paths $P[u_{i_1}, u_{j_1}]$ and $P[u_{i_2}, u_{j_2}]$ is a subpath of the other.*

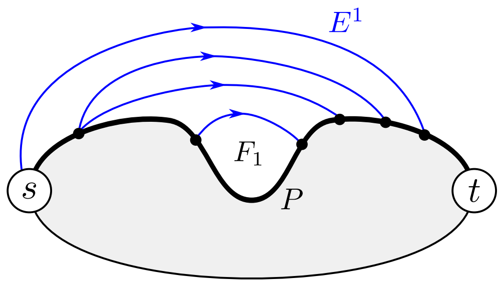

Figure 7.7: An $\{s,t\}$-equivalent embedding of $(G, s, t) \in \mathcal{D}^0_{1_\Gamma, \alpha}$ in Case (C).

Proof. Fix an embedding of G with the conditions of Case (C), and let P_0 and P_1 be the s–t paths along the boundary of the outer face F_0 of G whose labels are 1_Γ and α, respectively.

Let G^* be the dual graph of G (as an undirected graph), i.e., the vertex set of G^* is the face set $\mathcal{F}$ of G, the edge set of G^* coincides with the edge set of G, and each two faces whose boundaries share an edge $e \in E$ in G are connected by the same-named edge e in G^*. Take a shortest F_1–F_0 path Q in $G^* - E(P_0)$. We prove that the second condition holds with $E^1 = E(Q)$.

Note that $G'' := G - E(Q)$ is connected since Q is a shortest path without the corresponding edge to any arc in $E(P_0)$, and that G'' is balanced since F_1 is the unique unbalanced inner face. We then have $l(G''; s, t) = 1_\Gamma$ by Proposition 7.1. Hence, we may assume that $\psi_G(e) = 1_\Gamma$ for every arc $e \in E(G'')$ by shifting at some vertices $v \in V \setminus \{s, t\}$. Thus we obtain G' with the second condition, since $\psi_G(\mathrm{bd}(F)) = 1_\Gamma$ for any $F \in \mathcal{F} \setminus \{F_0, F_1\}$. $\square$

The following two lemmas are utilized to derive a contradiction by constructing an s–t path of label $\gamma \notin \Gamma \setminus \{\alpha, \beta\}$ in G.

Lemma 7.15. *For a triplet $(G, s, t) \in \mathcal{D}$, if G contains an unbalanced cycle C with $\psi_G(\bar{C}) = \psi_G(C)$, then there exist distinct elements $\alpha', \beta' \in l(G; s, t)$ with $\alpha'\beta'^{-1} = \beta'\alpha'^{-1}$.*

Proof. We first note that the equality $\psi_G(\bar{C}) = \psi_G(C)$ does not depend on the choices of the direction and the end vertex of the cycle C. Suppose that G contains such an unbalanced cycle C. By Menger's theorem (cf. the proof of Proposition 2.8), for some distinct vertices $x, y \in V(C)$, one can take an s–x path P and a y–t path Q in G so that $V(P) \cap V(C) = \{x\}$, $V(Q) \cap V(C) = \{y\}$, and $V(P) \cap V(Q) = \emptyset$, and choose y as the end vertex of C.

Let $\alpha'' := \psi_G(C[x, y])$ and $\beta'' := \psi_G(\bar{C}[x, y])$, which are distinct since C is unbalanced. We then have $\alpha''\beta''^{-1} = \psi_G(C) = \psi_G(\bar{C}) = \beta''\alpha''^{-1}$. By extending $C[x, y]$ and $\bar{C}[x, y]$ using P and Q, we obtain two s–t paths in G whose labels are $\alpha' := \psi_G(Q) \cdot \alpha'' \cdot \psi_G(P)$ and $\beta' := \psi_G(Q) \cdot \beta'' \cdot \psi_G(P)$, which are also distinct. Since $\alpha''\beta''^{-1} = \beta''\alpha''^{-1}$, we have $\alpha'\beta'^{-1} = \psi_G(Q) \cdot \alpha'' \cdot \beta''^{-1} \cdot \psi_G(Q)^{-1} = \psi_G(Q) \cdot \beta'' \cdot \alpha''^{-1} \cdot \psi_G(Q)^{-1} = \beta'\alpha'^{-1}$. $\square$

In particular, G contains no unbalanced cycle C with $\psi_G(\bar{C}) = \psi_G(C)$ if $l(G; s, t) = \{\alpha, \beta\}$ (recall that $\alpha\beta^{-1} \neq \beta\alpha^{-1}$) and $(G, s, t) \in \mathcal{D}$.

Lemma 7.16. *For a triplet $(G, s, t) \in \mathcal{D}$, if there exist two paths P_i $(i = 1, 2)$ in G with the following conditions (see Fig. 7.8), then $|l(G; s, t)| \geq 3$:*

- *P_i is from s to $x_i \in V \setminus \{s, t\}$ for $i = 1, 2$,*
- *$\psi_G(P_1) \neq \psi_G(P_2)$, and*
- *$\{\alpha', \beta'\} \subseteq l(G - (V(P_i) - x_i); x_i, t)$ $(i = 1, 2)$ for some $\alpha', \beta' \in \Gamma$ with $\alpha'\beta'^{-1} \neq \beta'\alpha'^{-1}$.*

Proof. For each $i = 1, 2$, by concatenating P_i and each of two x_i–t paths in $G - (V(P_i) - x_i)$ whose labels are α' and β', we construct four s–t paths whose labels are $\gamma_1 := \alpha' \cdot \psi_G(P_1)$, $\gamma_2 := \beta' \cdot \psi_G(P_1)$, $\gamma_3 := \alpha' \cdot \psi_G(P_2)$, and $\gamma_4 := \beta' \cdot \psi_G(P_2)$.

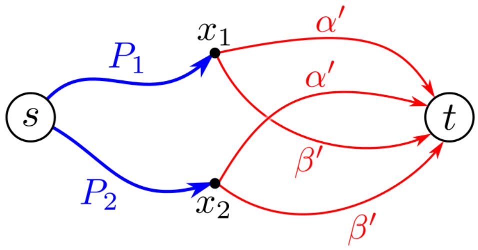

Figure 7.8: Combination of two labels leads to at least three labels.

Suppose to the contrary that $|l(G; s, t)| \leq 2$. Since $\gamma_1 \neq \gamma_2 \neq \gamma_4 \neq \gamma_3 \neq \gamma_1$, we must have $\gamma_1 = \gamma_4$ and $\gamma_2 = \gamma_3$. Hence, $\psi_G(P_1) = \alpha'^{-1} \cdot \beta' \cdot \psi_G(P_2)$ and $\psi_G(P_1) = \beta'^{-1} \cdot \alpha' \cdot \psi_G(P_2)$, which implies $\alpha'^{-1}\beta' = \beta'^{-1}\alpha'$. This is equivalent to $\alpha'\beta'^{-1} = \beta'\alpha'^{-1}$, a contradiction. □

7.4 Proof of Characterization

Here, we start a proof of "only if" part of Theorem 7.6. To derive a contradiction, suppose to the contrary that there exist distinct elements $\alpha, \beta \in \Gamma$ and a triplet $(G, s, t) \in \mathcal{D}$ such that $\alpha\beta^{-1} \neq \beta\alpha^{-1}$, $l(G; s, t) = \{\alpha, \beta\}$, and $(G, s, t) \notin \mathcal{D}_{\alpha,\beta}$. We choose such $\alpha, \beta \in \Gamma$ and $(G = (V, E), s, t) \in \mathcal{D}$ so that the value of $|V| + |E|$ is minimized. Note that we have $|V| \geq 3$ obviously, and we may assume $\beta = 1_\Gamma$ and $\alpha^{-1} \neq \alpha$ (i.e., $\alpha^2 \neq 1_\Gamma$) by Lemma 7.13-(1). By the minimality, G contains no contractible vertex set as follows.

Claim 7.17. *There is no 2-contractible vertex set in G.*

Proof. Suppose to the contrary that G contains a 2-contractible vertex set $X \subseteq V \setminus \{s, t\}$ with $N_G(X) = \{x, y\}$. Since $(G, s, t) \in \mathcal{D}$, we also have $(G[\![X]\!], x, y) \in \mathcal{D}$, where recall that $G[\![X]\!] := G[X \cup N_G(X)] - E(N_G(X))$. If $|l(G[\![X]\!]; x, y)| \geq 3$, then we also have $|l(G; s, t)| \geq 3$ (since G contains two disjoint paths between $\{s, t\}$ and $\{x, y\}$ by Proposition 2.1 and Menger's theorem), a contradiction. In the case that $l(G[\![X]\!]; x, y) = \{\alpha', \beta'\}$ for distinct $\alpha', \beta' \in \Gamma$ with $\alpha'\beta'^{-1} = \beta'\alpha'^{-1}$, there exists an unbalanced cycle C in $G[\![X]\!]$ (which is a subgraph of G) such that $\psi_G(\bar{C}) = \psi_G(C)$ by Proposition 7.5 (since $G[\![X]\!]$ is not balanced, and the label of any unbalanced cycle in $G[\![X]\!]$ is self-inversed by (7.1)), which contradicts Lemma 7.15.

Otherwise, i.e., if $|l(G[\![X]\!]; x, y)| = 1$ or $l(G[\![X]\!]; x, y) = \{\alpha', \beta'\}$ for some $\alpha', \beta' \in \Gamma$ with $\alpha'\beta'^{-1} \neq \beta'\alpha'^{-1}$, we can construct a smaller counterexample by the 2-contraction of X (by Definition 7.11 and Lemma 7.13-(3)), a contradiction. It should be noted that $(G[\![X]\!], x, y) \in \mathcal{D}_{\alpha',\beta'}$ if $l(G[\![X]\!]; x, y) = \{\alpha', \beta'\}$, since G is a minimal counterexample and $G[\![X]\!]$ is a proper subgraph of G by the definition of the term "2-contractible" (see Definition 7.8). □

Claim 7.18. *There is no 3-contractible vertex set in G.*

Proof. Suppose to the contrary that G contains a 3-contractible vertex set $X \subseteq V \setminus \{s,t\}$. The minimality of G implies $(G/_3X, s, t) \in \mathcal{D}_{\alpha,\beta}$, which means that there exists a sequence $G_0, G_1, \ldots, G_r = G/_3X$ such that $G_0 = (\{s,t\}, \{e_\alpha, e_\beta\})$, and G_{i-1} is obtained from G_i by some appropriate 2-contraction (cf. the proof of Lemma 7.13). We show that almost the same 2-contractions can be applied to G, which implies $(G, s, t) \in \mathcal{D}_{\alpha,\beta}$, a contradiction.

Let $j \in [r]$ be the maximum index such that $N_G(X) \cap (V(G_j) \setminus V(G_{j-1})) \neq \emptyset$ (note that $|V(G_0)| = 2$ and $|N_G(X)| = 3$). Then, we can apply to G the same 2-contractions as those to construct G_j from $G_r = G/_3X$. Let H_j be the resulting graph, $Y := V(G_j) \setminus V(G_{j-1})$ (i.e., $G_{j-1} = G_j/_2Y$), and $Z := X \cup Y \subseteq V(H_j)$. Then, $N_{H_j}(X) \subseteq Y \cup N_{G_j}(Y)$ since x and y are adjacent in G_j for any distinct $x, y \in N_{H_j}(X)$. Hence, X is 3-contractible also in $H_j[\![Z]\!]$, and we have $G_j[\![Y]\!] = H_j[\![Z]\!]/_3X - E(N_{G_j}(Y))$. This implies that the 2-contraction of Z in H_j does not violate the condition of $\mathcal{D}_{\alpha,\beta}$ (see Definition 7.11) since neither does that of Y in G_j, and $H_j/_2Z = G_j/_2Y$. Thus we have $(G, s, t) \in \mathcal{D}_{\alpha,\beta}$, a contradiction. $\square$

Fix an arbitrary arc $e_0 = sv_0 \in \delta_G^{\mathrm{out}}(s)$, and let $G' := G - e_0$. Note that G contains no arc between s and t by Claim 7.17, and hence $v_0 \neq t$. We next show the following claims, which lead to $(G', s, t) \in \mathcal{D}_{\alpha,\beta}$.

Claim 7.19. $(G', s, t) \in \mathcal{D}$.

Proof. By Proposition 2.1, it suffices to show that $G' + r + rs + rt$ is 2-connected. Suppose to the contrary that it is not 2-connected, i.e., there exists a 1-cut $w \in V$ separating some vertex from both s and t (possibly $w \in \{s,t\}$). If $w = s$, then $G - s$ is not connected, which contradicts $(G, s, t) \in \mathcal{D}$. Otherwise, $\{s, w\}$ is a 2-cut in G, and hence G contains a 2-contractible vertex set $X \subseteq V \setminus \{s,t\}$ with $N_G(X) = \{s, w\}$, which contradicts Claim 7.17. $\square$

Claim 7.20. $l(G'; s, t) = \{\alpha, \beta\}$.

Proof. Since each s–t path in G' is also in G, we have $l(G'; s, t) \subseteq l(G; s, t) = \{\alpha, \beta\}$. Suppose to the contrary that $|l(G'; s, t)| = 1$. Then, G' is balanced by Claim 7.19 and Proposition 7.1, and hence $G - s$ is also balanced. This implies that $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta} \subseteq \mathcal{D}_{\alpha,\beta}$ (in Case (A) in Definition 7.7), a contradiction. $\square$

By Claims 7.19 and 7.20 and the minimality of G, we have $(G', s, t) \in \mathcal{D}_{\alpha,\beta}$. We consider the following two cases separately: when $(G', s, t) \in \mathcal{D}^1_{\alpha,\beta}$ and when not. That is, the former case does not need any 2-contraction for G', and the latter involves some.

7.4.1 Case 1: Without 2-contraction

Suppose that $(G', s, t) \in \mathcal{D}^1_{\alpha,\beta}$. By Claim 7.18, if G' contains a 3-contractible vertex set $X \subseteq V \setminus \{s,t\}$, then X must contain the head v_0 of e_0. Hence, if we choose a maximal

3-contractible vertex set X, then we have $(G'/_3 X, s, t) \in \mathcal{D}^0_{\alpha,\beta}$. Define $\tilde{G} := G'/_3 X$ in this case, and $\tilde{G} := G'$ otherwise, so that $(\tilde{G}, s, t) \in \mathcal{D}^0_{\alpha,\beta}$. We discuss the three cases in Definition 7.7 separately. Recall that we may assume $\beta = 1_\Gamma$ and $\alpha^{-1} \neq \alpha$ (i.e., $\alpha^2 \neq 1_\Gamma$).

Case 1.1. When $(\tilde{G}, s, t)$ is in Case (A).

Note that any 3-contraction does not make an effect on this situation (i.e., either all unbalanced cycles in G' intersect s, or they do t) since it just replaces a balanced subgraph with a balanced triangle, and hence we may assume that $\tilde{G} = G'$ and G' satisfies the condition of Case (A) (by shifting at some vertices in $V \setminus \{s, t\}$ in advance of removing e_0 if necessary). Since G contains no 2-contractible vertex set, $G - \{s, t\}$ is connected, which implies that there exists a v_0–w path in $G - \{s, t\}$ for each neighbor $w \in N_G(t)$ (recall that $v_0 \neq t$). Therefore, if $e_0 = sv_0 \in \delta^{\mathrm{out}}_G(s)$ violates the condition of Case (A) (i.e., $\psi_G(e_0) \notin \{1_\Gamma, \alpha\}$ in the former case, and $\psi_G(e_0) \neq 1_\Gamma$ in the latter case), then it is easy to see that $|l(G; s, t)| \geq 3$ (see Figs. 7.9 and 7.10). Note that we use Lemma 7.16 in the latter case (let $P_1 := (s)$ and $P_2 := (s, e_0, v_0)$).

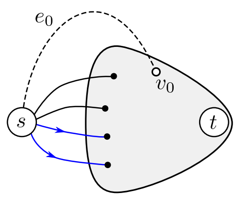

Figure 7.9: The former of Case (A).

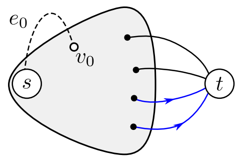

Figure 7.10: The latter of Case (A).

Case 1.2. When $(\tilde{G}, s, t)$ is in Case (B).

If $\tilde{G} = G'$, then it is easy to see $|l(G; s, t)| \geq 3$ by Lemma 7.16, since G contains no parallel arc with the same label (see Fig. 7.11). Otherwise, $\tilde{G} = G'/_3 X$ for some $X \subseteq V \setminus \{s, t\}$. If $N_{G'}(X) = \{s, v_1, v_2\}$, then $G[\![X]\!]$ is not balanced by Claim 7.18, and hence $|l(G; s, t)| \geq 3$ by Lemma 7.16 (e.g., we can take two s–v_1 paths P_1 and P_2 in $G[\![X]\!]$ with $\psi_G(P_1) \neq \psi_G(P_2)$).

Suppose that $N_{G'}(X) = \{v_3, v_4, t\}$ (see Fig. 7.12). If there exist two disjoint paths between $\{v_0, t\}$ and $\{v_3, v_4\}$ in $G[\![X]\!]$, then $|l(G; s, t)| \geq 3$ by Lemma 7.16 (e.g., we can take two s–v_1 paths P_1 and P_2 in $G[\![X + v_3]\!]$ with $\psi_G(P_1) \neq \psi_G(P_2)$ and $l(G - (V(P_i) - v_1); v_1, t) = \{1_\Gamma, \alpha\}$ $(i = 1, 2)$, if $G[\![X]\!]$ contains disjoint v_0–v_3 path and t–v_4 path). Otherwise, by Menger's theorem, $G[\![X]\!]$ contains a 1-cut $w \in X$ separating $\{v_0, t\}$ from $\{v_3, v_4\}$ (possibly $w = v_0$). In this case, $\{s, w\}$ is a 2-cut in G, which contradicts Claim 7.17.

Case 1.3. When $(\tilde{G}, s, t)$ is in Case (C).

Suppose that $\tilde{G} = (\tilde{V}, \tilde{E})$ is embedded with the conditions in Lemma 7.14 (we apply shifting at each vertex $v \in V \setminus \{s, t\}$ to G in advance of the construction of $\tilde{G}$ if necessary).

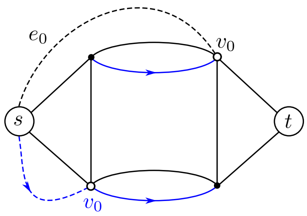

Figure 7.11: Case (B) ($\tilde{G} = G'$).

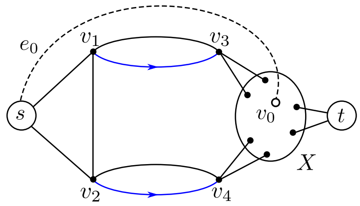

Figure 7.12: Case (B) ($\tilde{G} = G'/_3 X$).

Let $\tilde{E}^i \subseteq \tilde{E}$ be the arc set corresponding to $E^i \subseteq E$ in Lemma 7.14 for each $i \in \{0, 1\}$, and we refer to the path $P = (s = u_0, e_1, u_1, \ldots, e_l, u_l = t)$ along the outer boundary of $\tilde{G} - \tilde{E}^1$ as P itself.

In what follows, we derive a contradiction by showing that either $(G, s, t) \in \mathcal{D}_{1_\Gamma, \alpha}$, $\gamma \in l(G; s, t)$ for some $\gamma \in \Gamma \setminus \{1_\Gamma, \alpha\}$ (in particular, $\gamma = \alpha^2$ or α^{-1}), or G contains a contractible vertex set (which contradicts Claims 7.17 or 7.18). Note that $(\tilde{G}, s, t) \in \mathcal{D}$ follows from $(G', s, t) \in \mathcal{D}$, and hence $\tilde{G} - s$ is connected. Since every arc in $\tilde{E}^1$ connects two vertices on the path P in $\tilde{G} - \tilde{E}^1$, $\tilde{G} - \tilde{E}^1 - s$ is also connected. Hence, we have $\psi_G(e_0) \in l(G; s, t) = \{1_\Gamma, \alpha\}$, and consider the following two cases separately: when $\psi_G(e_0) = 1_\Gamma$, and when $\psi_G(e_0) = \alpha$.

Note that we have $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s) \neq \emptyset$. To see this, suppose that $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s) = \emptyset$. In this case, $G - s$ as well as $\tilde{G} - s$ is balanced, which implies that $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta} \subseteq \mathcal{D}_{\alpha,\beta}$ in Case (A) in Definition 7.7. We can also see $\tilde{E}^1 \setminus \delta_{\tilde{G}}(t) \neq \emptyset$ in the same way.

We first discuss the case when $\tilde{G} = G'$, and later explain that the case when $\tilde{G} = G'/_3 X$ for some $X \subseteq V \setminus \{s, t\}$ can be dealt with in almost the same way with the aid of Theorem 1.8. Assume $\tilde{G} = G' = G - e_0$, and let $\tilde{F}_0$ and $\tilde{F}'_0$ denote the outer faces of $\tilde{G}$ and $\tilde{G} - s$, respectively.

Case 1.3.1. When $\psi_G(e_0) = 1_\Gamma$.

Let us begin with an easy case: when $v_0 \in V(\mathrm{bd}(\tilde{F}'_0))$.

Case 1.3.1.1. Suppose that $v_0 \in V(\mathrm{bd}(\tilde{F}'_0)) \setminus V(P)$. In this case, we can embed $G = \tilde{G} + e_0$ on a plane by adding $e_0 = sv_0$ on $\tilde{F}_0$ so that (G, s, t) satisfies the conditions of Case (C), a contradiction.

Case 1.3.1.2 (Fig. 7.13)**.** Otherwise, $v_0 = u_h \in V(\mathrm{bd}(\tilde{F}'_0)) \cap V(P)$. Take an s–t path P' so that $(P' \cup P) - s$ forms the outer boundary of $\tilde{G} - \tilde{E}^1 - s$. Let j be the minimum index such that $E(P[u_j, t]) \subseteq E(P')$, and i the index such that $P[u_i, u_j] \cup P'[u_i, u_j]$ forms a cycle (i.e., they intersect only at u_i and u_j).

Take an arc $e' = u_{i'} u_{j'} \in \tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ so that $j' - i'$ is maximized. If $j' \leq i$, then G contains a 2-cut $\{s, u_i\}$ separating $u_{i-1} \neq s$ from $t \neq u_i$, which contradicts Claim 7.17. Hence, we have $i < j'$.

If $v_0 = u_h \in V(P) \cap V(P')$ or $h \leq i'$ ((a) in Fig. 7.13), then we can embed $e_0 = sv_0$ without violating the conditions of Case (C). Otherwise, we have $j' \leq h < j$ ((b) in Fig. 7.13) since $u_h = v_0 \in V(\mathrm{bd}(\tilde{F}_0')) \cap V(P)$. In this case, we can construct an s–t path of label $\alpha^{-1} \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G, a contradiction, e.g., by concatenating e_0, $\bar{P}[u_h, u_{j'}]$, $\bar{e}'$, $P[u_{i'}, u_i]$, $P'[u_i, u_j]$, and $P[u_j, t]$ if $0 < i' \leq i$.

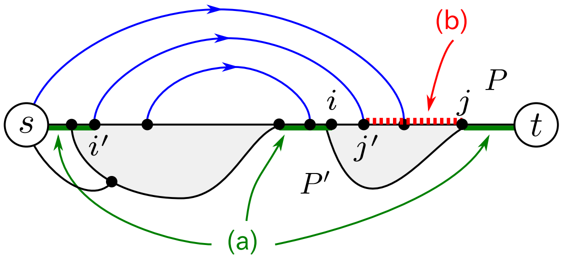

Figure 7.13: Case 1.3.1.2.

Otherwise, $v_0 \notin V(\mathrm{bd}(\tilde{F}_0'))$. Take a path Q in $\tilde{G} - \tilde{E}^1 - E(P) - s$ from $u_i \in V(P)$ to $u_j \in V(P)$ with $0 < i < j$ so that $Q \cup P[u_i, u_j]$ forms a cycle that encloses v_0 (possibly $v_0 \in V(P)$), i.e., $V(Q \cup P[u_i, u_j])$ separates v_0 from both of s and t in $\tilde{G}$ (or $v_0 = u_h \in V(P)$ with $i < h < j$). If there are multiple choices of Q, then choose Q so that the region enclosed by $Q \cup P[u_i, u_j]$ is maximized.

If $V(Q)$ separates v_0 from $V(P)$ in $\tilde{G}$, then G contains a 3-contractible vertex set $X \subseteq V \setminus V(P)$ such that $v_0 \in X$ and $N_G(X) = \{s, w_1, w_2\}$, which contradicts Claim 7.18, where $w_1, w_2 \in V(Q)$ are the vertices closest $u_i, u_j \in V(P) \cap V(Q)$, respectively, among those which are reachable from v_0 in $\tilde{G}$ without intersecting Q in between. Thus we can take a v_0–u_h path R in $\tilde{G} - V(Q)$ (possibly of length 0, i.e., $v_0 = u_h$) with $i < h < j$. If there are multiple choices of R, then choose R so that h is maximized under the condition that $V(R) \cap V(P) = \{u_h\}$.

Case 1.3.1.3 (Fig. 7.14)**.** Suppose that there is no arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ incident to an inner vertex on $P[u_i, u_j]$. If every arc in $\tilde{E}^1 \cap \delta_{\tilde{G}}(s)$ enters a vertex on $P[s, u_i] \cup P[u_j, t]$, then G contains a 3-contractible vertex set $X \subseteq V \setminus \{s, u_i, u_j\}$ such that $v_0 \in X$ and $N_G(X) = \{s, u_i, u_j\}$, a contradiction. Otherwise, every arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s) \neq \emptyset$ enters a vertex on $P[s, u_i]$. Then, G contains a 2-cut $\{s, u_i\}$ separating u_{i-1} from t (note that $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s) \neq \emptyset$ implies that $i > 1$), which contradicts Claim 7.17.

Case 1.3.1.4 (Fig. 7.15)**.** Suppose that there exists an arc $e' = u_{i'}u_{j'} \in \tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ such that $i' < h$ and $i < j' < j$. In this case, we can construct an s–t path of label $\alpha^{-1} \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G, a contradiction, e.g., by concatenating e_0, R, $P[u_h, u_{j'}]$, $\bar{e}'$, $\bar{P}[u_{i'}, u_i]$, Q, and $P[u_j, t]$ if $i \leq i'$ and $h \leq j'$.

Case 1.3.1.5 (Fig. 7.16)**.** Suppose that every arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ connects two vertices on $P[u_h, t]$. In this case, every arc in $\tilde{E}^1 \cap \delta_{\tilde{G}}(s)$ also enters a vertex on $P[u_h, t]$, and $v_0 \neq u_h$ since $v_0 \notin V(\mathrm{bd}(\tilde{F}'_0))$. Let w be the vertex closest to u_j among those on Q which are reachable from v_0 in $G - u_h$ without intersecting Q in between. By the maximality of j and h (i.e., the choice of Q and R), $\{s, u_h, w\}$ separates $v_0 \in V \setminus \{s, u_h, w\}$ from $V(P[u_h, t])$ in G, and hence G contains a 3-contractible vertex set $X \subseteq V \setminus \{s, u_h, w\}$ such that $v_0 \in X$ and $N_G(X) = \{s, u_h, w\}$, a contradiction.

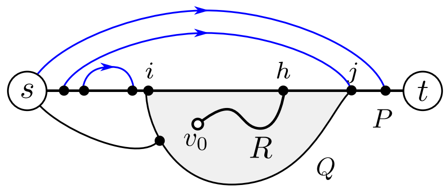

Figure 7.14: Case 1.3.1.3.

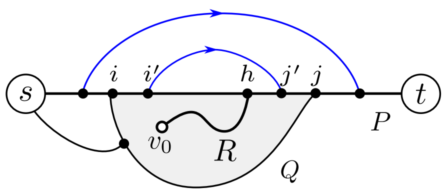

Figure 7.15: Case 1.3.1.4.

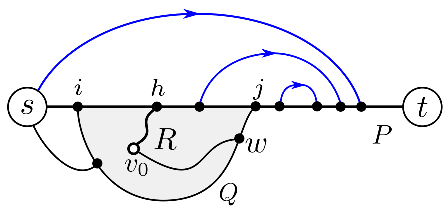

Figure 7.16: Case 1.3.1.5.

These three cases imply that there exists an arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ entering a vertex on $P[u_j, t]$. To see this, suppose to the contrary that every such arc enters a vertex on $P[u_1, u_{j-1}]$, and take $e' = u_{i'}u_{j'} \in \tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ so that $j' - i'$ is maximized. We may assume $i < j'$ by Case 1.3.1.3, and hence $h \leq i'$ by Case 1.3.1.4, which leads to the condition of Case 1.3.1.5, a contradiction. This implies also that no arc in $\tilde{E}^1 \cap \delta_{\tilde{G}}(s)$ enters a vertex on $P[u_1, u_{j-1}]$.

Case 1.3.1.6 (Fig. 7.17)**.** Suppose that all arcs in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ leave the same vertex $u_{i^*} \in V(P)$ with $i^* < h$. In this case, by Case 1.3.1.4, we may assume that every arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ enters a vertex on $P[u_j, t]$. Then, since $\{s, u_{i^*}, u_j\}$ separates $v_0 \in V \setminus \{s, u_{i^*}, u_j\}$ from $V(P[u_j, t])$ in G, there exists a 3-contractible vertex set $X \subseteq V \setminus \{s, u_{i^*}, u_j\}$ in G such that $v_0 \in X$ and $N_G(X) = \{s, u_{i^*}, u_j\}$, a contradiction.

Case 1.3.1.7 (Fig. 7.18)**.** Suppose that all arcs in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ enter the same vertex $u_{j^*} \in V(P)$ with $j \leq j^*$. In this case, $\{s, u_j, u_{j^*}\}$ separates $v_0 \in V \setminus \{s, u_j, u_{j^*}\}$ from $V(P[u_j, t])$ in G, and hence G contains a contractible vertex set $X \subseteq V \setminus \{s, u_j, u_{j^*}\}$ such that $v_0 \in X$ and $N_G(X) = \{s, u_j, u_{j^*}\}$, a contradiction.

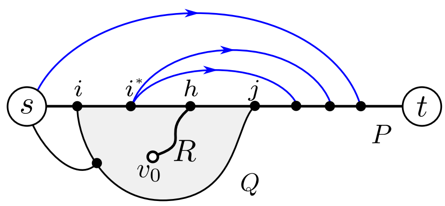

Figure 7.17: Case 1.3.1.6.

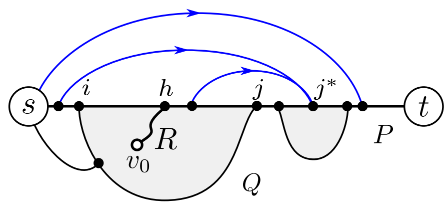

Figure 7.18: Case 1.3.1.7.

Otherwise, there exist two arcs $e_1 = u_{i_1}u_{j_1}$ and $e_2 = u_{i_2}u_{j_2}$ in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ such that $i_2 < i_1 < j_1 < j_2$ by Cases 1.3.1.6 and 1.3.1.7. We choose e_2 so that $j_2 - i_2$ is maximized. We then have $i_2 < h$ by Case 1.3.1.5, and $j \leq j_2$ by the argument just after Case 1.3.1.5. Since there exists an arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ incident to an inner vertex on $P[u_i, u_j]$ by Case 1.3.1.3, we can choose e_1 so that $i < i_1$ (which is obvious if $i \leq i_2$, and follows from Case 1.3.1.4 otherwise). We then have $h < j_1$, since otherwise we have $i < i_1 < j_1 \leq h < j$, which implies that e_1 satisfies the condition of Case 1.3.1.4. We choose e_1 so that i_1 is minimized under the condition that $i < i_1$.

Case 1.3.1.8 (Fig. 7.19)**.** Suppose that $j \leq i_1$. In this case, $\{s, u_{i_2}, u_j\}$ separates $v_0 \in V \setminus \{s, u_{i_2}, u_j\}$ from $P[u_j, t]$ in G, and hence G contains a 3-contractible vertex set $X \subseteq V \setminus \{s, u_{i_2}, u_j\}$ such that $v_0 \in X$ and $N_G(X) = \{s, u_{i_2}, u_j\}$, a contradiction.

Case 1.3.1.9 (Figs. 7.21 and 7.22)**.** Suppose that $j_2 = j$. We then have $h \leq i_1$ by $i < i_1 < j_1 < j_2 = j$ and Case 1.3.1.4. Let h^* be the maximum index such that there exists a w–u_{h^*} path R^* in $\tilde{G} - u_j$ for some $w \in (V(Q) \setminus V(P)) + v_0$ such that $V(R^*) \cap V(Q) \subseteq \{w\}$ and $V(R^*) \cap V(P) = \{u_{h^*}\}$. Note that $h \leq h^*$. If $i_1 < h^*$, then we have $h < h^*$ because of $h \leq i_1$. In this case (see Fig. 7.21), since R and R^* are disjoint by the maximality of h and h^*, we can construct an s–t path of label $\alpha^2 \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G, a contradiction, e.g., by concatenating e_0, R, $P[u_h, u_{i_1}]$, e_1, $\bar{P}[u_{j_1}, u_{h^*}]$, $\bar{R}^*$, $\bar{Q}[w, u_i]$, $\bar{P}[u_i, u_{i_2}]$, e_2, and $P[u_j, t]$ if $h^* \leq j_1$ and $i_2 \leq i$. Otherwise (i.e., if $h^* \leq i_1$), by the minimality of i_1 and the maximality of h^*, there exists a 2-cut $\{u_{h^*}, u_j\}$ separating u_{j_1} from u_i ($i < h \leq h^* \leq i_1 < j_1 < j_2 = j$) in G (see Fig. 7.22), a contradiction.

Case 1.3.1.10 (Fig. 7.20)**.** Otherwise, we have $i < i_1 < j < j_2$ (also recall that $i_2 < i_1 < j_1 < j_2$ and $i_2 < h < j_1$). In this case, we can construct an s–t path of label $\alpha^2 \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G, a contradiction, e.g., by concatenating e_0, R, $\bar{P}[u_h, u_{i_1}]$, e_1, $\bar{P}[u_{j_1}, u_j]$, $\bar{Q}$, $P[u_i, u_{i_2}]$, e_2, and $P[u_{j_2}, t]$ if $i_1 \leq h$, $j \leq j_1$, and $i \leq i_2$.

Case 1.3.2. When $\psi_G(e_0) = \alpha$.

This case is rather easier than Case 1.3.1. Note that, if there exists a v_0–t path of label α in $\tilde{G} = G' = G - e_0$, then we can construct an s–t path of label $\alpha^2 \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G,

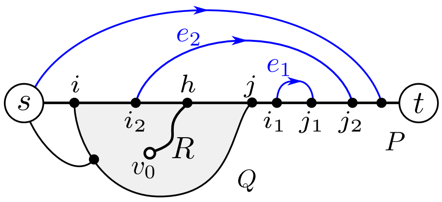

Figure 7.19: Case 1.3.1.8.

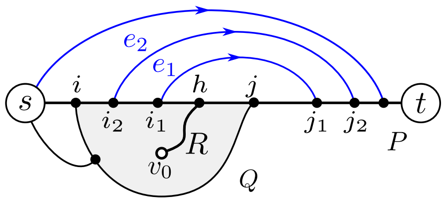

Figure 7.20: Case 1.3.1.10.

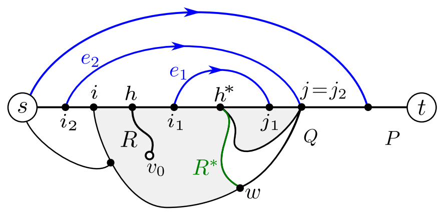

Figure 7.21: Case 1.3.1.9 (label α^2).

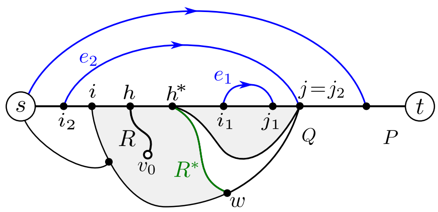

Figure 7.22: Case 1.3.1.9 (a 2-cut $\{u_{h^*}, u_j\}$).

a contradiction, by extending the v_0–t path using $e_0 = sv_0$. Hence, we may assume that $\tilde{G}$ contains no such path.

Case 1.3.2.1. Suppose that $v_0 = u_h \in V(P)$. If there exists an arc $e' = u_{i'}u_{j'} \in \tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ with $h < j'$, then we can construct a v_0–t path of label α, a contradiction, e.g., by concatenating $P[u_h, u_{i'}]$, e', and $P[u_{j'}, t]$ if $h \leq i'$. Otherwise, every arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s) \neq \emptyset$ connects two vertices on $P[u_1, u_h]$. Hence, we can embed $e_0 = su_h$ without violating the conditions of Case (C) in Definition 7.7 (cf. Lemma 7.14).

Case 1.3.2.2 (Fig. 7.23)**.** Otherwise, $v_0 \notin V(P)$. Let i and j be the minimum and maximum indices, respectively, such that there exist a v_0–u_i path Q and a v_0–u_j path R in $\tilde{G} - \tilde{E}^1 - s$ that do not intersect P in between. If there exists an arc $e' = u_{i'}u_{j'} \in \tilde{E}^1 \setminus \delta_{\tilde{G}}(s)$ with $i < j'$, then we can construct a v_0–t path of label α, a contradiction, e.g., by concatenating Q, $\bar{P}[u_i, u_{i'}]$, e', and $P[u_{j'}, t]$ if $i' \leq i$.

Otherwise, every arc in $\tilde{E}^1 \setminus \delta_{\tilde{G}}(s) \neq \emptyset$ connects two vertices on $P[u_1, u_i]$. Since G contains no 3-contractible vertex set (by Claim 7.18), there exists an arc from s to the connected component of $\tilde{G} - \{s, u_i, u_j\}$ that contains v_0 with label 1_Γ in $\tilde{G}$. Hence, because of the minimality of i and the planarity of $\tilde{G}$, there is no path from an inner vertex on $P[s, u_i]$ to a vertex on $P[u_j, t]$ in $\tilde{G} - \tilde{E}^1 - s$ which does not intersect P in between. This implies that G contains a 2-cut $\{s, u_i\}$ separating $u_1 \neq u_i$ from t, a contradiction.

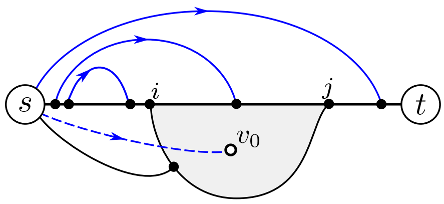

Figure 7.23: Case 1.3.2.2.

Case 1.3.3. When $\tilde{G} = G'/_3 X$ for some $X \subseteq V \setminus \{s,t\}$.

Recall that X must contain v_0 by Claim 7.18. Suppose that $N_{G'}(X) = \{y_1, y_2, y_3\}$. Since $\tilde{G}$ is embedded as Lemma 7.14, the resulting triangle $y_1y_2y_3$ of the 3-contraction of X (which is a balanced cycle by the definition) consists of either three arcs in $\tilde{E}^0$ or one arc in $\tilde{E}^0$ and two arcs in $\tilde{E}^1$. Without loss of generality (by the symmetry of y_1, y_2, y_3), assume that the arc between y_2 and y_3 is in $\tilde{E}^0$, i.e., $l(G'[\![X]\!]; y_2, y_3) = 1_\Gamma$. Then, by shifting at vertices in X in advance of removing $e_0 = sv_0$ from G if necessary, we may assume that the label of every arc in $G'[\![X]\!] - y_1$ is 1_Γ and in $\delta_{G'[\![X]\!]}(y_1)$ is γ, where γ is a fixed element in $\{1_\Gamma, \alpha, \alpha^{-1}\}$ and all arcs in $\delta_{G'[\![X]\!]}(y_1)$ are assumed to enter y_1 (recall that $G'[\![X]\!]$ is balanced by Definition 7.9).

Let $\tilde{G}'$ be the Γ-labeled graph obtained from G' by the following procedure:

- merge all vertices in X into v_0,
- identify parallel arcs with the same label as a single arc, and
- for each $\{i,j,k\} = \{1,2,3\}$, add an arc from y_j to y_k with label $l(G'[\![X]\!]; y_j, y_k)$ if there is no such arc and there are disjoint v_0–y_i path and y_j–y_k path in $G'[\![X]\!]$ (note that otherwise, by Theorem 1.8, $G'[\![X]\!]$ can be embedded on a plane so that v_0, y_j, y_i, y_k are on the outer boundary in this order).

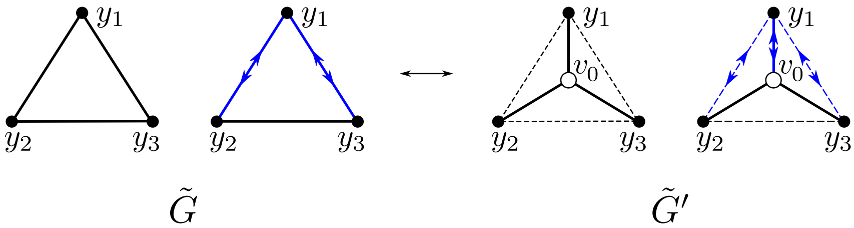

Figure 7.24: Corresponding parts of $\tilde{G}$ and $\tilde{G}'$.

Since $\tilde{G}$ is embedded as Lemma 7.14, we can naturally embed $\tilde{G}'$ so (see Fig. 7.24). By the same argument for $\tilde{G}'$ as Cases 1.3.1 and 1.3.2, we can derive a contradiction in this case. Note that, if we can construct an s–t path of label $\gamma \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in $\tilde{G}' + e_0$, then it can be expanded to one in $G = G' + e_0$ (which may use disjoint v_0–y_i path and y_j–y_k path in $G'[\![X]\!]$). Besides, if we can embed $\tilde{G}' + e_0$ as Lemma 7.14, then the embedding can be expanded to one of G without violating the conditions, since any embedding of $\tilde{G}'$ with v_0 exposed on the outer boundary can be expanded to one of G' so by Theorem 1.8. Note that, for any k-cut ($k \in \{2, 3\}$) separating some vertex set from $\{v_0, y_1, y_2, y_3\}$ in $G'[\![X]\!]$, we can perform the k-contraction, respectively, which emulates the operation in Condition 3 in Theorem 1.8, since $G'[\![X]\!]$ is balanced.

7.4.2 Case 2: Involving 2-contraction

Suppose that $(G', s, t) \in \mathcal{D}_{\alpha,\beta} \setminus \mathcal{D}^1_{\alpha,\beta}$. In this case, G' contains a 2-contractible vertex set $X \subseteq V \setminus \{s, t\}$ by the definition of $\mathcal{D}_{\alpha,\beta}$ (see Definition 7.11). Due to the previous section, we may assume that this situation occurs regardless of the choice of the arc $e_0 = sv_0 \in \delta^{\mathrm{out}}_G(s)$, which has at least two possibilities by Lemma 7.13-(2). We first show a useful claim about such a vertex set (in fact, slightly more general).

Claim 7.21. *Let $X \subseteq V \setminus \{s, t\}$ be a vertex set with $N_G(X) = \{s, x, y\}$ for some distinct vertices $x, y \in V$ (see Fig. 7.25). Then, $s \notin \{x, y\}$, $G[\![X]\!]$ is not balanced, and $(G[\![X]\!] - x, s, y) \in \mathcal{D}$. Moreover, if $|l(G[\![X]\!]; s, y)| = 1$, then $X = \{v\}$ for some $v \in V \setminus \{s, x, y\}$ and $G[\![X]\!]$ consists of the following four arcs (see Fig. 7.26): one between s and v, one between v and y, and two parallel arcs between v and x.*

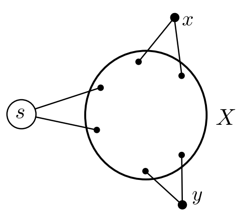

Figure 7.25: The situation of Claim 7.21.

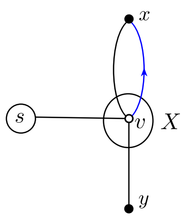

Figure 7.26: When $|l(G[\![X]\!]; s, y)| = 1$.

Proof. If $s \in \{x, y\}$, then X is 2-contractible in G, which contradicts Claim 7.17. Besides, if $G[\![X]\!]$ is balanced, then X is 3-contractible in G, which contradicts Claim 7.18.

Suppose to the contrary that $(G[\![X]\!] - x, s, y) \notin \mathcal{D}$. Then, $G[\![X]\!] - x + sy$ contains a 1-cut $w \in X \cup \{s, y\}$ by Proposition 2.1. The vertex set of the connected component of $G[\![X]\!] - \{w, x\} + sy$ that contains none of s and y is separated from both s and t by $\{w, x\}$ in G (possibly $t \in \{w, x\}$), and hence it is 2-contractible, a contradiction.

Moreover, suppose that $|l(G[\![X]\!]; s, y)| = 1$, which leads to $(G[\![X]\!], s, y) \notin \mathcal{D}$ by Proposition 7.1. Then, $G[\![X]\!] - x$ is balanced since $(G[\![X]\!] - x, s, y) \in \mathcal{D}$, which also implies that $G[\![X]\!] + sy$ contains a unique 1-cut $w \in X$. The 1-cut w separates x from the balanced component $G[\![X]\!] - x$, and hence there are two parallel arcs between w and x (which form an unbalanced cycle). Besides, if $X - w \neq \emptyset$, then G contains a contractible vertex set $Y \subseteq X - w$ with $N_G(Y) \subseteq \{s, w, y\}$, a contradiction. Thus we have done. $\square$

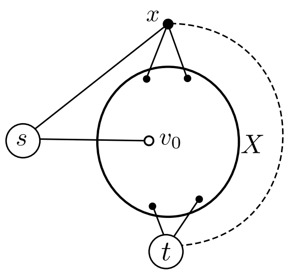

Figure 7.27: Case 2.1.1.

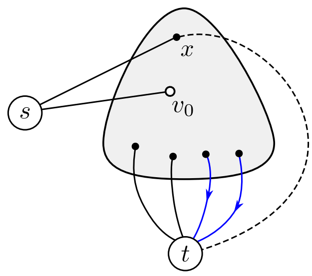

Figure 7.28: Case 2.1.1.1.

Choose a minimal 2-contractible vertex set X in G', and let $N_{G'}(X) = \{x, y\}$. We then have $v_0 \in X$ and $s \notin \{x, y\}$ by Claim 7.17 ($G = G' + e_0$ contains no 2-contractible vertex set), and $(G'[\![X]\!], x, y) \in \mathcal{D}^1_{\alpha',\beta'}$ for some distinct $\alpha', \beta' \in \Gamma$ with $\alpha'\beta'^{-1} \neq \beta'\alpha'^{-1}$ by Lemma 7.15 and Claims 7.19–7.21. Besides, $G[X]$ must be connected, since otherwise some connected component of $G[X]$ does not contain v_0 and hence its vertex set is contractible in G, which contradicts Claim 7.17 or 7.18.

Case 2.1. When $t \in \{x, y\}$.

Without loss of generality (by the symmetry of x and y), we may assume that $y = t$.

Case 2.1.1. When $V = X \cup \{s, x, t\}$.

Recall that G contains no arc between s and t. Hence, by Lemma 7.13-(2), G contains an arc between s and x, and there exists exactly one such arc $e = sx \in E$ (see Fig. 7.27), since $(G'[\![X]\!], x, y) \in \mathcal{D}^1_{\alpha',\beta'}$ and $|l(G; s, t)| = 2$. We assume $\psi_G(e) = 1_\Gamma$ by shifting at x if necessary. In the same way as the previous section, let $\tilde{G} := G'[\![X]\!]/_3 Y$ for a maximal 3-contractible vertex set $Y \subseteq X$ with $v_0 \in Y$ if exists, and $\tilde{G} := G'[\![X]\!]$ otherwise. We then have $(\tilde{G}, x, t) \in \mathcal{D}^0_{\alpha,\beta}$, and consider the three cases in Definition 7.7 separately.

Case 2.1.1.1. Suppose that $(\tilde{G}, x, t)$ is in the latter case of Case (A) (see Fig. 7.28). We may assume that the label of every arc in $E(X + x)$ is 1_Γ (by shifting at vertices in X if necessary). If $\psi_G(e_0) = 1_\Gamma$, then obviously $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$. Otherwise (i.e., if $\psi_G(e_0) \neq 1_\Gamma$), since $G[X]$ is connected, there exists a v_0–w path in $G'[\![X]\!]$ for each neighbor $w \in N_G(t)$, and hence $|l(G, s, t)| \geq 3$ by Lemma 7.16. Note that any 3-contraction does not make an effect on the above argument.

Case 2.1.1.2. Suppose that $(\tilde{G}, x, t)$ is in the former case of Case (A) (see Fig. 7.29). We may assume that the label of every arc in $E(X + t)$ is 1_Γ and in $\delta_G(x) - e$ leaving x is 1_Γ or α with $\alpha^{-1} \neq \alpha$ (recall that we may assume $\beta = 1_\Gamma$ by Lemma 7.13-(1)). Note again that any 3-contraction does not make an effect on whether $(\tilde{G}, x, t)$ is in Case (A) or not, and hence we may assume that $\tilde{G} = G'[\![X]\!]$.

Let H be the graph obtained from $G - s$ (which coincides with $G'[\![X]\!]$ if $xt \notin E$) by splitting x into two vertices x_0 and x_1 so that every arc leaving x in $G - s$ with label $\alpha^i \in \{1_\Gamma, \alpha\}$ leaves x_i in H for each $i = 0, 1$ (see Fig. 7.30).

Since $l(G; s, t) = \{1_\Gamma, \alpha\}$, either $\psi_G(e_0) = 1_\Gamma$ or $\psi_G(e_0) = \alpha$. Suppose that $\psi_G(e_0) = 1_\Gamma$. If H contains disjoint v_0–x_1 path P and x_0–t path Q, then we can construct an s–t path of label $\alpha^{-1} \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G by concatenating e_0, P, and Q with identifying $x_0, x_1 \in V(H)$ as $x \in V$. Otherwise, by Theorem 1.8, H can be embedded on a plane so that $v_0, x_0, x_1, t \in V(H)$ are on the outer boundary in this order (note that if there exists a vertex set $Y \subseteq V(H) \setminus \{v_0, x_0, x_1, t\} = V \setminus \{v_0, x, t\}$ such that $|N_H(Y)| \leq 3$, then either $|N_G(Y)| \leq 2$ or $|N_G(Y)| \leq 3$ and $G[\![Y]\!]$ is balanced, which contradicts Claim 7.17 or 7.18, respectively). This embedding can be easily extended to an embedding of G by merging $x_0, x_1 \in V(H)$ into $x \in V$ and by adding s, $e_0 = sv_0$, and $e = sx$, and the resulting embedding satisfies the conditions of Case (C) in Definition 7.7 (cf. Lemma 7.14), which implies $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$, a contradiction.

Otherwise, $\psi_G(e_0) = \alpha$. Also in this case, by a similar argument to the above, we can either construct an s–t path of label $\alpha^2 \in \Gamma \setminus \{1_\Gamma, \alpha\}$ in G by concatenating e_0 and disjoint v_0–x_0 path P and x_1–t path Q with identifying $x_0, x_1 \in V(H)$ as $x \in V$, or embed G so that (G, s, t) is in Case (C).

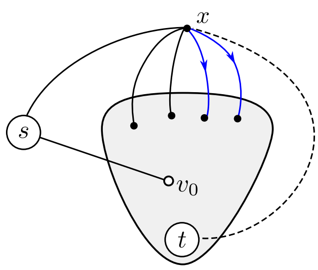

Figure 7.29: Case 2.1.1.2.

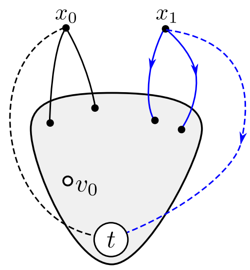

Figure 7.30: H in Case 2.1.1.2.

Case 2.1.1.3. Suppose that $(\tilde{G}, x, t)$ is in Case (B). If $\tilde{G} = G'[\![X]\!]$, it is easy to confirm that $\{x\}$ is 3-contractible in G (if there is no arc between x and t) or $|l(G; s, t)| \geq 3$ (otherwise, i.e., if $xt \in E$) by Lemma 7.16 (see Fig. 7.31).

Otherwise (i.e., if $\tilde{G} = G'[\![X]\!]/_3Y$ for some $Y \subseteq X$), we have either $N_{G'}(Y) = \{x, v_1, v_2\}$ or $N_{G'}(Y) = \{v_3, v_4, t\}$. Suppose that $N_{G'}(Y) = \{v_3, v_4, t\}$. In this case, we can derive a contradiction by Menger's Theorem in a similar way to Case 1.2. That is, $G'[\![Y]\!]$ contains either two disjoint paths between $\{v_0, t\}$ and $\{v_3, v_4\}$ or a 1-cut $w \in Y$ separating them (possibly $w = v_0$). In the former case, $|l(G; s, t)| \geq 3$ by Lemma 7.16, and in the latter case, G contains a 2-cut $\{x, w\}$ separating $\{v_3, v_4\}$ from $\{s, v_0, t\}$, which contradicts Claim 7.17.

Otherwise, $N_{G'}(Y) = \{x, v_1, v_2\}$ (see Fig. 7.32). If $xt \in E$, then we can similarly derive a contradiction by Menger's Theorem, i.e., either $|l(G; s, t)| \geq 3$ by Lemma 7.16 ($G'[\![Y]\!]$ contains two disjoint paths between $\{v_0, x\}$ and $\{v_1, v_2\}$) or G contains a 2-cut $\{w, t\}$ ($G'[\![Y]\!]$ contains a 1-cut $w \in Y + x$ separating $\{v_0, x\}$ and $\{v_1, v_2\}$). Otherwise, $N_G(Y + x) = \{s, v_1, v_2\}$. Since $Y + x$ is not 3-contractible in G by Claim 7.18, $G[\![Y + x]\!]$ is not balanced. If $|l(G[\![Y + x]\!]; s, v_1)| = 1$, then G contains a 3-contractible vertex set $Z \subseteq Y + x$ with $N_G(Z) = \{s, v_1, w\}$ for some $w \in Y$ (note that $G'[Y] = G[Y]$ is connected by Definition 7.9), a contradiction. Otherwise, i.e., if $|l(G[\![Y + x]\!]; s, v_1)| \geq 2$, we have $|l(G; s, t)| \geq 3$ by Lemma 7.16.

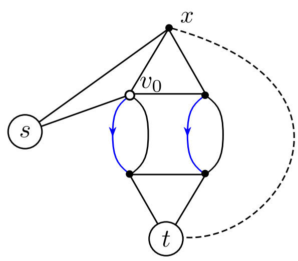

Figure 7.31: Case 2.1.1.3 ($\tilde{G} = G'[\![X]\!]$).

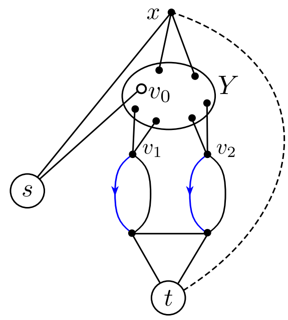

Figure 7.32: Case 2.1.1.3 ($\tilde{G} = G'[\![X]\!] /_3 Y$).

Case 2.1.1.4. Suppose that $(\tilde{G}, x, t)$ is in Case (C). In this case, by extending the x–t path P (in Lemma 7.14) to an s–t path using the arc $e = sx$, we can see that (G', s, t) (or $(G' /_3 Y, s, t)$ if $\tilde{G} = G'[\![X]\!] /_3 Y$) is also in Case (C) (see Fig. 7.33), which contradicts $(G', s, t) \notin \mathcal{D}^1_{\alpha,\beta}$.

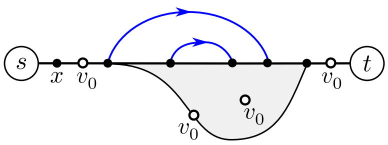

Figure 7.33: Case 2.1.1.4.

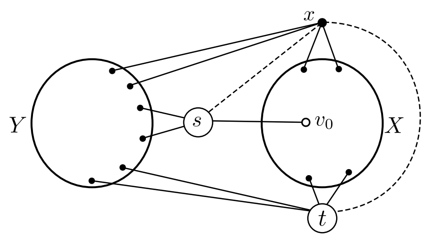

Figure 7.34: Case 2.1.2.

Case 2.1.2. When $V \setminus (X \cup \{s, x, t\}) \neq \emptyset$.

Let $Y := V \setminus (X \cup \{s, x, t\})$. Since Y is not 3-contractible in G by Claim 7.18, $G[\![Y]\!]$ is not balanced. We focus on $G - X - t$, which coincides with $G[\![Y]\!] - t$ if $sx \notin E$ (see Fig. 7.34). We have $(G[\![Y]\!] - t, s, x) \in \mathcal{D}$ by Claim 7.21, and hence $(G - X - t, s, x) \in \mathcal{D}$. Suppose that $G - X - t$ is not balanced. In this case, $|l(G - X - t; s, x)| \geq 2$ by Proposition 7.1, and hence $|l(G; s, t)| \geq 3$ by Lemma 7.16 (recall that $(G'[\![X]\!], x, t) \in \mathcal{D}_{\alpha',\beta'}$).

Otherwise, $G - X - t$ is balanced. By Claim 7.21 (with the symmetry of s and t), we have $(G - X - s, x, t) \in \mathcal{D}$ and hence $|l(G - X - s, x, t)| \geq 2$. This implies that $G[\![X]\!] - t$ is balanced, since otherwise $|l(G; s, t)| \geq 3$ by Lemma 7.16 (note that $(G[\![X]\!] - t, s, x) \in \mathcal{D}$ by Claim 7.21). In this case, by Proposition 2.5, we may assume that $\psi_G(e) = 1_\Gamma$ for every edge $e \in E \setminus (\delta_G(t) + e_0)$ (by shifting at each $v \in V \setminus \{s, t\}$ if necessary).

If $\psi_G(e_0) = 1_\Gamma$, then $G - t$ is also balanced, and hence $(G, s, t) \in \mathcal{D}^0_{\alpha,\beta}$ in the latter case of Case (A) in Definition 7.7, a contradiction. Otherwise, we have $|l(G; s, t)| \geq 3$ by Lemma 7.16 (we choose $P_1 := (s, e_0, v_0)$ and P_2 as an arbitrary s–x path in $G - X - t$, there are two arcs entering t from X with distinct labels since X is not 3-contractible in G by Claim 7.18, and recall that $G[X]$ is connected as discussed just before starting Case 2.1), a contradiction.

Case 2.2. When $t \notin \{x, y\}$.

Suppose that $V = X \cup \{s, x, y, t\}$ (see Fig. 7.35). Then, by the symmetry of x and y, we may assume that there exists an arc $e = sx \in \delta^{\text{out}}_G(s)$ such that $(G - e, s, t) \in \mathcal{D}_{\alpha,\beta} \setminus \mathcal{D}^1_{\alpha,\beta}$ (recall the discussion in the first paragraph of this section). Besides, t is adjacent to both of x and y since otherwise $\{s, y\}$ or $\{s, x\}$ is a 2-cut in G, which contradicts Claim 7.17. Hence, by choosing e instead of e_0, we can reduce this case to Case 2.1 (since x and t are adjacent, t must be a neighbor of any 2-contractible vertex set in $G - e$ that contains x).

In what follows, we assume that $Y := V \setminus (X \cup \{s, x, y, t\}) \neq \emptyset$ (see Fig. 7.36), and consider the following two cases separately: when $G - X - s$ is balanced and when not.

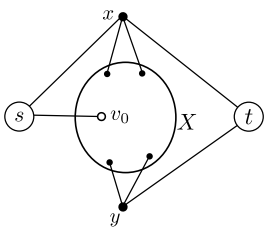

Figure 7.35: When $V = X \cup \{s, x, y, t\}$.

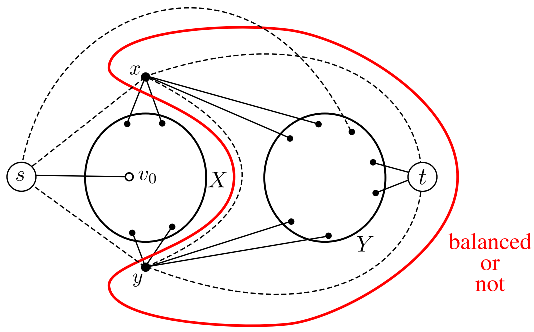

Figure 7.36: When $V \setminus (X \cup \{s, x, y, t\}) \neq \emptyset$.

Case 2.2.1. When $G - X - s$ is balanced.

Since Y is not 3-contractible in G by Claim 7.18, there exists an arc $e' = sv' \in \delta_G(s)$ with $v' \in Y$ such that $G - e'$ contains a 2-contractible vertex set $X' \subseteq V \setminus \{s, t\}$ with $v' \in X'$ and $N_{G'}(X') = \{x', y'\}$ for some distinct $x', y' \in V \setminus \{s, v'\}$ (recall that, if $G - e'$ contains no 2-contractible vertex set, then we can reduce this case to Case 1 by choosing e' instead of e_0). Choose minimal X'. If $\{x', y'\} \subseteq Y \cup \{x, y, t\}$, then $G[\![X']\!]$ is balanced and hence X' is 3-contractible in G, a contradiction. Besides, if $\{x', y'\} \subseteq X$, then $G - e'$ contains a smaller 2-contractible vertex set $X'' \subsetneq X'$ with $v' \in X''$ and $N_{G'}(X'') = \{x, y\}$.

Thus we have $|\{x', y'\} \cap X| = 1$, and assume $x' \in X$ and $y' \in Y \cup \{x, y, t\}$ (see Fig. 7.37). Let $Z \subseteq Y$ be the vertex set of the connected component of $G - \{x, y, y'\} - e'$ that contains v'. Then, since Z is not 3-contractible in G and v' is separated from both s and t by $\{x', y'\}$ in $G - e'$, we have $N_{G-e'}(Z) = \{x, y, y'\}$ and $y' \notin \{x, y\}$. If $y' = t$, then this case reduces to Case 2.1 by choosing e' instead of e_0. Otherwise, $\{s, y'\}$ is a 2-cut in G separating v' from t, which contradicts Claim 7.17.

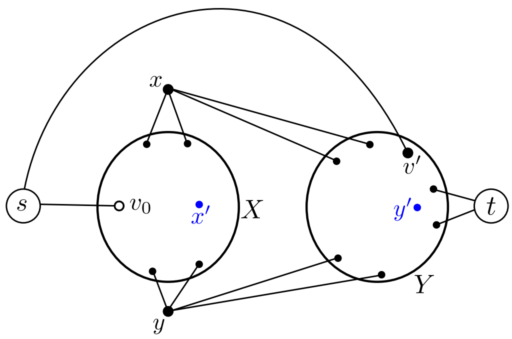

Figure 7.37: Case 2.2.1.

Case 2.2.2. When $G - X - s$ is not balanced.

Recall that $G[X]$ is connected (discussed just before starting Case 2.1). Suppose that $G[\![X]\!] - x$ and $G[\![X]\!] - y$ are balanced. Then, by Proposition 2.5, we may assume that $\psi_G(e) = 1_\Gamma$ for every $e \in E(G[\![X]\!])$ by shifting at each $v \in X \cup \{x, y\}$ if necessary. This implies that $G[\![X]\!]$ is also balanced, which contradicts Claim 7.18.

Thus, at least one of $G[\![X]\!] - x$ and $G[\![X]\!] - y$ is not balanced. By Claim 7.21 and the symmetry of x and y, $(G[\![X]\!] - x, s, y) \in \mathcal{D}$ and $(G[\![X]\!] - y, s, x) \in \mathcal{D}$. Hence, we may assume that $|l(G[\![X]\!] - y; s, x)| \geq 2$ by Proposition 7.1. Note that $G - X - s - y$ contains an x–t path (otherwise, $\{s, y\}$ is a 2-cut in G separating x from t, which contradicts Claim 7.17). This implies $|l(G[\![X]\!] - y; s, x)| = 2$ since $|l(G; s, t)| = 2$.

Let $Z \subseteq Y \cup \{x, y, t\}$ be the set of vertices that are contained in some x–t path in $G - X - s$ (i.e., the vertex set of the 2-connected component of $(G - X - s) + r + rx + rt$ that contains both of x and t, except for r, by Proposition 2.1). Then, $(G[Z], x, t) \in \mathcal{D}$. If $G[Z]$ is not balanced, then $|l(G[Z]; x, t)| \geq 2$ by Proposition 7.1, and hence we derive $|l(G; s, t)| \geq 3$ from $|l(G[\![X]\!] - y; s, x)| = 2$ by Lemma 7.16 (note that there exist $\alpha', \beta' \in l(G[Z]; x, t)$ such that $\alpha' \beta'^{-1} \neq \beta' \alpha'^{-1}$ by Lemma 7.15). Hence, we assume that $G[Z]$ is balanced, which implies that $Z \neq Y \cup \{x, y, t\}$ (note that $G[Y \cup \{x, y, t\}] = G - X - s$).

Case 2.2.2.1. Suppose that $y \in Z$. Let $W := Y \setminus Z \neq \emptyset$. Since $G[Z] + r + rx + rt$ is a 2-connected component of $G - X - s + r + rx + rt$, we have $|N_{G-s}(W)| \leq 1$ (see Fig. 7.38). This implies that W is 2-contractible in G, which contradicts Claim 7.17.

Case 2.2.2.2. Suppose that $Z = Y \cup \{x, t\}$. Note that $G - X - s - x$ contains a y–t path since $\{s, x\}$ is not a 2-cut in G. Hence, G contains no arc between x and y, and there uniquely exists a neighbor $z \in N_{G-X-s}(y)$ with $z \neq x$. Recall that $G - X - s$ is not balanced, which implies that there are parallel arcs between y and z (see Fig. 7.39). By the definition of Z, $G[Z] - x$ contains a z–t path (possibly of length 0). Hence, by Lemma 7.16, we may assume that $|l(G[\![X]\!]; s, y)| = 1$.

In this case, $X = \{v_0\}$ and $G'[\![X]\!]$ consists of an arc between v_0 and y and two parallel arcs between v_0 and x, by Claim 7.21. Suppose that there exists an arc e' from s to $z' \in Z$ in G. If $G[Z]$ contains two disjoint paths between $\{z', t\}$ and $\{x, z\}$ (possibly of length 0, e.g., $z' = z$), then we derive $|l(G; s, t)| \geq 3$ by Lemma 7.16 (e.g., if $G[Z]$ contains disjoint z'–z path and x–t path, then let P_1, P_2 be two s–y paths obtained from by extending the z'–z path using $e' = sz'$ and the parallel arcs between y and z). Otherwise, by Menger's Theorem, $G[Z]$ contains a 1-cut $w \in Z - t$ separating them, which implies that $\{s, w\}$ is a 2-cut in G, contradicting Claim 7.17.

Thus we have $s \notin N_G(Z)$. Since $G[Z]$ is balanced and contains no contractible vertex set, $Z = \{x, z, t\}$ (note that $z \notin \{x, t\}$ since $Y \neq \emptyset$). By Lemma 7.13-(2), there must be single arcs between s and y and between x and t, which leads to Case (B) in Definition 7.7. Note that the labels of arcs are easily confirmed according to $l(G; s, t) = \{1_\Gamma, \alpha\}$.

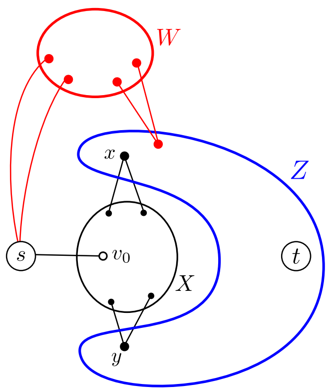

Figure 7.38: Case 2.2.2.1.

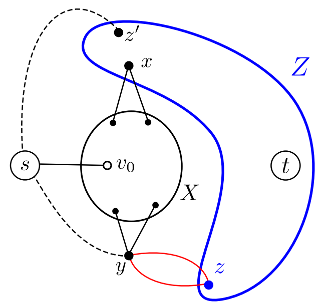

Figure 7.39: Case 2.2.2.2.

Case 2.2.2.3. Otherwise, $Z \subsetneq Y \cup \{x, t\}$. Let $W := Y \setminus Z \neq \emptyset$. By the definition of Z, we have $N_{G-s}(W) \subseteq \{y, z\}$ for some $z \in Z - x$. Since G contains no 2-contractible vertex set by Claim 7.17, we have $N_G(W) = \{s, y, z\}$ (see Fig. 7.40). If $|l(G[\![W]\!] - z; s, y)| \geq 2$, then we derive $|l(G; s, t)| \geq 3$ from $(G'[\![X]\!], x, y) \in \mathcal{D}^1_{\alpha', \beta'}$ by Lemma 7.16.

Hence, suppose that $|l(G[\![W]\!]-z;s,y)|=1$. Then, $(G[\![W]\!]-z,s,y)\notin\mathcal{D}$ or $G[\![W]\!]-z$ is balanced by Proposition 7.1. In the former case, $G[\![W]\!]-z$ contains a 1-cut $w\in W$, which implies that $\{w,z\}$ is a 2-cut in G separating some vertex from both s and t, contradicting Claim 7.17. In the latter case, there are parallel arcs between z and $w\in W$, since $G[\![W]\!]-s$ is not balanced (recall that $G-X-s$ is not balanced and $G[Z]$ is balanced). If $W\neq\{w\}$, then G contains a contractible vertex set $W'\subsetneq W$ with $N_G(W')\subseteq\{s,w,y\}$ (see Fig. 7.41), which contradicts Claim 7.18. Besides, by Claim 7.21, we have $X=\{v_0\}$ since otherwise we derive $|l(G;s,t)|\geq 3$ from $|l(G[\![X]\!];s,y)|\geq 2$. Thus, $\{y\}$ is contractible or there are parallel arcs from s to y, which also leads to $|l(G;s,t)|\geq 3$ by Lemma 7.16.

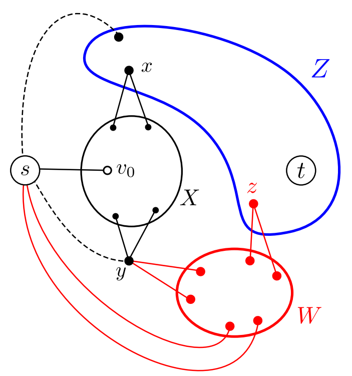

Figure 7.40: Case 2.2.2.3 (general).

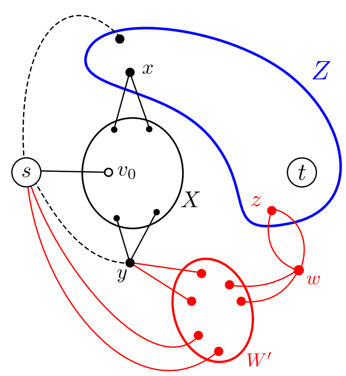

Figure 7.41: Case 2.2.2.3 (3-contractible).

Chapter 8

Conclusion

In this thesis, we have mainly studied paths in group-labeled graphs.

We have shown a reduction of the packing non-zero A-paths problem to the matroid matching problem in Chapter 4. Through our reduction, extending Lovász' work [59], we have given alternative proofs for the min-max duality due to Chudnovsky et al. [7] and the polynomial-time solvability. Furthermore, we have presented a possible extension of our reduction to the axiomatic model with the weak triple exchange property, which does not necessarily lead to a min-max duality or tractability.

In Chapter 5, we have clarified when the subgroup-forbidden model of packing A-paths admits a reasonable reduction to the linear matroid parity problem, which extends Schrijver's reduction [78, Section 73.1a] of Mader's $\mathcal{S}$-paths. Our reduction leads to fast algorithms for the subgroup-forbidden model via linear matroid parity, and the speeding-up technique of Cheung et al. [5, Section 5.1.3] is also applicable. We have also shown a variety of examples of underlying groups which satisfy our necessary and sufficient condition. As a result, it has turned out that a large class of the subgroup-forbidden model can be solved in $\mathrm{O}(|V|^{\omega})$ time.

We have discussed possible extensions of the above two reductions to a weighted situation in Chapter 6. It has shown that the problem of minimizing the total length of a designated number of vertex-disjoint non-zero A-paths can be reduced to the weighted matroid matching problem, and moreover to the weighted linear matroid parity problem when the same necessary and sufficient condition given in Chapter 5 is satisfied. Since Iwata [40] and Pap [75] announced polynomial-time algorithms for the weighted linear matroid parity problem, this reduction provides the first polynomial-time algorithm for such a weighted setting of packing non-zero A-paths, even for the same setting of Mader's $\mathcal{S}$-paths problem.

In Chapter 7, we have characterized group-labeled graphs with exactly two possible labels of s–t paths, which leads to the first nontrivial classification of them in terms of the number of possible labels of s–t paths. Based on our characterization, we have also proposed an efficient algorithm for finding an s–t path with arbitrary two labels forbidden.

It is easy to forbid just one label (or all labels in a fixed subgroup), and this is the first step to reveal the tractability of label-forbidding constraints.

There are several open problems related to our work as follows.

As shown in Section 3.5, packing non-zero A-paths can be extended to the axiomatic model, which seems on the border of tractability (cf. Theorems 3.13, 3.14, and 4.14). While the weak triple exchange property leads to several nice properties, it is still unknown whether it does a good characterization or a polynomial-time solvability or not.

Our reduction of the shortest disjoint non-zero A-paths problem to the weighted linear matroid parity problem in Chapter 6 leads to a polynomial-time algorithm for some restricted cases with the aid of the weighted linear matroid parity algorithms of Iwata and Pap. If there exists a polynomial-time algorithm for a larger class of the weighted matroid matching problem like Lovász' algorithm for matroid matching, then one can solve the problem with no restriction by an analogous argument in Chapter 4.

Regarding the problem of finding an s–t path with label-forbidding constraint, it is quite nontrivial to handle the situation in which we forbid three labels. The reduction of the k-disjoint paths problem shown in Section 1.4.2 can be reformulated to the situation in which we forbid $\mathrm{O}(2^k)$ labels. This maybe suggests that it is not so difficult as the k-disjoint paths problem to find an s–t path in a group-labeled graph with k labels forbidden.

Bibliography

[1] C. Berge: Sur le couplage maximum d'un graphe. *Comptes Rendus Hebdomadaires des Séances de l'Académie de Sciences*, **247** (1958), pp. 258–259.

[2] G. Birkhoff: Tres observaciones sobre el algebra lineal. *Revista Facultad de Ciencias Exactas, Puras y Aplicadas Universidad Nacional de Tucuman, Ser. A*, **5** (1946), pp. 147–151.

[3] A. Björklund, T. Husfeldt: Shortest two disjoint paths in polynomial time. *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming* (*ICALP 2014*), pp. 211–222, 2014.

[4] P. M. Camerini, G. Galbiati, F. Maffioli: Random pseudo-polynomial algorithms for exact matroid problems. *Journal of Algorithms*, **13** (1992), pp. 258–273.

[5] H. Y. Cheung, L. C. Lau, K. M. Leung: Algebraic algorithms for linear matroid parity problems. *ACM Transactions on Algorithms*, **10** (2014), No. 10.

[6] M. Chudnovsky, W. H. Cunningham, J. Geelen: An algorithm for packing non-zero A-paths in group-labelled graphs. *Combinatorica*, **28** (2008), pp. 145–161.

[7] M. Chudnovsky, J. Geelen, B. Gerards, L. Goddyn, M. Lohman, P. Seymour: Packing non-zero A-paths in group-labelled graphs. *Combinatorica*, **26** (2006), pp. 521–532.

[8] R. Churchley, B. Mohar, H. Wu: Weak duality for packing edge-disjoint odd (u, v)-trails. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms* (*SODA 2016*), pp. 2086–2094, 2016.

[9] E. Cohen, N. Megiddo: Recognizing properties of periodic graphs. *Applied Geometry and Discrete Mathematics*, **4** (1991), pp. 135–146.

[10] L. Collatz: Spektren periodischer Graphen. *Resultate der Mathematik*, **1** (1978), pp. 42–53.

[11] R. Diestel: *Graph Theory*, 4th ed., Springer-Verlag, Heidelberg, 2010.

[12] T. A. Dowling: A class of geometric lattices based on finite groups. *Journal of Combinatorial Theory, Ser. B*, **15** (1973), pp. 61–86.

[13] J. Edmonds: Paths, trees, and flowers. *Canadian Journal of Mathematics*, **17** (1965), pp. 449–467.

[14] J. Edmonds: Submodular functions, matroids, and certain polyhedra. *Proceedings of the Calgary International Conference on Combinatorial Structures and Their Applications*, pp. 69–87, 1970. (Also in *Combinatorial Optimization — Eureka, You Shrink!*, pp. 11–26, 2003).

[15] S. Even, A. Itai, A. Shamir: On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, **4** (1976), pp. 691–703.

[16] L. R. Ford, Jr., D. R. Fulkerson: Maximal flow through a network. *Canadian Journal of Mathematics*, **8** (1956), pp. 399–404.

[17] S. Fortune, J. Hopcroft, J. Wyllie: The directed subgraph homeomorphism problem. *Theoretical Computer Science*, **10** (1980), pp. 111–121.

[18] A. Frank: *Connections in Combinatorial Optimization*, Oxford University Press, 2011.

[19] F. G. Frobenius: Über Matrizen aus nicht negativen Elementen. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften zu Berlin*, pp. 456–477, 1912.

[20] F. G. Frobenius: Über zerlegbare Determinanten. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften zu Berlin*, pp. 274–277, 1917.

[21] N. Fu: A strongly polynomial time algorithm for the shortest path problem on coherent planar periodic graphs. *Proceedings of the 23rd International Symposium Algorithms and Computation* (*ISAAC 2012*), pp. 392–401, 2012.

[22] H. N. Gabow, M. Stallmann: An augmenting path algorithm for linear matroid parity. *Combinatorica*, **6** (1986), pp. 123–150.

[23] T. Gallai (T. Grünwald): Ein neuer Beweis eines Mengerischen Satzes. *Journal of the London Mathematical Society*, **13** (1938), pp. 188–192.

[24] T. Gallai: Maximum-minimum Sätze und verallgemeinerte Faktoren von Graphen. *Acta Mathematica Academiae Scientiarum Hungaricae*, **12** (1961), pp. 131–173.

[25] T. Gallai: Kritische Graphen II. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, **8** (1963), pp. 373–395.

[26] M. R. Garey, D. S. Johnson: *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, 1979.

[27] T. Gallai: Maximale Systeme unabhängiger Kanten. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, **9** (1964), pp. 401–413.

[28] A. V. Goldberg, A. V. Karzanov: Maximum skew-symmetric flows and matchings. *Mathematical Programming, Ser. A*, **100** (2004), pp. 537–568.

[29] P. Hall: On representatives of subsets. *The Journal of the London Mathematical Society*, **10** (1935), pp. 26–30.

[30] F. Harary: On the notion of balance of a signed graph. *Michigan Mathematics Journal*, **2** (1953), pp. 143–146.

[31] N. J. A. Harvey: Algebraic algorithms for matching and matroid problems. *SIAM Journal on Computing*, **39** (2009), pp. 679–702.

[32] A. Hatcher: *Algebraic Topology*, Cambridge University Press, 2002.

[33] H. Hirai, G. Pap: Tree metrics and edge-disjoint S-paths. *Mathematical Programming, Ser. A*, **147** (2014), pp. 81–123.

[34] F. Höfting, E. Wanke: Minimum cost paths in periodic graphs. *SIAM Journal on Computings*, **24** (1995), pp. 1051–1067.

[35] J. Hopcroft, R. Tarjan: Efficient algorithm for graph manipulation. *Communications of the ACM*, **16** (1973), pp. 372–378.

[36] J. Hopcroft, R. Tarjan: Efficient planarity testing. *Journal of the ACM*, **21** (1974), pp. 549–568.

[37] T. Huynh: *The Linkage Problem for Group-Labelled Graphs*, Ph.D. Thesis, Department of Combinatorics and Optimization, University of Waterloo, Ontario, 2009.

[38] K. Iwano: Some problems on doubly periodic infinite graphs. *Technical Reports*, CS-TR-078-87, Princeton University, 1987.

[39] K. Iwano, K. Steiglitz: Planarity testing of doubly periodic infinite graphs. *Networks*, **18** (1988), pp. 205–222.

[40] S. Iwata: A weighted linear matroid parity algorithm. *Proceedings of the 8th Japanese-Hungarian Symposium on Discrete Mathematics and Its Applications*, pp. 251–259, 2013.

[41] P. Jensen, B. Korte: Complexity of matroid property algorithms. *SIAM Journal on Computing*, **11** (1982), pp. 184–190.

[42] R. M. Karp: Reducibility among combinatorial problems. *Proceedings of a Symposium on the Complexity of Computer Computations*, pp. 85–103, 1972.

[43] R. M. Karp: On the computational complexity of combinatorial problems. *Networks*, **5** (1975), pp. 45–68.

[44] R. M. Karp, R. E. Miller, S. Winograd: The organization of computations for uniform recurrence equations. *Journal of the ACM*, **14** (1967), pp. 563–590.

[45] A. V. Karzanov: Edge-disjoint T-paths of minimum total cost. *Technical Report*, STAN-CS-92-1465, Department of Computer Science, Stanford University, 1993.

[46] A. V. Karzanov: Multiflows and disjoint paths of minimum total cost. *Mathematical Programming*, **78** (1997), pp. 219–242.

[47] K. Kawarabayashi, Y. Kobayashi, Y. Yamaguchi: *Personal Communication*, 2015.

[48] K. Kawarabayashi, P. Wollan: Non-zero disjoint cycles in highly connected group labelled graphs. *Journal of Combinatorial Theory, Ser. B*, **96** (2006), pp. 296–301.

[49] Y. Kawase, Y. Kobayashi, Y. Yamaguchi: Finding a path in group-labeled graphs with two labels forbidden. *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming* (*ICALP 2015*), pp. 797–809, 2015.

[50] Y. Kobayashi, S. Toyooka: Finding a shortest non-zero path in group-labeled graphs. *Mathematical Engineering Technical Reports*, METR 2015-19, University of Tokyo, 2015.

[51] D. Kőnig: Vonalrendszerek és determinánsok, *Mathematikai és Természettudományi Értesitő*, **33** (1915), pp. 221–229.

[52] D. Kőnig: Graphok és alkalmazásuk a determinánsok és a halmazok elméletére, *Mathematikai és Természettudományi Értesitő*, **34** (1916), pp. 104–119.

[53] D. Kőnig: Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre, *Mathematische Annalen*, **77** (1916), pp. 453–465.

[54] D. Kőnig: Graphok és matrixok, *Mathematikai és Fizikai Lapok*, **38** (1931), pp. 116–119.

[55] E. L. Lawler: Matroid intersection algorithms, *Mathematical Programming*, **9** (1975), pp. 31–56.

[56] E. L. Lawler: *Combinatorial Optimization: Networks and Matroids*, Holt, Rinehart and Winston, 1976.

[57] F. Le Gall: Powers of tensors and fast matrix multiplication. *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation* (*ISSAC 2014*), pp. 296–303, 2014.

[58] L. Lovász: On deteminants, matchings and random algorithms. *Proceedings of the International Symposia on Fundamentals of Computation Theory* (*FCT 1979*), pp. 565–574, 1979.

[59] L. Lovász: Matroid matching and some applications. *Journal of Combinatorial Theory, Ser. B*, **28** (1980), pp. 208–236.

[60] L. Lovász: Selecting independent lines from a family of lines in a space. *Acta Scientiarum Mathematicarum*, **42** (1980), pp. 121–131.

[61] L. Lovász: The matroid matching problem. *Algebraic Methods in Graph Theory,* II: *Colloquia Mathematica Societatis János Bolyai*, **25** (1981), pp. 495–517.

[62] L. Lovász, M. D. Plummer: *Matching Theory*, Akadémiai Kiadó, 1986.

[63] W. Mader: Über die Maximalzahl krezungsfreier H-Wege. *Archiv der Mathematik*, **31** (1978), pp. 387–402.

[64] K. Menger: Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, **10** (1927), pp. 96–115.

[65] S. Micali, V. V. Vazirani: An $\mathrm{O}(\sqrt{V}E)$ algorithm for finding maximum matching in general graphs. *Proceedings of the 21st Annual IEEE Symposium on Foundations of Computer Science* (*FOCS 1980*), pp. 17–27, 1980.

[66] M. Mucha, P. Sankowski: Maximum matchings via Gaussian elimination. *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science* (*FOCS 2004*), pp. 248–255, 2004.

[67] J. von Neumann: A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games Volume* II, pp. 5–12, 1953.

[68] J. B. Orlin: Some problems on dynamic/periodic graphs. *Progress in Combinatorial Optimization*, pp. 273–293, Academic Press, 1984.

[69] J. B. Orlin: A fast, simpler algorithm for the matroid parity problem. *Proceedings of the 13th Conference on Integer Programming and Combinatorial Optimization* (*IPCO 2008*), pp. 240–258, 2008.

[70] J. Oxley: *Matroid Theory*, 2nd ed., Oxford University Press, 2011.

[71] G. Pap: *A constructive approach to matching and its generalizations*, Ph.D. Thesis, Institute of Mathematics, Eötvös Loránd University, Budapest, 2006.

[72] G. Pap: Packing non-returning A-paths. *Combinatorica*, **27** (2007), pp. 247–251.

[73] G. Pap: Packing non-returning A-paths algorithmically. *Discrete Mathematics*, **308** (2007), pp. 1472–1488.

[74] G. Pap: A polynomial time algorithm for weighted node-disjoint S-paths. *Proceedings of the 7th Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications*, pp. 322–331, 2011.

[75] G. Pap: Weighted linear matroid matching. *Proceedings of the 8th Japanese-Hungarian Symposium on Discrete Mathematics and Its Applications*, pp. 411–413, 2013.

[76] N. Robertson, P. D. Seymour: Graph minors. XIII. the disjoint paths problem. *Journal of Combinatorial Theory, Ser. B*, **63** (1995), pp. 65–110.

[77] A. Schrijver: A short proof of Mader's $\mathcal{S}$-paths theorem. *Journal of Combinatorial Theory, Ser. B*, **85** (2001), pp. 319–321.

[78] A. Schrijver: *Combinatorial Optimization — Polyhedral and Efficiency*, Springer-Verlag, 2003.

[79] A. Sebő, L. Szegő: The path-packing structure of graphs. *Proceedings of the 10th International Conference on Integer Programming and Combinatorial Optimization* (*IPCO 2004*), pp. 256–270, 2004.

[80] P. D. Seymour: Disjoint paths in graphs. *Discrete Mathematics*, **29** (1980), pp. 293–309.

[81] Y. Shiloach: A polynomial solution to the undirected two paths problem. *Journal of the ACM*, **27** (1980), pp. 445–456.

[82] S. Tanigawa: Matroids of gain graphs in applied discrete geometry. *Transactions of the American Mathematical Society*, **367** (2015), pp. 8597–8641.

[83] S. Tanigawa, Y. Yamaguchi: Packing non-zero A-paths via matroid matching. *Mathematical Engineering Technical Reports*, METR 2013-08, University of Tokyo, 2013.

[84] C. Thomassen: 2-linked graphs. *European Journal of Combinatorics*, **1** (1980), pp. 371–378.

[85] W. T. Tutte: The factorization of linear graphs. *The Journal of the London Mathematical Society*, **22** (1947), pp. 107–111.

[86] W. T. Tutte: Menger's theorem for matroids. *Journal of Research National Bureau of Standards Section B*, **71** (1965), pp. 49–53.

[87] V. V. Vazirani: A theory of alternating paths and blossoms for proving correctness of the $O(\sqrt{V}E)$ general graph maximum matching algorithm. *Combinatorica*, **14** (1994), pp. 71–109.

[88] V. V. Vazirani: A proof of the MV matching algorithm. A preprint, 2014. http://www.cc.gatech.edu/~vazirani/new-proof.pdf

[89] E. Wanke: Paths and cycles in finite periodic graphs. *Proceedings of the 18th Symposium on Mathematical Foundations of Computer Science* (*MFCS 1993*), pp. 751–760, 1993.

[90] P. Wollan: Packing cycles with modularity constraint. *Combinatorica*, **31** (2011), pp. 95–126.

[91] M. A. Woodbury: Inverting modified matrices. *Memorandum Report*, **42** (1950), p. 106.

[92] Y. Yamaguchi: *Structures and Algorithms for Path-Packing Problems*, Master's Thesis, Department of Mathematics and Mathematical Sciences, Kyoto University, Kyoto, 2013.

[93] Y. Yamaguchi: Shortest disjoint non-zero A-paths via weighted matroid matching. *Methematical Engineering Technical Reports*, METR 2015-20, University of Tokyo, 2015.

[94] Y. Yamaguchi: Packing A-paths in group-labelled graphs via linear matroid parity. *SIAM Journal on Discrete Mathematics*, to appear.

[95] Y. Yamaguchi, A. Ogawa, A. Takeda, S. Iwata: Cyber security analysis of power networks by hypergraph cut algorithms. *IEEE Transactions on Smart Grid*, **6** (2015), pp. 2189–2199.

[96] T. Zaslavsky: Biased graphs. I. Bias, balance and gains. *Journal of Combinatorial Theory, Ser. B*, **47** (1989), pp. 32–52.

[97] T. Zaslavsky: Biased graphs. II. The three matroids. *Journal of Combinatorial Theory, Ser. B*, **51** (1991), pp. 46–72.

[98] T. Zaslavsky: A mathematical bibliography of signed and gain graphs and allied areas. *Electronic Journal of Combinatorics*, **5** (1998), Dynamic Surveys 6th Ed., 124 pp. (The newest version, 9th Ed., 426 pp., is available on the author's website: http://www.math.binghamton.edu/zaslav/Bsg/)