

論文 / 著書情報
Article / Book Information

題目(和文)	
Title(English)	Space Efficient Algorithms for Planar Separator Theorem and Grid Graph Reachability Problem
著者(和文)	芦田亮
Author(English)	Ryo Ashida
出典(和文)	学位:博士(理学), 学位授与機関:東京工業大学, 報告番号:甲第11250号, 授与年月日:2019年9月20日, 学位の種別:課程博士, 審査員:渡辺 治,南出 靖彦,田中 圭介,伊東 利哉,鹿島 亮,脇田 建
Citation(English)	Degree:Doctor (Science), Conferring organization: Tokyo Institute of Technology, Report number:甲第11250号, Conferred date:2019/9/20, Degree Type:Course doctor, Examiner:,,,,,
学位種別(和文)	博士論文
Type(English)	Doctoral Thesis

Space Efficient Algorithms for Planar Separator Theorem and Grid Graph Reachability Problem

Ryo Ashida

A dissertation submitted in partial fulfillment
of the requirement for the degree of

DOCTOR OF SCIENCE

Department of Mathematical and Computing Science
Tokyo Institute of Technology

July 2019

Abstract

The graph reachability or the st-connectivity problem asks, for a given graph G and two of its vertices s and t , whether there is a path in G from s to t or not. This problem is one of the fundamental computational problems in both algorithm design and in complexity theory. In particular, this problem is known to be complete (under logspace reductions) for the complexity class NL—nondeterministic logspace. Thus an $O(\log n)$ -space algorithm for graph reachability implies that the complexity class NL equals L, and resolves one of the fundamental open problems in computational complexity theory. Interestingly, for its undirected version, that is, for the undirected graph reachability problem, Reingold gave a remarkable $O(\log n)$ -space algorithm. Note that, by using standard algorithmic techniques such as BFS or DFS we can design an algorithm that runs in almost linear-space and in almost linear-time for solving reachability over directed graphs. Savitch gave an algorithm that solves reachability using $O((\log n)^2)$ -space. However Savitch’s algorithm takes super polynomial-time. While the BFS/DFS algorithm is time-efficient, Savitch’s algorithm is space-efficient. Thus it is natural to ask whether there exists an algorithm, for the directed graph reachability problem, that is efficient in both time and space. In a survey article on the graph reachability and related problems, Wigderson asked whether there is a sublinear-space and polynomial-time algorithm for the graph reachability problem. For this question, the best known answer is the algorithm given by Barnes et al. that runs in $O(n/2^{\sqrt{\log n}})$ -space and polynomial-time. Unfortunately, though the bound $O(n/2^{\sqrt{\log n}})$ is “sublinear”, it is quite close to being linear. If we view $O(n^{1-\epsilon})$ (for some $\epsilon > 0$), as “sublinear”, it has been open whether such a sublinear-space and polynomial-time algorithm exists for reachability problem.

Recently, there have been some advancements on this sublinear-space and polynomial-time computability for restricted graph classes. The first break through was given by Asano and Doerr, who showed an algorithm that solves the reachability problem on directed grid graphs in $O(n^{1/2+\epsilon})$ -space and polynomial-time. Inspired by this work, Imai, Nakagawa, Pavan, Vinodchandran, and Watanabe gave an $O(n^{1/2+\epsilon})$ -space and polynomial-time algorithm for the reachability problem on directed planar graphs. Later Asano et al. gave an improved $\tilde{O}(\sqrt{n})$ -space algorithm. In this doctoral dissertation, we propose an $\tilde{O}(n^{1/3})$ and polynomial-time algorithm for the grid graph reachability problem. Note

that, in all of the above algorithms, the input graphs are planar and they all critically rely on the existence of $O(\sqrt{n})$ -size separator.

The Separator Theorem states that any planar graph G with n vertices has a separator of size $O(\sqrt{n})$, that is, a set S of $O(\sqrt{n})$ vertices of G such that by removing S , G is split into disconnected subgraphs of almost equal size, say, each having more than $n/3$ vertices. In fact, in their seminal work that first proved the Separator Theorem, Lipton and Tarjan gave an efficient separator algorithm, an algorithm for computing an $O(\sqrt{n})$ -size separator for planar graphs. The notion of planar separator and the algorithm of Lipton and Tarjan have found several applications in designing efficient algorithms for planar graphs. Since the work of Lipton and Tarjan, several versions of separator algorithms have been proposed, and they have been applied to design various algorithms for planar graphs.

The above mentioned reachability algorithms use a separator algorithm proposed by Imai et al. that computes an $O(\sqrt{n})$ -size separator in polynomial-time and $\tilde{O}(\sqrt{n})$ -space for any given planar graph with n vertices. In their paper, it is claimed that such a separator algorithm can be designed by modifying the algorithm of Gazit and Miller that uses Miller’s separator algorithm as a basic subroutine. While a key modification idea is given in their paper, one needs to supply nontrivial details to design a desired algorithm claimed in their paper. In this doctoral dissertation, we completely reconsider the design of a separator algorithm that runs in $\tilde{O}(\sqrt{n})$ -space and polynomial-time on planar graphs. While we borrow many ideas from Gazit and Miller and Imai et al, our algorithm in this dissertation differs from both the algorithms in many technical details. Through this work we also clarify that the separator construction problem is log-space reducible to the problem of constructing a BFS tree for a given undirected graph. That is, we show that a key to get yet more space efficient separator algorithm is to improve the space complexity of the BFS tree construction problem.

Acknowledgement

I would like to express the deepest appreciation to professor Osamu Watanabe. Professor Watanabe has led me as my supervisor since I joined his laboratory in 2013. He gave me chances to meet and discuss with many excellent researchers, and these experiences greatly helped my research. Professor Watanabe took office as executive vice president for research of Tokyo Institute of Technology in 2018. Although professor Watanabe was very busy, he spent time on discussion for our research. I cannot thank professor Watanabe enough.

I am deeply grateful to professor Toshiya Itoh. I joined Itoh laboratory and the seminar of the lab after Watanabe laboratory closed in 2018. Thanks to professor Itoh, I could meet graduation requirements.

Special thanks also to the other examiners professor Yasuhiko Minamide, professor Keisuke Tanaka, associate professor Ryo Kashima, and associate professor Ken Wakita.

I would like to thank Shoshisha Foundation for a full scholarship that made it possible to work on research for three years.

I would like to thank HEROZ, Inc. where I am working. The company supported me in finishing my doctorate.

I would like to express my gratitude to professor Akinori Kawachi currently at Mie University, and assistant professor Ryuhei Mori. They were assistant professors of Watanabe laboratory, and they gave me a lot of helpful comments in the seminar. I would also like to show my appreciation to Yuko Hirauchi, Tamami Watanabe and Kazuyo Kawaguchi who are/were the secretaries of our floor and have cheered me up a lot.

I would like to thank my family and my colleagues for their support and encouragements. I thank doctor Kotaro Nakagawa, doctor Tatsuya Imai, doctor Yoshiki Nakamura, Kazumasa Shinagawa, Mikito Nanashima and all other students in Watanabe laboratory.

Contents

Abstract	i
Acknowledgement	iii
1 Introduction	1
1.1 Planar separator theorem	1
1.2 Graph reachability problem	3
1.3 Contribution	5
1.4 Outline of this dissertation	6
2 Preliminaries	7
2.1 Basic notations and notion	7
2.2 Computational model	10
3 Modification of Miller’s Separator Algorithm	13
3.1 Outline of the algorithm	13
3.2 Face level tree and frontier cycle tree	15
3.3 Pruning a frontier cycle tree	18
3.4 Constructing a spanning tree	21
3.5 Finding a separator	23
3.6 Log-space reduction to BFS tree construction	26
4 A Sublinear-space and Polynomial-time Planar Separator Algorithm	29
4.1 Outline of the algorithm	29
4.2 Base Graphs G and $\tilde{G}$	42
4.3 Voronoi Region	46
4.4 Multiple-Dual-Cycle (m.d.-cycle)	51
4.5 Dividing Voronoi Regions	55
4.6 Frame-Graph	64
4.6.1 Preprocessing pre-frame-cycles	64
4.6.2 Properties of the frame-graph.	68
4.7 Floor and Ceiling Modification	72

5	$\tilde{O}(n^{1/3})$-Space Algorithm for the Grid Graph Reachability Problem	87
5.1	Outline of the algorithm	87
5.2	Graph transformation	88
5.2.1	Circle graph	88
5.2.2	Gadget graph	91
5.3	Apply PlanarReach to a gadget graph	117
6	Conclusion	121

Chapter 1

Introduction

1.1 Planar separator theorem

A *planar graph* is a graph that can be embedded in a plane without any edge intersections. There are several significant and useful characterizations and properties of planar graphs. (e.g., Euler's formula, Kuratowski's and Wagner's theorems, etc. [41, 63, 64, 46, 24, 55, 61]) Above all, graph theoretically, the *Separator Theorem* is important for illustrating one of the key characteristics of planar graphs. This theorem claims that any planar graph G with n vertices has an $O(\sqrt{n})$ -size *separator*, that is, a set S of $O(\sqrt{n})$ vertices of G such that by removing S , G is split into disconnected subgraphs of almost equal size, say, each having more than $n/3$ vertices. For example, consider a $\sqrt{n} \times \sqrt{n}$ grid graph G (see Figure 1.1). The central column S is a $\sqrt{n}$ -size separator. In fact, G is divided into left and right hand sides of the same size by removing S . A set of $k < \sqrt{n}$ vertices in G can enclose at most $k(k-1)/2$ vertices in G . Thus, we need at least $\sqrt{2n/3}$ vertices for dividing G such that disconnected subgraphs have




Figure 1.1: A separator for a grid graph. S divides the grid graph into two subgraphs A and B of the same size.

more than $n/3$ vertices. Therefore, the $O(\sqrt{n})$ -size in the theorem is optimal.

In their seminal work that first proved the separator theorem, Lipton and Tarjan [45] gave an efficient separator algorithm, an algorithm for computing an $O(\sqrt{n})$ -size separator for planar graphs. For an input planar graph G , they first perform a breadth first search, starting from an arbitrary vertex v , and classify the vertices of G by distances from v . A level- ℓ layer is a set of vertices whose distances from v are equal to ℓ . If the level- ℓ layer is removed, G is split into level- $< \ell$ layers and level- $> \ell$ layers. They showed that we can compute two layers such that a union of them becomes an $O(\sqrt{n})$ -size separator in linear time. Precisely speaking, the size of the separator computed in this way is at most $\sqrt{8n}$. Djidjev [28] improved the constant factor from $\sqrt{8}$ to $\sqrt{6} \approx 2.45$; he also proved a lower bound of $1.56\sqrt{n}$, which is still the best known. Another approach was proposed by Miller [47]. For an input planar graph G , Miller's strategy first computes a spanning tree T of G such that the diameter of T is $O(\sqrt{n})$. If we fix an edge e which does not appear in T , a path on T connecting end points of e is determined. This path and the non-tree edge e form a simple cycle. The simple cycle divides G into its inside and outside, and the size of this cycle is $O(\sqrt{n})$ since the diameter of T is $O(\sqrt{n})$. Miller showed that we can construct such a cycle that each of its inside and outside has more than $n/3$ vertices, thus the cycle is in fact an $O(\sqrt{n})$ -size separator. Moreover, many other separator algorithms have been proposed; e.g., [5, 27, 48, 56, 4, 57, 35, 50, 30].

The separator theorem has been used heavily in computer science for designing efficient algorithms for planar graphs. For solving many kinds of problems, *divide and conquer* is a strong and useful method. A divide and conquer algorithm breaks down a problem into two or more subproblems. The subproblems are solved by using the method recursively. Then the solutions of these subproblems are combined, and a solution of the original problem is obtained. For designing efficient divide and conquer algorithms, decomposed subproblems should be significantly smaller than the original. One basic way to guarantee this is dividing a problem into subproblems of roughly the same size. The separator theorem divides a planar graph into two disconnected subgraphs of almost equal size. Thus, the theorem can be used for designing efficient divide and conquer algorithms for problems on planar graphs. In fact, fast algorithms based on divide and conquer using separators have been proposed for many problems on planar graphs; e.g., shortest path problem [31, 37, 59, 40], shortest cycle problem [42, 20], construction of nearest neighbor graphs [32], polygon triangulation [36] and point location problem [44]. In addition to these linear or polynomial-time solvable problems, this strategy can be applied to NP-hard problems, and various NP-hard problems on planar graphs are solvable in time $2^{O(\sqrt{n})}$ or $2^{O(\sqrt{n} \log n)}$; e.g., maximum independent set [44], Steiner tree problem [16], Hamiltonian cycle [25] and traveling salesperson problem [44]. On general graphs, $2^{o(n)}$ -time algorithms for these problems are not known. Besides exact solutions of NP-hard problems, approximation algorithms of NP-hard optimization problems can be obtained by controlling the recursion depth of divide and conquer. A polynomial-time approximation schemes (PTAS) is an algorithm that takes an instance of an optimization problem and a parameter $\varepsilon > 0$ as

1. INTRODUCTION

input, and computes a solution whose value is at most $(1 + \varepsilon) \cdot OPT$ (for a minimization problem) or at least $(1 - \varepsilon) \cdot OPT$ (for a maximization problem) in polynomial-time, where OPT is the optimal value of the problem. For instance, for the traveling salesperson problem, a PTAS could output a tour of length less than $(1 + \varepsilon) \cdot L$, where L is the length of the shortest tour. Several PTAS's for NP-hard optimization problems on planar graphs have been proposed; e.g., maximum independent set [44, 13], minimum vertex cover [14, 21] and traveling salesperson problem [7].

The separator theorem is useful for constructing not only time efficient algorithms but also space efficient algorithms. For the graph reachability problem, some space efficient algorithms have been proposed. (The details about the graph reachability problem will be described in the next section.) Asano and Dorr [8] showed an algorithm that solves the reachability problem on grid graphs in $O(n^{1/2+\varepsilon})$ -space. Inspired by this work, Imai, Nakagawa, Pavan, Vinodchandran and Watanabe [38] gave an $O(n^{1/2+\varepsilon})$ -space algorithm for the reachability problem on planar graphs in general. Later Asano et al. [9] gave an improved $\tilde{O}(\sqrt{n})$ -space¹ algorithm. In all these algorithms, using $O(\sqrt{n})$ -size separator is crucial. Furthermore, these algorithms except for the first one use a separator algorithm proposed in [38] that produces an $O(\sqrt{n})$ -size separator in $\tilde{O}(\sqrt{n})$ -space. This separator algorithm is based on an algorithm of Gazit and Miller [33] that is designed for computing a separator efficiently on a parallel computation model. Since space-bounded computation and parallel computation share common features, one can expect that the algorithm of Gazit and Miller can be modified naturally for, say, $\tilde{O}(\sqrt{n})$ -space bounded computation. Unfortunately, it was not so simple and there have been several nontrivial details left unexplained in [38]. In this doctoral dissertation we completely reconsider the design of a separator algorithm that runs in $\tilde{O}(\sqrt{n})$ -space and polynomial-time on planar graphs [10].

1.2 Graph reachability problem

The *graph reachability problem*, for a graph $G = (V, E)$ and two distinct vertices $s, t \in V$, is to determine whether there exists a path from s to t . This problem characterizes many important complexity classes. The *directed* graph reachability problem is a canonical complete problem for the non-deterministic log-space class, NL. SL is the class of problems solvable by *symmetric*, non-deterministic log-space computation. (Symmetric means that, in non-deterministic computation by a Turing machine, a configuration A yields a configuration B if and only if B yields A .) The inclusion relation $L \subseteq SL \subseteq NL$ holds, where L is the deterministic log-space computable class. The *undirected* graph reachability problem is SL-complete [43]. There have been progresses of *deterministic* algorithms for the undirected graph reachability problem; an $O(\log^{3/2} n)$ -space algo-

¹In this dissertation “ $\tilde{O}(s(n))$ -space” means $O(s(n))$ -words intuitively and precisely $O(s(n) \log n)$ -space.

rithm by Nisan, Szemerédi and Wigderson [49], an $O(\log^{4/3} n)$ -space algorithm by Armoni, Ta-Shma, Wigderson and Zhou [6] and an $O(\log n \log \log n)$ -space algorithm by Trifonov [60]. Reingold [53] gave a remarkable $O(\log n)$ -space algorithm solving this problem. This result implies that $L = SL$ and the undirected graph reachability problem characterizes the class L . As with P vs. NP problem, whether $L = NL$ or not is a major open problem in computational complexity theory. This problem is equivalent to whether the directed graph reachability problem is solvable in deterministic log-space. There exist two fundamental solutions for the directed graph reachability problem, breadth first search, denoted as BFS, and Savitch's algorithm [54]. BFS runs in $O(n)$ -space and $O(m)$ -time, where n and m are the number of vertices and edges, respectively. For Savitch's algorithm, we use only $O(\log^2 n)$ -space but require $\Theta(n^{\log n})$ -time. BFS needs short time but large space. Savitch's algorithm uses small space but super polynomial-time. A natural question is whether we can make an efficient deterministic algorithm in both space and time for the directed graph reachability problem. In particular, Wigderson [65] proposed a problem that does there exist an algorithm for the directed graph reachability problem that uses polynomial-time and $O(n^\varepsilon)$ -space, for some $\varepsilon < 1$?, and this question is still open. The best known polynomial-time algorithm, shown by Barnes, Buss, Ruzzo and Schieber, uses $O(n/2^{\sqrt{\log n}})$ -space [15]. Note that, though the bound $O(n/2^{\sqrt{\log n}})$ is "sub-linear", it is quite close to linear and improving this bound remains a significant open question. In fact, improving this bound is considered to be difficult because there exists a tight lower bound for this problem on the node-named JAG (NNJAG) model [29]. The *jumping automation for graphs*, or JAG, introduced by Cook and Rackoff [23] is a computational model which moves a set of pebbles on the graph. The operations of JAG are moving a pebble along an edge in the graph and making a pebble jump a node occupied by another pebble. NNJAG is an extension of JAG, where the computation is allowed to depend on the name of the node on which a pebble is located. Though NNJAG is a restricted model compared to Turing machine, many known algorithms for the directed reachability can be implemented in NNJAG without significant blow up in time and space [23, 51].

For some restricted graph classes, better results are known. Stolee and Vinodchandran [58] showed that for any $0 < \varepsilon < 1$, the reachability problem for directed acyclic graph with $O(n^\varepsilon)$ sources and embedded on a surface with $O(n^\varepsilon)$ genus can be solved in polynomial-time and $O(n^\varepsilon)$ -space. Chakraborty and Tewari [18] showed that reachability in directed layered planar graphs can be decided in polynomial-time and $O(n^\varepsilon)$ -space for any $\varepsilon > 0$. Kannan, Khanna and Roy [39] gave an $O(n^\varepsilon)$ -space and polynomial-time algorithm for solving the reachability problem in unique path graphs. Planar graphs are a natural topological restriction of general graphs, and grid graphs are a subclass of planar graphs, where a grid graph is a graph whose vertices are located on grid points, and whose vertices are adjacent only to their immediate horizontal or vertical neighbors. The *planar* and *grid* graph reachability problems have been well-studied. They are in the *unambiguous* log-space class, UL [17], which is a subclass of NL . (Unambiguous means that problems are decided by a non-

deterministic machine with at most one accepting computation.) Furthermore, They are equivalent under log-space reduction [3], and log-space reducible to their complements [2], and hard for NC^1 under AC^0 reduction [2]. (NC^1 consists of all families of circuits of depth $O(\log n)$ and polynomial size, with constant bounded fan-in boolean gates, and AC^0 consists of all families of circuits of depth $O(1)$ and polynomial size, with unbounded fan-in boolean gates.) Thus, in fact, they are in $\text{UL} \cap \text{co-UL}$, but are not known to be hard for L under AC^0 reduction. Asano and Doerr [8] showed an algorithm that solves the reachability problem on grid graphs in $O(n^{1/2+\epsilon})$ -space. Inspired by this work, Imai et al. [38] gave an $O(n^{1/2+\epsilon})$ -space algorithm for the reachability problem on planar graphs. In both algorithms, due to the recursive structure of the algorithms, the exponent of polynomial of time complexity grows in proportion to $1/\epsilon$. Asano et al. [9] devised an efficient way to control the recursion, and proposed a polynomial-time and $\tilde{O}(\sqrt{n})$ -space algorithm for the planar graph reachability problem. In this doctoral dissertation we propose an $\tilde{O}(n^{1/3})$ -space and polynomial-time algorithm for the grid graph reachability problem [12].

1.3 Contribution

Space efficient computation has been one of the important research targets in theory of computing from both theoretical and practical view points. For example, constructing a polynomial-time and sublinear-space algorithm for the directed graph reachability problem has been recently investigated actively [8, 38, 9], and several interesting sublinear-space algorithms have been found for the problem on grid and planar graphs. These algorithms make use of a small size separator of a given planar graph, and it is important to compute such a separator in sublinear-space. More precisely, these algorithms use a separator algorithm proposed in [38] that computes an $O(\sqrt{n})$ -size separator in polynomial-time and $\tilde{O}(\sqrt{n})$ -space for any given planar graph with n vertices. In [38], it is claimed that such a separator algorithm can be designed by modifying the algorithm of Gazit and Miller [33] that uses Miller's algorithm [47] as a basic subroutine. While a key modification idea is given in [38], one needs to supply nontrivial details to design a desired algorithm claimed in [38]. In this doctoral dissertation, we give a complete separator algorithm that runs in $\tilde{O}(\sqrt{n})$ -space and polynomial-time on planar graphs [10]. For the sake of this, we give nontrivial modifications of Miller's algorithm [11] and the algorithm of Gazit and Miller [10].

Miller's algorithm is a linear-time algorithm for computing an $O(\sqrt{n})$ -size separator for 2-connected planar graphs which in fact can be implemented as an NC algorithm (i.e., a polylogarithmic parallel time algorithm with a polynomial number of processors). We modify it to a space efficient version for using as a subroutine of the main algorithm. More specifically, we show a polynomial-time algorithm that computes, for any weighted plane graph G with n vertices and f faces such that it is 2-connected and all its faces have weight $\leq 2/3$ and size $\leq d$, a weighted cycle separator of size $O(\sqrt{dn})$ in $O(f \log n)$ -space. Through

this work we also clarify that the separator construction problem is log-space reducible to the problem of constructing a breadth first search tree (BFS tree) for a given undirected graph. That is, we show that a key to get yet more space efficient separator algorithm is to improve the space complexity of the BFS tree construction problem.

As mentioned in the previous section, Asano et al. [9] gave a polynomial-time and $\tilde{O}(\sqrt{n})$ -space algorithm for the planar graph reachability problem. Since grid graphs are special cases of planar graphs, we can solve the grid graph reachability problem in $\tilde{O}(\sqrt{n})$ -space and polynomial-time. In this doctoral dissertation, we give an $\tilde{O}(n^{1/3})$ -space and polynomial-time algorithm for the grid graph reachability problem [12]. Since this algorithm uses the algorithm of Asano et al. as a subroutine, which crucially uses an $O(\sqrt{n})$ -size separator, this algorithm also highly depends on an $O(\sqrt{n})$ -size separator. This result is interesting in this sense because its space bound is much less than the separator size.

1.4 Outline of this dissertation

In the next chapter, some formal definitions are shown to explain our algorithms. Some basic notations and notion for discussing planar graphs are given in the first section, and we discuss our computational model in the second section.

In the third chapter, we modify Miller's separator algorithm to a space efficient version. The outline of our algorithm is described in the first section. The algorithm consists of four steps, and section 2 to 5 correspond to those steps. In the sixth section, we show that the separator construction problem is log-space reducible to the problem of constructing a BFS tree.

In the fourth chapter, a complete construction of an $O(\sqrt{n})$ -size separator with $\tilde{O}(\sqrt{n})$ -space and polynomial-time is proposed. The outline of the algorithm is described in the first section. The algorithm consists of three steps, and Step 2 consists of four sub-steps. Step 3 is contained in the outline section, and the other step and sub-steps correspond to the following sections.

In the fifth chapter, the $\tilde{O}(n^{1/3})$ -space and polynomial-time algorithm for the grid graph reachability problem is proposed. The algorithm transforms an input grid graph to a specific planar graph, and applies the planar graph reachability algorithm of Asano et al. [9] to the modified planar graph. The graph transformation algorithm is shown in the second section. In the third section, we explain how to apply the reachability algorithm of [9].

Finally the conclusion is discussed in the last chapter.

Chapter 2

Preliminaries

2.1 Basic notations and notion

Throughout this dissertation, for any set X , $|X|$ denotes the number of elements in X . We refer to the maximum and minimum elements of X as $\max X$ and $\min X$, respectively. For any graph G , we formally use $V(G)$ and $E(G)$ to denote the set of vertices of G and the set of edges of G respectively. On the other hand, we will mainly discuss with some fixed graph G , and in this case, we simply use V and E to denote $V(G)$ and $E(G)$. For any vertex $v \in V$, let $N_G(v)$ denote the set of its *neighbors*, namely, vertices adjacent to v in G . For any $U \subseteq V$ (resp., $D \subseteq E$) we use $G[U]$ (resp., $G[D]$) to denote a subgraph of G *induced* by U (resp., by D).

We assume that an input graph is *simple*, that is, each pair of vertices has at most one edge. By a *path* of $G = (V, E)$, we mean a “simple” path, that is, a sequence adjacent edges of E where no vertex is visited more than once. Similarly, a *cycle* of $G = (V, E)$ is a sequence adjacent edges of E that comes back to the first vertex where no vertex except for the first one is visited more than once. We simply represent a path and a cycle by a sequence $(v_1, \dots, v_m)$ of vertices of V such that every $\{v_i, v_{i+1}\}$ (and also $\{v_m, v_1\}$ for a cycle) is an edge of E . Although the same sequence could be used for representing both a path and a cycle, the difference should be clear from the context. We may also use this sequence representation for indicating the “direction” of a path or a cycle.

While our argument should be mathematically precise, we would like to avoid rigorous but tedious topological treatments of plane graphs. Thus, following the previous work in the literature (see, e.g., [47]) we prepare a framework for discussing plane graphs in a combinatorial way.

Definition 2.1 (Planar graph, plane graph, and combinatorial embedding). A graph G is *planar* if it has a *planar embedding*, a way to arrange the vertices of $N_G(v)$ on the plane (i.e., $\mathbb{R}^2$) for all $v \in V$ so that no pair of edges intersect each other except for their endpoints. A planar embedding is specified by a

combinatorial embedding, that is, a set $\{\pi_v : v \in V\}$ of lists, where each π_v is an enumeration of vertices of $N_G(v)$ in the clockwise order around v under the embedding.

Remark. Throughout this dissertation, we will use a *plane graph* to mean a planar graph that is embedded in the plane following one of such combinatorial embeddings.

Consider any plane graph G embedded in the plane under a combinatorial embedding $\{\pi_v : v \in V\}$. One of the keys for discussing plane graphs is a way to define the faces of G . Intuitively, the graph G (embedded in the plane) separates the plane into “subplanes,” each of which is a “face” of G . Here we formally define the notion of “face” as follows.


First define the notion of “left-traverse” of G . Consider any edge $\{v_1, v_2\}$ of G , and fix its direction as, e.g., (v_1, v_2) , which we regard as the first directed edge e_1 (of the left-traverse). Consider the enumeration π_{v_2} of adjacent vertices of v_2 , and let v_3 be the next vertex of v_1 in the enumeration, that is, v_3 is clockwise the first vertex adjacent to v_2 after v_1 . Then define $e_2 = (v_2, v_3)$ as the second directed edge. Continue this process of identifying directed edges (based on edges of G) until we come back to the first directed edge e_1 . We call this process *left-traverse process*, and a sequence of vertices visited during the process *left-traverse* of G started from (v_1, v_2) . Precisely speaking, we do not include the last vertex, i.e., the starting vertex v_1 in the sequence; see Figure 2.1 for examples. Although a left-traverse is a sequence of vertices, this can be regraded as the sequence of the directed edges in the order that they are identified by the left-traverse process. We remark here the following fact.

Lemma 2.1. Consider any graph G . For any of its left-traverse $(v_1, v_2, \dots)$, no directed edge (v_i, v_{i+1}) appears more than once. Thus, no infinite loop occurs in any left-traverse process.

*Proof.*¹ Assume to the contrary that some directed edge (v_i, v_{i+1}) appears more than once. Let $e_{i-1} = (v_{i-1}, v_i)$ and $e_i = (v_i, v_{i+1})$ be the directed edges identified in the left-traverse process when v_i appears for the first time. From our assumption, we have another (v_j, v_i) for some $j > i + 1$ right after which e_i appears for the second time in the process. Then since v_{i+1} is clockwise the next vertex of v_{i-1} among all vertices adjacent to v_i , it follows that v_j must be located between v_{i+1} and v_{i-1} in the clockwise enumeration π_{v_i} of $N_G(v_i)$; that is, v_{i+1} cannot be clockwise the next of v_j in π_{v_i} because at least v_{i-1} comes before v_{i+1} after v_j ; see Figure 2.1 (c). Thus, (v_i, v_{i+1}) cannot be next to (v_j, v_i) in the process; a contradiction. $\square$

Consider any plane graph G with at least two faces. Then any left-traverse $t = (v_1, \dots, v_m)$ of G separates the plane where G is embedded into at least two “subplanes.” From the choice of directed edges, it is clear that there is no vertex

¹The lemma may be clear since the same vertex cannot be the “next” vertex of two vertices in an enumeration of a combinatorial embedding. But in order to give a clear image on the left-traverse notion, we give a bit careful proof here.



(a) An example of the first two edges identified by a left-traverse process. As the next directed edge of $e_1 = (v_1, v_2)$, an edge $e_2 = (v_2, v_3)$ is selected, where v_3 is the clockwise next vertex of v_1 among all adjacent vertices of v_2 . By the choice of these edges, no vertex adjacent to v_2 exists in the left of $e_1 \rightarrow e_2$, which guarantees that the left of each left-traverse is a “face.” (b) A connected but not 2-connected graph. A left-traverse from edge $(3, 4)$ is $(3, 4, 5, 6, 3, 2, 1, 2)$, which defines the outer face. On the other hand, the face inside the square is defined by, e.g., a left-traverse from edge $(4, 3)$, that is, $(4, 3, 6, 5)$. (c) An example for the proof of Lemma 2.1.

Figure 2.1: Left-traverse

of G in the “subplane” left of the directed edges of t ; that is, a “face” is identified by this left-traverse. In fact, it has been known that for any plane graph, all its faces are identified by some left-traverse as stated in Proposition 2.1 below. In this dissertation, we regard this proposition as our definition of the notion of face; throughout this dissertation, we will identify each face by its corresponding left-traverse, which is called the *boundary* of the face.

Proposition 2.1. Consider any plane graph G . Then for any left-traverse of G , there is a unique subplane, i.e., face, located left of the traverse w.r.t. its direction. On the other hand, for any face of G , there is a left-traverse of G that defines the face as a subplane located left of the left-traverse w.r.t. its direction.

Remark. Since each left-traverse is determined by a starting directed edge, there are more than one essentially the same left-traverses. Thus, the correspondence between faces and left-traverses is not one-to-one but one-to-many.

Note that a left-traverse is not a cycle in general; for example, Figure 2.1 (b). But if G is 2-connected, then any left-traverse of G is a cycle; see, e.g., [26]. (A graph is *2-connected* if every pair $\{u, v\}$ of its vertices are connected by two paths not sharing vertices except for $\{u, v\}$, or equivalently, it has no *cut vertex*.)

Proposition 2.2. Consider any plane graph G that is also 2-connected. Then any left-traverse is a cycle. That is, every face of G has a boundary consisting of one cycle, which we call a *face boundary cycle*.

Definition 2.2 (Face boundary representation and face size). For any 2-connected plane graph G , for any face, its *face boundary representation* is its face boundary cycle c , or more precisely, a sequence of vertices of c in the order so that the face is located left of c w.r.t. the order. The *size* of a face is the number of vertices of its face boundary cycle.

Definition 2.3 (Incidence and edge-incidence). Two faces are *incident* (resp., *edge-incident*) if their boundaries share some vertex (resp., some edge).

Remark. In general, we say that two graph objects are *incident* if they share some vertex. For example, we say that two paths are incident if they share some vertex (and possibly more).

Definition 2.4 (Complete face information). A *complete face information* of a graph G is a list of all face boundary representations of G and lists of respectively all incident and edge-incident pairs of faces of G .

While it is not essential for our discussion, we note here that a complete face information of any plane graph with n vertices can be expressed in $\tilde{O}(n)$ -space, which is easy to show by the Euler's formula on the number of vertices, edges, and faces of a plane graph.

We introduce a generalized separator notion, which will be mainly used in this dissertation. First we define the standard separator notion formally.

Definition 2.5 (Standard separator). For any graph G with n vertices, and for any parameter $\alpha \in (0, 1)$, an α -separator of G is a set S of vertices of G such that the removal of S creates two disconnected subgraphs each of which has at least αn vertices.

A generalized separator notion is defined for weighted graphs. Here a weighted graph² is a plane graph G for which we additionally give a weight to each face of G . Formally, we may represent a weighted graph $G = (V, E)$ by V , E , its combinatorial planar embedding, and its complete *weighted* face information where each face representation is also given its weight. By normalizing weights, we may assume that the total weight is always 1.

Definition 2.6 (Cycle weighted separator). For any 2-connected weighted graph, and for any $\rho \in (0, 1)$, a *cycle weighted ρ -separator* of the graph is a cycle C that creates two subplanes each of which has weight $\geq \rho$.

Remark. The weight of a subplane is simply the sum of the weights of all faces in the subplane.

2.2 Computational model

For discussing sublinear-space algorithms formally, we use the standard multi-tape Turing machine model. A multi-tape Turing machine consists of a read-only input tape, a write-only output tape, and a constant number of work tapes.

²Our explanation is simplified from [47] by assuming zero weight to all vertices and edges.

2. PRELIMINARIES

The space complexity of this Turing machine is measured by the total number of cells that can be used as its work tapes. Throughout this paper, we will use n to denote the number of vertices of an input graph and use it as our complexity parameter; that is, we will measure the space and time complexity in terms of n only (unless stated explicitly in other ways). As a unit for space complexity, we will often use “word”, by which we mean $c_0 \log n$ cells, where c_0 is a constant large enough to keep an input graph vertex index (i.e., a number between 1 and n) in one word in our Turing machine model. We use $\tilde{O}(s(n))$ to mean $O(s(n))$ words. Thus, for example, by a “ $\tilde{O}(s(n))$ -space algorithm” we formally mean a Turing machine that implements the algorithm by using $O(s(n) \log n)$ -size work tapes, i.e., work tapes consisting of $O(s(n) \log n)$ cells in total, for processing an input graph with n vertices.

For the sake of explanation, we will follow a standard convention and give a sublinear-space algorithm by a sequence of constant number of sublinear-space subroutines $A_1, \dots, A_k$ such that each A_i computes, from its given input, some output that is passed to A_{i+1} as input. Note that some of these outputs cannot be stored in a sublinear-size work tape; nevertheless, there is a standard way to design a sublinear-space algorithm based on these subroutines. The key idea is to compute intermediate inputs every time when they are necessary. For example, while computing A_i , when it is necessary to see the j th bit of the input to A_i , simply execute A_{i-1} (from the beginning) until it yields the desired j th bit on its work tape, and then resume the computation of A_i using this obtained bit. It is easy to see that this computation can be executed in sublinear-space. Furthermore, while a large amount of extra computation time is needed, we can show that the total running time can be polynomially bounded if all subroutines run in polynomial-time. In this dissertation, we in fact do not calculate the exact value of the exponent of polynomial of time complexity, but, of course, we will prove that our algorithms run in polynomial-time by using the above argument.

Chapter 3

Modification of Miller's Separator Algorithm

Miller [47] proposed a linear-time algorithm for computing small separators for 2-connected planar graphs. In this chapter, we explain his algorithm and present a way to modify it to a space efficient version [11]. Our algorithm can be regarded as a log-space reduction from the separator construction to the breadth first search tree construction.

3.1 Outline of the algorithm

We introduce notations used in this chapter, and then state our main result in this chapter and the outline of our algorithm.

We consider simple, undirected, 2-connected, and weighted plane graphs. Consider a plane graph $G = (V, E)$. We will use, e.g., $\vec{c}$ to denote a directed cycle, and use $(\vec{c})^{-1}$ to denote its reverse directed cycle. (We will also use, e.g., $\vec{e}$ and $(\vec{e})^{-1}$ to denote a directed edge and its reversal.) For a directed cycle $\vec{c}$, we regard its left hand side w.r.t. its direction as the *inside* of $\vec{c}$, and the right hand side is called the *outside* of $\vec{c}$. We use $\text{ins}(\vec{c})$ and $\text{out}(\vec{c})$ to denote respectively the inside and the outside of $\vec{c}$. Precisely speaking, the *boundary cycle* $\vec{c}$ is not included in $\text{ins}(\vec{c})$ or $\text{out}(\vec{c})$; that is, the graph is divided into three disjoint parts, namely, $\text{ins}(\vec{c})$, $\vec{c}$, and $\text{out}(\vec{c})$. By $\text{ins}^+(\vec{c})$ (resp., $\text{out}^+(\vec{c})$) we denote the subgraph induced by the vertices of $\text{ins}(\vec{c})$ (resp., $\text{out}(\vec{c})$) and $\vec{c}$. In general, for any set C of directed cycles, we define $\text{ins}(C)$ as the intersection of the insides of all directed cycles of C , which is simply called the *inside* of C . Note that determining whether a given vertex is located inside of a given directed cycle under an assumed planar embedding is $O(\log n)$ -space computable by using the undirected graph reachability algorithm of Reingold [53].

In this chapter, we assume that an input graph G is a weighted graph; that is, each face is assigned a positive value as its *weight*. For each face $\vec{f}$, we use $w(\vec{f})$ to denote the weight of $\vec{f}$ (w.r.t. the graph considered in each context). We

assume that the sum of the weights of all faces is 1. (For simplicity, we would use integers for specifying the weights of faces considering that their normalized values divided by their total are real weights.)

In summary, we assume that an input to our separator algorithm is a plane weighted graph $G = (V, E)$ with n vertices that is specified by (i) its combinatorial embedding (i.e., a list of cyclic order enumerations of adjacent vertices of $v \in V$), and (ii) a list of all faces, (i.e., their boundary cycles) and their weights (in integers). We assume that each integer weight can be expressed in $O(\log n)$ bits so that we may assume that the input can be encoded in $O(n \log n)$ bits. Note that there is a $O(\log n)$ -space (and hence, polynomial-time) algorithm that determines whether a given graph is planar and (if it is a planar) computes one of its planar embeddings [3]. Furthermore, it is easy to enumerate all boundary cycles for faces in $O(\log n)$ -space from a combinatorial planar embedding. Thus, if there is some simple way to define the weight of a face, then we can prepare the above input data from any standard undirected graph representation in $O(\log n)$ -space.

Consider any weighted plane graph $G = (V, E)$. A cycle of G (or more specifically, a directed cycle $\vec{c}$) is called a *weighted cycle separator* of G if both the weight of the inside $\vec{c}$ and that of the outside of $\vec{c}$ is at most $2/3$. Here the weight of the inside (resp., the outside) of $\vec{c}$ is the sum of the weights of all faces located in the inside (resp., the outside) of $\vec{c}$. The *size* of a separator $\vec{c}$ is simply the number of vertices of $\vec{c}$.

We state the theorem of Miller [47] and our result.

Theorem 3.1 ([47]¹). For any weighted plane graph G with n vertices, if it is 2-connected and all its faces have weight $\leq 2/3$ and size $\leq d$, then there exists a weighted cycle separator of size $O(\sqrt{dn})$. Furthermore, we have an algorithm that, for a given weighted plane graph G satisfying the above conditions, computes such a weighted cycle separator in $\tilde{O}(n)$ -time (and hence $\tilde{O}(n)$ -space).

In this chapter, based on the above algorithm of Miller, we propose an algorithm that is space efficient in the following sense.

Theorem 3.2 ([11]). We have an algorithm that, for a given weighted plane graph G satisfying the conditions of the above theorem, computes a weighted cycle separator of size $O(\sqrt{dn})$ in $\tilde{O}(f)$ -space and polynomial-time w.r.t. n , where f is the number of faces of G .

We explain the outline of our algorithm. Some technical terms (i.e., quoted ones) are used without definition; see the corresponding sections for the details. The outline of our algorithm is essentially the same as that of Miller's algorithm. Let G be a given weighted plane graph satisfying the conditions of the theorem. We consider the following four steps, each of which is conducted by a space efficient algorithm. (Since the polynomial-time computability of these

¹Miller [47] gave a more precise separator size bound by specifying also a constant factor. He also showed a polylogarithmic parallel time algorithm. But we omit such discussions in this dissertation.

3. MODIFICATION OF MILLER'S SEPARATOR ALGORITHM

algorithms is clear, we will omit mentioning their running time in the following discussion.)

1. In this step we reorganize G as a breadth first search tree (in short, BFS tree) $\mathcal{F}_G$ of the faces of G and construct a tree $\mathcal{T}_G$ consisting of cycles relevant to the faces in $\mathcal{F}_G$. We will refer $\mathcal{F}_G$ as a “face level tree” and $\mathcal{T}_G$ as a “frontier cycle tree.” (In order to distinguish from the input graph, we will use *node* to refer a vertex of a frontier cycle tree.)

We consider two algorithms for achieving the computation of this step. The first one is to compute a face level tree $\mathcal{F}_G$ for G , and the second one computes a frontier cycle tree $\mathcal{T}_G$ from $\mathcal{F}_G$. It should be noted here that the first one needs $\tilde{O}(f)$ -space, which is only the part of our algorithm that needs more than $\tilde{O}(1)$ -space, i.e., more than constant number of working variables. The other parts of the algorithm can be computed in $\tilde{O}(1)$ -space.

2. We modify the frontier cycle tree $\mathcal{T}_G$. This step consists of two substeps. First prune some branches and an ancestor part of $\mathcal{T}_G$. An algorithm for this substep creates a pruned frontier cycle tree $\mathcal{T}_{H'}$, which defines a subgraph H' of G . Secondly add back some nodes of $\mathcal{T}_G$ to $\mathcal{T}_{H'}$ to obtain a frontier cycle tree $\mathcal{T}_H$ so that all its leaf nodes correspond to small size cycles in G . A subgraph H of G is also produced from $\mathcal{T}_H$.

3. In this step we construct a small diameter spanning tree of H . More specifically, we create a spanning tree T of H such that any pair of vertices of H has a path of length $O(\sqrt{dn})$ in T .

Our algorithm for this task is different from the one given by Miller [47] that uses breadth first search in H starting from all but one edge in the root cycle of $\mathcal{T}_H$, for which we need $\tilde{O}(n)$ -space. Our algorithm produces a tree whose diameter may become at most four times larger; but it is simpler, and it can be computed in $\tilde{O}(1)$ -space.

4. Finally construct a desired weighted cycle separator $\vec{c}_*$. Roughly speaking, an algorithm for this step computes $\vec{c}_*$ by collecting faces in an appropriate cycle $\vec{c}_{\vec{e}}$ formed by a non-tree edge $\vec{e} = (u, v)$ of H and the path connecting v and u in T . We show a way to implement this computation in $\tilde{O}(1)$ -space.

In the following sections, we explain each step and show the corresponding algorithms. Throughout the explanation, let $G = (V, E)$ denote an input plane graph that satisfies the conditions of the theorem. Let F_G denote the set of the faces of G . For this G , as mentioned in the above theorems, we will use n , d , and f to denote respectively the number of its vertices, an upper bound for the size of its faces, and the number of its faces.

3.2 Face level tree and frontier cycle tree

We reorganize a graph as a face level tree and a frontier cycle tree for the subsequent algorithms. For a pair of faces, we say that they are *incident* if their boundary cycles share at least one vertex. We extend this notion to a relation

3.2. FACE LEVEL TREE AND FRONTIER CYCLE TREE

including cycles in general; we say that, for example, a face is *incident* to a cycle if the boundary cycle of the face has a common vertex with the cycle. Note here that the incidence relation can be checked in $\tilde{O}(1)$ -space, i.e., by using only constant number of words, for a given pair of (boundary) cycles.

A *face level tree* $\mathcal{F}_G$ and a *frontier cycle tree* $\mathcal{T}_G$ are defined by a breadth first search process w.r.t. the distance induced naturally by the incidence relationship among the faces of G . Specifically, $\mathcal{F}_G$ and $\mathcal{T}_G$ are defined as follows. Fix any face f_0 of G , and let $\vec{c}_0$ be its boundary cycle. We regard this f_0 as a *level zero* face and as the root face of $\mathcal{F}_G$. The directed cycle $\vec{c}_0$ is the root node of $\mathcal{T}_G$. We then consider the faces in $\text{out}(\vec{c}_0)$ that are incident to $\vec{c}_0$, which are called *level one* faces (see Figure 3.1), i.e., the faces of the next level of $\mathcal{F}_G$. Let F_1 be the set of all such level one faces, and let $\cup F_1$ denote the subgraph of G consisting of the boundary cycles of the faces of F_1 . Then from the 2-connectivity of G , it follows that $\cup F_1$ is also 2-connected (Lemma 3.3) and its boundary is defined by a set of boundary cycles with disjoint outsides. (Note that the cycle $(\vec{c}_0)^{-1}$, the reversal of $\vec{c}_0$, is one of these boundary cycles.) These cycles except for $(\vec{c}_0)^{-1}$ are called *frontier cycles* and regarded as the child nodes of the root $\vec{c}_0$ in the frontier cycle tree $\mathcal{T}_G$. In this way, level one nodes of $\mathcal{F}_G$ and $\mathcal{T}_G$ are defined.

Now consider any one of these frontier cycles, and denote it as $\vec{c}_1$. Our breadth first search process explores the outside of $\vec{c}_1$, i.e., $\text{out}(\vec{c}_1)$. We identify faces in $\text{out}(\vec{c}_1)$ that are incident to $\vec{c}_1$ (if they exist) as level two faces of $\mathcal{F}_G$. Then the child nodes of $\vec{c}_1$ in the frontier cycle tree $\mathcal{T}_G$ are defined as the boundary cycles of these level two faces (except for $(\vec{c}_1)^{-1}$). The face level tree and the frontier cycle tree of G are defined in this way. This process is justified by the following lemma.


Lemma 3.3. For any directed cycle $\vec{c}$ of G , let F be the set of faces of G outside of $\vec{c}$ that are incident to $\vec{c}$, and let $\cup F$ be the subgraph of G consisting of the boundary cycles of F . Then $\cup F$ is 2-connected. Hence, there exists a set C of directed cycles such that $\text{ins}^+(C) = \cup F$. Note that C has at least $(\vec{c})^{-1}$.

Proof. We need to show the 2-connectivity of $\cup F$ since the rest of the lemma is from well-known properties of 2-connected graphs; see, e.g., [26]. Thus, we consider any two distinct vertices $u, v \in \cup F$ and show that there exist two vertex disjoint paths between u and v on $\cup F$.

When both u and v are on $\vec{c}$, there exist two paths that are parts of the cycle $\vec{c}$ from u to v in clockwise and anti-clockwise directions.

Next consider the case where $u \in \vec{c}$ and $v \notin \vec{c}$. Let $\vec{f} \in F$ be a face on which v lies, and let $\vec{p}_v$ denote a path following the boundary $\vec{f}$ that begins and ends respectively with a vertex of $\vec{c}$ and that contains no edge whose reversal is in $\vec{c}$. (Intuitively, two endpoints a and b of $\vec{p}_v$ are vertices of $\vec{c}$ closest to v following the boundary $\vec{f}$. We call such a path a *face boundary path* (w.r.t. $\vec{f}$) attached to $\vec{c}$.) If $a \neq b$, then we have two disjoint paths from v to u that use $\vec{p}_v$ and $\vec{c}$; one goes through a , and the other goes through b . If $a = b$, then we use a face $\vec{f}' \in F$ whose boundary has a vertex w that is reachable from v without going through a . Let $\vec{p}_w$ be a face boundary path w.r.t. $\vec{f}'$ attached to $\vec{c}$ containing w . We may assume that it has an endpoint c on $\vec{c}$ that is different from a . If there

3. MODIFICATION OF MILLER'S SEPARATOR ALGORITHM



An example of the breadth first search process of the faces of a given graph. In this figure, the process is started by choosing a face numbered (0) as a level zero face. The next level faces, i.e., level one faces are faces numbered (1), which are all incident to the level zero face. It is easy to see that a subgraph consisting of boundary cycles of these level one faces is 2-connected and that its boundary is a set of five boundary cycles (those started with arrowed edges labeled by A, B, C, D, and E), including the reversal of the boundary cycle of the level zero face (the one started with edge C).

Figure 3.1: Level one faces and their boundary cycles

were no such face like $\vec{f}'$, then G has a face whose boundary is not a simple cycle contradicting the fact that a boundary of a face of a 2-connected graph is always a simple cycle (and the assumption that G is 2-connected). Thus we have two disjoint paths from v to u . One goes through a , and the other goes through w and then c .

Finally we consider the case where neither u nor v are on $\vec{c}$. Consider two face boundary paths attached to $\vec{c}$; one denoted as $\vec{p}_u$ containing u and the other denoted as $\vec{p}_v$ containing v . (If $\vec{p}_u = \vec{p}_v$, then we clearly have two disjoint paths between u and v because they are on the same face boundary cycle.) If both $\vec{p}_u$ and $\vec{p}_v$ have two endpoints on $\vec{c}$, then we can find two disjoint paths between u and v . In the case where $\vec{p}_u$ (respectively, $\vec{p}_v$) has only one endpoint, we consider, as before, another face boundary path attached to $\vec{c}$ that has a different endpoint on $\vec{c}$ and that is reachable from u (respectively, v), which can be used as the second path to $\vec{c}$ from u (respectively, from v). $\square$

We now explain the algorithm for computing the face level tree $\mathcal{F}_G$ of G and the frontier cycle tree $\mathcal{T}_G$ of G . We achieve this task by two substeps computed by algorithms $A_{1,1}$ and $A_{1,2}$. Algorithm $A_{1,1}$ computes $\mathcal{F}_G$, and then by traversing $\mathcal{F}_G$, $A_{1,2}$ produces the set of frontier cycles, i.e., the nodes of $\mathcal{T}_G$. It also computes (i) parent-child relationship among the frontier cycles in $\mathcal{T}_G$, and (ii) for each node of $\mathcal{T}_G$ (i.e., a boundary cycle $\vec{c}$), the information on which faces of $\mathcal{F}_G$ are the faces of $\text{out}(\vec{c})$ and incident to $\vec{c}$.


The algorithm $A_{1,1}$ first chooses one face $\vec{f}_0$, say, the first face of F_G given in the input, and it outputs $\vec{f}_0$ as a root and the level zero face of $\mathcal{F}_G$. Then it identifies the level one faces that are incident to $\vec{f}_0$, and it outputs identified faces as the children of $\vec{f}_0$ (with their level). The algorithm uses an array $\mathbf{F}$ to keep such identified faces. In general, the algorithm repeats the following: (i) select the first unmarked face $\vec{f}$ in $\mathbf{F}$, (ii) mark it, (iii) output its incident faces not in $\mathbf{F}$ as the child faces of $\vec{f}$ with the next level of $\vec{f}$, and (iv) store these newly identified faces in $\mathbf{F}$ unmarked. In this way, $A_{1,1}$ produces the face level tree $\mathcal{F}_G$ of G in the breadth first way in $\tilde{O}(f)$ -space.

Algorithm $A_{1,2}$ computes the frontier cycles of $\mathcal{T}_G$. For this, at each level, the algorithm checks all faces of this level to identify their boundary cycles. Let F and $\cup F$ denote respectively the set of such faces (as well as their boundary cycles) and the subgraph of G consisting of (the undirected version of) the edges of the boundary cycles of F . We say that a directed edge of F is *solitary* if its reversal does not appear in F . It is easy to see (e.g., in Figure 3.2) that the set of solitary edges forms the boundary cycles of $\cup F$. Thus, we can find a boundary cycle $\vec{c}$ of $\cup F$ by traversing adjacent solitary directed edges from any one of the solitary directed edges of $\vec{c}$ until coming back to the start vertex. When traversing, if there are more than one solitary directed edges adjacent to the current edge $\vec{e} = (u, v)$, then we choose the one next to $\vec{e}$ in the counterclockwise order in the solitary directed edges leaving from v . In this way the algorithm $A_{1,2}$ identifies the boundary cycles of $\cup F$ and outputs those starting from the solitary directed edge with the lowest index in each boundary cycle (so that each boundary cycle is produced exactly once). During its computation for producing a boundary cycle $\vec{c}$, the algorithm can also determine which faces in the next level in $\mathcal{F}_G$ are the faces in $\text{out}(\vec{c})$ and incident to $\vec{c}$. By using this information, the algorithm can also compute the required additional information on the structure of $\mathcal{T}_G$ and $\mathcal{F}_G$.

Note that Lemma 3.3 guarantees that the traversed boundary edges form a cycle; thus, for identifying one boundary cycle in the above computation, the algorithm needs to keep only the start vertex and the currently traversed boundary edge in its work memory. Then it is easy to see that the computation of $A_{1,2}$ as a whole can be done in $\tilde{O}(1)$ -space.

3.3 Pruning a frontier cycle tree

We show two algorithms $A_{2,1}$ and $A_{2,2}$ for computing a subgraph H of G by merging the faces of G so that H has small diameter and yet it consists of faces of small size. The construction is based on the frontier cycle tree $\mathcal{T}_G$ obtained by the previous step. The algorithm $A_{2,1}$ prunes subtrees under some frontier cycles to shorten the depth of $\mathcal{T}_G$ to make a tree $\mathcal{T}_{H'}$. Note that for each frontier cycle $\vec{c}$, i.e., a node of $\mathcal{T}_G$, a subtree of $\mathcal{T}_G$ under $\vec{c}$ represents a subgraph of G located outside of $\vec{c}$. Thus, pruning this subtree means to merge the faces of this part to create a new face with $(\vec{c})^{-1}$ as its boundary cycle. The weight of the new face is the sum of the weights of the merged faces, which is denoted by



An example of boundary cycles of a frontier cycle tree. The boundary cycles $\vec{c}$, $\vec{c}_i$, and $\vec{c}_{i+1}$ are the cycles bounding the corresponding labeled areas. The figure shows the situation when the outside of $\vec{c}$ is explored and $\vec{c}_i$ and $\vec{c}_{i+1}$ are identified as child cycles of $\vec{c}$. Note, for example, $(\vec{c}_i)^{-1}$, one of the frontier cycles, consists of the solitary boundary edges of the faces $\vec{f}_1, \vec{f}_2, \vec{f}_3$ that are incident to $\vec{c}$.

Figure 3.2: Child boundary cycles $\vec{c}_i$ and $\vec{c}_{i+1}$ of $\vec{c}$

$w(\text{out}(\vec{c}))$, extending our notation for face weights. Recall that we assume that the size of every face of G , i.e., the number of vertices of its boundary cycle is at most d . This pruning process may create faces with large size. The algorithm $A_{2,2}$ then adds back some removed nodes of $\mathcal{T}_G$ to $\mathcal{T}_{H'}$ to extend a tree from each new leaf node of $\mathcal{T}_{H'}$ until the obtained subtree consists of only leaf frontier cycles with small size; the obtained subtree of $\mathcal{T}_G$ is denoted by $\mathcal{T}_H$. We can show that such a subtree exists within our desired depth.

Now we explain these two algorithms in detail. We first state the goal of these algorithms precisely. Below by, e.g., $w(\text{out}(\vec{c}))$ and $|\vec{c}|$, we mean respectively the sum of the weights of the faces outside of $\vec{c}$ and the size, i.e., the number of vertices, of $\vec{c}$.

Lemma 3.4. We have $\tilde{O}(1)$ -space algorithms $A_{2,1}$ and $A_{2,2}$ that compute from a given frontier cycle tree $\mathcal{T}_G$ (together with an original input) its subtree $\mathcal{T}_H$ such that (i) its depth h is $O(\sqrt{dn})$, (ii) the cycle $\vec{c}_i$ corresponding to its root node satisfies $w(\text{ins}(\vec{c}_i)) \leq 2/3$ and $|\vec{c}_i| = O(\sqrt{dn})$, and (iii) for every cycle $\vec{c}$ corresponding to one of its leaf nodes, we have $w(\text{out}(\vec{c})) \leq 2/3$ and $|\vec{c}| = O(\sqrt{dn})$.

Remark. From (ii) and (iii) it follows that the subgraph H of G induced by the cycles of $\mathcal{T}_H$ and the faces between them consists of faces of weight $\leq 2/3$ and size $O(\sqrt{dn})$. (Note that $d \leq \sqrt{dn}$.)

Consider the algorithm $A_{2,1}$. It first computes the weight of (the outside of) every frontier cycle, which is regarded as the weight of the corresponding node of

$\mathcal{T}_G$. Note that a frontier cycle $\vec{c}$ corresponding to a leaf node of $\mathcal{T}_G$ is a directed cycle with no vertex in its outside; that is, it is a reversal of the boundary cycle of some face of G , and the weight of this face is the weight of (the outside of) $\vec{c}$. Thus, we can compute the weights of all frontier cycles from the leaf nodes. Clearly, this substep can be computed in $\tilde{O}(1)$ -space. The algorithm $A_{2,1}$ then selects “trunk” nodes. Starting from the root node the algorithm chooses, as the next *trunk node*, a node with the largest weight among all child nodes of the current trunk node. Let $\vec{c}_0, \dots, \vec{c}_t$ denote the sequence of selected trunk nodes, which forms the *trunk path* of $\mathcal{T}_G$ from its root $\vec{c}_0$ to one of its leaf nodes $\vec{c}_t$. Again this trunk path is $\tilde{O}(1)$ -space computable. Now the algorithm cuts this trunk path further to get a subtree $\mathcal{T}_{H'}$ of $\mathcal{T}_G$ with a small diameter. That is, the algorithm computes a subpath $\vec{c}_i, \vec{c}_{i+1}, \dots, \vec{c}_j$ of the trunk path, and $\mathcal{T}_{H'}$ is the subtree of $\mathcal{T}_G$ consisting of this subpath and the child nodes of each trunk node. For this, the algorithm searches the trunk path from $\vec{c}_0$ for the first $\vec{c}_j$ such that $w(\text{ins}(\vec{c}_j)) \geq 1/3$. It then searches from $\vec{c}_j$ for an ancestor $\vec{c}_i$ with small height $h := j - i$ from $\vec{c}_j$ and small size; more precisely, the root node $\vec{c}_i$ of $\mathcal{T}_{H'}$ that satisfies (i) and (ii) of Lemma 3.4. The existence of such a cycle is shown in [47]; but for the sake of completeness, we restate the proof here.

Claim 3.1. Let n' be the number of vertices of the subgraph of G consisting of the cycles of the trunk subpath $\vec{c}_0, \dots, \vec{c}_j$. Then there exists a cycle $\vec{c}_i$ in the trunk path above $\vec{c}_j$ such that $2d'h + |\vec{c}_i| \leq 2\sqrt{d'n'}$, where $d' = \lfloor d/2 \rfloor$ and $h = j - i$.

Proof. We prove by contradiction. Assume $|\vec{c}_i| > 2\sqrt{d'n'} - 2d'(j - i)$ for every $i \leq j$. When $j - i \leq \sqrt{n'/d'}$, we have $2\sqrt{d'n'} - 2d'(j - i) \geq 0$. Now,

$$\begin{aligned} n' &= \sum_{i=0}^j |\vec{c}_i| > \sum_{i=j-\lfloor \sqrt{n'/d'} \rfloor}^j 2\sqrt{d'n'} - 2d'(j - i) \\ &= 2\sqrt{d'n'} \cdot \left(\left\lfloor \sqrt{\frac{n'}{d'}} \right\rfloor + 1 \right) - 2d' \sum_{h=0}^{\lfloor \sqrt{n'/d'} \rfloor} h \\ &= \left(2\sqrt{d'n'} - d' \cdot \left\lfloor \sqrt{\frac{n'}{d'}} \right\rfloor \right) \cdot \left(\left\lfloor \sqrt{\frac{n'}{d'}} \right\rfloor + 1 \right) \\ &> \sqrt{d'n'} \cdot \sqrt{\frac{n'}{d'}} = n'. \end{aligned}$$

□

Note that we have $w(\text{ins}(\vec{c}_i)) < 1/3$ since $\vec{c}_j$ is the first one from $\vec{c}_0$ that satisfies $w(\text{ins}(\vec{c}_j)) \geq 1/3$. On the other hand, each non-trunk child $\vec{c}$ of $\vec{c}_k$, $i \leq k \leq j - 1$, satisfies $w(\text{out}(\vec{c})) \leq 1/2$ as the trunk node $\vec{c}_{k+1}$ is a child node of $\vec{c}_k$ with the largest weight. Hence, including $\vec{c}_j$, for every leaf node $\vec{c}$ of $\mathcal{T}_{H'}$, we have $w(\text{out}(\vec{c})) \leq 2/3$; thus, $\mathcal{T}_{H'}$ satisfies (i) and (ii) of Lemma 3.4. This tree is the output of the algorithm $A_{2,1}$.

3. MODIFICATION OF MILLER'S SEPARATOR ALGORITHM

It may be possible that a cycle $\vec{c}$ corresponding to some leaf of $\mathcal{T}_{H'}$ has large size. The task of the algorithm $A_{2,2}$ is to add back some nodes of the subtree below $\vec{c}$ in $\mathcal{T}_G$ up to some “level” so that all leaf nodes of this level are cycles of small size. Here we say that a descendant node of $\vec{c}$ is of *level* k if there is a path of length k from $\vec{c}$ to this node in $\mathcal{T}_G$, and by “the size of the level k cycles” we mean the size of the largest level k cycle. By using the same argument as above, Miller showed the following.

Claim 3.2. Consider any cycle $\vec{c}$ corresponding to a leaf node of $\mathcal{T}_{H'}$ of size more than $\sqrt{2dn}$, and let n' be the number of vertices of the subgraph of G outside of $\vec{c}$. Then there exists some level k such that $d'k + s_k \leq 2\sqrt{d'n'}$, where $d' = \lfloor d/2 \rfloor$ and s_k is the size of level k cycles.

For each leaf cycle $\vec{c}$, the algorithm $A_{2,2}$ adds all frontier cycles of $\mathcal{T}_G$ under $\vec{c}$ up to level k such that $d'k + s_k \leq 2\sqrt{d'n'}$ holds. It produces the obtained tree as $\mathcal{T}_H$. Clearly, this is $\tilde{O}(1)$ -space computable. Furthermore, it is easy to see from our discussion that $\mathcal{T}_H$ satisfies the conditions (i) to (iii) of Lemma 3.4.

3.4 Constructing a spanning tree

We show an algorithm for constructing a small diameter spanning tree of H . Recall that H is the induced subgraph of G obtained by the second step; it consists of the vertices appearing in the frontier cycles of $\mathcal{T}_H$ and the faces between these frontier cycles. The set F_H of the faces of H (and their weights) is also obtained in the second step. Note that a face of H is either a face with a boundary cycle corresponding to the root or a leaf node² of $\mathcal{T}_H$, or a face in the original graph G . Then from the properties of $\mathcal{T}_H$ (Lemma 3.4) and the condition on the input graph G , we may assume that every face has weight $\leq 2/3$, and size $O(\sqrt{dn})$.

Lemma 3.5. We have an $\tilde{O}(1)$ -space algorithm A_3 that computes, for a given pair of $\mathcal{T}_H$ and F_H , a spanning tree T of H with diameter $O(\sqrt{dn})$.

Our algorithm A_3 is different from the one given in [47] for computing a spanning tree of H . Miller constructed a spanning tree of H using breadth first search starting from all but one edge of the root frontier cycle of $\mathcal{T}_H$. We avoid the $O(n)$ -space incurred by breadth first search and use computations that are local to the faces behind each frontier cycle. The diameter of our spanning tree can be larger by a factor of up to 4 compared with Miller's; but it can be constructed in $\tilde{O}(1)$ -space.

We first establish some terminologies. A face of H is called an *inner face* if it does not correspond to a node of $\mathcal{T}_H$; that is, it is a face from G between frontier cycles. Consider any inner face f . Note that it must be incident to at least 2 frontier cycles of $\mathcal{T}_H$. Among these cycles, its *parent cycle* is the cycle

²Precisely speaking, in the case of a cycle corresponding to a leaf node, its reversal is a boundary cycle of the corresponding face.

$\vec{c}$ that is closest to the root. Let $P_{\vec{f}}$ denote the set of directed face boundary paths w.r.t. $\vec{f}$ attached to $\vec{c}$, where their direction is the same as $\vec{f}$. Note that every edge of $\vec{f}$ whose reversal is not in $\vec{c}$ occurs in exactly one of the paths in $P_{\vec{f}}$. For $\vec{p} \in P_{\vec{f}}$, let $\ell(\vec{p})$ be the length of $\vec{p}$ and let p_i be the i th vertex of $\vec{p}$, i.e., $\vec{p} = (p_0, \dots, p_{\ell(\vec{p})})$. For any frontier cycle $\vec{c}$ of $\mathcal{T}_H$, let $P_{\vec{c}} = \bigcup \{P_{\vec{f}} \mid \vec{f} \text{ is an inner face of } H \text{ with parent cycle } \vec{c}\}$. Note that all vertices of frontier cycles corresponding to child nodes of $\vec{c}$ appear in $P_{\vec{c}}$. In the following, we often regard a path or a cycle as a set of vertices on the path or the cycle, and we write, e.g., $p_i \in \vec{p}$ for a path $\vec{p}$. We also write $p_i \in P_{\vec{f}}$.

The algorithm A_3 computes a set R of edges and then derives $T := H \setminus R$ as its output. It first adds one edge (say the largest one) from the root cycle to R . For each non-leaf node $\vec{c}$ of $\mathcal{T}_H$, the algorithm executes as follows.

For every path $\vec{p} \in P_{\vec{c}}$, the algorithm first decides *path directions* of the edges in $\vec{p}$. For any i with $0 \leq i < \lfloor \ell(\vec{p})/2 \rfloor$, the edge $\{p_i, p_{i+1}\}$ has path direction (p_i, p_{i+1}) , and for any i with $\lfloor \ell(\vec{p})/2 \rfloor \leq i < \ell(\vec{p})$, the edge $\{p_i, p_{i+1}\}$ has path direction (p_{i+1}, p_i) . Then it defines the *path distance* of an edge $\{p_i, p_{i+1}\}$ on $\vec{p}$. The path distance of $\{p_i, p_{i+1}\}$ is defined as $\min(i, \ell(\vec{p}) - i - 1)$. Now, for every edge e belonging to some path in $P_{\vec{c}}$, it decides the *actual direction* of e as follows. If e occurs only once on a path in $P_{\vec{c}}$, the actual direction of e is oriented as its path direction on that path. Consider the case where e occurs on two paths $\vec{p}, \vec{p}' \in P_{\vec{c}}$ as $\{p_i, p_{i+1}\}$ and $\{p'_j, p'_{j+1}\}$ respectively. If the path directions of e on $\vec{p}$ and $\vec{p}'$ are the same, the actual direction of e follows it. If the path directions of e on $\vec{p}$ and $\vec{p}'$ are different, then the actual direction of e follows the path direction on the path on which the path distance of e is less than the other. In the case where the path distances are equal, it follows the path direction on the path that is lexicographically smaller. Finally, the algorithm considers every vertex $v \in P_{\vec{c}}$ and all in-coming edges to v following their actual directions. When the in-degree of v is greater than 1, it leaves only one edge whose path distance is the smallest and adds the other edges to R . If there are several edges with smallest path distance, it leaves the one of them that is lexicographically smallest.

Let us now examine the correctness of this algorithm. Recall that T is obtained by removing edges in R from H . For showing T is a tree, we will show that the edges of T can be oriented such that every vertex has in-degree at most 1 and at least one vertex has in-degree 0. This means that the number of the edges is less than that of the vertices. Moreover, we will show that T is connected. Thus, T is a tree because a connected graph that has fewer edges than vertices must be a tree.

Claim 3.3. In the computed graph T , the number of the edges is less than that of the vertices.

Proof. We orient edges in T . Note that one edge of the root cycle of $\mathcal{T}_H$ is in R (and hence removed from H to define T). We orient the other edges of the root cycle to form a directed path. Orient the other edges of T by their actual directions. For any frontier cycle $\vec{c}$ and for any $v \in \vec{c}$, consider

3. MODIFICATION OF MILLER'S SEPARATOR ALGORITHM

$T[P_{\vec{c}}]$, a subgraph of T restricted on $P_{\vec{c}}$. We can observe, from the definition of path/actual direction, that no edge $\{u, v\}$ in $T[P_{\vec{c}}]$ gets oriented as (u, v) . From this observation, we can show that all vertices of T except for the root vertex of T , which has in-degree 0, have in-degree 1 in T . Hence, the claim follows. $\square$

Claim 3.3 and the following claim imply that T is a tree.

Claim 3.4. For every vertex $v \in P_{\vec{c}}$, there is a vertex $u \in \vec{c}$ such that the distance from u to v in $T[P_{\vec{c}}]$ is at most $\lceil d/2 \rceil$.

Proof. We show that if a vertex $v \in P_{\vec{c}}$ is incident to an edge e that has path distance k on some path $\vec{p} \in P_{\vec{c}}$, then v has distance at most $k + 1$ in $T[P_{\vec{c}}]$ to some vertex $u \in \vec{c}$. This shows the claim, as every edge in $\bigcup P_{\vec{c}}$ has path distance at most $\lceil d/2 \rceil - 1$.

The proof is by induction on k . Suppose that $e = \{p_k, p_{k+1}\}$ and $v = p_{k+1}$ (the case where $e = \{p_{\ell(\vec{p})-(k+1)}, p_{\ell(\vec{p})-k}\}$ and $v = p_{\ell(\vec{p})-(k+1)}$ is symmetric). If $k = 0$, the vertex p_0 itself lies on $\vec{c}$. We first consider the case where the edge e is present in T . Obviously, p_1 has distance 1 from $\vec{c}$. If the edge e is not present in T , the construction of T ensures that there is a unique incoming edge e' at v according to actual orientations, and that this edge e' has path distance 0. Thus p_1 has distance 1.

If $k > 0$, we apply the inductive hypothesis to the edge $e_{\text{prev}} = \{p_{k-1}, p_k\}$ and p_k has distance at most k to some vertex $u \in \vec{c}$. When the edge e is present in T , it follows that $v = p_{k+1}$ has distance at most $k + 1$ to the vertex u . If the edge e is not present in T , as with the case $k = 0$, the construction of T ensures that there is a unique incoming edge e' at v according to actual orientations, and that this edge e' has path distance at most k . Applying the above argument to the occurrence of v in e' , it follows again that v has distance at most $k + 1$ to some vertex $u \in \vec{c}$. $\square$

Using this claim, an easy induction by depth over the nodes of $\mathcal{T}_H$ shows that all vertices on a face whose parent cycle has depth k in $\mathcal{T}_H$ (this includes all vertices on depth $k + 1$ nodes in $\mathcal{T}_H$) have distance at most $2d(k + 1)$ to some vertex on the root cycle. From Claim 3.1 and 3.2, we can bound $2d(k + 1) = O(\sqrt{dn})$. As the size of the root cycle is bounded by $O(\sqrt{dn})$, this implies that the diameter of T is $O(\sqrt{dn})$.

The algorithm can clearly be implemented in $\tilde{O}(1)$ -space. In the first substep, it computes the set $P_{\vec{c}}$ for each frontier cycle $\vec{c}$ of $\mathcal{T}_H$. To compute R in the second substep, it just needs a constant number of variables for computing path directions, path distances and actual directions. Finally, the tree $T := H \setminus R$ is computed as the third substep.

3.5 Finding a separator

Finally we show an algorithm for constructing a weighted cycle separator for H of small size. Recall that H is the induced subgraph of G that is obtained

by merging the faces (and their weights) of G in the regions corresponding to the nodes pruned from $\mathcal{T}_G$ to create $\mathcal{T}_H$. Thus, the obtained weighted cycle separator for H is also a weighted cycle separator for G . Recall also that the algorithm can use the frontier cycle tree $\mathcal{T}_H$, the set F_H of faces of H (and their weights), and the spanning tree T of H with diameter $O(\sqrt{dn})$ obtained in the previous steps.

Lemma 3.6. We have an $\tilde{O}(1)$ -space algorithm A_4 that computes a cycle weighted separator $\vec{c}_*$ of H with size $O(\sqrt{dn})$ from given $\mathcal{T}_H$, F_H (with weight information), and T .

This step is almost the same as Miller's algorithm. We show here how Miller's algorithm for this step can be implemented in $\tilde{O}(1)$ -space. We first introduce some notation. For any non-tree edge $\{u, v\}$ (i.e., an edge of H but not of T), consider its directed version $\vec{e} = (u, v)$. Let $\vec{f}_{\vec{e}}$ be the face of H that has $\vec{e}$ on its boundary cycle, and let $\vec{c}_{\vec{e}}$ be the cycle in H that consists of $\vec{e}$ and the directed path from v to u in T . Note that $\vec{f}_{\vec{e}}$ belongs to the inside of $\vec{c}_{\vec{e}}$. Let $R_{\vec{e}}$ denote the tree that has the faces inside $\vec{c}_{\vec{e}}$ as nodes, where two faces are adjacent if they share a non-tree edge. (Note that $R_{\vec{e}}$ is indeed a tree: If it were disconnected, then T would contain a cycle, and if $R_{\vec{e}}$ contained a cycle, then T would not be a spanning tree of H .)

The algorithm A_4 uses a subroutine $A_{4,1}$ that computes, for a given directed non-tree edge $\vec{e} = (u, v)$, the cycle $\vec{c}_{\vec{e}}$, its weight $w(\vec{c}_{\vec{e}})$, and $R_{\vec{e}}$. We first explain this subroutine $A_{4,1}$. It works as follows for a given $\vec{e} = (u, v)$. (We again explain as if things obtained at each step are passed to the next step as input. The first two steps use ideas from the logspace reachability algorithm for undirected forests given by Cook and McKenzie [22].)

1. Compute first a directed walk $\vec{w}_{\vec{e}}$ from v to u in T : Starting from $v_0 = v$, pick a neighbor v_1 of v in T (say the lexicographically least one). For $i \geq 1$, pick v_{i+1} as the neighbor of v_i that is next after v_{i-1} in the counterclockwise order around v_i under the given planar embedding of H (inherited from G). Stop once u is reached; this always happens, as the above rules ensure that the whole tree T would be explored if the process was not aborted.
2. Compute $\vec{c}_{\vec{e}}$ as the sequence of edges that are used only in one direction in $\vec{w}_{\vec{e}}$ in the order they appear.
3. Compute a tree $R_{\vec{e}}$ that consists of faces located inside of $\vec{c}_{\vec{e}}$ as nodes and edges between two faces sharing an edge (in different directions) in $H \setminus T$. Then traverse this tree from the face $\vec{f}_{\vec{e}}$ to compute the weight $w(\vec{c}_{\vec{e}})$. Compute also $C_{\vec{e}} = \{\vec{e}' \mid \vec{f}_{\vec{e}'}$ is a child of $\vec{f}_{\vec{e}}$ in $R_{\vec{e}}\}$. (Since we know that $R_{\vec{e}}$ is a tree, traversing all nodes in $R_{\vec{e}}$ is $\tilde{O}(1)$ -space computable.)

Now we explain our algorithm A_4 . It first checks whether there is a face $\vec{f}$ of H with $1/3 \leq w(\vec{f}) \leq 2/3$; if so, then output $\vec{f}$ as the desired separator $\vec{c}_*$. It then uses $A_{4,1}$ to compute $\vec{c}_{\vec{e}}$ and $w(\vec{c}_{\vec{e}})$ for each directed non-tree edge $\vec{e}$, and if there is a non-tree edge $\vec{e}$ such that $1/3 \leq w(\vec{c}_{\vec{e}}) \leq 2/3$, then output $\vec{c}_{\vec{e}}$ as

3. MODIFICATION OF MILLER'S SEPARATOR ALGORITHM

$\vec{c}_*$. Note that the size of this separator $\vec{c}_*$ (i.e., the number of vertices of $\vec{c}_*$) is at most 1 plus the diameter of T , that is, $O(\sqrt{dn})$.


Otherwise, the algorithm fixes any directed non-tree edge $\vec{e} = (u, v)$ such that $w(\vec{c}_{\vec{e}}) > 2/3$ and $w(\vec{c}_{\vec{e}'}) < 1/3$ for all $\vec{e}' \in C_{\vec{e}}$, the set of directed edges corresponding to child faces of $f_{\vec{e}}$ in $R_{\vec{e}}$. It remains to select $S \subset C_{\vec{e}}$ so that the region $f_{\vec{e}} \cup \bigcup \{\vec{c}_{\vec{e}'} | \vec{e}' \in S\}$ has an appropriate weight and small enough boundary. The weight condition could be satisfied by a simple greedy strategy, but ensuring a small boundary requires more care.

Let $(v_1, \dots, v_\ell)$ be the traversal of $\vec{f}_{\vec{e}}$ with $v_1 = v$ and $v_\ell = u$. For $\vec{e}' \in C_{\vec{e}} \cup \{\vec{e}\}$, let

$$\begin{aligned} \text{rightmost}(\vec{e}') &= \min\{i \in \{1, \dots, \ell\} \mid v_i \in \vec{c}_{\vec{e}'}\}, \text{ and} \\ \text{leftmost}(\vec{e}') &= \max\{i \in \{1, \dots, \ell\} \mid v_i \in \vec{c}_{\vec{e}'}\}. \end{aligned}$$

For $\vec{e}_1, \vec{e}_2 \in C_{\vec{e}} \cup \{\vec{e}\}$, we say that $\vec{e}_1$ is *dominated* by $\vec{e}_2$ (written as $\vec{e}_1 \leq_d \vec{e}_2$) if both $\text{rightmost}(\vec{e}_2) \leq \text{rightmost}(\vec{e}_1)$ and $\text{leftmost}(\vec{e}_1) \leq \text{leftmost}(\vec{e}_2)$ hold.

The algorithm computes the rooted tree $D_{\vec{e}}$ on the vertex set $C_{\vec{e}} \cup \{\vec{e}\}$ that is the transitive reduction of $\leq_d$. (Acyclicity is guaranteed because if one edge was dominated by two incomparable edges, then it would follow that T contains a cycle.) Then, for $\vec{e}' \in C_{\vec{e}} \cup \{\vec{e}\}$, the algorithm computes $W_{\vec{e}'} := \sum_{\vec{e}'' \leq_d \vec{e}'} w(\vec{c}_{\vec{e}''})$. If there is $\vec{e}' \in C_{\vec{e}}$ such that $1/3 \leq w(\vec{f}_{\vec{e}}) + W_{\vec{e}'} \leq 2/3$, then the algorithm outputs the cycle as $\vec{c}_*$ that consists of the path from $v_{\text{rightmost}(\vec{e}')}$ to $v_{\text{leftmost}(\vec{e}')}$ in T and the path from $v_{\text{leftmost}(\vec{e}')$ to $v_{\text{rightmost}(\vec{e}')$ in $\vec{f}_{\vec{e}}$. Clearly, its length is bounded by the diameter of T plus the size of $\vec{f}_{\vec{e}}$. See Figure 3.3 for an example.



Edges in T are drawn as thick lines, and edges on $\vec{c}_*$ are drawn as arrowed edges (except for $\vec{e}'$, whose arrow shows a direction of $\vec{f}_{\vec{e}'}$).

Figure 3.3: An example where $1/3 \leq w(\vec{f}_{\vec{e}}) + W_{\vec{e}'} \leq 2/3$


Otherwise there is $\vec{e}' \in C_{\vec{e}} \cup \{\vec{e}\}$ such that $W_{\vec{e}'} + w(\vec{f}_{\vec{e}}) > 2/3$ and $W_{\vec{e}''} + w(\vec{f}_{\vec{e}}) < 1/3$ for each child $\vec{e}''$ of $\vec{e}'$ in $D_{\vec{e}}$. The algorithm greedily selects a subset S of the children of $\vec{e}'$ in $D_{\vec{e}}$ so that the reversals of the edges in $S \cup \{\vec{e}'\}$ are consecutive on $\vec{f}_{\vec{e}}$ among the non-tree edges and $1/3 \leq w(\vec{f}_{\vec{e}}) + \sum_{\vec{e}'' \in S} W_{\vec{e}''} \leq$

3.6. LOG-SPACE REDUCTION TO BFS TREE CONSTRUCTION

2/3. Let S_R and S_L be the partition of S into edges whose reversals occur before and after $\vec{e}$ on $\vec{f}_{\vec{e}}$, respectively. Further, let

$$\begin{aligned} i_L &= \min\{\text{rightmost}(e'') \mid e'' \in S_L\}, \\ j_L &= \max\{\text{leftmost}(e'') \mid e'' \in S_L\}, \\ i_R &= \min\{\text{rightmost}(e'') \mid e'' \in S_R\}, \text{ and} \\ j_R &= \max\{\text{leftmost}(e'') \mid e'' \in S_R\}. \end{aligned}$$

The algorithm outputs as $\vec{c}_*$ the cycle that consists of the paths on $\vec{f}_{\vec{e}}$ from v_{j_R} to v_{i_L} and from v_{j_L} to v_{i_R} , and of the paths in T from v_{i_L} to v_{j_L} and from v_{i_R} to v_{j_R} . Clearly, its length is bounded by the size of $\vec{f}_{\vec{e}}$ plus twice the diameter of T , which is bounded by $O(\sqrt{dn})$. See Figure 3.4 for an example.



Edges in T are drawn as thick lines, and edges on $\vec{c}_*$ are drawn as arrowed edges (except for $\vec{e}$, whose arrow shows a direction of $\vec{f}_{\vec{e}}$).

Figure 3.4: The remaining case

It is not hard to see that each step of the algorithm can be implemented in $\tilde{O}(1)$ -space.

3.6 Log-space reduction to BFS tree construction

As mentioned in Section 2.2, almost all parts of the algorithm are $\tilde{O}(1)$ -space computable; that is, they can be implemented as $O(\log n)$ -space Turing machines. The only exception is the algorithm $A_{1,1}$ for computing a face level tree from the face incidence relation of an input graph. This task is essentially to compute a BFS tree from a given graph. Thus, from the view point of the space-limited computability, we can claim that the task of computing BFS tree construction is the key for computing a separator.

We state this observation formally. For this, we define the BFS tree construction problem as the following decision problem.

3. MODIFICATION OF MILLER'S SEPARATOR ALGORITHM

n' -size BFS tree construction problem

(decision version)

input: A graph $G' = (V', E')$ with n' vertices,
a root vertex $r \in V'$, and
a pair of vertices $u, v \in V'$.
task: Determine whether u is a child of v in
a BFS tree rooted from r under any fixed
ordering among child vertices.

Note that a BFS tree for given graph G' and root vertex r may differ depending on the order of visiting child vertices of a parent vertex. On the other hand, the distance of each vertex from r in a BFS tree, which is important for our algorithm, is independent from such an ordering. Also note that we can design a log-space Turing machine that indeed produces a BFS tree by using an oracle that solves this decision problem. That is, the task of a BFS tree construction is log-space Turing reducible to this decision problem. Now as a corollary of our algorithm design, we have, for example, the following relation.

Theorem 3.7. The task of computing a separator of size $O(\sqrt{n})$ for a given planar graph with n vertices is $O(\log n)$ -space Turing reducible to the $O(n)$ -size BFS tree construction problem (decision version).

Remark. As shown in the proof, we consider a triangulated version of a given graph, which is clearly bi-connected (except for the trivial case). Thus, this theorem holds in general for any planar graph.

Proof. We may assume that a given graph G is nontrivial one and has more than three vertices.

Consider a process for obtaining a separator for a given planar graph G with n vertices. First we compute [3] one of its combinatorial embeddings to the plane in $O(\log n)$ -space. Then from this embedding, we convert G to a *triangulated graph* $\hat{G}$ (i.e., a graph consisting of triangle faces) by adding edges, which is easy to do in $O(\log n)$ -space. Note that $\hat{G}$ is bi-connected and that it has at most $3n$ faces from Euler's formula. Now a desired separator of G is obtained by applying our algorithm of Theorem 5.2 to $\hat{G}$ assuming the unit weight to every triangle face because a weighted separator of $\hat{G}$ can be regarded as a separator of the original graph G . As mentioned above, we only need $O(\log n)$ -space to execute our algorithm by using an oracle solving the $\hat{f}$ -size BFS tree construction problem (for some $\hat{f} \leq 3n$). Therefore, this whole process can be regarded as a log-space Turing reduction to the $O(n)$ -size BFS tree construction problem. $\square$

3.6. LOG-SPACE REDUCTION TO BFS TREE CONSTRUCTION

Chapter 4

A Sublinear-space and Polynomial-time Planar Separator Algorithm

The purpose of this chapter is to give a detail construction of an $\tilde{O}(\sqrt{n})$ -space and polynomial-time algorithm for computing an $O(\sqrt{n})$ -size separator for a given planar graph with n vertices [10]. We use the separator algorithm obtained in Chapter 3 as a subroutine.

4.1 Outline of the algorithm

We recall two separator algorithms that will be used to compute a separator in our algorithm.

The first one is the standard algorithm of Lipton and Tarjan [45].

Proposition 4.1 (Lipton-Tarjan separator algorithm). For any planar graph with n vertices, there exists a $1/3$ -separator of size $2\sqrt{2n}$. Furthermore, there exists a polynomial-time algorithm that finds one of such separators by a breadth-first search. The algorithm can be modified so that it uses a given BFS (breadth-first search) tree of G and it runs in polynomial-time w.r.t. n and in $\tilde{O}(k)$ -space w.r.t. the depth k of the BFS tree.

The second one is based on an algorithm proposed by Miller in [47]. In that paper, Miller showed a linear-time algorithm computing a cycle weighted $1/3$ -separator for a given 2-connected plane and weighted graph.

Proposition 4.2. Consider any plane and 2-connected weighted graph with n vertices such that each of its faces has weight $\leq 2/3$ and size $\leq d$. Then there exists a cycle weighted $1/3$ -separator of size $2\sqrt{2\lceil d/2 \rceil n}$. Furthermore, there exists an algorithm that finds one of such separators in $O(n)$ -time (under the unit cost).

As remarked in Section 1.2 of [33], for any graph with n vertices and n_{face} faces of size $\leq d$, we have $n \leq 2dn_{\text{face}}$, which implies that $2\sqrt{2}\lfloor d/2 \rfloor n \leq \sqrt{8} \cdot d\sqrt{n_{\text{face}}}$. Also looking into the proof of this proposition, we see that the above algorithm operates only on faces of a given weighted graph if its complete face information is given, from which we can expect that it can be modified to run in $\tilde{O}(n_{\text{face}})$ -space. This has been indeed proved formally in the previous chapter([11]) as follows; in this chapter, we will refer this algorithm as *Miller's algorithm*.

Proposition 4.3 (Modification of Miller's separator algorithm [11]). Consider any plane and 2-connected weighted graph with n vertices and n_{face} faces such that each of its faces has weight $\leq 2/3$ and size $\leq d$. Then there exists a weighted cycle 1/3-separator of size $O(d\sqrt{n_{\text{face}}})$. Furthermore, there exists an algorithm that finds one of such separators in $\tilde{O}(n_{\text{face}})$ -space if a combinatorial planar embedding and a complete face information are given.

We also recall the log-space algorithm of Reingold [53] for the *undirected graph reachability problem*. More specifically, we have an $O(\log n)$ -space algorithm that determines, for a given graph G and a pair of vertices u and v of G , whether u is reachable to v in $O(\log n)$ -space. This algorithm will be used several places in the following; we will call it *Reingold's algorithm* for the undirected graph reachability test.

The main theorem of this chapter is the following.

Theorem 4.1 ([10]). For some constant $\alpha > 0$, there exists an algorithm that takes an undirected planar graph with n vertices as input and outputs its α -separator of size $O(\sqrt{n})$ in polynomial-time and $\tilde{O}(\sqrt{n})$ -space.

We explain the outline of our separator algorithm and state necessary claims to show the correctness of the algorithm. These claims will be restated and proved in the subsequent sections.

Following the outline proposed by Gazit and Miller [33] we design a separator algorithm that consists of the following steps:

1. Check the planarity condition etc. of an input graph, make some modifications on the input graph to create our base graphs, and prepare supplementary information.
2. Transform the graph obtained at Step 1 to generate a frame-graph. (This step is separated into four sub-steps. In some cases, a desired separator of the base dual-graph is created during some sub-step; the algorithm of Lipton and Tarjan is used there.)
3. By using Miller's algorithm, compute a desired weighted separator of the frame-graph, which can be naturally transformed to a separator for the base dual-graph. Then compute a desired separator for the original input graph.

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

Below we first explain Step 1, a preprocessing step to create two graphs that we will consider in the subsequent steps. We then introduce the notion of “frame-graph”, and state important transformation sub-steps of Step 2 to obtain a frame-graph. Finally, we give the proof of our separator theorem by showing how to achieve Step 3 for obtaining desired separators for the base dual-graph and the input graph. Following the convention explained in Preliminaries section, we design an algorithm for each step and each sub-step so that it gets data as input from the previous step and computes data as output passed to the next step in $\tilde{O}(\sqrt{n})$ -space and polynomial-time.

Notational Remark. We would like to use symbols G and $\tilde{G}$ to denote the base graphs that we will mainly consider in this chapter, where G is obtained by transforming the input graph and $\tilde{G}$ is its dual-graph. Thus, we will use (if necessary) G_{org} to denote the original input graph. We also use n to denote the number of vertices of G (and $\tilde{n}$ for the number of vertices of $\tilde{G}$). Again we use n_{org} to denote the number of vertices of G_{org} . As we will see, we have $n = \Theta(n_{\text{org}}) = \Theta(\tilde{n})$; hence, there is no difference among n , n_{org} , and $\tilde{n}$ as a complexity parameter, and we will simply use n for discussing the complexity of algorithms.

Step 1 and our base graphs (Corresponding technical section: Section 4.2)

In Step 1 we check the conditions assumed for the input graph G_{org} and prepare two base graphs and some additional information for these graphs. Here Reingold’s algorithm for the reachability test plays a key role for achieving this preprocessing in $O(\log n)$ -space (and hence, in polynomial-time).


First consider a sub-step for checking input graph conditions. For the graph G_{org} given by its adjacency matrix, we check whether it is connected by using Reingold’s algorithm. If G_{org} is not connected, the algorithm considers its largest connected component. If it has less than $30n/31$ vertices, then the algorithm can claim the empty set as a desired separator. Otherwise, the algorithm can work on a subgraph G'_{org} induced by the largest component to obtain its $1/31$ -separator, which is clearly a $1/31$ -separator of the original input graph G_{org} (by adding the remaining connected components of G_{org} to a smaller separated subgraph of G'_{org}). Note that identifying the largest connected component and checking its size can be done in $O(\log n)$ -space by using Reingold’s algorithm.

We then apply the algorithm of Allender and Mahajan [3] for the planarity test. Precisely speaking, they showed an $O(\log n)$ -space algorithm for the planarity test by using the undirected graph reachability test as a subroutine, which can be used as a standard $O(\log n)$ -space algorithm with Reingold’s algorithm. In the same paper, they also gave an $O(\log n)$ -space algorithm (by using Reingold’s algorithm) that produces a combinatorial planar embedding (if a given graph is planar). Thus, we use their algorithms to check whether the input graph is planar and (if it is so) to compute one of its combinatorial planar embeddings.

In the rest of this chapter, we assume that a given input graph G_{org} is connected and planar and that the algorithm is also given one of its combinatorial

planar embeddings. (From some technical reason, we also assume that G_{org} is not a trivial single triangle.)

Next we transform the input graph G_{org} into a well-structured graph G . From the above sub-step, we have the combinatorial planar embedding of G_{org} . Then it is not so hard to compute its triangulation in $O(\log n)$ -space, that is, a combinatorial planar embedding for its triangulated graph (see Section 4.2 for more explanation). Note here that a triangulated graph is 2-connected, and this 2-connectivity will be kept in the following sub-steps. As we will see below it would be easier if we can assume that a graph is three-regular. So we transform the obtained triangulated graph to a three-regular graph by vertex expansion shown in Figure 4.1. Formally, the vertex expansion is defined as follows: Consider any vertex v of a triangulated graph. Since the graph is triangulated (and if it is not a trivial single triangle), we may assume that its degree $\deg(v)$ is at least 3. The *vertex expansion* on v is to replace v with a $\deg(v)$ -size face, that is, a face having a boundary cycle consisting of $\deg(v)$ vertices. Each of the $\deg(v)$ edges incident to v is now connected to each vertex on the face boundary cycle following their order under the assumed combinatorial planar embedding. We expand all vertices of the triangulated G_{org} in this way, and the obtained graph is our first *base graph* G .



(a) An example of vertex expansion. Each vertex of the left graph is expanded to respectively a triangle, a square, and a pentagon in the right graph. (b) An example of a base graph G_{org} (only its subgraph). All faces are either a face obtained by the vertex expansion or a polygon (triangle $\sim$ hexagon) corresponding to a triangle of G_{org} .


Figure 4.1: Vertex expansion

Here we point out some properties that are easy to see from Figure 4.1. That is, (i) G is three-regular, and (ii) all faces of G are either a face obtained by expanding the corresponding vertex of G_{org} or a polygon (triangle $\sim$ hexagon) corresponding to a triangle of G_{org} . The other properties are summarized below.

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

Claim 4.1 (Base graph G ; restated as Lemma 4.3). Hereafter let G denote a *base graph*, a graph obtained by applying the vertex expansion on all vertices of the triangulated input graph G_{org} (under its combinatorial planar embedding). Let n denote the number of vertices of G (whereas we use n_{org} to denote the number of vertices of G_{org}). Then G satisfies the following: (i) it is three-regular, (ii) all faces of G are either a face obtained by expanding the corresponding vertex of G_{org} or a polygon (triangle $\sim$ hexagon) corresponding to a triangle of G_{org} , (iii) $n = \Theta(n_{\text{org}})$, and (iv) G has a planar embedding naturally defined from the one for G_{org} .

Gazit and Miller [33] considered a dual-like graph of an input graph; their separator algorithm was explained by using both an input graph and its dual-like graph. We follow this style, but here we consider the standard dual-graph. (See Figure 4.2 for an example of a dual-graph; it is easy to see from the figure that a dual-graph is triangulated if its base graph is three-regular.)



An example of a three-regular base graph (only its sub-graph) and its dual-graph. The base graph is indicated by black nodes (for vertices) and solid lines (for edges), while its dual-graph is indicated by white nodes (for dual-vertices) and dashed lines (for dual-edges). As shown here a dual of a three-regular graph is triangulated.


Figure 4.2: A three-regular graph and its dual-graph

Definition 4.1 (Base dual-graph $\tilde{G}$). Throughout this chapter, let $\tilde{G} = (\tilde{V}, \tilde{E})$ denote a dual-graph of G (under its planar embedding) where $\tilde{G}$ is the set of faces of G and $\tilde{E}$ is the set of dual-edges, i.e., pairs of dual-vertices whose corresponding faces share at least one edge.

Remark. We identify a dual-vertex of $\tilde{G}$ and its corresponding face of G , and we sometimes say that a dual-vertex $\tilde{v}$ has an edge e of G , meaning that e is one of the boundary edges of the face corresponding $\tilde{v}$.

Notational Remark. We will use a prefix “dual-” to denote a vertex, an edge, a path, and a cycle in $\tilde{G}$ to distinguish them from those in G , and we will attach “~” to variables denoting these objects; for example, a vertex of $\tilde{G}$ is called a *dual-vertex* and it is denoted by, e.g., $\tilde{v}$. (There are some exceptions. In case we add some other prefix such as “boss-”, “frame-”, etc. for referring an object of the dual-graph $\tilde{G}$, we may omit adding “dual-.” Also since subplanes and faces are notions involving a topological interpretation on $\mathbb{R}^2$, we will not use “dual-” to denote them even if they are defined w.r.t. the dual-graph $\tilde{G}$.)

We remark here that the base dual-graph $\tilde{G}$ may not be a simple graph though the base graph G is simple (by assuming the input graph is a simple graph); see Figure 4.3 for an example. We assume any reasonable way to identify non-simple dual-edges.



An example of the case where a non-simple dual-graph is created from a simple base graph. As the previous figure, i.e., Figure 4.2, a base graph is indicated by black nodes and solid lines, while some (not all) of dual-vertices and dual-edges are indicated by white nodes and dashed lines.

Figure 4.3: An example of the creation of a non-simple dual-graph

It is easy to see that the number $\tilde{n}$ of dual-vertices of $\tilde{G}$ is linearly related to n . Furthermore, we have the following relation.

Claim 4.2 (Restated as Lemma 4.4). Assume that n is large enough. If $\tilde{G}$ has a $1/10$ -separator of $\tilde{G}$ of size $O(\sqrt{n})$, then G_{org} has a $1/31$ -separator of size $O(\sqrt{n_{\text{org}}})$.

A frame-graph and four sub-steps of Step 2

A “frame-graph” is a weighted subgraph of the base dual-graph $\tilde{G}$. Roughly speaking, we transform $\tilde{G}$ by removing its dual-vertices, which creates new faces. The weight of these faces are defined proportional to the number of removed dual-vertices. In order to be able to use Miller’s separator algorithm, we require the following conditions to a frame-graph.

Definition 4.2 (Frame-graph). A *frame-graph* is a weighted subgraph $\tilde{H}$ of the base dual-graph $\tilde{G}$ that satisfies the following conditions (we let $k = \sqrt{n}$ and we will fix this usage hereafter):

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

- (F1) $\tilde{H}$ is a weighted subgraph of $\tilde{G}$ induced by some subset of dual-edges of $\tilde{G}$. The weight of each face¹ f of $\tilde{H}$ is proportional² to the number $\tilde{n}_f$ of dual-vertices of $\tilde{G}$ located in the face (which are removed from $\tilde{G}$), and it is less than $1/3$.
- (F2) $\tilde{H}$ is 2-connected.
- (F3) $\tilde{H}$ contains $O(\tilde{n}/k)$ faces.
- (F4) The size of each face of $\tilde{H}$ is $O(\sqrt{k})$; that is, for each face of $\tilde{H}$, its boundary dual-cycle consists of $O(\sqrt{k})$ dual-vertices.

Step 2 of our algorithm is to obtain a frame-graph $\tilde{H}$ from the base dual-graph $\tilde{G}$. Roughly speaking, we identify “frame-cycles”, dual-cycles that define faces of $\tilde{H}$, and we compose $\tilde{H}$ as a collection of such dual-cycles. Step 2 is divided into the following four sub-steps. (Here we only specify the outputs of these sub-steps; some of the technical terms will be explained below. We may assume that the input to these sub-steps is all data prepared as output in Step 1 and the previous sub-steps.)

- 2.1. Identify Voronoi regions and compute their boss-vertices that are used as a root of a BFS dual-tree of each Voronoi region.

Output: The description of the Voronoi regions. That is, a set $\tilde{I}$ of the boss-vertices, and for each $\tilde{b} \in \tilde{I}$, the description of a BFS dual-tree rooted at $\tilde{b}$, and a set of boundary cycles for the corresponding Voronoi region.

- 2.2. Identify ridge edges (if needed), and compute connectors and branch vertices. Then identify “preliminary” frame-cycles³.

Output: The set of branch vertices, connectors, and the description of preliminary frame-cycles.

- 2.3. Preprocess preliminary frame-cycles and compute frame-cycles, from which a frame-graph $\tilde{H}$ is defined. (During this preprocessing step, there are cases where a desired separator of $\tilde{G}$ is obtained and the algorithm terminates by transforming it to a desired $1/31$ -separator of G_{org} . The algorithm of Lipton and Tarjan is used in one of such cases.)

Output: The description of $\tilde{H}$, that is, its combinatorial planar embedding and its complete weighted face information.

- 2.4. Identify floor- and ceiling-cycles, and modify $\tilde{H}$ by adding floor- and ceiling-cycles as new face boundary dual-cycles.

Output: The description of the modified frame-graph $\tilde{H}$, that is, its combinatorial planar embedding and its complete face information.

¹It is clear that $\tilde{H}$ is a plane graph w.r.t. the embedding naturally inherited from $\tilde{G}$. Thus, precisely speaking, we assume this embedding for discussing the faces of $\tilde{H}$.

²It is simply $\tilde{n}_f$ divided by the total number $\tilde{n}^-$ of dual-vertices removed from $\tilde{G}$ for defining $\tilde{H}$. In fact, we will use, for simplicity, $\tilde{n}_f$ instead of $\tilde{n}_f/\tilde{n}^-$ in our algorithm.

³In our technical discussion, we will define these preliminary frame-cycles formally and call them “pre-frame-cycles.”

We explain sub-steps a bit more in detail so that we can state important properties on their outputs. (Note that some notions are used below without a definition or with an only rough and not precise definition; see the corresponding section for their definitions. Though omitted below, we clearly need to show that these sub-steps can be done in $\tilde{O}(\sqrt{n})$ -space and polynomial-time, which will be proved in the subsequent technical sections.)

Step 2.1. (Corresponding technical section: Section 4.3)

We borrow the idea of Gazit and Miller [33] and introduce the notion of Voronoi region. A *Voronoi region* is defined by a set of connected dual-vertices within distance $2k$ from some dual-vertex, which we call a *boss-vertex*. Boss-vertices are selected in a greedy way, and it can be shown that we can cover $\tilde{G}$ by selecting $\tilde{n}/k$ boss-vertices. We use $\tilde{I}$ to denote the set of all selected boss-vertices, and for each boss-vertex $\tilde{b}$, we use $\text{Vr}(\tilde{b})$ to denote a Voronoi region having $\tilde{b}$ as a boss-vertex. Precisely speaking, these Voronoi regions are subplanes dividing the plane where the base plane graph G is drawn. We can show that each Voronoi region is defined by boundary cycles consisting of G 's edges; see Figure 4.4 for an example. The following properties are important.


Claim 4.3 (Voronoi region; follows from Lemmas 4.6 and 4.7). We have $|\tilde{I}| \leq \tilde{n}/k$; that is, there are at most $\tilde{n}/k$ Voronoi regions. Consider any $\tilde{b} \in \tilde{I}$ and its Voronoi region $\text{Vr}(\tilde{b})$. Its boundary is a set of cycles of G . For any dual-vertex $\tilde{v} \in \text{Vr}(\tilde{b})$, there exists a dual-path from $\tilde{b}$ to $\tilde{v}$ of length $\leq 2k$ consisting of only dual-vertices in $\text{Vr}(\tilde{b})$. Thus, $\text{Vr}(\tilde{b})$ has an $O(k)$ -depth BFS dual-tree rooted at $\tilde{b}$.

See Figure 4.4. There are vertices where three Voronoi regions meet, which we call *branch vertices*, and a dual-triangle consisting of three dual-edges incident to a branch vertex is called a *branch-triangle*. As shown in the figure, we can define a dual-cycle in a pair of Voronoi regions sharing a boundary edge that consists of four dual-paths and two dual-edges of two branch-triangles. The dual-cycle following a tour of dual-vertices labeled with $1 \sim 6$ is an example. The four dual-paths are those in the BFS dual-trees of the corresponding Voronoi regions. This dual-cycle is a basis of a frame-cycle; in fact, during this outline explanation, we regard it as a preliminary frame-cycle. From the property of Voronoi regions, we may assume an $O(k)$ -depth BFS dual-tree inside of such a preliminary frame-cycle.

Step 2.2. (Corresponding technical section: Section 4.5)

Dual-cycles like the one in Figure 4.4 can be used as frame-cycles for defining the target frame-graph $\tilde{H}$. There is unfortunately a situation where some Voronoi region has more than one boundary cycle, and in this situation, we may have a more complicated dual-cycle as illustrated in Figure 4.5 whose size cannot be bounded appropriately.

This situation can be resolved by introducing some more branch vertices by which we can divide a large dual-cycle into small ones. Here again we borrow




An example of Voronoi regions and dual-cycles that are bases for frame-cycles. Solid lines are the boundaries of Voronoi regions, and black nodes (which are called “branch vertices”) are vertices of G that are incident to three Voronoi regions. White nodes are dual-vertices (whose corresponding faces are) incident to these black nodes. In particular, three dual-vertices incident to a branch vertex forms a dual-triangle, which we call “branch-triangle.” The boss-vertex of each Voronoi region (which is also a dual-vertex) is indicated by a star. Six dual-vertices labeled with 1 $\sim$ 6 connected with dashed lines that indicate dual-paths (some of which is a dual-edge) form a dual-cycle; this may be regarded as a “frame-cycle.” Note that this dual-cycle contains a part of Voronoi boundary shared by two Voronoi regions that connects two branch vertices.

Figure 4.4: Voronoi regions and an example of “frame-cycle”

the idea from Gazit and Miller to select some of the edges of G as “ridge edges”; intuitively, by adding ridge edges to Voronoi boundary edges, some Voronoi boundary cycles are connected and new branch vertices are created. More formally, let B and R denote respectively the set of Voronoi boundary edges and that of ridge edges. Then we can show that $B \cup R$ forms a connected component in G . We regard a degree three vertex in $B \cup R$ as a *branch vertex* (whereas a branch vertex was previously a degree three vertex in B). By introducing new branch vertices, large dual-cycles are divided into small and “standard” dual-cycles, i.e., preliminary frame-cycles.

A path connecting a pair of branch vertices are called a *connector*. This is a generalization of the shared Voronoi boundary that appears in the frame-cycle of Figure 4.4. In fact, centering each connector, we can define a preliminary frame-cycle. The main task of this sub-step is to obtain all such preliminary frame-cycles of $\tilde{G}$ from connectors. This one-to-one correspondence between connectors and preliminary frame-cycles is important for a later discussion.

Step 2.3. (Corresponding technical section: Section 4.6)



An example of the case where there exists a Voronoi region with multiple Voronoi boundary cycles. Solid lines are Voronoi boundaries. We omit here showing vertices, dual-vertices, and branch-triangles except for boss-vertices. The outmost solid cycle is one of the boundary cycles of a Voronoi region $Vr(\tilde{b})$. This $Vr(\tilde{b})$ contains four Voronoi regions and it has two boundary cycles between them (i.e., two large squares consisting of two subsquares). Dashed lines are branch dual-paths. There are some “standard” dual-cycles that could be used as frame-cycles; for example, two dual-cycles indicated by dashed lines surrounding shadow parts have the same structure as the preliminary frame-cycle we saw in Figure 4.4. On the other hand, the graph has a dual-cycle consisting of twelve dual-paths labeled by numbers (where branch-triangles are omitted). The problem here is that we cannot bound the length of such a dual-cycle; thus, it cannot be used for a face boundary of a frame-graph as it is for satisfying the condition (F4).

Figure 4.5: A Voronoi region with multiple boundary cycles

Some preliminary frame-cycles may not be appropriate for frame-cycles from the following reasons: (i) there may be a preliminary frame-cycle that defines a large subplane having too many dual-vertices of $\tilde{G}$, which creates a face of weight more than $1/3$ in $\tilde{H}$; (ii) there may be a preliminary frame-cycle that splits $\tilde{G}$ into multiple disconnected subplanes, due to which $\tilde{H}$ cannot be 2-connected by removing dual-vertices of $\tilde{G}$ not participating to the frame-cycles. Thus, in this sub-step, we first preprocess the obtained preliminary frame-cycles to modify them to appropriate forms.

During this preprocessing step, we may have a situation where some preliminary frame-cycle itself is a desired separator of $\tilde{G}$, say, $1/10$ -separator of $\tilde{G}$. Or we can apply the algorithm of Lipton-Tarjan to the subgraph that is inside of some preliminary frame-cycle to obtain a desired separator of $\tilde{G}$. Note that the latter one is the case where some preliminary frame-cycle is large enough to cover many dual-vertices of $\tilde{G}$. Since (the inside of) each preliminary frame-

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

cycle is covered by (at most) two Voronoi regions, and each Voronoi region has an $O(k)$ -depth BFS dual-tree, we can show by Proposition 4.1 that the algorithm of Lipton-Tarjan runs in $\tilde{O}(k)$ -space. Once a desired separator of $\tilde{G}$ is obtained, the algorithm can terminate by yielding our target $1/31$ -separator of G_{org} by Claim 4.2 (i.e., by the method stated in its proof).

After this preprocessing step (and a separator is not yet obtained), we now have *frame-cycles* and we can guarantee that each frame-cycle defines a single subplane that does not have so many dual-vertices of $\tilde{G}$. Thus, we can define our frame-graph $\tilde{H}$ as a collection of faces defined⁴ by these frame-cycles. We should note, on the other hand, that the actual step of our algorithm is designed in a slightly different way. This step is designed to yield a frame-graph $\tilde{H}$ as an induced dual-graph consisting of dual-edges of all frame-cycles. Thus, we need to prove the following claim. For its proof the one-to-one correspondence between connectors and (preliminary) frame-cycles plays an important role.

Claim 4.4 (Face of $\tilde{H}$; restated as Lemmas 4.15 and 4.16). Each frame-cycle defines exactly one face in $\tilde{H}$. Furthermore, there is no face in $\tilde{H}$ other than ones defined by either a frame-cycle or a branch-triangle.

Based on this claim, for a complete weighted face information of $\tilde{H}$, the algorithm needs to output lists of frame-cycles, their weights, and their incidence relations. Note that the weight of each frame-cycle is simply the number of dual-vertices of $\tilde{G}$ located its inside. From the preprocessing step, we can guarantee that it is less than $\tilde{n}^-/3$, where $\tilde{n}^-$ is the total number of dual-vertices removed from $\tilde{G}$ to define $\tilde{H}$. Thus, the obtained $\tilde{H}$ satisfies the condition (F1).

We also prove that $\tilde{H}$ satisfies (F2) and (F3).

Claim 4.5 (Restated as Lemmas 4.13). $\tilde{H}$ is 2-connected.

Claim 4.6 (Number of connectors and frame-cycles; restated as Lemma 4.17). The numbers of connectors is bounded by $O(\tilde{n}/k)$. Hence, the number of faces of $\tilde{H}$, that is, frame-cycles and branch-triangles is bounded by $O(\tilde{n}/k)$.


Step 2.4. (Corresponding technical section: Section 4.7)

Let us see again the frame-cycle illustrated in Figure 4.4. It consists of four BFS dual-paths and two branch-triangle edges. While we can bound the length of each BFS dual-path by $O(k)$, this bound is not enough for satisfying the condition (F4) that requires an $O(\sqrt{k})$ bound for the length of frame-cycles. We again borrow the idea of Gazit and Miller to introduce the notion of “floor-cycle” and “ceiling-cycle” to overcome this problem.

Floor- and ceiling-cycles are defined by using dual-vertices of the same level. A dual-vertex is given a *level* that is the distance from its boss-vertex (which is a temporal definition just for an explanation here). We consider connected components of dual-vertices of the same level; in particular, we are interested in

⁴Precisely speaking, faces are defined also by branch-triangles, but in the following explanation, we sometimes omit mentioning branch-triangles for simplicity.

connected components of size $\leq \sqrt{k}$. A *neck* of level ℓ is a connected component of level ℓ dual-vertices whose size is bounded by $\sqrt{k}$. In such a neck we can find (at least) two dual-cycles, one facing dual-vertices of level $\ell + 1$ and another facing dual-vertices of level $\ell - 1$. A *floor-cycle* is a dual-cycle of the former type that has enough number of dual-vertices (i.e., $2\tilde{n}/3$ dual-vertices) in its “outside”, the side the dual-cycle where level $\ell + 1$ dual-vertices are located. Similarly, a *ceiling-cycle* is a dual-cycle of the latter type that has more than $2\tilde{n}/3$ dual-vertices in its “outside”, the side the dual-cycle where level $\ell - 1$ dual-vertices are located. If multiple nested such dual-cycles exist, then we take the one with the largest level (resp., the smallest level) as the floor-cycle (resp., the ceiling-cycle).



Cycles indicated by dashed lines surrounding shadow parts are floor- or ceiling-cycles. By the modification of $\tilde{H}$ in sub-step 2.4, their “insides” (i.e., these shadow parts) are removed so that floor- and ceiling-cycles become are new face boundary dual-cycles; then the length of each dual-path of a frame-cycle from a boss-vertex to the next face boundary dual-cycle is reduced to $O(\sqrt{k})$.

Figure 4.6: Floor- and ceiling-cycles

Consider any frame-cycle $\tilde{c}$ of $\tilde{H}$. We can identify one floor-cycle containing its boss-vertex “inside.” On the other hand, a ceiling-cycle (if it exists) is located around its branch-triangle containing some dual-vertices of the branch-triangle “inside” (Figure 4.6). By removing all dual-vertices inside of these floor- and ceiling-cycles (thereby creating new faces), we can reduce the length of each dual-path of the frame-cycle $\tilde{c}$ remained “outside” of these new faces to $O(\sqrt{k})$.

Claim 4.7 (Reducing dual-path length of frame-cycles; restated as Lemma 4.26). Consider any frame-cycle $\tilde{c}$, and let $\tilde{p}$ be any dual-path between its boss-vertex and the next dual-vertex of a branch-triangle used in $\tilde{c}$. There exist one floor-cycle and at most one ceiling-cycle crossing $\tilde{p}$; let $\tilde{p}'$ be the part of $\tilde{p}$ located outside of these cycles. Then the length of $\tilde{p}'$ is $O(\sqrt{k})$.

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

In sub-step 2.4, for a given $\tilde{G}$ (and $\tilde{H}$) we identify floor- and ceiling-cycles and modify the frame-graph $\tilde{H}$ by introducing new faces by removing all dual-vertices of $\tilde{H}$ from the inside of these floor- and ceiling-cycles. We refer the obtained dual-graph $\tilde{H}'$ as a *modified frame-graph*. The weight of its faces is defined in the same way as $\tilde{H}$; that is, the weight of each face is defined by the number of removed dual-vertices of $\tilde{G}$ in the face. Then from the property of the original frame-graph $\tilde{H}$ and by our definition of floor- and ceiling-cycles, we can show that the modified frame-graph $\tilde{H}'$ satisfies the condition (F1) $\sim$ (F3). Furthermore, from the above claim (and since the length of each floor- and ceiling-cycle is $O(\sqrt{k})$ by definition), we have an $O(\sqrt{k})$ bound for the length of each face boundary dual-cycle. Therefore, the obtained modified frame-graph $\tilde{H}'$ satisfies all the required conditions (F1) $\sim$ (F4). For simplifying our notation, we will use from now on $\tilde{H}$ to denote this modified frame-graph.

Step 3 and the proof of our separator theorem

Now that a modified frame-graph $\tilde{H}$ satisfying the conditions (F1) $\sim$ (F4) is obtained. Then the last step of our algorithm is to apply Miller's algorithm to obtain a weighted separator for $\tilde{H}$ and compute our desired separators for $\tilde{G}$ and G_{org} .

Let us first prepare some parameters needed for the analysis. We use $\tilde{n}_{\tilde{H}}$ and $\tilde{n}_{\tilde{H}_{\text{face}}}$ to denote respectively the number of dual-vertices and faces of $\tilde{H}$, and let d be its max. face size. Then from the conditions (F3) and (F4), we have $\tilde{n}_{\tilde{H}_{\text{face}}} = O(\tilde{n}/k) = O(n^{1/2})$ and $d = O(\sqrt{k}) = O(n^{1/4})$ respectively; hence, $\tilde{n}_{\tilde{H}} \leq d \times \tilde{n}_{\tilde{H}_{\text{face}}} = O(n^{3/4})$. Thus, the number of dual-vertices removed from $\tilde{G}$ to obtain $\tilde{H}$, which we denote by $\tilde{n}^-$, is $\tilde{n} - \tilde{n}_{\tilde{H}} = \Omega(n) - O(n^{3/4})$; we assume that n is large enough so that $\tilde{n}^- \geq 3\tilde{n}/4$ holds.

We consider the application of Miller's algorithm to $\tilde{H}$. This creates a cycle $1/3$ -separator of $\tilde{H}$. From the condition (F4) and Proposition 4.3, it follows that the size of this separator is $O(d\sqrt{\tilde{n}_{\tilde{H}}}) = O(\sqrt{n})$, and that Miller's algorithm can be executed in $\tilde{O}(\tilde{n}_{\tilde{H}_{\text{face}}})$ -space ($= \tilde{O}(\sqrt{n})$ -space). Then since the obtained separator is a dual-cycle, it can be used as a $1/4$ -separator of $\tilde{G}$ as the following lemma claims. (Cf. A separator of $\tilde{H}$ is not necessarily a separator of $\tilde{G}$ in general.) Thus, in Step 3, we first apply Miller's algorithm to the frame graph $\tilde{H}$ to obtain a weighted $1/3$ -separator (that is indeed a $1/4$ -separator of $\tilde{G}$) in $\tilde{O}(\sqrt{n})$ -space, and then for the obtained separator for $\tilde{G}$, we apply the method stated in the proof of Claim 4.2 to compute a desired $1/31$ -separator of G_{org} in $\tilde{O}(1)$ -space.

Lemma 4.2. Any weighted dual-cycle $1/3$ -separator of a modified frame-graph $\tilde{H}$ is a dual-cycle $1/4$ -separator of $\tilde{G}$.

Proof. Let $\tilde{C}$ denote a weighted dual-cycle $1/3$ -separator of $\tilde{H}$. Recall that $\tilde{H}$ is a dual-edge induced subgraph of $\tilde{G}$. Thus, any dual-cycle of $\tilde{H}$ is also a dual-cycle of $\tilde{G}$; thus, the dual-cycle separator $\tilde{C}$ of $\tilde{H}$ is a dual-cycle and

hence a separator of $\tilde{G}$. Recall that the weight of each face of $\tilde{H}$ is the number of dual-vertices of $\tilde{G}$ that are located in the face and that are removed when defining $\tilde{H}$ normalized by the total number $\tilde{n}^-$ of removed dual-vertices. Thus, since $\tilde{C}$ is a weighted $1/3$ -separator, two subplanes separated by $\tilde{C}$ have weight $\geq 1/3$; hence, the numbers of dual-vertices of these subplanes in $\tilde{G}$ are at least $\tilde{n}^-/3 \geq \tilde{n}/4$ (as we checked before this lemma). Therefore, $\tilde{C}$ is a $1/4$ -separator of $\tilde{G}$. $\square$

4.2 Base Graphs G and $\tilde{G}$

We explain Step 1 of the algorithm, a preprocessing step for computing the base graphs G and $\tilde{G}$ and supplementary information. Here we prove the key properties of these base graphs claimed in Outline section. We also give some basic algorithms for achieving this step.

We begin by defining the notion of “triangulation” formally.

Definition 4.3 (Triangulation). A plane graph is *triangulated* (w.r.t. its planar embedding) if addition of any edge results in a nonplanar graph. For a plane graph, its *triangulation* means to add edges to the plane graph until it gets triangulated w.r.t. this planar embedding.

As explained in Outline section, we may assume, after the preprocessing sub-step of Step 1, that the input graph G_{org} is planar and connected and that the algorithm can use one of its combinatorial embeddings. Then in the rest of Step 1, the algorithm conducts the following tasks: (i) triangulating the input graph, (ii) applying the vertex expansion to this triangulated graph to obtain the base graph G , and (iii) computing its dual as the base dual-graph $\tilde{G}$.

We explain the key properties of the base graphs. First one is for the base graph G .

Lemma 4.3 (Restatement of Claim 4.1). The following holds for G and G_{org} : (i) G is three-regular, (ii) all faces of G are either a face obtained by expanding the corresponding vertex of G_{org} or a polygon (triangle $\sim$ hexagon) corresponding to a triangle of G_{org} , (iii) $n = \Theta(n_{\text{org}})$, and (iv) it has a planar embedding naturally defined from the one for G_{org} .

Proof. Properties (i) and (ii) are clear from the definition; see Figure 4.1 and the explanation there. Thus, we prove (iii) and (iv) here. Recall that the vertex expansion on v is to replace v with a $\deg_{\text{org}}(v)$ -size face, by which we introduce $\deg_{\text{org}}(v)$ new vertices replacing v . Hence, by applying the vertex expansion to all vertices of $G_{\text{org}} = (V_{\text{org}}, E_{\text{org}})$, we create $\sum_{v \in V_{\text{org}}} \deg_{\text{org}}(v) = 2|E_{\text{org}}|$ new vertices, which is at least n_{org} and at most $6n_{\text{org}}$ since $|E_{\text{org}}| \leq 3n_{\text{org}}$ due to the planarity of G_{org} and the Euler’s formula. This proves (iii). For showing (iv), consider any vertex $v \in V_{\text{org}}$. Let $e_0, \dots, e_{d-1}$ be edges incident to v indexed in the order of the given combinatorial embedding of G_{org} . We may consider that the vertex expansion on v introduces a new vertex v'_i corresponding to

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

each edge e_i , where v'_i has an edge e'_1 corresponding to e_i and two edges e'_0 and e'_2 connecting to $v'_{(i-1) \bmod d}$ and $v'_{(i+1) \bmod d}$. Then it is easy to see that the order e'_0, e'_1, e'_2 is a combinatorial embedding for v'_i under the embedding of G naturally inherited from G_{org} . $\square$

Next consider the base dual-graph $\tilde{G}$. Recall that it is defined a dual-graph of G (Definition 4.1) and it is triangulated thanks to the three-regularity of G (Figure 4.2). Here we show a way to get a separator for G_{org} from a separator for $\tilde{G}$.

Lemma 4.4 (Restatement of Claim 4.2). *If $\tilde{G}$ has a $1/10$ -separator of $\tilde{G}$ of size $O(\sqrt{n})$, then G_{org} has a $1/31$ -separator of size $O(\sqrt{n_{\text{org}}})$.*

Proof. First, we introduce a way to define a set of vertices of G (resp., G_{org}) that corresponds to a given dual vertex of $\tilde{G}$. Consider any vertex $\tilde{v}$ of $\tilde{G}$. Note that $\tilde{v}$ is a face in G ; we define $\text{inv}(\tilde{v})$ by a set of vertices of G that are incident to this face. Then define $\text{inv}_{\text{org}}(\tilde{v})$ by a set of vertices whose vertex expansion intersects with $\text{inv}(\tilde{v})$. For $\text{inv}_{\text{org}}(\tilde{v})$ the following more direct interpretation would be helpful. Note that $\tilde{v}$ is a face in G ; on the other hand, the above lemma claims (as property (ii)) that every face of G is either a face obtained by expanding a vertex v of G_{org} or a polygon (triangle $\sim$ hexagon) face that corresponds to a triangle face $\{v_1, v_2, v_3\}$ of G_{org} . Hence, we have $\text{inv}_{\text{org}}(\tilde{v}) = \{v\}$ for the former case and $\text{inv}_{\text{org}}(\tilde{v}) = \{v_1, v_2, v_3\}$ for the latter case. For any set $\tilde{W}$ of dual vertices of $\tilde{G}$, we define $\text{inv}(\tilde{W})$ and $\text{inv}_{\text{org}}(\tilde{W})$ by $\text{inv}(\tilde{W}) = \bigcup_{\tilde{v} \in \tilde{W}} \text{inv}(\tilde{v})$ and $\text{inv}_{\text{org}}(\tilde{W}) = \bigcup_{\tilde{v} \in \tilde{W}} \text{inv}_{\text{org}}(\tilde{v})$, respectively.

Now for a given separator $\tilde{S}$, we show that $S := \text{inv}_{\text{org}}(\tilde{S})$ is a separator of G_{org} with a desired property. Note first that we have $|S| \leq 3|\tilde{S}|$ from the above interpretation of inv_{org} . Thus, from the assumption, we have $|S| = O(\sqrt{n})$ ($= O(\sqrt{n_{\text{org}}})$). Below we use $\tilde{H}_1$ and $\tilde{H}_2$ to denote sets of vertices of $\tilde{G}$ separated by $\tilde{S}$. Also for each $i \in \{1, 2\}$, we let $H_i = \text{inv}_{\text{org}}(\tilde{H}_i) \setminus S$. These symbols are used also to denote their corresponding induced graphs.

We show that S is indeed a separator of G separating H_1 and H_2 . For this we consider $S' := \text{inv}(\tilde{S})$ and $H'_i := \text{inv}(\tilde{H}_i) \setminus S'$ for each $i \in \{1, 2\}$, and show that H'_1 and H'_2 are separated by S' . Suppose otherwise, and assume that some $v'_1 \in H'_1$ and $v'_2 \in H'_2$ are connected after removing S' . Then there should be a path p' connecting v'_1 and v'_2 that has no vertex in S' . Consider a set P' of faces of G that share an edge with p' . Then a set of vertices of $\tilde{G}$ corresponding the faces of P' forms a connected component in $\tilde{G}$. This is due to the fact that any two dual-vertices of $\tilde{G}$ corresponding faces of G incident to some common vertex of G must have an edge between them, thanks to the three-regularity of G . Also it is clear that P' has no dual-vertex in $\tilde{S}$, and that there are two dual-vertices whose corresponding faces are incident to v'_1 and v'_2 that are in $\tilde{H}_1$ and $\tilde{H}_2$ respectively. This contradicts that $\tilde{S}$ separates $\tilde{H}_1$ and $\tilde{H}_2$.

Next we show that S is a $1/31$ -separator. Without losing generality we may assume that $|H_1| \leq |H_2|$, and here we show that $|H_1| \geq n/31$ for sufficiently

large n .

Consider first vertices of $\tilde{H}_1$ and $\tilde{G}$. As explained above, each vertex of $\tilde{G}$ corresponds to either a vertex of G_{org} or a triangle face of G_{org} ; we call a vertex of $\tilde{G}$ a *v-vertex* for the former case and a *f-vertex* for the latter case. The converse relation also holds; that is, each vertex of G_{org} and each triangle face of G_{org} corresponds to some vertex of $\tilde{G}$. Thus, we have $\tilde{n} = n_{\text{org}} + F$, where $\tilde{n}$ is the number vertices of $\tilde{G}$ and F is the number of faces of G_{org} . Let V_1 and F_1 denote respectively the number of v-vertices and f-vertices of $\tilde{H}_1$. Then we have from our assumption that

$$V_1 + F_1 \geq \tilde{n}/10 \geq n_{\text{org}}/10. \quad (4.1)$$

Now consider a graph H_1^+ induced by $\text{inv}_{\text{org}}(\tilde{H}_1)$. (Recall that $H_1 = \text{inv}_{\text{org}}(\tilde{H}_1) \setminus S$.) Let n_1^+ , e_1^+ , and f_1^+ denote the number of vertices, edges, and faces of H_1^+ respectively. We let $x = |H_1^+ \cap S|$. Then we have $n_1^+ = V_1 + x$, $f_1^+ \leq F_1 + x$ and $2e_1^+ \geq 3F_1$. The two inequalities hold because H_1^+ might have faces that are f-vertices of $\tilde{S}$ and the number of them is at most x . Then apply the Euler's formula for plane graphs to H_1^+ , we have $n_1^+ = e_1^+ - f_1^+ + 2$, which derives $V_1 - 0.5F_1 \geq -2x + 2$. Hence, by using (4.1) we have

$$3V_1 \geq n_{\text{org}}/10 - 4x + 4 \geq n_{\text{org}}/10 - 4x.$$

Recall that $H_1 = H_1^+ \setminus S$. Thus, we have

$$|H_1| \geq n_1^+ - |S| = V_1 + x - |S| \geq n_{\text{org}}/30 - x/3 - |S|.$$

Then the desired bound holds for sufficiently large n_{org} since $x \leq |S|$ and $|S| = O(\sqrt{n_{\text{org}}})$. $\square$


Algorithms for Step 1

We discuss algorithms for achieving Step 1. As explained in Outline section, thanks to the previous work of Reingold [53] and Allender and Mahajan [3], we have $O(\log n)$ -space algorithms for checking the connectivity and planarity and for computing one of the combinatorial embeddings of the input graph. Thus, our remaining algorithmic tasks are (i) triangulating the input graph G_{org} , (ii) computing a three-regular graph (i.e., the base graph G) by applying the vertex-expansion to the triangulated graph, and (iii) computing the dual-base graph $\tilde{G}$. (By “computing a graph” we formally mean to compute its embedding and its complete face information.) Here we explain how to achieve each task in $\tilde{O}(1)$ -space.

First consider the triangulation. The task of triangulation is essentially to introduce edges to separate non-triangle faces into triangles. Hence, for the triangulation task, we first need to identify each face, i.e., each traverse of G_{org} . Consider any directed edge (u, v) for any undirected edge $\{u, v\}$ of G_{org} , and consider the task of computing the left-traverse of G_{org} from (u, v) . As

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

we remarked in Preliminary section (Lemma 2.1), no same directed edge is identified during the left-traverse process (until coming back to (u, v)). Thus, for conducting the left-traverse process, we only need to remember (u, v) besides the directed edge currently identified in the process; hence, only some constant number of variables are needed for computing each left-traverse. For an obtained left-traverse $t = (v_1, \dots, v_m)$, the triangulation of the face corresponding to the left-traverse is easy. Since the face corresponding t is the subspace located from (v_m, v_1) to (v_1, v_2) clockwise, we simply add edges $\{v_1, v_3\}, \dots, \{v_1, v_{m-1}\}$ for triangulating this face. For the combinatorial embedding π_{v_1} of vertices adjacent to v_1 , we need to insert $v_3, \dots, v_{m-1}$ between v_m and v_1 in the reverse order. Note that multiple edges may be created, which should be treated as different edges; see Figure 4.7.



An example of triangulating a face corresponding to a left-traverse. The original plane graph is a path whose edges are indicated by solid lines, while newly added edges are indicated by dotted lines. Here we consider a left-traverse $(1, 2, 3, 4, 5, 4, 3, 2)$, which defines a single face in the original graph. This face is divided into triangle faces by adding edges $\{1, 3\}$, $\{1, 4\}$, $\{1, 5\}$, $\{1, 4'\}$, and $\{1, 3'\}$. The combinatorial embedding π_1 for $N_G(1)$ is modified from (2) to $(2, 3', 4', 5, 4, 3)$. We remark here that the graph becomes a simple graph again after the vertex expansion.

Figure 4.7: An example of triangulating a face

Algorithmically, it is easier if we consider the triangulation by two sub-steps. First enumerate all left-traverses of G_{org} . Here for the second sub-step, we design the algorithm for this sub-step so that it does output a traverse $t = (v_1, \dots)$ only if v_1 has the smallest index among all vertices in t . Then in the second sub-step, the algorithm modifies the combinatorial embedding π_v for each vertex v of G_{org} . For this, the algorithm searches through all left-traverses produced in the first sub-step for left-traverses in which v appears. In the case where v appears in a left-traverse t as the first vertex (say, $t = (v, v_2, \dots, v_m)$), a sequence $v_{m-1}, \dots, v_3$ is inserted between v_m and v_2 in π_v . On the other hand, in the case where v appears somewhere middle of t (say, $t = (v_1, \dots, v_i, v, v_{i+2}, \dots)$), v_1 is inserted between v_i and v_{i+2} in π_v . Note that each left-traverse causes the insertion in a different part of π_v . The algorithm produces the updated π_v 's for all $v \in G_{\text{org}}$, and then update $E(G_{\text{org}})$ by adding the introduced edges.

Now for the tasks (ii) and (iii), the algorithm first enumerates all triangle face boundary cycles of the triangulated G_{org} , which in fact can be done easily when introducing edges during the above triangulation. Once all triangle face boundary cycles are obtained, the rest of tasks (ii) and (iii) can be done easily

by using some constant number of variables because local information is enough for the vertex expansion for obtaining the base graph G and also for computing the dual-graph $\tilde{G}$. Note that the computation of face boundary dual-cycles of $\tilde{G}$ is nothing but that of left-traverses of $\tilde{G}$.

4.3 Voronoi Region

In this section, we formally define the notion of Voronoi region and explain the properties of Voronoi regions. We then explain Step 2.1 of our algorithm for computing Voronoi regions.

We start with introducing the notion of “region”, an important notion throughout this chapter.

Definition 4.4 (Region and boundary). A *region* is a set R of dual-vertices of $\tilde{G}$ such that any two dual-vertices of R is connected by some dual-path consisting of only dual-vertices of R . In other words, R is a region if $\tilde{G}[R]$, the subgraph of $\tilde{G}$ induced by R , is connected. We also consider region R as a set of faces of G , and a set of edges defined by

$$\{ e \in E \mid \text{there exists exactly one face in } R \text{ that contains } e \}$$

is the boundary of R .

Remark. Though a region is defined as a set of dual-vertices of $\tilde{G}$, (formally speaking) it should be regarded as a subplane defined in G by the corresponding set of faces of G . Thus, we will use symbols without $\sim$ for regions.

We recall the following important fact from [47] that is provable for three-regular graphs.


Proposition 4.4. Any region of a three-regular graph (e.g., our base graph G) has a boundary consisting of simple cycles.

We introduce the notion of “Voronoi region”, a specific family of regions. For any dual-vertices $\tilde{u}$ and $\tilde{v}$, by $\text{dist}(\tilde{u}, \tilde{v})$ we mean the distance between $\tilde{u}$ and $\tilde{v}$, that is, the length of a shortest dual-path between $\tilde{u}$ and $\tilde{v}$. Consider any dual-vertex $\tilde{v}$. Technically, it would be easier if we can identify the “nearest” dual-vertex to $\tilde{v}$; for this purpose, we introduce a total order (denoted by $<_{\tilde{v}}$) in $\tilde{V}$ based on the distance from $\tilde{v}$. For any dual-vertices $\tilde{u}$ and $\tilde{w}$, we say that $\tilde{u}$ is *nearer* to $\tilde{v}$ than $\tilde{w}$ (written as $\tilde{u} <_{\tilde{v}} \tilde{w}$) if we have either

- $\text{dist}(\tilde{u}, \tilde{v}) < \text{dist}(\tilde{w}, \tilde{v})$, or
- $\text{dist}(\tilde{u}, \tilde{v}) = \text{dist}(\tilde{w}, \tilde{v})$, and $\tilde{u}$ has a smaller index⁵ than $\tilde{w}$.

For any dual-vertex $\tilde{v}$ and any integer $d \geq 0$, let $L(\tilde{v}, d)$ denote the set of dual-vertices whose distance from $\tilde{v}$ is d . We introduce two types of neighborhoods.

⁵We assume that dual-vertices are indexed in a certain way.



An example of a region and its boundary. Here black nodes and solid lines are those in G ; each face is regarded as a dual-vertex of $\tilde{G}$. The shaded part is a region. The boundary of this region consists of two cycles. In general Miller showed that in three-regular graph, the boundary of any region is a set of simple cycles.

Figure 4.8: A region and its boundary

(*Remark.* Since $L(\tilde{v}, d)$'s and the neighborhoods defined below are also regions, we do not use $\sim$ for their symbols. Recall that k is the parameter that is set $\sqrt{n}$ throughout this chapter.)

Definition 4.5 (k -neighborhood). For any dual-vertex $\tilde{v}$, let $d_{\text{nb}}(\tilde{v})$ be the largest integer d such that $|\cup_{0 \leq i \leq d} L(\tilde{v}, i)| < k$ holds. Then we define $N_k(\tilde{v})$ and $N_k^+(\tilde{v})$ by

$$N_k(\tilde{v}) = \bigcup_{0 \leq i \leq d_0} L(\tilde{v}, i), \text{ and } N_k^+(\tilde{v}) = \bigcup_{0 \leq i \leq d_0+1} L(\tilde{v}, i),$$

where $d_0 = d_{\text{nb}}(\tilde{v})$. We call $N_k(\tilde{v})$ and $N_k^+(\tilde{v})$ a k -neighborhood and a k -neighborhood⁺ respectively.

Remark. By $N_k(\tilde{v})$ and $N_k^+(\tilde{v})$ we also mean subgraphs of $\tilde{G}$ induced by sets $N_k(\tilde{v})$ and $N_k^+(\tilde{v})$.

The following fact is immediate from the above definition.

Lemma 4.5. For any dual-vertex $\tilde{v}$, we have the following: (i) $N_k(\tilde{v}) \subseteq N_k^+(\tilde{v})$, (ii) both $N_k(\tilde{v})$ and $N_k^+(\tilde{v})$ are regions, (iii) $|N_k(\tilde{v})| < k$ and $|N_k^+(\tilde{v})| \geq k$, and (iv) the distance between $\tilde{v}$ and any vertex in $N_k^+(\tilde{v})$ is at most k .

We select representative k -neighborhoods, which will be used for defining Voronoi regions.

Definition 4.6. A subset $\tilde{I}$ of $\tilde{V}$ is a *k-maximal independent set* if it satisfies following conditions:

- For any two dual-vertices $\tilde{b}_1, \tilde{b}_2 \in \tilde{I}$, $N_k^+(\tilde{b}_1) \cap N_k^+(\tilde{b}_2) = \emptyset$.
- For any dual-vertex $\tilde{v} \notin \tilde{I}$, there exists a dual-vertex $\tilde{b} \in \tilde{I}$ such that $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) \neq \emptyset$.

The following properties are immediate from the definition.

Lemma 4.6. Consider any *k*-maximal independent set $\tilde{I}$ of $\tilde{G}$. Then we have (i) $|\tilde{I}| \leq \tilde{n}/k$, and (ii) for any dual-vertex $\tilde{v} \in \tilde{V}$, we have some dual-vertex $\tilde{b} \in \tilde{I}$ and some dual-vertex $\tilde{u}$ such that $\tilde{u} \in N_k^+(\tilde{v}) \cap N_k^+(\tilde{b})$, and we have $\text{dist}(\tilde{v}, \tilde{b}) \leq 2k$.

Notational Remark. (The set $\tilde{I}$)

Though we explain our algorithm design for Voronoi regions later, we briefly mention that a standard greedy method can be used to design an $\tilde{O}(\sqrt{n})$ -space and polynomial-time algorithm that computes a *k*-maximal independent set. From now on, let us use $\tilde{I}$ to denote the *k*-maximal independent set obtained by this algorithm for the base dual-graph $\tilde{G}$.

Now we define Voronoi regions w.r.t. the above defined $\tilde{I}$. Intuitively, for each $\tilde{b} \in \tilde{I}$, we would like to define the “Voronoi region” $\text{Vr}(\tilde{b})$ of $\tilde{b}$ by a set of dual-vertices for which $\tilde{b}$ is closest among all dual-vertices in $\tilde{I}$. But such a dual-vertex in $\tilde{I}$ that is closest to $\tilde{v}$ may not be computable within our desired space bound. Thus, we use “closest *k*-neighborhood” to determine our Voronoi regions, by which we can identify Voronoi regions in $O(k + \tilde{n}/k)$ -space.

Definition 4.7 (Voronoi region and boss-vertex). For any set W of dual-vertices, $\text{nrst}_{\tilde{v}}(W)$ denotes a dual-vertex $\tilde{u}$ in W such that $\tilde{u} <_{\tilde{v}} \tilde{w}$ of all $\tilde{w} \in W \setminus \{\tilde{u}\}$. For any sets W and W' of dual-vertices, we write $W <_{\tilde{v}} W'$ if $\text{nrst}_{\tilde{v}}(W) <_{\tilde{v}} \text{nrst}_{\tilde{v}}(W')$. For any dual-vertex $\tilde{v}$, the boss-vertex of $\tilde{v}$ (denoted by $\text{boss}(\tilde{v})$) is a dual-vertex $\tilde{b} \in \tilde{I}$ such that $N_k^+(\tilde{b}) <_{\tilde{v}} N_k^+(\tilde{b}')$ for any $\tilde{b}' \in \tilde{I} \setminus \{\tilde{b}\}$. The Voronoi region of $\tilde{b}$ (denoted by $\text{Vr}(\tilde{b})$) is a set of dual-vertices $\tilde{v}$ such that $\text{boss}(\tilde{v}) = \tilde{b}$.

Remark. In a nutshell, $\tilde{b}$ is the boss-vertex of $\tilde{v}$ if and only if its *k*-neighborhood is nearest to $\tilde{v}$ among all dual-vertices in $\tilde{I}$. It may be the case that there is some other dual-vertex in $\tilde{I}$ that is nearer to $\tilde{v}$.

By definition, for every dual-vertex $\tilde{v} \in \tilde{V}$, there exists a unique Voronoi region $\text{Vr}(\tilde{b})$ that contains $\tilde{v}$. Furthermore, for any Voronoi region $\text{Vr}(\tilde{b})$, we show below that it is indeed a region and that every dual-vertex in $\text{Vr}(\tilde{b})$ is reachable by a dual-path in $\text{Vr}(\tilde{b})$ of length $\leq 4k$.

Lemma 4.7. For any dual-vertex $\tilde{b} \in \tilde{I}$ and any dual-vertex $\tilde{v} \in \text{Vr}(\tilde{b})$, there exists a dual-path from $\tilde{b}$ to $\tilde{v}$ of length $\leq 2k$ consisting of only dual-vertices in $\text{Vr}(\tilde{b})$. (Hence, any pair of dual-vertices in $\text{Vr}(\tilde{b})$ is connected by a dual-path in $\text{Vr}(\tilde{b})$ of length $\leq 4k$.)

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

Proof. Consider any dual-vertex $\tilde{v} \in \text{Vr}(\tilde{b})$. By Lemma 4.6 (ii), there exists some $\tilde{b}' \in \tilde{I}$ such that $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}') \neq \emptyset$ holds. Clearly, $\tilde{b}$ also satisfies $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) \neq \emptyset$ because otherwise $\tilde{b}$ cannot be the boss-vertex of $\tilde{v}$. Thus, there is a dual-path from $\tilde{v}$ to $\tilde{b}$ of length $\leq 2k$.

For the lemma, it suffices to show that there exists a dual-path of length $\leq 2k$ from $\tilde{v}$ to $\tilde{b}$ consisting of only vertices in $\text{Vr}(\tilde{b})$. For this, consider $\tilde{u} := \text{nrst}_{\tilde{v}}(N_k^+(\tilde{b}))$, i.e., the dual-vertex in $N_k^+(\tilde{b})$ nearest to $\tilde{v}$, which in fact is the dual-vertex in the whole $\cup_{\tilde{b}' \in \tilde{I}} N_k^+(\tilde{b}')$ nearest to $\tilde{v}$ witnessing that $\tilde{v} \in \text{Vr}(\tilde{b})$. Clearly, any shortest dual-path from $\tilde{u}$ to $\tilde{b}$ belongs to $N_k^+(\tilde{b})$ and its length is at most k . (It is easy to see that $N_k^+(\tilde{b}) \subseteq \text{Vr}(\tilde{b})$.) Consider any shortest dual-path $\tilde{v}$ to $\tilde{u}$. Again it is clear that its length is at most k . We show that all dual-vertices on the dual-path belong to $\text{Vr}(\tilde{b})$. Suppose otherwise; that is, there exists some dual-vertex $\tilde{v}'$ on the dual-path that belongs to the other $\text{Vr}(\tilde{b}')$. Then for this $\tilde{v}'$, there exists some $\tilde{u}' \in \text{Vr}(\tilde{b}')$ that is nearer to $\tilde{v}'$ than $\tilde{u}$. But since $\tilde{v}'$ is on one of the shortest dual-paths from $\tilde{v}$ to $\tilde{u}$, this means that $\tilde{u}'$ is also nearer to $\tilde{v}$ than $\tilde{u}$, a contradiction. Therefore the lemma follows. $\square$

The lemma shows that each Voronoi region $\text{Vr}(\tilde{b})$ has a spanning BFS dual-tree on (the subgraph induced by) the Voronoi region with $\tilde{b}$ as a root and that the depth of such trees is at most $2k$. Furthermore, it follows from Lemma 4.6, the number of Voronoi regions (i.e., the size of $\tilde{I}$) is bounded by $\tilde{n}/k$. This proves the properties stated in Claim 4.3. Before showing the algorithms for Step 2.1, we show one more lemma.

Lemma 4.8. For any dual-vertex $\tilde{b} \in \tilde{I}$ and any dual-vertex $\tilde{v} \in \text{Vr}(\tilde{b})$, there exists an algorithm that computes $\text{dist}(\tilde{v}, \tilde{b})$ in $\tilde{O}(k)$ -space.

Proof. For any dual-vertex $\tilde{u}$, we can enumerate all dual-vertices in $N_k^+(\tilde{u})$ in $\tilde{O}(k)$ -space (the details will be described in the Algorithms for Step 2.1). Thus, if $\tilde{v}$ is in $N_k^+(\tilde{b})$, we can compute $\text{dist}(\tilde{v}, \tilde{b})$ in $\tilde{O}(k)$ -space. Assume $\tilde{v}$ does not belong to $N_k^+(\tilde{b})$. In this case, $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) \neq \emptyset$ holds, and let $\tilde{u}$ be a vertex in $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b})$. We suppose that there is a dual-vertex $\tilde{w} \notin N_k^+(\tilde{v}) \cup N_k^+(\tilde{b})$ in one of the shortest dual-paths between $\tilde{v}$ and $\tilde{b}$. The distances $\text{dist}(\tilde{v}, \tilde{u})$ and $\text{dist}(\tilde{b}, \tilde{u})$ are at most $d_{\text{nb}}(\tilde{v}) + 1$ and $d_{\text{nb}}(\tilde{b}) + 1$ respectively, and the shortest dual-paths are clearly included in their k -neighborhoods⁺. Since $\tilde{w} \notin N_k^+(\tilde{v}) \cup N_k^+(\tilde{b})$, the distances $\text{dist}(\tilde{v}, \tilde{w})$ and $\text{dist}(\tilde{b}, \tilde{w})$ are more than $d_{\text{nb}}(\tilde{v}) + 1$ and $d_{\text{nb}}(\tilde{b}) + 1$ respectively. Thus $\tilde{w}$ cannot become a dual-vertex in the shortest dual-path, a contradiction. Therefore, we can compute $\text{dist}(\tilde{v}, \tilde{b})$ by calculating $\text{dist}(\tilde{v}, \tilde{u}) + \text{dist}(\tilde{b}, \tilde{u})$ for any $\tilde{u} \in N_k^+(\tilde{v}) \cap N_k^+(\tilde{b})$ and taking their minimum. $\square$

Algorithms for Step 2.1

We discuss algorithms for achieving Step 2.1. For algorithms we show below, it is rather easy to see that they run within polynomial-time in n ; thus, we

omit explaining their time complexity. We also explain algorithms as if one algorithm can pass a large amount of data to another algorithm; but as explained in Preliminaries section this process can be realized without using a space for keeping such data.

The task of Step 2.1 is to compute information on the Voronoi regions of G . More specifically, we consider an algorithm that outputs

- (a) the set $\tilde{I}$ of boss-vertices;
- (b) a family $\{N_k(\tilde{b})\}_{\tilde{b} \in \tilde{I}}$ of k -neighborhoods;
- (c) a list of $(\tilde{v}, \text{boss}(\tilde{v}), \tilde{v}_{\text{pre}})$ for all $\tilde{v} \in \tilde{V}$, where $\text{boss}(\tilde{v})$ is the boss-vertex of the Voronoi region that $\tilde{v}$ belongs to and $\tilde{v}_{\text{pre}}$ is the dual-vertex that is a parent in a BFS dual-tree of $\text{Vr}(\text{boss}(\tilde{v}))$; and
- (d) a list of boundary cycles of all Voronoi regions and lists of pairs of Voronoi regions sharing a vertex (resp., an edge).

A key algorithmic tool we use here is a BFS algorithm traversing at most t dual-vertices from a given source dual-vertex $\tilde{v}$. In this chapter we consider specific BFS that is convenient for designing space efficient algorithms. We define a tree T satisfying the following condition as a *BFS tree* (from a source dual-vertex $\tilde{v}$): $\tilde{w}$ is a child of $\tilde{u}$ in T if and only if $\tilde{u}$ has the smallest index among dual-vertices that are adjacent to $\tilde{w}$ and satisfies $\text{dist}(\tilde{v}, \tilde{w}) - \text{dist}(\tilde{v}, \tilde{u}) = 1$. For computing this BFS tree, we keep all vertices having the same level, namely the same distance from the root, and collect the vertices in the next level by processing smaller indexed vertex earlier; *Cf.* In a standard BFS algorithm, we process vertices in the order added to a queue regardless of the indices. It is easy to see that the algorithm runs in $\tilde{O}(t)$ -space. Then we have an $\tilde{O}(k)$ -space algorithm that enumerates, for a given $\tilde{v}$, all elements of $N_k(\tilde{v})$, the set of the first k closest dual-vertices to $\tilde{v}$ in $\tilde{G}$. We may assume that the distance to $\tilde{v}$ is also computed and kept for all enumerated dual-vertices. Recall that $d_{\text{nb}}(\tilde{v}) = \max_{\tilde{u} \in N_k(\tilde{v})} \text{dist}(\tilde{u}, \tilde{v})$. By using this algorithm, we can design an algorithm that decides whether a given dual-vertex $\tilde{u}$ is in $N_k^+(\tilde{v})$, where $N_k^+(\tilde{v})$ contains additionally all dual-vertices whose distance from $\tilde{v}$ is $d_{\text{nb}}(\tilde{v}) + 1$. Though we cannot keep $N_k^+(\tilde{v})$ since it could become very large, we can determine whether $\tilde{u} \in N_k^+(\tilde{v})$ by checking whether it is adjacent to some $\tilde{w} \in N_k(\tilde{v})$ whose distance from $\tilde{v}$ is $\leq d_{\text{nb}}(\tilde{v})$. Then by using this algorithm, we can design an $\tilde{O}(k)$ -space algorithm that determines, for given $\tilde{v}$ and $\tilde{b}$, whether $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) \neq \emptyset$; it simply search in $\tilde{V}$ a dual-vertex that belongs to both $N_k^+(\tilde{v})$ and $N_k^+(\tilde{b})$. We may assume that the algorithm can also identify the dual-vertex in $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) \neq \emptyset$ that is nearest to $\tilde{v}$, namely, $\text{nrst}_{\tilde{v}}(N_k^+(\tilde{b}))$, and compute its distance to $\tilde{v}$.

Armed with these algorithmics we can compute (a) $\sim$ (d) in $\tilde{O}(k + \tilde{n}/k)$ -space. First, for computing $\tilde{I}$ as the output (a), we simply collect dual-vertices to $\tilde{I}$ in a greedy way until there is no dual-vertex $\tilde{v}$ such that $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) = \emptyset$ for all so far obtained boss-vertices $\tilde{b}$. Note that the working memory for keeping the

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

obtained boss-vertices is needed here to achieve the greedy algorithm. Hence, we need $\tilde{O}(k + \tilde{n}/k)$ -space since we can bound $|\tilde{I}| \leq \tilde{n}/k$. We may assume that $N_k(\tilde{b})$ is computed (and produced as the output (b)) every time a new boss-vertex $\tilde{b}$ is obtained.

Next let us fix any dual-vertex $\tilde{v}$, and consider how to compute $(\tilde{v}, \text{boss}(\tilde{v}), \tilde{v}_{\text{pre}})$ as the output (c). For computing $\tilde{b}_{\tilde{v}} := \text{boss}(\tilde{v})$, we enumerate all $\tilde{b} \in \tilde{I}$ such that $N_k^+(\tilde{v}) \cap N_k^+(\tilde{b}) \neq \emptyset$; then $\tilde{b}_{\tilde{v}}$ is obtained as $\tilde{b}$ such that $\text{nrst}_{\tilde{v}}(N_k^+(\tilde{b}))$ is the nearest to $\tilde{v}$. According to our BFS tree definition, $\tilde{v}_{\text{pre}}$ needs to satisfy that it is adjacent to $\tilde{v}$ and $\text{dist}(\tilde{b}, \tilde{v}) - \text{dist}(\tilde{b}, \tilde{v}_{\text{pre}}) = 1$ holds. By Lemma 4.8, we can enumerate all such dual-vertices in $\tilde{O}(k)$ -space, and $\tilde{v}_{\text{pre}}$ is the dual-vertex having the smallest index among the candidates.


Finally, consider the computation for (d), that is, a list of boundary cycles of all Voronoi regions and lists of their incidence and edge-incidence relations. Consider any Voronoi region $\text{Vr}(\tilde{b})$ (specified its boss-vertex $\tilde{b}$). Since it is easy to determine, for a given edge of G , whether it is a boundary edge of $\text{Vr}(\tilde{b})$, we can enumerate all boundary edges in $\tilde{O}(1)$ -space. Then since the boundary of $\text{Vr}(\tilde{b})$ is a collection of cycles of G , we can simply use the $\tilde{O}(1)$ -space algorithm for computing a left-traverse for identifying all boundary cycles of $\text{Vr}(\tilde{b})$ from the set of its boundary edges. Thus, a list of boundary cycles of all Voronoi regions and lists of pairs of Voronoi regions sharing a vertex (resp., an edge) are $\tilde{O}(1)$ -space computable.

4.4 Multiple-Dual-Cycle (m.d.-cycle)

In this section we introduce the notion of “multiple dual-cycle with an orientation” (in short, m.d.-cycle), which will be used to discuss faces defined by dual-cycles. This section is a preliminary section for our later discussion, and it does not correspond to any algorithmic step.

An m.d.-cycle $\tilde{c}$ is a sequence of dual-vertices connected with dual-edges (with a direction implied by the sequence) that defines nonoverlapping subplane(s) under the assumed planar embedding. Intuitively, an m.d.-cycle is simply a collection of incident dual-cycles and paths; see Figure 4.9 for examples. Formally we have the following definition.

Definition 4.8 (m.d.-cycle). An *m.d.-cycle* $\tilde{c}$ (under the clockwise/anti-clockwise order) is a sequence $(\tilde{v}_1, \tilde{v}_2), (\tilde{v}_2, \tilde{v}_3), \dots, (\tilde{v}_{m-1}, \tilde{v}_m)$ of directed dual-edges satisfying the following conditions: (i) $\tilde{v}_1 = \tilde{v}_m$, that is, it starts and ends with the same dual-vertex; (ii) every directed dual-edge $(\tilde{v}_i, \tilde{v}_{i+1})$ of $\tilde{c}$ is based on a dual-edge $\{\tilde{v}_i, \tilde{v}_{i+1}\}$ of $\tilde{G}$; (iii) no directed dual-edge appears more than once in $\tilde{c}$ (while it is possible that some directed dual-edge and its reverse may both appear in $\tilde{c}$); and (iv) for every pair $(\tilde{v}_{i-1}, \tilde{v}_i)$ and $(\tilde{v}_i, \tilde{v}_{i+1})$ of consecutive directed dual-edges, $(\tilde{v}_i, \tilde{v}_{i+1})$ must be the clockwise (resp., anti-clockwise) next directed dual-edge from $\tilde{v}_i$ appearing in $\tilde{c}$. (We regard the reverse directed dual-edge $(\tilde{v}, \tilde{u})$ of $(\tilde{u}, \tilde{v})$ the furthest dual-edge from $(\tilde{u}, \tilde{v})$ both clockwise and



This figure shows examples of m.d.-cycles. White nodes are dual-vertices that are connected by dual-edges indicated by dashed lines. We may specify an m.d.-cycle by $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_4, \tilde{v}_5, \tilde{v}_6, \tilde{v}_7, \tilde{v}_8, \tilde{v}_5, \tilde{v}_4, \tilde{v}_3, \tilde{v}_1)$, which traverses a set of directed dual-edges under the anti-clockwise order (i.e., turning always right). Similarly, $(\tilde{v}_1, \tilde{v}_3, \tilde{v}_8, \tilde{v}_4, \tilde{v}_5, \tilde{v}_6, \tilde{v}_7, \tilde{v}_6, \tilde{v}_5, \tilde{v}_4, \tilde{v}_3, \tilde{v}_1)$ is a m.d.-cycle traversing the same set of dual-vertices (with a different set of directed dual-edges, though) under the clockwise order (i.e., turning always left). Note that non-duplicate dual-edges form dual-cycles, and duplicate dual-edges form a dual-path.

An example of the cases that one needs to be careful is the difference between sequences $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_8, \tilde{v}_4, \tilde{v}_5, \tilde{v}_6, \tilde{v}_7, \tilde{v}_6, \tilde{v}_5, \tilde{v}_4, \tilde{v}_3, \tilde{v}_1)$ and $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_8, \tilde{v}_4, \tilde{v}_5, \tilde{v}_6, \tilde{v}_7, \tilde{v}_6, \tilde{v}_5, \tilde{v}_4, \tilde{v}_3, \tilde{v}_1)$; the difference is only the last three (or two) dual-vertices. Both sequences traverse all dual-vertices, but the latter one is not an m.d.-cycle while the former one is an m.d.-cycle. The difference is the usage of four directed dual-edges at $\tilde{v}_3$.

Figure 4.9: M.d.-cycles


anti-clockwise.)

Remark. We use a sequence $(\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_m)$ of dual-vertices for specifying a m.d.-cycle $\tilde{c}$, by which we mean a sequence $(\tilde{v}_1, \tilde{v}_2), \dots, (\tilde{v}_{m-1}, \tilde{v}_m)$ of dual-edges of $\tilde{c}$. We also regard $\tilde{c}$ as a directed dual graph induced by its directed dual-edges.

Definition 4.9 (Duplicate/nonduplicate dual-edge and line part). For any m.d.-cycle $\tilde{c}$, a directed dual-edge $(\tilde{u}, \tilde{v})$ of $\tilde{c}$ is called a *duplicate dual-edge* if its reverse $(\tilde{v}, \tilde{u})$ also appears in $\tilde{c}$. Otherwise, it is called a *nonduplicate dual-edge*.

Remark. It is easy to see (Figure 4.9) that duplicate directed dual-edges of $\tilde{c}$ form dual-path(s) in $\tilde{c}$; this motivate us to call the set of duplicate directed dual-edges of $\tilde{c}$ the *line part* of $\tilde{c}$. On the other hand, nonduplicate dual-edges of $\tilde{c}$ form dual-cycle(s). We will ignore the case where $\tilde{c}$ consists of only duplicate dual-edges, and in the following, we may assume that $\tilde{c}$ has at least one dual-cycle.

Consider any m.d.-cycle $\tilde{c}$. As mentioned above, we assume that it has at least one dual-cycle. Hence, the plane (where $\tilde{G}$ is embedded) is separated by $\tilde{c}$. We would like to classify these subplanes as “inside” or “outside” of $\tilde{c}$.



Examples for showing the condition of m.d.-cycles. (a) A combination of directed dual-edges around a dual-vertex $\tilde{v}$ that cannot appear “syntactically” in any m.d.-cycle. (Here we consider this particular combination around $\tilde{v}$. There exist some m.d.-cycle that have these directed dual-edges (and more) around $\tilde{v}$. By “syntactically” we mean that this combination does not appear as a part of any m.d.-cycle in any order.) (b) A set of directed dual-edges that may appear in some m.d.-cycle depending on their order; that is, in the order of (1),(2), ..., (3),(4) under the anti-clockwise order, and in the order of (1),(4), ..., (3),(2) under the clockwise order. (c) A set of directed dual-edges that does not appear in any m.d.-cycle in the order of (1),(2), ..., (3),(4), or in the order of (3),(4), ..., (1),(2). Let us call this property the *noncrossing property* of m.d.-cycle, which we will use several times later.


Figure 4.10: Examples of invalid/valid directed dual-edge sets for m.d.-cycles

Intuitively speaking, we define the inside/outside of $\tilde{c}$ by the left/right of the directed dual-edges of $\tilde{c}$.

We define the left/right notion formally, and show that one can indeed determine for a given dual-vertex, whether it is located left or right of $\tilde{c}$ uniquely. In the following explanation, we abuse the notation $\tilde{G} \setminus \tilde{c}$ to denote a subgraph of $\tilde{G}$ obtained by “removing” $\tilde{c}$; precisely, $\tilde{G} \setminus \tilde{c}$ is $\tilde{G}[\tilde{V} \setminus V(\tilde{c})]$, a subgraph of $\tilde{G} = (\tilde{V}, \tilde{E})$ induced by $\tilde{V} \setminus V(\tilde{c})$. Also in the following, for any dual-edge $\tilde{e}$ of $\tilde{G}$, by a *dual-triangle upon $\tilde{e}$* and its *top-vertex*, we mean respectively a dual-triangle that uses $\tilde{e}$ as one of its three dual-edges and its dual-vertex opposite to $\tilde{e}$.

We first introduce the notion of “left/right adjacent.” Consider any directed dual-edge $\tilde{e}$ of $\tilde{c}$. Since $\tilde{G}$ is triangulated, there are two dual-triangles upon $\tilde{e}$. We regard the one located left (resp., right) of $\tilde{e}$ w.r.t. the direction of $\tilde{e}$ as *left adjacent-triangle* (resp., *right adjacent-triangle*). Then the top-vertex of the left (resp., right) adjacent-triangle is called the *left adjacent-vertex* (resp., *right adjacent-vertex*) of $\tilde{e}$. While these notions are defined for any dual-edge of $\tilde{c}$, we consider only nonduplicate dual-edges to define the notion of “left/right side” of $\tilde{c}$. For any dual-vertex we say it is in the *left* (resp., *right*) of $\tilde{c}$ if it is connected in $\tilde{G} \setminus \tilde{c}$ to the left (resp., right) adjacent-vertex of some nonduplicate dual-edge of $\tilde{c}$. The following lemma shows that the side of $\tilde{c}$ is uniquely determined in this way.

Lemma 4.9. Consider any m.d.-cycle $\tilde{c}$ of $\tilde{G}$. W.r.t. $\tilde{c}$, the side of every dual-vertex of $\tilde{G} \setminus \tilde{c}$ is uniquely determined as left or right. For any duplicate dual-edge of $\tilde{c}$, its left and right adjacent-vertices are in the same side; that is, they are both in the left (or in the right) of $\tilde{c}$.



The m.d.-cycle $\tilde{c}$ considered in the proof of the lemma is indicated by bold dashed line with arrows showing its direction.

Figure 4.11: An example for the proof of Lemma 4.9

Proof. We give a proof following Figure 4.11. Consider some m.d.-cycle $\tilde{c}$, and suppose that there exists a dual-vertex $\tilde{v}$ that is in the both left and right of $\tilde{c}$. Let $(\tilde{v}_i, \tilde{v}_{i+1})$ and $(\tilde{v}_j, \tilde{v}_{j+1})$ be directed nonduplicate dual-edges of $\tilde{c}$ witnessing that $\tilde{v}$ is in the left/right of $\tilde{c}$ respectively. Thus, there is a dual-path to $\tilde{v}$ from the right (resp., left) adjacent-vertex of $(\tilde{v}_i, \tilde{v}_{i+1})$ (resp., $(\tilde{v}_j, \tilde{v}_{j+1})$) that has no common dual-vertex with $\tilde{c}$. Without losing generality, we assume that $i < j$; that is, $(\tilde{v}_i, \tilde{v}_{i+1})$ appears earlier than $(\tilde{v}_j, \tilde{v}_{j+1})$ in $\tilde{c}$. Consider two subsequences of $\tilde{c}$, one from $\tilde{v}_{i+1}$ to $\tilde{v}_j$ and another from $\tilde{v}_{j+1}$ to $\tilde{v}_1$. They must cross at some dual-vertex $\tilde{u}$ as shown in Figure 4.11, contradicting the noncrossing property of m.d.-cycle (Figure 4.10). $\square$

Armed with this lemma, we can formally define the side of an m.d.-cycle. Intuitively, the inside of an m.d.-cycle $\tilde{c}$ is subplane(s) located left of the non line part of $\tilde{c}$.

Definition 4.10 (Inside/outside of an m.d.-cycle). Consider any m.d.-cycle $\tilde{c}$. A dual-vertex is *inside* (resp., *outside*) of $\tilde{c}$ if it is in the left (resp., right) of $\tilde{c}$. A dual-triangle is *inside* (resp., *outside*) of $\tilde{c}$ if its top-vertex is inside (resp., outside) of $\tilde{c}$. In general, the *inside* (resp., *outside*) of $\tilde{c}$ is the set of subplanes

of the plane dual graph $\tilde{G}$ consisting of faces (not including $\tilde{c}$ itself) defined by dual-triangles inside (resp., outside) of $\tilde{c}$.

4.5 Dividing Voronoi Regions

In this section we introduce the notion of “ridge edge” and define the notion of “pre-frame-cycle” that is a base for a frame-cycle and a frame-graph. For our algorithm, this section corresponds to Step 2.2, and we explain algorithms for computing ridge edges and pre-frame-cycles.

While a dual-cycle shown in Figures 4.4 is a good candidate for a frame-cycle, there is a situation where a more complicated dual-cycle is obtained similarly as illustrated in Figure 4.5, and we cannot bound in general the size of such dual-cycles. Such a situation occurs if some Voronoi region has multiple boundary cycles. Here we introduce “ridge edges” for adding more branch vertices, thereby dividing Voronoi regions and large dual-cycles appropriately.

Consider any $\tilde{b} \in \tilde{I}$ such that Voronoi region $\text{Vr}(\tilde{b})$ has more than one boundary cycle. Let us fix these $\tilde{b}$ and $\text{Vr}(\tilde{b})$ for a while. Let $\tilde{T}$ be the BFS dual-tree of $\text{Vr}(\tilde{b})$ that is obtained at the step 2.1 of our algorithm. Recall that it is an $O(k)$ -depth spanning tree of dual-vertices in $\text{Vr}(\tilde{b})$. Consider any two adjacent dual-vertices $\tilde{u}$ and $\tilde{v}$ that are not connected in $\tilde{T}$. Then by connecting $\tilde{u}$ and $\tilde{v}$ by $\tilde{e} = \{\tilde{u}, \tilde{v}\}$, a dual-cycle is created⁶ with two dual-paths of $\tilde{T}$. This dual-cycle divides the plane into two subplanes. We consider the case where this separation divides the set of boundary cycles of $\text{Vr}(\tilde{b})$. Below we say that a dual-edge $\tilde{e} = \{\tilde{u}, \tilde{v}\}$ *crosses* an edge e (and in parallel, e *crosses* $\tilde{e}$) if two faces corresponding to $\tilde{u}$ and $\tilde{v}$ share the edge e .


Definition 4.11 (Ridge edge). Consider $\text{Vr}(\tilde{b})$ as a subgraph of G . An edge e of $\text{Vr}(\tilde{b})$ is a ridge edge if the dual-edge $\tilde{e} = \{\tilde{u}, \tilde{v}\}$ of $\text{Vr}(\tilde{b})$ crossing e satisfies the following (below we let $\tilde{c}(\tilde{e})$ denote the dual-cycle induced by $\tilde{e}$ and two dual-paths of $\tilde{T}$ from the boss-vertex respectively to $\tilde{u}$ and $\tilde{v}$):

- $\tilde{e}$ is not a dual-edge of the BFS dual-tree $\tilde{T}$ of $\text{Vr}(\tilde{b})$, and
- each of two subplanes separated by $\tilde{c}(\tilde{e})$ contains at least one boundary cycle of $\text{Vr}(\tilde{b})$.

Below we use $R_{\tilde{b}}$ to denote the set of ridge edges in $\text{Vr}(\tilde{b})$. Use $B_{\tilde{b}}$ to denote the set of boundary edges of $\text{Vr}(\tilde{b})$. These sets are sometimes regarded as subgraphs of G .

We first point out a simple but important property that is immediate from the definition: namely, no BFS dual-path of $\tilde{T}$ crosses $R_{\tilde{b}} \cup B_{\tilde{b}}$. We then show the following two properties of $R_{\tilde{b}}$ and $B_{\tilde{b}}$.

⁶Precisely speaking, the created one should be considered as an m.d.-cycle instead of a simple cycle because two dual-paths of $\tilde{T}$ may be merged before the boss-vertex. But since it does not have more than one dual-cycle, we discuss below, for simplicity, as if it were a simple cycle with an orientation so that we can specify its inside/outside.



An example of a ridge edge. A bold dashed line is a dual-edge connecting two dual-vertices $\tilde{u}$ and $\tilde{v}$ that creates a dual-cycle with BFS dual-paths containing a Voronoi boundary cycle (indicated by a solid line) in its both sides. Then an edge crossing this dual-edge $\{\tilde{u}, \tilde{v}\}$ (indicated by a bold line) is a ridge edge.

Figure 4.12: Ridge edge

Lemma 4.10. $R_{\tilde{b}}$ as a subgraph of G is a forest. Furthermore, every leaf of $R_{\tilde{b}}$ is incident to some boundary edge of $B_{\tilde{b}}$.

Proof. Consider any connected component K of $R_{\tilde{b}}$. We first show that K does not contain a cycle. If K had a cycle, then the cycle divides $\text{Vr}(\tilde{b})$ into more than two subplanes each of which contains at least one dual-vertex (since dual-vertices correspond to faces of G). From the definition of a ridge edge, any dual-edge of the BFS dual-tree $\tilde{T}$ cannot cross a ridge edge; hence, there must be a subplane whose dual-vertices are not in the BFS dual-tree. This contradicts to the fact that $\tilde{T}$ is a spanning tree of $\text{Vr}(\tilde{b})$.

Next we show that every leaf of K is incident to the boundary of $\text{Vr}(\tilde{b})$. Here we assume otherwise; that is, there is some leaf v of a connected component of K that is not incident to the boundary of $\text{Vr}(\tilde{b})$ and lead a contradiction.

Let e be the ridge edge that is incident to v . We consider three faces incident to v that are regarded as dual-vertices $\tilde{v}_0$, $\tilde{v}_1$, and $\tilde{v}_2$ indexed in clockwise starting from the one adjacent to e so that a dual-edge $\{\tilde{v}_0, \tilde{v}_2\}$ crosses e . Let e_1 and e_2 are edges incident to v that are crossed by edges $\{\tilde{v}_0, \tilde{v}_1\}$ and $\{\tilde{v}_1, \tilde{v}_2\}$ respectively (Figure 4.13 (a)). Note that neither e_1 nor e_2 is a ridge edge (because v is a leaf of some connected component consisting of ridge edges). If $\tilde{v}_1$ does not belong to $\text{Vr}(\tilde{b})$, e_1 and e_2 become boundary edges since $\tilde{v}_0$ and $\tilde{v}_2$ are in $\text{Vr}(\tilde{b})$, and v is incident to the boundary, a contradiction. Thus $\tilde{v}_1$ is in $\text{Vr}(\tilde{b})$.

Consider a dual-cycle $\tilde{c}$ induced by (two dual-paths of) the BFS dual-tree $\tilde{T}$ and a dual-edge $\{\tilde{v}_0, \tilde{v}_2\}$ (Figure 4.13 (b)). The cycle $\tilde{c}$ divides the plane into two subplanes; let W and W' denote them. Note that both contain at least one boundary cycle of $\text{Vr}(\tilde{b})$ since e is a ridge edge. Without losing generality, we may assume that W is the subplane containing v . Note that W contains $\tilde{v}_0$, $\tilde{v}_1$,




Figure 4.13: Examples for the proof of Lemma 4.10

and $\tilde{v}_2$.

We have three cases:

- (i) The dual-edges $\{\tilde{v}_0, \tilde{v}_1\}$ and $\{\tilde{v}_1, \tilde{v}_2\}$ are both in the BFS dual-tree $\tilde{T}$: In this case, W becomes the triangle consisting of $\tilde{v}_0$, $\tilde{v}_1$ and $\tilde{v}_2$. Obviously, W has no boundary cycle of $\text{Vr}(\tilde{b})$.
- (ii) Either one of $\{\tilde{v}_0, \tilde{v}_1\}$ and $\{\tilde{v}_1, \tilde{v}_2\}$ is in $\tilde{T}$: Without loss of generality, we assume that $\{\tilde{v}_0, \tilde{v}_1\}$ is in $\tilde{T}$ and $\{\tilde{v}_1, \tilde{v}_2\}$ is not in $\tilde{T}$. Consider a dual-cycle $\tilde{c}$ induced by $\tilde{T}$ and a dual-edge $\{\tilde{v}_1, \tilde{v}_2\}$, and let the divided subplane by $\tilde{c}$ that does not contain v be U . The difference of W and U , namely $W \setminus U$, is the triangle consisting of $\tilde{v}_0$, $\tilde{v}_1$ and $\tilde{v}_2$. Thus U contains a boundary cycle of $\text{Vr}(\tilde{b})$; hence, e_2 becomes a ridge edge. A contradiction.
- (iii) Neither $\{\tilde{v}_0, \tilde{v}_1\}$ nor $\{\tilde{v}_1, \tilde{v}_2\}$ is in $\tilde{T}$: Consider a dual-cycle $\tilde{c}_1$ induced by $\tilde{T}$ and a dual-edge $\{\tilde{v}_0, \tilde{v}_1\}$, and let W_1 denote one of the two subplanes divided by $\tilde{c}_1$ that is entirely contained in W (Figure 4.13 (c)). Note that W_1 has no boundary cycle of $\text{Vr}(\tilde{b})$ because otherwise the other side of $\tilde{c}_1$, which contains W' , has no boundary of $\text{Vr}(\tilde{b})$ (since e_1 is not a ridge edge), contradicting the above mentioned fact that W' has at least one boundary cycle of $\text{Vr}(\tilde{b})$. Similarly, the subplane W_2 defined by a dual-cycle $\tilde{c}_2$ induced by $\tilde{T}$ and an edge $\{\tilde{v}_1, \tilde{v}_2\}$ (Figure 4.13 (d)) has no boundary cycle of $\text{Vr}(\tilde{b})$. Note, however, v is only one vertex of G that is in $W \setminus (W_1 \cup W_2)$. Thus, W has no boundary cycle of $\text{Vr}(\tilde{b})$. A contradiction. $\square$

Lemma 4.11. $B_{\tilde{b}} \cup R_{\tilde{b}}$ as a graph is connected.

Proof. For proving by contradiction, suppose that $B_{\tilde{b}} \cup R_{\tilde{b}}$ is not connected. Let C and C' be two connected components of $B_{\tilde{b}} \cup R_{\tilde{b}}$. Here we note that both C and C' consist of some Voronoi boundary cycles and ridge edges connecting them, which is immediate from the above lemma.

Let $\tilde{c}$ denote a set of dual-vertices of $\text{Vr}(\tilde{b})$ incident to edges of C . Due to the three regularity of G , it is easy to see that $\tilde{c}$ forms an m.d.-cycle under the anti-clockwise⁷ order whose orientation is determined so that C is located its outside. Note here the following two properties of $\tilde{c}$: Firstly, no dual-vertex of $\text{Vr}(\tilde{b})$ belongs to the outside of $\tilde{c}$; if otherwise, i.e., if there were some $\tilde{v} \in \text{Vr}(\tilde{b})$ outside of $\tilde{c}$, then the BFS dual-path from $\tilde{b}$ to $\tilde{v}$ should cross some edge of C , which would not occur by the basic property of $B_{\tilde{b}}$ and $R_{\tilde{b}}$. Thus, all dual-vertices of $\text{Vr}(\tilde{b})$ (including dual-vertices of $\tilde{c}$) are located inside of $\tilde{c}$. Then clearly, C' is also included in the inside of $\tilde{c}$. Secondly, no edge of $B_{\tilde{b}} \cup R_{\tilde{b}}$ crosses $\tilde{c}$, which is immediate from the first property.

Among directed dual-edges of $\tilde{c}$, there must be some dual-edges⁸ that are

⁷The choice of the anti-clockwise order is necessary to define an m.d.-cycle consisting of *all* vertices of $\tilde{c}$ for the case where $\tilde{c}$ has a line part like the one shown in Figure 4.14 (a). Since C is located in the outside (i.e., the right) of the m.d.-cycle, no line part of the m.d.-cycle is created in its left. Thus, choosing dual-vertices in the anti-clockwise order is necessary and sufficient to collect all dual-vertices of $\tilde{c}$.

⁸These dual-edges are not from the line part of $\tilde{c}$; see the explanation of Figure 4.14 (a). Thus, its direction in $\tilde{c}$ is uniquely determined.

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

not dual-edges of the BFS dual-tree $\tilde{T}$. Let $\tilde{e}_1, \dots, \tilde{e}_h$ denote such directed dual-edges of $\tilde{c} \setminus \tilde{T}$ listed in the order of the m.d.-cycle $\tilde{c}$ (e.g., Figure 4.14 (b)). For each $\tilde{e}_i$, consider a pair of *directed* dual-paths of $\tilde{T}$ from the boss-vertex $\tilde{b}$ to the tail of $\tilde{e}_i$ and the head of $\tilde{e}_i$ to $\tilde{b}$, with the direction consistent with the directed dual-edge $\tilde{e}_i$. Then this pair of dual-paths and $\tilde{e}_i$ define a dual-cycle⁹, which we denote by $\tilde{a}_i$. Also denote by a_i the subplane(s) defined as the inside of $\tilde{a}_i$. By our choice of the direction, $a_1, \dots, a_h$ are included in the inside of $\tilde{c}$; also it is clear that there is no overlap between them. Furthermore, every vertex of G located inside of $\tilde{c}$ is in some a_i ; see the explanation of Figure 4.14 (b). In this sense, the set of subplanes $a_1, \dots, a_h$ is a partition of the inside of $\tilde{c}$. In particular, there must be some $\tilde{a}_i$ that contains some and hence all vertices of C' because no edges of C' can cross the dual-edges of $\tilde{a}_i$. Then $\tilde{e}_i$ satisfies the condition given in Definition 4.11. That is, the edge that crosses $\tilde{e}_i$ is a ridge edge, which contradicts the second property of $\tilde{c}$ confirmed above. $\square$

Now we are ready to define notions of “branch vertex” and “connector” formally.

Definition 4.12 (B , R , and branch vertex). Hereafter we use B and R to denote respectively the set of Voronoi boundary edges and the set of ridge edges defined for G and $\tilde{G}$. Regard $B \cup R$ as an induced subgraph of G . A *branch vertex* is a degree three vertex in this graph. For each branch vertex v , a dual-triangle consisting of three dual-vertices incident to v is a *branch-triangle*.


Remark. From the above two lemmas on $B_{\tilde{b}} \cup R_{\tilde{b}}$, it is easy to see that $B \cup R$, as a graph, is a connected graph, and that every vertex of $B \cup R$ has degree either 2 or 3.

Definition 4.13 (Connector). Two branch vertices are *adjacent* if they are connected by edges in $B \cup R$ with no other branch vertex on it. The path connecting adjacent branch vertices is a *connector*.

Remark. It is easy to see that each connector consists of Voronoi boundary edges only or ridge edges only. Note that it is possible that some branch vertex is adjacent to itself; that is, it is connected to itself by a cycle connector. We fix some simple way to give a direction to each connector, and in the following, we may assume that each connector has this direction.

We are ready to define the notion of “pre-frame-cycle”, which will be used as a basis for defining “frame-cycles” and “frame-graph.” We prepare some notation. Consider any connector p (see Figure 4.15). Following the remark of Definition 4.13, we may assume that p has a direction, and we use v_{fst} and v_{last} be the start and end vertices of p w.r.t. this direction. Note that v_{fst} and v_{last} are an adjacent pair of branch vertices. Let us call an edge of p having v_{fst} as its end point the *first edge*, and similarly, an edge of p having v_{last} as its end point the *last edge*. We will refer a part of p removing the first and last edges as a *body*

⁹Again formally speaking, this may not be a real cycle because two dual-paths of the BFS dual-tree may merge before the boss-vertex. But for simplicity, we regard it as a simple dual-cycle.



Examples for the m.d.-cycle $\tilde{c}$ considered in the proof of Lemma 4.11. (a) An example for the case where an m.d.-cycle $\tilde{c}$ has a line part. Solid lines indicate edges of C , while dashed lines and white nodes are dual-edges and dual-vertices respectively of the m.d.-cycle $\tilde{c}$ incident to C . Note that dual-edges of a line part, i.e., $\tilde{e}_1, \tilde{e}_2, \tilde{e}_3$, are all dual-edges of the BFS dual-tree $\tilde{T}$ from $\tilde{b}$. Suppose otherwise and, say, $\tilde{e}_2$ were non-tree dual-edge; then at least one of the two BFS dual-paths from $\tilde{b}$ to $\tilde{e}_2$ must cross C , a contradiction. (b) Thin dashed lines are dual-paths of the BFS dual-tree $\tilde{T}$ from $\tilde{b}$. Bold dashed lines are dual-edges $\tilde{e}_1, \dots, \tilde{e}_4$ of $\tilde{c}$ that do not belong to $\tilde{T}$, where their order and directions are consistent with $\tilde{c}$. For each $\tilde{e}_i$, an m.d.-cycle is defined by two BFS dual-paths to the two end points of $\tilde{e}_i$ and $\tilde{e}_i$, with the direction consistent with that of $\tilde{e}_i$, which define a subplane a_i as its inside. Clearly, a_i is located inside of $\tilde{c}$, and furthermore, $a_1, \dots, a_4$ are the partition of the inside of $\tilde{c}$. To see this, for each $\tilde{e}_i$, consider $\tilde{v}_i$ (resp., $\tilde{v}_{i+1}$) that is the first dual-vertex in the BFS dual-path from $\tilde{b}$ to the tail (resp., the head) of $\tilde{e}_i$. Then it is easy to see that each pair a_i and a_{i+1} of adjacent faces share the same boundary from $\tilde{b}$ to $\tilde{v}_{i+1}$; no situation like the one crossed by X occurs because thin dashed lines are dual-tree edges. Thus, all vertices located inside $\tilde{c}$ must be in some face a_i .

Figure 4.14: Examples for the proof of Lemma 4.11

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

of p and denote it by p' . Also let $\tilde{e}_{\text{fst}} = \{\tilde{v}_{\text{fst.r}}, \tilde{v}_{\text{fst.l}}\}$ and $\tilde{e}_{\text{last}} = \{\tilde{v}_{\text{last.r}}, \tilde{v}_{\text{last.l}}\}$ be dual-edges that cross the first and the last edges respectively, where $\tilde{v}_{\text{fst.l}}$ (resp., $\tilde{v}_{\text{last.l}}$) is the one located left of the first (resp., the last) edge, and $\tilde{v}_{\text{fst.r}}$ (resp., $\tilde{v}_{\text{last.r}}$) is the one located right of the first (resp., the last) edge.




Figure 4.15: Notation on a connector

We can easily see that both $\tilde{v}_{\text{fst.l}}$ and $\tilde{v}_{\text{last.l}}$ (similarly, both $\tilde{v}_{\text{last.r}}$ and $\tilde{v}_{\text{fst.r}}$) have the same boss dual-vertex.

Lemma 4.12. $\text{boss}(\tilde{v}_{\text{fst.l}}) = \text{boss}(\tilde{v}_{\text{last.l}})$ and $\text{boss}(\tilde{v}_{\text{fst.r}}) = \text{boss}(\tilde{v}_{\text{last.r}})$.

Proof. Consider dual-vertices incident to p that are located on the left side of p . We name them $\tilde{v}_0 = \tilde{v}_{\text{fst.l}}$, $\tilde{v}_1, \dots, \tilde{v}_{t-1}$, $\tilde{v}_t = \tilde{v}_{\text{last.l}}$ so that $\tilde{v}_i$ and $\tilde{v}_{i+1}$ are adjacent. Assume that there exists $\tilde{v}_i$ such that $\text{boss}(\tilde{v}_i) \neq \text{boss}(\tilde{v}_{i+1})$, then an edge e that crosses $\{\tilde{v}_i, \tilde{v}_{i+1}\}$ must be an edge of a Voronoi boundary. But since e is incident to p , the end point of e that is on p must be a branch vertex. This contradicts that p is a connector. Therefore, there is no i such that $\text{boss}(\tilde{v}_i) \neq \text{boss}(\tilde{v}_{i+1})$, and hence $\text{boss}(\tilde{v}_{\text{fst.l}}) = \text{boss}(\tilde{v}_{\text{last.l}})$. With a similar argument, we can show $\text{boss}(\tilde{v}_{\text{fst.r}}) = \text{boss}(\tilde{v}_{\text{last.r}})$. $\square$

Then the notion of “pre-frame-cycle” is defined as follows. (This notion will be revised in the next section. Thus, let us call the one defined here as “the first version.”)

Definition 4.14 (pre-frame-cycle, the first version). Use notation defined in Figure 4.15; also see Figure 4.16. For any connector p , a *pre-frame-cycle* (w.r.t. p) is a “directed dual-cycle” that consists of (1) a BFS dual-path from $\tilde{v}_{\text{fst.r}}$ to its boss-vertex, (2) a BFS dual-path from this boss-vertex to $\tilde{v}_{\text{last.r}}$, (3) a branch-triangle dual-edge $\tilde{e}_{\text{last}}$ from $\tilde{v}_{\text{last.r}}$ to $\tilde{v}_{\text{last.l}}$, (4) a BFS dual-path from $\tilde{v}_{\text{last.l}}$ to its boss-vertex, (5) a BFS dual-path from this boss-vertex to $\tilde{v}_{\text{fst.l}}$, and (6) a branch-triangle dual-edge $\tilde{e}_{\text{fst}}$ from $\tilde{v}_{\text{fst.l}}$ to $\tilde{v}_{\text{fst.r}}$. The direction of this dual-cycle is defined naturally from the above order of dual-edges and dual-paths, from which the connector p is located in the left, i.e., the inside, of the pre-frame-cycle.

Remark. As shown in Figure 4.16, a pre-frame-cycle is not a simple dual-cycle in general; it may have a line part and/or it may consist of two dual-cycles. Thus, it seems better to regard it as an m.d.-cycle, more specifically, an m.d.-cycle under the clockwise order from Figure 4.16. Unfortunately, though, there are cases where an important part of a pre-frame-cycle is missed if we consider an m.d.-cycle under the clockwise order; see, e.g., Figure 4.17 (b),

(c). Furthermore, we will eliminate the cases where more than two dual-cycles appear as Figure 4.16 (b). Thus, though some pre-frame-cycle may have a line part, we regard it as a simple directed dual-cycle.

A dual-cycle we saw in Figure 4.4 is in fact a pre-frame-cycle given as Figure 4.16 (a) (1). On the other hand, a dual-cycle like the one in Figure 4.5 is now divided into several pre-frame-cycles thanks to the introduction of new branch vertices; see Figure 4.18.

Algorithms for Step 2.2

The task of this step is to compute information of the pre-frame-cycles of G . We again explain as if each step receives the output of the previous step as input, and omit analysis of time complexity. Note that we can use the outputs of Step 2.1 as input; the set of boss-vertices $\tilde{I}$, the BFS tree $\tilde{T}_{\tilde{b}}$ of every Voronoi region $\text{Vr}(\tilde{b})$ ($\tilde{b} \in \tilde{I}$), a list of boundary cycles of all Voronoi regions and a list of pairs of Voronoi regions sharing a vertex. We consider an algorithm that outputs

- (a) the set $B \cup R$ of Voronoi boundary edges and ridge edges;
- (b) the set of branch vertices of $B \cup R$;
- (c) a list of connectors; and
- (d) a list of pre-frame-cycles¹⁰.


First, for computing $B \cup R$ as the output of (a), we compute $R_{\tilde{b}}$ for each $\tilde{b} \in \tilde{I}$. Note that the boundaries $B_{\tilde{b}}$ are the input of the algorithm. If $\text{Vr}(\tilde{b})$ has only one boundary cycle, $R_{\tilde{b}}$ is empty from the definition of ridge edge. Assume $\text{Vr}(\tilde{b})$ has at least two boundary cycles. For each non-tree edge $\tilde{e} = \{\tilde{u}, \tilde{v}\}$ of $\text{Vr}(\tilde{b})$, we compute the path from $\tilde{u}$ to $\tilde{v}$ on $\tilde{T}_{\tilde{b}}$ with a standard DFS (depth-first search) algorithm. Since the diameter of $\tilde{T}_{\tilde{b}}$ is $O(k)$, this runs in $\tilde{O}(k)$ -space. We denote the dual-cycle¹¹ consisting of this path and $\tilde{e}$ as $\tilde{c}_{\tilde{e}}$, where its direction is determined in any appropriate way. For each boundary cycle of $B_{\tilde{b}}$, fix one dual-vertex $\tilde{v}_B$ of $\text{Vr}(\tilde{b})$ adjacent to $B_{\tilde{b}}$, and run a DFS algorithm from $\tilde{v}_B$ on $\tilde{T}_{\tilde{b}}$ until getting to some vertex of $\tilde{c}_{\tilde{e}}$, and check the side of $\tilde{v}_B$ (and hence, the boundary cycle) is located, which is again computed in $\tilde{O}(k)$ -space. If both sides of $\tilde{c}_{\tilde{e}}$ have at least one boundary cycle, we add the edge crossing $\tilde{e}$ to $R_{\tilde{b}}$.

Next consider how to compute the branch vertices. For any vertex v of G , we simply check whether v appears in $B \cup R$ three times or not.

Now we compute connectors as output (c). For each branch vertex v , we traverse on $B \cup R$ from v to all three directions until reaching another branch vertex u (u may be equal to v). The path from v to u is a connector. We


¹⁰This was not the output of Step 2.2 explained in Outline section; but we added this for the sake of explanation in the next section

¹¹Here again we explain, for simplicity, as if $\tilde{c}_{\tilde{e}}$ is a simple dual-cycle, though it may have a line part.



This figure shows four typical pre-frame-cycles. pre-frame-cycles are defined by bold dashed lines following the order indicated by numbers. Two boss-vertices are used for a pre-frame-cycle in type (a), and only one boss-vertex is used in type (b). Note that two subplanes are created as the inside/outside of the pre-frame-cycles in (a), whereas three subplanes are created in (b). Also note that a connector, i.e., a solid line with an arrow, is included (except its end points) in the inside subplane P_0 of a pre-frame-cycle for every type. A pre-frame-cycle of type (a) (2) has a line part, and more examples of pre-frame-cycles with a line part are shown in Figure 4.17.

Figure 4.16: Typical pre-frame-cycles



pre-frame-cycles may have line parts when it has two BFS dual-paths (with the same boss-vertex) merges before reaching to their boss-vertex. Here are some typical examples. Figures (b) and (c) show the cases where the line part indicated by thin dashed lines (and hence the corresponding boss-vertex) cannot be included if we regard them as m.d.-cycles under the clockwise order.

Figure 4.17: Examples of line parts of pre-frame-cycles

outputs the path only if v has smaller or equal index than u for not outputting a connector twice. The algorithm needs $\tilde{O}(1)$ -space.

Finally, we compute pre-frame-cycle as output (d). For each connector, we compute and output BFS dual-paths and branch-triangle dual-edges according to the definition of pre-frame-cycle. The dual-paths can be computed by DFS in $\tilde{O}(k)$ -space. Then a sequence (similar to an m.d.-cycle) representing a pre-frame-cycle can be easily obtained from these dual-paths and branch-triangle dual-edges in $\tilde{O}(k)$ -space.


4.6 Frame-Graph

We define the notion of “frame-graph” formally and show that an obtained frame-graph satisfying the condition (F1) $\sim$ (F3). This section corresponds to Step 2.3 of our algorithm.

Roughly speaking, our frame-graph is defined by removing all dual-vertices (and dual-edges) not participating in any pre-frame-cycle or branch-triangle, and these pre-frame-cycle and branch-triangles define all the faces of the frame-graph. The actual situation is a bit more complicated in general because of the type (b) of Figure 4.13 in which the 2-connectivity may not be guaranteed by removing unused dual-edges. Fortunately, we can show that this situation can be avoided by preprocessing pre-frame-cycles and merging two subplanes for type (b) graphs (if necessary).

4.6.1 Preprocessing pre-frame-cycles

We show a preprocessing step to simplify pre-frame-cycles so that each simplified pre-frame-cycle becomes a simple dual-cycle by removing (at most two) line



An example of new branch vertices that introduce new pre-frame-cycles. We use the Voronoi region of Figure 4.5. Thin solid lines are five boundary cycles of the Voronoi region. Bold solid lines are paths consisting of ridge edges connecting these three boundary cycles. Black vertices are new branch vertices, from which new branch-triangles (not shown here) are created. Then new branch dual-paths are introduced; for visibility, only three of them are shown as bold dashed lines in (a) while all of them are shown as bold dashed lines in (b). In (a), a pre-frame-cycle is indicated by its component four dual-paths given numbers in the order of its direction. Note that this directed dual-cycle contains one connector. Similarly, we can see in (b) (though a bit busy figure) that the long dual-cycle of Figure 4.5 is now separated into pre-frame-cycles.

Figure 4.18: New branch vertices and obtained pre-frame-cycles

parts. Thus, we will simply call such a simplified pre-frame-cycles as a dual-cycle although it may contain a line part.

In this preprocessing step, we go through all pre-frame-cycles, and for each pre-frame-cycle $\tilde{c}$, we check the number of dual-vertices in each subplane defined as the inside of $\tilde{c}$ and then take one of the following actions depending on these numbers: (i) do nothing, that is, keep $\tilde{c}$ as it is (only for the type (a)), (ii) use the algorithm of Lipton and Tarjan (Proposition 4.1) to compute a separator and use it as a desired separator of $\tilde{G}$, (iii) merge two subplanes (only for the type (b)), or (iv) use the pre-frame-cycle $\tilde{c}$ as a desired separator of $\tilde{G}$. Recall that we can compute a target separator of G_{org} from a separator of $\tilde{G}$ as explained in Outline section (i.e., Claim 4.2). Thus, the whole algorithm terminates with a desired output if either (ii) or (iv) is executed.

Now we explain this outline in detail. Consider any pre-frame-cycle $\tilde{c}$, and let us use the symbols given in Figure 4.16 for $\tilde{c}$. In particular, P_0 denotes a subplane defined as the inside of $\tilde{c}$ that has the body of a connector p , and let P_1 (resp., $P_{1,1}$ and $P_{1,2}$) a subplane(s) located outside of $\tilde{c}$. Let $\tilde{n}_0$ (resp., $\tilde{n}_1, \tilde{n}_{1,1}, \tilde{n}_{1,2}$) be the number of dual-vertices in P_0 (resp., $P_1, P_{1,1}, P_{1,2}$), not including dual-vertices on its boundary $\tilde{c}$. Let $\tilde{n}_{\tilde{c}}$ denote the number of dual-vertices of $\tilde{c}$; recall that $\tilde{n}_{\tilde{c}} = O(k) (= O(\sqrt{n}))$, which can be assumed negligible compared with $\tilde{n}$.

First consider the type (a). We have the following three cases (here we do not have the case (iii) mentioned above):

- (i) If $\tilde{n}_0 < \tilde{n}/4$: do nothing and use $\tilde{c}$ as a pre-frame-cycle.
- (ii) If $\tilde{n}_0 \geq 2\tilde{n}/3$: Note first that P_0 , that is, the inside of $\tilde{c}$ is included in the union of two Voronoi regions. Consider one of these Voronoi regions having at least half of dual-vertices in P_0 . Let $\tilde{G}_0$ be an induced subgraph of $\tilde{G}$ consisting of dual-vertices of $\tilde{G}$ located in this part; note that it has at least $\tilde{n}/3$ dual-vertices. Recall that we have a BFS dual-tree covering dual-vertices of $\tilde{G}_0$ of depth $O(k) (= O(\sqrt{n}))$, and also that this dual-tree (Section 4.3) and the list of dual-vertices inside of $\tilde{c}$ (Section 4.5) are given as input. Thus, we may assume that a BFS dual-tree of $\tilde{G}_0$ is given. Then we can apply the algorithm of Lipton and Tarjan to $\tilde{G}_0$ to obtain a separator of $\tilde{G}_0$ that separates $\tilde{G}_0$ into two subgraphs each of which has at least $(\tilde{n}/3)/3 = \tilde{n}/9$ dual-vertices. Since the work space needed for the algorithm of Lipton and Tarjan is linearly bounded by the depth of a BFS dual-tree of a given graph, this computation can be executed in $\tilde{O}(\sqrt{n})$ -space. Clearly, the obtained separator is also a $1/9$ -separator for $\tilde{G}$, that is, our desired separator.
- (iv) Otherwise: In this case we have $\tilde{n}_0 \geq \tilde{n}/4$ and $\tilde{n}_1 \geq \tilde{n} - \tilde{n}_0 - \tilde{n}_{\tilde{c}}$, which is larger than, say, $\tilde{n}/4$. Thus, we can output $\tilde{c}$ as a $1/4$ -separator of $\tilde{G}$.

Next consider the type (b). We have the following three cases (here we do not have the case (i) mentioned above):

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

- (ii) If $\tilde{n}_0 \geq \tilde{n}/3$: In this case, for an induced subgraph $\tilde{G}_0$ of $\tilde{G}$ consisting of dual-vertices of $\tilde{G}$ in P_0 , we take the same action as (ii) for the type (a).
- (iii) If either $\tilde{n}_0 + \tilde{n}_{1,1} < \tilde{n}/4$ or $\tilde{n}_0 + \tilde{n}_{1,2} < \tilde{n}/4$: Suppose that $\tilde{n}_0 + \tilde{n}_{1,1} < \tilde{n}/4$. Then we merge P_0 and $P_{1,1}$, and use a dual-cycle $\tilde{c}'$ bounding this subplane as a replacement of $\tilde{c}$ and *all* pre-frame-cycles in P_0 and $P_{1,1}$. More precisely, $\tilde{c}'$ is a sequence of dual-edges following $(5) \rightarrow (6) \rightarrow (1)$ of (b) in Figure 4.16. Then remove all pre-frame-cycles located left of $\tilde{c}'$, i.e., in the subplane $P_0 \cup P_{1,1}$, from the list of pre-frame-cycles; that is, a new pre-frame-cycle $\tilde{c}'$ becomes the boundary dual-cycle of the subplane $P_0 \cup P_{1,1}$ and no pre-frame-cycle exist in it. The case where $\tilde{n}_0 + \tilde{n}_{1,2} < \tilde{n}/4$ is handled in a similar way, and use a dual-cycle following $(2) \rightarrow (3) \rightarrow (4)$ to replace $\tilde{c}$ and all pre-frame-cycles in the subplane $P_0 \cup P_{1,2}$.
- (iv) Otherwise: In this case, we have $\tilde{n}_{1,1} + \tilde{n}_{1,2} \geq \tilde{n} - \tilde{n}_0 - \tilde{n}_{\tilde{c}}$, which is larger than, say, $\tilde{n}/2$ (since $\tilde{n}_0 < \tilde{n}/3$). Thus, either $\tilde{n}_{1,1}$ or $\tilde{n}_{1,2}$ must be larger than $\tilde{n}/4$. Let us assume that $\tilde{n}_{1,1} > \tilde{n}/4$. Then since $\tilde{n}_0 + \tilde{n}_{1,2} \geq \tilde{n}/4$, $\tilde{c}$ is a 1/4-separator separating $\tilde{G}$ into subgraphs located in $P_{1,1}$ and $P_0 \cup P_{1,2}$. Thus, we output $\tilde{c}$ as a desired separator of $\tilde{G}$.

In the rest of this chapter, we consider the situation where the algorithm does not terminate during this preprocessing and we still need to compute a separator. Also we revise our notion of pre-frame-cycle, and consider these dual-cycles obtained by this preprocessing as *pre-frame-cycles*. We should remark here that the inside of each pre-frame-cycle has less than $\tilde{n}/4$ dual-vertices of $\tilde{G}$.

Now we define “frame-graph” formally as follows.

Definition 4.15 (Frame-graph). For each pre-frame-cycle $\tilde{c}$ of $\tilde{G}$, let $\tilde{E}_{\tilde{c}}$ be the set of dual-edges that appear once in $\tilde{c}$. Let $\tilde{E}_1$ be the union of $\tilde{E}_{\tilde{c}}$ for all pre-frame-cycles $\tilde{c}$, and let $\tilde{E}_2$ be the set of all branch-triangle edges. Then a *frame-graph* of $\tilde{G}$ is a subgraph $\tilde{H} = (\tilde{U}, \tilde{D})$ of $\tilde{G}$, where $\tilde{D} = \tilde{E}_1 \cup \tilde{E}_2$ and $\tilde{U}$ is the set of all dual-vertices that are end points of dual-edges of $\tilde{D}$. A frame-graph is a plane graph under the embedding of $\tilde{G}$ restricted to $\tilde{U}$ and $\tilde{D}$.

Remark. From the definition of “pre-frame-cycle” (i.e., Figure 4.16 and our preprocessing step), a pre-frame-cycle $\tilde{c}$ is one dual-cycle that may have at most two line part(s) if dual-paths merge before its boss-vertices; hence, the set $\tilde{E}_{\tilde{c}}$ of dual-edges that appear once is exactly the dual-cycle part of $\tilde{c}$. Thus, we call it a *frame-cycle* and denote it by $(\tilde{c})^-$ if we want to consider it as a directed dual-cycle whose direction is consistent with $\tilde{c}$ (while $\tilde{E}_{\tilde{c}}$ is simply a set of dual-edges).

As explained in Outline section, a frame-graph $\tilde{H}$ needs to be a weighted subgraph of $\tilde{G}$. Here to meet the condition (F1), we define the weight of each face of $\tilde{H}$ as follows.

Definition 4.16 (Frame-graph and face weight). For each face of a frame-graph $\tilde{H}$, its *weight* is the number of dual-vertices of $\tilde{G}$ located in the face (that are removed to define $\tilde{H}$) divided by the total number $\tilde{n}^-$ of dual-vertices removed from $\tilde{G}$.

Algorithms for Step 2.3

The task of this step is to compute $\tilde{H}$, that is, its complete weighted face information, which consists of the list of branch-triangles and frame-cycles (as face boundary dual-cycles), the weights of faces defined by them, and their incidence relations. Since we have already computed pre-frame-cycles in Step 2.2 and the preprocessing step, we only need to consider a way to remove line parts from each pre-frame-cycle to obtain the corresponding frame-cycle. From the definition of pre-frame-cycle, Note that a line part exists in a pre-frame-cycle only if two dual-paths from two branch-triangles merge before the boss-vertex; such a line part can be identified by a standard DFS using $\tilde{O}(k)$ -space. Thus, we can compute the list of all frame-cycles of $\tilde{H}$ in $\tilde{O}(k)$ -space. From the obtained list of frame-cycles, it is easy to compute their incidence relations. Note that the number of removed dual-vertices from the inside of each frame-cycle has been already computed in the preprocessing step.

4.6.2 Properties of the frame-graph.

We show that the frame-graph $\tilde{H}$ satisfies the conditions (F1) $\sim$ (F3) given in Outline section.

We start with showing the 2-connectivity of $\tilde{H}$, that is, the condition (F2).

Lemma 4.13. $\tilde{H} = (\tilde{U}, \tilde{D})$ is 2-connected.

Proof. For any two distinct dual-vertices $\tilde{u}, \tilde{v} \in \tilde{U}$, we show that there exists two vertex disjoint dual-paths from $\tilde{u}$ to $\tilde{v}$ in $\tilde{H}$. When both $\tilde{u}$ and $\tilde{v}$ are on the same frame-cycle or the same branch dual-triangle, there exist two dual-paths that are parts of the dual-cycle from $\tilde{u}$ to $\tilde{v}$ clockwise and anti-clockwise. Suppose $\tilde{u}$ and $\tilde{v}$ are on two frame-cycles $\tilde{c}$ and $\tilde{c}'$ respectively. Let p_0 and p_* be the connectors whose bodies are contained in $\tilde{c}$ and $\tilde{c}'$ respectively. Since $B \cup R$ is connected, there is a sequence of connectors $p_0 = p_1, p_2, \dots, p_k = p_*$ such that p_i and p_{i+1} share one endpoint for each i , $1 \leq i < k$. Note also that for each i , $1 < i < k$, the body of p_i is contained in a frame-cycle $\tilde{c}_i$ of the type (a)(1) of Figure 4.16. We re-define the directions of p_i so that we can traverse from p_1 to p_k . For every connector p_i , it is obvious that two dual-paths on $\tilde{c}_i$ one from $\tilde{v}_{\text{fst.l}}$ to $\tilde{v}_{\text{last.l}}$ another from $\tilde{v}_{\text{fst.r}}$ to $\tilde{v}_{\text{last.r}}$ share no dual-vertex. We call these two dual-paths left- and right-path respectively. Then we can construct a path $\tilde{p}_L$ (resp., $\tilde{p}_R$) from $\tilde{v}_{\text{fst.l}}$ (resp., $\tilde{v}_{\text{fst.r}}$) of p_2 to $\tilde{v}_{\text{last.l}}$ (resp., $\tilde{v}_{\text{last.r}}$) of p_{k-1} by using only left- (resp., right-) paths and dual-edges of the corresponding branch-triangles; clearly, these two dual-paths $\tilde{p}_L$ and $\tilde{p}_R$ share no dual-vertex. That is, they are two vertex disjoint dual-paths from $\tilde{u}$ to $\tilde{v}$. (We need to consider

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

the cases where $\tilde{u}$ or $\tilde{v}$ is on a branch-triangle or on a frame-cycle based on a pre-frame-cycle created in the preprocessing step; these cases can be treated similarly, and we omit detail explanation for these cases.) $\square$

Next we consider the conditions (F1) and (F3). For this we need to identify faces of $\tilde{H}$. Intuitively, it is almost clear that every face of $\tilde{H}$ is defined by either a branch-triangle or a frame-cycle. We prove below this intuition. Recall that we define our frame-graph $\tilde{H}$ by collecting necessary dual-edges instead of by removing dual-vertices from the inside of each pre-frame-cycle. Thus, we need to prove that there is no dual-vertex of $\tilde{H}$ inside of each frame-cycle to guarantee that each frame-cycle indeed is a boundary of a face of $\tilde{H}$. We also need to prove that there is no other face other than those defined by branch-triangles or frame-cycles.

We first discuss (almost) one-to-one correspondence between pre-frame-cycles and connectors in $\tilde{G}$. Consider any pre-frame-cycle $\tilde{c}$ that is not among those introduced as a boundary dual-cycle of two merged subplanes in the preprocessing step. From the definition of pre-frame-cycle, it is clear that the inside of $\tilde{c}$ has the body of some connector. We show that it indeed does not contain any other edge of $B \cup R$.

Lemma 4.14. Consider $\tilde{G}$ and its pre-frame-cycle $\tilde{c}$ that is not introduced in the preprocessing step. Let p' be the body of a connector p that is located inside of $\tilde{c}$. Then there is no edge of $B \cup R$ except p' that is inside of $\tilde{c}$.


Proof. We use the notation of Figure 4.15 here with p and $\tilde{c}$ of the lemma. Since the BFS dual-tree of each Voronoi region does not cross any edge of $R \cup B$, only two dual-edges $\tilde{e}_{\text{fst}}$ and $\tilde{e}_{\text{last}}$ of $\tilde{c}$ are crossing edges of $R \cup B$. Therefore, $R \cup B$ is parted into at most three connected components by $\tilde{c}$, and each component is in either in the left side of $\tilde{c}$ (that is, inside of $\tilde{c}$) or in the right side of $\tilde{c}$ (that is, not outside of $\tilde{c}$). Clearly, one component of them is p' and it is inside of $\tilde{c}$. The others¹² contain the branch vertices of p , i.e., two end points of p that are not in the pre-frame-cycle $\tilde{c}$; hence, those components have no edge in the inside of $\tilde{c}$. $\square$

We now show that there is no dual-vertex of $\tilde{H}$ inside of each pre-frame-cycle. Thus, each frame-cycle is indeed a boundary of a face of $\tilde{H}$. In the following, we consider $\tilde{H}^+$ that is defined as a subgraph of $\tilde{G}$ consisting of all dual-vertices and dual-edges of branch-triangles and pre-frame-cycles. Note that $\tilde{H}$ is obtained from $\tilde{H}^+$ by removing dual-vertices and dual-edges not participating in frame-cycles or branch-triangles.

Lemma 4.15. Consider any pre-frame-cycle $\tilde{c}$. No dual-edge of $\tilde{H}^+ \setminus \tilde{c}$ is located inside of $\tilde{c}$. Thus, the inside of $(\tilde{c})^-$ has no dual-vertex of $\tilde{H}$.

Remark. It is possible that the inside of $(\tilde{c})^-$ has some dual-vertex of $\tilde{c}$, which is a dual-vertex of a line part of $\tilde{c}$ that is removed in $\tilde{H}$ (Figure 4.19).

¹²There is a case where p forms a cycle from a branch vertex (Figure 4.16 (a)(2)), in which case $B \cup R$ has only one component other than p' ; this case can be treated similarly.



Typical examples of line parts of pre-frame-cycles that are removed from $\tilde{H}^+$ when defining $\tilde{H}$. Dashed lines (both bold and thin) are pre-frame-cycles (where the arrows indicate the directions for defining their insides), bold dashed lines are their line parts removed in $\tilde{H}$. Figure (c) indicates the case where a new pre-frame-cycle is created by merging two suplanes in the preprocessing step.

Figure 4.19: Examples of line parts removed in $\tilde{H}$

Proof. Consider any pre-frame-cycle $\tilde{c}$. If it is one of those introduced in the preprocessing step, then the lemma is immediate because the inside of $\tilde{c}$ has no pre-frame-cycle (due to the preprocessing step), and hence all dual-vertices are removed from there when defining $\tilde{H}$. Thus, in the following, we consider a pre-frame-cycle $\tilde{c}$ of the type (a) of Figure 4.16. In particular, we assume for simplicity that it is of the type (a)(1); the type (a)(2) can be treated similarly. Let p' be the body of a connector located inside of $\tilde{c}$.

Assume to the contrary that some dual-edge $\tilde{e}$ of the other pre-frame-cycle $\tilde{c}'$ exists inside of $\tilde{c}$. It cannot be a dual-edge of a branch-triangle because if so, the inside of $\tilde{c}$ would have an edge of $B \cup R$ other than p' , contradicting to Lemma 4.14 above. Hence, it must be a part of a dual-path $\tilde{p}$ of a BFS dual-tree. Since there is no Voronoi boundary (besides p'), the boss of this BFS dual-tree (and one end of the dual-path $\tilde{p}$) is the boss-vertex of $\tilde{c}$ located in the same side of p' as $\tilde{e}$; let us denote it $\tilde{b}$. On the other hand, the other end of this dual-path in $\tilde{c}'$ must be a branch-triangle that is located outside of $\tilde{c}$ because the corresponding branch vertex must be outside of $\tilde{c}$. Thus, the dual-path $\tilde{p}$ must cross $\tilde{c}$, and let $\tilde{w}$ denote a dual-vertex of the crossing point; that is, $\tilde{w}$ is a dual-vertex in both $\tilde{p}$ and $\tilde{c}$. Clearly, its boss-vertex is $\tilde{b}$ because $\tilde{p}$ cannot cross p' . Thus, we have two dual-paths from $\tilde{w}$ to $\tilde{b}$, one following $\tilde{c}$ and another following $\tilde{p}$, contradicting the fact that each dual-vertex has a unique BFS dual-path to its boss-vertex. $\square$

From this lemma, we can claim that each frame-cycle is a boundary of a face in $\tilde{H}$ because it is a dual-cycle (see Remark of Lemma 4.15) having no dual-vertex in its inside in $\tilde{H}$.

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

Lemma 4.16. The boundary of any face of frame-graph $\tilde{H}$ is either a frame-cycle or a branch-triangle.

Proof. Consider any face of $\tilde{H}$. Since $\tilde{H}$ is 2-connected, the face has a single boundary dual-cycle $\tilde{c}$ (Proposition 2.1). Let us assume that this is not a branch-triangle and show that it indeed is a frame-cycle.

Note first that $\tilde{c}$ has at least one dual-edge from some branch-triangle. This is because $\tilde{H}$ consists of dual-edges of disjoint BFS dual-trees of $\tilde{G}$ (from disjoint Voronoi regions) and branch dual-triangles and we need at least one dual-edge of some branch-triangle since no dual-cycle is formed by dual-edges of disjoint dual-trees. Let $\tilde{e}$ and $\tilde{t}$ be respectively such a dual-edge and the branch-triangle with this dual-edge. Note also that $\tilde{c}$ and $\tilde{t}$ are (the boundaries of) adjacent faces sharing dual-edge $\tilde{e}$.

Here we overlap the original plane graph G with $\tilde{H}$. Then by definition there exists a branch vertex in branch-triangle $\tilde{t}$ where three connectors of $B \cup R$ are merged. Clearly, there should be one connector that starts with an edge crossing dual-edge $\tilde{e}$. Thus, $\tilde{c}$ contains this connector. On the other hand, there must be a frame-cycle $\tilde{c}'$ whose inside contains the body of this connector (Definition 4.14¹³). Note that both $\tilde{c}$ and $\tilde{c}'$ are boundaries of some faces of $\tilde{H}$. Therefore, $\tilde{c}$ must be the same as $\tilde{c}'$ since two faces cannot overlap; that is, $\tilde{c}$ is a frame-cycle. $\square$

Now that we have identified all faces of $\tilde{H}$, we consider the conditions (F1) and (F3) for $\tilde{H}$. Recall that the condition (F1) requires that each face weight is less than $1/3$. Since faces are defined by either a branch-triangle or a frame-cycle, and the former case is trivial, we need to consider faces defined a frame-cycle. Note that the preprocessing step guarantees that the number of dual-vertices of $\tilde{G}$ inside of each frame-cycle is less than $\tilde{n}/4$. Furthermore, for the total number $\tilde{n}^-$ of removed dual-vertex, we will see later (in Section 9) that $\tilde{n}^- = \tilde{n} - O(\sqrt{\tilde{n}})$; hence, we have $\tilde{n}/4 < \tilde{n}^-/3$ for sufficiently large n , satisfying the weight requirement of (F1).

For the condition (F3), that is, for showing that $\tilde{H}$ has $O(n/k)$ faces, we count the number of connectors and branch vertices.

Lemma 4.17. The number of connectors is $O(n/k)$. Also, the number of branch vertices is $O(n/k)$.

Proof. We consider a plane graph $G' = (V', E')$ defined from $B \cup R$, where V' is the set of branch vertices of $B \cup R$, and E' is the set of edges between adjacent branch vertices in $B \cup R$; thus, each edge of E' corresponds to some connector. We assume the planar embedding following that of G . Since G' is plane, it satisfies Euler's formula $|F'| + |V'| = |E'| + 2$, where F' is the set of faces of G' . Since the degree of any branch vertex is 3, we have $2|E'| = 3|V'|$. By substituting this to the above Euler's formula, we have $3|F'| + 2|E'| = 3|E'| + 6$, implying $|E'| = 3|F'| - 6 = O(|F'|)$.

¹³Precisely speaking, we also need to consider frame-cycles introduced in the preprocessing step; but it is easy to see that this fact remains true even after the preprocessing step.

Here we note that there is one to one correspondence between faces of G' and Voronoi regions. To see this, consider a graph consisting of only B edges, i.e., Voronoi boundary edges. Then the one-to-one correspondence is clear. On the other hand, adding ridge edges does not divide any Voronoi region; otherwise, the BFS dual-tree of the divided Voronoi region crosses some ridge edge, contradicting the definition of ridge edges. Hence, no new face is created by adding R edges. Thus, we have $|F'| = O(n/k)$. Therefore, we have $|E'|$, which is the number of connectors, is $O(n/k)$. Also we can bound the number of branch vertices by $O(n/k)$ because $2|E'| = 3|V'|$. $\square$

Then the following corollary is immediate from the one-to-one correspondence respectively between branch-triangles and branch vertices and between frame-cycles and connectors.

Corollary 4.18. The number of faces of $\tilde{H}$ is $O(n/k)$.

4.7 Floor and Ceiling Modification

In this section, we formally define the notion of floor- and ceiling-cycles, and explain their properties. We then explain Step 2.4 of our algorithm for computing floor- and ceiling-cycles and modified $\tilde{H}$ as a final output of the whole Step 2.

We begin with defining the notion of “core”, which is used in case no neck appropriate for a floor-cycle exists. Recall that $d_{\text{nb}}(\tilde{v})$ is the largest d such that $|\cup_{0 \leq i \leq d} L(\tilde{v}, i)| < k$ holds.

Definition 4.17 (Core). For any boss-vertex $\tilde{b} \in \tilde{I}$, let $d_{\text{core}}(\tilde{b})$ denote the largest $d \leq d_{\text{nb}}(\tilde{b})$ such that $|L(\tilde{b}, d)| \leq \sqrt{k}$. The *core* of $\tilde{b}$ (denoted by $\text{Core}(\tilde{b})$) is defined by

$$\text{Core}(\tilde{b}) = \bigcup_{0 \leq i \leq d_{\text{core}}(\tilde{b})} L(\tilde{b}, i).$$

We first note the following relation.

Lemma 4.19. For any $\tilde{b} \in \tilde{I}$, we have $d_{\text{nb}}(\tilde{b}) - d_{\text{core}}(\tilde{b}) \leq \sqrt{k}$.

Proof. Suppose otherwise, that is, $d_{\text{core}}(\tilde{b}) + \sqrt{k} < d_{\text{nb}}(\tilde{b})$ holds. Then we have


$$\bigcup_{i=1}^{\sqrt{k}} L(\tilde{b}, d_{\text{core}}(\tilde{b}) + i) \subseteq N_k(\tilde{b}).$$

On the other hand, by definition, for any d such that $d_{\text{core}}(\tilde{b}) < d \leq d_{\text{nb}}(\tilde{b})$, we have $|L(\tilde{b}, d)| > \sqrt{k}$, from which we have $|N_k(\tilde{b})| > k$, contradicting the definition of k -neighborhood. $\square$

Consider any boss-vertex $\tilde{b} \in \tilde{I}$ and its core $\text{Core}(\tilde{b})$. By definition $\text{Core}(\tilde{b})$ is a subset of $N_k(\tilde{b})$. Hence, $k' := |\text{Core}(\tilde{b})| < k$; that is, the core contains less than

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

k dual-vertices. It is also a region. Thus, its boundary is a set of disjoint cycles of G . Let c be one of such boundary cycles. We regard the subplane separated by c not containing any dual-vertex of the core as the *outside* of c . The cycle c is called a *core boundary cycle* if its outside contains at least $2\tilde{n}/3$ dual-vertices. It may be the case that no core boundary cycle exists. In this case, by removing the core, we would have separated subgraphs $\tilde{G}_1, \dots, \tilde{G}_t$ from which we can define $\cup_{1 \leq i \leq t_0} \tilde{G}_i$ and $\cup_{t_0 < i \leq t} \tilde{G}_i$ for some t_0 such that both parts have at least $(\tilde{n}/3) - k'$ dual-vertices. Hence, the core itself is, say, a $1/4$ -separator of $\tilde{G}$ for sufficiently large $k = \sqrt{n}$ and $\tilde{n}$, and its size is $k' \leq \sqrt{n}$; thus, the core itself can be used as a desired separator from which we can construct a target separator of G_{org} using the method stated in Lemma 4.4. Therefore, in the following we may assume that each core has a core boundary cycle. Note also that more than one core boundary cycles do not exist for each boss-vertex because the outside of such cycles are disjoint and no disjoint two sets can contain $2\tilde{n}/3$ dual-vertices in each. Then for such a core boundary cycle, consider a set $\tilde{c}$ of dual-vertices in $\text{Core}(\tilde{b})$ (regarded as faces of G) sharing a boundary edge with the core boundary cycle. We show in Lemma 4.20 below that a subgraph of $\tilde{G}$ induced by $\tilde{c}$ in fact forms a dual-cycle. This observation leads the following definition.



A bold solid line indicates the core boundary cycle c stated in the lemma, and some of its vertices are shown as black nodes. White nodes are some of dual-vertices that share a face boundary edge with c . In particular, dual-vertices $\tilde{v}_1$ and $\tilde{v}_2$ are those in $\text{Core}(\tilde{b})$ sharing boundary edges with c , and edges e_1 and e_2 are representatives of such boundary edges discussed in the proof. It is easy to see that such dual-vertices form an m.d.-cycle $\tilde{c}$ under the anti-clockwise order that is indicated by a bold dashed line with arrows showing its direction.


Figure 4.20: An example of core boundary cycle and core-cycle for the proof of Lemma 4.20 (1)

Definition 4.18 (Core boundary cycle and core-cycle). For each core $\text{Core}(\tilde{b})$, its *core boundary cycle* is a cycle such that (i) it is one of the boundary cycles

of the region $\text{Core}(\tilde{b})$, and (ii) its outside has at least $2\tilde{n}/3$ dual-vertices. The *core-cycle* of $\text{Core}(\tilde{b})$ is a directed dual-cycle induced by the set of dual-vertices in $\text{Core}(\tilde{b})$ sharing an edge with the core boundary cycle. The *inside* of the core-cycle is the side with the boss-vertex $\tilde{b}$.

Lemma 4.20. Consider any boss-vertex $\tilde{b}$ and its core $\text{Core}(\tilde{b})$, and let c be its core boundary cycle. Then we can define an m.d.-cycle $\tilde{c}$ under the anticlockwise order consisting of all dual-vertices of $\text{Core}(\tilde{b})$ sharing an edge with c . This $\tilde{c}$ in fact is a dual-cycle. (We may assume that c is directed so that its inside, i.e., the left side w.r.t. the direction, contains $\text{Core}(\tilde{b})$. The direction of $\tilde{c}$ is also determined consistent with that of c , by which $\tilde{b}$ is located the inside of $\tilde{c}$.)

Remark. From the relation between c and $\tilde{c}$, we refer $\tilde{c}$ as the m.d.-cycle *left-adjacent* to c .



A graph consisting of black vertices connected by solid edges is a part of the core boundary cycle c stated in the lemma. A part of the m.d.-cycle $\tilde{c}$ is indicated by bold dashed directed edges connecting white vertices as $\tilde{v}_1, \tilde{v}, \tilde{v}_2, \dots, \tilde{v}_3$, and $\tilde{v}, \dots$. Its sub m.d.-cycle $\tilde{c}'$ considered in the proof corresponds to the part $\tilde{v}, \tilde{v}_2, \dots, \tilde{v}_3$, and $\tilde{v}$. Here $\tilde{w}_1, \tilde{w}_2$, and $\tilde{w}_3$ are dual-vertices located in the opposite side, i.e., the outside of c from, e.g., $\tilde{v}$.

Figure 4.21: An example for the proof of Lemma 4.20 (2)

Proof. Let $d_0 = d_{\text{core}}(\tilde{b})$; hence, $\text{Core}(\tilde{b}) = \cup_{d \leq d_0} L(\tilde{b}, d)$. We first show the existence of the m.d.-cycle $\tilde{c}$ stated in the lemma (see Figure 4.20). Consider any edge e_1 of the boundary cycle c , and let $\tilde{v}_1$ be a dual-vertex in $\text{Core}(\tilde{b})$ (regarded as a face of G) that has e_1 as one of its face boundary edges. Clearly, c has edges that are not boundary edges of $\tilde{v}_1$, and let e_2 be the first such edge from e_1 following the direction so that $\tilde{v}_1$ is located left. Then we have some dual-vertex $\tilde{v}_2$ in $\text{Core}(\tilde{b})$ having e_2 (as its face boundary edge). Similarly, we can find e_3 and $\tilde{v}_3$, and so on until coming back to e_1 . Let $\tilde{c}$ be the obtained sequence of dual-vertices. Then $\tilde{c}$ is an m.d.-cycle under the anti-clockwise order that has all dual-vertices of $\text{Core}(\tilde{b})$ in its inside. This is because there is no dual-vertex of $\text{Core}(\tilde{b})$ from $\tilde{v}_{i-1}$ to $\tilde{v}_{i+1}$ anti-clockwise among all dual-vertices

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

adjacent to $\tilde{v}_i$. In fact, it is also easy to see that all $\tilde{v}_i$'s are in $L(\tilde{b}, d_0)$ facing $\tilde{w}_j$'s in $L(\tilde{b}, d_0 + 1)$.

We then prove that $\tilde{c}$ is indeed a dual-cycle (see Figure 4.21). For this, it suffices to show that the degree of every dual-vertices of $\tilde{c}$ is two in $\tilde{c}$. Suppose otherwise; that is, we had a dual-vertex $\tilde{v}$ connecting more than two dual-vertices in $\tilde{c}$. Hence, the m.d.-cycle $\tilde{c}$ would be expressed as a sequence $(\tilde{v}_1, \tilde{v}, \tilde{v}_2, \dots, \tilde{v}_3, \tilde{v}, \tilde{v}_4, \dots)$. (It may be the case that $\tilde{v}_4$ is the same as $\tilde{v}_3$.) We can split this into two subsequences: $(\tilde{v}, \tilde{v}_2, \dots, \tilde{v}_3)$ and the remaining one, i.e., $(\tilde{v}_1, \tilde{v}, \tilde{v}_4, \dots)$. Note that both are m.d.-cycles, and clearly, one of them has the boss-vertex $\tilde{b}$ in its inside while the other one does not. Without losing generality, we assume that $\tilde{c}' = (\tilde{v}, \tilde{v}_2, \dots, \tilde{v}_3)$ does not have the boss-vertex $\tilde{b}$ in its inside. Now consider, e.g., the dual-vertex $\tilde{v}_2$. Since $\tilde{v}_2 \in \text{Core}(\tilde{b})$, its distance from $\tilde{b}$ is $\leq d_0$; hence, there must be some dual-vertex $\tilde{u} \in L(\tilde{b}, d_0 - 1)$ inside of $\tilde{c}'$ adjacent to $\tilde{v}_2$. Thus, there must be a path of length $d_0 - 1$ from $\tilde{b}$ to $\tilde{u}$. On the other hand, this path should go through $\tilde{v}$ because $\tilde{b}$ is in the inside of the other m.d.-cycle. But then the distance of $\tilde{v}$ becomes less than d_0 , contradicting the assumption that $\tilde{v}$ shares a boundary edge with c , which means that there exists a dual-vertex $\tilde{w}_1$ that shares this boundary edge with $\tilde{v}$ in the outside of $\text{Core}(\tilde{b})$, i.e., in $L(\tilde{b}, d_0 + 1)$. $\square$

We state here the following size bounds, which are immediate from the definition.

Lemma 4.21. Each core-cycle has at most $\sqrt{k}$ dual-vertices and at most $\tilde{n}/3$ dual-vertices in its inside.


We define the notions of neck, floor, and ceiling. For this we introduce the notion of “level.” The main difference from the intuitive explanation in Outline section is that it is defined in terms of the distance from *all* k -neighborhoods (instead of a single boss-vertex). For any $\ell \geq 1$, let $\tilde{L}_{\text{nb}}(\ell)$ denote a set of dual-vertices $\tilde{v}$ whose distance from its nearest k -neighborhood in $\{N_k(\tilde{b})\}_{\tilde{b} \in \tilde{I}}$ is ℓ . More formally, it is defined by

$$\tilde{L}_{\text{nb}}(\ell) = \{ \tilde{v} \mid \text{dist}(\tilde{v}, \tilde{v}_{\text{nrst}}) = \ell, \text{ where } \tilde{v}_{\text{nrst}} = \text{nrst}_{\tilde{v}}(N_k(\text{boss}(\tilde{v}))) \}$$

Clearly, a family $\{\tilde{L}_{\text{nb}}(\ell)\}_{\ell \geq 1}$ is a partition of $\tilde{V}' := \tilde{V} \setminus \cup_{\tilde{b} \in \tilde{I}} N_k(\tilde{b})$. For any dual-vertex $\tilde{v} \in \tilde{V}'$ (resp., any set $U \subseteq \tilde{V}'$ of dual-vertices), the *level* of $\tilde{v}$ (resp., U) is ℓ such that $\tilde{v} \in \tilde{L}_{\text{nb}}(\ell)$ holds (resp., $U \subseteq \tilde{L}_{\text{nb}}(\ell)$ holds). We use $\tilde{L}_{\text{nb}}(\ell)$ also to denote a subgraph of $\tilde{G}$ induced by $\tilde{L}_{\text{nb}}(\ell)$, which we call a *sliced graph* of level ℓ .

We first note a property of sliced graphs that allows us to define interior and exterior boundary cycles.

Lemma 4.22 (Interior and exterior boundary cycles). For any $\ell \geq 1$, consider any connected component C of sliced graph $\tilde{L}_{\text{nb}}(\ell)$ with a dual-vertex adjacent to some dual-vertex of level $\ell + 1$. Then C has at least two disjoint boundary



An example of a connected component C of some level graph $\tilde{L}_{nb}(\ell)$. A solid line indicates its exterior boundary cycle c , and a dashed line indicates the m.d.-cycle $\tilde{c}$ in C right-adjacent to c ; note that the edges of $\tilde{c}$ are traversed under the clockwise order. The level of both $\tilde{v}_1$ and $\tilde{v}_2$ (and in fact all dual-vertices of $\tilde{c}$) is ℓ , which is determined by the distance from either $N_k(\tilde{b}_1)$ (e.g., for $\tilde{v}_1$) or $N_k(\tilde{b}_2)$ (e.g., for $\tilde{v}_2$).

Figure 4.22: Exterior boundary and its right-adjacent m.d.-cycle

cycles: *interior boundary cycle* and *exterior boundary cycle* (see the proof for their definition). Furthermore, for these boundary cycles, we can define m.d.-cycles like core-cycles. More precisely, consider any one of such boundary cycles c . Then a set of dual-vertices of C sharing a boundary edge with c forms an m.d.-cycle $\tilde{c}$ under the clockwise order.

Remark. It may be the case that the above $\tilde{c}$ is not a simple dual-cycle; see Figure 4.22. (Cf. A core-cycle is a simple dual-cycle.) Below we will refer this $\tilde{c}$ as the m.d.-cycle in C *right-adjacent* to c . On the other hand, a cycle c is called the *base* of $\tilde{c}$.

Proof. By definition C is a region, and hence, its boundary is a disjoint union of cycles of G . Since C consists of only level ℓ dual-vertices, each boundary cycle edge is an edge between a pair of dual-vertices of level either $\ell - 1$ and ℓ or ℓ and $\ell + 1$. Let us call the former one an *interior edge* and the latter one an *exterior edge*. Note that an interior edge must exist to define level ℓ dual-vertices of C , and that an exterior edge must exist because of the assumption of the lemma. We show that no boundary cycle has a vertex incident with both interior and exterior boundary edges. Suppose otherwise; that is, some vertex v exists that is incident with both interior and exterior boundary edges. Since G is three regular, this v is incident to three faces, one corresponding to a level ℓ dual-vertex (of C), one corresponding to a level $\ell - 1$ dual-vertex, and one corresponding to a level $\ell + 1$ dual-vertex. But this contradicts the definition of level. Thus, C , as a region, has at least the following two disjoint boundary cycles: an *interior boundary cycle* consisting of interior edges and an *exterior boundary cycle* consisting of exterior edges.

Now consider any one of such boundary cycles of C , and denote it by c . Then by the same argument as the proof of Lemma 4.20, we can define an m.d.-cycle $\tilde{c}$ consisting of all and only dual-vertices of C that share boundary edges with c . (We may assume that the direction of $\tilde{c}$ is determined so that vertices of c

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

are located its inside (Figure 4.23), which is different from core-cycles. From this difference, we need to consider the clockwise order when identifying an m.d.-cycle right-adjacent to a given interior/exterior boundary cycle.) $\square$




Figure 4.23: Interior and exterior boundary cycles and their right-adjacent m.d.-cycles

Based on this lemma, we introduce the notion of interior- and exterior-cycles. (Recall here that every m.d.-cycle is a collection of dual-cycles incident or connected by a line part.)

Definition 4.19 (Interior- and exterior-cycles). Let c be any interior (resp., exterior) boundary cycle of some connected component C of $\tilde{L}_{nb}(\ell)$. Let $\tilde{c}$ be the m.d.-cycle in C right-adjacent to c , and let $\tilde{D}_{\tilde{c}}$ be the set of dual-cycles obtained from $\tilde{c}$ by removing all its line parts. An *interior-cycle* (resp., *exterior-cycle*) is a dual-cycle in $\tilde{D}_{\tilde{c}}$.

Remark. Note that $\tilde{D}_{\tilde{c}}$ is defined for each boundary cycle. There may be more than one interior (resp., exterior) boundary cycles, but we focus on one boundary cycle c for defining $\tilde{c}$. We assume that each interior- and exterior-cycle is directed following its base m.d.-cycle $\tilde{c}$; the notions of inside/outside are the same as $\tilde{c}$ of Figure 4.23.

We show some basic properties of interior- and exterior-cycles. Below we say that two dual-cycles *have an overlap* if each has a dual-vertex of the other dual-cycle in its inside (or intuitively, they intersect at least at two dual-vertices).

Lemma 4.23. There is no overlap among interior-, exterior-, and core-cycles. The inside of each interior-cycle contains at least one boss-vertex.

Proof. Recall that every interior- or exterior-cycle is a subgraph of sliced graph $\tilde{L}_{nb}(\ell)$ for some ℓ , and that $\{\tilde{L}_{nb}(\ell)\}_{\ell \geq 1}$ is a partition of $\tilde{V} \setminus \cup_{\tilde{b} \in \tilde{I}} N_k(\tilde{b})$. On the other hand, every core-cycle is in $N_k(\tilde{b})$ for some $\tilde{b}$. Hence, in order to show that there is no overlap among interior-, exterior-, and core-cycles, it suffices to show that any two m.d.-cycles of level ℓ that could become interior- or exterior-cycles have no overlap. Suppose otherwise; that is, suppose that we have m.d.-cycles $\tilde{c}$ and $\tilde{c}'$ that are right-adjacent to some interior or exterior boundary cycles,

say, c and c' respectively of some connected component C of some sliced graph $\tilde{L}_{nb}(\ell)$. (Note that $\tilde{c}$ and $\tilde{c}'$ must be in the same connected component because they share some dual-vertex.) But by the construction of $\tilde{c}$ and $\tilde{c}'$, c and c' must have a common vertex, contradicting that all boundary cycles of any region are disjoint.

The second claim of the lemma follows from the following two facts: (i) every interior-cycle $\tilde{c}$ of level ℓ has a dual-vertex $\tilde{v}$ of level $\ell - 1$ in its inside, and (ii) there should be $N_k(\tilde{b})$ witnessing the level of $\tilde{v}$ and there should be a dual-path from $\tilde{v}$ to the boss-vertex $\tilde{b}$ that does not intersect with $\tilde{c}$. $\square$

Lemma 4.24. Consider any boss-vertex $\tilde{b} \in \tilde{I}$, dual-vertex $\tilde{u} \in \text{Vr}(\tilde{b})$, and the BFS dual-path $\tilde{p}$ from $\tilde{b}$ to $\tilde{u}$. Let $(\tilde{v}_1, \dots, \tilde{v}_t)$ be the tail subpath of $\tilde{p}$ where $\tilde{v}_t = \tilde{u}$, and $\tilde{v}_1$ is the first dual-vertex on $\tilde{p}$ that is not in $N_k(\tilde{b})$. For each i , $1 \leq i < t$, there exist an interior-cycle (resp., an exterior-cycle) that contains $\tilde{v}_i$ and that has $\tilde{b}$ in its one side and $\tilde{v}_t$ in the other side.

Proof. We fix any i , $1 \leq i < t$, and let C be a connected component of a sliced graph containing $\tilde{v}_i$, and let ℓ be the level of $\tilde{v}_i$ and C . We first show that $\tilde{v}_i$ appears in one of the exterior-cycles of C that satisfies the condition of the lemma. (The argument for the interior-cycles of C is similar.)

Note that the level of $\tilde{v}_{i+1}$ is $\ell + 1$; hence, an edge e_0 crossing the dual-edge $\{\tilde{v}_i, \tilde{v}_{i+1}\}$ is an edge of some exterior boundary cycle c of C . We focus on this cycle; see Figure 4.24. Consider the enumeration of dual-vertices adjacent to $\tilde{v}_i$ anti-clockwise from $\tilde{v}_{i+1}$ to $\tilde{v}_{i-1}$. Here we note that the level of the dual-vertices in the enumeration is either $\ell - 1$, ℓ , or $\ell + 1$, and that $(*)$ the levels of every consecutive pair of dual-vertices differs at most one because they are connected since $\tilde{G}$ is triangulated. Let $\tilde{x}_{j-1}$ is the last dual-vertex in the enumeration that is of level $\ell + 1$ and the edge e_k crossing $\tilde{v}_i$ and $\tilde{x}_{j-1}$ is an edge of the exterior boundary cycle c . Such $\tilde{x}_{j-1}$ exists because (i) there exists a dual-vertex of level $\ell + 1$ in the adjacent dual-vertex enumeration whose boundary edge between $\tilde{v}_i$ is an edge of c because at least we have $\tilde{v}_{j+1}$ that satisfies this condition, and (ii) there should be the last one from the fact $(*)$ stated above. Now from the property of $\tilde{x}_{j-1}$ and also from $(*)$ it follows that the next one $\tilde{x}_j$ in the adjacent dual-vertex enumeration is of level ℓ . Then the boundary edge e_{k+1} between $\tilde{x}_{j-1}$ and $\tilde{x}_j$ is a part of *some* exterior boundary cycle, which in fact must be the cycle c . This is because (i) both e_k and e_{k+1} must share the unique vertex w in the triangle $\tilde{v}_i, \tilde{x}_{j-1}$, and $\tilde{x}_j$ as their end point, and (ii) the third edge e' incident to w is not an exterior boundary edge.

Below let us denote $\tilde{x}_i$ simply by $\tilde{x}$. We can symmetrically show the existence of a dual-vertex $\tilde{y}$ with the same property in the clockwise enumeration of adjacent dual-vertices of $\tilde{v}_i$ from $\tilde{v}_{i+1}$ to $\tilde{v}_{i-1}$.

Consider the m.d.-cycle $\tilde{c}$ right adjacent to c . Note first that a directed dual-edge $(\tilde{x}, \tilde{v}_i)$ appears in $\tilde{c}$. This is because (i) a consecutive pair of boundary edges e_{k+1} and e_k are shared by $\tilde{x}$ and $\tilde{v}_i$ respectively, and (ii) the direction is consistent with the requirement that c is located inside of $\tilde{c}$. A similar argument shows that $(\tilde{v}_i, \tilde{y})$ appears in $\tilde{c}$. Our goal is to show that this pair of dual-edges

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM


$(\tilde{x}, \tilde{v}_i)$ and $(\tilde{v}_i, \tilde{y})$ appear in one of the dual-cycles of $\tilde{D}_{\tilde{c}}$, that is, they are part of some exterior-cycle $\tilde{c}'$ (see Definition 4.19 for $\tilde{D}_{\tilde{c}}$). This goal clearly suffices for the lemma because $\tilde{v}_{i-1}$ (and hence $\tilde{b}$) is in one side of $\tilde{c}'$ while $\tilde{v}_{i+1}$ (and hence $\tilde{t}$) is in the other side. (Note that $\tilde{v}_i$ is the unique point where the exterior-cycle $\tilde{c}'$ crosses the BFS dual-path from $\tilde{b}$ to $u = \tilde{v}_t$.)

For our goal, it suffices to show that $(*)$ after $(\tilde{v}_i, \tilde{y})$ in $\tilde{c}$, the dual-edge $(\tilde{x}, \tilde{v}_i)$ must be used when the m.d.-cycle $\tilde{c}$ comes back to $\tilde{v}_i$ next, that is, no other dual-edge to $\tilde{v}_i$ is used. Once this is proved, it is clear that the subsequence of $\tilde{c}$ from $(\tilde{v}_i, \tilde{y})$ to $(\tilde{x}, \tilde{v}_i)$ has a directed dual-cycle having both $(\tilde{v}_i, \tilde{y})$ and $(\tilde{x}, \tilde{v}_i)$ because $\tilde{v}_i$ appears in only these two dual-edges in this subsequence.

Now we show $(*)$. For this, consider the m.d.-cycle $\tilde{c}$ starting from $\tilde{v}_i$ with dual-edge $(\tilde{v}_i, \tilde{y})$. Also we consider an upper part (i.e., the part from $(\tilde{x}, \tilde{v}_i)$ to $(\tilde{v}_i, \tilde{y})$ anti-clockwise) and a lower part (i.e., the part from $(\tilde{v}_i, \tilde{y})$ to $(\tilde{x}, \tilde{v}_i)$ anti-clockwise), and show that no dual-edge to $\tilde{v}_i$ in these parts is used for coming back to $\tilde{v}_i$ for the first time. For the upper part, we can in fact easily show that there is no directed dual-edge of $\tilde{c}$ to $\tilde{v}_i$ in this part. Suppose otherwise and that some directed dual-edge of $\tilde{c}$ coming into $\tilde{v}_i$. This should come from a dual-vertex adjacent to $\tilde{v}_i$ that is of level ℓ sharing an edge of c that is a boundary edge to a dual-vertex of level $\ell + 1$. Clearly, this contradicts to the choice of $\tilde{x}$ and $\tilde{y}$. Consider next the lower part. Suppose that the m.d.-cycle $\tilde{c}$, after $(\tilde{v}_i, \tilde{y})$, comes back to $\tilde{v}_i$ for the first time from the lower part by some dual-edge $(\tilde{z}, \tilde{v}_i)$. Then after $(\tilde{z}, \tilde{v}_i)$, all dual-edges going out from $\tilde{v}_i$ or coming back to $\tilde{v}_i$ must be located between $(\tilde{z}, \tilde{v}_i)$ and $(\tilde{v}_i, \tilde{y})$ clockwise because the m.d.-cycle $\tilde{c}$ is defined under the clockwise order; otherwise, some subsequence of $\tilde{c}$ must cross the subsequence of $\tilde{c}$ from $\tilde{y}$ to $\tilde{z}$, contradicting the noncrossing property of m.d.-cycle (Figure 4.10). Then $(\tilde{x}, \tilde{v}_i)$ cannot appear in $\tilde{c}$, a contradiction. $\square$

Finally, we define the notion of neck, floor, and ceiling. Consider any interior- or exterior-cycle. It is called a *neck* if it has at most $\sqrt{k}$ dual-vertices, and an interior (resp., exterior) neck is a neck interior-cycle (resp., neck exterior-cycle). A neck is called *biased* if it has less than $\tilde{n}/3$ dual-vertices in its inside. Note that if there is any neck that has at least $\tilde{n}/3$ dual-vertices in both inside and outside, then we can use the neck itself as a $1/4$ -separator of $\tilde{G}$ (when n is sufficiently large) from which we can construct a target separator for G_{org} . Therefore, in the following discussion, we assume that a neck is biased or it has less than $\tilde{n}/3$ dual-vertices in its outside.

Roughly speaking, a *floor-cycle* (resp., *ceiling-cycle*) is a biased neck interior-cycle (resp., exterior-cycle). But some more conditions are needed. As explained in Outline section, we modify the frame graph $\tilde{H}$ by covering each boss-vertex by a floor-cycle and a branch-triangle by a ceiling-cycle in order for reducing the length of each dual-path from a boss-vertex to (a dual-vertex of) a branch-triangle. Thus, we only need the “largest” one covering each boss-vertex or branch-triangle. Also we do not need an exterior-cycle that has no intersection with any branch-triangle as a ceiling-cycle. On the other hand, it may be the case where some boss-vertex is contained in no floor-cycle. For such



A bold line (with black nodes) indicates a part of the exterior boundary cycle c focused in the proof. In the proof, for the anti-clockwise enumeration of dual-vertices adjacent to $\tilde{v}_i$ from $\tilde{v}_{i+1}$ to $\tilde{v}_{i-1}$, we consider $\tilde{x}_{j-1}$, the last one in the enumeration that is of level $\ell + 1$ and whose boundary edge, i.e., e_k , is a part of c . We then show that directed dual-edges $(\tilde{x}, \tilde{v}_i)$ and $(\tilde{v}_i, \tilde{y})$ are part of the m.d.-cycle $\tilde{c}$ right adjacent to c , and in fact they are part of the same dual-cycle component of $\tilde{c}$, i.e., an exterior-cycle defined from $\tilde{c}$.

Figure 4.24: An example for the proof of Lemma 4.24

boss-vertices, we would use their core-cycles as floor-cycles. From these considerations, we define the notion of floor- and ceiling-cycles as follows.

Definition 4.20 (Floor- and ceiling-cycles). A *floor-cycle* is a biased neck interior-cycle that is not contained in the inside of any other biased neck interior-cycle. For any boss-vertex $\tilde{b}$ that is contained in no floor-cycle, we regard the core-cycle of $\text{Core}(\tilde{b})$ also as a *floor-cycle*. A *ceiling-cycle* is a biased neck exterior-cycle that is not contained in the inside of any other biased neck cycle-cycle and that has at least one dual-vertex of some branch-triangle.

Remark. When a floor-cycle (resp., ceiling-cycle) is a special case of interior-cycles (resp., exterior-cycles), its inside/outside is the same as the corresponding interior-cycles (resp., exterior-cycles). Similarly, for a floor-cycle defined as a core-cycle, its inside/outside follow those of the core-cycle.

We show that floor- and ceiling-cycles have desired properties. First we state the following size bounds, which are immediate from the definition.

Lemma 4.25. Each floor- and ceiling-cycle has at most $\sqrt{k}$ dual-vertices and less than $\tilde{n}/3$ dual-vertices in its inside.

Next one is the key property of floor- and ceiling-cycles, which is the main reason of introducing these structures. (One can find a similar argument in the proof of, e.g., Lemma 3.3 of [33].)

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

Lemma 4.26. Consider any pre-frame-cycle (after the preprocessing), and let $\tilde{p}$ be any dual-path between its boss-vertex $\tilde{b}$ and the “next” dual-vertex $\tilde{u}$ of a branch-triangle used in the pre-frame-cycle¹⁴. There exist one floor-cycle $\tilde{c}_f$ and at most one ceiling-cycle $\tilde{c}_c$ crossing $\tilde{p}$. Let $\tilde{u}_f$ and $\tilde{u}_c$ denote dual-vertices at the crossing point of respectively $\tilde{p}$ and $\tilde{c}_f$, and $\tilde{p}$ and $\tilde{c}_c$. Then $\tilde{p}$ is separated into at most three parts $\tilde{p}_1, \tilde{p}_2$, and $\tilde{p}_3$ by $\tilde{c}_f$ and $\tilde{c}_c$: that is, dual-path $\tilde{p}_1$ between $\tilde{b}$ and $\tilde{u}_f$, dual-path $\tilde{p}_3$ between $\tilde{u}$ and $\tilde{u}_c$, and the remaining dual-path $\tilde{p}_2$. (In the case where $\tilde{c}_c$ does not exist, $\tilde{p}_2$ is simply a dual-path defined by $\tilde{p} \setminus \tilde{p}_1$.) Furthermore, the length of $\tilde{p}_2$ is less than $4\sqrt{k}$.

Proof. Recall that $\tilde{p}$ is a dual-path of the BFS dual-tree from $\tilde{b}$. The existence of the unique floor-cycle $\tilde{c}_f$ and at most one ceiling-cycle $\tilde{c}_c$ follows from Lemmas 4.23 and 4.24, and Definition 4.20. Furthermore, it is easy to see that all dual-vertices of $\tilde{p}$ between $\tilde{b}$ and $\tilde{u}_f$ (except $\tilde{u}_f$ itself) are inside of the floor-cycle $\tilde{c}_f$ and that all dual-vertices of $\tilde{p}$ between $\tilde{u}_c$ (except $\tilde{u}_c$ itself) and $\tilde{u}$ (if they exist) are inside of the ceiling-cycle $\tilde{c}_c$. Thus, $\tilde{p}$ is divided into at most three parts $\tilde{p}_1, \tilde{p}_2$, and $\tilde{p}_3$ as stated in the lemma, and $\tilde{p}_2$ is the part of $\tilde{p}$ located outside of both $\tilde{c}_f$ and $\tilde{c}_c$.

We show that the length of $\tilde{p}_2$ is less than $4\sqrt{k}$. For our discussion, let us regard $\tilde{p}_2$ as a *directed* dual-path based on the distance from $\tilde{b}$. Consider first the head part of $\tilde{p}_2$ consisting of dual-vertices in $N_k(\tilde{b})$. This part exists only if the floor-cycle $\tilde{c}_f$ is defined by a core-cycle, and the length of this part is by definition at most $d_{nb}(\tilde{b}) - d_{core}(\tilde{b})$ that is bounded by $\sqrt{k}$ (Lemma 4.19). Thus, the length of this head part of $\tilde{p}_2$ is either 0 or at most $\sqrt{k}$.

Now consider the remaining tail part of $\tilde{p}_2$, and let $\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_t$ be the enumeration of dual-vertices of this part following the direction of $\tilde{p}_2$. For the lemma, it suffices to show that the length of this part of $\tilde{p}_2$, i.e., t is less than $3\sqrt{k}$. Assume to the contrary that $t \geq 3\sqrt{k}$.

Considering the fact that $\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_t$ are all in $Vr(\tilde{b}) \setminus N_k(\tilde{b})$ and also on one dual-path of the BFS dual-tree from $\tilde{b}$, we have $\tilde{v}_i \in \tilde{L}_{nb}(\ell_1 + (i - 1))$ for each i , $1 \leq i \leq t$ where ℓ_1 is the level of $\tilde{v}_1$, which fact will be referred as (*) later.

Consider any i , $1 \leq i \leq t$. Lemma 4.24 guarantees the existence of an interior-cycle $\tilde{c}_i^{in}$ and an exterior-cycle $\tilde{c}_i^{ex}$ containing $\tilde{v}_i$ (where $\tilde{c}_i^{in}$ and $\tilde{c}_i^{ex}$ might be the same dual-cycle). We argue that either $\tilde{c}_i^{in}$ or $\tilde{c}_i^{ex}$ is not a neck; that is, $|\tilde{c}_i^{in}| > \sqrt{k}$ or $|\tilde{c}_i^{ex}| > \sqrt{k}$. Note first that $\tilde{c}_i^{in}$ (resp., $\tilde{c}_i^{ex}$) contains $\tilde{c}_f$ (resp., $\tilde{c}_c$ if it exists) in its inside (Lemmas 4.23 and 4.24). Hence, neither $\tilde{c}_i^{in}$ nor $\tilde{c}_i^{ex}$ is a biased neck (Definition 4.20). Next assume that both $\tilde{c}_i^{in}$ and $\tilde{c}_i^{ex}$ are necks. Since $\tilde{c}_i^{in}$ is not biased, $\tilde{c}_i^{in}$ would have less than $\tilde{n}/3$ dual-vertices in its outside. On the other hand, the inside of $\tilde{c}_i^{ex}$ is contained in the outside of $\tilde{c}_i^{in}$ (Figure 4.23). Hence, the inside of $\tilde{c}_i^{ex}$ would have less than $\tilde{n}/3$ dual-vertices, that is, $\tilde{c}_i^{ex}$ is biased; a contradiction. Therefore, either $\tilde{c}_i^{in}$ or $\tilde{c}_i^{ex}$ is not a neck; let $\tilde{c}_i$ denote this non-neck dual-cycle. We note here that $|\tilde{c}_i| > \sqrt{k}$.

¹⁴More precisely, the “next” dual-vertex $\tilde{u}$ is the first/last dual-vertex of a branch-triangle from/to the boss-vertex appearing in the frame m.d.-cycle following its direction.

4.7. FLOOR AND CEILING MODIFICATION

Now consider any i , $1.5\sqrt{k} \leq i \leq 2.5\sqrt{k}$, and let $\tilde{C}_i$ be a set of $\sqrt{k} + 1$ dual-vertices of $\tilde{c}_i$ that are nearest to $\tilde{v}_i$ (including $\tilde{v}_i$ itself). (Here we can use the fact that $|\tilde{c}_i| > \sqrt{k}$ to ensure the existence of $\tilde{C}_i$.) Let $\tilde{C} = \cup_{1.5\sqrt{k} \leq i \leq 2.5\sqrt{k}} \tilde{C}_i$. Then we have $|\tilde{C}| > k$. Note also that the distance between any dual-vertex of $\tilde{C}_i$ and $\tilde{v}_i$ is at most $\sqrt{k}$. (Recall that $\tilde{c}_i$ is a dual-cycle.) Here consider a set $\tilde{B}$ of dual-vertices whose distance from $\tilde{v}_{2\sqrt{k}}$ is at most $1.5\sqrt{k}$. Then we have $\tilde{B} \supset \tilde{C}$, and hence we have $|\tilde{B}| > k$. Note here that $N_k^+(\tilde{b})$ is the nearest k -neighborhood⁺ of $\tilde{v}_{2\sqrt{k}}$ and the distance between $\tilde{v}_{2\sqrt{k}}$ and $N_k^+(\tilde{b})$ is at least $2\sqrt{k}$ (from $(*)$ mentioned above). Hence, no dual-vertex of $\tilde{B}$ belongs to $N_k^+(\tilde{b}')$ for any $\tilde{b}' \in \tilde{I}$; because otherwise, the distance of $\tilde{v}_{2\sqrt{k}}$ from $N_k^+(\tilde{b}')$ would be at most $1.5\sqrt{k}$, contradicting that $\tilde{b}$ is the boss-vertex of $\tilde{v}_{2\sqrt{k}}$. That is, $N_k^+(\tilde{v}_{2\sqrt{k}})$ has no intersection with $\cup_{\tilde{b}' \in \tilde{I}} N_k^+(\tilde{b}')$, contradicting the choice of $\tilde{I}$. $\square$

Finally, we define the notion of “modified frame-graph” for our final graph to which Miller’s algorithm is applicable.

Definition 4.21 (Modified frame-graph). Let $\tilde{H}$ be the frame-graph defined from $\tilde{G}$ in the previous section. Let $\tilde{F}$ and $\tilde{C}$ be respectively a set of floor-cycles and ceiling-cycles having at least one dual-vertex of $\tilde{H}$ in their insides. We let $\tilde{E}'_1$ be the set of dual-edges that appear in some cycle in $\tilde{F} \cup \tilde{C}$ and $\tilde{E}'_2$ be the set of dual-edges of $\tilde{H}$ that are not in the inside of any cycle of $\tilde{F} \cup \tilde{C}$. A graph $\tilde{H}' = (\tilde{U}', \tilde{D}')$ is a *modified frame-graph*, where $\tilde{D}' = \tilde{E}'_1 \cup \tilde{E}'_2$ and $\tilde{U}'$ is the set of all dual-vertices that are end points of dual-edges of $\tilde{D}'$.

Note that faces of $\tilde{H}'$ are defined by one of the following boundary dual-cycles: (i) a floor-cycle, (ii) a ceiling-cycle, (iii) a branch-triangle, or (iv) a frame-cycle of $\tilde{H}$ modified by floor- and ceiling-cycles. We call these boundary dual-cycles $\tilde{H}'$ -*face-cycles*. Based on this observation, we define weights of faces of $\tilde{H}'$ in the same way as $\tilde{H}$.

Definition 4.22. For each face of a modified frame-graph $\tilde{H}'$, its weight is the number of dual-vertices of $\tilde{G}$ located in the face (i.e., in the inside of its $\tilde{H}'$ -face-cycle) divided by the total number of dual-vertices removed from $\tilde{G}$.

We show that the modified frame graph $\tilde{H}'$ satisfies the conditions (F1) $\sim$ (F4).

First the condition (F1) follows from the definition of a frame-graph and modified frame-graph. In particular, the condition that the weight of each face is less than $1/3$ follows from the fact that $\tilde{H}$ satisfies this condition. Next consider the condition (F3); that is, the number of faces is bounded by $O(\tilde{n}/k)$. Again this follows easily from the fact that $\tilde{H}$ satisfies this condition. As mentioned above, new faces introduced by our modification are those defined by floor- or ceiling-cycles. By definition, the number of these dual-cycles are bounded by either the number of boss-vertices or that of (dual-vertices of) branch-triangles, which is bounded by $O(\tilde{n}/k)$. Note that a face defined a frame-cycle of $\tilde{H}$ could

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

be divided by ceiling-cycles; but it is easy to see that each face is divided at most some constant number of faces because the number of floor- and ceiling-cycles overlapping each frame-cycle is constant, say, at most six. From these observations, we can bound the number of faces of $\tilde{H}'$ by $O(\tilde{n}/k)$.

We argue that $\tilde{H}'$ keeps the condition (F2) as follows.

Lemma 4.27. $\tilde{H}'$ is 2-connected.

Proof. For any two distinct dual-vertices $\tilde{u}$ and $\tilde{v}$ in $\tilde{U}'$, we show that there exist two vertex disjoint dual-paths between $\tilde{u}$ and $\tilde{v}$. If $\tilde{u}$ and $\tilde{v}$ are located on the same $\tilde{H}'$ -face-cycle, it is obvious. Assume that $\tilde{u}$ and $\tilde{v}$ are in $\tilde{H}$ and they survived the modification; namely, they are on branch-triangles or frame-cycles of $\tilde{H}$ remained after the modification of Definition 4.22. There are two vertex disjoint dual-paths $\tilde{p}_L$ and $\tilde{p}_R$ in $\tilde{H}$ since $\tilde{H}$ is 2-connected (Lemma 4.13). In $\tilde{H}'$, some floor- and ceiling-cycles are introduced. When some of these dual-cycles touches only one of $\tilde{p}_L$ and $\tilde{p}_R$, then it is obvious that two disjoint dual-paths exist in $\tilde{H}'$. Suppose that a floor- or ceiling-cycle $\tilde{c}$ touches both $\tilde{p}_L$ and $\tilde{p}_R$. Let $\tilde{u}_L$ and $\tilde{v}_L$ (resp., $\tilde{u}_R$ and $\tilde{v}_R$) respectively denote the first and the last dual-vertices on $\tilde{p}_L$ (resp., $\tilde{p}_R$) on which $\tilde{p}_L$ (resp., $\tilde{p}_R$) intersects $\tilde{c}$. These four dual-vertices $\tilde{u}_L$, $\tilde{v}_L$, $\tilde{u}_R$ and $\tilde{v}_R$ divide $\tilde{c}$ into four subsequences. We can get two disjoint subsequences from $\tilde{u}_L$ to $\tilde{v}_L$ or $\tilde{v}_R$ and from $\tilde{u}_R$ to $\tilde{v}_R$ or $\tilde{v}_L$. If otherwise; for example, if $\tilde{u}_L$ is adjacent to $\tilde{u}_R$ and $\tilde{v}_L$, and $\tilde{u}_R$ is adjacent to $\tilde{u}_L$ and $\tilde{v}_L$ on $\tilde{c}$, then $\tilde{v}_R$ never appears on $\tilde{c}$, a contradiction. Therefore, there exist two disjoint dual-paths consisting of three subpaths; namely, the subpaths from $\tilde{u}$ to $\tilde{u}_L$ (resp., $\tilde{u}_R$), the disjoint subsequences of $\tilde{c}$ from $\tilde{u}_L$ (resp., $\tilde{u}_R$) to $\tilde{v}_L$ or $\tilde{v}_R$, and the subpath from $\tilde{v}_L$ (resp., $\tilde{v}_R$) to $\tilde{v}$.

Finally, assume that $\tilde{u}$ and $\tilde{v}$ are on different floor- or ceiling-cycles $\tilde{c}_{\tilde{u}}$ and $\tilde{c}_{\tilde{v}}$. From the definition, each of $\tilde{c}_{\tilde{u}}$ and $\tilde{c}_{\tilde{v}}$ has at least one dual-vertex of $\tilde{H}$ in its inside; let $\tilde{w}_{\tilde{u}}$ and $\tilde{w}_{\tilde{v}}$ denote them. There exist two disjoint dual-paths $\tilde{p}_L$ and $\tilde{p}_R$ between $\tilde{w}_{\tilde{u}}$ and $\tilde{w}_{\tilde{v}}$ in $\tilde{H}$. Note that $\tilde{p}_L$ and $\tilde{p}_R$ have distinct intersections with $\tilde{c}_{\tilde{u}}$; let $\tilde{u}_L$ (resp., $\tilde{u}_R$) be the last intersection of $\tilde{p}_L$ (resp., $\tilde{p}_R$) and $\tilde{c}_{\tilde{u}}$ from $\tilde{w}_{\tilde{u}}$ to $\tilde{w}_{\tilde{v}}$. Similarly, let $\tilde{v}_L$ (resp., $\tilde{v}_R$) be the first intersection of $\tilde{p}_L$ (resp., $\tilde{p}_R$) and $\tilde{c}_{\tilde{v}}$. Clearly, there exist two disjoint dual-paths from $\tilde{u}$ to $\tilde{u}_L$ and $\tilde{u}_R$ on $\tilde{c}_{\tilde{u}}$; one can find them by going along $\tilde{c}_{\tilde{u}}$ clockwise and anti-clockwise. Similarly, we have two disjoint dual-paths from $\tilde{v}_L$ and $\tilde{v}_R$ to $\tilde{v}$. By combining them, there exist two disjoint dual-paths from $\tilde{u}$ to $\tilde{v}$. Note that there may be some floor- or ceiling-cycle touching these disjoint dual-paths. Even in this case, by an argument similar to the above, we can show two disjoint dual-paths between $\tilde{u}$ and $\tilde{v}$. $\square$

Finally, we confirm as a corollary of Lemma 4.25 that the condition (F4) holds.

Corollary 4.28. The size of each face of $\tilde{H}'$ is $O(\sqrt{k})$.

Proof. As mentioned above, the faces of $\tilde{H}'$ are defined by boundary dual-cycles of type (i) $\sim$ (iv); see the comment after Definition 4.22. For those defined by

floor-, ceiling-cycles, and branch-triangles, the size bound of the lemma clearly holds by definition. Thus, we below consider faces defined by modified frame-cycles.

Consider any face defined by a modified frame-cycle, and let $(\tilde{c})^-$ and $\tilde{c}$ denote respectively the original frame-cycle and the pre-frame-cycle from which $(\tilde{c})^-$ is defined (Definition 4.15). Consider any dual-path $\tilde{p}$ of $\tilde{c}$ connecting a boss-vertex of $\tilde{c}$ and its “next” dual-vertex of a branch-triangle used in $\tilde{c}$. By our modification, a part $\tilde{p}'$ of $\tilde{p}$ that is in the outside of the corresponding floor-cycle and ceiling-cycle (if it exists) could be used as a component of the modified frame-cycle, and its length is bounded by $4\sqrt{k}$ by Lemma 4.26. Note that the floor-cycle may not be used in $\tilde{H}'$ if it only intersects with the line part of $\tilde{c}$ that is removed for defining $(\tilde{c})^-$ (e.g., Figure 4.19 (a)). In this case, however, only a part of $\tilde{p}'$ is used for the modified frame-cycle, which is even shorter. Thus, the modified frame-cycle consists of (at most) four such reduced dual-paths, (a part of) two floor-cycles, (a part of) four ceiling-cycles (i.e., could be two for each branch-triangle), and two dual-edges from two branch-triangles, and their total length is $O(\sqrt{k})$. $\square$

Algorithms for Step 2.4

The task of this step is to compute information of the modified frame-graph $\tilde{H}'$. We again explain as if each step receives the output of the previous step as input, and omit analysis of time complexity. Note that we can use the outputs of Step 2.3, i.e., the frame-graph $\tilde{H}$, as input. We consider an algorithm that outputs

- (a) a list of core-cycles;
- (b) a list of interior- and exterior-cycles;
- (c) a list of floor- and ceiling-cycles; and
- (d) the complete weighted face information of $\tilde{H}'$.

First we point out that a cycle, a sequence of vertices representing a (directed) cycle, is $\tilde{O}(1)$ -space computable provided we have a way to distinguish a set of edges of the cycle (*and*, more precisely, it is guaranteed that this set of edges indeed forms a cycle). Also for a given cycle c of G , the task of computing an m.d.-cycle consisting of dual-vertices left/right adjacent to c can be done locally and hence in $\tilde{O}(1)$ -space; see the proof of Lemma 4.20 and Lemma 4.22.

Once these algorithmic techniques are clear, what remains to show here for computing the above (a) and (b) is a way to identify core/interior/exterior boundary cycle edges.

For (a), we first compute all boundary cycles of $\text{Core}(\tilde{b})$ for all $\tilde{b} \in \tilde{I}$. For any $\tilde{b} \in \tilde{I}$, consider $\text{Core}(\tilde{b})$. Note that it is a subset of $N_k(\tilde{b})$; thus, we have an $\tilde{O}(k)$ -space algorithm that determines whether a given dual-vertex belongs to $\text{Core}(\tilde{b})$ or not by using the same BFS algorithm of Step 2.1. Then by using this

4. A SUBLINEAR-SPACE AND POLYNOMIAL-TIME PLANAR SEPARATOR ALGORITHM

procedure, we can identify edges of each boundary cycle of $\text{Core}(\tilde{b})$; basically, what we need is to check, for a given candidate edge of G , whether one of its two adjacent dual-vertices belongs to $\text{Core}(\tilde{b})$ and the other does not. Next we select *core* boundary cycles from all enumerated boundary cycles. For a given boundary cycle c , we can use Reingold's algorithm to determine whether a given dual-vertex is in the outside of c ; hence, we can count the number of dual-vertices outside of c and determine whether c is a core boundary cycle. Thus, we can enumerate all core boundary cycle, from which we can enumerate all core-cycle as the output (a).

For (b), i.e., interior- and exterior-cycles, we first compute the level of all dual-vertices of $\tilde{G}$. For each dual-vertex $\tilde{v}$ and its boss-vertex $\tilde{b}$, we have an algorithm that computes $\text{dist}(\tilde{v}, \tilde{b})$ in $\tilde{O}(k)$ -space (Lemma 4.8). The level of all dual-vertices is computable by this procedure. Based on this level information, we can identify connected components of dual-vertices with the same level (outside of $\cup_{\tilde{b} \in \tilde{I}} \tilde{N}_k(\tilde{b})$). Again this procedure is enough for enumerating interior and exterior boundary cycles, and m.d.-cycles right-adjacent to them. Once we obtain a list of all m.d.-cycles right-adjacent to interior or exterior boundary cycles, the remaining task for computing the output (b) is to remove all line parts from these m.d.-cycles and decompose them into dual-cycles, which can be done in $\tilde{O}(1)$ -space.

By using the core-, interior- and exterior-cycles obtained above as input, it is easy to list floor- and ceiling-cycles; that is, the output (c). We compute the size of each cycle and check whether it is a neck, and count the number of dual-vertices in its inside and check whether it is biased. By these computations, biased neck interior- and exterior-cycles are obtained in $\tilde{O}(1)$ -space. Furthermore, we can check inclusion relations of these cycles and branch-triangles in $\tilde{O}(1)$ -space, from which we can obtain a list of floor- and ceiling-cycles, i.e., the output (c).

Now we can use $\tilde{H}$ and the list of floor- and ceiling-cycles as input for computing the output (d). For this, we start with the list of frame-cycles (of $\tilde{H}$) as an initial list of $\tilde{H}'$ -face-cycles. For each floor- and ceiling-cycle, we check whether the dual-cycle has at least one dual-vertex of $\tilde{H}$; if it does, then it is added to the list, and the corresponding frame-cycles of $\tilde{H}$ in the list are modified appropriately (if necessary). In this way we can compute a list of all $\tilde{H}'$ -face-cycles in $\tilde{O}(1)$ -space. Once we have the list, the rest of the task is similar to the previous sub-steps.

4.7. FLOOR AND CEILING MODIFICATION

Chapter 5

$\tilde{O}(n^{1/3})$ -Space Algorithm for the Grid Graph Reachability Problem

The directed graph reachability problem takes as input an n -vertex directed graph $G = (V, E)$, and two distinguished vertices s and t . The problem is to determine whether there exists a path from s to t in G . This is a canonical complete problem for class NL. The main result of this chapter is to show that the directed graph reachability problem restricted to grid graphs can be solved in polynomial time using only $\tilde{O}(n^{1/3})$ space [12].

5.1 Outline of the algorithm

Asano et al. proposed an $\tilde{O}(\sqrt{n})$ -space and polynomial time algorithm for the directed grid and planar graph reachability problem. Our algorithm uses this algorithm as a subroutine.

Theorem 5.1 ([9]). There exists an algorithm that decides directed planar graph reachability in polynomial time and $\tilde{O}(\sqrt{n})$ space. (We refer to this algorithm by `PlanarReach`.)

Recall that a grid graph is a graph whose vertices are located on grid points, and whose vertices are adjacent only to their immediate horizontal or vertical neighbors. We refer to a vertex on the boundary of a grid graph as a *rim vertex*. For any grid graph G , we denote the set of the rim vertices of G as R_G . For any $u, v \in V$, a directed edge e from u to v is denoted as $e = (u, v)$; on the other hand, the tail u and the head v of e are denoted as $t(e)$ and $h(e)$, respectively.

We show the outline of our algorithm. We assume both $\sqrt{n}$ and $n^{1/3}$ are integers for simplicity. Let G be an input $\sqrt{n} \times \sqrt{n}$ grid graph with n vertices.

1. Separate G into $n^{1/3} \times n^{1/3}$ small grid graphs, or “blocks”. There are $n^{1/3}$ blocks, and each block contains $n^{2/3}$ vertices.
2. Transform each block B into a special planar graph, “gadget graph”, with $O(n^{1/3})$ vertices. The reachability among the vertices in R_B should be unchanged. The total number of vertices in all blocks becomes $O(n^{2/3})$.
3. We apply the algorithm **PlanarReach** to the transformed graph of size $O(n^{2/3})$, then the reachability is computable in $\tilde{O}\left(\sqrt{n^{2/3}}\right) = \tilde{O}(n^{1/3})$ space.

In step 1 and 2, we reduce the number of vertices in the graph G while keeping the reachability between the rim vertices of each block so that we can solve the reachability problem of the original graph. Then to this transformed graph we apply **PlanarReach** in step 3, which runs in $\tilde{O}(n^{1/3})$ space.

Theorem 5.2 ([12]). There exists an algorithm that computes the grid graph reachability in polynomial-time and $\tilde{O}(n^{1/3})$ space.

The start vertex s (resp., the end vertex t) may not be on the rim of any block. In such a situation, we make an additional block so that s (resp., t) would be on the rim of the block. This operation would not increase the time and space complexity. We assume that s (resp., t) is on the rim of some block.

5.2 Graph transformation

In this section, we explain an algorithm that modifies each block and analyze time and space complexity of the algorithm. Throughout this section, we let a directed graph $G_0 = (V_0, E_0)$ denote a block of the input grid graph, and let V_0^{rim} denote the set of its rim vertices. We use N to denote the number of vertices of the input grid graph and n to denote $|V_0^{\text{rim}}|$, which is $O(N^{1/3})$; note, on the other hand, that we have $|V_0| = O(n^2) = O(N^{2/3})$. Our task is to transform this G_0 to a plane “gadget graph”, an augmented plane graph, $\tilde{G}_p$ with $O(n) = O(N^{1/3})$ vertices including V_0^{rim} so that the reachability among vertices in V_0^{rim} on G_0 remains the same on $\tilde{G}_p$.

There are two steps for this transformation. We first transform G_0 to a circle graph G_0^{cir} , and then obtain $\tilde{G}_p$ from the circle graph.

5.2.1 Circle graph

We introduce the notion of “circle graph”. A *circle graph* is a graph embedded on the plane so that all its vertices are placed on a cycle and all its edges are drawn inside of the cycle. Note that a circle graph may not have an edge between a pair of adjacent vertices on the cycle. We introduce some basic notions on circle graphs. Consider any circle graph $G = (V, E)$, and let C be a cycle on which all vertices of V are placed. For any $u, v \in V$, a *clockwise tour* (resp., *anti-clockwise tour*) is a part of the cycle C from u to v in a clockwise direction (resp., in an




Figure 5.1: An example of the notions on chords. (a) a figure showing a chord, arcs, a lower area, an upper area, (b) a figure showing crossing chords (e_1 and e_2) and semi-crossing chords (e_3 and e_4) and (c) separating chords (e_3 separates e_1 and e_2).

anti-clockwise direction). We use $C^{\text{cl}}[u, v]$ (resp., $C^{\text{acl}}[u, v]$) to denote this tour (Figure 5.1(a)). When we would like to specify the graph G , we use $C_G^{\text{cl}}[u, v]$ (resp., $C_G^{\text{acl}}[u, v]$). The tour $C^{\text{cl}}[u, v]$, for example, can be expressed canonically as a sequence of vertices $(v_1, \dots, v_k)$ such that $v_1 = u$, $v_k = v$, and $v_2, \dots, v_{k-1}$ are all vertices visited along the cycle C clockwise. We use $C^{\text{cl}}(u, v)$ and $C^{\text{cl}}[u, v]$ (resp., $C^{\text{acl}}(u, v)$ and $C^{\text{acl}}[u, v]$) to denote the sub-sequences $(v_2, \dots, v_{k-1})$ and $(v_1, \dots, v_{k-1})$ respectively. Note here that it is not necessary that G has an edge between adjacent vertices in such a tour. The length of the tour is simply the number of vertices on the tour. An edge (u, v) of G is called a *chord* if u and v are not adjacent on the cycle C . For any chord (u, v) , we may consider two arcs, namely, $C^{\text{cl}}[u, v]$ and $C^{\text{acl}}[u, v]$; but in the following, we will simply use $C[u, v]$ to denote one of them that is regarded as the arc of the chord (u, v) in the context. When necessary, we will state, e.g., “the arc $C^{\text{cl}}[u, v]$ ” for specifying which one is currently regarded as the arc. A *gap- d* (resp., *gap- d^+*) *chord* is a chord (u, v) whose arc $C[u, v]$ is of length $d + 2$ (resp., length $\geq d + 2$). For any chord (u, v) , the subplane inside of the cycle C surrounded by the chord (u, v) and the arc $C[u, v]$ is called the *lower area* of the chord; on the other hand, the other side of the chord within the cycle C is called the *upper area* (see Figure 5.1(a)). A *lowest gap- d^+* chord is a gap- d^+ chord that has no other gap- d^+ chord in its lower area. We say that two chords (u_1, v_1) and (u_2, v_2) *cross* if they cross in the circle C in a natural way (see Figure 5.1(b)). Formally, we say that (u_1, v_1) *crosses* (u_2, v_2) if either (i) u_2 is on the tour $C^{\text{cl}}(u_1, v_1)$ and v_2 is on the tour $C^{\text{acl}}(u_1, v_1)$, or (ii) v_2 is on the tour $C^{\text{cl}}(u_1, v_1)$ and u_2 is on the tour $C^{\text{acl}}(u_1, v_1)$. Also, we say that (u_1, v_1) *semi-crosses* (u_2, v_2) if either (i) u_2 is on the tour $C^{\text{cl}}[u_1, v_1]$ and v_2 is on the tour $C^{\text{acl}}[u_1, v_1]$, or (ii) v_2 is on the tour $C^{\text{cl}}[u_1, v_1]$ and u_2 is on the tour $C^{\text{acl}}[u_1, v_1]$ (see Figure 5.1(b)). Note that clearly crossing implies semi-crossing. In addition, we say that a chord (u_1, v_1)

separates two chords (u_2, v_2) and (u_3, v_3) if the endpoints of two chords v_2 and v_3 are separated by the chord (u_1, v_1) (see Figure 5.1(c)). Formally, (u_1, v_1) *separates* (u_2, v_2) and (u_3, v_3) if either (i) v_2 is on the tour $C^{\text{cl}}[u_1, v_1]$ and v_3 is on the tour $C^{\text{acl}}[u_1, v_1]$, or (ii) v_3 is on the tour $C^{\text{cl}}[u_1, v_1]$ and v_2 is on the tour $C^{\text{acl}}[u_1, v_1]$. We say that k chords $(u_1, v_1), (u_2, v_2), \dots, (u_k, v_k)$ are *traversable* if the following two conditions are satisfied:

1. (u_1, v_1) semi-crosses (u_2, v_2) ,
2. $\forall i \in [3, k], \exists p, q < i, (u_i, v_i)$ separates (u_p, v_p) and (u_q, v_q) .

Now for the graph $G_0 = (V_0, E_0)$, we define the circle graph $G_0^{\text{cir}} = (V_0^{\text{cir}}, E_0^{\text{cir}})$ by

$$\begin{aligned} V_0^{\text{cir}} &= V_0^{\text{rim}}, \text{ and} \\ E_0^{\text{cir}} &= \{ (u, v) \mid \exists \text{ path from } u \text{ to } v \text{ in } G_0 \}, \end{aligned}$$

where we assume that the rim vertices of $V_0^{\text{cir}} (= V_0^{\text{rim}})$ are placed on a cycle C_0 as they are on the rim of the block in the grid graph. Then it is clear that G_0^{cir} keeps the same reachability relation among vertices in $V_0^{\text{cir}} = V_0^{\text{rim}}$. Recall that G_0 has $O(n^2)$ vertices. Thus, by using `PlanarReach`, we can show the following lemma.

Lemma 5.3. G_0^{cir} keeps the same reachability relation among vertices in $V_0^{\text{cir}} = V_0^{\text{rim}}$. That is, for any pair u, v of vertices of V_0^{cir} , v is reachable from u in G_0^{cir} if and only if it is reachable from u in G_0 . There exists an algorithm that transforms G_0 to G_0^{cir} in $O(n)$ -space and polynomial-time in n .

The notion of traversable is a key for discussing the reachability on G_0^{cir} . Based on the following lemma, we use a traversable sequence of edges for characterizing the reachability on the circle graph G_0^{cir} .

Lemma 5.4. For a circle graph $G_0^{\text{cir}} = (V_0^{\text{cir}}, E_0^{\text{cir}})$ obtained from a block grid graph G_0 , if there are traversable edges $(u_1, v_1), (u_2, v_2), \dots, (u_k, v_k) \in E_0^{\text{cir}}$, then $(u_1, v_k) \in E_0^{\text{cir}}$.

Proof. We show that v_k is reachable from u_1 in G_0 by induction on k . First, we consider the case $k = 2$, namely (u_1, v_1) semi-crosses (u_2, v_2) . G_0 contains a path p_{u_1, v_1} which goes from u_1 to v_1 . Also, G_0 contains a path p_{u_2, v_2} which goes from u_2 to v_2 . Since G_0 is planar and u_1, v_1, u_2 , and v_2 are the rim vertices and the edges are semi-crossing, there exists a vertex w which is common in p_{u_1, v_1} and p_{u_2, v_2} in G_0 . Since w is reachable from u_1 and v_2 is reachable from w , there exists a path from u_1 to v_2 .

Next, we assume that the lemma is true for all sequences of traversable edges of length less than k . By the definition, there exist two edges (u_p, v_p) and (u_q, v_q) that the edge (u_k, v_k) separates $(p, q < k)$. We have two paths p_{u_1, v_p} from u_1 to v_p and p_{u_1, v_q} from u_1 to v_q in G_0 by the induction hypothesis. Also we have a path p_{u_k, v_k} from u_k to v_k . Since (u_k, v_k) separates (u_p, v_p) and (u_q, v_q) , v_p and v_q are on the different sides of arcs of the edge (u_k, v_k) . If u_1 and v_p are on the




Figure 5.2: A common vertex w' of a path from u_k to v_k and a path from u_1 to v_q or v_p for some $p, q < k$.

same arc of (u_k, v_k) , the paths p_{u_1, v_q} and p_{u_k, v_k} have a common vertex w' (see Figure 5.2(a)). On the other hand, if u_1 and v_q are on the same arc of (u_k, v_k) , the paths p_{u_1, v_p} and p_{u_k, v_k} have a common vertex w' (see Figure 5.2(b)). Thus there exists a path from u_1 to v_k via w' in G_0 . $\square$

5.2.2 Gadget graph

We introduce the notion of “gadget graph”. A gadget graph is a graph that is given a “label set” to each edge.

Definition 5.1. A *gadget graph* $\tilde{G}$ is a graph defined by a tuple $(\tilde{V}, \tilde{E}, \tilde{K}, \tilde{L})$, where $\tilde{V}$ is a set of vertices, $\tilde{E}$ is a set of edges, $\tilde{K}$ is a *path function* that assigns an edge or $\perp$ to each edge, and $\tilde{L}$ is a *level function* that assigns a *label set* to each edge. A label set is a set $\{i_1 \rightarrow o_1, i_2 \rightarrow o_2, \dots, i_k \rightarrow o_k\}$ of labels where each label $i_j \rightarrow o_j$, $i_j, o_j \in \mathbb{R} \cup \{\infty\}$, is a pair of *in-level* and *out-level*.

Remark. For an edge $(u, v) \in \tilde{E}$, we may use expressions $\tilde{K}(u, v)$ and $\tilde{L}(u, v)$ instead of $\tilde{K}((u, v))$ and $\tilde{L}((u, v))$ for simplicity.

Our goal is to transform a given circle graph (obtained from a block grid graph) $G_0^{\text{cir}} = (V_0^{\text{cir}}, E_0^{\text{cir}})$ in which all vertices in V_0^{cir} are placed on a cycle C to a *plane* gadget graph $\tilde{G}_p = (\tilde{V}_p^{\text{out}} \cup \tilde{V}_p^{\text{in}}, \tilde{E}_p, \tilde{K}_p, \tilde{L}_p)$ where $\tilde{V}_p^{\text{out}}$ is the set of *outer vertices* that are exactly the vertices of V_0^{cir} placed in the same way as G_0^{cir} on the cycle C , and $\tilde{V}_p^{\text{in}}$ is the set of *inner vertices* placed inside of C . All edges of $\tilde{E}_p$ are also placed inside of C under our embedding. The inner vertices of $\tilde{V}_p^{\text{in}}$ are used to replace crossing points of edges of E_0^{cir} to transform to a planar graph (see Figure 5.3). We would like to keep the “reachability” among vertices in $\tilde{V}_p^{\text{out}}$ in $\tilde{G}_p$ while bounding $|\tilde{V}_p^{\text{in}}| = O(n)$.




Figure 5.3: An example of the transformation from a circle graph to a gadget graph.

We explain how to characterize the reachability on a gadget graph. Consider any gadget graph $\tilde{G} = (\tilde{V}, \tilde{E}, \tilde{K}, \tilde{L})$, and let x and y be any two vertices of $\tilde{V}$. Intuitively, the reachability from x to y is characterized by a directed path on which we can send a token from x to y . Suppose that there is a directed path $p = (e_1, \dots, e_m)$ from x to y . We send a token through this path. The token has a level, which is initially ∞ when the token is at vertex x . (For a general discussion, we use a parameter ℓ_s for the initial level of the token.) When the token reaches the tail vertex $t(e_j)$ of some edge e_j of p with level ℓ , it can “go through” e_j to reach its head vertex $h(e_j)$ if $\tilde{L}(e_j)$ has an *available label* $i_j \rightarrow o_j$ such that $i_j \leq \ell$ holds for its in-level i_j . If the token uses a label $i_j \rightarrow o_j$, then its level becomes the out-level o_j at the vertex $h(e_j)$. If there are several available labels, then we naturally use the one with the highest out-level. If the token can reach y in this way, we consider that a “token tour” from x to y is “realized” by this path p . Technically, we introduce $\tilde{K}$ so that some edge can specify the next edge. We consider only a path $p = (e_1, \dots, e_m)$ as “valid” such that $e_{i+1} = \tilde{K}(e_i)$ for all e_i such that $\tilde{K}(e_i) \neq \perp$. We characterize the reachability from x to y on gadget graph $\tilde{G}$ by using a valid path realizing a token tour from x to y .

Definition 5.2. For any gadget graph $\tilde{G} = (\tilde{V}, \tilde{E}, \tilde{K}, \tilde{L})$, and for any two vertices x, y of $\tilde{V}$, there exists a *token tour* from x to y with initial level ℓ_s if there exists a sequence of edges $(e_1, \dots, e_m)$ that satisfies

1. $x = t(e_1)$ and $y = h(e_m)$,
2. $h(e_i) = t(e_{i+1})$ ($1 \leq i < m$),
3. if $\tilde{K}(e_i)$ is not $\perp$ ($1 \leq i < m$), then $e_{i+1} = \tilde{K}(e_i)$,
4. there exist labels $i_1 \rightarrow o_1 \in \tilde{L}(e_1), \dots, i_m \rightarrow o_m \in \tilde{L}(e_m)$ such that $\ell_s \geq i_1$ and $o_t \geq i_{t+1}$ for all $1 \leq t < m$.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM




Figure 5.4: An initial transformation step from $\tilde{G}_0$ to $\tilde{G}_1$.

At the beginning of our algorithm, we obtain a gadget graph $\tilde{G}_0 = (\tilde{V}_0, \tilde{E}_0, \tilde{K}_0, \tilde{L}_0)$ whose base graph is equal to G_0^{cir} , and $\tilde{K}_0(e) = \perp$, $\tilde{L}_0(e) = \{0 \rightarrow \infty\}$ for every $e \in \tilde{E}_0$. It is obvious that G_0^{cir} and $\tilde{G}_0$ have the same reachability. Namely, there exists a token tour from x to y for $x, y \in \tilde{V}_0$ in $\tilde{G}_0$ if and only if there exists an edge $(x, y) \in \tilde{E}_0$.

We explain first the outline of our transformation from $\tilde{G}_0$ to $\tilde{G}_p$. We begin by finding a chord $e_* = (u, v)$ with $\text{gap} \geq 2$ having no other $\text{gap-}2^+$ chord in its lower area, that is, one of the lowest $\text{gap-}2^+$ chords. (If there is no $\text{gap-}2^+$ chord, then the transformation is terminated.) For this e_* and its lower area, we transform them into a planar part and reduce the number of crossing points as follows (see Figure 5.4): (i) Consider all edges of $\tilde{G}_0$ crossing this chord e_* (e_1, e_2 and e_3 in Figure 5.4). Create a new inner vertex v_* of $\tilde{G}_p$ on the chord and bundle all crossing edges going through this vertex v_* ; that is, we replace all edges crossing e_* by edges between their end points in the lower area of e_* and v_* , and edges between v_* and their end points in the upper area of e_* . (ii) Introduce new inner vertices for edges crossing $\text{gap-}1$ chords in the lower area of e_* (w in Figure 5.4). (iii) Add appropriate label sets to those newly introduced edges so that the reachability is not changed by this transformation. At this point we regard the lower area of e_* as processed, and remove this part from the circle graph part of $\tilde{G}_0$ by replacing the arc $C[u, v]$ by a tour (u, v_*, v) to create a new circle graph part of $\tilde{G}_1$. We then repeat this transformation step on the circle graph part of $\tilde{G}_1$. In the algorithm, U_t is the vertices of the circle graph part of $\tilde{G}_t$, thus $\tilde{G}_t[U_t]$ indicates the circle graph part of $\tilde{G}_t$. Note that e_* is not removed and becomes a $\text{gap-}1$ chord in the next step.

We explain step (ii) for $\tilde{G}_0$ in more detail. Since e_* is a $\text{gap-}2^+$ chord, there exist only $\text{gap-}1$ chords or edges whose one end point is v_* in the lower area of e_* . If there are two edges e_0 and e_1 that cross each other, we replace the crossing point by a new inner vertex u (see Figure 5.5(a), (b)). The edge e_i becomes two edges $(t(e_i), u)$ and $(u, h(e_i))$ ($i = 0, 1$), and we set $\tilde{K}_1(t(e_i), u) = (u, h(e_i))$.




Figure 5.5: Examples of vertices made by MakePlanar.

The edges might be divided into more than two segments (see Figure 5.5(c)). We call the edge of $\tilde{G}_0$ *original edge* of the divided edges. By the path function, we must move along the original edge. An edge e might have a reverse direction edge $\bar{e} = (h(e), t(e))$ (see Figure 5.5(d)). In this case, e and $\bar{e}$ share a new vertex for resolving crossing points. For $\tilde{G}_t[U_t]$ ($t > 0$), we process the lower area in the same way. We refer to this algorithm as **MakePlanar**, and the new inner vertices created by **MakePlanar** in step t as V_{MP}^t .

The detailed process of step (iii) is written in Algorithm 2, and Algorithm 1 describes the entire process of step (i), (ii) and (iii). The following lemma shows that an output graph of Algorithm 1 has small size.

Lemma 5.5. Algorithm 1 terminates creating a planar graph of size $O(n)$.

Proof. In the beginning of the algorithm, $|U_0| = n$ and $|U_t|$ decreases by at least 1 for each iteration since the picked edge e_*^t is a gap-2⁺ chord. Hence the algorithm stops after at most n iterations and the number of the new inner vertices made at line 7, or v_*^t , is also at most n . If a gap- k chord is picked, we make at most $2k - 1$ new inner vertices by **MakePlanar**, namely $|V_{\text{MP}}^t| \leq 2k - 1$, since there exist only gap-1 chords in the lower area of the picked edge. The total number of inner vertices becomes at most

$$n + \sum_{i=1}^t (2k_i - 1) = n + 2 \sum_{i=1}^t k_i - t \leq n + 2 \times 2n = 5n$$

where t is the number of iterations and k_i means that a gap- k_i chord was picked in the i -th iteration. After all, $|\tilde{V}_{\text{p}}^{\text{out}} \cup \tilde{V}_{\text{p}}^{\text{in}}| \leq n + 5n = 6n$. $\square$

Now we explain Algorithm 2 describing how to assign labels to $\tilde{G}_{t+1}$ constructed in Algorithm 1. For each outer vertex $v \in \tilde{V}^{\text{out}}$, we keep three attributes $p^t(v)$, $\ell_{\text{in}}^t(v)$ and $\ell_{\text{out}}^t(v)$, and we call them *parent*, *in-level* and *out-level* respectively. We calculate these values from line 2 to 7 and line 25 to 27. $p^t(v)$ is a vertex belonging to the circle graph part of $\tilde{G}_t$, namely $p^t(v) \in U_t$. From the algorithm, we can show that there are token tours from v to $p^t(v)$ and/or

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM

Algorithm 1

Input: A circle graph $G_0^{\text{cir}} = (V_0^{\text{cir}}, E_0^{\text{cir}})$ obtained from a block graph.

Task: Output a plane gadget graph $\tilde{G}_p = (\tilde{V}_p^{\text{out}} \cup \tilde{V}_p^{\text{in}}, \tilde{E}_p, \tilde{K}_p, \tilde{L}_p)$ which satisfies $\tilde{V}_p^{\text{out}} = V_0^{\text{cir}}$ and the reachability among vertices in $\tilde{V}_p^{\text{out}}$ in $\tilde{G}_p$ is the same as G_0^{cir} .

- 1: initialize $t = 0$ // loop counter
 - 2: $\tilde{G}_0 = (\tilde{V}_0^{\text{out}} \cup \tilde{V}_0, \tilde{E}_0, \tilde{K}_0, \tilde{L}_0)$ where $\tilde{V}_0^{\text{out}} \leftarrow V_0^{\text{cir}}, \tilde{V}_0 \leftarrow \emptyset, \tilde{E}_0 \leftarrow E_0^{\text{cir}}, \tilde{K}_0(e) \leftarrow \perp, \tilde{L}_0(e) \leftarrow \{0 \rightarrow \infty\}$ for each $e \in E_0^{\text{cir}}$, and $U_0 \leftarrow \tilde{V}_0^{\text{out}}$
 - 3: for every $v \in \tilde{V}_0^{\text{out}}, \ell_{in}^0(v) \leftarrow 0, \ell_{out}^0(v) \leftarrow \infty, p^0(v) \leftarrow v$.
 - 4: **while** $\tilde{G}_t[U_t]$ has a lowest gap- 2^+ chord **do**
 - 5: pick a lowest gap- 2^+ chord e_*^t
 - 6: make a new vertex v_*^t
 - 7: $\tilde{V}_{t+1} \leftarrow \tilde{V}_t \cup \{v_*^t\}$
 - 8: $\tilde{E}_{t+1} \leftarrow (\tilde{E}_t \cup \{(t(e), v_*^t), (v_*^t, h(e)) \mid e \text{ crosses } e_* \text{ or } e = e_*\}) \setminus \{e \mid e \text{ crosses } e_*\}$
 - 9: $U_{t+1} \leftarrow (U_t \cup \{v_*^t\}) \setminus C_{\tilde{G}_t[U_t]}(t(e_*^t), h(e_*^t))$
 - 10: use **MakePlanar** to make the lower area of e_*^t planar and update $\tilde{V}_{t+1}, \tilde{E}_{t+1}$ and $\tilde{K}_{t+1}$.
 - 11: change the labels by using Algorithm 2 for keeping reachability
 - 12: output $\tilde{G}_{t+1}[C_{\tilde{G}_t[U_t]}[t(e_*^t), h(e_*^t)] \cup \{v_*^t\} \cup V_{\text{MP}}^t]$, which is the lower area of e_*^t .
 - 13: $t \leftarrow t + 1$
 - 14: **end while**
 - 15: use **MakePlanar** to make $\tilde{G}_t[U_t]$ planar and assign labels by line 17-24 of Algorithm 2.
 - 16: output $\tilde{G}_t[U_t \cup V_{\text{MP}}^t]$
-

Algorithm 2

Task: Set $\tilde{L}_{t+1}$ so that $\tilde{G}_{t+1}$ has the same reachability as $\tilde{G}_t$

- 1: For every edge e appearing in both $\tilde{G}_t$ and $\tilde{G}_{t+1}$, let $\tilde{L}_{t+1}(e) = \tilde{L}_t(e)$.
- 2: S^ℓ (resp., S^u) $\leftarrow \{v \in U_t \mid \exists e \in \tilde{E}_t \text{ s.t. } e \text{ crosses } e_*^t, t(e) = v \text{ or } h(e) = v, \text{ and } v \text{ is at the lower (resp., upper) area of } e_*^t\}$
- 3: T^ℓ (resp., T^u) $\leftarrow \{v \in V_0^{\text{cir}} \mid p^t(v) \in S^\ell \text{ (resp. } S^u)\}$
- 4: Fix any vertices $x', y' \in V_0^{\text{cir}}$ such that $p^t(x') = t(e_*^t), p^t(y') = h(e_*^t)$.
- 5: Set an order to T^ℓ according to the order appearing in $C_{G_0^{\text{cir}}}[y', x']$. We regard T^ℓ as a sequence $(t_1^\ell, t_2^\ell, \dots, t_{|T^\ell|}^\ell)$ (see Figure 5.6(b)).
- 6: Set an order to T^u in the same way as T^ℓ but according to the tour along the other arc. We also regard T^u as a sequence $(t_1^u, t_2^u, \dots, t_{|T^u|}^u)$ (see Figure 5.6(b)).
- 7: Use Algorithm 3 for calculating $\ell_{in}^{t+1}(v)$ and $\ell_{out}^{t+1}(v)$ for all $v \in T^\ell$.
- 8: **for** $u \in S^\ell$ **do**
- 9: $\tilde{L}_{t+1}(u, v_*^t) \leftarrow \{\ell_{in}^t(v) \rightarrow \ell_{in}^{t+1}(v) \mid p^t(v) = u\}$
- 10: $\tilde{L}_{t+1}(v_*^t, u) \leftarrow \{\ell_{out}^{t+1}(v) \rightarrow \ell_{out}^t(v) \mid p^t(v) = u\}$
- 11: **end for**
- 12: **for** $u \in S^u$ **do**
- 13: $\tilde{L}_{t+1}(u, v_*^t) \leftarrow \{\ell_{in}^t(t_i^u) \rightarrow \max_{t^\ell \in T^\ell} \{\ell_{out}^{t+1}(t^\ell) \mid (t_i^u, t^\ell) \in E_0^{\text{cir}}\} \mid t_i^u \in T^u \text{ and } p^t(t_i^u) = u\}$
- 14: $\tilde{L}_{t+1}(v_*^t, u) \leftarrow \{\min_{t^\ell \in T^\ell} \{\ell_{in}^{t+1}(t^\ell) \mid (t^\ell, t_i^u) \in E_0^{\text{cir}}\} \rightarrow \ell_{out}^t(t_i^u) \mid t_i^u \in T^u \text{ and } p^t(t_i^u) = u\}$
- 15: **end for**
- 16: $\tilde{L}_{t+1}(t(e_*^t), v_*^t) \leftarrow \{\infty \rightarrow 0\}, \tilde{L}_{t+1}(v_*^t, h(e_*^t)) \leftarrow \{\infty \rightarrow 0\}$
- 17: **for all** edge e created by **MakePlanar** **do**
- 18: Let e' be the original edge of e
- 19: **if** $t(e) = t(e')$ **then**
- 20: $\tilde{L}_{t+1}(e) \leftarrow \{a \rightarrow b \mid a \rightarrow b \in \tilde{L}_t(e')\}$
- 21: **else**
- 22: $\tilde{L}_{t+1}(e) \leftarrow \{b \rightarrow a \mid a \rightarrow b \in \tilde{L}_t(e')\}$
- 23: **end if**
- 24: **end for**
- 25: **for** $v \in \{u \in V_0^{\text{cir}} \mid \exists w \in U_t \text{ s.t. } w \text{ is at the lower area of } e_*^t \text{ and } p^t(u) = w\}$ **do**
- 26: $p^{t+1}(v) = v_*^t$
- 27: **end for**
- 28: Unchanged $\ell_{in}^t(\cdot), \ell_{out}^t(\cdot)$ and $p^t(\cdot)$ will be taken over to $\ell_{in}^{t+1}(\cdot), \ell_{out}^{t+1}(\cdot)$ and $p^{t+1}(\cdot)$.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH REACHABILITY PROBLEM

Algorithm 3

Task: Calculate $\ell_{in}^{t+1}(v)$ and $\ell_{out}^{t+1}(v)$ for all $v \in T^\ell$.

```

1: for  $i \in [1, |T^\ell|]$  do
2:    $\ell_{in}^{t+1}(t_i^\ell) \leftarrow \max\{j \mid (t_i^\ell, t_j^u) \in E_0^{\text{cir}}, t_j^u \in T^u\} + i/n$ 
3:    $\ell_{out}^{t+1}(t_i^\ell) \leftarrow \min\{j \mid (t_j^u, t_i^\ell) \in E_0^{\text{cir}}, t_j^u \in T^u\} + i/n$ 
4: end for
5: for  $i = 1$  to  $|T^\ell|$  do
6:    $\Delta \leftarrow \max(0, \max\{\ell_{out}^{t+1}(t_j^\ell) - j/n \mid 1 \leq j < i\} - (\ell_{in}^{t+1}(t_i^\ell) - i/n))$ 
7:   for  $k \in [i, |T^\ell|]$  do
8:      $\ell_{in}^{t+1}(t_k^\ell) \leftarrow \ell_{in}^{t+1}(t_i^\ell) + \Delta$ 
9:      $\ell_{out}^{t+1}(t_k^\ell) \leftarrow \ell_{out}^{t+1}(t_i^\ell) + \Delta$ 
10:  end for
11: end for

```

from $p^t(v)$ to v . For the token tour from v to $p^t(v)$, the final level of the token becomes $\ell_{in}^t(v)$. On the other hand, for the token tour from $p^t(v)$, it is enough to have $\ell_{out}^t(v)$ as an initial level to reach v . We will show these facts implicitly in the proof of Lemma 5.7.

At the beginning of each iteration of Algorithm 1, we choose a lowest gap-2⁺ chord e_*^t . We collect vertices in U_t which are endpoints for some edges crossing with e_*^t , and we refer to the vertices among them which are in the lower area of e_*^t as S^ℓ and the vertices in the upper area of e_*^t as S^u (see Figure 5.6(a) and line 2). Next we collect vertices whose parents are in S^ℓ (resp., S^u), and we denote them by T^ℓ (resp., T^u) (line 3). Let x' and y' be vertices whose parents are $t(e_*^t)$ and $h(e_*^t)$ respectively. We assign indices to the vertices in T^u and T^ℓ such that the nearer to x' a vertex is located, the larger index the vertex has (see Figure 5.6(b)). We regard T^ℓ as a sequence $(t_1^\ell, t_2^\ell, \dots, t_{|T^\ell|}^\ell)$, and T^u as a sequence $(t_1^u, t_2^u, \dots, t_{|T^u|}^u)$. For each vertex t_i^ℓ in T^ℓ , we calculate $\ell_{in}^{t+1}(t_i^\ell)$ and $\ell_{out}^{t+1}(t_i^\ell)$ in Algorithm 3. From line 1 to 4, we decide temporary values of $\ell_{in}^{t+1}(t_i^\ell)$ and $\ell_{out}^{t+1}(t_i^\ell)$ according to reachability among vertices in T^ℓ and T^u in G_0^{cir} . When t_j^u has the maximum index among vertices that t_i^ℓ can reach in T^u , we let $\ell_{in}^{t+1}(t_i^\ell) = j + i/n$. When t_j^ℓ has the minimum index among vertices which can reach t_i^ℓ in T^u , we let $\ell_{out}^{t+1}(t_i^\ell) = j + i/n$. The term i/n is for breaking ties. See Figure 5.7: $T^\ell = \{t_1^\ell, t_2^\ell, t_3^\ell\}$, $T^u = \{t_1^u, t_2^u, t_3^u\}$ and the edges are derived from E_0^{cir} . The vertex t_3^ℓ can reach t_1^u , t_2^u and t_3^u . Thus $\ell_{in}^{t+1}(t_3^\ell) = \max(1, 2, 3) + 3/n = 3 + 3/n$. The vertices t_2^ℓ and t_3^ℓ can reach t_2^ℓ . Thus $\ell_{out}^{t+1}(t_2^\ell) = \min(2, 3) + 2/n = 2 + 2/n$. In the next for-loop, we change the in- and out-levels so that the in-level of the larger indexed vertex is larger than the out-level of the smaller indexed vertex. If there exists a vertex t_j^ℓ such that $i > j$ and $\ell_{out}^{t+1}(t_j^\ell) > \ell_{in}^{t+1}(t_i^\ell)$, then we let $\Delta = (\ell_{out}^{t+1}(t_j^\ell) - j/n) - (\ell_{in}^{t+1}(t_i^\ell) - i/n)$ and add Δ to $\ell_{in}^{t+1}(t_i^\ell)$ and $\ell_{out}^{t+1}(t_i^\ell)$. For preserving the magnitude relationship between in- and out-levels of t_i^ℓ and those of t_k^ℓ ($k > i$), we also add Δ to $\ell_{in}^{t+1}(t_k^\ell)$ and $\ell_{out}^{t+1}(t_k^\ell)$. In Figure 5.7, we have $\ell_{in}^{t+1}(t_2^\ell) < \ell_{out}^{t+1}(t_1^\ell)$. Thus we

add $1 = (\ell_{out}^{t+1}(t_1^\ell) - 1/n) - (\ell_{in}^{t+1}(t_2^\ell) - 2/n)$ to $\ell_{in}^{t+1}(t_2^\ell)$. Moreover, we add 1 to $\ell_{out}^{t+1}(t_2^\ell)$, $\ell_{in}^{t+1}(t_3^\ell)$ and $\ell_{out}^{t+1}(t_3^\ell)$ so that we keep the magnitude relationship. Here we state a lemma.

Lemma 5.6. For any t , if $i > j$, then $\ell_{in}^{t+1}(t_i^\ell) > \ell_{out}^{t+1}(t_j^\ell)$.

Back to Algorithm 2. From line 8 to 16, we assign labels to edges newly appearing in $\tilde{G}_{t+1}$. Figure 5.8 is an example of how to assign label sets based on Figure 5.7. The vertex a is the parent of t_3^u , b is the parent of t_1^u and t_2^u , c is the parent of t_2^ℓ and t_3^ℓ and d is the parent of t_1^ℓ . Let v be any vertex in T^ℓ . For edges in the lower area of e_*^t , the edge $(p^t(v), v_*)$ has a label $\ell_{in}^t(v) \rightarrow \ell_{in}^{t+1}(v)$ (line 9), and the edge $(v_*, p^t(v))$ has a label $\ell_{out}^{t+1}(v) \rightarrow \ell_{out}^t(v)$ (line 10). In Figure 5.8, the edge (c, v_*) has labels $\ell_{in}^t(t_2^\ell) \rightarrow \ell_{in}^{t+1}(t_2^\ell)$ and $\ell_{in}^t(t_3^\ell) \rightarrow \ell_{in}^{t+1}(t_3^\ell)$. The edge (v_*, c) has a label $\ell_{out}^{t+1}(t_2^\ell) \rightarrow \ell_{out}^t(t_2^\ell)$ and the edge (v_*, d) has a label $\ell_{out}^{t+1}(t_1^\ell) \rightarrow \ell_{out}^t(t_1^\ell)$. Consider edges in the upper area of e_*^t . Let v be any vertex in T^u . The edge $(p^t(v), v_*)$ has a label $\ell_{in}^t(v) \rightarrow \ell_{max}$ where ℓ_{max} is the maximum in-level of vertices in T^ℓ that can reach v (line 13). The edge $(v_*, p^t(v))$ has a label $\ell_{min} \rightarrow \ell_{out}^t(v)$ where ℓ_{min} is the minimum out-level of vertices in T^ℓ that v can reach (line 14). In Figure 5.8, the edge (a, v_*) has a label $\ell_{in}^t(t_3^u) \rightarrow \ell_{out}^{t+1}(t_2^\ell)$ since t_3^u can reach t_1^ℓ and t_2^ℓ , and $\ell_{out}^{t+1}(t_1^\ell) < \ell_{out}^{t+1}(t_2^\ell)$ (see Figure 5.7). The edge (v_*, b) has a label $\ell_{in}^{t+1}(t_2^\ell) \rightarrow \ell_{out}^t(t_1^u)$ since t_2^ℓ and t_3^ℓ can reach t_1^u , and $\ell_{in}^{t+1}(t_2^\ell) < \ell_{in}^{t+1}(t_3^\ell)$ (see Figure 5.7). The edges $(t(e_*^t), v_*)$ and $(v_*, h(e_*^t))$ have only one label $\infty \rightarrow 0$, which prohibits using these edges (line 16).

From line 17 to 24, we assign labels to edges made by **MakePlanar**. For every edge (u, v) in the lower area of e_*^t , the edge (u, v) might be divided into some edges, for instance $(u, w_1), (w_1, w_2), \dots, (w_k, v)$ by **MakePlanar**. In this case, when (u, v) has a label $a \rightarrow b$, (u, w_1) has a label $a \rightarrow b$ and the other edges have labels $b \rightarrow b$ (see Figure 5.9).

From line 25 to 27, we update the parents of the vertices whose parents are in the lower area of e_*^t . For each vertex v in V_0^{cir} which $p^t(v)$ is in the lower area of e_*^t , we let $p^{t+1}(v) = v_*$.

For a gadget graph $\tilde{G} = (\tilde{V}, \tilde{E}, \tilde{K}, \tilde{L})$, we use $(v_1, \ell_1) \Rightarrow (v_2, \ell_2) \Rightarrow \dots \Rightarrow (v_m, \ell_m)$ to denote a token tour from v_1 to v_m in $\tilde{G}$ with having a level ℓ_i at $v_i \in \tilde{V}$ for any $1 \leq i \leq m$. Needless to say, if such a tour exists, $(v_i, v_{i+1}) \in \tilde{E}$ and $\ell'_i \rightarrow \ell_{i+1} \in \tilde{L}(v_i, v_{i+1})$ where $\ell_i \geq \ell'_i$ for any $1 \leq i < m$. In addition, when we would like to show which available labels we used, we write, for example, $(v_i, \ell_i; \ell'_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+1}, \ell_{i+1})$, which means the available label $\ell'_i \rightarrow \ell_{i+1}$ was used. The following lemma shows that paths in $\tilde{G}_0$ remain in $\tilde{G}_t$ for every t .

Lemma 5.7. For any t in Algorithm 1, if there exists an edge from x toward y in $\tilde{G}_0$, then there exists a token tour from x to y in $\tilde{G}_t$ whose length is at most $2t + 1$.

Proof. We prove the lemma by showing that if $(x, y) \in \tilde{E}_0$ then one of the following two statements holds in $\tilde{G}_t$ for any t :

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM




Figure 5.6: (a) An example of S^ℓ and S^u , (b) An example of T^ℓ and T^u .

- (i) there exists a token tour of length at most $2t + 1$ from x to y which uses no chords appearing in $\tilde{G}_t[U_t]$ (see Figure 5.10(i)).
- (ii) there exists a token tour $t_{x,y} = (x, \infty) \Rightarrow \dots \Rightarrow (p^t(x), \ell_{in}^t(x); \ell_{in}^t(x)^- \rightarrow \ell_{out}^t(y)^+) \Rightarrow (p^t(y), \ell_{out}^t(y)^+; \ell_{out}^t(y)^- \rightarrow \ell) \Rightarrow \dots \Rightarrow (y, \infty)$ where $\ell_{in}^t(x)^- \leq \ell_{in}^t(x)$, $\ell_{out}^t(y)^+ \geq \ell_{out}^t(y)$ and $\ell_{out}^t(y)^- \leq \ell_{out}^t(y)$. In addition, this tour uses no chords appearing in $\tilde{G}_t[U_t]$ except $(p^t(x), p^t(y))$, and its length is at most $2t + 1$ (see Figure 5.10(ii)).

We prove by induction on t . We have a tour $(x, \infty; 0 \rightarrow \infty) \Rightarrow (y, \infty)$ in $\tilde{G}_0$. Thus $\tilde{G}_0$ satisfies the statement (i) if x and y are consecutive on the cycle, and otherwise satisfies the statement (ii).

Assume that the statement (i) holds in $\tilde{G}_t$. The tour from x to y appears also in $\tilde{G}_{t+1}$ and satisfies the statement (i) in $\tilde{G}_{t+1}$. Now, we suppose the statement (ii) holds in $\tilde{G}_t$. We first consider the case that the chord $(p^t(x), p^t(y))$ does not cross e_*^t . When $(p^t(x), p^t(y))$ is in the lower area of e_*^t , the tour $t_{x,y}$ satisfies statement (i) in $\tilde{G}_{t+1}$. When $(p^t(x), p^t(y))$ is in the upper area of e_*^t or equal to e_*^t , the tour $t_{x,y}$ satisfies statement (ii) in $\tilde{G}_{t+1}$. Next, we assume that the chord $(p^t(x), p^t(y))$ crosses e_*^t . There are two cases:

- (I) $p^t(x) \in S^\ell$ and $p^t(y) \in S^u$: We have $x \in T^\ell$ and $y \in T^u$. There exists a label $\ell_{in}^t(x) \rightarrow \ell_{in}^{t+1}(x) \in \tilde{L}^{t+1}(p^t(x), v_*^t)$ (cf. line 9 of Algorithm 2). There also exists a label $\ell_{min} \rightarrow \ell_{out}^t(y) \in \tilde{L}^{t+1}(v_*^t, p^t(y))$ where $\ell_{min} = \min_{t^\ell \in T^\ell} \{\ell_{in}^{t+1}(t^\ell) \mid (t^\ell, y) \in E_0^{\text{cir}}\}$ (cf. line 14 of Algorithm 2). Since $x \in T^\ell$ and $(x, y) \in E_0^{\text{cir}}$, we have $\ell_{min} \leq \ell_{in}^{t+1}(x)$. Thus, in $\tilde{G}_{t+1}$, there exists a token tour $(x, \infty) \Rightarrow \dots \Rightarrow (p^t(x), \ell_{in}^t(x); \ell_{in}^t(x) \rightarrow \ell_{in}^{t+1}(x)) \Rightarrow (v_*^t, \ell_{in}^{t+1}(x); \ell_{min} \rightarrow \ell_{out}^t(y)) \Rightarrow (p^t(y), \ell_{out}^t(y)) \Rightarrow \dots \Rightarrow (y, \infty)$. If the




Figure 5.7: How to calculate in and out levels.




Figure 5.8: How to assign labels to edges.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM




Figure 5.9: How to assign labels to edges made by MakePlanar.




Figure 5.10: Two cases of token tours in the proof of Lemma 5.7.

edge $(p^t(x), v_*^t)$ has a crossing point in the lower area of e_*^t , we have to modify the part $(p^t(x), \ell_{in}^t(x); \ell_{in}^t(x) \rightarrow \ell_{in}^{t+1}(x)) \Rightarrow (v_*^t, \ell_{in}^{t+1}(x))$ to $(p^t(x), \ell_{in}^t(x); \ell_{in}^t(x) \rightarrow \ell_{in}^{t+1}(x)) \Rightarrow (u, \ell_{in}^{t+1}(x); \ell_{in}^{t+1}(x) \rightarrow \ell_{in}^{t+1}(x)) \Rightarrow (v_*^t, \ell_{in}^{t+1}(x))$ where u is a vertex created by **MakePlanar**.

- (II) $p^t(x) \in S^u$ and $p^t(y) \in S^\ell$: We have $x \in T^u$ and $y \in T^\ell$. There exists a label $\ell_{in}^t(x) \rightarrow \ell_{max} \in \tilde{L}^{t+1}(p^t(x), v_*^t)$ where $\ell_{max} = \max_{t^\ell \in T^\ell} \{\ell_{out}^{t+1}(t^\ell) \mid (x, t^\ell) \in E_0^{\text{cir}}\}$ (cf. line 13 of Algorithm 2). Since $y \in T^\ell$ and $(x, y) \in E_0^{\text{cir}}$, we have $\ell_{max} \geq \ell_{out}^{t+1}(y)$. There also exists a label $\ell_{out}^{t+1}(y) \rightarrow \ell_{out}^t(y) \in \tilde{L}^{t+1}(v_*^t, p^t(y))$ (cf. line 10 of Algorithm 2). Thus, in $\tilde{G}_{t+1}$, there exists a token tour $(x, \infty) \Rightarrow \dots \Rightarrow (p^t(x), \ell_{in}^t(x); \ell_{in}^t(x) \rightarrow \ell_{max}) \Rightarrow (v_*^t, \ell_{max}; \ell_{out}^{t+1}(y) \rightarrow \ell_{out}^t(y)) \Rightarrow (p^t(y), \ell_{out}^t(y)) \Rightarrow \dots \Rightarrow (y, \infty)$ in $\tilde{G}_{t+1}$. If the edge $(v_*^t, p^t(y))$ has a crossing point in the lower area of e_*^t , we have to modify the part $(v_*^t, \ell_{max}; \ell_{out}^{t+1}(y) \rightarrow \ell_{out}^t(y)) \Rightarrow (p^t(y), \ell_{out}^t(y))$ to $(v_*^t, \ell_{max}; \ell_{out}^{t+1}(y) \rightarrow \ell_{out}^t(y)) \Rightarrow (u, \ell_{out}^t(y); \ell_{out}^t(y) \rightarrow \ell_{out}^t(y)) \Rightarrow (p^t(y), \ell_{out}^t(y))$ where u is a vertex created by **MakePlanar**.

In both cases, the length of the new tour is longer than that of $t_{x,y}$ by at most 2, thus it is at most $2(t+1)+1$. We have $p^{t+1}(x) = v_*^t$ in case (I) and $p^{t+1}(y) = v_*^t$ in case (II). Thus the new tour has only one chord $(p^{t+1}(x), p^{t+1}(y))$ appearing in $\tilde{G}_{t+1}[U_{t+1}]$, and the chord has a label $\ell_{in}^t(x)^- \rightarrow \ell_{out}^t(y)^+$. Therefore the new tour satisfies statement (ii). $\square$

The following lemma shows the other direction: if there exists a token tour from x to y in the gadget graph, then there exists a path from x to y in the circle graph. From Lemma 5.4, it is enough to prove the following Lemma.

Lemma 5.8. For any t and $x, y \in V_0^{\text{cir}}$, if there exists a token tour from x to y in $\tilde{G}_t$, then there exists a traversable edge sequence $(e_1, \dots, e_k)$ in G_0^{cir} such that $t(e_1) = x$ and $h(e_k) = y$.

Before proving this, we prepare several lemmas, i.e., Lemma 5.9 to Lemma 5.13.

We refer to temporal in- and out-levels for $t_i^\ell \in T^\ell$ calculated at line 2 and 3 in Algorithm 3 as $t\ell_{in}^{t+1}(t_i^\ell)$ and $t\ell_{out}^{t+1}(t_i^\ell)$. Namely,

$$\begin{aligned} t\ell_{in}^{t+1}(t_i^\ell) &= \max\{j \mid (t_i^\ell, t_j^u) \in E_0^{\text{cir}}, t_j^u \in T^u\} + i/n, \\ t\ell_{out}^{t+1}(t_i^\ell) &= \min\{j \mid (t_j^u, t_i^\ell) \in E_0^{\text{cir}}, t_j^u \in T^u\} + i/n. \end{aligned}$$

Lemma 5.9. In Algorithm 3 of step t , if $i > j$ then $t\ell_{in}^{t+1}(t_i^\ell) > t\ell_{in}^{t+1}(t_j^\ell)$, $\ell_{in}^{t+1}(t_i^\ell) > \ell_{in}^{t+1}(t_j^\ell)$, $t\ell_{out}^{t+1}(t_i^\ell) > t\ell_{out}^{t+1}(t_j^\ell)$ and $\ell_{out}^{t+1}(t_i^\ell) > \ell_{out}^{t+1}(t_j^\ell)$. Moreover, if $i > j$ and $t\ell_{io_1}^{t+1}(t_i^\ell) > t\ell_{io_2}^{t+1}(t_j^\ell)$ then $\ell_{io_1}^{t+1}(t_i^\ell) > \ell_{io_2}^{t+1}(t_j^\ell)$ ($io_1, io_2 \in \{in, out\}$).

Proof. Let $i' = \max\{k \mid (t_i^\ell, t_k^u) \in E_0^{\text{cir}}, t_k^u \in T^u\}$ and $j' = \max\{k \mid (t_j^\ell, t_k^u) \in E_0^{\text{cir}}, t_k^u \in T^u\}$. Namely, $t\ell_{in}^t(t_i^\ell) = i' + i/n$ and $t\ell_{in}^t(t_j^\ell) = j' + j/n$. Assume $i' < j'$. Now we have $i > j$ and $i' < j'$. Thus the edges $(t_i^\ell, t_{i'}^u)$ and $(t_j^\ell, t_{j'}^u)$ are crossing (see Figure 5.11(a)). From Lemma 5.4, the edge $(t_i^\ell, t_{j'}^u)$ is in E_0^{cir} . This




Figure 5.11: Examples of wrong positions of vertices.

is contrary to the fact that i' is the maximum index. Thus $t\ell_{in}^t(t_i^\ell) = i' + i/n > j' + j/n = t\ell_{in}^t(t_j^\ell)$ holds.

Let $i' = \min\{k \mid (t_k^u, t_i^\ell) \in E_0^{\text{cir}}, t_k^u \in T^u\}$ and $j' = \min\{k \mid (t_k^u, t_j^\ell) \in E_0^{\text{cir}}, t_k^u \in T^u\}$. Namely, $t\ell_{out}^t(t_i^\ell) = i' + i/n$ and $t\ell_{out}^t(t_j^\ell) = j' + j/n$. Assume $i' < j'$. Now we have $i > j$ and $i' < j'$. Thus the edges $(t_{i'}^u, t_i^\ell)$ and $(t_{j'}^u, t_j^\ell)$ are crossing (see Figure 5.11(b)). From Lemma 5.4, the edge $(t_{i'}^u, t_j^\ell)$ is in E_0^{cir} . This is contrary to the fact that j' is the minimum index. Thus $t\ell_{out}^t(t_i^\ell) = i' + i/n > j' + j/n = t\ell_{out}^t(t_j^\ell)$ holds.

The in- and out-levels $\ell_{in}^t(t_i^\ell)$, $\ell_{in}^t(t_j^\ell)$, $\ell_{out}^t(t_i^\ell)$ and $\ell_{out}^t(t_j^\ell)$ might be larger than $t\ell_{in}^t(t_i^\ell)$, $t\ell_{in}^t(t_j^\ell)$, $t\ell_{out}^t(t_i^\ell)$ and $t\ell_{out}^t(t_j^\ell)$ since some positive integer Δ might be added (see at line 8 and 9 in Algorithm 3). However, when Δ is added to $\ell_{in}^t(t_j^\ell)$ (resp., $\ell_{out}^t(t_j^\ell)$), Δ is also added to $\ell_{in}^t(t_i^\ell)$ (resp., $\ell_{out}^t(t_i^\ell)$) since i is not less than j . Thus $\ell_{in}^t(t_i^\ell) > \ell_{in}^t(t_j^\ell)$ and $\ell_{out}^t(t_i^\ell) > \ell_{out}^t(t_j^\ell)$. By the same argument, the second statement holds. $\square$

Lemma 5.10. In Algorithm 2 in step t , for $t_p^\ell, t_q^\ell \in T^\ell$, if $t\ell_{in}^{t+1}(t_p^\ell) \geq t\ell_{out}^{t+1}(t_q^\ell)$, then $\ell_{in}^{t+1}(t_p^\ell) \geq \ell_{out}^{t+1}(t_q^\ell)$.

Proof. When $p > q$, this lemma holds from Lemma 5.9. Consider the case $p \leq q$. Let $t\ell_{in}^{t+1}(t_p^\ell) = i + p/n$ and $t\ell_{out}^{t+1}(t_q^\ell) = j + q/n$. Since $t\ell_{in}^{t+1}(t_p^\ell) \geq t\ell_{out}^{t+1}(t_q^\ell)$ and $p \leq q$, we have $i \geq j$. Assume $\ell_{in}^{t+1}(t_p^\ell) < \ell_{out}^{t+1}(t_q^\ell)$. In order that $\ell_{in}^{t+1}(t_p^\ell) < \ell_{out}^{t+1}(t_q^\ell)$ holds, some positive integer Δ should be added to $\ell_{out}^{t+1}(t_q^\ell)$ at line 9, and not added to $\ell_{in}^{t+1}(t_p^\ell)$ at line 8 of Algorithm 3. Thus, there should be a vertex t_r^ℓ such that $q \geq r \geq p$ and $\Delta = \max\{t\ell_{out}^{t+1}(t_k^\ell) - k/n \mid 1 \leq k < r\} - (t\ell_{in}^{t+1}(t_r^\ell) - r/n) > 0$. Since Δ is positive, there exists a vertex t_s^ℓ such that $r > s$ and $t\ell_{in}^{t+1}(t_r^\ell) < t\ell_{out}^{t+1}(t_s^\ell)$. Since $r \geq p$ and $q \geq s$, we have $t\ell_{in}^{t+1}(t_r^\ell) \geq t\ell_{in}^{t+1}(t_p^\ell)$ and $t\ell_{out}^{t+1}(t_q^\ell) \geq t\ell_{out}^{t+1}(t_s^\ell)$ from Lemma 5.9. Now $t\ell_{in}^{t+1}(t_p^\ell) \geq t\ell_{out}^{t+1}(t_q^\ell)$ holds, thus $t\ell_{in}^{t+1}(t_r^\ell) \geq t\ell_{out}^{t+1}(t_s^\ell)$ and Δ becomes non-positive. This is a contradiction. Thus $\ell_{in}^{t+1}(t_p^\ell) \geq \ell_{out}^{t+1}(t_q^\ell)$ holds. $\square$

For every label in $\tilde{G}_t$ and $k \leq t$, there are three types: (i) $\ell_{in}^k(x) \rightarrow \ell_{in}^{k+1}(x)$ (cf. line 9), (ii) $\ell_{out}^{k+1}(x) \rightarrow \ell_{out}^k(x)$ (cf. line 10) and (iii) $\ell_{in}^k(x) \rightarrow \ell_{out}^k(y)$ (cf. line 13, 14) for $x, y \in V_0^{\text{cir}}$. We define a *source vertex* and a *sink vertex* for any types of labels.

- (i) source vertex is x . When $i = \max\{j \mid (x, t_j^u) \in E_0^{\text{cir}}, t_j^u \in T^u\}$, sink vertex is t_i^u .
- (ii) sink vertex is x . When $i = \min\{j \mid (t_j^u, x) \in E_0^{\text{cir}}, t_j^u \in T^u\}$, source vertex is t_i^u .
- (iii) source vertex is x and sink vertex is y .

For a label L , we refer to an edge in G_0^{cir} from L 's source vertex to L 's sink vertex as a *source edge* of L . It is obvious that any source edge exists in G_0^{cir} .

Lemma 5.11. We consider any token tour of length 2 going through v_*^t : $(x, a'; a \rightarrow b) \Rightarrow (v_*^t, b; c \rightarrow d) \Rightarrow (y, d)$ in $\tilde{G}_{t'}$ where $t < t'$.

1. (x, v_*^t) is in upper area, and (v_*^t, y) is in upper area of e_*^t : Let (t_i^u, t_p^ℓ) be $a \rightarrow b$'s source edge, and we let (t_q^ℓ, t_j^u) be $c \rightarrow d$'s source edge. We have $p \geq q$.
2. (x, v_*^t) is in upper area, and (v_*^t, y) is in lower area of e_*^t : Let (t_i^u, t_p^ℓ) be $a \rightarrow b$'s source edge, and we let (t_j^u, t_q^ℓ) be $c \rightarrow d$'s source edge. We have $p \geq q$.
3. (x, v_*^t) is in lower area, and (v_*^t, y) is in upper area of e_*^t : Let (t_p^ℓ, t_i^u) be $a \rightarrow b$'s source edge, and we let (t_q^ℓ, t_j^u) be $c \rightarrow d$'s source edge. We have $i \geq j$ and $p \geq q$.
4. (x, v_*^t) is in lower area, and (v_*^t, y) is in lower area of e_*^t : Let (t_p^ℓ, t_i^u) be $a \rightarrow b$'s source edge, and we let (t_j^u, t_q^ℓ) be $c \rightarrow d$'s source edge. We have (i) $i \geq j$ and $p \geq q$, (ii) $i \geq j$ and $p < q$ or (iii) $i < j$ and $p \geq q$.

The indices i, j, p and q are based on the sequences T^u and T^ℓ made in Algorithm 2 in step t .

Proof.

1. We have $b = \ell_{out}^{t+1}(t_p^\ell)$ and $c = \ell_{in}^{t+1}(t_q^\ell)$. From the rule of token tours, $\ell_{out}^{t+1}(t_p^\ell) \geq \ell_{in}^{t+1}(t_q^\ell)$ holds. If $p < q$, we have $\ell_{out}^{t+1}(t_p^\ell) < \ell_{in}^{t+1}(t_q^\ell)$ from Lemma 5.6. Thus we have $p \geq q$.
2. We have $b = \ell_{out}^{t+1}(t_p^\ell)$ and $c = \ell_{out}^{t+1}(t_q^\ell)$. From the rule of token tours, $\ell_{out}^{t+1}(t_p^\ell) \geq \ell_{out}^{t+1}(t_q^\ell)$ holds. If $p < q$, we have $\ell_{out}^{t+1}(t_p^\ell) < \ell_{out}^{t+1}(t_q^\ell)$ from Lemma 5.9. Thus we have $p \geq q$.




Figure 5.12: Relations of vertices and their parents.

3. We have $b = \ell_{in}^{t+1}(t_p^\ell)$ and $c = \ell_{in}^{t+1}(t_q^\ell)$. From the rule of token tours, $\ell_{in}^{t+1}(t_p^\ell) \geq \ell_{in}^{t+1}(t_q^\ell)$ holds. If $p < q$, we have $\ell_{in}^{t+1}(t_p^\ell) < \ell_{in}^{t+1}(t_q^\ell)$ from Lemma 5.9. Thus we have $p \geq q$. From the definition of source edge, t_i^u has the maximum index among vertices that t_p^ℓ can reach. Assume $i < j$. Now we have $p \geq q$ and $i < j$. Thus the edges (t_p^ℓ, t_i^u) and (t_q^ℓ, t_j^u) are semi-crossing. From Lemma 5.4, the edge (t_p^ℓ, t_j^u) is in E_0^{cir} . This is contrary to the fact that i is the maximum index. Thus $i \geq j$ holds.
4. We have $b = \ell_{in}^{t+1}(t_p^\ell)$ and $c = \ell_{out}^{t+1}(t_q^\ell)$. From the rule of token tours, $\ell_{in}^{t+1}(t_p^\ell) \geq \ell_{out}^{t+1}(t_q^\ell)$ holds. We will show that $i < j$ and $p < q$ do not hold simultaneously. Assume $i < j$ and $p < q$. We have $i + p/n = t\ell_{in}^{t+1}(t_p^\ell) < t\ell_{out}^{t+1}(t_q^\ell) = j + q/n$. From Lemma 5.9, we have $\ell_{in}^{t+1}(t_p^\ell) < \ell_{out}^{t+1}(t_q^\ell)$ since $p < q$, and we cannot follow the tour. Thus, there are three possible relationships: (i) $i \geq j$ and $p \geq q$, (ii) $i \geq j$ and $p < q$, (iii) $i < j$ and $p \geq q$.

□

Lemma 5.12. For any three vertices $u, v, w \in U_t$, if (u, v, w) is in clockwise (resp., anti-clockwise) order in $\tilde{G}_t[U_t]$, then (x, y, z) is also in clockwise (resp., anti-clockwise) order in G_0^{cir} for any $x, y, z \in V_0^{cir}$ such that $p^t(x) = u, p^t(y) = v$ and $p^t(z) = w$ (see Figure 5.12(a)).

Proof. We prove by induction on t . Since the parent of every vertex is itself in step 0, the Lemma is true in step 0. Let u, v and w be vertices such that (u, v, w) is in clockwise (resp., anti-clockwise) order in $\tilde{G}_{t+1}[U_{t+1}]$. Fix any three vertices x, y and z such that $p^{t+1}(x) = u, p^{t+1}(y) = v$ and $p^{t+1}(z) = w$. When none of $u,$

v or w is v_*^t , (u, v, w) is in clockwise (resp., anti-clockwise) order also in $\tilde{G}_t[U_t]$. Thus (x, y, z) is in clockwise (resp., anti-clockwise) order from the induction hypothesis. Let $u = v_*^t$. Since $p^t(x)$ is in the lower area of e_*^t , $(p^t(x), u, v)$ is in clockwise (resp., anti-clockwise) order in $\tilde{G}_t[U_t]$ (see Figure 5.12(b)). From the induction hypothesis, (x, y, z) is in clockwise (resp., anti-clockwise) order in G_0^{cir} . In cases $v = v_*^t$ or $w = v_*^t$, the Lemma is proved in the same way. $\square$

Lemma 5.13. For any $k \leq t$ such that $v_*^k \in U_t$, let x and y be vertices such that $p^k(x) = t(e_*^k)$ and $p^k(y) = h(e_*^k)$ respectively. $p^t(x)$, v_*^k and $p^t(y)$ are consecutive in $\tilde{G}_t[U_t]$.

Proof. Fix any k . We prove by induction on t . When $t = k$, it is obvious that $p^t(x)$, v_*^k and $p^t(y)$ are consecutive. Assume the Lemma is true for a fixed t . If v_*^k is not an endpoint of e_*^t , we have $p^t(x) = p^{t+1}(x)$ and $p^t(y) = p^{t+1}(y)$. Thus $p^{t+1}(x)$, v_*^k and $p^{t+1}(y)$ are consecutive in $\tilde{G}_{t+1}[U_{t+1}]$ from the induction hypothesis. When v_*^k is an endpoint of e_*^t and $p^t(x)$ (resp., $p^t(y)$) is on $C_{\tilde{G}_t[U_t]}(t(e_*^t), h(e_*^t))$, we have $p^{t+1}(x) = v_*^t$ (resp., $p^{t+1}(y) = v_*^t$). Since v_*^k and v_*^t are adjacent, $p^{t+1}(x)$, v_*^k and $p^{t+1}(y)$ are consecutive in $\tilde{G}_{t+1}[U_{t+1}]$. $\square$

Now we prove Lemma 5.8.

Lemma 5.14 (restated). For any t and $x, y \in V_0^{\text{cir}}$, if there exists a token tour from x to y in $\tilde{G}_t$, then there exists a traversable edge sequence $(e_1, \dots, e_k)$ in G_0^{cir} such that $t(e_1) = x$ and $h(e_k) = y$.

Proof. Let $t_{x,y}$ be a token tour $(v_1, \ell_1; f_1 \rightarrow \ell_2) \Rightarrow (v_2, \ell_2; f_2 \rightarrow \ell_3) \Rightarrow \dots \Rightarrow (v_m, \ell_m)$ such that $v_1 = x$ and $v_m = y$. We modify $t_{x,y}$. If the edges $(v_i, v_{i+1}), \dots, (v_{i+d-1}, v_{i+d})$ are made by **MakePlanar** and they have the same original edge, namely $\tilde{K}_t(v_j, v_{j+1}) = (v_{j+1}, v_{j+2})$ ($i \leq j < i + d - 1$), we change the partial tour $(v_i, \ell_i; f_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+1}, \ell_{i+1}; f_{i+1} \rightarrow \ell_{i+2}) \Rightarrow \dots \Rightarrow (v_{i+d}, \ell_{i+d})$ to $(v_i, \ell_i; f_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+d}, \ell_{i+d})$. Note that ℓ_{i+1} is equal to ℓ_{i+d} . Next, we remove redundant moves. Consider a partial tour of length 2 $(v_i, \ell_i; f_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+1}, \ell_{i+1}; f_{i+1} \rightarrow \ell_{i+2}) \Rightarrow (v_{i+2}, \ell_{i+2}; f_{i+2} \rightarrow \ell_{i+3})$. When $v_i = v_{i+2}$ and $f_i \geq \ell_{i+2}$, we regard this move as a redundant move. We have $\ell_i \geq f_i$ and $\ell_{i+2} \geq f_{i+2}$. If the move is redundant, $\ell_i \geq f_{i+2}$ holds. We change the partial tour $(v_i, \ell_i; f_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+1}, \ell_{i+1}; f_{i+1} \rightarrow \ell_{i+2}) \Rightarrow (v_{i+2}, \ell_{i+2}; f_{i+2} \rightarrow \ell_{i+3}) \Rightarrow (v_{i+3}, \ell_{i+3})$ to $(v_i, \ell_i; f_{i+2} \rightarrow \ell_{i+3}) \Rightarrow (v_{i+3}, \ell_{i+3})$. Again, we let the changed token tour be $t_{x,y} = (v_1, \ell_1; f_1 \rightarrow \ell_2) \Rightarrow (v_2, \ell_2; f_2 \rightarrow \ell_3) \Rightarrow \dots \Rightarrow (v_m, \ell_m)$.

We construct a traversable edge sequence. We put source edges of the labels appearing in $t_{x,y}$. Let this edge sequence be $(e_1, \dots, e_{m-1})$ where e_i is a source edge of the label $f_i \rightarrow \ell_{i+1}$. For each i ($2 \leq i < m$), v_i corresponds to v_*^k for some $1 \leq k < t$ since we removed vertices made by **MakePlanar** from the tour. For every i ($2 \leq i < m$), we take an edge $e'_i \in E_0^{\text{cir}}$ such that $p^k(t(e'_i)) = t(e_*^k)$ and $p^k(h(e'_i)) = h(e_*^k)$ where k is derived from $v_i = v_*^k$. There exist several ways to choose e'_i . We show that if we select e'_i 's appropriately, the edge sequence $(e_1, e'_2, e_2, e'_3, e_3, \dots, e_{m-2}, e'_{m-1}, e_{m-1})$ becomes traversable.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM




Figure 5.13: Relations of two source edges.

It is obvious that e_1 crosses e'_2 . We have to show that all edges except for e_1 and e'_2 separate two edges appearing before themselves. By induction, suppose we fixed e'_j ($i+2 \leq j < m$). We show which pair of edges e_{i+1} separates. Consider the partial tour of length 2 $(v_i, \ell_i; f_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+1}, \ell_{i+1}; f_{i+1} \rightarrow \ell_{i+2}) \Rightarrow (v_{i+2}, \ell_{i+2})$, and let $v_{i+1} = v_*^k$. We suppose that this partial tour corresponds to case 1 of Lemma 5.11, namely the edge (v_i, v_{i+1}) is in upper area and the edge (v_{i+1}, v_{i+2}) is in upper area of e_*^k . We let (t_i^u, t_p^l) be $f_i \rightarrow \ell_{i+1}$'s source edge, and (t_q^l, t_j^u) be $f_{i+1} \rightarrow \ell_{i+2}$'s source edge. From Lemma 5.11, we have $p \geq q$. Thus $e_{i+1} = (t_q^l, t_j^u)$ separates e'_{i+1} and $e_i = (t_i^u, t_p^l)$ for any choice of e'_{i+1} (see Figure 5.13 1-(i), 1-(ii)). The horizontal edge corresponds to e'_{i+1} . In both cases (i) $i \geq j$ and (ii) $i < j$, t_p^l and $h(e'_{i+1})$ are on the opposite side of (t_q^l, t_j^u) . When the partial tour corresponds to case 2, 3, 4-(i) or 4-(ii), e_{i+1} separates e'_{i+1} and e_i for any choice of e'_{i+1} (see Figure 5.13).

Suppose the partial tour corresponds to case 4-(iii). Assume $v_i = v_{i+2}$. Let v_i and v_{i+1} be v_*^k and v_*^t respectively. Let (t_p^l, t_i^u) be $f_i \rightarrow \ell_{i+1}$'s source edge and (t_j^u, t_q^l) be $f_{i+1} \rightarrow \ell_{i+2}$'s source edge. The indices i, j, p and q are based on the sequences T^u and T^l made in Algorithm 2 in step t . Recall $i < j$ and $p \geq q$. We will show that $f_i = \ell_{in}^{k+1}(t_p^l) \geq \ell_{out}^{k+1}(t_q^l) = \ell_{i+2}$. Suppose e_*^k and e_*^t has the same direction, namely t_p^l and t_q^l had indices p' and q' respectively in step k such that $p' \geq q'$. In this case, we have $\ell_{in}^{k+1}(t_p^l) \geq \ell_{out}^{k+1}(t_q^l)$ from Lemma 5.6. Assume e_*^k and e_*^t has the opposite direction, namely t_i^u, t_j^u, t_p^l and t_q^l had indices i', j', p' and q' respectively in step k such that $i' > j'$ and $p' \leq q'$. Since $i' > j'$, $\ell_{in}^{k+1}(t_p^l) - p/n \geq i'$ and $\ell_{out}^{k+1}(t_q^l) - q/n \leq j'$, we have $\ell_{in}^{k+1}(t_p^l) \geq \ell_{out}^{k+1}(t_q^l)$. From Lemma 5.10, $\ell_{in}^{k+1}(t_p^l) \geq \ell_{out}^{k+1}(t_q^l)$. Thus, when




Figure 5.14: case 4-(iii): $p^t(h(e'))$ is on $C_{\tilde{G}_t[U_t]}[t(e_*^t), v_{i+2}]$.

$v_i = v_{i+2}$, this is a redundant move and does not appear in $t_{x,y}$.

Again, let v_i and v_{i+1} be v_*^k and v_*^t respectively. Let e' be an edge such that $p^k(t(e')) = t(e_*^k)$ and $p^k(h(e')) = h(e_*^k)$. Now we consider the case $v_i \neq v_{i+2}$, thus there are two cases for a location of $p^t(h(e'))$ from Lemma 5.13.

1. $p^t(h(e'))$ is on $C_{\tilde{G}_t[U_t]}[t(e_*^t), v_{i+2}]$: Consider the four vertices $p^t(h(e'))$, v_{i+2} , $h(e_*^t)$ and $p^t(t(e_{i+1}))$. A possible order on U_t of the four vertices is $(p^t(h(e')), v_{i+2}, h(e_*^t), p^t(t(e_{i+1})))$. From Lemma 5.12, e_{i+1} separates e' and e'_{i+1} for any choice of e' (see Figure 5.14).
2. $p^t(h(e'))$ is equal to v_{i+2} : If e_{i+1} separates e' and e'_{i+1} , we set e' as e'_i (see Figure 5.15 (a)). However, e_{i+1} might not separate e' and e'_{i+1} . In this case, e' crosses e_{i+1} . From Lemma 5.4, there exists an edge $(t(e'), h(e_{i+1}))$ in E_0^{cir} (see Figure 5.15 (b)). We choose $(t(e'), h(e_{i+1}))$ as e'_i and e_{i+1} separates e'_i and e'_{i+1} in G_0^{cir} .

Suppose we fixed e'_j ($i+2 \leq j < m$). We show how to select e'_{i+1} and which pair of edges e'_{i+2} separates. Consider any partial token tour of length 2 $(v_i, \ell_i; f_i \rightarrow \ell_{i+1}) \Rightarrow (v_{i+1}, \ell_{i+1}; f_{i+1} \rightarrow \ell_{i+2}) \Rightarrow (v_{i+2}, \ell_{i+2})$. Assume $v_i = v_*^s$, $v_{i+1} = v_*^k$ and $v_{i+2} = v_*^t$. Note that e_i and e_{i+1} are source edges of $f_i \rightarrow \ell_{i+1}$ and $f_{i+1} \rightarrow \ell_{i+2}$ respectively, and e'_{i+2} is an edge such that $p^t(t(e'_{i+2})) = t(e_*^t)$ and $p^t(h(e'_{i+2})) = h(e_*^t)$. We let e' be a source edge of a label of e_*^k . When $t > k$:

1. v_*^s is on U_t : Consider the four vertices $p^t(t(e'_{i+2}))$, v_*^k , $p^t(h(e'_{i+2}))$ and $p^t(h(e_{i+1}))$. A possible order on U_t of the four vertices is $(p^t(t(e'_{i+2})), v_*^k, p^t(h(e'_{i+2})), p^t(h(e_{i+1})))$. From Lemma 5.12, e'_{i+2} separates e_{i+1} and e_i (see Figure 5.16). We select e' as e'_{i+1} .
2. v_*^s is not on U_t : From Lemma 5.13, there are three cases for a location of $p^t(h(e'))$.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM






Figure 5.15: case 4-(iii): $p^t(h(e'))$ is equal to v_{i+2} .

- (i) $p^t(h(e'))$ is neither $h(e_{i+1}^t)$ nor $t(e_{i+1}^t)$: Consider the four vertices $p^t(t(e'_{i+2}))$, $p^t(h(e'))$, $p^t(h(e'_{i+2}))$ and $p^t(h(e_{i+1}))$. A possible order on U_t of the four vertices is $(p^t(t(e'_{i+2})), p^t(h(e')), p^t(h(e'_{i+2})), p^t(h(e_{i+1})))$. From Lemma 5.12, e'_{i+2} separates e_{i+1} and e' (see Figure 5.17). We select e' as e'_{i+1} .
- (ii) $p^t(h(e'))$ is equal to $h(e_{i+1}^t)$: When e'_{i+2} separates e_{i+1} and e' (see Figure 5.18(a)), we select e' as e'_{i+1} . e'_{i+2} might not separate e_{i+1} and e' (see Figure 5.18(b)). In this case, e' crosses e'_{i+2} . From Lemma 5.4, there exists an edge $(t(e'), h(e'_{i+2}))$ in E_0^{cir} . We choose $(t(e'), h(e'_{i+2}))$ as e'_{i+1} and e'_{i+2} separates e_{i+1} and e'_{i+1} in G_0^{cir} .
- (iii) $p^t(h(e'))$ is equal to $t(e_{i+1}^t)$: When e'_{i+2} separates e_{i+1} and e' (see Figure 5.19(a)), we select e' as e'_{i+1} . e'_{i+2} might not separate e_{i+1} and e' (see Figure 5.19(b)). In step k , $h(e_i)$ and $h(e'_{i+2})$ was in T^u . Let i and j be indices of $h(e_i)$ and $h(e'_{i+2})$ respectively, that is, $h(e_i) = t_i^u$ and $h(e'_{i+2}) = t_j^u$. If $i < j$, then e_i crosses e'_{i+2} , and $t(e_i)$ can reach $h(e'_{i+2})$ (see Figure 5.19(c)). Since t_i^u has the maximum index among


 Figure 5.16: $t > k$: v_*^s is on U_t .

 Figure 5.17: $t > k$: v_*^s is not on U_t , $p^t(h(e'))$ is neither $h(e_*^t)$ nor $t(e_*^t)$.

vertices $t(e_i)$ can reach, this is a contradiction. Therefore we have $i \geq j$. Thus e'_{i+2} separates e_i and e' (see Figure 5.19(d)). We select e' as e'_{i+1} .

When $k > t$:

1. (v_*^s, v_*^k) and (v_*^k, v_*^t) are in the same side of e_*^k : $p^k(t(e'_{i+2}))$, $p^k(h(e_{i+1}))$ and $p^k(h(e'_{i+2}))$ are consecutive on U_k from Lemma 5.13. The parent of $h(e_i)$ is on the other side of e_*^k . From Lemma 5.12, e'_{i+2} separates e_{i+1} and e_i (see Figure 5.20). We select e' as e'_{i+1} .
2. (v_*^s, v_*^k) and (v_*^k, v_*^t) are in the opposite side of e_*^k : From Lemma 5.13, there are three cases for a location of $p^k(h(e'_{i+2}))$ and $p^k(t(e'_{i+2}))$.
 - (i) Neither $p^k(h(e'_{i+2}))$ nor $p^k(t(e'_{i+2}))$ is $h(e_*^k)$: Consider the four vertices $p^k(t(e'_{i+2}))$, $p^k(h(e'_{i+2}))$, $p^k(h(e_{i+1}))$ and $p^k(e')$. A possible

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM




Figure 5.18: $t > k$: v_*^s is not on U_t , $p^t(h(e'))$ is equal to $h(e_*^t)$.




Figure 5.19: $t > k$: v_*^s is not on U_t , $p^t(h(e'))$ is equal to $t(e_*)$.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH
REACHABILITY PROBLEM




Figure 5.20: $k > t$: (v_*^s, v_*^k) and (v_*^k, v_*^t) are in the same side of e_*^k .

order on U_k of the four vertices is

$(p^k(t(e'_{i+2})), p^k(h(e_{i+1})), p^k(h(e'_{i+2})), p^k(e'))$ (see Figure 5.21). From Lemma 5.12, e'_{i+2} separates e_{i+1} and e' . We select e' as e'_{i+1} .

- (ii) $p^k(h(e'_{i+2}))$ is equal to $h(e_*^k)$: When e'_{i+2} separates e_{i+1} and e' (see Figure 5.22(a)), we select e' as e'_{i+1} . e'_{i+2} might not separate e_{i+1} and e' (see Figure 5.22(b)). In this case, e' crosses e'_{i+2} . From Lemma 5.4, there exists an edge $(t(e'), h(e'_{i+2}))$ in E_0^{cir} . We choose $(t(e'), h(e'_{i+2}))$ as e'_{i+1} , and e'_{i+2} separates e_{i+1} and e'_{i+1} in G_0^{cir} .
- (iii) $p^k(t(e'_{i+2}))$ is equal to $h(e_*^k)$: When e'_{i+2} separates e_{i+1} and e' (see Figure 5.23(a)), we select e' as e'_{i+1} . e'_{i+2} might not separate e_{i+1} and e' (see Figure 5.23(b)). Assume the edge (v_*^s, v_*^k) is in the upper area of e_*^k . In step k , $h(e_i)$ and $h(e'_{i+2})$ was in T^ℓ . Let p and q be indices of $h(e_i)$ and $h(e'_{i+2})$ respectively, that is, $h(e_i) = t_p^\ell$ and $h(e'_{i+2}) = t_q^\ell$. If $p < q$, then e_i crosses e'_{i+2} , and $t(e_i)$ can reach $h(e'_{i+2})$ (see Figure 5.23(c)). Since t_p^ℓ has the maximum index among vertices $t(e_i)$ can reach, this is a contradiction. Therefore we have $p \geq q$. Thus e'_{i+2} separates e_i and e' (see Figure 5.23(d)). We select e' as e'_{i+1} . In the case (v_*^s, v_*^k) is in the lower area of e_*^k , we could show that e'_{i+2} separates e_i and e' in the almost same way.

When we consider which pair of edges e_{i+1} separates, we might choose a specific e'_i (see Figure 5.15(b)). When we consider which pair of edges e'_{i+1} separates, we also might choose specific e'_i (see Figure 5.18(b) and 5.22(b)). If the cases of Figure 5.15(b) and Figure 5.18(b) occur simultaneously, $h(e_*^t)$ must be v_{i+2} , but the edge $(v_*^t, t(e_*^t))$ has no available label. In the case of Figure 5.22(b), the edge e_i is in the upper area of e_*^k . Thus, these cases never occur simultaneously. From the above, the constructed edge sequence is traversable. $\square$

We analyze the space and time complexity of Algorithm 1. Note that, for




Figure 5.21: $k > t$: (v_s^s, v_s^k) and (v_s^k, v_s^t) are in the opposite side of e_*^k , neither $p^k(h(e'_{i+2}))$ nor $p^k(t(e'_{i+2}))$ is $h(e_*^k)$.

saving computation space, we do not implement the Algorithm straightforwardly in some points. We begin with the space complexity. We regard the circle graph $G_0^{cir} = (V_0^{cir}, E_0^{cir})$ as the input. For every $v \in V_0^{cir}$, we keep three attributes $\ell_{in}^t(v)$, $\ell_{out}^t(v)$ and $p^t(v)$ in step t . The in- and out-levels are rational numbers that have the form of $i + j/n$. Thus we keep two integers i and j for each in- and out-level. We use $\tilde{O}(n)$ space for preserving them. In step t , we also keep U_t by using $\tilde{O}(n)$ space. We need $\tilde{G}_t[U_t]$, but we do not keep $\tilde{E}_t$ explicitly. For $u, v \in U_t$, whether there exists an edge (u, v) in $\tilde{G}_t[U_t]$ is equivalent to whether there exists an edge (x, y) in E_0^{cir} such that $p^t(x) = u$ and $p^t(y) = v$. Since E_0^{cir} is included in the input, we could calculate it with $\tilde{O}(1)$ space. We keep no other information throughout the Algorithm. The number of edges in $\tilde{G}_t[U_t]$ is at most $2|U_t|^2 = O(n^2)$. Thus, for line 4 and 5, we can find a lowest gap-2⁺ chord by using $\tilde{O}(1)$ space. For line 7 and 9, we use only $\tilde{O}(1)$ space for updating $\tilde{V}_t$ and U_t . For line 8, we ignore the edges in the upper area of e_*^t (these edges belong to $\tilde{G}_{t+1}[U_{t+1}]$, thus we have no need to keep them). For the edges in the lower area of e_*^t , since there exist only gap-1 chords in the area, the number of edges in the area is $O(n)$. We use $\tilde{O}(n)$ space for temporarily keeping them. In **MakePlanar** (line 10), we look through them, and find crossing points and resolve them and set $\tilde{K}_{t+1}(\cdot)$ by using $\tilde{O}(n)$ space.

Now we consider Algorithm 2. The number of edges in $\tilde{G}_t[U_t]$ is at most $2|U_t|^2 = O(n^2)$. Thus, for line 2, we can find S^ℓ and S^u by using $\tilde{O}(1)$ space, and we use $\tilde{O}(n)$ space for keeping them. For line 3 to 6, since $|T^\ell|, |T^u| = O(n)$, we also use $\tilde{O}(n)$ for keeping T^ℓ and T^u . In addition, we use $\tilde{O}(n)$ space for calculating $\ell_{in}^{t+1}(v)$ and $\ell_{out}^{t+1}(v)$ for all $v \in T^\ell$. In Algorithm 3, we use $\tilde{O}(1)$ space for each operation and the length of for-loops is $O(n)$. Thus we use $\tilde{O}(1)$ space in all. For line 8 to 11, we only refer to in- and out-levels that we are

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH REACHABILITY PROBLEM




Figure 5.22: $k > t$: (v_*^s, v_*^k) and (v_*^k, v_*^t) are in the opposite side of $e_*^k, p^k(h(e'_{i+2}))$ is equal to $h(e_*^k)$.

keeping. For line 12 to 15, we do not keep and ignore the labels belonging to edges in the upper area. For line 16, we use $\tilde{O}(1)$ space. For line 17 to 24, we check whether an edge in the lower area was divided by **MakePlanar** and we use additional $\tilde{O}(1)$ space. For line 25 to 27, we can find all vertices in the lower area of e_*^t by using $\tilde{O}(n)$ space, and we use additional $\tilde{O}(1)$ space for updating $p^{t+1}(\cdot)$.

We go back to Algorithm 1. For line 12, we output the information of the vertices, edges, labels and values of the path function in the lower area of e_*^t . Here we have to calculate the labels on the gap-1 chords (other information is preserved now). Let the gap-1 chord be (v_*^p, v_*^q) . If $p < q$, this edge was made in step q and the labels on the edge were calculated at line 13 of Algorithm 2. Thus, for any $v \in V_0^{cir}$ such that $p^t(v) = v_*^p$, we calculate $\ell_{out} = \max_{t^\ell \in V_0^{cir}, p^t(t^\ell) = v_*^q} \{\ell_{out}^t(t^\ell) \mid (v, t^\ell) \in E_0^{cir}\}$, and $\ell_{in}^t(v) \rightarrow \ell_{out}$ becomes one of the labels on the edge (if the vertex v is not in T^u , ℓ_{out} is not




Figure 5.23: $k > t$: (v_*^s, v_*^k) and (v_*^k, v_*^t) are in the opposite side of $e_*^k, p^k(t(e'_{i+2}))$ is equal to $h(e_*^k)$.

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH REACHABILITY PROBLEM

defined and a label for v does not exist). On the other hand, if $p > q$, this edge was made in step p and the labels on the edge were calculated at line 14 of Algorithm 2. Thus, for any $v \in V_0^{\text{cir}}$ such that $p^t(v) = v_*^q$, we calculate $\ell_{in} = \min_{t^\ell \in V_0^{\text{cir}}, p^t(t^\ell) = v_*^p} \{\ell_{in}^t(t^\ell) \mid (t^\ell, v) \in E_0^{\text{cir}}\}$, and $\ell_{in} \rightarrow \ell_{out}^t(v)$ becomes one of the labels on the edge (if the vertex v is not in T^u , ℓ_{in} is not defined and a label for v does not exist). We use additional $\tilde{O}(1)$ space for these calculation. For line 15, we trace line 10 to 12. In total, we use $\tilde{O}(n)$ space.

Next consider the time complexity. In Lemma 5.5, we proved that the while-loop at line 4 stops after at most n steps. Since the sizes of U_t , S^ℓ , S^u , T^ℓ and T^u are all $O(n)$, every operation in the Algorithms takes $\text{poly}(n)$ time. Thus this algorithm runs in polynomial time.

Lemma 5.15. Algorithm 1 runs in polynomial time with using $\tilde{O}(n)$ space.

From Lemma 5.7, 5.8 and 5.15, we can obtain desired $\tilde{G}_p$ with $\tilde{O}(n) = \tilde{O}(N^{1/3})$ space and polynomial time.


5.3 Apply PlanarReach to a gadget graph

By applying PlanarReach to the obtained plane gadget graph $\tilde{G}_p$ with $O(N^{2/3})$ vertices, we can prove Theorem 5.2. In this section, we explain how to apply PlanarReach to a plane gadget graph, which has labels in edges. We have to modify PlanarReach slightly. We now describe the outline of the algorithm PlanarReach. The notion of “separator” is central to the algorithm.

Definition 5.3. For any undirected graph $G = (V, E)$ and for any constant ρ , $0 < \rho < 1$, a subset of vertices S is called a ρ -separator if (i) removal of S disconnects G into two subgraphs A and B , and (ii) the number of vertices of any component is at most $\rho \cdot |V|$. The size of separator is the number of vertices in the separator.

It is well known that every planar graph with n vertices has a $(2/3)$ -separator of size $O(\sqrt{n})$ [33, 45], and we refer an algorithm which obtains such a separator as **Separator**.

Let $G = (V, E)$, s and t be the given input; that is, G is a directed graph, and s and t are the start and goal vertices in V . We assume that t is reachable from s in G , and explain that the algorithm confirms it. We use $\underline{G}$ to denote an underlying undirected graph of G . The algorithm first uses **Separator** to compute a separator S of size $O(\sqrt{n})$ for $\underline{G}$, and suppose G is divided into two subgraphs $G[V_0]$ and $G[V_1]$ by S ($V_0 \cap V_1 = \emptyset$, $V_0 \cup V_1 \cup S = V$). Let us fix a path p from s to t . The path p is divided into some k subpaths $p_1, p_2, \dots, p_k$ by S . Note that the end vertex u_i of p_i is on S and whole path p_i is in either one of $G[V_0 \cup S]$ and $G[V_1 \cup S]$. Suppose p_1 is in $G[V_0 \cup S]$. By searching in $G[V_0 \cup S]$, we can find u_1 is reachable from s . The algorithm remembers it and searches $G[V_1 \cup S]$ from u_1 in the next step. Then we can find p_2 , namely u_2 is reachable from u_1 and s . By repeating this procedure, we can confirm u_i is reachable from s for


 Figure 5.24: An example of a separator S and separated paths.

any i . More precisely, for each vertex $v \in S$, we keep a boolean value which represents reachability from s . In each searching step, we start the search from vertices in S whose boolean values are true. We use this reachability algorithm recursively when searching $G[V_b \cup S]$ ($b \in \{0, 1\}$). Algorithm 4 is a pseudo code for this algorithm. In the actual algorithm, we have to control the recursion more carefully, but this is enough for explaining where to modify the algorithm for gadget graphs.

Algorithm 4 PlanarReach($G = (V, E), V_s, R[V_s], V_t$)

Input: A planar graph G , start vertices V_s , a boolean array $R[V_s]$ for V_s , end vertices V_t .

Task: Return a boolean array $R[V_t]$ for V_t . For any $v \in V_t$, $R[v]$ is true if and only if v is reachable from some vertex $u \in V_s$ such that $R[u]$ is true.

```

1: if the size of  $V$  is small enough then
2:   use a standard BFS algorithm and compute  $R[V_t]$ .
3:   return  $R[V_t]$ 
4: else
5:   Run Separator and obtain a separator  $S$  ( $G$  is divided into  $G[V_0]$  and  $G[V_1]$ ).
6:    $R[S] = \text{PlanarReach}(G[V_0 \cup S \cup V_s], V_s, R[V_s], S)$ 
7:   while unsearched paths remain do
8:      $R[S] = \text{PlanarReach}(G[V_0 \cup S], S, R[S], S)$ 
9:      $R[S] = \text{PlanarReach}(G[V_1 \cup S], S, R[S], S)$ 
10:  end while
11:  return  $\text{PlanarReach}(G[V_1 \cup S \cup V_t], S, R[S], V_t)$ 
12: end if
    
```

Now, we explain where to modify. Let $\tilde{G}_p = (\tilde{V}_p, \tilde{E}_p, \tilde{K}_p, \tilde{L}_p)$ be an input plane gadget graph of PlanarReach and N be the number of vertices of an input grid graph of Algorithm 1. Consider a gadget graph $\tilde{G}'_p = (\tilde{V}'_p, \tilde{E}'_p, \tilde{K}'_p, \tilde{L}'_p)$ which

5. $\tilde{O}(N^{1/3})$ -SPACE ALGORITHM FOR THE GRID GRAPH REACHABILITY PROBLEM

is a subgraph of $\tilde{G}_p$. While we execute **PlanarReach**, for every $v \in S$, we have to keep a boolean value whether v is reachable from s with using $O(|S|)$ space. For $\tilde{G}'_p$, instead of the boolean value, we keep the maximum level that a token starting from s could have at v . When v is equal to $t(\tilde{K}'_p(e))$ for some edge e , we should keep a specific level that a token can have at v when the token used the edge e last. Such a vertex is made by **MakePlanar**, and we should keep at most two specific levels for a vertex. Thus we use $\tilde{O}(|S|)$ space for preserving them, and we can still obtain an $\tilde{O}(N^{1/3})$ space algorithm.

For $\tilde{G}'_p$, we use Algorithm 5 like Bellman-Ford algorithm instead of BFS. Algorithm 5 takes as input $\tilde{G}'_p$, a start vertex s , an initial level ℓ_s and an edge restriction $r \in \tilde{E}_p \cup \{\perp\}$. For any $v \in \tilde{V}'_p$, the algorithm computes the maximum level that a token starting from s with a level ℓ_s can have at v . When v is equal to $t(\tilde{K}'_p(e))$ for some edge e , the algorithm calculates the maximum level that a token can have at v when the token used the edge e last. In Algorithm 5, $A[v_e]$ means that the maximum level that a token can have at v with using the edge e last, and $A[v_\perp]$ means that the maximum level that a token can have at v with using an edge e last such that $\tilde{K}'_p(e) = \perp$. At the end of t -th while-loop, for any $v \in \tilde{V}'_p$, $A[v_*]$ has the maximum level which we can have at v within t steps by starting from s with level ℓ_s . At line 4, we use two mappings k and $\tilde{K}^{-1}$, and they are defined as follows:

$$k(e) = \begin{cases} \perp & \text{if } \tilde{K}'_p(e) = \perp \\ e & \text{otherwise} \end{cases}, \quad \tilde{K}^{-1}(e) = \begin{cases} e' & \exists e', \tilde{K}'_p(e') = e \\ \perp & \text{otherwise} \end{cases}$$

Since the value $A[\cdot]$ changes no more than two times with the same label, the while-loop will terminate in $|\tilde{L}'_p|$ steps where $|\tilde{L}'_p| = |\bigcup_{e \in \tilde{E}'_p} \tilde{L}'_p(e)|$. Thus the computation time for Algorithm 5 is $O(|\tilde{L}'_p|^2)$. An edge has at most $O(N^{1/3})$ labels, thus the algorithm runs in polynomial time of N .

Algorithm 5

Input: A gadget graph $\tilde{G}'_p = (\tilde{V}'_p, \tilde{E}'_p, \tilde{K}'_p, \tilde{L}'_p)$, start vertex s , initial level ℓ_s , edge restriction $r \in \tilde{E}_p \cup \{\perp\}$.

- 1: initialize $A[v_\perp] = A[v_e] = -1$ for every $v \in \tilde{V}'_p$ and $e \in \tilde{E}'_p$ such that $h(e) = v$ except for s and let $A[s_r] = \ell_s$
- 2: **while** A was changed in the previous loop **do**
- 3: **for all** $e \in \tilde{E}$ **do**
- 4: $A[h(e)_{k(e)}] \leftarrow \max(A[h(e)_{k(e)}], \max\{b \mid a \rightarrow b \in \tilde{L}'_p(e), A[t(e)_{\tilde{K}^{-1}(e)}] \geq a\})$
- 5: **end for**
- 6: **end while**
- 7: output A

5.3. APPLY PLANARREACH TO A GADGET GRAPH

Chapter 6

Conclusion

In the third and fourth chapters, we presented an $\tilde{O}(\sqrt{n})$ -space and polynomial-time algorithm computing an $O(\sqrt{n})$ -size separator. As mentioned in the third chapter, the separator construction problem is log-space reducible to the problem of constructing a BFS tree. Thus, if we could construct a sub-linear space algorithm of a BFS tree, then a more space efficient separator algorithm can be obtained. Though DFS search (decision version) is P-complete [52], BFS search (decision version) is solvable by a circuit in NC^3 , namely a circuit of depth $O(\log^3 n)$ and polynomial size [34]. Since BFS can be highly parallelized, we might be able to devise a $o(n)$ -space BFS algorithm.

In the fifth chapter, we presented an $\tilde{O}(n^{1/3})$ -space algorithm for the grid graph reachability problem. The most natural question is whether we can apply our algorithm to the planar graph reachability problem. Although the directed planar reachability is reduced to the directed reachability on grid graphs [1], the reduction blows up the size of the graph by a large polynomial factor and hence it is not useful. Moreover, it is known that there exist planar graphs that require quadratic grid area for embedding [62]. However we do not have to stick to grid graphs. We can apply our algorithm to graphs which can be divided into small blocks efficiently. For instance we can use our algorithm for king's graphs [19]. More directly, for using our algorithm, it is enough to design an algorithm that divides a planar graph into small blocks efficiently.

Bibliography

- [1] Eric Allender, David A. Mix Barrington, Tanmoy Chakraborty, Samir Datta, and Sambuddha Roy. Planar and grid graph reachability problems. *Theory of Computing Systems*, 45(4):675–723, 2009.
- [2] Eric Allender, T. Chakraborty, DAM Barrington, Samir Datta, and Sambuddha Roy. Grid graph reachability problems. In *21st Annual IEEE Conference on Computational Complexity (CCC'06)*, pages 15–pp. IEEE, 2006.
- [3] Eric Allender and Meena Mahajan. The complexity of planarity testing. *Information and Computation*, 189(1):117, 2004.
- [4] Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society*, 3(4):801–808, 1990.
- [5] Noga Alon, Paul Seymour, and Robin Thomas. Planar separators. *SIAM Journal on Discrete Mathematics*, 7(2):184–193, 1994.
- [6] Roy Armoni, Amnon Ta-Shma, Avi Wigderson, and Shiyu Zhou. An $O(\log^{4/3} n)$ space algorithm for (s, t) connectivity in undirected graphs. *Journal of the Association for Computing Machinery*, 47(2):294–294, 2000.
- [7] Sanjeev Arora, Michelangelo Grigni, David R. Karger, Philip N. Klein, and Andrzej Woloszyn. A polynomial-time approximation scheme for weighted planar graph tsp. In *SODA*, volume 98, pages 33–41, 1998.
- [8] Tetsuo Asano and Benjamin Doerr. Memory-constrained algorithms for shortest path problem. In *CCCG 2011*, pages 315–319. CCCG. CA, 2011.
- [9] Tetsuo Asano, David Kirkpatrick, Kotaro Nakagawa, and Osamu Watanabe. $\tilde{O}(\sqrt{n})$ -space and polynomial-time algorithm for planar directed graph reachability. In *International Symposium on Mathematical Foundations of Computer Science*, pages 45–56. Springer, 2014.
- [10] Ryo Ashida, Tatsuya Imai, Kotaro Nakagawa, Aduri Pavan, N. Variyam Vinodchandran, and Osamu Watanabe. A sublinear-space and polynomial-time separator algorithm for planar graphs. *ECCC*, TR19-091, 2019.

-
- [11] Ryo Ashida, Sebastian Kuhnert, and Osamu Watanabe. A space-efficient separator algorithm for planar graphs. *IEICE Transactions on Fundamentals*, E102-A(9), 2019.
 - [12] Ryo Ashida and Kotaro Nakagawa. $\tilde{O}(n^{1/3})$ -space algorithm for the grid graph reachability problem. In *34th International Symposium on Computational Geometry (SoCG 2018)*, volume 99, pages 5:1–5:13, 2018.
 - [13] Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.
 - [14] Reuven Bar-Yehuda and Shimon Even. On approximating a vertex cover for planar graphs. In *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, pages 303–309. ACM, 1982.
 - [15] Greg Barnes, Jonathan F. Buss, Walter L. Ruzzo, and Baruch Schieber. A sublinear space, polynomial time algorithm for directed st connectivity. *SIAM Journal on Computing*, 27(5):1273–1282, 1998.
 - [16] Marshall Bern. Faster exact algorithms for Steiner trees in planar networks. *Networks*, 20(1):109–120, 1990.
 - [17] Chris Bourke, Raghunath Tewari, and NV. Vinodchandran. Directed planar reachability is in unambiguous log-space. *ACM Transactions on Computation Theory (TOCT)*, 1(1):4, 2009.
 - [18] Diptarka Chakraborty and Raghunath Tewari. An $O(n^\epsilon)$ space and polynomial time algorithm for reachability in directed layered planar graphs. *ACM Transactions on Computation Theory (TOCT)*, 9(4):19, 2018.
 - [19] Gerard Jennhwa Chang. Algorithmic aspects of domination in graphs. *Handbook of Combinatorial Optimization*, pages 221–282, 2013.
 - [20] Hsien-Chih Chang and Hsueh-I Lu. Computing the girth of a planar graph in linear time. *SIAM Journal on Computing*, 42(3):1077–1094, 2013.
 - [21] Norishige Chiba, Takao Nishizeki, and Nobuji Saito. Applications of the Lipton and Tarjan’s planar separator theorem. *J. Inf. Process*, 4(4):203–207, 1981.
 - [22] Stephen A. Cook and Pierre McKenzie. Problems complete for deterministic logarithmic space. *Journal of Algorithms*, 8(3):385–394, 1987.
 - [23] Stephen A. Cook and Charles W. Rackoff. Space lower bounds for maze threadability on restricted machines. *SIAM Journal on Computing*, 9(3):636–652, 1980.
 - [24] Hubert de Fraysseix and Pierre Rosenstiehl. A depth-first-search characterization of planarity. *Annals of Discrete Mathematics*, 13:75–80, 1982.

BIBLIOGRAPHY

- [25] Vladimir G. Dei, Bettina Klinz, Gerhard J. Woeginger, et al. Exact algorithms for the Hamiltonian cycle problem in planar graphs. *Operations Research Letters*, 34(3):269–274, 2006.
- [26] Reinhard Diestel. *Graph Theory (Graduate Texts in Mathematics)*. Springer, 2005.
- [27] Hristo N. Djidjev and Shankar M. Venkatesan. Reduced constants for simple cycle graph separation. *Acta Informatica*, 34(3):231–243, 1997.
- [28] Hristo Nicolov Djidjev. On the problem of partitioning planar graphs. *SIAM Journal on Algebraic Discrete Methods*, 3(2):229–240, 1982.
- [29] Jeff Edmonds, Chung Keung Poon, and Dimitris Achlioptas. Tight lower bounds for st-connectivity on the nnjag model. *SIAM Journal on Computing*, 28(6):2257–2284, 1999.
- [30] David Eppstein, Gary L. Miller, and Shang-Hua Teng. A deterministic linear time algorithm for geometric separators and its applications. *Fundamenta Informaticae*, 22(4):309–329, 1995.
- [31] Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing*, 16(6):1004–1022, 1987.
- [32] Alan M. Frieze, Gary L. Miller, and Shang-Hua Teng. Separator based parallel divide and conquer in computational geometry. In *SPAA*, volume 92, pages 420–30, 1992.
- [33] Hillel Gazit and Gary L. Miller. A parallel algorithm for finding a separator in planar graphs. In *Foundations of Computer Science, 1987., 28th Annual Symposium on*, pages 238–248. IEEE, 1987.
- [34] Hillel Gazit and Gary L. Miller. An improved parallel algorithm that computes the bfs numbering of a directed graph. *Information Processing Letters*, 28(2):61–65, 1988.
- [35] Hillel Gazit and Gary L. Miller. Planar separators and the euclidean norm. In *Algorithms*, pages 338–347. Springer, 1990.
- [36] Michael T. Goodrich. Planar separators and parallel polygon triangulation. *Journal of Computer and System Sciences*, 51(3):374–389, 1995.
- [37] Monika R. Henzinger, Philip Klein, Satish Rao, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. *journal of computer and system sciences*, 55(1):3–23, 1997.
- [38] Tatsuya Imai, Kotaro Nakagawa, Aduri Pavan, NV. Vinodchandran, and Osamu Watanabe. An $O(n^{1/2+\epsilon})$ -space and polynomial-time algorithm for directed planar reachability. In *Computational Complexity (CCC), 2013 IEEE Conference on*, pages 277–286. IEEE, 2013.

-
- [39] Sampath Kannan, Sanjeev Khanna, and Sudeepa Roy. Stcon in directed unique-path graphs. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2008.
 - [40] Philip Klein, Shay Mozes, and Oren Weimann. Shortest paths in directed planar graphs with negative lengths: A linear-space $O(n \log^2 n)$ -time algorithm. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 236–245. SIAM, 2009.
 - [41] Casimir Kuratowski. Sur le probleme des courbes gauches en topologie. *Fundamenta mathematicae*, 15(1):271–283, 1930.
 - [42] Jakub Łącki and Piotr Sankowski. Min-cuts and shortest cycles in planar graphs in $O(n \log \log n)$ time. In *European Symposium on Algorithms*, pages 155–166. Springer, 2011.
 - [43] Harry R. Lewis and Christos H. Papadimitriou. Symmetric space-bounded computation. *Theoretical Computer Science*, 19(2):161–187, 1982.
 - [44] Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 162–170. IEEE, 1977.
 - [45] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
 - [46] Saunders Mac Lane. *A combinatorial condition for planar graphs*. Seminarium Matemat., 1936.
 - [47] Gary L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and system Sciences*, 32(3):265–279, 1986.
 - [48] Gary L. Miller, Shang-Hua Teng, William Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *Journal of the ACM (JACM)*, 44(1):1–29, 1997.
 - [49] Noam Nisan, Endre Szemerédi, and Avi Wigderson. Undirected connectivity in $O(\log^{1.5} n)$ space. In *Proceedings., 33rd Annual Symposium on Foundations of Computer Science*, pages 24–29. IEEE, 1992.
 - [50] Christos H. Papadimitriou and Martha Sideri. The bisection width of grid graphs. *Mathematical systems theory*, 29(2):97–110, 1996.
 - [51] Chung Keung Poon. Space bounds for graph connectivity problems on node-named jags and node-ordered jags. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 218–227. IEEE, 1993.
 - [52] John H. Reif. Depth-first search is inherently sequential. *Information Processing Letters*, 20(5):229–234, 1985.

BIBLIOGRAPHY

- [53] Omer Reingold. Undirected connectivity in log-space. *Journal of the ACM (JACM)*, 55(4):17, 2008.
- [54] Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4(2):177–192, 1970.
- [55] Walter Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323–343, 1989.
- [56] Daniel A. Spielman and Shang-Hua Teng. Disk packings and planar separators. In *Symposium on Computational Geometry*, pages 349–358, 1996.
- [57] Daniel A. Spielmat and Shang-Hua Teng. Spectral partitioning works: Planar graphs and finite element meshes. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 96–105. IEEE, 1996.
- [58] Derrick Stolee and NV. Vinodchandran. Space-efficient algorithms for reachability in surface-embedded graphs. In *Computational Complexity (CCC), 2012 IEEE 27th Annual Conference on*, pages 326–333. IEEE, 2012.
- [59] Siamak Tazari and Matthias Müller-Hannemann. Shortest paths in linear time on minor-closed graph classes, with an application to steiner tree approximation. *Discrete Applied Mathematics*, 157(4):673–684, 2009.
- [60] Vladimir Trifonov. An $O(\log n \log \log n)$ space algorithm for undirected st-connectivity. *SIAM Journal on Computing*, 38(2):449–483, 2008.
- [61] William T. Tutte. Toward a theory of crossing numbers. *Journal of Combinatorial Theory*, 8(1):45–53, 1970.
- [62] Leslie G. Valiant. Universality considerations in VLSI circuits. *IEEE Transactions on Computers*, 100(2):135–140, 1981.
- [63] Klaus Wagner. Über eine eigenschaft der ebenen komplexe. *Mathematische Annalen*, 114(1):570–590, 1937.
- [64] Hassler Whitney. Non-separable and planar graphs. *Transactions of the American Mathematical Society*, 34(2):339–339, 1932.
- [65] Avi Wigderson. The complexity of graph connectivity. In *International Symposium on Mathematical Foundations of Computer Science*, pages 112–132. Springer, 1992.