

Cycle Bases in Graphs

Characterization, Algorithms, Complexity, and Applications

Telikepalli Kavitha* Christian Liebchen† Kurt Mehlhorn‡
Dimitrios Michail Romeo Rizzi§ Torsten Ueckerdt¶ Katharina A. Zweig||

July 18, 2009

Abstract

Cycles in graphs play an important role in many applications, e.g., analysis of electrical networks, analysis of chemical and biological pathways, periodic scheduling, and graph drawing. From a mathematical point of view, cycles in graphs have a rich structure. Cycle bases are a compact description of the set of all cycles of a graph. In this paper, we survey results on cycle bases and prove new ones. We introduce different kinds of cycle bases, characterize them in terms of their cycle matrix, and prove structural results about them, in particular, a-priori length bounds. We give polynomial algorithms for the minimum cycle basis problem for some of the classes and prove \mathcal{APX} -hardness for others. We also discuss three applications and show that they require different kinds of cycle bases.

Contents

| | | |
|----------|---------------------------------------|-----------|
| 1 | Introduction | 3 |
| 2 | Definitions | 5 |
| 3 | Classification of Cycle Bases | 10 |
| 3.1 | Existence. | 10 |
| 3.2 | Characterizations. | 11 |
| 3.3 | Simple Examples | 15 |
| 3.4 | Variants of the MCB Problem | 17 |
| 3.5 | Directed and $GF(p)$ -Bases | 19 |
| 3.6 | Circuits versus Cycles. | 20 |
| 3.7 | Reductions. | 24 |
| 3.8 | Selected Properties. | 25 |

*Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India

†DB Schenker Rail Deutschland AG, Mainz, Germany. Part of this work was done within the DFG Research Center MATHEON in Berlin

‡Max-Planck-Institut für Informatik, Saarbrücken, Germany

§Dipartimento di Matematica ed Informatica (DIMI), Università degli Studi di Udine, Udine, Italy

¶Technische Universität Berlin, Berlin, Germany, Supported by . . .

||Department of Biological Physics, Eötvös Loránd University, Budapest, Hungary

| | | |
|----------|---|-----------|
| 4 | Length and Weight of Cycle Bases | 28 |
| 4.1 | Weakly Fundamental Bases | 29 |
| 4.2 | Fundamental Bases | 31 |
| 5 | Polynomial Time Algorithms for Minimum Cycle Bases | 36 |
| 5.1 | The Greedy Algorithm and the Horton Set | 36 |
| 5.2 | De Pina's Approach | 39 |
| 5.3 | Maintaining the Orthogonal Space | 41 |
| 5.4 | Computing the Circuits | 43 |
| 5.5 | Computing Shortest Path Trees | 45 |
| 5.6 | Putting it Together | 45 |
| 5.7 | A Greedy Algorithm for Integral Cycle Bases? | 46 |
| 5.8 | Planar Graphs | 48 |
| 5.9 | Approximation | 50 |
| 5.9.1 | Approximate Shortest Paths | 51 |
| 5.9.2 | 2-approximation | 54 |
| 5.10 | Algorithm Engineering | 55 |
| 5.11 | Relevant Cycles | 57 |
| 6 | Hardness Results | 58 |
| 6.1 | Strictly Fundamental Cycle Bases | 58 |
| 6.2 | Weakly Fundamental Cycle Basis | 61 |
| 7 | Applications | 70 |
| 7.1 | Kirchhoff's Voltage Law | 70 |
| 7.2 | Periodic Scheduling in Traffic Planning | 70 |
| 7.3 | Graph Drawing | 72 |

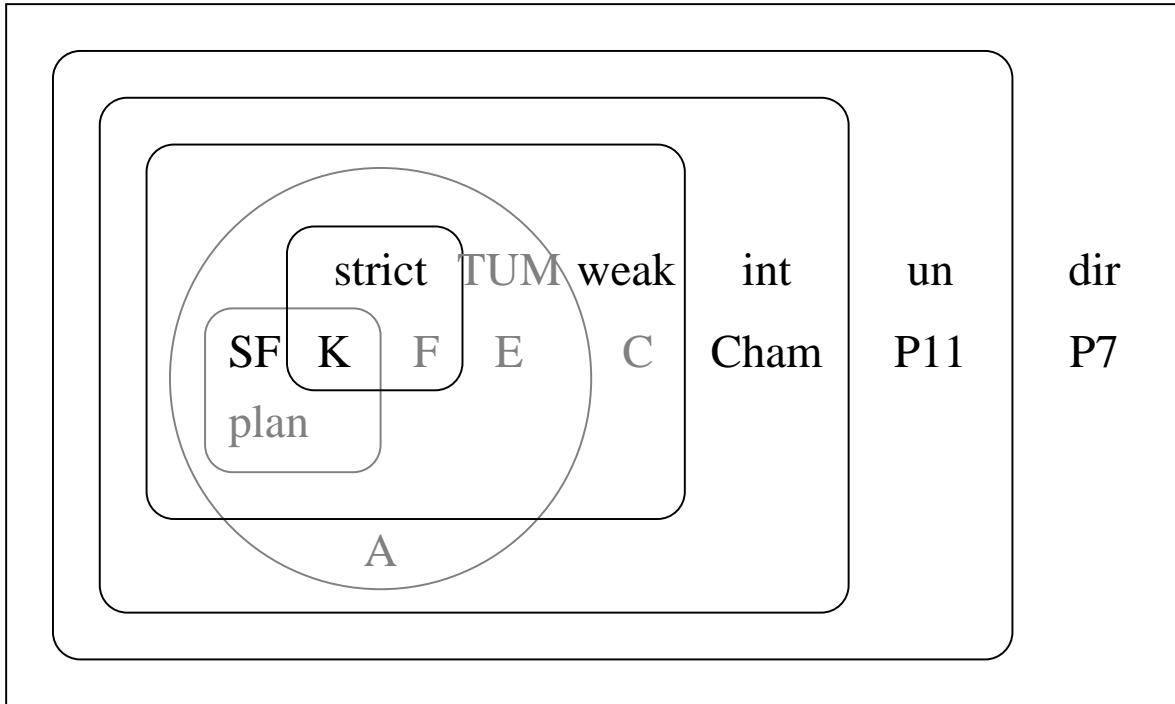


Figure 1: The inclusion diagram of cycle bases and the complexity status of their minimum weight cycle basis problems.

1 Introduction

Cycles in graphs play an important role in many applications, e.g., analysis of electrical networks, analysis of chemical and biological pathways, periodic scheduling, and graph drawing. From a mathematical point of view, cycles in graphs have a rich structure. Cycle bases are a compact description of the set of all cycles of a graph and cycle bases consisting of short cycles or, in weighted graphs, of small weight cycles are interesting mathematically and from an application view point. In the applications above, sparse descriptions are to be preferred.

The study of cycle bases dates back to the early days of graph theory; MacLane (1937) gave a characterization of planar graphs in terms of cycle bases. In the last decade, many new results on cycle bases appeared, most notably a classification of different kinds of cycle basis, structural results, a-priori bounds on the length and weight of minimum cycle bases, polynomial time algorithms for constructing exact or approximate minimum cycle bases of some kinds and hardness results for other kinds of minimum cycle bases.

In this paper, we survey these results and prove new ones. Figure 1 shows the landscape of cycle bases. We will introduce different kinds of cycle bases in Sections 2 and 3: Directed cycle bases, undirected cycle bases, integral cycle bases, weakly fundamental cycle basis, totally unimodular cycle bases, strictly fundamental cycle bases, and 2-bases. In section 3, we will characterize the different kinds in terms of properties of their cycle matrix. For example, undirected cycle bases are characterized by the fact that the determinant of their cycle matrix is odd and integral cycle basis are characterized by the fact that their determinant is ± 1 . We will establish the inclusion map and show that the different classes lead to different minimum cycle basis problems. We will also establish many structural results.

Section 4 deals with a-priori length and weight bounds on minimum cycle bases. We will prove results of the following flavor: every graph of n nodes and m edges has a weakly fundamental cycle basis of length $O(m \log m / \log(m/n))$. We will also show that there are graphs that do not have a shorter basis.

In Section 5, we will give polynomial time algorithms for constructing minimum weight directed, undirected and planar cycle basis. We will also discuss approximation algorithms.

Section 6 treats hardness results; in particular, we will show that the minimum cycle basis problem is \mathcal{APX} -hard for weakly fundamental and strictly fundamental bases. Figure 1 summarizes the complexity results. For two classes the complexity is open.

Finally, Section 7 discusses three applications of cycle bases; we will see that they require different kinds of cycle bases. The analysis of electrical circuits can do with any kind of cycle basis, periodic scheduling requires integral cycle bases, and graph drawing needs strictly fundamental bases.

The paper surveys mostly known results, but it also contains several new results. In particular, we give additional structural and characterization results, we obtain tight length bounds for weakly fundamental cycle bases for the full spectrum of graph densities, we give a simplified algorithmic treatment of directed cycle bases, and we present the first algorithms for minimum cycle bases in the presence of negative edges. In all sections, we also pose open problems.

2 Definitions

An (*undirected*) graph is a pair $G = (V, E)$, where V is a finite set, and E is a family of unordered pairs of elements of V . The elements of V are called *vertices* or *nodes* and the elements of E are called *edges*. An edge $e = \{v, w\}$ is *incident to* the vertices v and w and v and w are the endpoints of e . The same pair $\{v, w\}$ may occur several times in E ; we refer to a pair occurring more than once as a *multiple edge*. Graphs without multiple edges are called *simple*. An edge of the form $\{v, v\}$ is called a *loop*. The *degree* $\deg(v)$ of a vertex v is the number of times v occurs as an endpoint of an edge. Observe, that a loop $\{v, v\}$ contributes two to the degree of v . We use $\delta(v)$ to denote the set of edges incident to v ; a loop $\{v, v\}$ appears twice in $\delta(v)$.

A (*directed*) graph is a pair $D = G = (V, A)$, where V is a finite set, and A is a family of ordered pairs of elements of V . The elements of V are called the vertices or nodes of G , and the elements of A are called the (*directed*) edges or *arcs* of G . We use $G = (V, E)$ to denote directed and undirected graphs and $D = (V, A)$ to denote directed graphs. The vertices v and w are called the *tail* and *head* of the arc $e = (v, w)$, respectively; e is said to leave v and to enter w ; it is incident to v and w . Every directed graph gives rise to a unique undirected graph, by ignoring the orientation of the edges. The notions *multiple edge*, *simple*, and *loop* are defined analogously as for undirected graphs. The *outdegree* $\text{outdeg}(v)$ and *indegree* $\text{indeg}(v)$ of a vertex v is the number of times v occurs as the tail, respectively, head of an edge. Observe, that a loop (v, v) contributes one to the indegree and the outdegree of v . We use $\delta^+(v)$ and $\delta^-(v)$ for the edges leaving and entering v , respectively.

We use n and m to denote the number of nodes and edges or arcs, respectively, i.e., $n = |V|$ and $m = |E|$ or $m = |A|$. We use the notation $e = vw$ to denote both directed and undirected edges, i.e., the notation stands for the directed edge (v, w) and the undirected edge $\{v, w\}$. Every undirected graph G gives rise to a directed graph D by orienting the edges arbitrarily; we call D an *orientation* of G . In this way, we can view every graph as directed.

A *subgraph* $G' = (V', E')$ of G is a graph with $V' \subseteq V$ and $E' \subseteq E$. If V' is a subset of V , $G - V'$ denotes the graph obtained by removing all vertices in V' and their incident edges from G . A *path* P from v to w in G is a subgraph of G with $V' = \{v_0 = v, v_1, \dots, v_k = w\}$ with $v_i \neq v_j$ and $E' = \{x_0x_1, x_1x_2, \dots, x_{k-1}x_k\}$. We write $P(v, w)$ if we want to emphasize that P is a path from v to w . The *length of a path* is the number of its edges. An undirected graph is *connected* if there exists a path from any vertex to every other vertex. A vertex v in a connected graph G is called an *articulation point*, or *cut vertex*, if $G - v$ is disconnected. An undirected graph is *biconnected* if it has no articulation point. A *directed graph is connected* if the underlying undirected graph is connected. Any maximal connected subgraph of G is called a *connected component*. A graph T is a *tree* if it is connected and has $n - 1$ edges. A subgraph G' of a connected graph G is called a *spanning tree* if it constitutes a tree on all vertices in G . If G is not connected, any union of spanning trees for each connected component is called a *spanning forest*.

A *cycle* in an undirected graph is a vector $C \in \mathbb{Z}_2^E$ such that $\sum_{e \in \delta(v)} C_e = 0$ for any vertex v ; here C_e denotes the component of C indexed by e , $\mathbb{Z}_2 = GF(2) = \{0, 1\}$ is the field of two elements, and addition is addition in \mathbb{Z}_2 . A cycle may alternatively be viewed as a set of edges; e belongs to C iff $C_e = 1$. We use C to denote the vector in \mathbb{Z}_2^E , the corresponding subset of E , and also the subgraph (V', C) where V' is the set of vertices having at least one edge in E incident to it. Instead of C_e , we will also write $C(e)$. We prefer the latter notation when $C = C_i$ belongs to an indexed family of cycles. A cycle is an *even* or *Eulerian* subgraph,

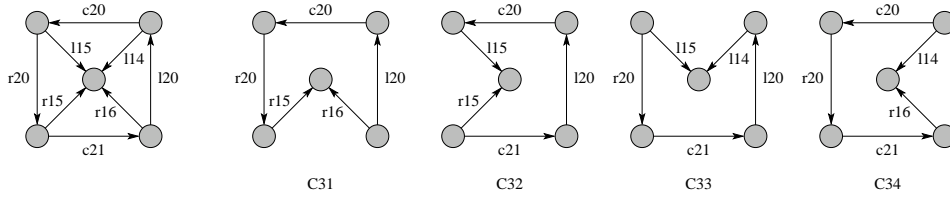


Figure 2: The figure shows an orientation D of the undirected wheel graph W_5 , and four circuits C_1 to C_4 in D . The edges of D are numbered from e_1 to e_8 . The circuit C_1 uses the edges e_1 , e_2 , e_3 , and e_5 in forward direction and the edge e_8 in backward direction. Thus $C_1 = (1, 1, 1, 0, 1, 0, 0, -1)$. The cycles C_1 to C_4 form a directed cycle basis of D . The cycle C consisting of edges 1 to 4 is represented as: $C = (1, 1, 1, 1, 0, 0, 0, 0) = (C_1 + C_2 + C_3 + C_4)/3$. Let G be the underlying undirected graph, let $\pi(C_i)$ be the undirected cycle corresponding to C_i , and let $\pi(C)$ be the undirected cycle corresponding to C . Then $\pi(C_1) = (1, 1, 1, 0, 1, 0, 0, 1)$ and $\pi(C) = \pi(C_1) \oplus \pi(C_2) \oplus \pi(C_3) \oplus \pi(C_4)$. The circuits $\pi(C_1)$ to $\pi(C_4)$ form an undirected cycle basis of G .

The set $\{C_1, C_2, C_3, 2C_4\}$ is also a directed cycle basis of D . However, $\pi(2C_4) = \mathbf{0}$ and hence $\{\pi(C_1), \pi(C_2), \pi(C_3), \pi(2C_4)\}$ is *not* an undirected cycle basis of G . There are less trivial reasons for a directed cycle basis not projecting into an undirected cycle basis.

i.e., every vertex has even degree in C . Conversely, any even subgraph is a cycle. A cycle is a *circuit* if it is connected and every one of its vertices has degree two. The set

$$\mathcal{UC}_G := \{C \mid C \text{ is a cycle of } G\}$$

forms a vector space over $GF(2)$; \mathcal{UC}_G is called the (*undirected*) *cycle space* of G . Addition in this vector space corresponds to symmetric set difference. An *undirected cycle basis* of G is a set of circuits forming a basis of the cycle space of G .

A (*directed*) *cycle* C in a directed graph is a vector in \mathbb{Z}^E such that for any vertex v we have $\sum_{e \in \delta^+(v)} C_e = \sum_{e \in \delta^-(v)} C_e$. Here, addition denotes addition in \mathbb{Q} , the field of rational numbers. In other contexts, cycles are sometimes referred to as *circulations* and the constraint $\sum_{e \in \delta^+(v)} C_e = \sum_{e \in \delta^-(v)} C_e$ is called flow conservation. Cycles in directed graphs may use arcs in *forward* ($C_e > 0$) or *backward* ($C_e < 0$) direction. If any arc is replaced by C_e copies of itself and, in addition, the direction of all arcs e with $C_e < 0$ is reversed, then we end up with a digraph in which the indegree of every vertex is equal to its outdegree. Observe that if C is a cycle, then $-C$ is a cycle, too, though a different one. The set

$$\mathcal{DC}_G := \{\lambda C \mid \lambda \in \mathbb{Q} \text{ and } C \text{ is a directed cycle of } G\}$$

forms a vector space over \mathbb{Q} ; \mathcal{DC}_G is called the (*directed*) *cycle space* of G . A cycle is *simple* if $C_e \in \{-1, 0, +1\}$ for all e , and a simple cycle is an *circuit* if it is connected and for any v , there are exactly two edges¹ e incident to v with $C_e \neq 0$. A *directed cycle basis* is a set of circuits forming a basis of the cycle space.

Let D be a directed graph and let G be the underlying undirected graph. We may refer to G as $G(D)$. For any directed cycle C of D , let $\pi(C) := (C_e \bmod 2)_{e \in E}$. Then $\pi(C)$ is an undirected cycle in G . We call $\pi(C)$ the *projection* of C .

¹Recall, that a self-loop counts twice.

Figure 2 illustrates all these definitions. In addition, it provides a first example showing that directed cycle bases do not necessarily project onto undirected cycle basis. However, a set of dependent cycles projects into a set of dependent cycles. Let $C_i, i \in I$, be a family of dependent directed cycles. Then $\sum_{i \in I} \lambda_i C_i = \mathbf{0}$ with $\lambda_i \in \mathbb{Q}$, not all zero. Here $\mathbf{0}$ denote the zero-vector in \mathbb{Z}^E . We may assume $\lambda_i \in \mathbb{Z}$ and not all even. Then $\sum_{i \in I} (\lambda_i \bmod 2) \pi(C_i) = \mathbf{0} \bmod 2$ and at least one coefficient $\lambda_i \bmod 2$ is nonzero. Thus the $\pi(C_i), i \in I$, are dependent.

We use $+$ and Σ to denote addition in \mathbb{Q} and in $GF(2)$ (and also in $GF(p)$ for prime p). The distinction will usually be clear from the context. If both fields occur in the same argument, as in the paragraph above, we will emphasize the difference by the additional operator $\bmod 2$.

We may also lift undirected cycles from an undirected graph G to an orientation D of it. Let C' be a connected undirected cycle² consisting of k edges. Since an undirected cycle is a Eulerian subgraph of G , there is a closed traversal (e_0, \dots, e_{k-1}) of the edges of C' , i.e., $e_i = \{v_i, v_{i+1}\}$ for $0 \leq i < k$ and $v_0 = v_k$. This traversal defines a simple cycle C in D ; we have $C_e = 0$ if C' does not contain e and $C_e = +1$ (-1) if the traversal uses e in forward (backward) direction. We call C a *lifting* or *orientation* of C' . Observe that the lifting is not unique; it depends on the choice of walk. For a circuit C' the lifting is unique up to sign. Clearly, if C' lifts to C then C projects to C' .

A *weighted graph* is a graph together with a weight $w : E \rightarrow \mathbb{R}$. If the graph is unweighted, we set $w : E \rightarrow 1$ and call w the *uniform weight function*. The weight of a set of edges is the sum of the weights of its members. The *weight of a cycle* C is

$$w(C) := \sum_e |C_e| w(e)$$

and the *length of a cycle* is

$$|C| := \sum_e |C_e| .$$

In an unweighted graph, weight and length are identical. The *weight of a cycle basis* B is the sum of the weights of its cycles, i.e.,

$$w(B) = \sum_{C \in B} w(C) .$$

A *minimal cycle basis*, or *MCB*, of G is a cycle basis with minimal weight. We assume that there are no simple cycles of negative weight; such weight functions are called *conservative*. For most of our algorithms, we need to assume that weights are nonnegative, i.e., $w : E \rightarrow \mathbb{R}^+$.

We close this section with a first theorem. Every graph has a (directed and undirected) cycle basis and the dimension of the (directed or undirected) cycle space is given by the graph's *cyclomatic number* $\nu := m - n + CC$, where CC denotes the number of connected components of G . On the way, we get to know a particularly simple set of cycles, the fundamental cycles with respect to a spanning forest. Let G be an (undirected or directed) graph and let T be a spanning forest of G . For any nontree edge e , let $C_T(e)$ be the circuit consisting of e and the tree path connecting the endpoints of e . In the case of a directed graph, we use e in forward direction and traverse the tree path from the head of e to the tail of e , see Figure 3. We call $C_T(e)$ the *fundamental circuit* defined by T and e .

Lemma 2.1. *Let G be a graph and let T be any spanning forest of G . Let C be a cycle that uses only edges in T , i.e., $C_e = 0$ for $e \notin T$. Then $C = \mathbf{0}$.*

²Components are lifted independently.

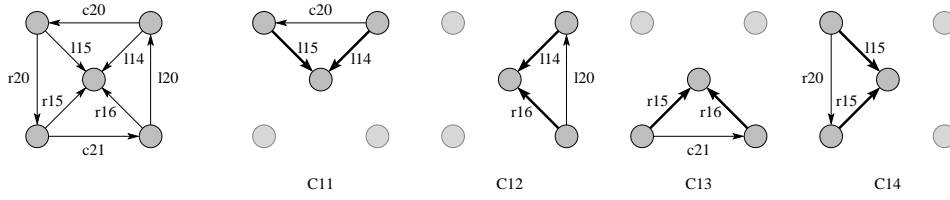


Figure 3: The figure shows an orientation D of the undirected wheel graph W_5 and four circuits C_1 to C_4 in D . The edges of are numbered from e_1 to e_8 . The **bold** edges $\{e_5, e_6, e_7, e_8\}$ form a spanning tree T of D . The circuit C_1 is induced by the nontree edge e_2 and uses the edges e_2 and e_6 in forward direction and the edge e_7 in backward direction. Thus $C_1 = (0, 1, 0, 0, 0, 1, -1, 0)$. As the cycles C_2 , C_3 , and C_4 are obtained in an analogous way, the set $\{C_1, C_2, C_3, C_4\}$ is a strictly fundamental cycle basis of D .

Proof. View G and C as undirected. Then C is an even subgraph of G . Since C uses only edges in T , it is an even subgraph of T . Even subgraphs of forests are empty. \square

Lemma 2.2. *Let B be a set of cycles in G and let T be any spanning forest of G . For any cycle $C \in B$, let C' be its restriction to $N := E \setminus T$. The cycles are linearly independent if and only if their restrictions to N are linearly independent.*

Proof. Clearly, linear dependence of the cycles implies linear dependence of their restrictions. Conversely, assume that there is a nontrivial linear combination of the restrictions that yields the zero vector, i.e., $\sum_{C \in B} \lambda_C C' = \mathbf{0}_N$. Here $\mathbf{0}_N$ denotes the zero vector over index set N . Then $\sum_{C \in B} \lambda_C C$ is a cycle that uses only tree edges and hence is equal to $\mathbf{0}$. \square

Thus, we may restrict attention to the restricted incidence vectors when discussing questions of linear independence.

Theorem 2.3 (Dimension of the Cycle Space of a Graph). *The dimension of the undirected and directed cycle space of a graph G is given by its cyclomatic number*

$$\nu = m - n + CC,$$

where CC denotes the number of connected components of G . Moreover, if T is any spanning forest of G , the set of fundamental circuits with respect to T forms a basis.

Proof. The number of fundamental circuits is equal to ν , because a connected component with m' edges and n' vertices contributes $m' - (n' - 1)$ fundamental cycles. Let N be the set of nontree edges. The fundamental cycles are clearly independent since any edge $e \in N$ is contained precisely in $C_T(e)$. It remains to prove that the set of fundamental circuits spans all cycles. Let C be an arbitrary directed cycle and consider the cycle

$$\tilde{C} := \sum_{e \in N} C_e C_T(e) .$$

We claim that $C = \tilde{C}$. Indeed, for any $e \in N$, we have $\tilde{C}_e = C_e$ and hence $C - \tilde{C}$ is a cycle using only edges of T . Thus $C - \tilde{C} = \mathbf{0}$. \square

The following Lemma is a first step towards clarifying the relationship of directed and undirected cycle bases.

Lemma 2.4. *Let D be a directed graph, let $B = \{C_1, \dots, C_\nu\}$ be a set of circuits in D , let G be the underlying undirected graph, and let $\pi(B) = \{\pi(C_1), \dots, \pi(C_\nu)\}$. If $\pi(B)$ is an undirected cycle basis of G then B is a directed cycle basis of D .*

Proof. We have already shown that a set of dependent cycles projects into a set of dependent cycles. Hence $\pi(B)$ being an undirected cycle basis implies that the cycles in B are independent. Also, ν must be equal to the cyclomatic number of D , since $\pi(B)$ is a basis. \square

3 Classification of Cycle Bases

We present seven classes of cycle bases and provide characterizations for them. We will show that each class gives rise to its own minimum cycle basis problem. The complexity of the minimum cycle basis problem differs widely. For three classes the problem is polynomial time, for two classes the problem is NP-complete, and for two classes the status is unknown. This section is mainly based on Liebchen and Rizzi (2007); the missing proofs can be found there.

Definition 3.1 (classes of cycle bases). *A directed cycle basis (D-basis) $B = \{C_1, \dots, C_\nu\}$ of a directed graph D is called a*

1. *undirected or U-basis, if the projections $\pi(C_i)$ of the basic circuits C_i onto the underlying undirected graph $G(D)$ constitute a cycle basis of $G(D)$;*
2. *integral or I-basis, if each cycle C of D can be written as an integer linear combination of circuits in B , i.e.*

$$\exists \lambda_i \in \mathbb{Z} : C = \lambda_1 C_1 + \dots + \lambda_\nu C_\nu;$$

3. *zero-one or³ TUM-basis, if each cycle C' of $G(D)$ has an orientation C that can be written as a linear combination with coefficients in $\{-1, 0, +1\}$ of circuits in B , i.e.*

$$\exists \lambda_i \in \{-1, 0, +1\} : \gamma_C = \lambda_1 C_1 + \dots + \lambda_\nu C_\nu;$$

4. *weakly fundamental or W-basis, if there exists some permutation σ such that*

$$C_{\sigma(i)} \setminus (C_{\sigma(1)} \cup \dots \cup C_{\sigma(i-1)}) \neq \emptyset, \forall i = 2, \dots, \nu;$$

5. *strictly fundamental or F-basis, if there exists some spanning forest $T \subseteq E$ such that $B = \{C_T(e) \mid e \in E \setminus T\}$, where $C_T(e)$ denotes the unique circuit in $T \cup \{e\}$; and*
6. *planar, or 2-basis, if each arc is contained in at most two basic circuits and the basis is undirected.*

Figure 4 visualizes the relationship between these classes: Valid inclusions are established in Theorem 3.4, and examples for the non-emptiness of the regions will be provided in Section 3.4.

3.1 Existence.

Except for 2-bases, every graph has a basis of each type. This follows from the fact that every graph has a strictly fundamental cycle basis and that all other classes generalize fundamental cycle bases. In contrast, MacLane (1937) established that a graph has a 2-basis, if and only if it is planar.

³It will become clear in Theorem 3.4 why zero-one bases are called totally unimodular (TUM).

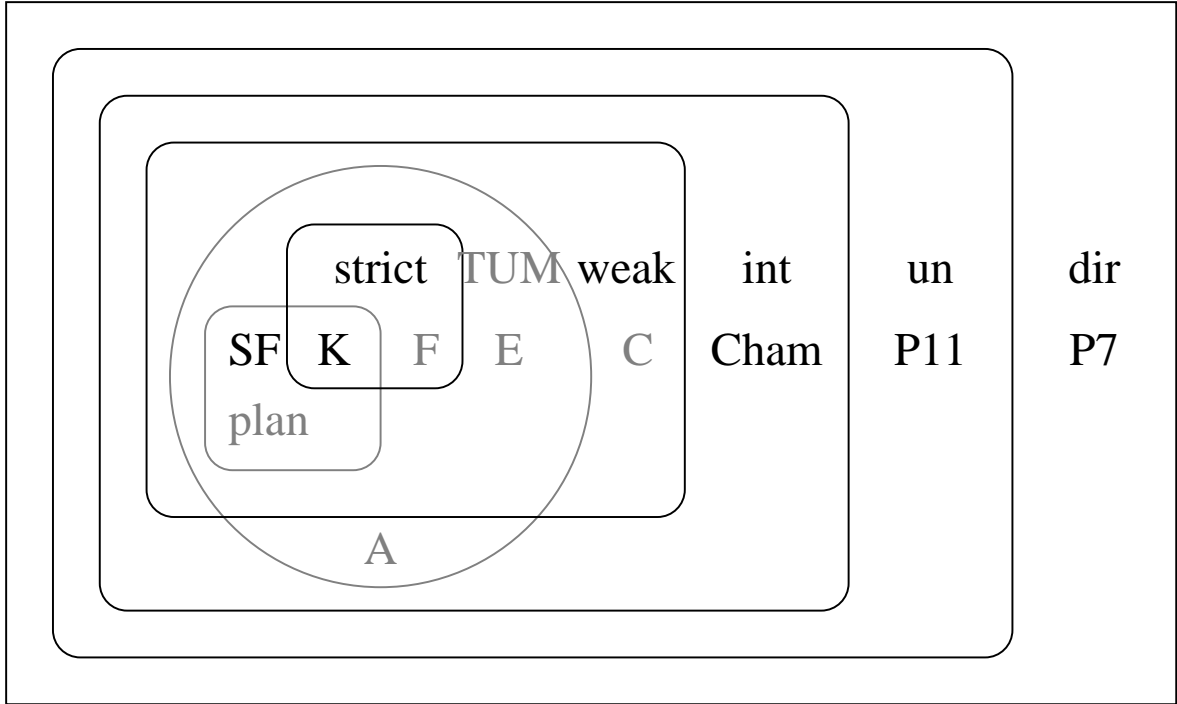


Figure 4: Map of directed cycle bases: Ex. 3.X refers to an example that is discussed in detail later in this section, $K_{3,3}$ refers to a weighted version of the complete bipartite graph on 3×3 vertices, $P_{7,2}$ is a weighted version of a generalized Petersen graph, V_8 is Wagner’s graph (cf. Section 6), $F_{3,2}$ is a fan graph on five vertices, and G_1 is a simple graph on eight vertices, see Liebchen and Rizzi (2007).

3.2 Characterizations.

We define the cycle matrix corresponding to a basis and show that the different classes of cycle bases can be characterized in terms of simple properties of this matrix. An important property will be the determinant of the cycle basis. The *cycle matrix* corresponding to a D-basis B of D is a $m \times \nu$ matrix whose columns are the incidence vectors of the basic circuits. The cycle matrix is determined up to a permutation of the rows and columns.

The cycle matrix Φ of a fundamental basis has a particularly simple form. Let T be a spanning forest and let N be the set of co-tree arcs. Then, for a suitable permutation of the columns, the $\nu \times \nu$ submatrix Φ' selected by the rows corresponding to co-tree arcs is the identity matrix.

Lemma 3.1 (Liebchen (2003)). *Let B be a directed cycle basis of a directed graph G and let Γ be the corresponding cycle matrix. A $\nu \times \nu$ submatrix Γ' of Γ is nonsingular if and only if the rows of Γ' correspond to the co-tree arcs of some spanning forest of D .*

Proof. To prove sufficiency, consider a spanning forest T of D , and let Φ be the cycle matrix of the fundamental basis with respect to T . As B is a directed cycle basis, any fundamental cycle is a linear combination of cycles in B . Thus there is a matrix $R \in \mathbb{Q}^{\nu \times \nu}$ with $\Phi = \Gamma R$. The restriction of Φ to the co-tree arcs of T is the identity matrix. Hence, R is the inverse of Γ' .

Conversely, assume that the rows which are *not* in Γ' do not form a spanning forest. Then there is a circuit C consisting only of such arcs. As B is a D-basis, we have $C = \Gamma x_C$ for some x_C ; clearly $x_C \neq \mathbf{0}$. Restricting to the rows indexing Γ' yields $\mathbf{0} = \Gamma' x_C$, and hence Γ' is singular. \square

Lemma 3.2 (Liebchen (2003)). *Let B be a D-basis, let Γ be its cycle matrix, and let A_1 and A_2 be two nonsingular $\nu \times \nu$ submatrices of Γ . Then $\det A_1 = \pm \det A_2$.*

Proof. By Lemma 3.1, the rows of A_1 correspond to the co-tree arcs of some spanning forest T . Let Φ be the cycle matrix of the fundamental basis with respect to T . Then $\Phi A_1 = \Gamma$, cf. Berge (1962) or see the proof of Theorem 2.3. Considering only the rows of A_2 , we obtain $\Phi' A_1 = A_2$; here Φ' is the submatrix of Φ selected by the rows of A_2 . Since Φ is totally unimodular (Schrijver (1986)), we have $\det \Phi' = \pm 1$ and hence $\det A_1 = \pm \det A_2$. \square

The above lemma allows us to define the determinant of a directed cycle basis.

Definition 3.2 (Determinant of a set of ν oriented circuits). *Let B denote a set of ν circuits in a directed graph D . Consider the matrix Γ with the incidence vectors of B as columns. Let Γ' be the $\nu \times \nu$ submatrix of Γ that arises when deleting the arcs of some spanning forest of D . We define*

$$\det B := |\det \Gamma'|.$$

The determinants of directed cycle bases are positive integers. The value of the determinant is invariant under reorienting arcs of D or reorienting circuits of B , because this simply translates to multiplying a row or column by minus one. Thus, starting with a cycle basis of an undirected graph G , orienting the edges of G arbitrarily, and choosing one of the two orientations for each circuit, always results in the same determinant.

There are directed cycle bases with quite large determinant although the number of vertices of the corresponding graphs is comparatively small. More precisely there is an infinite family of graphs with cycle bases with determinants that grow linearly in the number of vertices.

Lemma 3.3. *Consider the generalized Petersen graph $P_{n,2}$ with $n \geq 5$ odd. Let \mathcal{C} denote the set of circuits each of them containing exactly one inner edge, $n - 2$ outer edges and two spokes. \mathcal{C} together with the inner circuit C_I forms a cycle basis of $P_{n,2}$ and its determinant equals $n - 2$.*

Proof. $P_{n,2}$ consists of $2n$ vertices and $3n$ edges. Therefore every cycle bases has to consist of $n + 1$ cycles which is indeed the number of considered circuits. Additionally it should be mentioned that the inner circuit C_I is indeed a simple cycle since n is odd.

Now let T be a spanning tree of $P_{n,2}$ made up of all but one inner edges and all spokes. Consider the square submatrix Γ' of the cycle matrix Γ obtained by deleting the rows corresponding to T . The co-tree edges and the circuits in $\mathcal{C} \cup \{C_I\}$ can be oriented and permuted

such that

$$\Gamma' = \left(\begin{array}{cccccc|c} 1 & \cdots & \cdots & 1 & 0 & 0 & 0 \\ 0 & 1 & \ddots & \ddots & 1 & 0 & \\ 0 & 0 & 1 & \ddots & \ddots & 1 & \vdots \\ 1 & 0 & 0 & 1 & \ddots & \vdots & \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots & \\ 1 & \cdots & 1 & 0 & 0 & 1 & 0 \\ \hline * & & \cdots & & & * & 1 \end{array} \right)$$

where the last column and the last row corresponds to the inner circuit and the inner edge, respectively. The determinant of Γ' equals the determinant of its $n \times n$ submatrix obtained by deleting the last row and column. The resulting matrix is a *circulant matrix* whose first row has $n - 2$ consecutive ones followed by two zeros. The entries of every other row result from the row above by a circular shift to the right. We have

$$\det \Gamma' = n - 2,$$

see Ueckerdt (2008) for the calculation of the determinant. \square

Theorem 3.4. *Let B be a directed cycle basis with cycle matrix Γ .*

1. *B is undirected, if and only if $\det B$ is odd.*
2. *B is integral, if and only if $\det B$ is one.*
3. *B is zero-one if and only if Γ is totally unimodular.*
4. *B is weakly fundamental, if and only if Γ can be permuted as to have a regular upper triangular $\nu \times \nu$ matrix in its last ν rows.*
5. *B is strictly fundamental, if and only if Γ can be permuted as to have the $\nu \times \nu$ unit matrix in its last ν rows.*
6. *B is a 2-basis, if and only if B is an undirected cycle basis and Γ has at most two nonzero entries per row.*

Proof. Case 1. The projections $\pi(C_i)$ of the basic circuits are linearly independent if $\pi(\Gamma)$ has full rank, i.e., if there is a square submatrix $\pi(\Gamma')$ with nonzero determinant over $\text{GF}(2)$. The value of the determinant is $(\det \Gamma') \bmod 2$. We conclude that B is undirected if and only if $\det B$ is odd.

Case 2. Let T be some spanning forest, and let Γ' be the square submatrix of Γ indexed by the co-tree arcs of T .

Let Φ be the cycle matrix of the fundamental basis with respect to T . Since B is integral, there is an integral $\nu \times \nu$ matrix R such that $\Phi = \Gamma R$. Restriction to the co-tree arcs of T yields $I = \Gamma' R$. We have $\det \Gamma' \in \mathbb{Z}$ and $\det R \in \mathbb{Z}$, because both matrices are integral. Thus $(\det \Gamma') \cdot (\det R) = 1$ and hence $\det \Gamma' = \pm 1$.

Let C be an arbitrary circuit. The representation x_C of C in terms of B satisfies $C = \Gamma x_C$. Restriction to the co-tree arcs of T yields $C' = \Gamma' x_C$ or $x_C = (\Gamma')^{-1} C'$. The inverse of Γ' is integral, by Cramer's rule and since $\det \Gamma' = \pm \det B = \pm 1$. Thus $x_C \in \mathbb{Z}^\nu$.

Case 3. A matrix is totally unimodular if and only if for any subset I of its columns there are coefficients $\lambda_i \in \{-1, +1\}$ such that $\sum_{i \in I} \lambda_i C_i$ is a vector with entries in $\{-1, 0, +1\}$, see (Schrijver, 1986, Theorem 19.3).

Assume first that B is a zero-one basis. In particular, B is an integral cycle basis and hence $\{\pi(C_i) \mid C_i \in B\}$ is an undirected basis of $G(D)$. Let I be an arbitrary subset of the columns of Γ and consider the \mathbb{Z}_2 -sum of the projections of the circuits in I , and call the resulting cycle C' ,

$$\sum_{i \in I} \pi(C_i) = C'.$$

Since B is a zero-one basis, C' has an orientation C that can be written as a linear combination with coefficients $\lambda_i \in \{-1, 0, +1\}$ of the circuits in B , i.e.,

$$\sum_{i=1}^{\nu} \lambda_i C_i = C.$$

Projecting this equation onto \mathbb{Z}_2 , we obtain

$$\sum_{i=1}^{\nu} |\lambda_i| \pi(C_i) = C'.$$

Since the representation of C' with respect to the basis of the projections of the circuits in B is unique, we have that λ_i is nonzero if and only if $i \in I$. Thus, in the TUM characterization, C is the desired linear combination of the columns selected by I .

Assume conversely that Γ is totally unimodular. Then $\det B = 1$ and hence $\{\pi(C_i) \mid C_i \in B\}$ is a basis of $G(D)$. Let C' be any cycle in $G(D)$. Then $C' = \sum_{i \in I} \pi(C_i) \pmod{2}$ for some index set $I \subseteq \{1, \dots, \nu\}$. Since Γ is totally unimodular, there are coefficients $\lambda_i \in \{-1, +1\}$ such that $\sum_{i \in I} \lambda_i C_i$ is a vector C with components in $\{-1, 0, +1\}$. Clearly, $\pi(C) = C'$ and hence C is the desired orientation of C' .

Case 4. Order the columns of Γ such that $C_{\sigma(i)}$ is in the i -th column for $1 \leq i \leq \nu$. Order the rows of Γ such that an arc a with $a \in C_{\sigma(i)} \setminus (C_{\sigma(1)} \cup \dots \cup C_{\sigma(i-1)})$ corresponds to row $\nu - 1 + i$.

Case 5. This is nothing but a reformulation of Sysłó's Sysłó (1979) characterization of a strictly fundamental cycle basis B , namely any circuit in the basis contains an arc that is contained in no other circuit of the basis.

Case 6. This is nothing but a reformulation of the definition of 2-bases. □

The determinant of a set of ν circuits can be computed over any field k . For directed bases the determinant is nonzero in \mathbb{Q} , for undirected bases the determinant is nonzero in $GF(2)$. We therefore call directed bases also \mathbb{Q} -bases and undirected bases $GF(2)$ -bases. We call a directed basis a $GF(p)$ -basis, where p is a prime, if its determinant is nonzero modulo p .

Theorem 3.4 establishes most inclusions of Figure 4: Every fundamental basis is both weakly fundamental and totally unimodular, every weakly fundamental or totally unimodular basis is integral, every integral basis is undirected, and every undirected basis is directed. We shall next relate 2-bases to the other classes.

Lemma 3.5. *Every 2-basis is totally unimodular and weakly fundamental.*

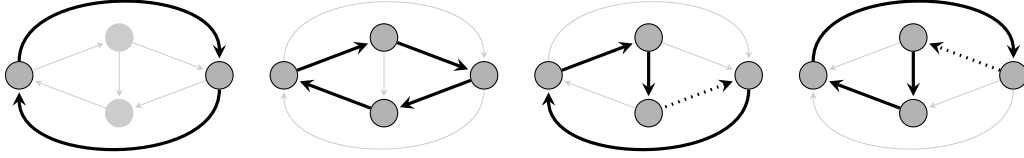


Figure 5: A graph and a directed cycle basis. For each of the four circuits, the arcs belonging to the circuit are shown in bold. Arcs used in reversed direction are shown dotted. Every arc is used in exactly two circuits. The determinant of this basis is two. Thus the basis is not totally unimodular. Also, since each arc is used in exactly two circuits, the basis is not weakly fundamental.

Proof. Let $B = \{C_1, \dots, C_\nu\}$ be a 2-basis of G . MacLane (1937) showed that a graph having a 2-basis is planar and that moreover the basic circuits correspond to the bounded face cycles of some planar embedding of G . Orient the edges of G arbitrarily and let the C_i 's correspond to counterclockwise traversals of the face cycles. Then every row of Γ has at most two nonzero entries; if there are two nonzero entries, one is $+1$ and one is -1 . Thus Γ is totally unimodular ((Schrijver, 1986, page 274)).

We next show that B is weakly fundamental. Let $C = \{e_1, \dots, e_k\}$ be the boundary of the infinite face of G . For $i = 1, 2, \dots, k$, denote by C_{e_i} the unique circuit in B that contains $e_i \in C$. In the first iteration, we define

$$C_{\sigma(v)} = C_{e_1}, \quad C_{\sigma(v-1)} = C_{e_2}, \quad \dots, \quad C_{\sigma(v-k+1)} = C_{e_k}.$$

Then, we remove the edges of C from G and proceed in the same way for the 2-connected components of the remaining graph. \square

We required a 2-basis to use every arc at most twice and to be undirected. Figure 5 shows a graph and a directed basis that uses every arc exactly twice and is neither totally unimodular nor weakly fundamental (Tomasz Jurkiewicz, personal communication).

Open Problem 1. *The definition of zero-one bases may seem strange. It would be equally natural to require that every circuit (every simple cycle) is a linear combination of the basic circuits with coefficients in $\{-1, 0, +1\}$. How do these definitions relate?*

3.3 Simple Examples

Figure 6 presents three cycle bases for the wheel graph W_5 : the strictly fundamental cycle basis $B_1 = \{C_{11}, C_{12}, C_{13}, C_{14}\}$, which is also a 2-basis, the weakly fundamental cycle basis $B_2 = \{C_{21}, C_{22}, C_{23}, C_{24}\}$, and the undirected basis $B_3 = \{C_{31}, C_{32}, C_{33}, C_{34}\}$; the latter is *not* integral. The strictly fundamental cycle basis B_1 corresponds to the spanning tree $T = \{e_1, e_2, e_3, e_4\}$. The corresponding cycle matrices are as follows:

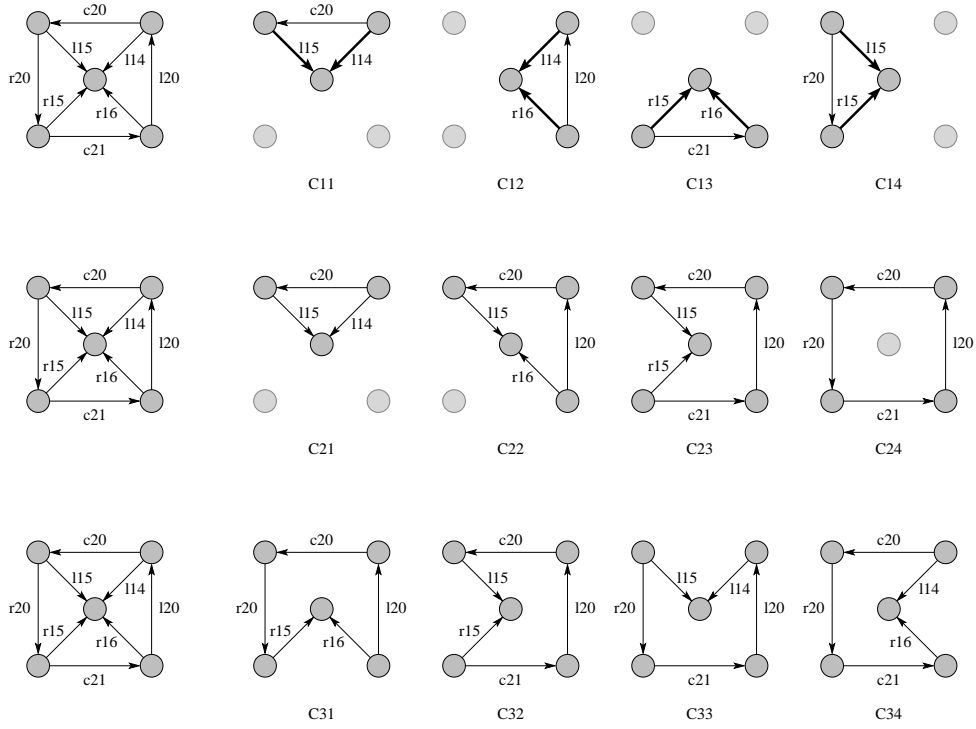


Figure 6: Examples of a strictly fundamental cycle basis that is also a 2-basis, a weakly fundamental cycle basis, and a non-integral cycle basis in the wheel graph W_5 . The latter originates from Hartvigsen and Zemel (1989).

$$\Gamma_1 = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 1 & 0 & 0 & -1 \\ 0 & 0 & -1 & 1 \\ 0 & -1 & 1 & 0 \\ \hline 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \Gamma_2 = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ \hline 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \Gamma_3 = \begin{pmatrix} 0 & 0 & 1 & -1 \\ 0 & 1 & -1 & 0 \\ 1 & -1 & 0 & 0 \\ -1 & 0 & 0 & 1 \\ \hline 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix}.$$

The first four rows correspond to the arcs of T and the last four rows correspond to the co-tree arcs. In Γ_1 , every row has at most two nonzero entries and the last four rows constitute a 4×4 unit matrix. Thus B_1 is a 2-basis and is strictly fundamental. In Γ_2 , the last four rows constitute a regular upper triangular matrix and so B_2 is weakly fundamental. Finally, in Γ_3 the determinant of the submatrix formed by the last four rows has determinant three. Hence B_3 is undirected but *not* integral. As a consequence, it cannot be weakly fundamental either, and thus its rows and columns cannot be permuted as to provide a triangular matrix. A direct demonstration that B_3 is not integral is provided by the representation of the circuit C_{24} as a linear combination of the basis B_3 , namely

$$C_{24} = \frac{1}{3}C_{31} + \frac{1}{3}C_{32} + \frac{1}{3}C_{33} + \frac{1}{3}C_{34}.$$

3.4 Variants of the MCB Problem

Each of our classes of cycle bases induces its own variant of the MCB problem. Let D be a directed graph and let \mathcal{B} be a class of cycle bases of D . A minimum (weight) cycle basis of class \mathcal{B} is a basis $B' \in \mathcal{B}$ such that

$$w(B') = \min\{w(B) \mid B \in \mathcal{B}\}.$$

For instance, in the minimum strictly fundamental cycle basis (MFCB) problem we aim at finding a spanning forest in D such that the sum of the weights of its induced fundamental circuits is as small as possible.

Our seven classes define seven different minimum cycle basis problems, i.e., for any two distinct classes \mathcal{B}_1 and \mathcal{B}_2 there is a directed graph D and a weight function w such that

$$\min\{w(B) \mid B \in \mathcal{B}_1\} \neq \min\{w(B) \mid B \in \mathcal{B}_2\}.$$

In the sequel, we show some of these differences; for the others, we refer our readers to Liebchen and Rizzi (2007) and to Fig. 4. In each case, we will exhibit a graph, a weight function, and a basis B , argue that the basis belongs to class \mathcal{B}_1 , and finally show that every basis of class \mathcal{B}_2 must have larger weight. The three graphs that we present next differentiate between the following pairs of the MCB problem:

1. strictly fundamental cycle bases vs. 2-bases and weakly fundamental cycle bases
2. weakly fundamental cycle bases vs. integral cycle bases
3. integral cycle bases vs. undirected cycle bases

A graph that distinguishes between the MCB problems for weakly fundamental and totally unimodular cycle bases is given in Figure 29 of Section 6.

Example 3.1 (F-bases vs. 2-bases and W-bases).

The sunflower graph SF(3) in Fig. 7 contains precisely four circuits with three edges. These are independent and hence constitute its unique minimum cycle basis B . Obviously, B is a 2-basis. And, by Lemma 3.5, B is also weakly fundamental.

However, B is not strictly fundamental, as the center triangle contains no edge that is not contained in any other circuit of the basis; cf. case 5 of Theorem 3.4. This example was inspired by Hubicka and Sysło (1975).

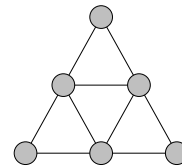


Figure 7: The sunflower graph SF(3) has a unique minimum cycle basis that is a 2-basis.

Example 3.2 (W-bases vs. I-bases).

Champetier (1987) introduced the graph shown in Fig. 8. The graph is specified as a node labelled planar graph. The nodes sharing a label are to be identified. The resulting simple G_{Ch} has 17 vertices and 52 edges. There are precisely 36 triangles in G_{Ch} ; they correspond to the finite faces of the underlying planar graph.

Claim 3.1 (Gleiss (2001b)). *The 36 triangles in G_{Ch} constitute the unique minimum cycle basis B of G_{Ch} . B is integral but not weakly fundamental.*

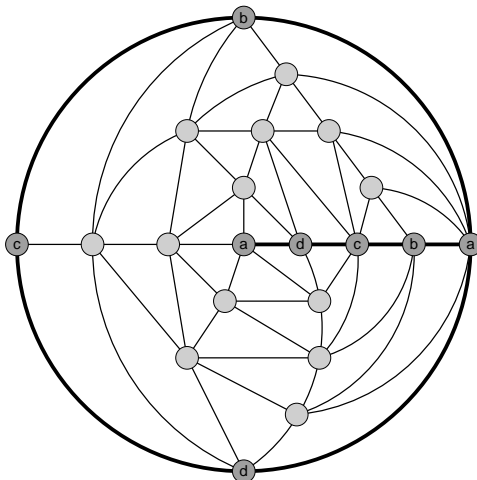


Figure 8: The minimum integral cycle basis of Champetier’s graph Champetier (1987) is unique and not weakly fundamental. Nodes with the same label are to be identified.

Proof. Consider some orientation D of G_{Ch} and orient the circuits in B clockwise, with respect to Fig. 8. Consider the sum C' over \mathbb{Q} of all the triangles, $C' = \sum_{C \in B} C$. In G_{Ch} , all edges except for the ones shown bold in Fig. 8, are part of two triangles. The bold edges belong to three triangles. Thus C' is the 4-circuit that links the labeled vertices. In the visualization of Fig. 8, this translates to following the outer bold circuit clockwise, or following its representation as a path from left to right.

We now construct a new basis B' by replacing an arbitrarily chosen circuit of B by C' . Let Γ and Γ' be the corresponding cycle matrices and consider the transformation matrix R such that $\Gamma' = \Gamma R$. With $R = [r_1, \dots, r_\nu]$, we have $r_i = \mathbf{1}$ for some $i \in \{1, \dots, \nu\}$, and $r_j = e_j$ for all $j \neq i$. Hence, R constitutes a unimodular transformation and thus B and B' have the same determinant.

The cycle basis B' is weakly fundamental; as in the proof of Lemma 3.5 one can construct a suitable ordering of its circuits. Thus $\det B' = 1$ and hence $\det B = 1$. We conclude that B is an integral basis. However B is *not* weakly fundamental as every arc is part of two or three triangles. \square

We mention that the minimum cycle basis B of Champetier’s graph is *not* totally unimodular, cf. Liebchen and Rizzi (2007).

Example 3.3 (U-bases vs I-basis).

Consider the generalized Petersen graph $P_{11,4}$ (cf. Figure 9) with the following weight function

$$w_{ij} = \begin{cases} 4, & \text{if } i \text{ and } j \text{ are outer vertices,} \\ 5, & \text{if } i \text{ and } j \text{ are inner vertices,} \\ 12, & \text{otherwise.} \end{cases}$$

Claim 3.2. $(P_{11,4}, w)$ has precisely 12 circuits of weight at most 44. These constitute the unique minimum cycle basis.

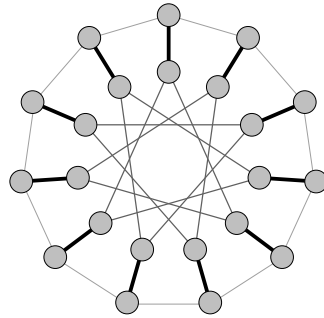


Figure 9: A weighted version of the generalized Petersen graph $P_{11,4}$ has a unique minimum cycle basis that is not integral.

Proof. Any cycle basis consists of $\nu = 33 - 22 + 1 = 12$ circuits. We call the edges e with $w_e = 12$ *spokes* and observe that every circuit contains an even number of spokes. There are only two circuits with no spokes; the outer circuit has weight 44 whereas the inner circuit has weight 55. Any circuit with at least four spokes has weight at least 48.

We classify the circuits that contain two spokes according to their number of outer edges. As there are always two possible choices for the path through the inner edges, we only consider the shorter one in Table 1. Similarly, we may restrict attention to circuits that use at most $5 = \lfloor \frac{11}{2} \rfloor$ outer edges.

| | | | | | |
|-------------------------------|----|----|----|----|----|
| Number of outer edges | 1 | 2 | 3 | 4 | 5 |
| Number of inner edges | 3 | 5 | 2 | 1 | 4 |
| Weight of the shorter circuit | 43 | 57 | 46 | 45 | 64 |

Table 1: Weights of the circuits in $(P_{11,4}, w)$ that use two spokes.

Let B consist of the outer circuit and of the 11 circuits that use precisely one outer edge. We claim that B is an undirected cycle basis. Assume otherwise. Then, there exists a nontrivial linear combination yielding the zero vector, over $\text{GF}(2)$. If such a combination made use of any of the 11 circuits that use precisely one outer edge, then it has to use each of these circuits in order to cancel out the spokes. The sum of these 11 circuits is the outer circuit plus the inner circuit. Thus there is no nontrivial linear combination yielding the zero vector. \square

It remains to show that B is *not* an integral cycle basis. Indeed, its determinant is 3 as a simple calculation shows.

3.5 Directed and $\text{GF}(p)$ -Bases

We show that the computation of minimum directed cycle bases can be reduced to the computation of minimum $\text{GF}(p)$ -bases for suitable primes p .

Lemma 3.6. *Let B be a minimum weight directed cycle basis and let p be a prime. The weight of a minimum weight $\text{GF}(p)$ -basis is no smaller than the weight of B . If p does not divide the determinant of B , B is also a minimum weight $\text{GF}(p)$ -basis.*

Proof. Linear dependence over \mathbb{Q} implies linear dependence over $GF(p)$ for any p . Therefore, any $GF(p)$ -basis is a directed basis. If the determinant of B is not divided by p , $\det B \bmod p \neq 0$ and B is a $GF(p)$ -basis. \square

In order to apply the preceding lemma, we need a bound on the determinant of a directed cycle basis. Consider any directed cycle basis B . Its determinant is the determinant of a $\nu \times \nu$ matrix with entries in $\{-1, 0, +1\}$. Moreover, each column of this matrix contains at most n nonzero entries.

Lemma 3.7. *The determinant of a directed cycle basis is an integer bounded by $n^{m/2}$.*

Proof. The determinant is a sum of $\nu!$ terms; each term has absolute value at most one. This gives a bound of $\nu! \leq \nu^\nu$. Hadamard's inequality yields a slightly better bound. The absolute value of the determinant is bounded by the product of the ℓ_2 -norms of the column vectors. The norm of each column vector is at most \sqrt{n} and hence we have the bound \sqrt{n}^ν . \square

Combining the two preceding lemmas, we obtain a characterization of minimum directed basis in terms of minimum $GF(p)$ -bases.

Theorem 3.8. *Let P be a set of m primes each of value at least n . For each $p \in P$, let B_p be a minimum $GF(p)$ -basis, and let p_0 be such that B_{p_0} has minimum weight among the bases B_p .*

1. B_{p_0} is a minimum weight directed basis.
2. Let $p \in P$ be chosen uniformly at random. Then B_p is a minimum weight directed basis with probability at least $1/2$.

The primes in P can be chosen in $O(m \log m)$.

Proof. Let B be any minimum directed basis. No more than $m/2$ primes in P can divide the determinant of B .

For an integer s , let $\pi(s)$ be the number of primes less than or equal to s . Then $s/(6 \log s) \leq \pi(s) \leq 8s/\log s$ (Apostol (1997)). Then there are most $8n/\log n$ primes less than n . If t is such that $t/(6 \log t) \geq 8n/\log n + m$ then there are at least m primes of value at least n less than t ; $t = O(m \log m)$ suffices. \square

If $p = O(m \log m)$ and hence $\log p = O(\log m)$, arithmetic in $GF(p)$ takes time $O(1)$.

3.6 Circuits versus Cycles.

We defined cycle bases as sets of circuits. Alternatively, we could have defined them as sets of cycles. Is there always a minimum weight basis that consists only of circuits? Is the minimum weight basis of a disconnected graph the union of minimum weight bases of the components? For some of our classes, the answers are yes. For some, the answers to these and related questions are not known.

Theorem 3.9. *Exchange Theorem (Horton (1987)) If B is a D or U-basis of G , $C \in B$ and $C = C_1 + C_2$, then either $B \setminus \{C\} \cup \{C_1\}$ or $B \setminus \{C\} \cup \{C_2\}$ is also a cycle basis of G .*

Proof. Let Γ be the cycle matrix for B and let Γ_i be the cycle matrix for $B - C + C_i$, $i = 1, 2$. Let T be a spanning forest of G and let A and A_i be the respective square submatrices indexed by the arcs not in T . Then, using the linearity of the determinant function for the column which corresponds to C , we find that $0 \neq \det A = \det A_1 + \det A_2$. \square

The family of linearly independent cycles forms a matroid.

Theorem 3.10. *The set of (directed) cycles of a graph G forms a matroid. The bases of the matroid clearly coincide with the (directed) cycle bases of G .*

Proof. Let \mathcal{I} denote the system of all linear independent sets of cycles in G . It suffices to show the following.

- $\emptyset \in \mathcal{I}$.
- $A \in \mathcal{I}$ and $B \subset A$ implies $B \in \mathcal{I}$.
- For all sets $A, B \in \mathcal{I}$ with $|A| > |B|$ there exists an element $a \in A \setminus B$ such that $B \cup \{a\} \in \mathcal{I}$.

The listed properties hold since the (directed) cycle space of G forms a vector space. \square

We will show that Theorem 3.10 does not hold for integral cycle bases, i.e. the system of all subsets of all integral bases in G does not form a matroid. This will cause the computational approaches suitable for U-bases and D-bases to fail for I-bases. In section 5.7 will we discuss these issues. At this point we will examine the validity of Theorem 3.9 for K-bases with K neither D nor U.

Theorem 3.9 does not hold for totally unimodular bases.

Lemma 3.11 (T. Jurkiewicz, personal communication). *There is a graph G and a totally unimodular bases B of G containing a circuit C and a decomposition $C = C_1 + C_2$ of C such that neither $B \setminus \{C\} \cup \{C_1\}$ nor $B \setminus \{C\} \cup \{C_2\}$ is a totally unimodular bases.*

Proof. Figure 10 shows a graph and a TUM-basis of this graph. We invite the reader to verify that this basis is TUM. Figure 11 shows a decomposition of the first circuit into two circuits. Replacing the first circuit by either one of the two circuits shown in Figure 11 results in a basis that is not TUM. In both cases, the cycle matrix of the resulting basis contains a 2 by 2 submatrix of the form

$$\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

This matrix has determinant -2 ; in a TUM-basis, the determinants of all square submatrices must be in $\{-1, 0, +1\}$. \square

For weakly fundamental bases, we can show Theorem 3.9 under the additional assumption that C_1 and C_2 use only edges that are also used by C .

Lemma 3.12. *Let B be W-basis of G , let $C \in B$ and $C = C_1 + C_2$ where $|C_i(e)| \leq |C(e)|$ for all e and $C_i \neq C$ for $i = 1, 2$. Then at least one of $B - C + C_1$ or $B - C + C_2$ is a W-basis of G .*

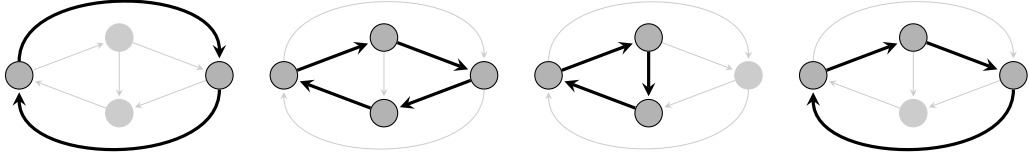


Figure 10: A graph and a totally unimodular cycle basis. For each of the four circuits, the edges belonging to the circuit are shown in bold.

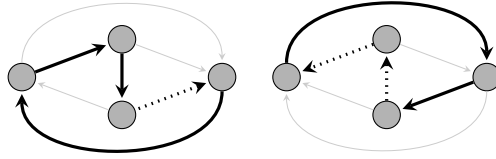


Figure 11: Decomposition for first element of bases from Figure 10. Edges that are used in reversed direction are shown dotted.

Proof. If B is a weakly fundamental basis, there is an ordering of the cycles in B such that every cycle introduces a nontree edge not used in any preceding cycle. Let e be the edge introduced by C . Then $C(e) \neq 0$ and hence at least one of $C_1(e)$ or $C_2(e)$ is nonzero, say the former. We replace C by C_1 . Since for any nonzero coefficient of C_1 , the corresponding coefficient of C is nonzero, the nontree part of the new cycle matrix is still lower triangular. \square

Observe that Lemma 3.12 is *not* true for strictly fundamental cycle bases. To see this, consider the sunflower graph $SF(3)$, see Fig. 7, and some minimum strictly fundamental cycle basis B of $SF(3)$. If we decompose the 4-circuit $C \in B$ into $C = C_1 + C_2$, where both C_1 and C_2 are triangles, then neither of the $B - C + C_i$ is a strictly fundamental cycle basis. In the next lemma we show that Lemma 3.12 does not hold for integral bases either.

Lemma 3.13. *There is a graph G and an integral basis B of G containing a non-circuit C such that for any decomposition $C = C_1 + C_2$ of C with $|C_i(e)| \leq |C(e)|$ for all e and $C_i \neq C$ for $i = 1, 2$, neither $B - C + C_1$ nor $B - C + C_2$ is an integral basis.*

Proof. The graph is $P_{7,3}$ as shown in Figure 12. It consists of two disjoint cycles of length 7, called the outer and the inner cycle, respectively. We use O_i and I_i , $0 \leq i < 7$, to denote the nodes on the outer and inner cycle, respectively. The outer and inner cycles have edges (I_i, I_{i+1}) and (O_i, O_{i+4}) , $0 \leq i < 7$, respectively. All indices are modulo 7. Furthermore, we have the edges (O_i, I_i) , $0 \leq i < 7$, called spokes. To summarize, $n = 14$, $m = 21$, and the cyclomatic number ν is thus eight.

The basis B consists of the following cycles. For $0 \leq i < 7$, we have the cycle C_i consisting of the edges (O_i, O_{i+1}) , (O_{i+1}, O_{i+2}) , (O_{i+2}, I_{i+2}) , (I_{i+2}, I_{i+6}) , (I_{i+6}, I_{i+10}) , (I_{i+10}, I_{i+14}) , and (I_{i+14}, O_i) . Observe that the sum of the C_i 's is the nonsimple cycle consisting of two copies of the outer circuit and three copies of the inner circuit. We also have the cycle $D_{a,b}$ that consists of a copies of the outer circuit and b copies of the inner circuit, where $a, b \in \mathbb{Z}$. We will fix a and b later.

We next determine the determinant of the set of cycles above as a function of a and b . We fix a spanning tree T consisting of the spokes and all inner edges except for edge (I_2, I_6) .

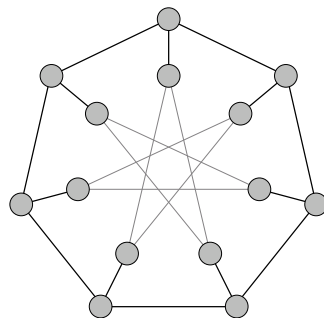


Figure 12: The generalized Petersen graph $P_{7,3}$. We provide an integral cycle basis of $P_{7,3}$ which features a nonsimple cycle, but cannot be decomposed into a basis which only consists of circuits.

We obtain the following square part of the cycle matrix:

| | C_0 | C_1 | C_2 | C_3 | C_4 | C_5 | C_6 | $D_{a,b}$ |
|--------------|-------|-------|-------|-------|-------|-------|-------|-----------|
| (O_0, O_1) | 1 | 1 | | | | | | a |
| (O_1, O_2) | | 1 | 1 | | | | | a |
| (O_2, O_3) | | | 1 | 1 | | | | a |
| (O_3, O_4) | | | | 1 | 1 | | | a |
| (O_4, O_5) | | | | | 1 | 1 | | a |
| (O_5, O_6) | | | | | | 1 | 1 | a |
| (O_6, O_1) | 1 | | | | | | 1 | a |
| (I_2, I_6) | 1 | | | 1 | 1 | | | b |

Observe that the edge (I_j, I_{j+4}) is used by the cycles C_{j-2} , C_{j-6} and C_{j-10} . The determinant of the matrix above is $2b - 3a$ as a little calculation, e.g., Gaussian elimination, shows. For $a = b = 1$, the determinant is -1 and hence the basis is integral. The cycle $D_{1,1}$ is not a circuit and uses the outer O and the inner circuit I in the forward direction. The only decomposition of D is $O + I$. The determinant of the basis $B - D + O$ is -3 (use $a = 1$ and $b = 0$ in the formula for the determinant) and the determinant of the basis $B - D + I$ is 2 (use $a = 0$ and $b = 1$ in the formula for the determinant). Thus neither basis is integral. \square

The next two lemmas provide us with properties of minimum cycle bases which are extremely valuable in practice. These properties are an immediate consequence of Lemma 3.12, and turn out to be true for strictly fundamental cycle bases, too.

Theorem 3.14. *For $K \in \{D, U, W, F\}$, any graph G has a minimum K -basis consisting only of circuits.*

Proof. The cycles in fundamental bases are circuits by definition. For any of the other K 's, consider a basis B containing a cycle C that is not a circuit. We may decompose C into a sum of circuits C_i , $1 \leq i \leq k$. By the preceding lemmas, some $B - C + C_i$ is a K -basis of G . Also $w(C_i) \leq w(C)$. \square

Theorem 3.15. *For $K \in \{D, U, W, F\}$, a minimum weight K -basis of a graph can be obtained as the union of minimum weight K -bases of its maximal 2-connected components.*

Proof. By the preceding theorem, there is a minimum weight K-basis consisting only of circuits. A circuit uses edges only from one 2-connected component. \square

Open Problem 2. *Does Theorem 3.14 or Theorem 3.15 hold for integral basis or totally unimodular bases? Does Lemma 3.12 hold for totally unimodular bases?*

3.7 Reductions.

We study some simplification rules. At first sight, all might appear quite natural. However, for certain classes of cycle bases, we do not know whether these rules are valid.

For example, is there a simple way to deal with parallel edges? Is there a simple way of handling edges of weight zero?

Let $g = (u, v)$ be a zero weight edge without parallel edges. Let G' be obtained from G by removing g and identifying u and v , i.e., replacing endpoints u and v in edges in $E(G)$ by a new vertex uv . The edges of G' correspond to the edges in $E - g$. Let B' be a basis of G' , where $K \in \{D, U, I, W, TUM, F\}$. Consider the following set B of cycles in G : for any $C' \in B'$ we add a cycle C to B that is obtained from C' by adding g with appropriate multiplicity; the appropriate multiplicity guarantees flow conservation at u and v .

Lemma 3.16. *Let G' be obtained from G by contracting an edge of cost zero not having any parallel edges, let B' is a minimum weight K-basis of G' , and let B be obtained from B' as described above. Then B is a minimum weight K-basis of G for $K \in \{D, U, I, W, TUM, F\}$.*

Proof. Let T' be a spanning forest of G' and let Γ' be the cycle matrix corresponding to B' . Let A' be the square submatrix selected by the rows not in T' . Then $T := T' + g$ is a spanning tree of G . We obtain the cycle matrix for B by adding a row for g and setting the entries in this row appropriately. Observe that A' stays the square matrix selected by the nontree edges. Thus B is a K-basis of G . The weight of B is the weight of B' .

Conversely, let B be any K-basis of G and let B' be obtained from B by identifying u and v . The matrix Γ' for B' is obtained from the matrix Γ for B by deleting the row corresponding to g .

Let C' be any cycle in G' . We lift C' to a cycle C in G . The representation of C with respect to B translates into a representation of C' with respect to B' . Thus B' is also of type K. Also $w(B') \leq w(B)$. \square

Lemma 3.17. *Let $K \in \{D, U\}$ and let e be any edge. For any minimum weight circuit F containing e , there is a minimum weight K-basis containing F . Any minimum weight K-basis contains a minimum weight circuit containing e .*

Proof. Let B be a minimum weight K-basis. Then $F = \sum_{C \in B} \lambda_C C$. Clearly, there must be a $C \in B$ such that $e \in C$ and $\lambda_C \neq 0$. Then $w(F) \leq w(C)$ and $B' := B \setminus C + F$ is a K-basis of weight no larger than the weight of B . Hence B' is a minimum weight K-basis. If $w(F) < w(C)$, B was not a minimum weight K-basis. \square

Lemma 3.17 is not true for strictly fundamental cycle bases, already for the sunflower graph SF(3). In particular, for the edge e which induces the only 4-circuit in a minimum such basis B , the unique shortest circuit through e is not contained in B .

Lemma 3.18. *Let g and f be parallel edges with $w(g) \leq w(f)$. For $K \in \{D, U, W\}$ a minimum weight K -basis of G can be obtained from a minimum weight K -basis B' of $G' := G - f$ by adding a cheapest circuit through f ; call it C .*

Proof. C is clearly independent from B' . Also, C introduces an edge that is not used in any of the other cycles. Thus, if B' is a K -basis of G' , $B := B' \cup C$ is a K -basis of G with $w(B) = w(B') + w(C)$. Assume, for the sake of a contradiction, that G has a K -basis \hat{B} with $w(\hat{B}) < w(B)$. We will show that this implies that G' has a K -basis of weight less than $w(B')$.

Assume first that $K \in \{D, U\}$. The cheapest circuit containing f is either the circuit⁴ $g \circ f^{-1}$ or has the form $f \circ P$ where P is a cheapest path connecting the endpoints of f in $G \setminus \{f, g\}$. In the latter case, $g \circ P$ is a cheapest cycle containing g . By Lemma 3.17 we may assume that B contains the circuit $g \circ f^{-1}$ in the first case or the circuits $g \circ P$ and $f \circ P$ in the latter case. Assume now, that \hat{B} contains another circuit, say $f \circ Q$, containing f . Replacing this circuit by $g \circ Q$ yields a basis of weight no larger than \hat{B} as $g \circ Q = f \circ Q + g \circ f^{-1}$ in the former case and $g \circ Q = f \circ Q + g \circ P - f \circ P$ in the latter case. We conclude that G has a basis of weight no larger than \hat{B} in which $g \circ f$, respectively $f \circ P$ is the only circuit containing f . Deleting this circuit from the basis gives us a basis of $G - f$.

For $K = W$, we have to argue differently. Let $\hat{\Gamma}$ be the cycle matrix for \hat{B} . We may assume that $\hat{\Gamma}$ has an upper triangular matrix in its last ν rows, and the arcs of some spanning forest T placed above.

Assume first that the row for g is above the row for f . Then f must be a co-tree arc, because otherwise g and f form a circuit in T . Hence there is a circuit $C_f \in \hat{B}$ “introducing” f , i.e., the diagonal entry in the row indexed by f belongs to C_f . We delete C_f from the basis and replace, in the other basic circuits, occurrences of f by g . Removing the row of f , too, we obtain the cycle matrix of a W -basis for G' of weight $w(\hat{B}) - w(C_f) < w(B) - w(C) = w(B')$, a contradiction.

Assume next that the row for f is above the row for g . Then g must be a co-tree arc and hence there is a circuit C_g introducing g ; f may be a tree arc or a co-tree arc. If f is a tree arc, we make g a tree arc, replace f by g in all circuits and delete f and C_g . If f is a co-tree arc, the circuit C_f introducing f does not use the arc g , because g was assumed to be arranged below f . We replace f by g in all circuits and delete f and C_g . The circuit C'_f obtained from C_f by replacing f by g now introduces g . In either case, we obtain the cycle matrix of a W -basis for G' of weight $w(\hat{B}) - w(C_g) < w(B) - w(C) = w(B')$, a contradiction. \square

Open Problem 3. *Extend statements 3.12 to 3.18 to types of cycle bases not covered by the statements.*

Open Problem 4. *Let $e = \{u, v\}$ be a non-metric edge of a biconnected graph G , i.e., $\text{dist}_G(u, v) < w(e)$. Each minimum K -basis B has precisely one circuit $C \in B$ with $e \in C$. This is true for $K \in \{D, U\}$. Is it true for any other type?*

3.8 Selected Properties.

We consider the sequence of weights of circuits in a minimal K -basis in non-decreasing order. Let $K \in \{D, U\}$ and B and B' be distinct K -bases of G both of minimal weight. Then their ordered sequences of weights coincide. This is not true for integral bases.

⁴We assume that g and f are oriented in same way; f^{-1} is the reversal of f and runs anti-parallel to g .

Lemma 3.19. For $K \in \{D, U\}$, let w and w' be the non-decreasing sequences of weights of circuits of two minimal K -bases B and B' respectively. Then $w = w'$.

Proof. This is true since both the undirected and the directed cycle space form a vector space over $GF(2)$ and \mathbb{Q} respectively. Hence the cycles together with linear independence form a *matroid* (cf. Theorem 3.10). Finally, it is a well known fact that the non-decreasing weight sequences of minimal bases in matroids coincide. \square

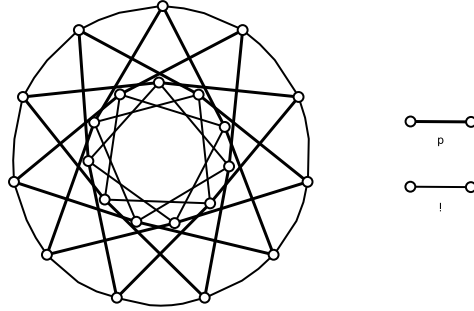


Figure 13: A graph G featuring minimal integral cycle bases with different weight sequences.

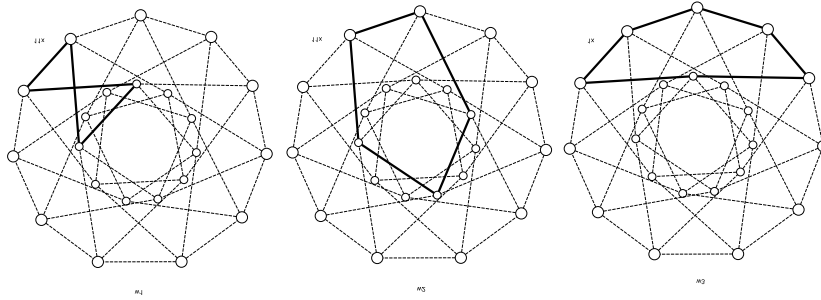


Figure 14: A minimal I-basis B of G with weight sequence $w = (38, \dots, 38, 42, \dots, 42, 42, 46)$.

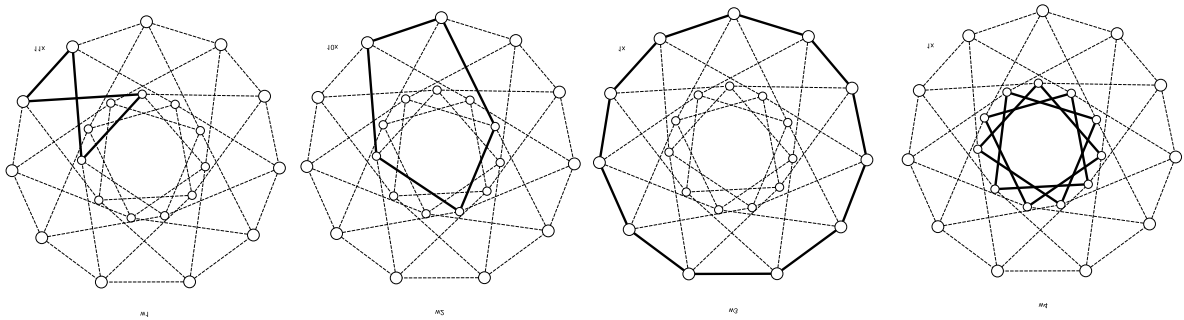


Figure 15: A minimal I-basis B' of G with weight sequence $w' = (38, \dots, 38, 42, \dots, 42, 44, 44)$.

Lemma 3.20. There is a graph G and two minimal integral bases B and B' of G whose non-decreasing weight sequences do not coincide.

Proof. Consider the graph G depicted in Figure 13. It arises from the generalized Petersen graph $P_{11,3}$ by the addition of an extra set of 11 spokes. We give weight four to all inner and outer edges and weight 15 to all spokes. Let B and B' be the two sets of circuits in G shown in Figures 14 and 15, respectively. Either set forms an integral cycle basis of G as the reader might check. To see minimality of B and B' note that the first 22 circuits in B are in fact the only ones in the graph whose weight does not exceed 42. Beside them there are only two more circuits, the inner one and the outer one, whose weight is at most 44. Replacing the last circuit in B by either of them yields a cycle basis that is not integral, i.e. its determinant equals 2 or 3, respectively. Every circuit in G other than the so far considered ones has weight at least 46. Hence there are exactly two sets of 23 circuits whose weight is less than 926. Both form a cycle basis of G , but neither is integral. \square

| Graph class | minimum W-basis | minimum F-basis |
|-------------------|--|--|
| Weighted | | |
| General | $O(W \log n)$, Thm 4.4 | |
| General | $O(n \cdot W(\text{MST}) + W)$, Thm 4.2 | $O(W \log^2 n \log \log n)$, Thm 4.11 |
| Planar | $\Theta(W)$ | |
| Unweighted | | |
| General | $O(m \log n \log(m/n))$, Thms 4.1 and 4.5 | $O(n^2)$, Thm 4.12 |
| Planar | | $O(n\sqrt{n\phi})$, Thm 4.7 |
| Outerplanar | | $\Theta(n)$ Reich (2007) |
| d -dim grids | $\Theta(n)$ | $\Theta(n \log n)$, Thm 4.8 |

Table 2: Bounds for minimum weight W- and F-bases. W denotes the total edge weight and $W(\text{MST})$ is the weight of a minimum spanning tree. Bounds for unweighted graphs are only stated if they are better than the bound derived for weighted graphs with $W = m$. In the bound for planar graphs, ϕ is the maximal size of any face.

4 Length and Weight of Cycle Bases

In this section we discuss a-priori bounds on the length and weight of minimum cycle bases. We state the bounds as functions of the number n of vertices, the number m of arcs, and the total weight W of the edges. Many applications benefit from small length or small weight bases as we will see in Section 7; algorithms for computing minimum or nearly minimum weight bases will be discussed in Section 5. Table 2 summarizes the results. It is interesting to note that all upper bounds have been shown for either weakly or strongly fundamental bases. Although we know that general bases are not always fundamental (see Example 3.3), it seems that fundamentality gives sufficient structure to the problem to make an analysis of their length achievable or at least easier than for general bases.

Open Problem 5. *Derive a-priori bounds on the weight (length) of directed, undirected, integral, and totally unimodular bases.*

Open Problem 6. *For $K, K' \in \{\text{D}, \text{U}, \text{I}, \text{TUM}, \text{W}, \text{F}\}$ and a graph G with weight function W , let*

$$r_{K,K'}(G) = \frac{\text{weight of a minimum } K\text{-basis}}{\text{weight of a minimum } K'\text{-basis}}$$

and

$$r_{K,K'}(n, m) = \max\{r_{K,K'}(G) \mid G \text{ is a graph with } n \text{ nodes and } m \text{ edges}\} .$$

Derive upper bounds on $r_{K,K'}(n, m)$. For example, $r_{\text{W},\text{D}}(n, m) = O(\log n)$ since every graph with n nodes has a W-basis of weight $O(W \log n)$ and since every D-basis has weight at least W .

The bounds given in Table 2 are obtained by different methods. There are essentially four approaches:

1. Use of special graph properties like planarity.
2. Induction.
3. Use of clusters, partitions, and spanners.
4. Results of extremal graph theory.

We start with some obvious bounds. Throughout this section, we restrict attention to biconnected graphs. There are $m - n + 1$ circuits in a basis and each circuit has length at most n . Thus any basis has length at most mn and weight at most mW . Throughout this section, $W = \sum_{e \in E} w(e)$ denotes the total weight of the edges. Obvious lower bounds are $\Omega(m)$ and $\Omega(W)$, since in biconnected graphs every edge has to belong to at least one circuit of any basis. Extremal graph theory provides a nontrivial lower bound.

Theorem 4.1. *For any integer $h = 2 \pmod{4}$ and $h \geq 6$, there is a graph $G_h(n)$ with n nodes and $m = hn/2$ edges, such that any cycle basis for $G_h(n)$ has length $\Omega(m \log n / \log(m/n))$. In particular, there is a graph family where $m = \Theta(n)$ and any basis has length $\Omega(m \log n)$, and for any integer k , there is graph family where $m = \Theta(n^{1+1/k})$ and any basis has length $\Omega(mk)$.*

Proof. For any integer h with $h = 2 \pmod{4}$ and $h \geq 6$, there exists an infinite family of h -regular graphs, i.e., $m = hn/2$, in which every cycle has length $\Omega(\log n / \log(m/n))$, see Lubotzky et al. (1988). Since a basis consists of $m - n + 1$ circuits, any basis has length $\Omega(n \log n / \log(m/n))$. For $h = 6$, we obtain graphs with $m = \Theta(n)$, for which every basis has length $\Omega(m \log n)$. For $h = n^{1/k}$, we obtain graphs with $m = \Theta(n^{1+1/k})$, for which every basis has length $\Omega(mk)$. \square

Open Problem 7. *Prove a nontrivial lower bound for weighted graphs.*

4.1 Weakly Fundamental Bases

The first result for general graphs was given by Horton in 1987; Liebchen observed that the construction yields a weakly fundamental basis and not only an undirected basis. We generalize Horton's proof to yield an upper bound for weighted graphs.

Theorem 4.2 (Horton (1987); Liebchen (2003)). *Every simple graph G has a W-basis of length at most $3(n-1)(n-2)/2$ and weight at most $2nW(\text{MST}) + 2W$, where $W(\text{MST})$ is the weight of a minimum spanning tree.*

Proof. We prove only the weight upper bound. For the case of uniform weights, we have $W(\text{MST}) = n - 1$ and $W = m$. This gives a bound of $2n^2 + 2n^2/2 = 3n^2$ for the uniform case.

Let T be a MST of G . The claim clearly holds for $n \leq 3$. So assume that G has more than three vertices and let v be a leaf of T . Our W-basis for G consists of two parts: first, a W-basis $B(G-v)$ of $G-v$ constructed recursively, and second, $d(v) - 1$ cycles passing through v . Observe, that a basis for $G-v$ has cardinality $m - d(v) - (n-2) = m - (n-1) - (d(v) - 1)$ and hence we are adding the right number of cycles. The graph consisting of T plus the $d(v) - 1$ non-tree edges incident to v is planar. We form $d(v) - 1$ circuits by taking all but one face cycle of this planar graph. The resulting set of circuits is weakly fundamental as an argument similar to the one used in the proof of Theorem 3.5 shows.

It remains to argue the weight bound. By the induction hypothesis, $B(G - v)$ has weight at most $2(n - 1) \cdot W(\text{MST} - v) + 2W(G - v)$. The circuits added in the induction step have weight at most $2W(\text{MST}) + 2W(v)$ where $W(v)$ denotes the sum of the weights of the edges incident to v . Thus the weight of our basis is at most $2nW(\text{MST}) + 2W$. \square

The upper bound is tight. Consider the complete graph on n nodes. It has $m = n(n - 1)/2$ edges. Since any circuit in any basis contains at least three edges, any cycle basis has length at least $3(m - n + 1) = 3(n - 1)(n - 2)/2$. For sparse graphs, a much better bound is possible. Rizzi (2007) proved that every graph has a weakly fundamental basis of length $O(m \log n)$. In fact, he even showed that every weighted graph has a W -basis of weight $O(W \log n)$. The proof given here is due to T. Kavitha and R. Rizzi. The proof makes use of the fact that every graph of minimum degree three contains a logarithmic length cycle.

Lemma 4.3 (Bollobás (1978)). *Any graph with minimal degree at least three contains a cycle of length at most $2 \lceil \log_2 n \rceil$. Moreover such a cycle can be found in time $O(n)$.*

Proof. Let G be our graph and let v be an arbitrary vertex. Grow a breadth-first search tree rooted at v . As long as only tree edges are encountered, every vertex has at least two children. Thus if $2^0 + 2^1 + \dots + 2^k > n$, there must at least one non-tree edge incident to a vertex of depth $k - 1$ and hence a cycle of length $2k$ exists. This proves the bound on the length of a shortest cycle. With respect to the time bound, we observe that the first non-tree edge encountered yields the desired cycle. \square

Theorem 4.4 (Rizzi (2007)). *Any weighted graph G with total weight W has a weakly fundamental basis of weight $O(W \log n)$. Such a basis can be determined in time $O(nm)$.*

Proof. We build the basis and a spanning tree concurrently. Initially, the basis and the spanning tree are empty. Let G be our current graph that is initially set to the input graph. If G is empty, we stop. If G has a vertex of degree zero, we delete the vertex, if G has a vertex of degree one, we delete the vertex and add the incident edge to the spanning tree. So assume that every vertex has degree two or more. We call a maximal path whose interior vertices have degree two a super-edge; an edge whose endpoints have both degree three or more is also a super-edge. The weight of a super-edge is the sum of the weights of the edges forming the super-edge. The endpoints of super-edges have degree three or more in G , see Figure 16. The graph consisting of the vertices of degree three or more and the super-edges joining them contains a circuit \overline{C} consisting of $O(\log n)$ super-edges. Let p be the heaviest super-edge in \overline{C} and let C be the cycle in G represented by \overline{C} . Then $w(C) = O(w(p) \log n)$. We add C to our basis. We also delete all edges belonging to p from G , designate an arbitrary edge of p as a non-tree edge, and add all other edges of p to T . If p consists of k edges, m decreases by k and n decreases by $k - 1$. So ν decreases by 1 as it should.

The basis constructed in this way is weakly fundamental because the edge of p designated as a non-tree edge is not used in any cycle constructed later. Also, its weight is $O(W \log n)$ as the cost of the cycle added in an iteration is at most $O(\log n)$ times the weight of the edges deleted in this iteration. \square

For graphs with $m = O(n)$ edges, the result is tight as Theorem 4.1 establishes the corresponding lower bound. Kaufmann and Michail (2008) have recently shown that the lower bound can also be matched for larger values of m . The improvement uses the fact that graphs with at least $n^{1+1/k}$ edges contain a cycle of length at most $2k$, see Bollobás (1978).

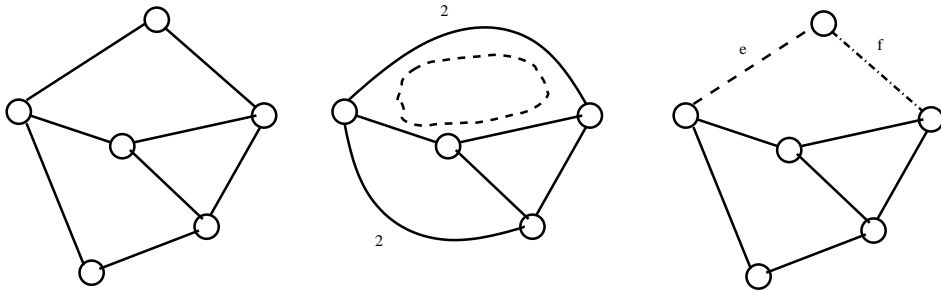


Figure 16: In the graph on the left, all edge weights are equal to one. In the graph in the middle, the two indicated super-edges have weight two. The dashed line indicates a short cycle. It consists of three super-edges and has weight four. The heaviest super-edges has weight two. We delete its edges from the graph and make e a non-tree edge and f a tree edge (or vice-versa).

We now proceed as follows. As long as $m \geq n^{1+1/k}$ for a constant k still to be determined, we find cycles of length at most $O(2k)$. We delete one of its edges and charge the cost of the cycle to it. As soon as $m \leq n^{1+1/k}$, we switch to the construction in Theorem 4.4. We construct cycles consisting of $O(\log n)$ super-edges, delete the edges in the heaviest super-edge and charge $O(\log n)$ to each edge removed. The total charge is

$$O(mk + n^{1+1/k} \log n) \stackrel{k=2(\log n)/\log(m/n)}{=} O\left(m \frac{\log n}{\log(m/n)}\right).$$

Theorem 4.5 (Kaufmann and Michail (2008)). *Every graph has a weakly fundamental basis of length $O(m \log n / \log(m/n))$. For $m = \Theta(n^{1+1/k})$, the bound is $O(mk)$ and for $m = n \log^c n$ and $c > 0$ constant, the bound is $O(m \log n / \log \log n)$. Finally, for $m = cn$, the bound gives $O(m \log n)$.*

We close our discussion of weakly fundamental bases with some remarks on planar graphs. Every planar graph has a 2-basis and 2-basis are weakly fundamental by Theorem 3.5. Thus every planar graph has a W-basis of length $O(n)$ and weight $O(W)$.

4.2 Fundamental Bases

Upper bounds for strictly fundamental bases are obtained by constructing spanning trees of small diameter or, more generally, spanning trees of small stretch. Clearly, a spanning tree T of diameter D or with $\sum_{e=(u,v) \in E} d_T(x,y)/m \leq D$ gives rise to an F-basis of length $O(Dm)$. Here, $d_T(x,y)$ is the length of the path in T connecting x and y . We review results for planar graphs and for general graphs. The constructions make use of graph separators and graph partitions with nice properties.

Definition 4.1. *A set $S \subset V$ is an (α, β) -separator if $|S| \leq \beta\sqrt{n}$ and any connected component of $G - S$ contains no more than αn vertices.*

Lemma 4.6 (Miller (1986)). *Any biconnected planar graph with n vertices, m edges and maximal face size ϕ has an (α, β) -separator with $\alpha = 2/3$ and $\beta = 2\sqrt{\phi/2}$. Moreover, the separator constitutes a simple cycle and is thus called a simple cycle separator (SCS).*

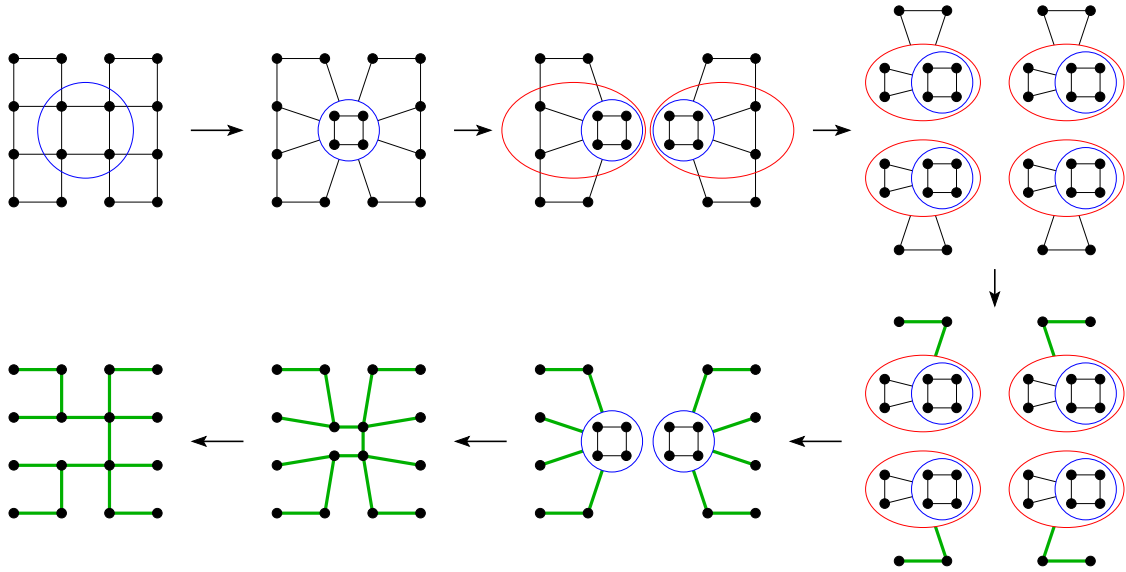


Figure 17: Stern's and Vavasis' construction of an F-basis for planar graphs.

Cycle separators are the basis for the following theorem.

Theorem 4.7 (Stern and Vavasis (1990)). *Any planar graph G with maximal face size ϕ has an F-basis of length $O(n\sqrt{\phi n})$.*

Proof. We may assume that G is biconnected. Figure 17 illustrates the construction. Let S be a simple cycle separator of size $\beta\sqrt{n}$ in G . We contract S into a single vertex v . Clearly, v becomes an articulation point in the resulting graph. Let G_1, G_2, \dots, G_k denote the components that would result if v were deleted. We make k copies v_1 to v_k of v , one for each component, and connect v_i with v 's neighbors in G_i . Each G_i has at most αn vertices and the maximal face size is no more than ϕ . A spanning tree of G is obtained by taking the cycle S minus one edge plus spanning trees of the components. The spanning trees of the components are constructed recursively. We stop when the components have constant size and use any spanning tree for them.

Let $D(n)$ be the diameter of the spanning tree constructed in this way. Then $D(n) \leq O(\sqrt{\phi n}) + D(\alpha n) = O(\sqrt{\phi n})$. \square

Outerplanar graphs have strictly fundamental bases of linear size (Reich (2007)).

Theorem 4.8. *For grid graphs of fixed dimension the minimal length of a fundamental basis is $\Theta(n \log n)$.*

The upper bound for two-dimensional grids was first shown by Stern and Vavasis (1990). A simplified construction that, in addition, applies to any fixed dimension was found by Alon et al. (1995). It is illustrated in Fig. 18 and yields a basis of length no more than $(4/3)n \log n$ as shown by Köhler et al. (2009). The lower bound was also established by Alon et al. (1995); Köhler et al. (2009) paid attention to the constant factor and proved, with a different method, that any strictly fundamental basis for the planar grid has length at least $(1/12)n \log_2 n - O(n)$.

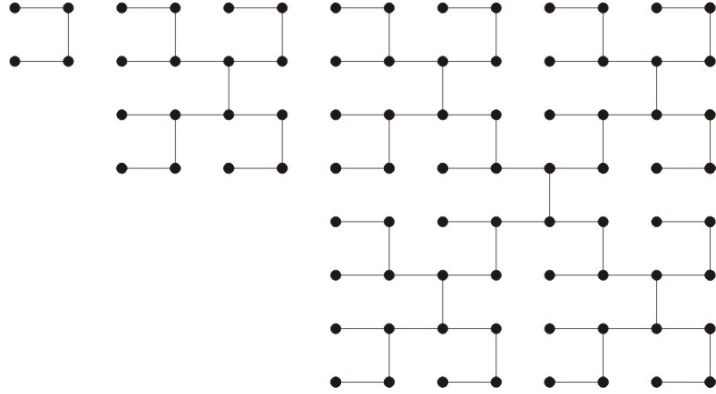


Figure 18: A spanning tree for d -dimensional grid graphs with length 2^i in all dimensions Alon et al. (1995). The construction is shown for $d = 2$. If $i = 1$, an optimal spanning tree for the structure is returned. If $i > 1$, the graph is partitioned into 2^d cubes of length 2^{i-1} and trees for the subgraphs are constructed recursively. The set of 2^d vertices in the center of the graph is connected such that they form the same tree that is used at the base of the recursion.

The first upper bound on the length of *strictly fundamental cycle bases* in general graphs was given by Alon et al. (1995). We follow the very descriptive explanation of their technique by Peleg (2000). The construction relies on partitioning a given graph into *clusters* such that the diameter of the clusters and the number of edges between clusters (*intercluster edges*) are controlled at the same time.

Lemma 4.9 ((Peleg, 2000, p.153)). *Given an unweighted graph $G = (V, E)$, $|V| = n$, and a parameter $x > 1$, there is a partition P of G into clusters C_i such that*

1. *the radius of each cluster is at most $x \ln m$, and*
2. *the number of intercluster edges is at most m/x .*

Proof. The clusters are grown one by one. As long as there is a vertex not assigned to any cluster, choose one such node and grow a cluster C around it in discrete steps. Initially, C consists only of the vertex. Let $E_{out}(C)$ be the set of edges with exactly one endpoint in C , let N_{out} be the endpoints outside C of the edges in $E_{out}(C)$, and let $E_{in}(C)$ be the edges with both endpoints in C . We add N_{out} to C if $|E_{out}(C)|/|E_{in}(C)| \geq 1/x$. If $|E_{out}(C)|/|E_{in}(C)| < 1/x$, the growth of C is stopped, C is added to the partition and deleted from G , and the next cluster is grown.

Clearly, any edge of G is contained in at most one cluster. Thus the number of intercluster edges is at most m/x . Consider the growth of any particular cluster C . We start with a single node v and no edge. In the first iteration all neighbors of v (that are not assigned to any previous cluster) are added to the cluster. Let m_i be the number of edges added in the i -th iteration. Then $m_i \geq (m_1 + \dots + m_{i-1})/x$. For the analysis of the growth of the m_i 's assume equality. Then $m_i - m_{i-1} = m_{i-1}/x$ and hence $m_i = (1 + x)m_{i-1}/x$. We conclude

$$m_1 + \dots + m_i \geq m_i = \Omega \left(\left(\frac{1+x}{x} \right)^i \right).$$

Thus $i \leq (\ln m)/\ln(1 + 1/x) \leq x \ln m$ and we have also established the first property. \square

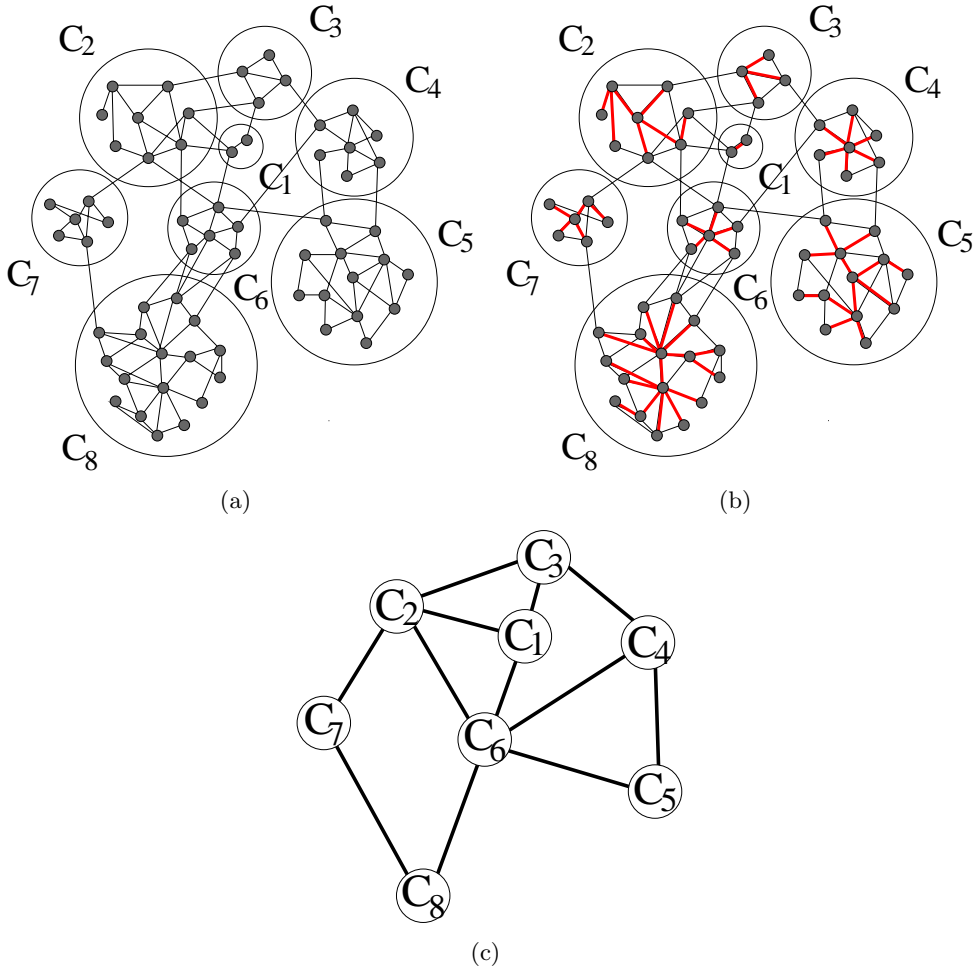


Figure 19: Alon et al.'s approach to constructing a spanning tree with average stretch in $\exp(O(\sqrt{\log n \log \log n}))$. Shown is the first level: (a) A partition in 8 clusters C_1 to C_8 with the properties described in Theorem 4.9; (b) Red edges denote the spanning tree with radius $\leq x \ln m$ for each cluster. All red edges are part of the resulting tree T ; (c) Each cluster is contracted to one vertex, possibly introducing multiple edges between clusters. The resulting graph has less than m/x edges. This graph is the starting point for the next level.

We now come to the construction of the small stretch spanning tree. Figure 19 illustrates the construction. Let P be a partition of $G_1 = G$ as described in the above theorem. For every cluster C_i let T_i be a spanning tree of diameter $2x \ln m$. Such a tree exists by construction. The union of the T_i form a forest F in G . Any intracluster edge, and there are at most m of them, will give rise to a fundamental circuit of length no larger than $1 + 2x \ln m$. Only the m/x intercluster edges can give rise to longer fundamental circuits.

We contract every cluster C_i to a single vertex v_i and obtain the multi-graph G_2 formed by the intercluster edges. We apply the theorem to G_2 and obtain spanning trees of diameter $2x \ln(m/x) \leq 2x \ln m$ for the clusters of G_2 . We add these spanning trees to the forest F . Consider any intracluster edge of G_2 . It gives rise to a cycle of length $1 + 2x \ln m$ in G_2 . With respect to F , this cycle may have length up to $(1 + 2x \ln m)^2$ as any vertex representing a

cluster of G_1 must be expanded to a path of length $1 + 2x \ln m$. We conclude that we might have m/x fundamental circuits of length $(1 + 2x \ln m)^2$. There are at most m/x^2 intercluster edges in G_2 .

The construction continues until graphs of constant size are obtained. The recursion depth is at most $\log_x m$. The total length of the fundamental circuits constructed in this way is

$$\sum_{0 \leq i \leq \log_x m} \frac{m}{x^i} (1 + 2x \ln m)^{i+1} \approx mx(2 \ln m)^{\log_x m}.$$

With $x = \exp(c(\sqrt{\ln n \ln \ln n}))$ for an appropriate constant c , we obtain:

Theorem 4.10 (Alon et al. (1995), (Peleg, 2000, p. 215)). *Every multi-graph has a strictly fundamental basis of length $m \exp(O(\sqrt{\log n \log \log n}))$.*

A much improved result was obtained recently.

Theorem 4.11 (Elkin et al. (2008)). *Every graph has a strictly fundamental cycle basis of length $O(W \log^2 n \log \log n)$.*

The key ingredient for the improved result is a more refined partitioning procedure, called *star-decomposition*. We refer the reader to Elkin et al. (2008) for details. We observe in passing that $\exp(O(\sqrt{\ln n \ln \ln n})) = o(n^\epsilon)$ for any $\epsilon > 0$ and hence even for planar graphs, the bounds given in Theorems 4.10 and 4.11 are better than the bound given in Theorem 4.7.

For dense graphs with $m = \Theta(n^2)$, optimal bounds can be achieved. Already in 1982, Deo et al. (1982) conjectured that every simple graph has a fundamental basis of length $O(n^2)$. It took 25 years to settle the conjecture.

Theorem 4.12 (Elkin et al. (2007)). *Every simple graph on n vertices has a fundamental cycle basis of length $O(n^2)$.*

Proof. Abraham et al. (2007) showed that any graph⁵ G with n vertices contains a spanning tree T with constant average stretch, averaged over all pairs of vertices, i.e.,

$$\sum_{x, y \in \binom{V}{2}} \frac{d_T(x, y)}{d_G(x, y)} = O(n^2).$$

Here $d_G(x, y)$ and $d_T(x, y)$ is the distance between x and y in G and T , respectively. Restricting the sum to the edges of G establishes

$$\sum_{e=(x,y) \in E} d_T(x, y) = O(n^2).$$

Since the length of a fundamental cycle closed by a non-tree edge $e = (x, y)$ is $d_T(x, y) + 1$, the theorem follows. \square

Open Problem 8. *Improve upon Theorem 4.11 or prove a lower bound that is asymptotically larger than $W \log n$ ($n \log n$ in the uniform case).*

⁵The result even holds for weighted graphs; we only need it for unweighted graphs here.

| Exact algorithms, nonnegative weights | | |
|---|--|------------------------------------|
| undirected bases | directed bases | |
| | deterministic | Monte-Carlo |
| $O(\frac{m^2n}{\log n} + mn^2)$ Thm 5.13 | $O(m^3n)$ Thm 5.14 | $O(m^2n)$ Thm 5.15 |
| Exact algorithms, conservative weights | | |
| undirected bases | directed bases | |
| | deterministic | Monte-Carlo |
| $O(n^3 \log n + \frac{m^2n}{\log n} + mn^2)$ Thm 5.16 | $O(m^3n)$ Thm 5.17 | $O(n^3 \log n + m^2n)$ Thm 5.18 |
| $(2k - 1)$ -approximation, integer $k > 1$, nonnegative weights | | |
| undirected bases | directed bases | |
| | deterministic | Monte-Carlo |
| $O(\frac{n^{3+2/k}}{\log n} + n^{3+1/k})$ and $O(kmn^{1+2/k} + mn^{(1+1/k)(\omega-1)})$ Thms. 5.27 and 5.29 | $O(n^{4+3/k})$ Thm. 5.30 | $O(n^{3+2/k})$ Thm. 5.30 |
| 2-approximation, nonnegative weights | | |
| undirected bases | directed bases | |
| expected $O(m^2\sqrt{n/\log n} + n^2m + m^\omega)$ Thm 5.34 | expected $O(m^2\sqrt{n\log n} + n^2m + m^\omega)$ Thm 5.34 | |

Table 3: Polynomial time algorithms for undirected and directed minimum cycle bases.

5 Polynomial Time Algorithms for Minimum Cycle Bases

We will present polynomial time algorithms for computing undirected and directed minimum cycle bases. All algorithms have running time $\Omega(m^2n/\log n + mn^2)$ and therefore cannot be used for very large graphs. Therefore, we will also present techniques for computing approximate minimum cycle bases at much improved running times. Table 3 contains a summary of the best running times. The hard variants of the minimum cycle basis problem will be discussed in Section 6.

Open Problem 9. *Most algorithms discussed in this chapter have space requirement $\Omega(m^2)$. Derive algorithms with reduced space requirement (and maybe increased running time). Derive algorithms for external memory.*

Recall that a directed basis is a set of ν circuits that are independent over \mathbb{Q} and that an undirected basis is a set of ν circuits that are independent over $GF(2)$. We use k to denote either \mathbb{Q} or $GF(p)$ where p is a prime and formulate most of the algorithms in terms of the field k .

5.1 The Greedy Algorithm and the Horton Set

A minimum (directed or undirected) cycle basis can be constructed by a simple greedy algorithm. This is almost a direct consequence of Thm 3.9. We start with an empty basis and

process the circuits of G in order of nondecreasing weight; ties are broken arbitrarily. We add a circuit to the partial cycle basis if it is linearly independent from the circuits in the partial basis. We continue until we have obtained ν linearly independent circuits. Checking linear independence can be easily done by Gaussian elimination.

Theorem 5.1. *The greedy algorithm constructs a minimum weight cycle basis.*

Proof. We could appeal to the fact the the greedy algorithm works for matroids (Korte and Vygen (2005)) and that the set of circuits of a graph form a matroid. We prefer to give a self-contained proof.

Assume otherwise and consider the first time in the execution of the algorithm that the partial basis cannot be extended to a minimum weight basis. Say, this happens after the addition of the circuit C . Before adding C , we had a partial basis B that can be extended to a minimum weight basis B_{opt} . Let us write C as a linear combination of the circuits in B_{opt} , say $C = \sum_{D \in B_{opt}} \lambda_D D$. Since C is linearly independent of B , there must be a $D \in B_{opt} \setminus B$ with $\lambda_D \neq 0$. Also, since this D is linearly independent of B , we must have $w(C) \leq w(D)$. Thus $B_{opt} - D + C$ is also a minimum weight basis, a contradiction. \square

As a graph may have an exponential number of circuits, the performance of the greedy algorithm in its basic form is miserable. Horton (1987) showed that the search for a basis can be restricted to a set of $O(nm)$ circuits. For a vertex v , let T_v be a shortest path tree in G rooted at v . We assume that the collection $(T_v)_{v \in V}$ is *consistent*, i.e., if a vertex x lies on the path to w in T_v , then the subpath from x to w appears in T_x . If shortest paths are unique, consistency is guaranteed.⁶ For any two nodes u and v , we use $p_{u,v}$ to denote the shortest path from u to v contained in T_u . A circuit C is called *isometric* if for any two vertices u and v on C , p_{uv} is contained in C .

Lemma 5.2 (Horton (1987)). *Cycles in minimum D- and U-bases are isometric.*

Proof. Let B be a minimum D- or U-basis and assume that B contains a circuit C that is not isometric. Then there are vertices u and v on C such that C does not contain p_{uv} . Split C at u and v to obtain a path p_1 from u to v and a path p_2 from v to u . Consider the cycles $C_1 = p_1 p_{vu}$ and $C_2 = p_2 p_{uv}$. Then $C = C_1 + C_2$ and hence $B \setminus C \cup C_1$ or $B \setminus C \cup C_2$ is also a basis. Both sets of cycles have lower cost than B . \square

Lemma 5.3 (Horton (1987)). *Let C be any isometric circuit and let x be an arbitrary vertex of C . Then there is an edge $e = (u, v)$ on C such that $C = p_{xu} e p_{vx}$. Conversely, if for any $x \in C$, there is such an edge, then C is isometric.*

Proof. Let $C = (x = v_0, v_1, \dots, v_k = x)$. Since the empty path is the shortest path from x to x and C is not the shortest path from x to x , there must be an i such that $p_{xv_i} = (v_0, v_1, \dots, v_i)$ but $p_{xv_{i+1}} \neq (v_0, v_1, \dots, v_i, v_{i+1})$. Then $p_{xv_{i+1}} = (v_k, v_{k-1}, \dots, v_{i+1})$ and hence $e = (v_i, v_{i+1})$ is the desired edge.

For the converse, consider any two nodes x and z on C and let $e = uv$ be such that $C = p_{xu} e p_{vx}$; z lies on one of the paths and hence the shortest path from x to z is contained in C . \square

⁶Shortest path can be made unique by a suitable tie breaking rule, e.g., by numbering the edges from 1 to m and by changing $w(e_i)$ into $w(e_i) + \epsilon^{2^i}$ where $\epsilon > 0$ is infinitesimal.

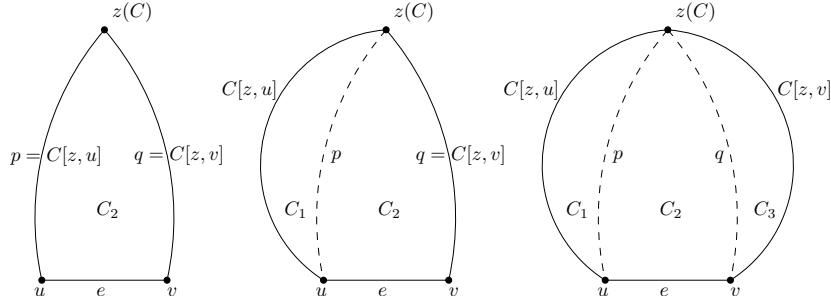


Figure 20: The three cases in the proof of Lemma 5.4 (not showing symmetrical cases).

Definition 5.1 (Horton (1987)). For a vertex v and an edge $e = (x, y)$, let $C[v, e]$ be the cycle consisting of the tree path from v to x in T_v , followed by e , followed by the reversal of the tree path from v to y . The Horton set \mathcal{H} consists of all cycles $C[v, e]$ such that the endpoints of e lie in different subtrees of T_v (in other words, the least common ancestor of the endpoints of e is the root v).

We will next show that \mathcal{H} contains an MCB. For a circuit C let $z(C) \in V \cap C$ be a vertex that minimizes the number of non-tree edges of C w.r.t T_v . We call $z(C)$ the *base node* of C .

Lemma 5.4 (Horton (1987); Mehlhorn and Michail (2008); Liebchen and Rizzi (2005)). \mathcal{H} contains a minimum cycle basis. Moreover, when the greedy algorithm is executed with \mathcal{H} , it extracts a minimum cycle basis.

Proof. Consider the greedy algorithm run on the set of all circuits. Circuits are ordered lexicographically according to

$$(\text{weight of } C, \text{ number of edges outside } T_{z(C)}, \text{ number of edges in } C) .$$

Observe that the circuits in \mathcal{H} have second coordinate equal to one and hence come first among cycles of equal weight.

Let C be the first circuit outside \mathcal{H} that is selected by the greedy algorithm. Let $z = z(C)$ and let $e = (u, v)$ be a nontree edge (with respect to T_z) on C . Write $C = C[z, u] \circ (u, v) \circ C[v, z]$ and let p and q be the tree path in T_z connecting z to u and v , respectively. The cost of p is at most the cost of either cycle path from z to u and the cost of q is at most the cost of either cycle path from z to v . Consider the cycles $C_1 = C[z, u] \circ p^{rev}$, $C_2 = p \circ e \circ q^{rev}$, $C_3 = q \circ C[v, z]$, see Figure 20. The weight of C_1 , C_2 and C_3 is at most the weight of C and $C = C_1 + C_2 + C_3$.

We now distinguish cases. Assume first that e is the only nontree edge on C . Then $C[z, u]$ and $C[v, z]$ are contained in T_z . Since C is a circuit and z lies on C , u and v must lie in distinct subtrees of T_z ; thus $C = C[z, e] \in \mathcal{H}$, a contradiction. Assume next that C contains more than one nontree edge. Then at least one of the cycles C_1 or C_3 is non-trivial. Also, with respect to T_z all three cycles have at least one fewer nontree edge than C and hence this is also true with respect to their respective base vertices.

Thus all three cycles are considered before C by the greedy algorithm. Also, at least one of them is independent of the current basis. So it was independent at the time it was considered and hence should have been added. This either contradicts our definition of C

(first cycle outside \mathcal{H} added to the basis) or the operation of the greedy algorithm (a cycle not added although it is independent). \square

For undirected cycle bases, Lemma 5.4 was first shown by Horton (1987). Mehlhorn and Michail (2008) observed that it suffices to consider a slightly smaller set of circuits. Let Z be a *feedback vertex set* of G , i.e., any circuit in G must contain at least one vertex in Z . Then it suffices to consider the circuits $C[z, e]$ where $z \in Z$ and the endpoints of e belong to different subtrees of T_z . Computing a minimum feedback vertex set is known to be APX-hard, however, a 2-approximation can be computed efficiently (Bafna et al. (1999)). Liebchen and Rizzi (2005) extended Lemma 5.4 to directed bases.

Lemma 5.4 implies polynomial time algorithms for finding a minimum undirected and directed cycle basis. We first construct \mathcal{H} by solving n single-source shortest path problems. In the case of non-negative weights, this amounts to n runs of Dijkstra’s algorithm and takes time $O(nm + n^2 \log n)$. We treat the case of conservative weights in Section 5.5. The Horton set consists of $O(mn)$ circuits and a partial basis consists of at most ν circuits. For any circuit in \mathcal{H} , we must decide whether it is independent of the current partial basis. Gaussian elimination performs this task with $O(\nu m) = O(m^2)$ arithmetic operations per circuit. Let Γ be the cycle matrix of the current basis. We keep the nontree part of Γ in upper triangular form. Then independence of a circuit can be checked with $O(\nu m)$ arithmetic operations and, in the case of independence, the cycle matrix can be extended by an additional column with the same number of arithmetic operations. We conclude that a minimum basis can be constructed with $O(m^3 n)$ arithmetic operations. The number of arithmetic operations can be reduced to $O(m^\omega n)$ (Golynski and Horton (2002); Liebchen and Rizzi (2005)), where ω denotes the exponent of matrix multiplication, i.e., $m \times m$ matrices can be multiplied with $O(m^\omega)$ arithmetic operations. It is known that $\omega < 2.376$.

Arithmetic operations over $GF(2)$ take constant time. We conclude that a minimum weight undirected cycle basis of a nonnegatively weighted graph can be computed in time $O(m^\omega n)$. For directed cycle bases the situation is more difficult. We appeal to Theorem 3.8. Let P be a set of m primes of value at least n . The primes $p \in P$ are in $O(m \log m)$ and hence arithmetic in $GF(p)$ takes constant time. Computing a minimum $GF(p)$ -basis for all $p \in P$ is guaranteed to find a minimum directed basis. This takes time $O(m^{1+\omega} n)$. Computing a minimum $GF(p)$ -basis for a random $p \in P$ takes time $O(m^\omega n)$. It finds a minimum directed basis with probability at least $1/2$.

5.2 De Pina’s Approach

We will describe an alternative approach for computing minimum cycle bases introduced by de Pina (1995) and later refined by Berger et al. (2004); Kavitha et al. (2004); Hariharan et al. (2006); Kavitha et al. (2007); Mehlhorn and Michail (2008). It operates in phases. Starting with an empty set of circuits, it adds one circuit per phase. It does *not* necessarily add the circuits in order of increasing weight. This increased flexibility results in faster running time.

For two vectors C and S , we use $\langle C, S \rangle$ to denote their inner product. Two vectors are *orthogonal* to each other, if their inner product is zero. The following theorem is the basis of de Pina’s approach; the version given here is due to Mehlhorn and Michail (2008). In the original version (de Pina (1995)), the third condition asked for a minimum weight circuit in G .

Algorithm 1 An algebraic framework for computing an minimum cycle basis.

- 1: **for** $i \leftarrow 1, \dots, \nu$ **do**
 - 2: Determine a nonzero vector S_i orthogonal to C_1 to C_{i-1} .
 - 3: Compute a minimum weight cycle $C_i \in \mathcal{H}$ with $\langle C_i, S_i \rangle \neq 0$.
 - 4: **end for**
-

Theorem 5.5 (de Pina (1995); Mehlhorn and Michail (2008)). *Circuits C_1, \dots, C_ν form a minimum k -basis, where $k = \mathbb{Q}$ or $k = GF(p)$, if there are vectors $S_1, \dots, S_\nu \in k^E$ such that for all i , $1 \leq i \leq \nu$:*

1. **Prefix Orthogonality:** $\langle C_j, S_i \rangle = 0$ for all $1 \leq j < i$.
2. **Non-Orthogonality:** $\langle C_i, S_i \rangle \neq 0$.
3. **Shortness:** C_i is a minimum weight circuit in \mathcal{H} with $\langle C_i, S_i \rangle \neq 0$.

Proof. We first show linear independence. Let $C := \sum_i \lambda_i C_i$ be a nontrivial linear combination and assume that i_0 is the largest index for which $\lambda_{i_0} \neq 0$. Then $\langle C, S_{i_0} \rangle = \lambda_{i_0} \langle C_{i_0}, S_{i_0} \rangle \neq 0$.

We next show that the circuits form a minimum cycle basis of G . Suppose not. Then consider the smallest i such that C_1, \dots, C_i are not contained in any minimum cycle basis consisting only of circuits in the Horton set \mathcal{H} . Let B be a minimum weight basis consisting of circuits in the Horton set that contains C_1 to C_{i-1} . We may write C_i as a linear combination of the circuits in B , $C_i = \sum_{C \in B} \lambda_C C$. Since $\langle C_i, S_i \rangle \neq 0$, there exists some $C \in B$ with $\langle C, S_i \rangle \neq 0$. Since C_i is a minimum weight cycle in \mathcal{H} with $\langle C_i, S_i \rangle \neq 0$, we have $w(C_i) \leq w(C)$. Also $C \neq C_j$ for $j < i$ since $\langle C_j, S_i \rangle = 0$ for $j < i$.

Let $B' = B \cup \{C_i\} \setminus \{C\}$; B' is a basis by Thm 3.9 and $w(B') \leq w(B)$. So B' is also a minimum cycle basis. It consists only of circuits in \mathcal{H} and contains C_1 to C_i , a contradiction. \square

Theorem 5.5 leads to Algorithm 1. The algorithm operates in ν phases. In each phase, a nonzero vector S orthogonal to all cycles in the partial basis is determined and then a shortest circuit $C \in \mathcal{H}$ with $\langle S, C \rangle \neq 0$ is computed and added to the basis. We still need to show that there is always a vector S of the desired form and a circuit to add.

Lemma 5.6. *Let T be a spanning tree of G . For each phase i , $1 \leq i \leq \nu$: There is a nonzero vector $S_i \in k^E$ such that $\langle S_i, C_j \rangle = 0$ for $j < i$ and $S_i(e) = 0$ for $e \in T$ and there is at least one cycle $C \in \mathcal{H}$ with $\langle C, S_i \rangle \neq 0$.*

Proof. Let C'_j be the restriction of C_j to $N := E \setminus T$. The space spanned by C'_1 to C'_{i-1} has dimension $i-1$ and $i-1 < \nu$. Thus there is a vector $S' \in k^N$ with $\langle C'_j, S' \rangle \neq 0$ for $j < i$. Define S_i by $S_i(e) = S'(e)$ for $e \in N$ and $S_i(e) = 0$ for $e \in T$.

Let e be any edge with $S_i(e) \neq 0$ and let C_e be the fundamental circuit defined by e . Then $\langle C_e, S_i \rangle = S_i(e) \neq 0$. Since the Horton set contains a basis, C_e can be written as a linear combination of circuits in \mathcal{H} . Thus there must be at least one circuit $C \in \mathcal{H}$ with $\langle C, S_i \rangle \neq 0$. \square

In the next sections we describe how to implement the two main steps of Algorithm 1.

Algorithm 2 Maintaining a Basis of the Orthogonal Space

- 1: Initialize S_j by $S_j(e_i) = \delta_{ij}$ for $1 \leq j \leq \nu$ and $1 \leq i \leq m$.
 - 2: **for** $i \leftarrow 1, \dots, \nu$ **do**
 - 3: Compute a minimum weight cycle $C_i \in \mathcal{H}$ with $\langle C_i, S_i \rangle \neq 0$
 - 4: **for** $j \leftarrow i + 1, \dots, \nu$ **do**
 - 5: $S_j = S_j - \frac{\langle C_i, S_j \rangle}{\langle C_i, S_i \rangle} S_i$
 - 6: **end for**
 - 7: **end for**
-

5.3 Maintaining the Orthogonal Space

The vector S_i is a nontrivial solution of the linear system $\langle C_j, S_i \rangle = 0$ for $1 \leq j < i$. The naive way would be to solve this linear system with Gaussian elimination with $O(m^\omega)$ arithmetic operations. Since we need to solve one linear system per phase, the total number of arithmetic operations required would be $O(m^{1+\omega})$.

However, the linear systems to be solved are not independent. Each phase adds one additional equality. de Pina (1995) and later Berger et al. (2004) observed that it pays off to maintain a basis of the solution space of this linear system. The basis is easily updated from one phase to the next.

Let T be an arbitrary spanning tree of G and let e_1 to e_ν be the nontree edges. We set $S_i(e_i) = 1$ and $S_i(e_j) = 0$ for $j \neq i$. This corresponds to the standard basis of the space k^N . At the beginning of phase i , we have $S_i, S_{i+1}, \dots, S_\nu$ that form a basis of the space \mathcal{C}^\perp orthogonal to the space \mathcal{C} spanned by cycles C_1, \dots, C_{i-1} . We use S_i to compute C_i (see Section 5.4) and update vectors $\{S_{i+1}, \dots, S_\nu\}$ to a basis $\{S'_{i+1}, \dots, S'_\nu\}$ of the subspace of \mathcal{C}^\perp that is orthogonal to C_i . The update step is as follows. For $i + 1 \leq j \leq \nu$, let

$$S'_j = S_j - \frac{\langle C_i, S_j \rangle}{\langle C_i, S_i \rangle} S_i.$$

Lemma 5.7. *The set $\{S'_{i+1}, \dots, S'_\nu\}$ forms a basis of the subspace orthogonal to $\{C_1, \dots, C_i\}$.*

Proof. We will first show that S'_{i+1}, \dots, S'_ν are orthogonal to C_1, \dots, C_i . Let $j \geq i + 1$ and $\ell \leq i$. We have

$$\langle S'_j, C_\ell \rangle = \langle S_j, C_\ell \rangle - \frac{\langle C_i, S_j \rangle}{\langle C_i, S_i \rangle} \langle S_i, C_\ell \rangle.$$

For $\ell < i$, $\langle S_j, C_\ell \rangle = \langle S_i, C_\ell \rangle = 0$. For $\ell = i$, the terms on the right hand side cancel.

Now we will show that S'_{i+1}, \dots, S'_ν are linearly independent. Consider a linear combination

$$\mathbf{0} = \sum_{j \geq i+1} \lambda_j S'_j = \sum_{j \geq i+1} \lambda_j S_j - \left(\sum_{j \geq i+1} \lambda_j \frac{\langle C_i, S_j \rangle}{\langle C_i, S_i \rangle} \right) S_i.$$

Since the S_j , $j \geq i$, are independent, we conclude that $\lambda_j = 0$ for all j . \square

Let us now bound the number of arithmetic operations. In each iteration, we update no more than ν vectors at a cost of $O(\nu)$ arithmetic operations each. Thus the total number of arithmetic operations is $O(\nu^3) = O(m^3)$. For undirected bases, this is also the running time.

The vector S_i is only needed in the i -th phase. In particular, the second half of the vectors is only needed in the second half of the computation. Let $k = \lfloor \nu/2 \rfloor$. Can we save time by

Algorithm 3 Maintaining a Basis of the Orthogonal Space with Bulk Updates

```

1: Initialize  $S_j$  by  $S_j(e_i) = \delta_{ij}$  for  $1 \leq j \leq \nu$  and  $1 \leq i \leq m$ .
2: MinimumCycleBasis( $1, \nu$ )
3: where
4: procedure MINIMUMCYCLEBASIS( $\ell, u$ )                                 $\triangleright$  Adds Circuits  $C_\ell$  to  $C_u$ 
5:   if  $\ell = u$  then
6:     compute a minimum weight cycle  $C_i \in \mathcal{H}$  with  $\langle C_i, S_i \rangle \neq 0$ ;
7:   else
8:      $k \leftarrow \lfloor (\ell + u)/2 \rfloor$ ;
9:     MinimumCycleBasis( $\ell, k$ );
10:     $C \leftarrow [C_\ell, \dots, C_k]$ ;
11:     $A \leftarrow (C^T[S_\ell, \dots, S_k])^{-1}C^T[S_{k+1}, \dots, S_u]$ ;
12:     $[S_{k+1}, \dots, S_u] \leftarrow [S_{k+1}, \dots, S_u] - [S_\ell, \dots, S_k]A$ ;     $\triangleright$  now  $C^T[S_{m+1}, \dots, S_u] = \mathbf{0}$ 
13:    MinimumCycleBasis( $m + 1, u$ );
14:   end if
15: end procedure

```

not updating the vectors S_{k+1} to S_ν in the first k phases at all and computing the cumulative effect of these update after phase k ? We will see that we can use fast matrix multiplication for the cumulative update. What is the effect of the first k phases on the vectors S_{k+1} to S_ν ?

For column vectors v_1 to v_ℓ , we use $[v_1, \dots, v_\ell]$ to denote the matrix with columns v_1 to v_ℓ . Let S_1 to S_ν denote our vectors before phase 1 and let S'_1 to S'_ν be the vectors after phase k . Then

$$[S'_{k+1}, \dots, S'_\nu] = [S_{k+1}, \dots, S_\nu] - [S'_1, \dots, S'_k]A$$

for some $k \times (\nu - k)$ matrix A . We want $\langle C_\ell, S'_j \rangle = 0$ for $1 \leq \ell \leq i$ and $i + 1 \leq j \leq \nu$. Let $C = [C_1, \dots, C_k]$. Then

$$\mathbf{0} = C^T[S'_{k+1}, \dots, S'_\nu] = C^T[S_{k+1}, \dots, S_\nu] - C^T[S'_1, \dots, S'_k]A$$

and hence

$$A = (C^T[S'_1, \dots, S'_k])^{-1}C^T[S_{k+1}, \dots, S_\nu].$$

Since $\langle C_\ell, S'_i \rangle = 0$ for $1 \leq \ell < i \leq k$ and $\langle C_i, S'_i \rangle \neq 0$, the matrix $C^T[S'_1, \dots, S'_k]$ is lower triangular with nonzero entries on the diagonal and hence invertible. We need to compute three matrix products and one matrix inversion. Each of them can be performed with $O(m^\omega)$ arithmetic operations. We conclude that the cumulative update of S_{k+1} to S_ν at the end of phase k requires only $O(m^\omega)$ arithmetic operations instead of the $\Theta(m^3)$ operations for the continuous update. We can carry this idea further by applying it recursively, for example, we do not update $S_{\lfloor k/2 \rfloor + 1}$ to S_k in the first $\lfloor k/2 \rfloor$ phases, but do a bulk update of these vectors after phase $\lfloor k/2 \rfloor$. We obtain Algorithm 3.

Consider a call of procedure MinimumCycleBasis which is not innermost and let $r = u - \ell + 1$, $s = k - \ell + 1$ and $t = u - k$. In the update of the vectors S_{k+1} to S_u we perform (s, m, s) , (s, m, t) , (s, s, t) , (m, s, t) matrix multiplications⁷, one inversion of an $s \times s$ matrix and one addition of two $m \times t$ matrices. If we split all matrices into blocks of $s \times s$ matrices and

⁷An (a, b, c) matrix multiplication multiplies a $a \times b$ matrix with a $b \times c$ matrix.

use fast matrix methods for the blocks, the update requires $O((m/s)s^\omega)$ arithmetic operations. The total number U of arithmetic operations for all updates follows the recursion

$$U(r) = \begin{cases} 0 & \text{if } r = 1 \\ O((ms^{\omega-1}) + U(s) + U(r-s)) & \text{if } r > 1 \text{ and } s = \lceil r/2 \rceil. \end{cases}$$

This recurrence solves to $U(r) = O(mr^{\omega-1})$. In our outermost call $r = \nu = O(m)$. We conclude that the total number of arithmetic operations in the update steps is $O(m^\omega)$.

Lemma 5.8. *The total number of arithmetic operations spent in lines 10 to 12 of Algorithm 3 is $O(m^\omega)$. In a computation over $GF(p)$ with $\log p = O(\log m)$, the time spent in lines 10 to 12 is $O(m^\omega)$.*

5.4 Computing the Circuits

We now come to the second main ingredient of the minimum cycle basis algorithm. Given a nonzero vector S , compute a minimum weight circuit C with $\langle S, C \rangle \neq 0$. We know from Theorem 5.5 that the search can be restricted to \mathcal{H} . We will exploit this fact in Section 5.4. Now, we will show how to find C without this additional knowledge.

We first consider the undirected case and nonnegative edge weights and reduce the computation to n shortest path computations. Over $GF(2)$, the vector S is zero-one and therefore corresponds to a subset of E ; $\langle S, C \rangle \neq 0$ if and only if C uses an odd number of edges in S . The following construction is well known (Barahona and Mahjoub (1986); Grötschel et al. (1988)). The *signed graph* G_S is defined from $G = (V, E)$ and S in the following manner. G_S has two copies for each vertex $v \in V$. Call them v^+ and v^- . Let $e = (u, v)$ be any edge of G . If $e \notin S$, we put the edges (v^+, u^+) and (v^-, u^-) into G_S and if $e \in S$, we put the edges (v^+, u^-) and (v^-, u^+) into G_S . In either case, the edges inherit the weight of e . Figure 21 illustrates the construction. The vertices of G_S naturally split into a $+$ level and a $-$ level. Edges of G_S corresponding to edges in $E \setminus S$ connect vertices in the same level, and edges corresponding to edges in S connect vertices in opposite levels. Each level we have edges of $E \setminus S_i$.

A path in G starting at a node v lifts to two paths in G_S , one starting in v^+ and one starting in v^- . The path ends in the other level if and only if it uses an odd number of edges in S . So a circuit passing through v and using an odd number of edges in S lifts to a simple path of the same weight connecting v^+ and v^- . The lifted path does not use both copies in G_S of an edge of G . Conversely, consider a path p connecting v^+ to v^- in G_S . It may use both copies of an edge of G . In our example, the path $\langle 3^+, 4^+, 1^-, 2^-, 4^-, 3^- \rangle$ uses both copies of $(3, 4)$. We split

$$p = \langle v^+, \dots, x^* \rangle \langle x^*, y^\dagger \rangle \langle y^\dagger \dots y^{-\dagger} \rangle \langle y^{-\dagger}, x^{-*} \rangle \langle x^{-*}, \dots, v^- \rangle$$

at the two copies of an edge, say (x, y) such that the “middle part” $q = \langle y^\dagger \dots y^{-\dagger} \rangle$ does not use both copies of any edge; q connects y^+ and y^- and $w(q) \leq w(p)$ since edge weights are nonnegative. We summarize the discussion in:

Lemma 5.9. *For each $v \in V$, let p_v be a minimum weight minimum cardinality path⁸ from v^+ to v^- in G_S . Let v_0 be such that p_{v_0} has minimum weight among the paths p_v . Break ties*

⁸A minimum weight minimum cardinality path from v^+ to v^- is a minimum weight path from v^+ to v^- . Among the minimum weight paths, it has a minimum number of edges.

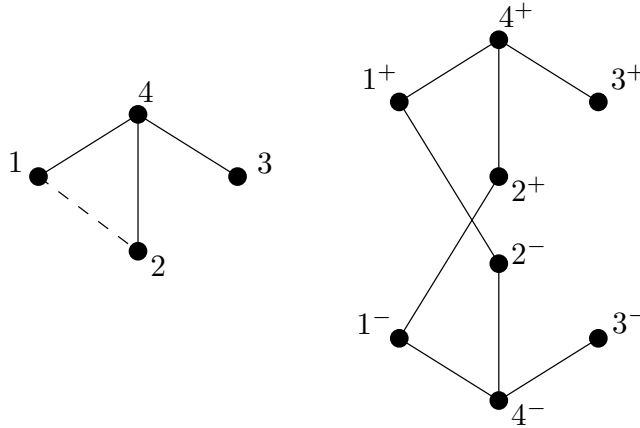


Figure 21: An example of the graph G_S , where $S = \{(1, 2)\}$. Since the edge $(1, 2)$ belongs to S we have the edges $(1^-, 2^+)$ and $(1^+, 2^-)$ going across the $-$ and $+$ levels. The edges not in S , i.e., $(1, 4)$, $(2, 4)$, and $(3, 4)$ have copies inside the $+$ level and the $-$ level.

in favor of the path containing fewer edges. Let $C = C_{v_0}$ be the projection of p_{v_0} into G . C is a minimum weight cycle in G using an odd number of edges in S .

The computation of the path p_{v_0} can be performed by computing n shortest (v^+, v^-) paths, one for each vertex $v \in V$, each by Dijkstra's algorithm in G_S and taking their minimum, or by one invocation of an all pairs shortest paths algorithm in G_S . This computation takes $O(n(m + n \log n))$. Note that depending on the relation between m and n , we may choose which shortest paths algorithm to use. For example, in the case when the edge weights are integers or the unweighted case it is better to use faster all pairs shortest paths algorithms than run Dijkstra's algorithm n times.

Computation over $GF(p)$: The signed graph technique extends to computations over $GF(p)$ (Kavitha and Mehlhorn (2005)). The entries of the vector S are now in $\{0, \dots, p-1\}$. Accordingly, we have p levels and p copies v^0 to v^{p-1} of each edge. An edge $e \in E$ with $s = S(v)$ gives rise to edges (v^i, v^{i+s}) for $0 \leq i < p$. Superscripts are to be read modulo p . Everything else is as before. Because of the larger graph, the cost of the shortest path computation is multiplied by p .

Hariharan et al. (2006) were able to remove the factor of p in the running time. Consider a shortest path computation starting at v^0 . The algorithm outlined in the previous paragraph computes for each w and each $i \in \{0, \dots, p-1\}$ a shortest path to w^i . The improved algorithm computes for any w only two paths. Let i_0 be such that the path from v^0 to w^{i_0} is no longer than to any w^i and let i_1 be such that the path from v^0 to w^{i_1} is no longer than to any w^i with $i \neq i_0$. The algorithm computes the paths to w^{i_0} and w^{i_1} . This can be done in Dijkstra-time.

We do not go into more detail as the following section presents a simpler and faster approach. Furthermore, the technique is the same for all $GF(p)$.

Labelled Trees: We know from Theorem 5.5 that the search for a shortest circuit C_i with $\langle C_i, S_i \rangle \neq 0$ in line 6 of Algorithm 3 may be restricted to the circuits in \mathcal{H} . A compact representation of the circuits in \mathcal{H} is given by the shortest path trees T_v , $v \in V$. For $v \in V$,

each edge $e = (x, y)$ connecting vertices in distinct subtrees of T_v gives rise to the circuit $C[v, e] \in \mathcal{H}$.

How can we compute $\langle C[v, e], S_i \rangle$ efficiently? The idea (Mehlhorn and Michail (2008)) is to precompute most of the inner product. For any v and w , let $p_{v,w}$ be the path from v to w in T_v . We label w in T_v with $\ell_{v,w} = \langle p_{v,w}, S_i \rangle$. For fixed v , the labels $\ell_{v,w}$ can be computed in $O(n)$ arithmetic operations. It takes $O(n^2)$ arithmetic operations to label all trees. Once, the labels are available, $\langle C[v, e], S_i \rangle$ can be computed with a constant number of arithmetic operations. If $e = (x, y)$,

$$\langle C[v, e], S_i \rangle = \ell_{v,x} + S_i(e) - \ell_{v,y}.$$

Lemma 5.10. *If the shortest path trees T_v , $v \in V$, are available, the minimum weight cycle $C \in \mathcal{H}$ with $\langle C, S_i \rangle \neq 0$ can be found with $O(nm)$ arithmetic operations.*

5.5 Computing Shortest Path Trees

For nonnegative edge weights, we use Dijkstra’s algorithm and obtain:

Lemma 5.11. *If edge weights are nonnegative, the shortest path trees T_v , $v \in V$, can be computed in time $O(n(m + n \log n))$.*

For conservative edge weights, heavier machinery needs to be used. It is known that computing all pairs shortest paths in undirected graphs with real edge weights but no negative cycles can be computed by solving a sequence of general weighted matching problems.

Lemma 5.12. *If edge weights are conservative, the shortest path trees T_v , $v \in V$, can be computed in time $O(n^2m + n^3 \log n)$.*

Proof. The single-sink-single-source shortest path problem in a conservatively weighted undirected graph reduces to a weighted perfect matching problem in a graph with $O(n)$ vertices and $O(m)$ edges ((Korte and Vygen, 2005, page 278)) and hence can be solved in time $O(n(m + n \log n))$ (Gabow (1990)). The construction of the perfect matching problem consists of n “searches”; each search takes time $O(m + n \log n)$. The all-pairs shortest path problem can be reduced to a perfect matching problem plus n^2 searches ((Korte and Vygen, 2005, page 279). \square

5.6 Putting it Together

We can now put the pieces together.

Theorem 5.13 (Kavitha et al. (2004); Mehlhorn and Michail (2008)). *For nonnegative weight functions, a minimum weight undirected cycle basis can be computed in time $O(m^2n/\log n + mn^2)$.*

Proof. It takes $O(nm + n^2 \log n)$ to compute the shortest path trees (Lemma 5.11, time $O(m^\omega)$ (Lemma 5.8 to compute the S_i , $1 \leq i \leq \nu$, and time $O(nm^2)$ to determine the cycles C_i , $1 \leq i \leq \nu$). The total running time is $O(m^2n)$.

Mehlhorn and Michail (2008) have shown that word-parallelism on words of length $O(\log n)$ can be used to extract the cycles in time $O(m^2n/\log n)$ at the cost of increasing preprocessing time to $O(mn^2)$. \square

Theorem 5.14 (Hariharan et al. (2006); Mehlhorn and Michail (2008)). *For nonnegative weight functions, a minimum weight directed cycle basis can be computed in time $O(m^3n)$.*

Proof. By Theorem 3.8, it suffices to compute the minimum $GF(p)$ -basis for m primes larger than n . The best such basis is a minimum weight directed cycle basis.

For each fixed p , it takes $O(nm+n^2 \log n)$ to compute the shortest path trees (Lemma 5.11), time $O(m^\omega)$ (Lemma 5.8) to compute the S_i , $1 \leq i \leq \nu$, and time $O(nm^2)$ to determine the cycles C_i , $1 \leq i \leq \nu$. The total running time is $O(m^2n)$ for each p and hence $O(m^3n)$, altogether. \square

Theorem 5.15 (Hariharan et al. (2006); Mehlhorn and Michail (2008)). *For nonnegative weight functions, a minimum weight directed cycle basis can be computed in time $O(m^2n)$ with probability at least $1/2$.*

Proof. By Theorem 3.8, it suffices to compute the minimum $GF(p)$ -basis for a prime p chosen randomly from a set of m primes larger than n . For such a prime the minimum $GF(p)$ -basis can be computed in time $O(m^2n)$. \square

Theorem 5.16. *For conservative weight functions, a minimum undirected cycle basis can be computed in time $O(n^3 \log n + m^2n/\log n + mn^2)$.*

Proof. Follows from Lemmas 5.12, 5.8, 5.10, and the remark made in the proof of Theorem 5.13. \square

Theorem 5.17. *For conservative weight functions, a minimum directed cycle basis can be computed in time $O(m^3n)$.*

Proof. Follows from Theorem 3.8, Lemmas 5.12, 5.8, and 5.10. \square

Theorem 5.18. *For conservative weight functions, a minimum directed cycle basis can be computed in time $O(n^3 \log n + m^2n)$ with probability at least $1/2$.*

Proof. Follows from Theorem 3.8, Lemmas 5.12, 5.8, and 5.10. \square

5.7 A Greedy Algorithm for Integral Cycle Bases?

Both the Greedy Algorithm (Section 5.1) and De Pina's Approach (Section 5.2) fundamentally rely on Theorem 3.10, namely the fact that all subsets of K -bases in G constitute a matroid for $K \in \{D, U\}$. This is not true for integral bases.

Theorem 5.19. (Liebchen and Rizzi (2005)) *The system of all subsets of integral cycle bases in G is not a matroid.*

Proof. We exhibit a graph with two integral cycle bases B_1 and B_2 and a circuit $C_1 \in B_1 \setminus B_2$ such that for no circuit $C_2 \in B_2 \setminus B_1$, $B_1 \setminus \{C_1\} \cup \{C_2\}$ is again an integral basis.

Consider the directed envelope graph shown in Figure 22 and the spanning tree T indicated by the bold edges. The bases B_1 and B_2 are given by the cycle matrices (only the parts corresponding to non-tree edges are shown)

$$\Gamma_1 = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \Gamma_2 = \begin{pmatrix} 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 \end{pmatrix}.$$

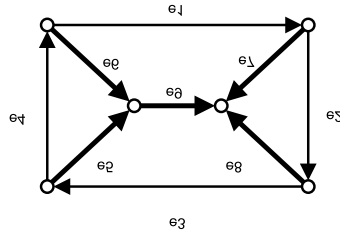


Figure 22: The envelope graph.

The bases are integral since $|\det \Gamma_1| = |\det \Gamma_2| = 1$. Now choose the circuit in the first column of Γ_1 , call it C_1 , to exit the basis. Of course neither the third nor the fourth circuit in B_2 can replace C_1 since both already appear in $B_1 \setminus \{C_1\}$. But adding the first or the second circuit of B_2 results in a cycle basis of determinant 2 and 3, respectively. \square

Theorem 5.19 does not yet imply the failure of the greedy algorithm nor of De Pina's approach, since the weights of cycles in G cannot be chosen independently for each cycle. A greedy algorithm for integral basis would consider circuits in order of increasing weight. It would maintain a partial basis that can be extended to an integral basis and add a circuit to the current basis if this property is maintained. It is not known how to implement this strategy efficiently. It would not work, anyhow.

Theorem 5.20. *The greedy algorithm may end up with a non-optimal cycle basis of G .*

Proof. We again consider the graph introduced in Lemma 3.20 together with the same two integral bases B_1 and B_2 depicted in Figure 23 and Figure 24. In contrast to Lemma 3.20 we assign other weights to the edges. Let every inner and outer edge have a weight of 5 whereas every spoke has a weight of 19. Then the first 22 cycles in B_1 are the only ones in G whose weights do not exceed 53. Moreover there are exactly two cycles, the inner cycle C_I and the outer cycle C_O , with weight 55 and the weight of every other cycle is at least 58. Under this assignment of weights B_1 has a total weight of 1169 whereas the weight of B_2 equals 1168.

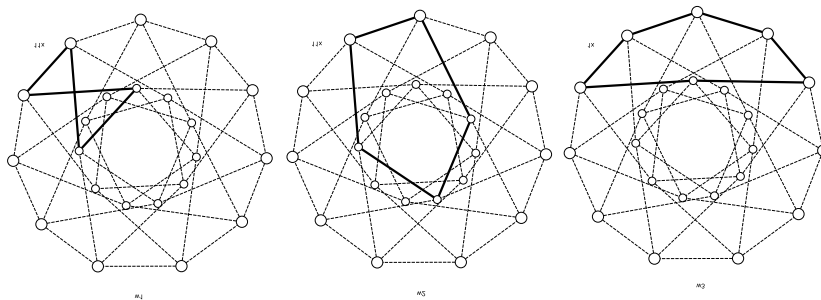


Figure 23: An integral cycle basis B_1 of G with total weight 1169.

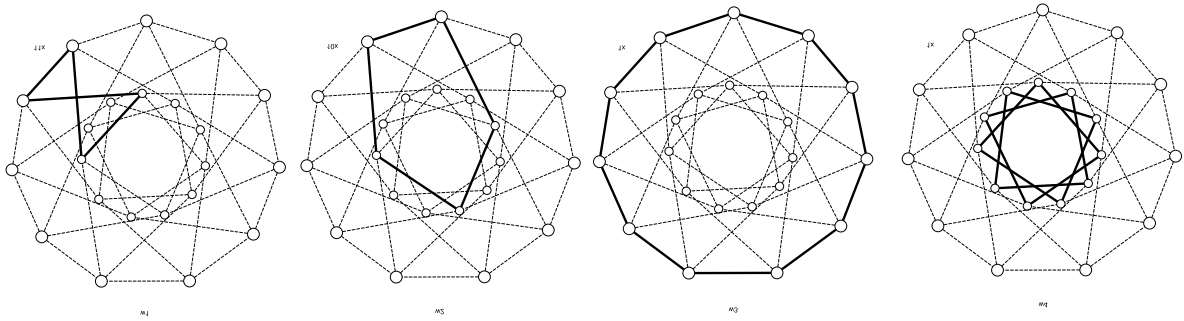


Figure 24: The (unique) minimum integral cycle basis B_2 of G with total weight 1168.

As a consequence, B_2 is the unique minimum integral cycle basis. On the other hand, the Greedy Algorithm performs by first picking the 22 cycles of weight at most 53. These cannot be extended to an integral basis by adding C_I nor C_O and hence the Greedy Algorithm will end up with a bases similar to B_1 and thus with a non-optimal basis. \square

Finally we observe that the basis B_1 in the preceding proof constitutes a non-optimal but *locally optimal* integral cycle basis of G , i.e. B_1 cannot be improved by an exchange of a single cycle in G . This is true since the only two exchanges which would decrease the weight of B_1 are the replacement of the 58-circuit by either C_I or C_O but both results in a non-integral basis. Hence a Local-Search-Procedure fails in general; de Pinas approach can be interpreted as such a Local-Search.

5.8 Planar Graphs

For planar graphs, a minimum undirected cycle basis in time $O(n^2 \log n)$, a minimum 2-basis can be computed in linear time, and the notions of minimum directed, undirected, integral, and weakly fundamental basis coincide.

Let G be a plane graph, i.e., a planar graph that is embedded into the plane. A plane graph divides the plane into maximal open connected sets of points that we call *faces*. Any circuit C divides the plane into two maximal open connected sets of points, one bounded and one unbounded. We use $interior(C)$ to denote the bounded set. If $interior(C)$ agrees with one of the faces of G , we call C a *face cycle*. A collection of circuits is called *nested* if for any two circuits C and D in the collection, the interiors are either disjoint or the interior of one is contained in the interior of the other.

For a collection B of circuits, let F_B be the face circuits that do not belong to B . We define the directed inclusion graph D_B with vertex set $B \cup F_B$ as follows. Let C and C' be circuits in $B \cup F_B$. We have an edge from C to C' if $interior(C) \subset interior(C')$ and there is no circuit $C'' \in B \cup F_B$ such that $interior(C) \subset interior(C'') \subset interior(C')$. The inclusion graph is acyclic; the sources of the inclusion graph are precisely the faces circuits of G . The inclusion graph is a forest if and only if B is nested.

Theorem 5.21 (Hartvigsen and Mardon (1994)). *Let G a (directed) plane graph. G has a minimum (directed or undirected) cycle basis that is nested. The number of isometric cycles is at most twice the number of facets of G .*

Proof. The circuits in a basis are isometric (Lemma 5.2). In a plane graph, any two isometric circuits have either disjoint interiors or the interior of one is contained in the interior of the other.

Let B be the set of all isometric circuits. The inclusion graph D_B is a forest with f leaves, where f is the number of facets of G . Each nonleaf has indegree at least two. Thus the number of nonleaves is at most $f - 1$. \square

Theorem 5.22 (Hartvigsen and Mardon (1994)). *Let G be (directed) plane graph. A nested collection B of circuits is a minimum (directed or undirected) cycle basis iff B is a minimum weight collection of circuits satisfying the following three properties:*

1. *the inclusion graph D_B is a forest,*
2. *every nonleaf in D_B has exactly one child in F_B , and*
3. *the circuits in F_B have parents in D_B .*

Proof. Assume first that B is a basis. We first observe that the number of circuits in B that are not face circuits is equal to the number of face circuits that do not belong to B since the face circuits form a basis and all bases have the same cardinality.

If B is nested, the inclusion graph is a forest. Consider any nonleaf C of D_B . If no child of C belongs to F_B , C is the sum of its children and B is not a basis. Thus any nonleaf C has at least one child in F_B . The nonleaves of the inclusion graph are precisely the circuits in B that are not face circuits. Thus, any nonleaf has exactly one child in F_B and every circuit in F_B must have a parent.

Conversely, assume that B is a minimum cost collection of circuits satisfying (1) to (3). Since D_B is a forest, B is nested. Since the circuits in F_B have parents in D_B and these parents are distinct, the number of nonleaves in D_B is exactly the number of circuits in F_B . So B has the right number of circuits for a basis. Finally, any face circuit is representable as a sum of circuits in B . This is obvious for the face circuits that belong to B . For the face circuits in F_B , it follows from (2) and (3). \square

We come to the algorithm for finding a minimum weight basis. We start by computing the shortest paths trees T_v for all vertices v ; this can certainly be done in time $O(n^2 \log n)$ (Frederickson (1987)). The Horton set consists of $O(n^2)$ circuits. The set of isometric cycles can be extracted from the Horton set in time $O(n^2 \log n)$ (Hartvigsen and Mardon (1994)). The number of isometric circuits is $O(n)$ (Theorem 5.21) and so sorting them by weight takes time $O(n \log n)$.

Next we construct the incidence matrix A between isometric circuits and the facets of G . The entry corresponding to a circuit C and a facet F is one if $F \subseteq \text{interior}(C)$. This matrix can clearly be computed in time $O(n^2)$.

We initialize the basis B to the empty set and set up the corresponding inclusion graph D_B . The vertices of D_B are the face circuits and there are no edges. As long as B does not have the right number of circuits and hence D_B does not satisfy (2) and (3), we do the following.

If there is a nonleaf node C that has two children in F_B (case 1), let R_1 and R_2 be two face circuits in F_B having C as their common parent. If there is no such nonleaf node, there must be a face circuit in F_B without a parent (case 2). Let R_1 be this face and let R_2 be the

unbounded face. In either case, we find the least weight circuit D containing exactly one of R_1 or R_2 in its interior. We can find D in time $O(n)$ by scanning the columns of A .

We add D to B and update D_B . If D is a face circuit, we only have to remove D from F_B . The inclusion graph stays the same. So assume that D is not a face circuit. Starting from the face circuits in $\text{interior}(D)$ (we can find them in matrix A), we determine the maximal subtrees of D_B that are contained in $\text{interior}(D)$. They become children of D . D either becomes a root (in case 2) or a child of C (in case 1). Updating D_B takes time $O(n)$. We conclude that we spend time $O(n)$ per base circuit for a total of $O(n^2)$.

Theorem 5.23 (Hartvigsen and Mardon (1994)). *A minimum (directed or undirected) circuit basis of a planar graph can be found in time $O(n^2 \log n)$.*

Theorem 5.24. *The algorithm by Hartvigson and Mardon constructs a minimum cycle basis that is weakly fundamental.*

Proof. The algorithm constructs a basis B that is nested. We show that B is weakly fundamental. Let D_B be the inclusion graph corresponding to B . If F_B is empty, every face circuit belongs to B , and the construction of Theorem 3.5 works. So assume that F_B is non-empty. Since every circuit in F_B has a parent, we have a non-leaf node C in D_B all of whose children are face circuits. One of these face circuits, say F belongs to F_B and the others belong to B . There must be at least one edge on the boundary of F that does not belong to C as, otherwise, $C = F$. Let F' be the other face circuit incident to e and let p be the maximal path containing e and having all interior vertices of degree two. We remove the edges of p from the graph, assign F' to e , delete F' from B , and add the edges of $p \setminus e$ to the spanning tree. Removal of p merges F and F' and $B \setminus F'$ is a basis for the modified graph. Continuing in this way constructs an elimination order for the edges. \square

We close this section with a discussion of 2-bases of planar graphs. Every planar graph has a 2-basis (take the boundary of each finite face), and, conversely, only planar graphs have 2-bases. A linear time algorithm for solving the minimum 2-basis problem on planar graphs, possibly with weights on the edges, was given in Liebchen and Rizzi (2007). This is based on the use of SPQR-trees, a data structure able to represent in a compact way all possible combinatorial embeddings of a 2-connected planar graph.

5.9 Approximation

Minimum directed and undirected cycle bases can be computed in polynomial time. However, the running times are fairly high degree polynomials, too high for applications, e.g., circuit analysis, that need to find cycle bases of graphs with several million vertices and edges. However, in these applications, a nearly optimal basis is almost as good as an optimal basis. It is therefore natural to explore approximation algorithms. The results presented in this section are based on Kavitha et al. (2004, 2007). We will present two approximation techniques. The first technique uses de Pina's approach but replaces shortest path computations by approximate shortest path computations. The second technique uses Horton's approach and replaces the Horton set \mathcal{H} by a smaller set of circuits that is guaranteed to contain a 2-approximate cycle basis. We start with lower bounds that we will use in our quality estimates.

Lemma 5.25 (de Pina (1995)). *Let R_1, \dots, R_ν be linearly independent vectors in k^ν and let A_i be a shortest cycle in G such that $\langle A_i, R_i \rangle \neq 0$. Then $\sum_{i=1}^\nu w(A_i)$ is a lower bound on the weight of any k -basis.*

Proof. Let $\{C_1, \dots, C_\nu\}$ be a k -basis. We may assume without loss of generality that the A_i 's and C_i 's are sorted by weight, that is, $w(A_1) \leq w(A_2) \leq \dots \leq w(A_\nu)$ and $w(C_1) \leq w(C_2) \leq \dots \leq w(C_\nu)$. The former may require a renumbering of the R_i 's. We will show $w(A_i) \leq w(C_i)$ for all i .

Consider a fixed i and observe that $\langle C_k, R_\ell \rangle \neq 0$ for some k and ℓ with $1 \leq k \leq i \leq \ell \leq \nu$. Otherwise, the $\nu - i + 1$ linearly independent vectors $R_i, R_{i+1}, \dots, R_\nu$ belong to the subspace orthogonal to C_1, \dots, C_i ; however, this subspace has dimension only $\nu - i$. Thus, $w(A_\ell) \leq w(C_k)$ since A_ℓ is a shortest cycle with $\langle A_\ell, R_\ell \rangle \neq 0$ and hence $w(A_i) \leq w(A_\ell) \leq w(C_k) \leq w(C_i)$. \square

Corollary 5.26. *Let G be a graph. For any edge e , let SC_e be the minimum weight cycle containing e . Then $\sum_{e \in E} w(SC_e)$ is a lower bound on the weight of any cycle basis.*

Proof. Let R_e be the unit vector whose entry corresponding to e is one. The vectors R_e , $e \in E$, are clearly independent (over \mathbb{Q} and over $GF(p)$) and $\langle R_e, SC_e \rangle = 1 \neq 0$. Clearly SC_e is the shortest cycle C with $\langle R_e, C \rangle \neq 0$. \square

5.9.1 Approximate Shortest Paths

De Pina's approach works in phases. In each phase, we compute a support vector S and a shortest circuit C with $\langle S, C \rangle \neq 0$. If instead of searching for a shortest circuit, we search for a t -approximation of it, we should obtain an t -approximate cycle basis. We next show how to realize this idea for any integer $k > 1$ and $t = 2k - 1$.

A t -spanner of an undirected graph G is a subgraph G' of G such that for any two vertices u and v , the distance from u to v in G' is at most t times their distance in G . Althöfer et al. (1993) showed that every weighted undirected graph on n vertices has a $(2k - 1)$ -spanner with $O(n^{1+1/k})$ edges. Such a spanner is easily constructed incrementally. We start with an empty graph G' and consider the edges of G in non-decreasing order of weight. When an edge is considered, we add it to G' , if its endpoints are not already connected by a path using at most $2k - 1$ edges of G' ; otherwise, we discard it. At any stage, G' is a $(2k - 1)$ -spanner of the edges already considered, and its unweighted girth⁹ is at least $2k + 1$, so it has only $O(n^{1+1/k})$ edges. The above procedure can be implemented to run in time $O(mn^{1+1/k})$. From now on, $G' = (V, E')$ denotes a t -spanner of G . Let $\lambda = |E \setminus E'|$ and $m' = |E'| = m - \lambda$. Observe that $\nu' = m' - n + 1 = m - n + 1 - \lambda = \nu - \lambda$.

For each edge $e = (v, w) \in E \setminus E'$, let C_e be the circuit consisting of e and the shortest path, say p , in G' connecting v and w . Then

$$w(C_e) = w(e) + w(p) \leq w(e) + t \operatorname{dist}_G(u, v) \leq t w(SC_e).$$

The circuits C_e , $e \in E \setminus E'$, are clearly independent and form the first λ circuits in our t -approximate basis. The cost of constructing these λ circuits is the cost of λ shortest path computations in G' and hence bounded by $O(\lambda \cdot (n^{1+1/k} + n \log n))$. Since $\lambda \leq m$ we can compute both the spanner and the λ circuits in time $O(mn^{1+1/k})$.

We need an additional $\nu - \lambda$ circuits for a basis. We outline one approach and then discuss a second approach in more detail. The first approach now switches to the recursive algorithm in Section 5.3. It first computes a basis $S_{\lambda+1}, \dots, S_\nu$ of the subspace orthogonal to C_e , $e \in E \setminus E'$ and then proceeds as in Section 5.3, see Algorithm 4. Instead of computing a shortest cycle in

⁹The girth of a graph is the minimum number of edges in any circuit.

Algorithm 4 Approximation algorithm. Best performance for sparse graphs.

- 1: **procedure** SPANNER-APPROX-SPARSE(Graph G)
 - 2: Construct a $(2k - 1)$ -spanner G' with $O(n^{1+1/k})$ edges. Let e_1, \dots, e_λ be the edges of $G \setminus G'$.
 - 3: For $1 \leq i \leq \lambda$ let $C_i = e_i + p_i$ where $e_i = (u_i, v_i)$ and p_i is a shortest path in G' from u_i to v_i .
 - 4: Find linearly independent $S_{\lambda+1}, \dots, S_\nu$ in the subspace orthogonal to cycles C_1, \dots, C_λ .
 - 5: Call the recursive algorithm in Section 5.3 with input: the graph G , sets $\{C_1, \dots, C_\lambda\}$, $\{S_{\lambda+1}, \dots, S_\nu\}$ and $\nu - \lambda$ to compute $(2k - 1)$ -approximate cycles $C_{\lambda+1}, \dots, C_\nu$.
 - 6: Return $\{C_1, \dots, C_\lambda\} \cup \{C_{\lambda+1}, \dots, C_\nu\}$.
 - 7: **end procedure**
-

Algorithm 5 Approximation algorithm. Best performance for dense graphs.

- 1: **procedure** SPANNER-APPROX-DENSE(Graph G)
 - 2: Construct a $(2k - 1)$ -spanner G' with $O(n^{1+1/k})$ edges. Let e_1, \dots, e_λ be the edges of $G \setminus G'$.
 - 3: For $1 \leq i \leq \lambda$ let $C_i = e_i + p_i$ where $e_i = (u_i, v_i)$ and p_i is the shortest path in G' from u_i to v_i .
 - 4: Call the best exact algorithm to find an MCB of G' . Let these cycles be $C_{\lambda+1}, \dots, C_\nu$.
 - 5: Return $\{C_1, \dots, C_\lambda\} \cup \{C_{\lambda+1}, \dots, C_\nu\}$.
 - 6: **end procedure**
-

each phase, it computes a t -approximate shortest path using the *approximate distance oracle* of Thorup and Zwick (2001a). This data structure answers $(2k - 1)$ -approximate shortest path queries in time $O(k)$. The structure requires space $O(kn^{1+1/k})$ and can be constructed in expected time $O(kmn^{1/k})$.

Theorem 5.27 (Kavitha et al. (2007); Mehlhorn and Michail (2008)). *For any integer $k \geq 2$, Algorithm 4 computes a $(2k - 1)$ -approximate undirected cycle basis in expected time $O(kmn^{1+2/k} + mn^{(1+1/k)(\omega-1)})$.*

The second approach is even simpler. We complete the basis by computing a minimum cycle basis of the t -spanner G' , see Algorithm 5. The dimension of the cycle space of G' is $\nu' = \nu - \lambda$ and thus we have the right number of circuits. Let $C_{\lambda+1}, \dots, C_\nu$ be a minimum cycle basis of G' . Circuits $\{C_1, \dots, C_\lambda\} \cup \{C_{\lambda+1}, \dots, C_\nu\}$ are by definition linearly independent and we are also going to prove that they form a t -approximation of an MCB of G .

For $1 \leq i \leq \lambda$, we have $C_i = e_i + p_i$, where p_i is a shortest path in G' between the endpoints of e_i . In order to show that cycles C_1, \dots, C_ν are a t -approximation of the MCB, we again define appropriate linearly independent vectors $S_1, \dots, S_\nu \in k^m$ and use Lemma 5.25. Consider the exact algorithm in Section 5.3 executing with the t -spanner G' as its input. Other than the cycles $C_{\lambda+1}, \dots, C_\nu$, the algorithm also returns the vectors $R_{\lambda+1}, \dots, R_\nu \in k^{m'}$ such that $\langle C_i, R_j \rangle = 0$ for $\lambda + 1 \leq i < j \leq \nu$ and C_i is a shortest cycle in G' such that $\langle C_i, R_i \rangle \neq 0$ for $\lambda + 1 \leq i \leq \nu$. Moreover, the $(\nu - \lambda) \times m'$ matrix whose j -th row is R_j is lower triangular with 1 in its diagonal. This implies that the R_j 's are linearly independent. Given any vector $S \in k^m$ let \tilde{S} be the projection of S onto its last m' coordinates. In other words, \tilde{S} is the restriction of S to the edge set of G' . We define S_j for $1 \leq j \leq \nu$ as follows. Let S_1, \dots, S_λ

be the first λ unit vectors of k^m . For $\lambda + 1 \leq j \leq \nu$ define S_j as:

$$S_j = (-\langle \tilde{C}_1, R_j \rangle, \dots, -\langle \tilde{C}_\lambda, R_j \rangle, R_{j,1}, R_{j,2}, \dots, R_{j,m'}),$$

where $R_{j,1}, \dots, R_{j,m'}$ are the coordinates of the vector $R_j \in k^{m'}$. Note that the vectors S_j for $1 \leq j \leq \nu$, defined above, are linearly independent. This is because the $\nu \times \nu$ matrix whose j -th row is S_j is lower triangular with nonzeros in its diagonal. The above definition of S_j 's is motivated by the property that for each $1 \leq i \leq \lambda$, we have $\langle C_i, S_j \rangle = -\langle \tilde{C}_i, R_j \rangle + \langle \tilde{C}_i, R_j \rangle = 0$, since the cycle C_i has 0 in all first λ coordinates, except the i -th coordinate, which is nonzero. Lemma 5.28, shown below, together with Lemma 5.25, implies the correctness of our approach.

Lemma 5.28. *Consider the above defined S_j for $1 \leq j \leq \nu$ and let D_j be a shortest cycle in G such that $\langle D_j, S_j \rangle \neq 0$. Cycle C_j returned by the algorithm in Figure 5 has weight at most t times the weight of D_j .*

Proof. This is obvious for $1 \leq j \leq \lambda$ since D_j is a shortest cycle in G which uses edge e_j and $C_j = e_j + p_j$, where p_j is a t -approximate shortest path between the endpoints of e_j . Consider now D_j for $\lambda + 1 \leq j \leq \nu$. If D_j uses any edge e_i for $1 \leq i \leq \lambda$ we replace it with the corresponding shortest path in the spanner. This is the same as saying consider the cycle $D_j - C_i$ instead of D_j . Let $D'_j = D_j - \sum_{1 \leq i \leq \lambda} (e_i \in D_j) C_i$ where $(e_i \in D_j)$ is 1 if $e_i \in D_j$ and 0 if $e_i \notin D_j$. Then

$$\langle D'_j, S_j \rangle = \langle D_j, S_j \rangle + \sum_{1 \leq i \leq \lambda} (e_i \in D_j) \langle C_i, S_j \rangle.$$

But recall that our definition of S_j ensures that $\langle C_i, S_j \rangle = 0$ for $1 \leq i \leq \lambda$. This implies that $\langle D'_j, S_j \rangle = \langle D_j, S_j \rangle \neq 0$. But D'_j by definition has 0 in the first λ coordinates and $\tilde{S}_j = R_j$, which in turn implies that

$$\langle \tilde{D}'_j, R_j \rangle = \langle \tilde{D}'_j, \tilde{S}_j \rangle = \langle D'_j, S_j \rangle \neq 0.$$

C_j is a shortest cycle in G' such that $\langle C_j, R_j \rangle \neq 0$. Thus, C_j has weight at most the weight of \tilde{D}'_j (which is the same cycle as D'_j), and by construction, D'_j has weight at most t times the weight of D_j . \square

Thus, we have shown that the cost of our approximate basis is at most t times the cost of an optimal basis. As a t -spanner we will again use a $(2k - 1)$ -spanner. The best time bound in order to compute a minimum undirected cycle basis is $O(m^2 n / \log n + mn^2)$ and since a $(2k - 1)$ -spanner has at most $O(n^{1+1/k})$ edges, the total running time becomes $O(n^{3+2/k} / \log n + n^{3+1/k})$.

Theorem 5.29 (Kavitha et al. (2007); Mehlhorn and Michail (2008)). *For any integer $k \geq 2$, Algorithm 5 computes a $(2k - 1)$ -approximate undirected cycle basis in time $O(n^{3+2/k} / \log n + n^{3+1/k})$.*

Directed graphs: Algorithm 5 readily extends to the directed case. For the spanner computation we view our directed graph G as undirected and we compute a $(2k - 1)$ -spanner G' . We then give to the edges of G' the orientation that they have in G .

As in the undirected case we return two sets of cycles. The first set is constructed as follows. For each edge $e_i \in E \setminus E'$ for $1 \leq i \leq \lambda$ we compute the cycle $e_i + p_i$ where p_i

is the shortest path in G' between the endpoints of e_i when G' is viewed as an undirected graph. Then, we traverse each such cycle in an arbitrary orientation and form our directed cycles based on the direction of the edges in G . The second set is simply the set of cycles of a directed MCB of G' . The resulting running time is $O(n^{4+3/k})$ deterministic and $O(n^{3+2/k})$ randomized, depending on the exact algorithm that we use.

Theorem 5.30 (Kavitha et al. (2007); Mehlhorn and Michail (2008)). *For any integer $k \geq 2$, a $(2k - 1)$ -approximate directed cycle basis can be computed in time $O(n^{4+3/k})$ by a deterministic and in time $O(n^{3+2/k})$ by a Monte Carlo algorithm.*

5.9.2 2-approximation

A direct consequence of the technique in Section 5.4 is that any reduction in size of the candidate collection \mathcal{H} would immediately imply better algorithmic bounds. In this section we show that a set of $O(m\sqrt{n \log n})$ cycles, which is a subset of \mathcal{H} , contains a 2-approximate minimum cycle basis. Again, the basis is extracted from the set by determining the least weight ν linearly independent cycles in it.

Definition 5.2. *For $v, x \in V$ and $S \subset V$, $\text{bunch}(v, S)$ consists of all vertices closer to v than to any vertex in S and $\text{cluster}(x, S)$ consists of all vertices v with $x \in \text{bunch}(v, S)$.*

Lemma 5.31 (Thorup and Zwick (2001b)). *Given a weighted graph $G = (V, E)$ and $0 < q < 1$, one can compute a set $S \subset V$ of size $O(nq \log n)$ in expected time $O(m/q \log n)$ such that $|\text{cluster}(x, S)| \leq 4/q$ for all $x \in V$.*

We take $q = 1/\sqrt{n \log n}$ and first compute, as given in Lemma 5.31, a set S of $O(\sqrt{n \log n})$ vertices. This takes expected time $O(m\sqrt{n \log n})$ and ensures that $\text{cluster}(v, S)$ has size $\sqrt{n \log n}$ for all $v \in V$. Also, $\text{bunch}(v, S)$ for all v can be computed in expected time $O(m/q)$ (Thorup and Zwick (2001a)), which is $O(m\sqrt{n \log n})$. We use two types of cycles:

- the $O(m\sqrt{n \log n})$ cycles $C[s, e]$ for all $s \in S$ and $e \in E$,
- the cycles $C[u, e]$ for each $u \in V$ and $e = (v, w) \in E$ and either v or w in $\text{bunch}(u, S)$. The number of such cycles is $\sum_{u \in V} \sum_{v \in \text{bunch}(u, S)} \deg(v)$. Rewriting this sum, we obtain $\sum_{v \in V} \deg(v) \cdot |\text{cluster}(v, S)|$, which in turn is at most $\sqrt{n \log n} \sum_{v \in V} \deg(v) = m\sqrt{n \log n}$.

Thus, our collection has $O(m\sqrt{n \log n})$ cycles. We need to show that it contains a 2-approximate cycle basis. Let B_1, \dots, B_ν be the minimum cycle basis of G determined by Horton's algorithm in order of non-decreasing weight, i.e., $w(B_1) \leq w(B_2) \leq \dots \leq w(B_\nu)$.

Lemma 5.32. *For all $1 \leq i \leq \nu$ we have $B_i = \sum_{C \in \mathcal{C}_i} C$ where \mathcal{C}_i is a subset of our collection and each cycle in \mathcal{C}_i has cost at most $2 \cdot w(B_i)$.*

Proof. Consider any B_i . If B_i belongs to our collection, we set $\mathcal{C}_i = \{B_i\}$. Otherwise, $B_i = C[u, e]$ where $e = (v, w)$ and neither v nor w is in $\text{bunch}(u, S)$. Let $s \in S$ be the nearest vertex in S to u . Then, $w(\text{SP}(s, u)) \leq w(\text{SP}(u, v))$ and $w(\text{SP}(s, u)) \leq w(\text{SP}(u, w))$.

For any edge $f \in B_i$, the cycle $C(s, f)$ is in our collection and $B_i = \sum_{f \in B_i} C(s, f)$ since the paths from s to the endpoints of the edges in B_i appear twice in this sum and cancel out. We set $\mathcal{C}_i = \{C(s, f) \mid f \in B_i\}$. It remains to show $w(C(s, f)) \leq 2w(B_i)$ for all $f \in B_i$. We distinguish cases.

Assume first that $f \neq e$. Then $f \in \text{SP}(u, v)$ or $f \in \text{SP}(u, w)$. We may assume w.l.o.g. that the former is the case. Then $w(C(s, f)) \leq w(\text{SP}(s, u)) + w(\text{SP}(u, v)) + w(\text{SP}(v, s))$ since $C(s, f)$ consists of f and the shortest paths from s to the endpoints of f and $w(\text{SP}(v, s)) \leq w(\text{SP}(s, u)) + w(\text{SP}(u, v))$ by the triangle inequality. Thus $w(C(s, f)) \leq 2(w(\text{SP}(s, u)) + w(\text{SP}(u, v))) \leq 2w(B_i)$ since $w(\text{SP}(s, u)) \leq w(\text{SP}(u, w))$.

Assume next that $f = e$. Then

$$\begin{aligned} w(C(s, f)) &= w(\text{SP}(s, v)) + c(e) + w(\text{SP}(w, s)) \\ &\leq w(\text{SP}(s, u)) + w(\text{SP}(u, v)) + c(e) + w(\text{SP}(s, u)) + w(\text{SP}(u, w)) \\ &\leq 2w(\text{SP}(u, v)) + c(e) + 2w(\text{SP}(u, w)) \\ &\leq 2w(B_i). \end{aligned}$$

□

Lemma 5.33. *The collection defined above contains ν linearly independent cycles A_1, \dots, A_ν with $w(A_i) \leq 2 \cdot w(B_i)$ for $i = 1, \dots, \nu$.*

Proof. The lemma follows from Lemma 5.32. Assume otherwise and let j be minimal such that $\cup_{i \leq j} \mathcal{C}_i$ contains less than j linearly independent vectors with $w(A_i) \leq 2 \cdot w(B_i)$ for $i = 1, \dots, j$. Then $j \geq 1$ and $\cup_{i \leq j-1} \mathcal{C}_i$ contains at least $j-1$ linearly independent vectors with $w(A_i) \leq 2 \cdot w(B_i)$ for $i = 1, \dots, j-1$. Also, $\cup_{i \leq j} \mathcal{C}_i$ spans $\{B_1, \dots, B_j\}$ and hence contains at least j linearly independent vectors. Thus, it contains a vector A_j linearly independent of $\{A_1, \dots, A_{j-1}\}$. Furthermore, $A_j \in \mathcal{C}_i$ for some $i \leq j$ and hence $w(A_j) \leq 2w(B_i) \leq 2w(B_j)$, a contradiction. □

It is now straightforward to extract the 2-approximate MCB using the techniques that we have discussed so far. This can be done in expected time $O(m^2 \sqrt{n \log n} + mn^2 + m^\omega)$. Similarly for directed graphs, we get the same running time if we do computation module some small random prime number. Moreover, in the undirected case we can achieve an $O(m^2 \sqrt{n/\log n} + mn^2 + m^\omega)$ running time by using the bit packing technique in Mehlhorn and Michail (2008).

Theorem 5.34 (Kavitha et al. (2007); Mehlhorn and Michail (2008)). *A 2-approximate undirected basis can be constructed in expected time $O(m^2 \sqrt{n/\log n} + mn^2 + m^\omega)$ and a 2-approximate directed basis can be constructed in expected time $O(m^2 \sqrt{n \log n} + n^2 m + m^\omega)$ with high probability.*

5.10 Algorithm Engineering

Exact and also the approximate algorithms for minimum cycle bases have a fairly large worst-case running time. In this section, we discuss heuristic improvements and algorithm engineering issues. The hope is that heuristics and algorithm engineering techniques improve upon the worst-case running time in many cases. We restrict attention to computing minimum undirected bases. Implementations of cycle basis algorithms are described in Gleiss (2001a); Huber (2003); Berger (2004); Bauer (2004); Mehlhorn and Michail (2006).

The first decision to be made is to choose between the two main approaches. Horton's approach first computes $O(nm)$ cycles and then uses Gaussian elimination to find an optimum basis. No heuristics are known that improve upon the worst-case in many cases. The situation

is different for the algebraic approach of Algorithm 1 where in each phase first a support vector and then a cycle is computed.

How should we represent cycles and support vectors, as sparse or as dense vectors? There are two arguments in favor of a sparse representation. The theoretical argument is that we know of the existence of bases of weight $O(W \log n)$; in such a basis, we expect most circuits to have $o(n)$ edges. The engineering argument is that a dense representation immediately introduces an $\Omega(m^2)$ lower bound; we are constructing m vectors of length m . Thus, the sparse representation is preferred.

The next major question is how to compute each cycle. In Algorithm 1 each cycle is computed after a support vector is found. However, there are two possible ways for doing this: (a) use the candidate set \mathcal{H} or some other collection that contains an minimum cycle basis and the labelled trees representation, or (b) use the signed graph approach.

Although the labelled trees approach is faster for sparse graphs by a logarithmic factor as well as for dense graphs in case the extra technique of bit-packing from Mehlhorn and Michail (2008) is used, it has a major practical drawback which needs to be addressed. It introduces a lower bound on the best case of the algorithm. The labelled trees approach maintains n shortest path trees. In each of the ν phases of the algorithm, each of these shortest path trees is traversed, in order to update the labels based on the current support vector. Thus, the technique introduces an $\Omega(mn^2)$ lower bound. For this reason we believe that the signed graph approach is better.

The signed graph approach constructs a graph $G_i(S_i)$ where S_i is the support vector during phase i of the algorithm. In this graph it executes n single source shortest path computations. There are however some heuristics that can be used to reduce the number of such computations. During phase i we might perform up to n shortest path computations in order to compute the shortest cycle C_i with an odd intersection with the vector S_i . We can use the shortest path found so far as an upper bound on the shortest path. This is implemented as follows; a node is only added in the priority queue of Dijkstra's implementation if its correct upper distance is not more than our current upper bound.

We come to the most important heuristic. In each of the ν phases we are performing n shortest path computations. This results to $\Omega(mn)$ shortest path computations. Let $S = \{e_1, e_2, \dots, e_k\}$ be a support vector at some point of the execution. We need to compute the shortest cycle C such that $\langle C, S \rangle = 1$. We can reduce the number of shortest path computations based on the following observation.

Let $C_{\geq i}$ be the shortest cycle in G such that $\langle C_{\geq i}, S \rangle = 1$, $C_{\geq i} \cap \{e_1, \dots, e_{i-1}\} = \emptyset$, and $e_i \in C_{\geq i}$. Then

$$C = \min_{i=1, \dots, k} C_{\geq i}.$$

We can compute $C_{\geq i}$ in the following way. We delete edges $\{e_1, \dots, e_i\}$ from G and the corresponding edges from the signed graph G_i . Let $e_i = (v, u) \in G$. Then we compute a shortest path in G_i from v^+ to u^+ . The path computed will have an even number of edges from the set S , and together with e_i an odd number. Since we deleted edges $\{e_1, \dots, e_i\}$, the resulting cycle does not contain any edges from $\{e_1, \dots, e_{i-1}\}$.

Using the above observation we can compute each cycle in time $O(k \cdot SP(n, m))$ when $|S| = k < n$ and in time $O(n \cdot SP(n, m))$ when $|S| \geq n$. Here $SP(n, m)$ is the time of a single-source shortest path computation in a graph with n nodes and m edges. In this way,

the total cost of computing the basic circuits becomes

$$SP(n, m) \cdot \sum_{i=1, \dots, \nu} \min(n, |S_i|).$$

Another issue that needs to be discussed is the use of fast matrix multiplication when computing the support vectors. Experiments in Mehlhorn and Michail (2006) with random graphs suggest, that the use of fast matrix multiplication is not necessary even for medium to large instances. The reason is that the cycles computation part of the algorithm dominates the running time, although in theory it is the other way around. The reason is that support vectors are typically sparse. Thus, the technique of Algorithm 2 which has a worst case bound of $O(m^3)$, is sufficient. Moreover, due to its simplicity, it is very easy to be implemented efficiently.

Moving to the approximation algorithms of Section 5.9, we note that they improve significantly the running times. This is not only a theoretical observation but is true in practice as well. Algorithm 5 reduces the computation of an approximate MCB to the computation of: (a) a spanner of the input graph, and (b) the MCB of a sparse graph (the previously computed spanner). Depending on the properties of the graph and how sparse spanners exist, this approximation algorithm is much faster than any exact algorithm. Moreover, the approximation algorithms do not really require fast matrix multiplication. Algorithm 5 requires $O(n^{3+2/k}/\log n + n^{3+1/k})$ time in order to compute a $(2k-1)$ -approximate MCB. If we do not use fast matrix multiplication the running time increases to $O(n^{3+3/k})$. We conclude that Algorithm 5 where the support vectors are maintained as in Algorithm 2 and the cycles are computed using the signed graph approach of Section 5.4 will be an effective way for computing approximate minimum cycle bases.

5.11 Relevant Cycles

In general, minimum cycle bases are not unique. In some applications, e.g. in chemistry (Gleiss (2001b)) it is useful to know *all* minimum cycle bases. A cycle that belong to some minimum cycle basis are called *relevant*. As their number could be exponential, the goal is to compute a set of prototype cycles from which all relevant cycles can then be derived easily. Vismara (1997) presented an algorithm that in a similar fashion to Horton's algorithm extracts these prototypes from a polynomially sized set of candidate cycles using Gaussian elimination. Vismara's algorithm runs in time $O(m^4)$. From these prototypes, relevant cycles can be computed in $O(n|\mathcal{C}_R|)$ time where \mathcal{C}_R denotes the set of relevant cycles.

6 Hardness Results

We will show that the minimization problems for strictly and weakly fundamental cycle bases are \mathcal{APX} -hard. A minimization problem belongs to class \mathcal{APX} if it has a $(1+\varepsilon)$ approximation algorithm for some $\varepsilon > 0$. A minimization problem is \mathcal{APX} -hard if any problem in \mathcal{APX} can be reduced to it by an L -reduction; \mathcal{APX} -hard problems do not have polynomial time approximation schemes unless $\mathcal{P} = \mathcal{NP}$ (Papadimitriou (1994); Trevisan (2004)).

An L -reduction from a optimization problem P_1 to a optimization problem P_2 consists of two polynomially computable functions t_1 and t_2 with the following properties:

- (i) t_1 maps instances of P_1 to instances of P_2 such that

$$\text{opt}_{P_2}(t_1(I)) \leq \beta_1 \text{opt}_{P_1}(I)$$

for any instance I of P_1 . Here $\text{opt}_{P_i}(X)$ denotes the optimum value for instance X of problem P_i and β_1 is some constant.

- (ii) t_2 associates to any instance I of P_1 and any feasible solution S' of the corresponding instance $I' := t_1(I)$ of P_2 a feasible solution $S := t_2(I, S')$ of I such that

$$|\text{opt}_{P_1}(I) - \text{val}_{P_1}(I, S)| \leq \beta_2 |\text{val}_{P_2}(I', S') - \text{opt}_{P_2}(I')|$$

Here, $\text{val}_{P_i}(X, Y)$ denotes the objective function value of the feasible solution Y for instance X of problem P_i and β_2 is some constant.

L -reductions preserve approximability. If S' is an ε -approximation to the optimum solution of I' , i.e., $|\text{opt}_{P_2}(I') - \text{val}_{P_2}(I', S')| \leq \varepsilon \cdot \text{opt}_{P_2}(I')$, then $|\text{opt}_{P_1}(I) - \text{val}_{P_1}(I, S)| \leq \beta_2 \cdot \varepsilon \cdot \text{opt}_{P_2}(I') \leq \beta_1 \cdot \beta_2 \cdot \varepsilon \cdot \text{opt}_{P_1}(I)$, i.e., S is a $\beta_1 \beta_2 \varepsilon$ -approximation to the optimum solution of I .

6.1 Strictly Fundamental Cycle Bases

Recall that a strictly fundamental cycle basis consists of the fundamental circuits with respect to some spanning tree. We saw in Theorem 4.11 that any graph has a strictly fundamental basis of weight $O(W \log^2 n \log \log n)$ and of length $O(n^2)$. Deo et al. (1982) showed the \mathcal{NP} -hardness of the minimum strictly fundamental cycle basis problem. We will now sketch a proof for its \mathcal{APX} -hardness (Galbiati et al. (2007)). The proof consists of an L -reduction from the following special case of the maximum satisfiability problem.

Max-3SAT-NAE-UN-9: Given a set $X = \{x_1, \dots, x_n\}$ of Boolean variables and a collection $\mathcal{C} = \{C_1, \dots, C_m\}$ of disjunctive clauses with exactly 3 variables per clause, where all variables appear unnegated and each variable occurs in at most 9 clauses, find a truth assignment to the variables maximizing the number of clauses containing both a *true* variable and a *false* variable.

MAX-3SAT-NAE-UN-9 is a Not All Equal version of MAX-3SAT restricted to instances with UNnegated variables, each variable having at most 9 occurrences. In Galbiati et al. (2007), MAX-3SAT-NAE-UN-9 was shown to be APX-hard by means of a sequence of standard L -reductions starting from MAX CUT-3, the problem of finding a cut containing the maximum number of edges in an undirected graph where all vertices have degree at most 3. In turn, MAX CUT-3 has been shown to be APX-hard in Alimonti and Kann (2000).

The Main Reduction. We now describe the L-reduction from MAX-3-SAT-NAE-UN-9 to MSFCB. Let q and M be integer constants. We will fix q and M later. Let I be an instance of MAX-3-SAT-NAE-UN-9 with variable set $X = \{x_1, \dots, x_n\}$ and clause collection $\mathcal{C} = \{C_1, \dots, C_m\}$. We construct an instance I' of MSFCB, i.e. a weighted graph G_I , as follows. The set of vertices is

$$V(G_I) = \{r, c_1^A, c_2^A, \dots, c_m^A, c_1^B, c_2^B, \dots, c_m^B, x_1^A, x_2^A, \dots, x_n^A, x_1^B, x_2^B, \dots, x_n^B\},$$

and the edges in $E(G_I)$ together with the corresponding weights are:

- for each $i = 1, 2, \dots, n$, we have one edge $\{r, x_i^A\}$ and one edge $\{r, x_i^B\}$ both of weight 1;
- for each $i = 1, 2, \dots, n$, we have $2q + 1$ parallel edges connecting vertices x_i^A and x_i^B , and all of them have weight 1;
- for each $j = 1, 2, \dots, m$ and for each variable x_i occurring in C_j , we have the edges $\{c_j^A, x_i^A\}$ and $\{c_j^B, x_i^B\}$ of weight M .

We remark that edges of weight M may be replaced by a path of length M and parallel edges may be split by an additional intermediate vertex. In other words, the graph G_I could also be constructed as an unweighted simple graph. The following easy-direction lemma indicates the intention behind the reduction.

Lemma 6.1. *Let $\Phi : \{x_1, x_2, \dots, x_n\} \mapsto \{true, false\}^n$ be a truth assignment such that there are t clauses in \mathcal{C} containing both a variable with value true and a variable with value false. Then G_I has a fundamental cycle basis of weight $n(4q + 3) + m(8M + 12) - t$.*

Proof. By relabeling the clauses, we can assume w.l.o.g. that, for $j = 1, 2, \dots, t$, clause C_j contains a variable with value true as well as a variable with value false. So, for $t < j \leq m$, all variables in C_j have the same truth value under Φ . Construct a spanning tree T as follows, see Figure 25:

- for each $i = 1, 2, \dots, n$, include in T a single parallel edge connecting x_i^A and x_i^B ;
- for each $i = 1, 2, \dots, n$, if $\Phi(x_i) = true$ then include in T the edge $\{r, x_i^A\}$; otherwise, if $\Phi(x_i) = false$, include in T the edge $\{r, x_i^B\}$;
- for each clause C_j , with $1 \leq j \leq t$, select one variable \bar{x} with $\Phi(\bar{x}) = true$ and one variable \tilde{x} with $\Phi(\tilde{x}) = false$, and include the edges $\{c_j^A, \bar{x}^A\}$ and $\{c_j^B, \tilde{x}^B\}$;
- for each clause C_j , with $j = t + 1, \dots, m$, select a single arbitrary variable x_i occurring in C_j and include in T both edges $\{c_j^A, x_i^A\}$ and $\{c_j^B, x_i^B\}$.

We next compute the costs of the fundamental cycles induced by the cotree edges. We distinguish the following cases.

- For each $i = 1, 2, \dots, n$, T contains exactly one of the two edges $\{r, x_i^A\}$ and $\{r, x_i^B\}$. The other edge induces a fundamental cycle of cost 3, for a total of $3n$.
- For each $i = 1, 2, \dots, n$, $2q$ edges connecting vertices x_i^A and x_i^B are not in T . Each of them induces a fundamental cycle of cost 2, for a total cost of $4qn$.

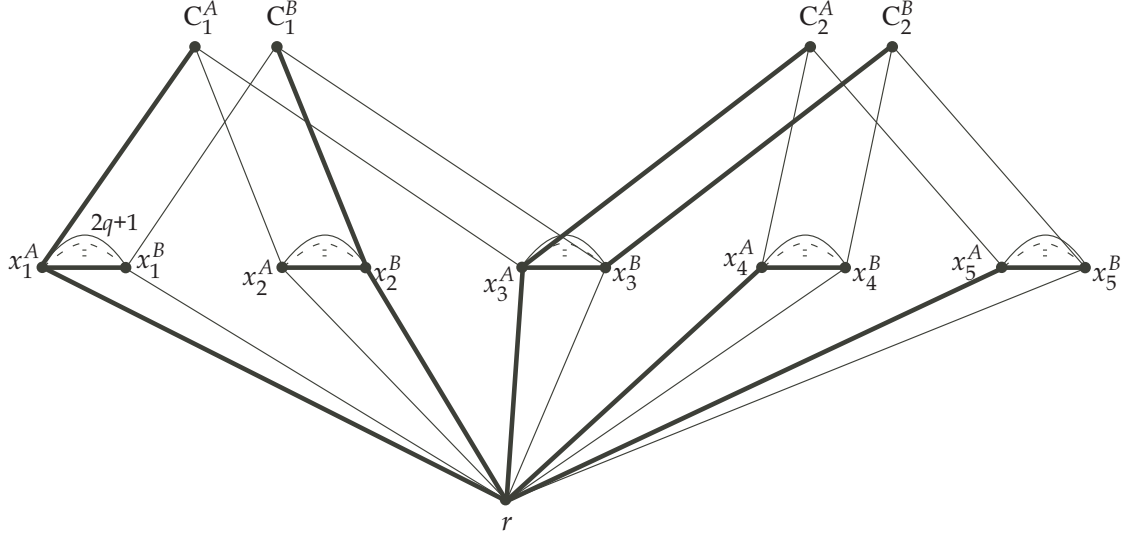


Figure 25: The graph G_I associated to the MAX-3-SAT-NAE-UN-9 instance I with variable set $X = \{x_1, x_2, x_3, x_4, x_5\}$ and clause collection $\mathcal{C} = \{x_1 \vee x_2 \vee x_3, x_3 \vee x_4 \vee x_5\}$. The spanning tree T of G_I derived from the truth assignment $\Phi = (true, false, true, true, true)$ is shown in bold.

- For each $j = 1, 2, \dots, t$, the four co-tree edges incident in c_j^A or c_j^B induce four cycles. Exactly one of these cycles has cost $2M + 2$, while the others have cost $2M + 3$. The corresponding costs sum to $t(8M + 11)$.
- For each $j = t + 1, \dots, m$, each one of the two co-tree edges incident in c_j^A induces a cycle of cost $2M + 2$ (respectively $2M + 4$) if all variables in C_j are *true* (respectively *false*). Analogously, each one of the two co-tree edges incident in c_j^B induces a cycle of cost $2M + 4$ (respectively of cost $2M + 2$) if all variables in C_j are *true* (respectively *false*). These costs sum to $(m - t)(8M + 12)$.

Therefore the fundamental cycle basis induced by T has a cost of

$$3n + 4qn + t(8M + 11) + (m - t)(8M + 12) = n(4q + 3) + m(8M + 12) - t.$$

□

A key property of the reduction is that the type of spanning tree considered in the lemma above, gives rise to a minimum strictly fundamental basis.

Definition 6.1. A spanning tree T of G_I is well-behaved if it satisfies the following properties:

1. for each $j = 1, 2, \dots, m$, the vertices c_j^A and c_j^B have degree 1 in T ,
2. for each $i = 1, 2, \dots, n$, exactly one edge of the $2q + 1$ edges $\{x_i^A, x_i^B\}$ belongs to T ,
3. for each $j = 1, 2, \dots, m$, either for some i both edges $\{c_j^A, x_i^A\}$ and $\{c_j^B, x_i^B\}$ belong to T or for some i_1 and i_2 , with $1 \leq i_1 \leq n$, $1 \leq i_2 \leq n$ and $i_1 \neq i_2$, both edges $\{c_j^A, x_{i_1}^A\}$ and $\{x_{i_1}^A, r\}$ as well as both edges $\{c_j^B, x_{i_2}^B\}$ and $\{x_{i_2}^B, r\}$ belong to T .

Lemma 6.2 (Galbiati et al. (2007)). *Assume $q \geq 9$ and $M \geq 4$. For any spanning tree T of graph G_I , we can in polynomial time derive from T a well-behaved spanning tree T' such that the weight of the basis induced by T' is no larger than the weight of the basis induced by T .*

We can now state the main result of the section.

Theorem 6.3 (Galbiati et al. (2007)). *The minimum strictly fundamental cycle basis problem is APX-hard.*

Proof. It suffices to verify that the reduction presented is an L-reduction. For any instance $I = (X, C)$ of MAX-3-SAT-NAE-UN-9, the corresponding instance I' of MSFCB can obviously be constructed in polynomial time.

The simple randomized argument implying that any MAX-SAT instance with m clauses admits a truth assignment satisfying at least $m/2$ clauses, is also valid for MAX-3-SAT-NAE-UN-9. Thus $\text{opt}(I) \geq m/2$.

According to Lemma 6.1, where $q = 9$, we have $\text{opt}(I') \leq n(4q + 3) + m(8M + 12)$. Since we may assume that $n \leq 3m$ (otherwise some variable would occur in no clause), $\text{opt}(I') \leq 3m(4q + 3) + m(8M + 12) = m(12q + 8M + 21) \leq m(24q + 16M + 42)\text{opt}(I)$. We set $\beta_1 = 24q + 16M + 42$.

By Lemma 6.2, from any spanning tree T of G_I we can derive a well-behaved spanning tree T' without increasing the weight of the associated fundamental cycle basis. Now the three properties characterizing well-behaved spanning trees make sure that it is possible to reverse the construction described in the proof of Lemma 6.1. Therefore, to derive a truth assignment Φ from any spanning tree T of G_I , it suffices first to derive a well-behaved spanning tree T' from T and then to set $\Phi(x_i) = \text{true}$ when $\{x_i^A, r\} \in T'$, and $\Phi(x_i) = \text{false}$ when $\{x_i^B, r\} \in T'$. Condition (ii) of an L-reduction is then satisfied with $\beta_2 = 1$. \square

We close this section with some open problems.

Open Problem 10. *Is there an $O(\log n)$ approximation algorithm for minimum F-bases? Is there one for planar graphs? The approximability of the bottleneck version, in which one looks for a strictly fundamental cycle basis where the weight of the maximum cycle is minimum, has been addressed in Galbiati (2003).*

Open Problem 11. *Is the minimum F-basis problem in \mathcal{APX} ?*

Open Problem 12. *What is the complexity of the minimum F-basis problem for planar graphs? The related problem of computing a spanning tree with shortest fundamental circuit is NP-complete for planar graphs (Fekete and Kremer (2001)).*

6.2 Weakly Fundamental Cycle Basis

We know from Thm 4.4 that any graph has a weakly fundamental cycle basis (W-basis) of weight $O(W \log n)$. Thus the weight of a minimum W-basis can be approximated within a factor of $O(\log n)$; no better approximation factor is known. Rizzi (2007) has shown that the minimum W-basis problem is \mathcal{APX} -hard and we will sketch his proof in this section.

Open Problem 13. *Is the minimum W-basis problem in \mathcal{APX} ?*

We first introduce a compact way for representing W -bases. For T a spanning tree of G , any ordering e_1, e_2, \dots, e_ν of the cotree edges is called a *removal sequence*.

Let \mathcal{C} be a W -basis of a connected graph G . If $\mathcal{C} \neq \emptyset$, that is, if G is not a tree, then there exists an edge e of G that is contained in precisely one circuit of \mathcal{C} . Let C_e be the only circuit in \mathcal{C} that contains e . Notice that $G \setminus e$ is connected and $\mathcal{C} \setminus C_e$ is a W -basis of $G \setminus e$. If this process is iterated over $G \setminus e$, we finally end up with a spanning tree T of G . Furthermore, if we label as e_i the i -th edge that has been removed in the process, then the sequence $s = e_1, e_2, \dots, e_\nu$ is a removal sequence having the edges in T as tree edges and certifying that \mathcal{C} is a W -basis according to Definition 3.1. We say that the spanning tree T , the ordering e_1, e_2, \dots, e_ν , and the W -basis \mathcal{C} from which we started, are *compatible*. Notice that at any iteration of the edge removal process it may be possible that more than one edge of the current graph is contained in precisely one circuit of \mathcal{C} . Actually, there is always some freedom of choice when removing the last edge. Thus, in general, a W -basis of G may be compatible with more than one spanning tree of G . Furthermore, even w.r.t. any particular tree T of G , a W -basis may be compatible with more than one ordering of the edges of $G \setminus T$. On the converse, any removal sequence e_1, e_2, \dots, e_ν of G , might be compatible with more than one W -basis of G . However, among the W -basis of G which are compatible with the ordering e_1, e_2, \dots, e_ν , we can efficiently find one of minimum cost by resorting to any shortest path algorithm as a subroutine. And we can also enforce this W -basis to be unique by possibly adopting a lexicographic scheme to resolve ties among circuits of the same weight. Indeed, given any removal sequence e_1, e_2, \dots, e_ν of G , the unique W -basis of G associated to the sequence e_1, e_2, \dots, e_ν is obtained as described in Algorithm W-DECODER.

Algorithm W-DECODER(e_1, e_2, \dots, e_ν)

start with $G' := G$, $\mathcal{C} = \emptyset$, and,

for $i = 1, 2, \dots, \nu$, do,

add to \mathcal{C} the *unique*¹⁰ cheapest circuit of G' containing edge e_i ;

remove edge e_i from G' .

return \mathcal{C} .

Notice that not every W -basis of G admits a removal sequence encoding it. We say that a W -basis \mathcal{C} of G is *locally-optimal* if there exists a removal sequence s of G such that the execution of Algorithm W-DECODER(s) produces \mathcal{C} . Indeed, the above remarks inspire a natural local search approach for the minimum W -basis problem, where, given any WFCB \mathcal{C} of G , we first obtain a removal sequence s compatible with \mathcal{C} and then substitute \mathcal{C} with the W -basis \mathcal{C}' produced by Algorithm W-DECODER(s). Notice that \mathcal{C}' is locally-optimal and its cost never exceeds the cost of \mathcal{C} .

A Fundamental Gadget: Our \mathcal{APX} -hardness proof is based on a single gadget. The gadget is derived from an graph first described by Liebchen and Rizzi (2007); every W -basis of this graph is strictly more expensive than the cheapest undirected cycle basis. Indeed, since the minimum U -basis problem is in \mathcal{P} , the graphs produced by a reduction from a generic \mathcal{APX} -hard optimization problem to the minimum W -basis problem are bound to involve such graphs.

Although, the inapproximability result also holds in the unweighted case, we find it convenient to allow the use of small natural weights (actually, all in $\{1, 2, 3\}$) in the constructions and in the gadgets to follow. Clearly, an edge of weight w may be replaced by a path of

¹⁰uniqueness is enforced by the adoption of a lexicographic scheme.

w edges and $w - 1$ new intermediate nodes without changing the essence of the cycle basis problem. The transformation is polynomial as long as the weights are polynomially bounded. So there is no harm in using small integer weights.

We start by describing a graph for which no minimum U-basis is weakly fundamental. Consider first the graph V_8 in Fig. 26. Here, $m = 12$, $n = 8$, $\nu = 5$ and the circuits $C_1 = 2 - 3 - 7 - 6$, $C_2 = 3 - 4 - 0 - 7$, $C_3 = 4 - 5 - 1 - 0$, $C_4 = 5 - 6 - 2 - 1$, $C_5 = 6 - 5 - 4 - 3 - 2$ form a W-basis of V_8 . Indeed, C_1 (resp. C_2, C_3, C_4) is the only circuit to contain the edge $\{6, 7\}$ (resp. $\{0, 7\}, \{0, 1\}, \{1, 2\}$). However, for later convenience, we also certify the independence of C_1, C_2, C_3, C_4, C_5 by giving five sets Σ_1 to Σ_5 such that $\langle \Sigma_i, C_j \rangle = 1$ if and only if $i = j$.

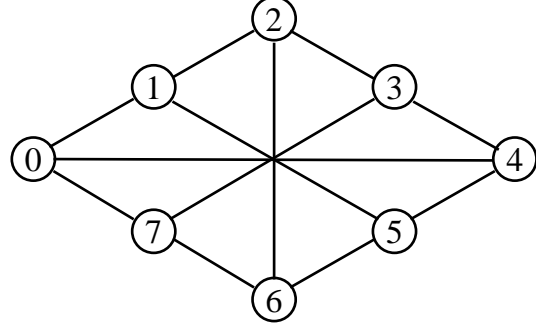


Figure 26: Wagner's graph V_8 .

$\Sigma_1 = \{\{2, 6\}, \{5, 6\}\}$ with odd intersection only with C_1 (see Fig. 27 on the left);

$\Sigma_2 = \{\{2, 6\}, \{5, 6\}, \{3, 7\}\}$ with odd intersection only with C_2 (see Fig. 27 on the right);

$\Sigma_3 = \{\{1, 5\}, \{2, 3\}, \{2, 6\}\}$ with odd intersection only with C_3 (mirror image of Σ_2);

$\Sigma_4 = \{\{2, 3\}, \{2, 6\}\}$ with odd intersection only with C_4 (mirror image of Σ_1);

$\Sigma_5 = \{\{2, 3\}, \{2, 6\}, \{5, 6\}\}$ with odd intersection only with C_5 (see Fig. 27 in the middle).

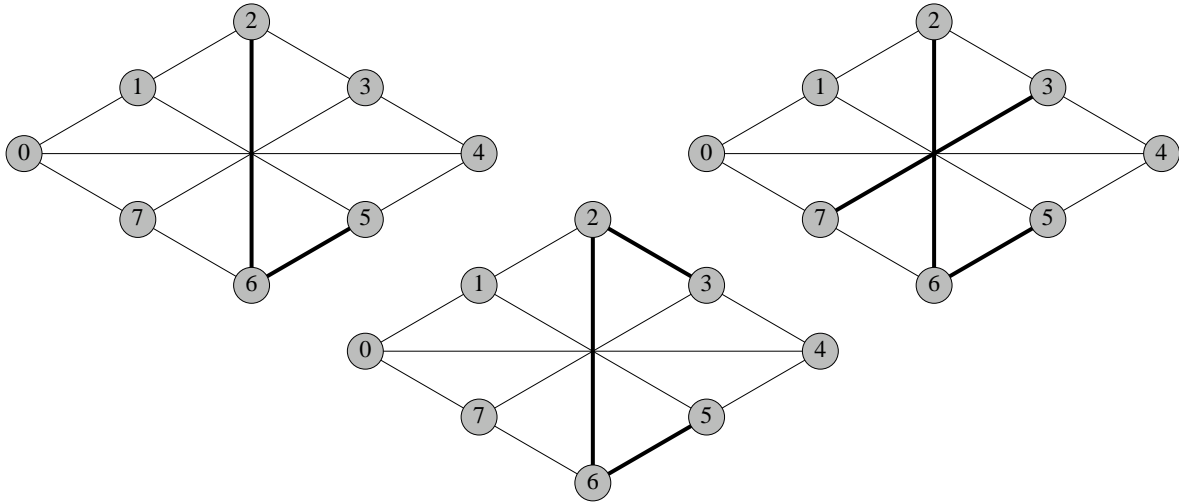


Figure 27: Certificates of independence for C_1 (left), C_2 (right), and C_5 (middle).

We call *petal* the weighted graph obtained from V_8 when edges $\{6, 7\}, \{7, 0\}, \{0, 1\}, \{1, 2\}$, also called *glue edges*, are assigned cost 2, and all other edges, called *internal edges*, receive cost 1. Notice that all circuits of the petal have cost at least 5, and C_1, C_2, C_3, C_4, C_5 are actually its only 5 circuits of cost precisely 5, whence they form the unique minimum basis

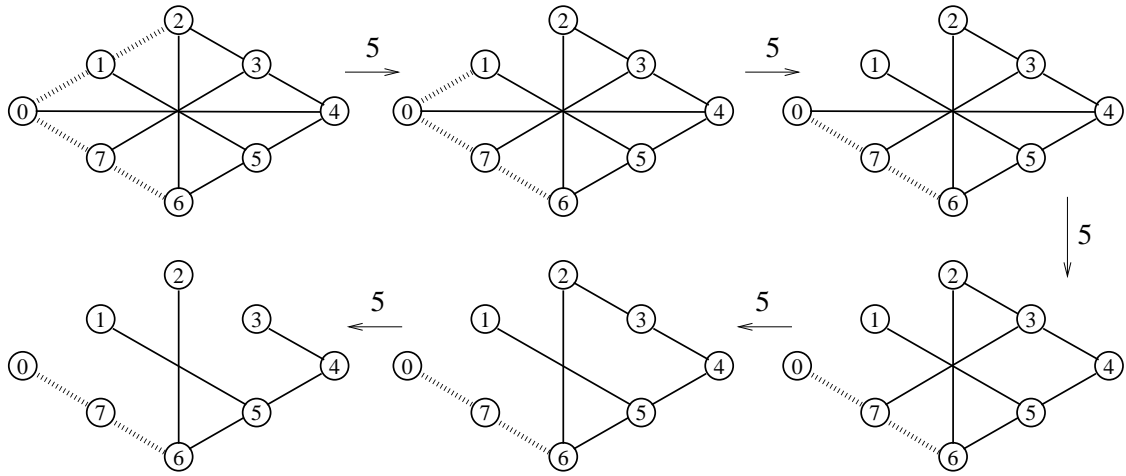


Figure 28: A removal sequence of a W-basis of weight 25.

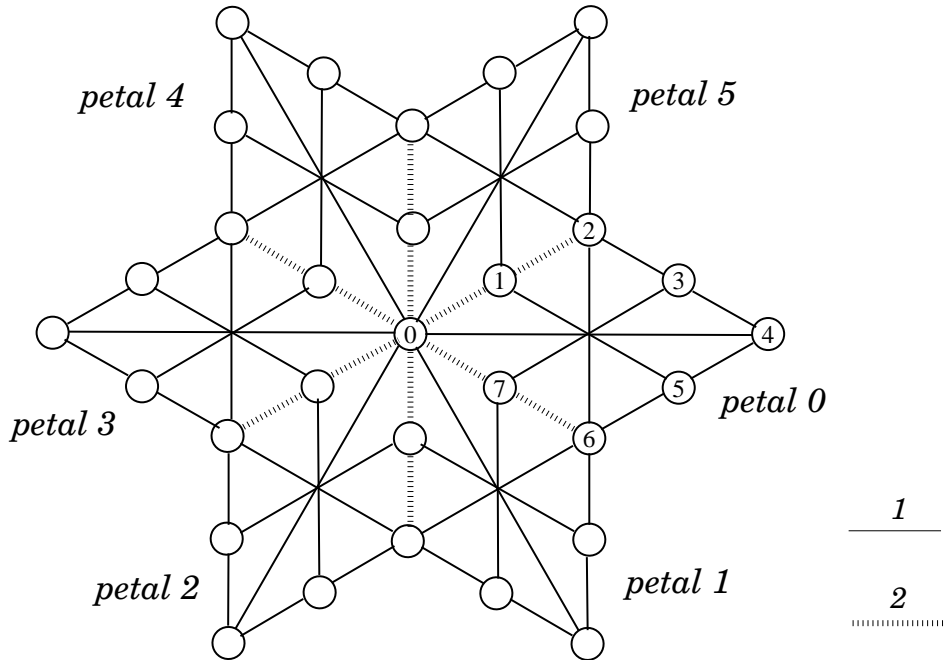


Figure 29: A weighted graph F whose unique minimum U-basis is not weakly fundamental.

of the petal. A removal sequence compatible with this locally-optimal W-basis of the petal is illustrated in Fig. 28.

Consider now the weighted graph F obtained by gluing together 6 distinct petals as shown in Fig. 29. Stated in words - let petal i , $0 \leq i \leq 5$, be on nodes $\{v_0^i, v_1^i, \dots, v_7^i\}$; we glue the petals through the following node identifications: $v_0^0 \leftrightarrow v_0^1 \leftrightarrow v_0^2 \leftrightarrow v_0^3 \leftrightarrow v_0^4 \leftrightarrow v_0^5$, and $v_1^{i+1 \bmod 6} \leftrightarrow v_7^i$ and $v_2^{i+1 \bmod 6} \leftrightarrow v_6^i$ for $0 \leq i \leq 5$. Edges that become parallel through this identification process are replaced by a single edge of weight 2.

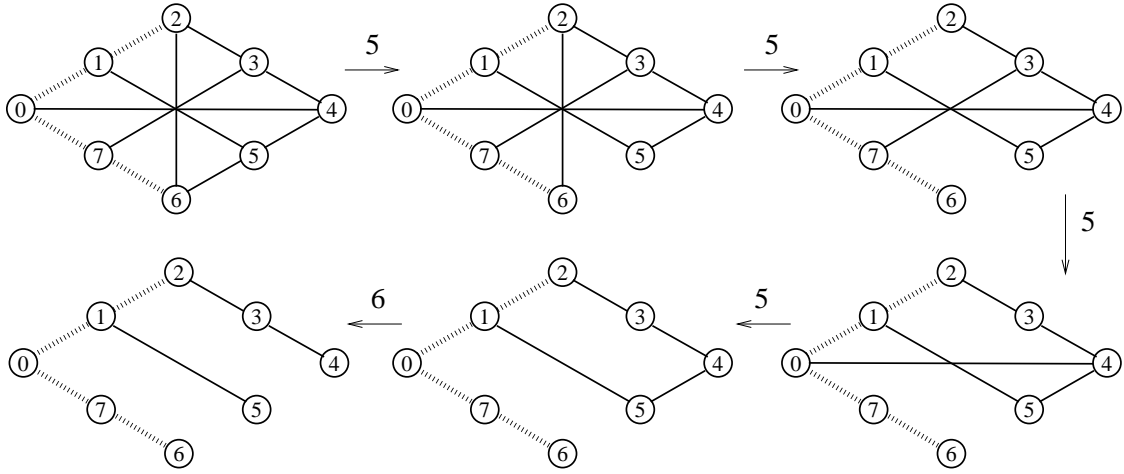


Figure 30: The removal sequence of a W-basis of weight 26. Notice that all glue edges are tree edges.

Clearly, $n_F = 31$, $m_F = 60$, and $\nu_F = 30$. The six copies of the circuits C_1, C_2, C_3, C_4, C_5 form a collection of 30 circuits in F whose independence can be established by taking the 6 corresponding copies of each of the odd sets $\Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5$. Hence, these 30 circuits form a U-basis of F . Each of these 30 circuits has weight 5. Every other circuit of F has cost at least 6; hence these 30 circuits form the unique minimum U-basis of F . This cycle basis is not weakly fundamental since each edge of F is contained in at least 2 of these circuits.

It is also relevant to our discussion to exhibit a cheap W-basis of F . In fact, F has a W-basis whose weight is only one larger than the weight of the unique minimum W-basis introduced above. Indeed, consider first one single petal of F and its W-basis as encoded by the removal sequence displayed in Fig. 30. This W-basis has weight 26 and leaves all glue edges as tree edges. It is easy to extend this removal sequence for one petal, say petal 0, to a removal sequence for F that encodes a W-basis of weight 151: simply append to it, for each one of the other 5 petals taken in clockwise order, a removal sequence like the one in Fig. 28 (see in the proof of Fact 6.1 for the details).

Fact 6.1. *There are precisely 30 circuits of cost 5 in F . These 30 circuits constitute the unique minimum U-basis of F . This basis has weight 150 and is not weakly fundamental. Furthermore, F admits a W-basis of weight 151.*

Proof. The W-basis associated by Algorithm WFCB-DECODER to the following removal sequence has weight 151: first, within petal 0, remove $v_5^0 v_6^0, v_2^0 v_6^0, v_3^0 v_7^0, v_0^0 v_4^0, v_4^0 v_5^0$ in this order, next, for $i = 1, 2, 3, 4, 5$ and in sequence, within petal i remove $v_1^i v_2^i, v_0^i v_1^i, v_0^i v_4^i, v_3^i v_7^i, v_2^i v_3^i$ in this order. \square

We are now ready to produce a weighted graph G (the gadget) with the following properties. (1) G contains 4 nodes x, y, z and w , and the edges wx, wy and wz , all of weight 1; (2) The minimum U-basis of G has weight B and G has a W-basis of weight B . (3) Let T be any spanning tree of G compatible with some W-basis of weight B . Then, the distance in T between any two of the 4 nodes x, y, z and w is at least 3; (4) G has a W-basis of weight $B + 1$, for which $wx, wy, wz \in T$.

The intended functioning of the gadget is as follows. Several copies of the gadget will be part of the graph G_H representing the minimum W-basis problem instance constructed by the L-reduction eventually proposed in Section 6.2. Each such gadget (copy) is attached to the rest of G_H by means of the 4 nodes x, y, z and w , and, actually, occurs as an induced subgraph of G_H . Each removal sequence for G_H contains, in a natural way, a removal sequence for each one of these gadgets. Indeed, a set of edges whose removal makes G_H acyclic also intersects all circuits of any given subgraph of G_H . Notice that, after the removal of the sole edges prescribed by a removal sequence for a gadget, the nodes of that gadget are still connected within the gadget. In order to intersect all circuits of G_H , a removal sequence for G_H will need to include further edges. When obtaining a short W-basis of G_H , the following boolean choice has to be taken for each one of these gadgets.

either locally pay $B + 1$: pay the extra price of $+1$ by locally applying a removal sequence avoiding the “cheap” edges wx, wy , and wz , hence allowing to disconnect cheaply the 4 nodes w, x, y , and z with later removals;

or locally pay B : pay just the minimum B , but then these 4 nodes will remain connected within the gadget since disconnecting them later would cost significantly more than than the $+1$. Indeed, by points (3) and (1) above, the cost of disconnecting any two nodes among the 4 nodes w, x, y will be at least $3 - 1 = 2$.

The above informal description will make full sense only later, after all the pieces of the proposed reduction will be in place. We now give our gadget. It is depicted in Fig. 31.

The weighted subgraph of G induced by the nodes a, b, c, y and w is called *the chamber*. The gadget graph G is similar to the flower F from Fig. 29 but it has 14 petals and 1 chamber, plus two edges wx and wz , dubbed the *jump edges*. In Fig. 31, the 14 petals are only hinted at for reasons of legibility. The numbers that label some of the nodes represent the distances from node w within the weighted graph $G \setminus \{wx, wz\}$ and, as such, certify the truth of the last two properties listed in the following lemma.

Lemma 6.4. *The graph G in Fig. 31 has 14 petals, one chamber, $n = 73$, $m = 147$, and $\nu = 75$. It has a unique minimum U-basis \mathcal{C}_{min} ; this basis is weakly fundamental and has weight $B := 377$. The edges wx, wy , and wz are all cotree edges w.r.t. any removal sequence encoding \mathcal{C}_{min} . There exists a W-basis of G of weight $B + 1$, and a removal sequence encoding this basis, w.r.t. which the edges wx, wy , and wz are all tree edges. The distance between w and y in $G \setminus yw$ is 4. The distance between w and x in $G \setminus xw$ is 5. The distance between w and z is 5 in $G \setminus zw$ and 6 in $G \setminus \{zw, xw\}$. The distance among any two nodes in $\{x, y, z\}$ is at least 4 in $G \setminus \{xw, yw, zw\}$.*

Proof. All claims in the first sentence are readily verified. As for the last four sentences, their truth can be readily verified through shortest path computations, and the distance values reported in Fig. 31 may partially support the reader in this task.

The remaining properties follow from the properties of the petals and from the structure of G . However, even if we hope that the reader might now possibly believe all these properties do hold, assessing them would be out of scope and we refer to Rizzi (2007) for more exhaustive arguments. \square

The source problem of the L-reduction: Hypergraphs generalize graphs. A hypergraph is a pair $H = (V, E)$, where V is a finite set of nodes and E is a finite set of *hyperedges*. A

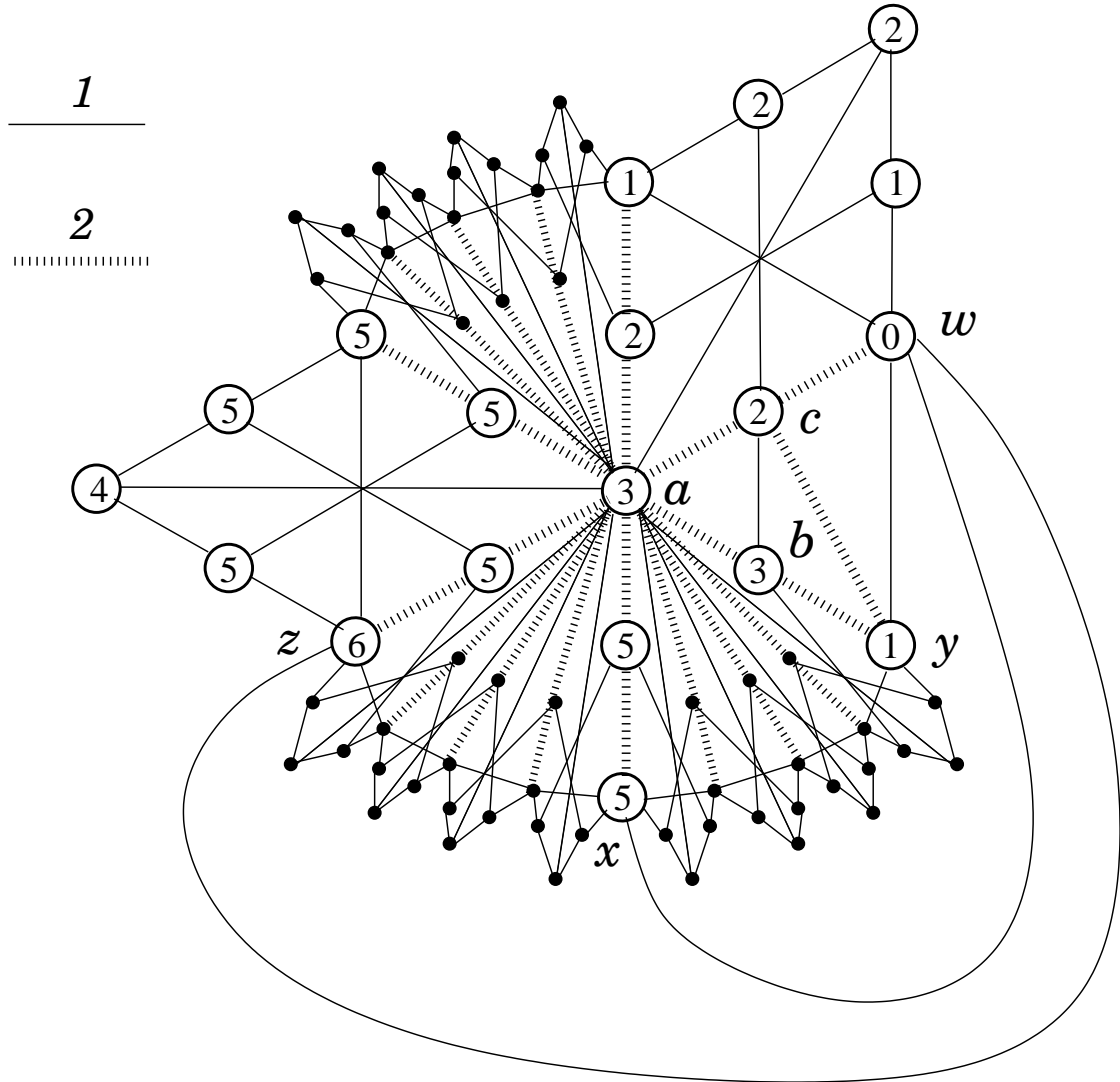


Figure 31: A gadget graph G with 11 petals, one chamber, and two jump edges wx and wz .

hyperedge is a set of nodes. When all hyperedges have size t , H is called t -uniform. Graphs are 2-uniform hypergraphs. A circuit in a hypergraph is an alternating sequence of nodes and edges $v_0, e_0, v_1, e_1, v_2, \dots, v_k, e_k$ such that, for every $i = 0, 1, 2, \dots, k$, we have $v_i \in e_i$ and $v_{i+1 \bmod k} \in e_i$. The *length* of the circuit is k , the number of edges comprising it. The *girth* γ_H of a hypergraph H is the minimum length of a circuit in H . A hypergraph is *acyclic* if it contains no circuit. A *feedback hyperedge set (FHS)* of a hypergraph $H = (V, E)$ is a set $F \subseteq E$ such that $(V, E \setminus F)$ is acyclic. Given a hypergraph H , the MINIMUM FEEDBACK HYPEREDGE SET (MFHS) problem seeks for a minimum cardinality FHS in H .

Lemma 6.5 (Rizzi (2007)). *There exists a constant $\alpha > 0$ such that the MFHS problem is APX-hard even when restricted to 4-uniform hypergraphs with $\gamma \geq 6$ in which a minimum cardinality FHS has size at least $\alpha |E|$.*

The main reduction: The main L-reduction is from the MFHS problem to the minimum W basis problem. Let $H = (V, E)$ be a 4-uniform hypergraph with $\gamma_H \geq 6$. Consider the weighted graph G_H obtained as follows.

1. start from node set V ;
2. add a further node r adjacent to all nodes in V through edges of weight 3;
3. for each hyperedge $e = \{v_1, v_2, v_3, v_4\} \in E$, add a new and private copy C_e of the gadget graph and perform the following 4 node identifications: $v_1 \leftrightarrow x$, $v_2 \leftrightarrow y$, $v_3 \leftrightarrow z$, and $v_4 \leftrightarrow w$.

The following lemma establishes that the above poly-time construction is an L-reduction from the MFHS problem restricted to instances conforming to the properties as in Lemma 6.5 to the minimum W-basis problem. As a consequence, the MWFCB problem is \mathcal{APX} -hard.

Lemma 6.6. *The hypergraph $H = (V, E)$ admits an FHS of size t iff G_H admits a W-basis of weight $(21 + B)m + t$, where $n = |V|$, $m = |E|$.*

Proof. Assume first that the hypergraph $H = (V, E)$ admits an FHS $F \subseteq E$ with $|F| = t$. We construct a W-basis \mathcal{C}_F of G_H and a removal sequence s encoding \mathcal{C}_F . We start with $\mathcal{C}_F := \emptyset$ and $s := \emptyset$ and set $G' := G_H$.

For each $e \in F$, we proceed as follows. First, at cost $(B + 1)$, we put in \mathcal{C}_F all circuits of the W-basis of C_e of weight $(B + 1)$. By Lemma 6.4, this basis can be encoded by a removal sequence with respect to which the edges xw , yw , and zw of C_e are all tree edges. Append this removal sequence to s meanwhile removing from G' the 75 edges it prescribes. After the removal of these edges C_e is acyclic. Furthermore, the only edges of C_e which are not bridges of G' are the edges xw , yw , and zw . Next, at cost $7 + 7 + 7 = 21$, remove from G' these three edges of component C_e , meanwhile appending them to s and adding to \mathcal{C}_F the three triangles they form together with node r . Each of these triangles costs $7 = 3 + 3 + 1$. Clearly, after the removal of these three edges, no circuit of G' can go through an edge of C_e . After this has been performed for each $e \in F$, we have paid $(B + 1 + 21)t = (21 + B)t + t$ in total.

At his point, the number of connected components of $(V, E \setminus F)$ is $n - 3(m - t)$; for each connected component C of $(V, E \setminus F)$, remove from G' all edges of the form rc , $c \in C$, except one. Each removal has cost $7 = 3 + 3 + 1$ (explained in more detail below) and adds a triangle through node r to \mathcal{C}_F . Once these edges have been removed, no circuit of G' contains r . We now explain in more detail how the removal of these edges can be performed within the claimed costs. First, select a node a of C and a spanning tree A of $G'[C]$. Then act as follows. Let $A' := A$. Consider A' as a tree rooted at a . While $A' \neq \{a\}$, consider any leaf q of A' and let p be the father of q in A' ; remove from G' the edge qr and append it to s , meanwhile inserting in \mathcal{C}_F the triangle $q - r - p$ (of cost $7 = 3 + 3 + 1$); remove node q from the rooted tree A' . In this way, we remove a total of $n - (n - 3(m - t)) = 3(m - t)$ edges, for a total cost of $21(m - t)$. Until here, we have paid $(B + 1)t + 21m$ in total.

Finally, for each $e \in E \setminus F$, put in \mathcal{C}_F all circuits of the W-basis of C_e of weight B . Also, append to s and remove from G' all cotree edges of C_e w.r.t. any removal sequence encoding this W-basis of C_e . After this, G' is a spanning tree of G_H . In particular, the acyclicity of G' follows from the acyclicity of $E \setminus F$. Thus, \mathcal{C}_F is a W-basis of G_H . In total, in this last step has cost $(m - t)B$ and hence the total cost of \mathcal{C}_F is $(21 + B)m + t$.

For the reverse direction, let \mathcal{C} be a W-basis of G_H of cost at most $(21 + B)m + t$. We may assume \mathcal{C} to be locally-optimal and encoded by means of a removal sequence s . Let T be the spanning tree of G_H made of the tree edges w.r.t. s . Let E' be the set of those hyperedges $e \in E$ such that $C_e \cap T$ is a spanning tree of C_e . Notice that the hypergraph (V, E') is acyclic since T is acyclic. Let $F := E \setminus E'$. It follows that F is an FHS of the hypergraph H . Let $f := |F|$. We will show $|F| \leq t$.

Let ν denote the cyclomatic number of G_H . By Lemma 6.4, the cyclomatic number of each gadget C_e is 75. Therefore $\nu = 75m + 3m = 78m$ since in order to make G_H acyclic, we need to remove $3m$ further edges after having made each C_e acyclic. Let $G^0 := G_H$, and, for $i = 1, 2, \dots, \nu$, let G^i be the weighted graph obtained from G^{i-1} by removing the i -th edge e_i of the removal sequence s . For every $e \in E$, we denote by V_e the nodes of the gadget C_e , and we say that edge e_i is *pertinent* to C_e if e_i is an edge of C_e and if the induced graphs $G^i[V_e]$ and $G^{i-1}[V_e]$ have the same number of connected components. Clearly, the removal sequence $s = e_1, e_2, \dots, e_\nu$ contains precisely $\nu_{C_e} = 75$ edges pertinent to C_e and the subsequence s_e of s comprising these 75 edges encodes a W-basis of C_e . Since the girth of H is at least 6, every circuit of G_H which is not a circuit of some C_e costs at least 7. As a consequence, for every $e \in E$, the total cost of the circuits in \mathcal{C} associated to the edges in s_e is at least B . Besides the $\nu = 75m$ removals considered until now (which in number are precisely enough to make each C_e gadget acyclic), we have $3m$ further removals. None of these further removals can cost less than 7, since none of the corresponding circuits in \mathcal{C} is entirely contained in one single C_e gadget. Furthermore, for every $e \in F$, the best edge removal for making $G_H[V_e]$ acyclic and disconnecting the subgraph $G_H[V_e] \cap T$ has cost $B + 1$. Indeed, for every $e \in F$, there exists an e_i in s such that $G^i[V_e]$ and $G^{i-1}[V_e]$ have the same number of connected components. As a consequence, the corresponding circuit C_i in \mathcal{C} contains an edge (the edge e_i) in C_e but is not entirely contained in C_e . Now, if e_i is neither the xw , nor the yw , nor the zw edge of C_e , then the cost of C_i is strictly bigger than 7 (actually, at least 10); otherwise, the total cost of the circuits in \mathcal{C} associated to the edges in s_e is at least $B + 1$. In total, the cost is at least $(m - f)B + f(B + 1) + 21m = mB + 21m + f$. Since we know that this cost is at most $(21 + B)m + t$, we conclude $f \leq t$. \square

7 Applications

Cycle bases arise in a wide range of engineering applications, we discuss three. They require different kinds of cycle basis. The analysis of electrical circuits can do with any kind of cycle basis, periodic scheduling problems in traffic planning require integral bases, and a graph drawing method requires strictly fundamental bases.

7.1 Kirchhoff's Voltage Law

Kirchhoff's circuits laws are the basis for the analysis of electrical circuits.

“The directed sum of the electrical potential differences around any closed circuit must be zero.”

(Wikipedia, page “Kirchhoff's circuit laws” as of Jan 04, 2009; Kirchhoff (1847))

Any system of differential (algebraic) equations which models an electrical system has to satisfy this property. Fortunately, it suffices to guarantee the zero-sum property for the elements of a directed cycle basis. Indeed, consider the cycle matrix Γ of some directed cycle basis and some arbitrary cycle C . Then $C = \Gamma\lambda$ for some coefficient vector λ . Now, if some vector x of potential differences satisfies Kirchhoff's law for every circuit in the basis, i.e., $x^T\Gamma = 0$, then $x^TC = x^T(\Gamma\lambda) = (x^T\Gamma)\lambda = \mathbf{0}^T\lambda = 0$. For a more detailed exposition of this application of cycle bases we refer to Bollobás (2002). An in-depth presentation of how cycle bases can be used for index reduction of differential algebraic systems is given in Bächle (2007).

7.2 Periodic Scheduling in Traffic Planning

Periodic scheduling problems arise frequently in traffic planning. Two examples are scheduling traffic lights and timetabling in public transport. They share a common mathematical model that can be traced back to early work by Gartner et al. (1975) and Rürger (1986) and was put into its final form by Serafini and Ukovich (1989).

In the *Periodic Event Scheduling Problem (PESP)* we are given a directed graph $D = (V, A)$, vectors ℓ and u on the arcs, and a scalar T called *cycle time* or *period*. For an arc a , $\ell(a)$ and $u(a)$ are lower and upper bounds for the travel time across a , respectively. In the feasibility version of the problem, the question is, whether a node potential π exists, such that

$$\ell_a \leq \pi_j - \pi_i + Tp_a \leq u_a, \quad \forall a = (i, j) \in A, \quad (1)$$

where p constitutes an *integral* vector on the arcs. The π_i is the event time at node i modulo the period T . It is a simple observation that PESP generalizes K-Vertex Colorability (Odijk (1996)). One may further add an objective function, which will depend on the application. We will next discuss the two applications in somewhat more detail and then make the connection to cycle basis.

Traffic Light Coordination. The task is to plan the red/green timings of traffic lights. We assume that the cycle time has already been determined, e.g., a 60 second cycle time is desired, and that minimum durations for the green phases of the individual signal groups (left turn lane, straight traffic etc.) have been derived from the traffic loads of the origin-destination pairs. It is then necessary to schedule the events for each signal group, i.e., when

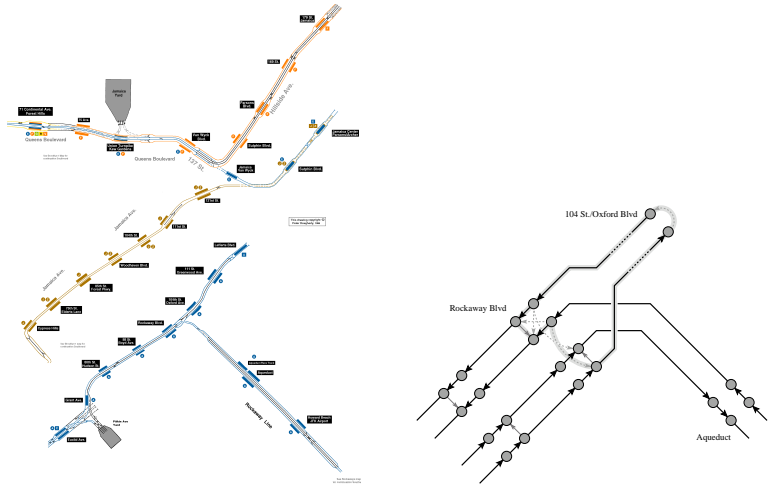


Figure 32: Excerpt of the New York City subway network in the area of JFK airport: The line plan, the track map, and the PESP graph. In the latter, light gray dotted arcs model passenger transfers or turnarounds of trains, and dark gray solid arcs model minimum headways and/or coordinated departures. A circuit in the PESP graph is highlighted.

signals turn from green to red and from red to green. Typical objectives are the minimization of the number of red-light stops of cars and the total travel time in the network. Wünsch (2008) discusses the traffic light coordination problem and related problems in detail.

Periodic Timetabling in Public Transport. The German train system runs essentially on either a one or a two hour period. Main lines runs on a one hour period, secondary lines operate on a two hour period. Of course, the period is not maintained during night hours. The Sunday schedule of the Berlin subway runs on a 10 minute period. Shorter periods are used on weekdays, in particular during rush hours. The process of timetabling in its full generality is highly complex, in particular when different train operation companies intend to use the same track for the very same time slot. We concentrate here on purely periodic schedules.

A periodic timetable assigns arrival and departure times to all pairs of lines and stations: Berlin metro line U9 leaves Zoo station southbound at minute 02, and arrives at the next station, Kurfürstendamm at minute 03. Many constraints have to be respected. These include minimum headways between two trains of different lines using the same track, collision-free service on a single track, and maximum durations for stops in intermediate stations of lines. Among the most important objectives are short transfer times for the passengers as well as short turnaround times for the trains in their terminus stations, where both have to respect certain minimum time durations, too. In Figure 32, these types of arcs are shown for a small part of New York City.

Cycle Bases for PESP. The practical performance of MIP solvers on PESP instances based on Equation (1) is rather poor. We next describe a more efficient problem formulation that makes use of integral cycle bases.

Given an (in-) feasible solution (π, p) of the PESP, consider the function $x(a)$ on the arcs,

$$x(a) := \pi_j - \pi_i + Tp_a, \quad a = (i, j). \quad (2)$$

The vector x is sometimes referred to as *periodic tension*, and models the time duration between its two events i and j . Summation of this equality for the arcs of any circuit C yields the *cycle periodicity property*

$$\frac{1}{T} \left(\sum_{a \in C^+} x_a - \sum_{a \in C^-} x_a \right) \in \mathbb{Z}; \quad (3)$$

the important observation is that the sum on the left must have an integral value for all circuits C . Nachtigall (1996) observed that the cycle periodicity property for the circuits in some strictly fundamental basis implies it for all circuits. Liebchen and Peeters (2009) generalized this result to integral cycle bases.

Theorem 7.1 (Liebchen and Peeters (2009); Nachtigall (1996)). *Let x be some vector on the arcs of a directed graph. There exists a pair (π, p) having (2), i.e., x is a periodic tension, if and only if the cycle periodicity property (3) is satisfied for the elements of some integral cycle basis.*

Indeed, assume that (3) holds for all the circuits in an integral basis and let C be an arbitrary circuit. Then $C = \sum_i \lambda_i C_i$, where the C_i 's are the basic circuits and the λ_i are integral. Then

$$\sum_a C(a)x_a = \sum_a \left(\sum_i \lambda_i C_i(a) \right) x_a = \sum_i \lambda_i \left(\sum_a C_i(a)x_a \right) = \sum_i \lambda_i q_{C_i} T$$

and hence the net travel time along C is an integral multiple of the period. For this argument to hold, it is essential that C is an *integral* linear combination of the basic circuits.

Practical Use. For both applications, the mathematical model sketched above made its way into practice – including the computation of short (fundamental) cycle bases as a pre-processing subroutine. For the traffic light scheduling problem, Wünsch (2008) reports that since 2008 the method is commercially available as a module in one of the major software suites in traffic planning. In periodic timetabling, Liebchen (2008) reports that the first mathematically optimized railway timetable went into service in 2005, for the Berlin subway network. About two years later, even a national railway company reports that their new timetable was designed with the help of combinatorial algorithms (Kroon et al. (2008)).

7.3 Graph Drawing

Graph drawing is concerned with embedding graphs into the plane in an aesthetically pleasing way. A position is assigned to each vertex and each edge is drawn as a (poly-)line. The goal is to obtain a clear, easily interpretable drawing of the graph. Lehmann and Kottler (2007) have shown that minimum or near minimum strictly fundamental cycle bases are very useful in this context.

They start with the observation that many real-world graphs, such as social networks, are *sparse* and simultaneously *clustered* in the sense, that the neighbors of a vertex are frequently

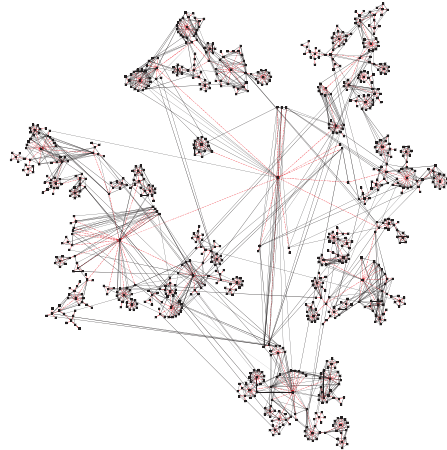


Figure 33: On the left side the graph is drawn by a refined version of a classic force-directed layout approach, on the right by computing an F -basis first that is then used to draw the whole graph as described in Lehmann and Kottler (2007). The edges of the spanning tree are marked in red.

also connected directly to each other. These edges will then form triangles. More generally, most edges of real-world graphs belong to triangles or at least short cycles. This is in contrast to sparse random graphs. However, there are usually also some edges that connect seemingly random vertices with each other (Watts and Strogatz (1998)). Edges of the first category are often called *local* edges and edges of the second category are called *global edges*. Although there is no clear definition for either of these categories, it is frequently desirable to show either the local structure or the global structure of the graph. The spanning tree underlying a (near) minimal cycle basis will provide the right scaffold. Moreover, it can easily be drawn in linear time with a tree drawing method (Kaufmann and Wagner (2001)).

With this spanning tree as a scaffold, global edges can now be defined as those edges that connect vertices with at least a given threshold distance in the tree. By adding them to the spanning tree the global structure of the graph can be emphasized. Analogously, by adding the other, non-global edges to the spanning tree, the local, clustered structure is prominently displayed. Thus, this method provides a neat way to show both, the local and global structure of a given graph, next to each other. Figure 33 shows an example.

References

- Abraham, I., Y. Bartal, and O. Neiman (2007). Embedding metrics into ultrametrics and graphs into spanning trees with constant average distortion. In N. Bansal, K. Pruhs, and C. Stein (Eds.), *SODA*, pp. 502–511. SIAM.
- Alimonti, P. and V. Kann (2000). Some APX-completeness results for cubic graphs. *Theor. Comput. Sci.* 237(1-2), 123–134.
- Alon, N., R. M. Karp, D. Peleg, and D. B. West (1995). A graph-theoretic game and its application to the k-server problem. *SIAM J. Comput.* 24(1), 78–100.
- Althöfer, I., G. Das, D. Dobkin, D. Joseph, and J. Soares (1993). On sparse spanners of weighted graphs. *Discrete Comput. Geom.* 9(1), 81–100.
- Apostol, T. (1997). *Introduction to Analytic Number Theory*. Springer.
- Bächle, S. (2007). *Numerical Solution of Differential-Algebraic Systems Arising in Circuit Simulation*. Ph.D. thesis, Technische Universität Berlin.
- Bafna, V., P. Berman, and T. Fujito (1999). A 2-approximation algorithm for the undirected feedback vertex set problem. *SIAM Journal on Discrete Mathematics* 12(3), 289–297.
- Barahona, F. and A. R. Mahjoub (1986). On the cut polytope. *Math. Programming* 36, 157–173.
- Bauer, U. (2004). Minimum cycle basis algorithms for the chemistry development toolkit. <http://geom.mi.fu-berlin.de/ubauer/CDK-Ringsearch.pdf>.
- Berge, C. (1962). *The Theory of Graphs and its Applications*. John Wiley & Sons, Inc.
- Berger, F. (2004). Kreisbasenbibliothek CyBaL. <http://www-m9.ma.tum.de/m9old/dm/projects/cycles/cybal/>.
- Berger, F., C. Flamm, P. M. Gleiss, J. Leydold, and P. F. Stadler (2004). Counterexamples in chemical ring perception. *Journal of Chemical Information and Computer Science* 44(2), 323–331.
- Berger, F., P. Gritzmann, and S. de Vries (2004). Minimum cycle bases for network graphs. *Algorithmica* 40, 51–62.
- Bollobás, B. (1978). *Extremal Graph Theory*. Academic Press, New York.
- Bollobás, B. (2002). *Modern Graph Theory*, Volume 184 of *Graduate Texts in Mathematics*. Springer. 2nd printing.
- Champetier, C. (1987). On the null-homotopy of graphs. *Discrete Mathematics* 64, 97–98.
- de Pina, J. C. (1995). *Applications of Shortest Path Methods*. Ph.D. thesis, Universiteit van Amsterdam.
- Deo, N., G. M. Prabhu, and M. Krishnamoorthy (1982). Algorithms for generating fundamental cycles in a graph. *ACM Transactions on Mathematical Software* 8(1), 26–42.

- Elkin, M., Y. Emek, D. A. Spielman, and S.-H. Teng (2008). Lower-stretch spanning trees. *SIAM J. Comput.* 38(2), 608–628.
- Elkin, M., C. Liebchen, and R. Rizzi (2007). New length bounds for cycle bases. *Inf. Process. Lett.* 104(5), 186–193.
- Fekete, S. P. and J. Kremer (2001). Tree spanners in planar graphs. *Discrete Applied Mathematics* 108(1-2), 85–103.
- Frederickson, G. (1987). Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Computing* 16, 1004–1022.
- Gabow, H. (1990). Data structures for weighted matching and nearest common ancestor with linking. In *SODA*, pp. 434–443.
- Galbiati, G. (2003). On finding cycle bases and fundamental cycle bases with a shortest maximal cycle. *Inf. Process. Lett.* 88(4), 155–159.
- Galbiati, G., R. Rizzi, and E. Amaldi (2007). On the approximability of the minimum strictly fundamental cycle bases problem. submitted.
- Gartner, N. H., J. D. Little, and H. Gabbay (1975). Optimization of traffic signal settings by mixed-integer linear programming, Part I: The network coordination problem. *Transportation Science* 9, 321–343.
- Gleiss, P. M. (2001a). Cycdeco. <http://www.tbi.univie.ac.at/pmg/cycdeco/cycdeco.html>.
- Gleiss, P. M. (2001b). *Short Cycles - Minimum Cycle Bases of Graphs from Chemistry and Biochemistry*. Ph. D. thesis, Fakultät für Naturwissenschaften und Mathematik der Universität Wien.
- Golynski, A. and J. D. Horton (2002). A polynomial time algorithm to find the minimum cycle basis of a regular matroid. In M. Penttonen and E. M. Schmidt (Eds.), *SWAT*, Volume 2368 of *Lecture Notes in Computer Science*, pp. 200–209. Springer.
- Grötschel, M., L. Lovász, and A. Schrijver (1988). *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag.
- Hariharan, R., T. Kavitha, and K. Mehlhorn (2006). A faster deterministic algorithm for minimum cycle bases in directed graphs. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener (Eds.), *ICALP (1)*, Volume 4051 of *Lecture Notes in Computer Science*, pp. 250–261. Springer.
- Hartvigsen, D. and R. Mardon (1994). The all-pairs min cut problem and the minimum cycle basis problem on planar graphs. *SIAM J. Discrete Math.* 7(3), 403–418.
- Hartvigsen, D. and E. Zemel (1989). Is every cycle basis fundamental? *Journal of Graph Theory* 13(1), 117–137.
- Horton, J. (1987). A polynomial-time algorithm to find the shortest cycle basis of a graph. *SIAM J. Comput.* 16, 359–366.

- Huber, M. (2003). Implementation of algorithms for sparse cycle bases of graphs. <http://www-m9.ma.tum.de/dm/cycles/mhuber>.
- Hubicka, E. and M. M. Sysło (1975). Minimal bases of cycles of a graph. In M. Fiedler (Ed.), *Recent advances in graph theory*, pp. 283–293. Academia Praha.
- Kaufmann, M. and D. Michail (2008). personal communication.
- Kaufmann, M. and D. Wagner (Eds.) (2001). *Drawing Graphs*. Number 2025 in Lecture Notes in Computer Science. Springer Verlag.
- Kavitha, T. and K. Mehlhorn (2005). A polynomial time algorithm for minimum cycle basis in directed graphs. In V. Diekert and B. Durand (Eds.), *STACS*, Volume 3404 of *Lecture Notes in Computer Science*, pp. 654–665. Springer.
- Kavitha, T., K. Mehlhorn, and D. Michail (2007). New approximation algorithms for minimum cycle bases of graphs. In W. Thomas and P. Weil (Eds.), *STACS*, Volume 4393 of *Lecture Notes in Computer Science*, pp. 512–523. Springer.
- Kavitha, T., K. Mehlhorn, D. Michail, and K. E. Paluch (2004). A faster algorithm for minimum cycle basis of graphs. In J. Díaz, J. Karhumäki, A. Lepistö, and D. Sannella (Eds.), *ICALP*, Volume 3142 of *Lecture Notes in Computer Science*, pp. 846–857. Springer.
- Kirchhoff, G. (1847). Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Vertheilung galvanischer Ströme geführt wird. *Annalen der Physikalischen Chemie* 72(12), 497–508.
- Köhler, E., C. Liebchen, R. Rizzi, and G. Wünsch (2009). Lower bounds for strictly fundamental cycle bases in grid graphs. *Networks* 53(2).
- Korte, B. and J. Vygen (2005). *Combinatorial Optimization*. Springer.
- Kroon, L., D. Huisman, E. Abbink, P.-J. Fioole, M. Fischetti, G. Maróti, L. Schrijver, A. Steenbeek, and R. Ybema (2008). The new dutch timetable: The OR revolution. Technical Report EI2008-19, Erasmus University Rotterdam, Econometric Institute. accepted for publication in Interfaces as winning paper of the 2008 INFORMS Franz Edelman Award.
- Lehmann, K. A. and S. Kottler (2007). Visualizing large and complex networks. In *Proceedings of the 14th International Symposium on Graph Drawing (GD'06)*.
- Liebchen, C. (2003). Finding short integral cycle bases for cyclic timetabling. In G. D. Battista and U. Zwick (Eds.), *ESA*, Volume 2832 of *Lecture Notes in Computer Science*, pp. 715–726. Springer.
- Liebchen, C. (2008). The first optimized railway timetable in practice. *Transportation Science* 42(4), 420–435.
- Liebchen, C. and L. Peeters (2009). Integral cycle bases for cyclic timetabling. *Discrete Optimization* 6(1), 98–109.
- Liebchen, C. and R. Rizzi (2005). A greedy approach to compute a minimum cycle basis of a directed graph. *Inf. Process. Lett.* 94(3), 107–112.

- Liebchen, C. and R. Rizzi (2007). Classes of cycle bases. *Discrete Applied Mathematics* 155(3), 337–355.
- Lubotzky, A., R. Phillips, and P. Sarnak (1988). Ramanujan graphs. *Combinatorica* 3, 261–277.
- MacLane, S. (1937). A combinatorial condition for planar graphs. *Fundamenta Mathematicae* 28, 22–32.
- Mehlhorn, K. and D. Michail (2006). Implementing minimum cycle basis algorithms. *ACM Journal of Experimental Algorithms* 11, 1–14.
- Mehlhorn, K. and D. Michail (2008). Minimum cycle bases: Faster and simpler. Accepted for publication in *ACM Transactions on Algorithms*.
- Miller, G. (1986). Finding small simple cycle separators for 2-connected planar graphs. *J. Comp. Syst. Sci.* 32, 265–279.
- Nachtigall, K. (1996). Periodic network optimization with different arc frequencies. *Discrete Applied Mathematics* 69(1-2), 1–17.
- Odiijk, M. A. (1996). A constraint generation algorithm for the construction of periodic railway timetables. *Transportation Research B* 30(6), 455–464.
- Papadimitriou, C. M. (1994). *Computational complexity*. Reading, Massachusetts: Addison-Wesley.
- Peleg, D. (2000). *Distributed Computing - A Locality Sensitive Approach*. SIAM Monographs on Discrete Mathematics and Applications.
- Reich, A. (2007). Streng fundamentale Kreisbasen planarer Graphen. Master’s thesis, Mathematisches Institut der Brandenburgischen Technischen Universität Cottbus.
- Rizzi, R. (2007). Minimum weakly fundamental cycle bases are hard to find. *Algorithmica online first*, doi: 10.1007/s00453-007-9112-8.
- Rüger, S. (1986). *Transporttechnologie städtischer öffentlicher Personenverkehr* (3rd ed.). Transpress, Verlag für Verkehrswesen. In German.
- Schrijver, A. (1986). *Theory of Linear and Integer Programming*. Wiley.
- Serafini, P. and W. Ukovich (1989). A mathematical model for periodic scheduling problems. *SIAM Journal on Discrete Mathematics* 2(4), 550–581.
- Stern, J. M. and S. A. Vavasis (1990). Nested dissection for sparse nullspace bases. Technical report, Department of Computer Science, Cornell University, Ithaca, NY 14853-7501.
- Sysło, M. M. (1979). On cycle bases of a graph. *Networks* 9, 123–132.
- Thorup, M. and U. Zwick (2001a). Approximate distance oracles. In *ACM Symposium on Theory of Computing*, pp. 183–192.
- Thorup, M. and U. Zwick (2001b). Compact routing schemes. In *Proceedings of 13th ACM Symposium on Parallel Algorithms and Architecture*, pp. 1–10.

- Trevisan, L. (2004). Inapproximability of combinatorial optimization problems. *Electronic Colloquium on Computational Complexity (ECCC)*.
- Ueckerdt, T. (2008). Berechnung kurzer ganzzahliger Kreisbasen in Graphen. Master's thesis, TU Berlin.
- Vismara, P. (1997). Union of all the minimum cycle bases of a graph. *Electronic Journal of Combinatorics* 4(1), 73–87.
- Watts, D. J. and S. H. Strogatz (1998, June). Collective dynamics of 'small-world' networks. *Nature* 393, 440–442.
- Wünsch, G. (2008). *Coordination of Traffic Signals in Networks and Related Graph Theoretical Problems on Spanning Trees*. Ph. D. thesis, Technische Universität Berlin, Institute of Mathematics.