Doubly Balanced Connected Graph Partitioning*

Saleh Soltan † Mihalis Yannakakis ‡ Gil Zussman§

Abstract

We introduce and study the Doubly Balanced Connected graph Partitioning (DBCP) problem: Let $G=(V,E)$ be a connected graph with a weight (supply/demand) function $p:V\rightarrow\{-1,+1\}$ satisfying $p(V)=\sum_{j\in V} p(j)=0$. The objective is to partition G into (V_1,V_2) such that $G[V_1]$ and $G[V_2]$ are connected, $|p(V_1)|,|p(V_2)|\leq c_p$, and $\max\{\frac{|V_1|}{|V_2|},\frac{|V_2|}{|V_1|}\}\leq c_s$, for some constants c_p and c_s. When G is 2-connected, we show that a solution with $c_p=1$ and $c_s=3$ always exists and can be found in polynomial time. Moreover, when G is 3-connected, we show that there is always a 'perfect' solution (a partition with $p(V_1)=p(V_2)=0$ and $|V_1|=|V_2|$, if $|V|\equiv 0(\bmod 4)$), and it can be found in polynomial time. Our techniques can be extended, with similar results, to the case in which the weights are arbitrary (not necessarily ± 1), and to the case that $p(V)\neq 0$ and the excess supply/demand should be split evenly. They also apply to the problem of partitioning a graph with two types of nodes into two large connected subgraphs that preserve approximately the proportion of the two types.

1 Introduction

Power Grid Islanding is an effective method to mitigate cascading failures in power grids [16]. The challenge is to partition the network into smaller connected components, called *islands*, such that each island can operate independently for a while. In order for an island to operate, it is necessary that the power supply and demand at that island are almost equal (if the supply and demand are not exactly equal but still relatively close, load shedding/generation curtailing can be used in order for the island to operate). Equality of supply and demand in an island, however, may not be sufficient for its independent operation. It is also important that the infrastructure in that island has the physical capacity to safely transfer the power from the supply nodes to the demand nodes. When the island is large enough compared to the initial network, it is more likely that it has enough capacity. This problem has been studied in the power systems community but almost all the algorithms provided in the literature are heuristic methods that have been shown to be effective only by simulations [8, 14–16].

Motivated by this application, we formally introduce and study the Doubly Balanced Connected graph Partitioning (DBCP) problem: Let $G=(V,E)$ be a connected graph with a weight (supply/demand) function $p:V\rightarrow\mathbb{Z}$ satisfying $p(V)=\sum_{j\in V} p(j)=0$. The objective is to partition V into (V_1,V_2) such that $G[V_1]$ and $G[V_2]$ are connected, $|p(V_1)|,|p(V_2)|\leq c_p$, and $\max\{\frac{|V_1|}{|V_2|},\frac{|V_2|}{|V_1|}\}\leq c_s$, for some constants c_p and c_s. We also consider the case that $p(V)\neq 0$, in which the excess supply/demand should be split roughly evenly.

The problem calls for a partition into two connected subgraphs that simultaneously balances two objectives, (1) the supply/demand within each part, and (2) the sizes of the parts. The connected partitioning problem with only the size objective has been studied previously. In the most well-known result, Lováz and Gyori [9, 13] independently proved, using different methods, that every k-connected graph can be partitioned into k arbitrarily sized connected subgraphs. However, neither of the proofs is constructive, and there are no known polynomial-time algorithms to find such a partition for $k>3$. For $k=2$, a linear time algorithm is provided in [17] and for $k=3$ an $O(|V|^2)$ algorithm is provided in [19].[1] The complexity of the problem with the size objective and related optimization problems have been studied in [3, 5, 6] and there are various NP-hardness and inapproximability results. Note that the size of the cut is not of any relevance here (so the extensive literature on finding balanced partitions, not necessarily

*This work was supported in part by DTRA grant HDTRA1-13-1-0021, DARPA RADICS under contract #FA-8750-16-C-0054, funding from the U.S. DOE OE as part of the DOE Grid Modernization Initiative, and NSF under grant CCF-1320654 and CCF-1423100. The work of G.Z. was also supported in part by the Blavatnik ICRC and the BSF. We thank Aliakbar Daemi for his helpful comments during our discussions.

†Department of Electrical Engineering, Columbia University. Email: saleh@ee.columbia.edu

‡Department of Computer Science, Columbia University. Email: mihalis@cs.columbia.edu

§Department of Electrical Engineering, Columbia University. Email: gil@ee.columbia.edu

[1] For $k=2$, a much simpler approach than the one in [17] is to use the st-numbering [11] for 2-connected graphs.

connected, that minimize the cut is not relevant.)

The objective of balancing the supply/demand alone, when all $p(i)$ are ± 1, can also be seen as an extension for the objective of balancing the size (which corresponds to $p(i){=}1$). Our bi-objective problem of balancing both supply/demand and size, can be seen also as an extension of the problem of finding a partition that balances the size for two types of nodes simultaneously: Suppose the nodes of a graph are partitioned into red and blue nodes. Find a partition of the graph into two large connected subgraphs that splits approximately evenly both the red and the blue nodes.

We now summarize our results and techniques. Since the power grids are designed to withstand a single failure ("$N{-}1$" standard) [1], and therefore 2-connected, our focus is mainly on the graphs that are at least 2-connected. We first, in Section 4, study the connected partitioning problem with only the supply/demand balancing objective, and show results that parallel the results for balancing size alone, using similar techniques: The problem is NP-hard in general. For 2-connected graphs and weights $p(i){=}\pm 1$, there is always a perfectly balanced partition and we can find it easily using an *st*-numbering. For 3-connected graphs and weights $p(i){=}\pm 1$, there is a perfectly balanced partition into three connected graphs, and we can find it using a nonseparating ear decomposition of 3-connected graphs [4] and similar ideas as in [19].

The problem is more challenging when we deal with both balancing objectives, supply/demand and size. This is the main focus and occupies the bulk of this paper. Our main results are existence results and algorithms for 2- and 3-connected graphs. It is easy to observe that we cannot achieve perfection in one objective ($c_p{=}0$ or $c_s{=}1$) without sacrificing completely the other objective. We show that allowing the supply/demand of the parts to be off balance by at most the weight of one node suffices to get a partition that is roughly balanced also with respect to size.

First, in Section 4.1, we study the case of 3-connected graphs since we use this later as the basis of handling 2-connected graphs. We show that if $\forall i,\ p(i){=}\pm 1$, there is a partition that is perfectly balanced with respect to both objectives, if $|V|{\equiv}0 (\text{mod } 4)$ (otherwise the sizes are slightly off for parity reasons); for general p, the partition is perfect in both objectives up to the weight of a single node. Furthermore, the partition can be constructed in polynomial time. Our approach uses the convex embedding characterization of k-connectivity studied by Linial, Lováz, and Wigderson [12]. We need to adapt it for our purposes so that the convex embedding also has certain desired geometric properties, and for this purpose we use the nonseparating ear decomposition of 3-connected graphs of [4] to obtain a suitable embedding.

Then, in Section 4.2, we analyze the case of 2-connected graphs. We reduce it to two subcases: either (1) there is a separation pair that splits the graph into components that are not very large, or (2) we can perform a series of contractions to achieve a 3-connected graph whose edges represent contracted subgraphs that are not too large. We provide a good partitioning algorithm for case (1), and for case (2) we extend the algorithms for 3-connected graphs to handle also the complications arising from edges representing contracted subgraphs. Finally, in Section 5, we briefly discuss the problem of finding a connected partitioning of a graph with two types of nodes that splits roughly evenly both types.

2 Definitions and Background

In this section, we provide a short overview of the definitions, notation, and tools used in our work. Most of the graph theoretical terms used in this paper are relatively standard and borrowed from [2] and [20]. All the graphs in this paper are loopless.

2.1 Terms from Graph Theory Cutpoints and Subgraphs:

A *cutpoint* of a connected graph G is a node whose deletion results in a disconnected graph. Let X and Y be subsets of the nodes of a graph G. $G[X]$ denotes the subgraph of G *induced* by X. We denote by $E[X,Y]$ the set of edges of G with one end in X and the other end in Y. The neighborhood of a node v is denoted $N(v)$.

Connectivity: The connectivity of a graph $G{=}(V,E)$ is the minimum size of a set $S{\subset}V$ such that $G \backslash S$ is not connected. A graph is k-connected if its connectivity is at least k.

2.2 *st*-numbering of a Graph

Given a pair of nodes $\{s,t\}$ in a 2-connected graph G, an *st-numbering* for G is a numbering for the nodes in G defined as follows [11]: the nodes of G are numbered from 1 to n so that s receives number 1, node t receives number n, and every node except s and t is adjacent both to a lower-numbered and to a higher-numbered node. It is shown in [7] that such a numbering can be found in $O(|V|{+}|E|)$.

2.3 Series-Parallel Graphs

A Graph G is *series-parallel*, with terminals s and t, if it can be produced by a sequence of the following operations:

1. Create a new graph, consisting of a single edge

between s and t.

2. Given two series parallel graphs, X and Y with terminals s_X, t_X and s_Y, t_Y respectively, form a new graph $G{=}P(X,Y)$ by identifying $s{=}s_X{=}s_Y$ and $t{=}t_X{=}t_Y$. This is known as the *parallel composition* of X and Y.

3. Given two series parallel graphs X and Y, with terminals s_X, t_X and s_Y, t_Y respectively, form a new graph $G{=}S(X,Y)$ by identifying $s{=}s_X, t_X{=}s_Y$ and $t{=}t_Y$. This is known as the *series composition* of X and Y.

It is easy to see that a series-parallel graph is 2-connected if, and only if, the last operation is a parallel composition.

2.4 Nonseparating Induced Cycles and Ear Decomposition Let H be a subgraph of a graph G. An *ear* of H in G is a nontrivial path in G whose ends lie in H but whose internal vertices do not. An ear decomposition of G is a decomposition $G{=}P_0{\cup}\cdots{\cup}P_k$ of the edges of G such that P_0 is a cycle and P_i for $i{\geq}1$ is an ear of $P_0{\cup}P_1{\cup}\cdots{\cup}P_{i-1}$ in G. It is known that every 2-connected graph has an ear decomposition (and vice-versa), and such a decomposition can be found in linear time.

A cycle C is a *nonseparating induced cycle* of G if $G\backslash C$ is connected and C has no chords. We say a cycle C avoids a node u, if $u{\notin}C$.

THEOREM 2.1. (TUTTE [18]) *Given a 3-connected graph $G(V,E)$ let $\{t,r\}$ be any edge of G and let u be any node of G, $r{\neq}u{\neq}t$. Then there is a nonseparating induced cycle of G through $\{t,r\}$ and avoiding u.*

Notice that since G is 3-connected in the previous theorem, every node in C has a neighbor in $G\backslash C$. Cheriyan and Maheshwari showed that the cycle in Theorem 2.1 can be found in $O(E)$ [4]. Moreover, they showed that any 3-connected graph G has a nonseparating ear decomposition $G{=}P_0{\cup}\cdots{\cup}P_k$ defined as follows: Let $V_i{=}V(P_0){\cup}V(P_1)\cdots{\cup}V(P_i)$, let $G_i{=}G[V_i]$ and $\overline{G}_i{=}G[V\backslash V_i]$. We say that $G{=}P_0{\cup}P_1{\cup}\cdots{\cup}P_k$ is an ear decomposition through edge $\{t,r\}$ and avoiding vertex u if the cycle P_0 contains edge $\{t,r\}$, and the last ear of length greater than one, say P_m, has u as its only internal vertex. An ear decomposition $P_0{\cup}P_1\cdots{\cup}P_k$ of graph G through edge $\{t,r\}$ and avoiding vertex u is a *nonseparating ear decomposition* if for all i, $0{\leq}i{<}m$, graph $\overline{G}_i$ is connected and each internal vertex of ear P_i has a neighbor in $\overline{G}_i$.

THEOREM 2.2. (CHERIYAN AND MAHESHWARI [4]) *Given an edge $\{t,r\}$ and a vertex u of a 3-connected graph G, a nonseparating induced cycle of G through $\{t,r\}$ and avoiding u, and a nonseparating ear decomposition can be found in time $O(|V|{+}|E|)$.*

2.5 Partitioning of Graphs to Connected Subgraphs The following theorem is the main existing result in partitioning of graphs into connected subgraphs and is proved independently by Lováz and Gyori [9, 13] by different methods.

THEOREM 2.3. (LOVÁZ AND GYORI [9, 13]) *Let $G{=}(V,E)$ be a k-connected graph. Let $n{=}|V|$, $v_1, v_2, \ldots, v_k{\in}V$ and let $n_1, n_2, \ldots, n_k$ be positive integers satisfying $n_1{+}n_2{+}\cdots{+}n_k{=}n$. Then, there exists a partition of V into $(V_1, V_2 \ldots, V_k)$ satisfying $v_i{\in}V_i$, $|V_i|{=}n_i$, and $G[V_i]$ is connected for $i{=}1,2,\ldots,k$.*

Although the existence of such a partition has long been proved, there is no polynomial-time algorithm to find such a partition for $k{>}3$. For $k{=}2$, it is easy to find such partition using st-numbering. For $k{=}3$, Wada and Kawaguchi [19] provided an $O(n^2)$ algorithm using the nonseparating ear decomposition of 3-connected graph.

2.6 Convex Embedding of Graphs In this subsection, we provide a short overview of the beautiful work by Linial, Lováz, and Wigderson [12] on convex embedding of the k-connected graphs. Let $Q{=}\{q_1, q_2, \ldots, q_m\}$ be a finite set of points in $\mathbb{R}^d$. The convex hull $\mathrm{conv}(Q)$ of Q is the set of all points $\sum_{i=1}^{m} \lambda_i q_i$ with $\sum_{i=1} \lambda_i{=}1$. The rank of Q is defined by $\mathrm{rank}(Q){=}1{+}\dim(\mathrm{conv}(Q))$. Q is in general position if $\mathrm{rank}(S){=}d{+}1$ for every $(d{+}1)$-subset $S{\subseteq}Q$. Let G be a graph and $X{\subset}V$. A convex X-embedding of G is any mapping $f{:}V{\to}\mathbb{R}^{|X|-1}$ such that for each $v{\in}V\backslash X$, $f(v){\in}\mathrm{conv}(f(N(v)))$. We say that the convex embedding is in general position if the set $f(V)$ of the points is in general position.

THEOREM 2.4. (LINIAL, LOVÁZ, AND WIGDERSON [12]) *Let G be a graph on n vertices and $1{<}k{<}n$. Then the following two conditions are equivalent:*

1. *G is k-connected*

2. *For every $X{\subset}V$ with $|X|{=}k$, G has a convex X-embedding in general position.*

Notice that the special case of the Theorem for $k{=}2$ asserts the existence of an st-numbering of a 2-connected graph. The proof of this theorem is inspired by physics. The embedding is found by letting the edges of the graph behave like ideal springs and letting its vertices settle. A formal summary of the proof ($1 \rightarrow 2$) is as follows (for more details see [12]). For each $v_i \in X$, define $f(v_i)$ arbitrary in $\mathbb{R}^{k-1}$ such that $f(X)$ is in general position. Assign to every edge $(u,v) \in E$ a positive elasticity coefficient c_{uv} and let $c \in \mathbb{R}^{|E|}$ be the vector of coefficients. It is proved in [12] that for almost any coefficient vector c, an embedding f that minimizes the potential function $P = \sum_{\{u,v\} \in E} c_{uv} \|f(u) - f(v)\|^2$ provides a convex X-embedding in general position ($\|.\|$ is the Euclidean norm). Moreover, the embedding that minimizes P satisfies the set of equations,

$$f(v) = \frac{1}{c_v} \sum_{u \in N(v)} c_{uv} f(u) \text{ for all } v \in V \setminus X,$$

where $c_v = \sum_{u \in N(v)} c_{uv}$. Hence, the embedding can be found by solving a set of linear equations, in at most $O(|V|^3)$ time (or matrix multiplication time).

3 Balancing the Supply/Demand Only

In this section, we study the single objective problem of finding a partition of the graph into connected subgraphs that balances (approximately) the supply and demand in each part of the partition, without any regard to the sizes of the parts. We can state the optimization problem as follows, and will refer to it as the Balanced Connected Partition with Integer weights (BCPI) problem.

Definition 3.1. *Given a graph $G=(V,E)$ with a weight (supply/demand) function $p : V \rightarrow \mathbb{Z}$ satisfying $\sum_{j \in V} p(j) = 0$. The BCPI problem is the problem of partitioning V into (V_1, V_2) such that*

1. *$V_1 \cap V_2 = \emptyset$ and $V_1 \cup V_2 = V$,*
2. *$G[V_1]$ and $G[V_2]$ are connected,*
3. *$|p(V_1)| + |p(V_2)|$ is minimized, where $p(V_i) = \sum_{j \in V_i} p(j)$.*

Clearly, the minimum possible value for $|p(V_1)| + |p(V_2)|$ that we can hope for is 0, which occurs iff $p(V_1) = p(V_2) = 0$. It is easy to show that the problem of determining whether there exists such a 'perfect' partition (and hence the BCPI problem) is strongly NP-hard. The proof is very similar to analogous results concerning the partition of a graph into two connected subgraphs with equal sizes (or weights, when nodes have positive weights) [3, 6].

Proposition 3.1. *(1) It is strongly NP-hard to determine whether there is a solution to the BCPI problem with value 0, even when G is 2-connected.*
(2) If G is not 2-connected, then this problem is NP-hard even when $\forall i, p(i) = \pm 1$.

Although it is NP-hard to tell whether there is a solution satisfying $p(V_1) = p(V_2) = 0$, even when $\forall i, p(i) = \pm 1$, in this case, if the graph G is 2-connected there is always such a solution. For general weights p, there is a solution such that $|p(V_1)|, |p(V_2)| \leq \max_{j \in V} |p(j)|/2$ and it can be found easily in linear time using the *st*-numbering between two nodes. (See the full paper for the proofs.)

Proposition 3.2. *Let G be a 2-connected graph and u, v any two nodes in V such that $p(u)p(v) > 0$.*
(1) There is a solution such that $u \in V_1$, $v \in V_2$, and $|p(V_1)| = |p(V_2)| \leq \max_{j \in V} |p(j)|/2$.
(2) If $\forall i, p(i) = \pm 1$, we can find a solution such that $u \in V_1$, $v \in V_2$, and $p(V_1) = p(V_2) = 0$.
In both cases, the solution can be found in $O(|E|)$ time.

Remark. The bound in Proposition 3.2 (1) is tight. A simple example is a cycle of length 4 like (v_1, v_2, v_3, v_4) with $p(v_1) = -p$, $p(v_2) = -p/2$, $p(v_3) = p$, and $p(v_4) = p/2$. It is easy to see that in this example $|p(V_1)| + |p(V_2)| = \max_{j \in V} |p(j)| = p$ is the best that one can do.

3.1 Connected Partitioning into Many Parts

The BCPI problem can be extended to partitioning a graph into $k=3$ or more parts. Let $G=(V,E)$ be a graph with a weight function $p : V \rightarrow \mathbb{Z}$ satisfying $\sum_{j \in V} p(j) = 0$. The BCPI$_k$ problem is the problem of partitioning G into $(V_1, V_2, \ldots, V_k)$ such that for any $1 \leq i \leq k$, $G[V_i]$ is connected and $\sum_{i=1}^{k} |p(V_i)|$ is minimized.

In the following proposition, we show that for $k=3$, if $p(i) = \pm 1, \forall i$, then there is always a perfect partition (i.e., with $p(V_1) = p(V_2) = p(V_3) = 0$) and it can be found efficiently. For general p, we can find a partition such that $|p(V_1)| + |p(V_2)| + |p(V_3)| \leq 2 \max_{j \in V} |p(j)|$. The proof and the algorithm use a similar approach as the algorithm in [19] for partitioning a 3-connected graph to three connected parts with prescribed sizes, using the nonseparating ear decomposition of 3-connected graphs as described in Subsection 2.4. (See the full paper for the proof.)

Proposition 3.3. *Let G be a 3-connected graph and u, v, w three nodes in V such that $p(u), p(v), p(w) > 0$ or $p(u), p(v), p(w) < 0$.*
(1) There is a solution such that $u \in V_1$, $v \in V_2$, $w \in V_3$, and $|p(V_1)| + |p(V_2)| + |p(V_3)| \leq 2 \max_{j \in V} |p(j)|$.
(2) If $\forall i, p(i) = \pm 1$, then there is a solution such that $u \in V_1$, $v \in V_2$, $w \in V_3$, and $|p(V_1)| = |p(V_2)| = |p(V_3)| = 0$.
In both cases, the solution can be found in $O(|E|)$ time.

4 Balancing Both Objectives

In this section, we formally define and study the Doubly Balanced Connected graph Partitioning (DBCP) problem.

DEFINITION 4.1. *Given a graph $G=(V,E)$ with a weight (supply/demand) function $p : V \to \mathbb{Z}$ satisfying $p(V)=\sum_{j\in V} p(j)=0$ and constants $c_p\geq 0$, $c_s\geq 1$. The* DBCP problem *is the problem of partitioning V into (V_1,V_2) such that*

1. $V_1\cap V_2=\emptyset$ *and* $V_1\cup V_2=V$,
2. $G[V_1]$ *and* $G[V_2]$ *are connected,*
3. $|p(V_1)|, |p(V_2)|\leq c_p$ *and* $\max\{\frac{|V_1|}{|V_2|}, \frac{|V_2|}{|V_1|}\}\leq c_s$, *where* $p(V_i)=\sum_{j\in V_i} p(j)$.

Remark. Our techniques apply also to the case that $p(V)\neq 0$. In this case, the requirement 3 on $p(V_1)$ and $p(V_2)$ is $|p(V_1)-p(V)/2|, |p(V_2)-p(V)/2|\leq c_p$, i.e., the excess supply/demand is split approximately evenly between the two parts.

We will concentrate on 2-connected and 3-connected graphs and show that we can construct efficiently good partitions. For most of the section we will focus on the case that $p(i)=\pm 1, \forall i\in V$. This case contains all the essential ideas. All the techniques generalize to the case of arbitrary p, and we will state the corresponding theorems.

We observed in Section 2 that if the graph is 2-connected and $p(i)=\pm 1, \forall i\in V$ then there is always a connected partition that is perfect with respect to the weight objective, $p(V_1)=p(V_2)=0$, i.e., (3) is satisfied with $c_p=0$. We know also from [9, 13] that there is always a connected partition that is perfect with respect to the size objective, $|V_1|=|V_2|$, i.e., condition 3 is satisfied with $c_s=1$. The following observations show that combining the two objectives makes the problem more challenging. If we insist on $c_p=0$, then c_s cannot be bounded in general, (it will be $\Omega(|V|)$), and if we insist on $c_s=1$, then c_p cannot be bounded. The series-parallel graphs of Figure 1 provide simple counterexamples.

OBSERVATION 1. *If $c_p=0$, then for any $c_s<|V|/2-1$, there exist a 2-connected graph G such that the DBCP problem does not have a solution even when $\forall i, p(i)=\pm 1$.*

Proof. In the graph depicted in Figure 1, set $t=0$.

OBSERVATION 2. *If $c_s=1$, then for any $c_p<|V|/6$, there exist a 2-connected graph G such that the DBCP problem does not have a solution even when $\forall i, p(i)=\pm 1$.*

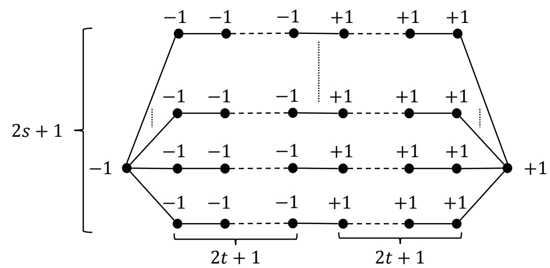

Figure 1: Series-parallel graphs with $2s+1$ paths of length $4t+2$ used in Observations 1 and 2.

Proof. In the graph depicted in Figure 1, set $s=1$.

Thus, c_p has to be at least 1 to have any hope for a bounded c_s. We show in this section that $c_p=1$ suffices for all 2-connected graphs. We first treat 3-connected graphs.

4.1 3-Connected Graphs Let $G=(V,E)$ be a 3-connected graph. Assume for most of this section that $\forall i, p(i)=\pm 1$ and $p(V)=0$ (we will state the results for general p at the end). We show that G has a partition that is essentially perfect with respect to both objectives, i.e., with $c_p=0$ and $c_s=1$. We say "essentially", because $p(V_1)=p(V_2)=0$ and $|V_1|=|V_2|$ imply that $|V_1|=|V_2|$ are even, and hence V must be a multiple of 4. If this is the case, then indeed we can find such a perfect partition. If $|V|\equiv 2 \pmod 4$ ($|V|$ has to be even since $p(V)=0$), then we can find an 'almost perfect' partition, one in which $|p(V_1)|=|p(V_2)|=1$ and $|V_1|=|V_2|$ (also one in which $p(V_1)=p(V_2)=0$ and $|V_1|=|V_2|+2$).

We first treat the case that G contains a triangle (i.e., cycle of length 3). In the following Lemma, we use the embedding for k-connected graphs introduced in [12] and as described in Subsection 2.6, to show that if G is 3-connected with a triangle and all weights are ± 1, then the DBCP problem has a perfect solution.

LEMMA 4.1. *If G is 3-connected with a triangle, $\forall i, p(i)=\pm 1$, and $|V|\equiv 0 \pmod 4$, then there exists a solution to the DBCP problem with $p(V_1)=p(V_2)=0$ and $|V_1|=|V_2|$. If $|V|\equiv 2 \pmod 4$, then there is a solution with $p(V_1)=p(V_2)=0$ and $|V_1|=|V_2|+2$. Moreover, this partition can be found in polynomial time.*

Proof. Assume that $|V|\equiv 0 \pmod 4$; the proof for the case $|V|\equiv 2 \pmod 4$ is similar. In [12] as described in Subsection 2.6, it is proved that if G is a k-connected graph, then for every $X\subset V$ with $|X|=k$, G has a convex X-embedding in general position. Moreover, this embedding can be found by solving a set of linear equations of size $|V|$. Now, assume $v, u, w\in V$ form a triangle in G. Set $X{=}\{v,u,w\}$. Using the theorem, G has a convex X-embedding $f{:}V{\to}\mathbb{R}^2$ in general position. Consider a circle $\mathcal{C}$ around the triangle $f(u), f(v), f(w)$ in $\mathbb{R}^2$ as shown in an example in Fig. 2. Also consider a directed line $\mathcal{L}$ tangent to the circle $\mathcal{C}$ at point A. If we project the nodes of G onto the line $\mathcal{L}$, since the embedding is convex and also $\{u,v\},\{u,w\},\{w,v\}{\in}E$, the order of the nodes' projection gives an st-numbering between the first and the last node (notice that the first and last nodes are always from the set X). For instance in Fig. 2, the order of projections give an st-numbering between the nodes u and v in G. Hence, if we set V_1 to be the $|V|/2$ nodes whose projections come first and V_2 are the $|V|/2$ nodes whose projections come last, then $G[V_1]$ and $G[V_2]$ are both connected and $|V_1|{=}|V_2|{=}|V|/2$. The only thing that may not match is $p(V_1)$ and $p(V_2)$. Notice that for each directed line tangent to the circle $\mathcal{C}$, we can similarly get a partition such that $|V_1|{=}|V_2|{=}|V|/2$. So all we need is a point D on the circle $\mathcal{C}$ such that if we partition based on the directed line tangent to $\mathcal{C}$ at point D, then $p(V_1){=}p(V_2){=}0$. To find such a point, we move $\mathcal{L}$ from being tangent at point A to point B (AB is a diameter of the circle $\mathcal{C}$) and consider the resulting partition. Notice that if at point A, $p(V_1){>}0$, then at point B since V_1 and V_2 completely switch places compared to the partition at point A, $p(V_1){<}0$. Hence, as we move $\mathcal{L}$ from being tangent at point A to point B and keep it tangent to the circle, in the resulting partitions, $p(V_1)$ goes from some positive value to a non-positive value. Notice that the partition (V_1,V_2) changes only if $\mathcal{L}$ passes a point D on the circle such that at D, $\mathcal{L}$ is perpendicular to a line that connects $f(i)$ to $f(j)$ for some $i,j{\in}V$. Now, since the embedding is in general position, there are exactly two points on every line that connects two points $f(i)$ and $f(j)$, so V_1 changes at most by one node leaving V_1 and one node entering V_1 at each step as we move $\mathcal{L}$. Hence, $p(V_1)$ changes by either ±2 or 0 value at each change. Now, since $|V|{\equiv}0(\text{mod } 4)$, $p(V_1)$ has an even value in all the resulting partitions. Therefore, as we move $\mathcal{L}$ from being tangent at point A to point B, there must be a point D such that in the resulted partition $p(V_1){=}p(V_2){=}0$.

It is also easy to see that since V_1 may change only when a line that passes through 2 nodes of graph G is perpendicular to $\mathcal{L}$, we can find D in at most $O(|V|^2)$ steps. $\square$

When G is a triangle-free 3-connected graph, however, the proof in Lemma 4.1 cannot be directly used anymore. The reason is if for example $\{u,v\}{\notin}E$ and we project the nodes of G onto the line $\mathcal{L}$, this time

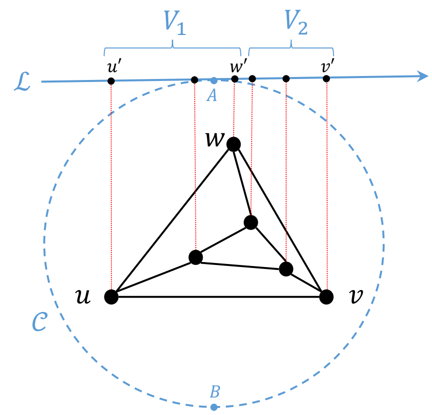

Figure 2: Proof of Lemma 4.1.

the order of the nodes projection does not give an st-numbering between the first and the last node if for example u and w are the first and last node, since some of the middle nodes may only be connected to v. To prove a similar result for triangle-free 3-connected case, we first provide the following two Lemmas. The main purpose of the following two Lemmas are to compensate for the triangle-freeness of G in the proof of Lemma 4.1. The idea is to show that in every 3-connected graph, there is a triple $\{u,w,v\}{\in}V$, such that $\{u,w\},\{w,v\}{\in}E$ and in every partition that we get by the approach used in the proof of Lemma 4.1, if u and v are in V_i, so is a path between u and v.

Lemma 4.2. *If G is 3-connected, then there exists a set $\{u,v,w\}{\subset}V$ and a partition of V into (V_1',V_2') such that:*

1. *$V_1'{\cap}V_2'{=}\emptyset$ and $V_1'{\cup}V_2'{=}V$,*
2. *$G[V_1']$ and $G[V_2']$ are connected,*
3. *$\{u,w\},\{v,w\}{\in}E$,*
4. *$w{\in}V_1'$, $u,v{\in}V_2'$,*
5. *$|V_2'|{\leq}|V|/2$.*

Moreover, such a partition and $\{u,v,w\}$ can be found in $O(|E|)$ time.

Proof. Using the algorithm presented in [4], we can find a non-separating cycle C_0 in G such that every node in C_0 has a neighbor in $G{\setminus}C_0$ in $O(|E|)$ time. Now, we consider two cases:

(i) If $|C_0|{\leq}|V|/2{+}1$, then select any three consecutive nodes (u,w,v) of C_0 and set $V_2'{=}C_0{\setminus}\{w\}$ and $V_1'{=}V{\setminus}V_2'$.

(ii) If $|C_0|{>}|V|/2{+}1$, since every node in C_0 has a neighbor in $G{\setminus}C_0$, there exists a node $w{\in}V{\setminus}C_0$

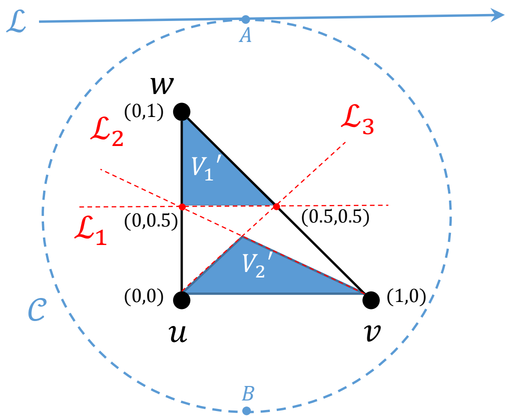

Figure 3: Proof of Lemma 4.3 and Theorem 4.1.

such that $|N(w) \cap C_0| \geq 2$. Select two nodes $u, v \in N(w) \cap C_0$. There exists a path P in C_0 from u to v such that $|P| < |V|/2+1$. Set $V_2' = P$ and $V_1' = V \setminus V_2'$.

□

Lemma 4.3. *Given a partition (V_1', V_2') of a 3-connected graph G with properties described in Lemma 4.2, G has a convex X-embedding in general position with mapping $f: V \to \mathbb{R}^2$ such that:*

1. *$X = \{u, w, v\}$, $f(u) = (0,0)$, $f(v) = (1,0)$, and $f(w) = (0,1)$,*
2. *Every node i in V_1' is mapped to a point $(f_1(i), f_2(i))$ with $f_2(i) \geq 1/2$,*
3. *Every node i in V_2' is mapped to a point $(f_1(i), f_2(i))$ with $f_1(i) \geq f_2(i)$ and $f_1(i) + 2f_2(i) \leq 1$.*

Moreover, such an embedding can be found in polynomial time.

Sketch of the proof. Set $X = \{v, u, w\}$. Using [12], G has a convex X-embedding in $\mathbb{R}^2$ in general position with mapping $f: V \to \mathbb{R}^2$ such that $f(u) = (0,0)$, $f(v) = (1,0)$, and $f(w) = (0,1)$. In the X-embedding of the nodes, we have a freedom to set the elasticity coefficient vector $\vec{c}$ to anything that we want (except a measure zero set of vectors). So for any edge $\{i,j\} \in G[V_1'] \cup G[V_2']$, set $c_{ij} = g$; and for any $\{i,j\} \in E[V_1', V_2']$, set $c_{ij} = 1$. Since both $G[V_1']$ and $G[V_2']$ are connected, as we increase g, nodes in V_1' get closer to w and nodes in V_2' get closer to the line uv (as $g \to \infty$, nodes in V_1' get in the same position as w and node in V_2' get on the line uv). Hence, intuitively there exists a value g (with polynomially many bits), for which all the nodes in V_1 are above line $\mathcal{L}_1$ and all the nodes in V_2' are below the lines $\mathcal{L}_2$ and $\mathcal{L}_3$ as depicted in Fig. 3. See the full paper for the detailed proof which shows also that a g with polynomially many bits suffices. □

Using Lemmas 4.2 and 4.3, we are now able to prove that for any 3-connected graph G such that all the weights are ± 1, the DBCP problem has a solution for $c_p = 0$ and $c_s = 1$. The idea of the proof is similar to the proof of Lemma 4.1, however, we use Lemma 4.2 to find a desirable partition (V_1', V_2') and then use this partition to find an embedding with properties as described in Lemma 4.3. By using this embedding, we can show that in every partition that we obtain by the approach in the proof of Lemma 4.1, if u and v are in V_i, so is a path between u and v. This implies then that $G[V_1]$ and $G[V_2]$ are connected. So we can use similar arguments as in the proof of Lemma 4.1 to prove the following theorem.

Theorem 4.1. *If G is 3-connected, $\forall i, p(i) = \pm 1$, and $|V| \equiv 0 (\text{mod } 4)$, then there exists a solution to the DBCP problem with $p(V_1) = p(V_2) = 0$ and $|V_1| = |V_2|$. If $|V| \equiv 2 (\text{mod } 4)$, then there is a solution with $p(V_1) = p(V_2) = 0$ and $|V_1| = |V_2| + 2$. Moreover, this partition can be found in polynomial time.*

It is easy to check for a 3-connected graph G, by using the same approach as in the proof of Lemma 4.1 and Theorem 4.1, that even when the weights are arbitrary (not necessarily ± 1) and also $p(V) \neq 0$, we can still find a connected partition (V_1, V_2) for G such that $|p(V_1) - p(V)/2|, |p(V_1) - p(V)/2| \leq \max_{i \in V} |p(i)|$ and $|V_1| = |V_2|$.

Corollary 4.1. *If G is 3-connected, then the DBCP problem (with arbitrary p, and not necessarily satisfying $p(V) = 0$) has a solution for $c_p = \max_{i \in V} |p(i)|$ and $c_s = 1$. Moreover, this solution can be found in polynomial time.*

4.2 2-Connected Graphs We first define a *pseudo-path* between two nodes in a graph as below. The definition is inspired by the definition of the *st*-numbering.

Definition 4.2. *A* pseudo-path *between nodes u and v in $G = (V, E)$, is a sequence of nodes $v_1, \ldots, v_t$ such that if $v_0 = u$ and $v_{t+1} = v$, then for any $1 \leq i \leq t$, v_i has neighbors v_j and v_k such that $j < i < k$. Note that the pseudo-path does not include the ending points u and v.*

Using the pseudo-path notion, in the following lemma we show that if G is 2-connected and has a separation pair such that none of the resulting components are too large, then the DBCP problem always has a solution for some $c_p = c_s = O(1)$. The idea used in the

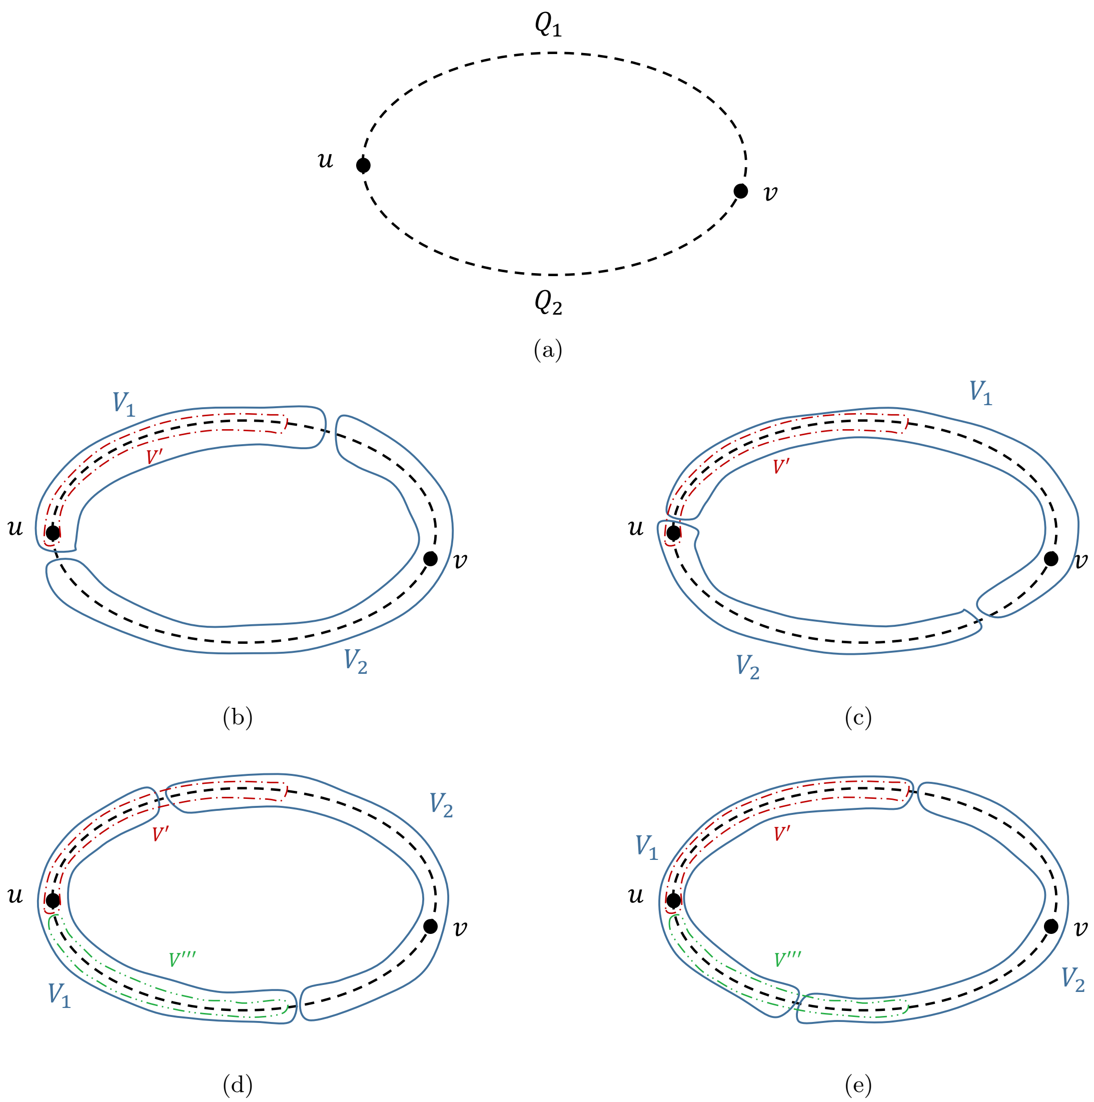

Figure 4: Proof of Lemma 4.4.

proof of this lemma is one of the building blocks of the proof for the general 2-connected graph case.

Lemma 4.4. *Given a 2-connected graph G and an integer $q{\geq}3$, if $\forall i{:}p(i){=}{\pm}1$ and G has a separation pair $\{u,v\}{\subset}V$ such that for every connected component $H_i{=}(V_{H_i},E_{H_i})$ of $G[V\setminus\{u,v\}]$, $|V_{H_i}|{<}(q{-}1)|V|/q$, then the DBCP problem has a solution for $c_p{=}1$, $c_s{=}q{-}1$, and it can be found in $O(|E|)$ time.*

Proof. Assume for simplicity that $|V|$ is divisible by q. There is a separation pair $\{u,v\}{\in}V$ such that if $H_1,\ldots,H_k$ are the connected components of $G\setminus\{u,v\}$, for any i, $|V_{H_i}|{<}(q{-}1)|V|/q$. Since G is 2-connected, $H_1,\ldots,H_k$ can be presented by pseudo-paths $P_1,\ldots,P_k$ between u and v. Assume $P_1,\ldots,P_k$ are in increasing order based on their lengths. There exists two subsets of the pseudo-paths S_1 and S_2 such that $S_1{\cap}S_2{=}\emptyset$, $S_1{\cup}S_2{=}\{P_1,\ldots,P_k\}$ and $\sum_{P_j\in S_i}|P_j|{\geq}|V|/q{-}1$ for $i{=}1,2$. The proof is very simple. Add greedily pseudo-paths in order to S_1 until its size becomes at least $|V|/q{-}1$. Let $S_1{=}\{P_1,\ldots,P_i\}$, and $S_2{=}\{P_{i+1},\ldots,P_k\}$. Since $|P_k|{<}(q{-}1)|V|/q$, we have $i{<}k$, and $S_2{\neq}\emptyset$. We have to show that $|S_2|{\geq}|V|/q{-}1$. If $|P_k|{\geq}|V|/q{-}1$, then also $|S_2|{\geq}|V|/q{-}1$. If $|P_k|{<}|V|/q{-}1$, then $|P_1|{+}\ldots{+}|P_{i-1}|{<}|V|/q{-}1$ and $|P_i|{\leq}|P_k|{<}|V|/q{-}1$ imply that $|S_2|{\geq}|V|/q{-}1$, since $q{\geq}3$.

Now, if we put all the pseudo-paths in S_1 back to back, they will form a longer pseudo-path Q_1 between u and v. Similarly, we can form another pseudo-path Q_2 from the pseudo-paths in S_2 (Fig. 4a). Without loss of generality we can assume $|Q_1|{\geq}|Q_2|$. From u, including u itself, we count $|V|/q$ of the nodes in Q_1 towards v and put them in a set V'. Without loss of generality, we can assume $p(V')\geq 0$. If $p(V')=0$, then $(V', V\backslash V')$ is a good partition and we are done. Hence, assume $p(V')>0$. We keep V' fixed and make a new set V'' by continuing to add nodes from Q_1 to V' one by one before we get to v. If $p(V'')$ hits 0 as we add nodes one by one, we stop and let $V_1=V''$ and $V_2=V\backslash V''$, then (V_1, V_2) is a good partition and we are done (Fig. 4b). So, assume $V''=Q_1\cup\{u\}$ and $p(V'')>0$. Since $|Q_2\cup\{v\}|\geq|V|/q$, $|V''|\leq(q-1)|V|/q$. If $|V''|<(q-1)|V|/q$, we add nodes from $Q_2\cup\{v\}$ one by one toward u until either $p(V'')=0$ or $|V''|=(q-1)|V|/q$. If we hit 0 first (i.e., $p(V'')=0$) and $|V''|\leq(q-1)|V|/q$, define $V_1=V''\backslash\{u\}$, then $(V_1, V\backslash V_1)$ is a good partition (Fig. 4c). So assume $|V''|=(q-1)|V|/q$ and $p(V'')>0$. Define $V'''=V\backslash V''$. Since $p(V'')>0$ and $|V''|=(q-1)|V|/q$, then $p(V''')<0$ and $|V'''|=|V|/q$. Also notice that $V'''\subseteq Q_2$. We consider two cases. Either $|p(V')|\geq|p(V''')|$ or $|p(V')|<|p(V''')|$.

If $|p(V')|\geq|p(V''')|$, then if we start from u and pick nodes one by one from Q_1 in order, we can get a subset V_1' of V' such that $|p(V_1')|=|p(V''')|$. Hence, if we define $V_1=V_1'\cup V'''$, then $(V_1, V\backslash V_1)$ is a good partition (Fig. 4d).

If $|p(V')|<|p(V''')|$, then we can build a new set V_1 by adding nodes one by one from Q_2 to V' until $P(V_1)=0$. It is easy to see that since $|p(V')|<|p(V''')|$, then $V_1\backslash V'\subseteq V'''$. Hence, $(V_1, V\backslash V_1)$ is a good partition (Fig. 4e). □

Corollary 4.2. *If G is a 2-connected series-parallel graph and $\forall i{:}p(i)=\pm 1$, then the DBCP problem has a solution for $c_p=1$, $c_s=2$, and the solution can be found in $O(|E|)$ time.*

The graph in Figure 1 with $s=1$ shows that these parameters are the best possible for series parallel graphs: if $c_p=O(1)$ then c_s must be at least 2.

To generalize Lemma 4.4 to all 2-connected graphs, we need to define the *contractible* subgraph and the *contraction* of a given graph as below.

Definition 4.3. *We say an induced subgraph H of a 2-connected graph G is* contractible, *if there is a separating pair $\{u, v\}\subset V$ such that $H=(V_H, E_H)$ is a connected component of $G[V\backslash\{u, v\}]$. Moreover, if we replace H by a weighted edge e' with weight $w(e')=|V_H|$ between the nodes u and v in G to obtain a smaller graph G', we say G is* contracted *to G'.*

Remark. Notice that every contractible subgraph of a 2-connected graph G can also be represented by a pseudo-path between its associated separating pair. We use this property in the proof of Theorem 4.2.

Using the notion of the graph contraction, in the following lemma, we show that to partition a 2-connected graph, we can reduce it to one of two cases: either G can be considered as a graph with a set of short pseudo-paths between two nodes, or it can be contracted into a 3-connected graph as illustrated in Fig. 5.

Lemma 4.5. *In every 2-connected graph $G=(V, E)$, given an integer $q\geq 3$, one of the following cases holds, and we can determine which in $O(|E|)$ time:*

1. *There is a separation pair $\{u, v\}\subset V$ such that if $H_1, \ldots, H_k$ are the connected components of $G[V\backslash\{u, v\}]$, for all i, $|V_{H_i}|<(q-1)|V|/q$.*
2. *After a set of contractions, G can be transformed into a 3-connected graph $G^*=(V^*, E^*)$ with weighted edges representing contracted subgraphs such that for every $e^*\in E^*$, $w(e^*)<|V|/q$.*

Proof. If there is no separation pairs in G, then G is 3-connected and there is nothing left to prove. So assume $\{u, v\}\subset V$ is a separation pair and $H_1, \ldots, H_k$ are the connected components of $G[V\backslash\{u, v\}]$. If $\forall i, |V_{H_i}|<(q-1)|V|/q$, we are done. So let's assume there is a connected component H_j such that $|V_{H_j}|\geq(q-1)|V|/q$. Then for every $i\neq j$, H_i can be contracted and represented by an edge of weight less than $|V|/q$ between u and v. Now, we repeat the process by considering the weight of the edges in the size of each connected component (a weighted edge can be contracted again as part of a new connected component and its weight will be added to the total number of nodes in that connected component). An example for each case is shown in Fig. 5 for $q=3$. We can find either a suitable separation pair as in case 1 or a suitable contracted graph G^* as in case 2 in linear time using the Hopcroft-Tarjan algorithm for finding the triconnected components [10]. □

Using Lemma 4.5 for $q=4$, then Lemma 4.4, and the idea of the proof for Theorem 4.1, we can prove that when G is 2-connected and all $p(i)=\pm 1$, the DBCP problem has a solution for $c_p=1$ and $c_s=3$. There are some subtleties in adapting Lemma 4.2 for this case to account for the fact that the edges of G^* are now weighted, and the partition (V_1', V_2') has to take into account the edge weights. We find a suitable convex embedding of the 3-connected graph G^* and then embed the nodes of the contracted pseudo-paths appropriately along the segments corresponding to the weighted edges. Some care is needed to carry out the argument of the 3-connected case, since as the line tangent to the circle rotates, the order of the projections of many nodes may change at once,

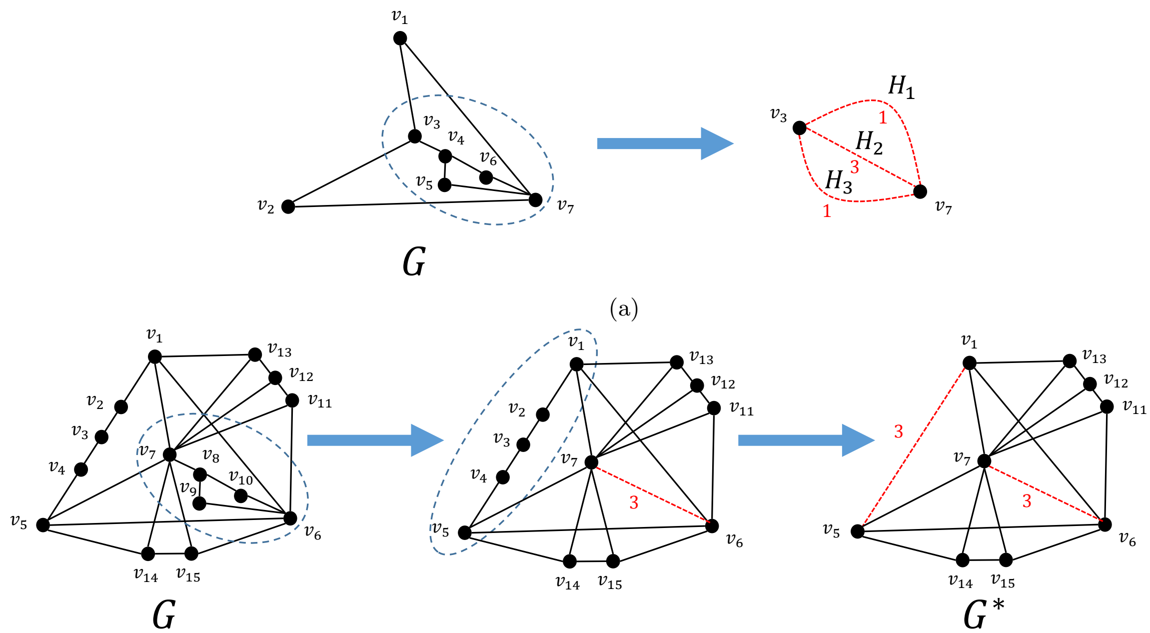

(a)

(b)

Figure 5: Lemma 4.5.

namely the nodes on an edge perpendicular to the rotating line. The details of the proof are given in the full paper. We have:

THEOREM 4.2. *If G is 2-connected, $\forall i, p(i)=\pm 1$, then the DBCP problem has a solution for $c_p=1$ and $c_s=3$. Moreover, this solution can be found in polynomial time.*

Similar to Corollary 4.1, the approach used in the proof of Theorem 4.2, can also be used for the case when the weights are arbitrary (not necessarily ± 1) and $p(V) \neq 0$. It can be shown that in this case, if G is 2-connected, the DBCP problem has a connected partition (V_1, V_2) such that $|p(V_1)-p(V)/2|, |p(V_2)-p(V)/2| \leq \max_{j \in V} |p(j)|$ and $|V_1|, |V_2| \geq |V|/4$.

COROLLARY 4.3. *If G is 2-connected, then the DBCP problem (with general p and not necessarily satisfying $p(V)=0$) has a solution for $c_p = \max_{j \in V} |p(j)|$ and $c_s=3$. Moreover, this solution can be found in polynomial time.*

5 Graphs with Two Types of Nodes

Assume G is a connected graph with nodes colored either red ($R \subseteq V$) or blue ($B \subseteq V$). Let $|V|=n$, $|R|=n_r$, and $|B|=n_b$. If G is 3-connected, set $p(i)=1$ if $i \in R$ and $p(i)=-1$ if $i \in B$. Corollary 4.1 implies then that there is always a connected partition (V_1, V_2) of V that splits both the blue and the red nodes evenly (assuming n_r and n_b are both even), i.e., such that $|V_1|=|V_2|$, $|R \cap V_1|=|R \cap V_2|$, and $|B \cap V_1|=|B \cap V_2|$. (If n_r and/or n_b are not even, then one side will contain one more red or blue node.)

COROLLARY 5.1. *Given a 3-connected graph G with nodes colored either red ($R \subseteq V$) or blue ($B \subseteq V$). There is always a partition (V_1, V_2) of V such that $G[V_1]$ and $G[V_2]$ are connected, $|V_1|=|V_2|$, $|R \cap V_1| = |R \cap V_2|$, and $|B \cap V_1|=|B \cap V_2|$ (assuming $|R|$ and $|B|$ are both even). Such a partition can be computed in polynomial time.*

Proof. Suppose without loss of generality that $n_r \geq n_b$ and let $n_r-n_b=2t$ and $n_r+n_b=n=2m$. Set $p(i)=1$ for $i \in R$ and $p(i)=-1$ for $i \in B$. Then $p(V)=2t$. From the equations, we have $n_r=m+t$ and $n_b=m-t$.

From Corollary 4.1 we can find a partition (V_1, V_2) such that $|V_1|=|V_2|$ and $|p(V_1)-p(V)/2|, |p(V_1)-p(V)/2| \leq 1$. Let $r_1=|R \cap V_1|$ and $b_1 = |B \cap V_1|$. We have $r_1+b_1=n/2=m$ and $t-1 \leq r_1-b_1 \leq t+1$. Therefore, $(m+t)/2-(1/2) \leq r_1 \leq (m+t)/2+(1/2)$. Since r_1 is an integer and $n_r=m+t$ is even, it follows that $r_1=(m+t)/2=n_r/2$. Hence, $b_1=(m-t)/2=n_b/2$. Therefore, V_2 also contains $n_r/2$ red nodes and $n_b/2$ blue nodes. □

If G is only 2-connected, we may not always get a perfect partition. Assume wlog that $n_r \leq n_b$. If for every $v \in R$ and $u \in B$, we set $p(v)=1$ and $p(u)=$

$-n_r/n_b$, Corollary 4.3 implies that there is always a connected partition (V_1, V_2) of V such that both $|(|R\cap V_1| - n_r/n_b|B\cap V_1|)| \leq 1$ and $|(|R\cap V_2| - n_r/n_b|B\cap V_2|)| \leq 1$, and also $\max\{\frac{|V_1|}{|V_2|}, \frac{|V_2|}{|V_1|}\} \leq 3$. Thus, the ratio of red to blue nodes in each side V_i differs from the ratio n_r/n_b in the whole graph by $O(1/n)$. Hence if the numbers of red and blue nodes are $\omega(1)$, then the two types are presented in both sides of the partition in approximately the same proportion as in the whole graph.

Corollary 5.2. *Given a 2-connected graph G with nodes colored either red ($R\subseteq V$) or blue ($B\subseteq V$), and assume wlog $|R|\leq|B|$. We can always find in polynomial time a partition (V_1, V_2) of V such that $G[V_1]$ and $G[V_2]$ are connected, $|V_1|, |V_2| \geq |V|/4$, and the ratio of red to blue nodes in each side V_i differs from the ratio $|R|/|B|$ in the whole graph by $O(1/n)$.*

6 Conclusion

In this paper, we introduced and studied the problem of partitioning a graph into two connected subgraphs that satisfy simultaneouly two objectives: (1) they balance the supply and demand within each side of the partition (or more generally, for the case of $p(V)\neq 0$, they split approximately equally the excess supply/demand between the two sides), and (2) the two sides are large and have roughly comparable size (they are both $\Omega(|V|)$). We showed that for 2-connected graphs it is always possible to achieve both objectives at the same time, and for 3-connected graphs there is a partition that is essentially perfectly balanced in both objectives. Furthermore, these partitions can be computed in polynomial time. This is a paradigmatic bi-objective balancing problem. We observed how it can be easily used to find a connected partition of a graph with two types of nodes that is balanced with respect to the sizes of both types. Overall, we believe that the novel techniques used in this paper can be applied to partitioning heterogeneous networks in various contexts.

There are several interesting further directions that suggest themselves. First, extend the theory and algorithms to find doubly balanced connected partitions to more than two parts. Second, even considering only the supply/demand objective, does the analogue of the results of Lováz and Gyori [9, 13] for the connected k-way partitioning of k-connected graphs with respect to size (which corresponds to $p(i)=1$) extend to the supply/demand case ($p(i)=\pm 1$) for $k>3$? And is there a polynomial algorithm that constructs such a partition? Finally, extend the results of Section 5 to graphs with more than two types of nodes, that is, can we partition (under suitable conditions) a graph with several types of nodes to two (or more) large connected subgraphs that preserve approximately the diversity (the proportions of the types) of the whole population?

References

[1] D. Bienstock. *Electrical Transmission System Cascades and Vulnerability: An Operations Research Viewpoint*, volume 22. SIAM, 2016.

[2] J. Bondy and U. Murty. Graph theory, 2nd printing, 2008.

[3] F. Chataigner, L. R. Salgado, and Y. Wakabayashi. Approximation and inapproximability results on balanced connected partitions of graphs. *DMTCS*, 9(1):177–192, 2007.

[4] J. Cheriyan and S. Maheshwari. Finding nonseparating induced cycles and independent spanning trees in 3-connected graphs. *J. Algorithms*, 9(4):507–537, 1988.

[5] J. Chlebíková. Approximating the maximally balanced connected partition problem in graphs. *Inform. Process. Lett.*, 60(5):225–230, 1996.

[6] M. Dyer and A. Frieze. On the complexity of partitioning graphs into connected subgraphs. *Discrete Applied Math.*, 10:139–153, 1985.

[7] S. Even and R. E. Tarjan. Computing an st-numbering. *Theoretical Computer Science*, 2(3):339–344, 1976.

[8] N. Fan, D. Izraelevitz, F. Pan, P. M. Pardalos, and J. Wang. A mixed integer programming approach for optimal power grid intentional islanding. *Energy Systems*, 3(1):77–93, 2012.

[9] E. Gyori. On division of graphs to connected subgraphs. In *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976)*, volume 1, pages 485–494, 1976.

[10] J. E. Hopcroft and R. E. Tarjan. Dividing a graph into triconnected components. *SIAM J. Comput.*, 2(3):135–158, 1973.

[11] A. Lempel, S. Even, and I. Cederbaum. An algorithm for planarity testing of graphs. In *Theory of graphs: International symposium*, volume 67, pages 215–232. Gordon and Breach, New York, 1967.

[12] N. Linial, L. Lovasz, and A. Wigderson. Rubber bands, convex embeddings and graph connectivity. *Combinatorica*, 8(1):91–102, 1988.

[13] L. Lovász. A homology theory for spanning tress of a graph. *Acta Mathematica Hungarica*, 30(3-4):241–251, 1977.

[14] S. Pahwa, M. Youssef, P. Schumm, C. Scoglio, and N. Schulz. Optimal intentional islanding to enhance the robustness of power grid networks. *Physica A*, 392(17):3741–3754, 2013.

[15] R. J. Sánchez-García, M. Fennelly, S. Norris, N. Wright, G. Niblo, J. Brodzki, and J. W. Bialek. Hierarchical spectral clustering of power

grids. *IEEE Trans. Power Syst.*, 29(5):2229–2237, 2014.

[16] K. Sun, Q. Zhao, D.-Z. Zheng, J. Ma, and Q. Lu. A two-phase method based on obdd for searching for splitting strategies of large-scale power systems. In *Proc. IEEE PowerCon'02*, 2002.

[17] H. Suzuki, N. Takahashi, and T. Nishizeki. A linear algorithm for bipartition of biconnected graphs. *Inform. Process. Lett.*, 33(5):227–231, 1990.

[18] W. T. Tutte. How to draw a graph. *Proc. London Math. Soc*, 13(3):743–768, 1963.

[19] K. Wada and K. Kawaguchi. Efficient algorithms for tripartitioning triconnected graphs and 3-edge-connected graphs. In *Graph-Theoretic Concepts in Computer Science*, pages 132–143. Springer, 1994.

[20] D. B. West et al. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River, 2001.