Upper Degree-Constrained Partial Orientations

Harold N. Gabow *

July 15, 2005

Abstract

We wish to orient as many edges as possible in an undirected graph, subject to upper bounds on the in-degree and out-degree of each vertex. Frank and Gyárfás [2] solve this problem in polynomial time when there are no in-degree bounds, and when every edge can be oriented within the given bounds. However we show that in general the problem is MAXSNP-hard. When viewed as a 3-dimensional matching problem the local improvement algorithm of Hurkens and Schrijver [4] achieves approximation ratio $2/3 - \epsilon$; we believe is the best previous bound for our problem. We give an LP-rounding algorithm that achieves approximation ratio $3/4$.

1 Introduction

An *orientation* of an undirected graph makes the graph directed by assigning a unique direction to each edge. A *partial orientation* assigns a unique direction to a subset of the edges, leaving the remaining edges unoriented. Consider an undirected graph G, with two integers $\mathrm{in}(v), \mathrm{out}(v)$ associated with each vertex v. (All graphs in this paper allow parallel edges.) We seek a partial orientation of G where each vertex v has in-degree at most $\mathrm{in}(v)$, out-degree at most $\mathrm{out}(v)$, and as many edges as possible are oriented.

This orientation problem can be viewed as the following matching problem. In a *symmetric bipartite graph* [6] $S = (V, E)$, the vertex set V consists of two copies of a set X, $V = \{x_0, x_1 : x \in X\}$; the edge set is symmetric, i.e., for every $x, y \in X$, $(x_0, y_1) \in E \iff (y_0, x_1) \in E$; and no $(x_0, y_1) \in E$ has $x = y$. Given such a graph we seek a maximum cardinality matching that contains at most one version of each edge (i.e., at most 1 of (x_0, y_1), (y_0, x_1) is matched). More generally the graph S is given with integers $u(v)$, $v \in V$, and we seek a subgraph containing at most one version of each edge, where each $v \in V$ has degree at most $u(v)$, and there are as many edges as possible.

This matching problem came to our attention as a subproblem in scheduling the University of Colorado High School Honors Institute. This program offers high school students a two day introduction to various engineering disciplines. Call the days Monday and Tuesday. Each engineering discipline presents two identical classes, one on each day, e.g., Computer Science Monday, Computer Science Tuesday. Each student specifies two favorite disciplines, e.g., Computer Science and Mechanical Engineering, in the hope of attending a class on each. The order of the classes is irrelevant, so our student can be scheduled in Computer Science Monday and Mechanical Engineering Tuesday, or vice versa. The number of students in each class is limited by room size. We wish to schedule as many students as possible. (A student must be scheduled in his two requested classes – he cannot attend just one class.)

*Department of Computer Science, University of Colorado at Boulder, Boulder, Colorado 80309-0430, USA. e-mail: hal@cs.colorado.edu

The matching formulation gives rise to a third view of our problem, as a 3-dimensional matching. For each edge pair $(x_0, y_1), (y_0, x_1)$ in the above symmetric graph S, introduce a vertex xy in the 3rd dimension, and place triplets (x_0, y_1, xy) and (y_0, x_1, xy) in the 3-dimensional matching problem. We seek a matching with the greatest number of triplets. The construction easily generalizes to arbitrary upper bounds $u(v)$.

Previous Work

Although we are unaware of work on our specific problem there is closely related work. The books of Bang-Jensen and Gutin [1] and Schrijver [7] have chapters on orientation, with Sections [1, 8.7] and [7, 61.1] devoted to degree-constrained orientation. The story begins with Hakimi [3] who characterized graphs that can be oriented so that every in-degree is at most a given upper bound. This was generalized in independent work by Frank and Gyárfás [2] to a characterization of graphs that can be oriented so that every in-degree is between given upper and lower bounds. The proofs of [2] give elegant augmenting path type algorithms for finding the desired orientations or showing they do not exist. Alternatively these theorems and algorithms can be derived from network flow theory.

[3] solves our problem when the out-degrees are unconstrained. [2] solves our problem when every edge can be oriented. (When every edge can be oriented in our problem, an upper bound on an out-degree is equivalent to a lower bound on the corresponding in-degree.) When all the upper bounds are 1, our problem is equivalent to finding a collection of paths and cycles with as many edges as possible; this can be done by matching techniques. We are unaware of characterizations or exact algorithms for other natural subcases of our problem.

The best approximation algorithm we could find results from viewing the problem as 3-dimensional matching. Hurkens and Schrijver [4] give an approximation algorithm for packing k-sets. We just take their general bound in the case $k = 3$. The result is an approximation algorithm that finds a 3-dimensional matching of size at least $2/3 - \epsilon$ times the maximum, for any $\epsilon > 0$. The algorithm is based on local search. The neighborhood size grows as ϵ approaches 0, and the running time is polynomial in $n^{\log 1/\epsilon}$.

Our Results

Our main result is an approximation algorithm achieving performance bound 3/4. The algorithm is based on iterated rounding – it repeatedly rounds a basic feasible solution to the natural LP corresponding to the problem. We also justify the approximation algorithm approach by showing the problem is MAXSNP-hard.

Section 2 illustrates our approach by giving a 5/7 approximation algorithm. Section 3 improves the approximation ratio to 3/4, at the expense of a much more involved algorithm. Section 4 proves MAXSNP-hardness; it also gives a simple upper bound of 4/5 on the integrality ratio of our LP.

We note that our results do not rule out the possibility of good combinatorial approximation algorithms (although all our attempts at such were essentially superseded by [4]) or better LP-rounding algorithms (our example for the integrality gap does not match our approximation factor).

An edge is denoted as (v, w) if it is undirected, and vw if it is directed. We use standard graph notation, e.g., $\rho(\cdot)$ and $\delta(\cdot)$ denote in-degree and out-degree respectively.

2 The Basic Algorithm

2.1 Approach

This subsection presents the overall strategy of the approximation algorithm. Consider our orientation problem on undirected graph $G = (V, E)$ with degree bounds $\text{in}(v), \text{out}(v)$. We use the natural formulation of the problem as an integer linear program. Each edge $(v, w) \in E$ has two variables x_{vw} and x_{wv}, with $x_{vw} = 1$ and $x_{wv} = 0$ when the edge is oriented from v to w.

$$\begin{array}{llcll}
\text{maximize} & \sum_{(v,w)\in E} x_{vw} + x_{wv} & & & \\
\text{subject to} & x_{vw} + x_{wv} & \leq & 1 & (v,w) \in E \\
 & \sum_{(v,w)\in E} x_{wv} & \leq & \text{in}(v) & v \in V \\
 & \sum_{(v,w)\in E} x_{vw} & \leq & \text{out}(v) & v \in V \\
 & x_{vw} & \in & \{0,1\} & (v,w) \in E
\end{array}$$

Let (LP) denote the linear relaxation of this program, i.e., we replace the fourth constraint by *nonnegativity constraints* $x_{vw} \geq 0$. We call the the first constraint of (LP) the *orientation constraint*; the next two are the *in-degree* and *out-degree constraints*, or together the *degree constraints*. We often refer to x_{vw} as the *weight* of directed edge vw.

Our algorithm will require a basic feasible solution (bfs) to (LP). Note that G is a multigraph, and (LP) can have superpolynomial size when edge multiplicities are large. This does not present a problem since an edge with multiplicity t can be represented by the constraint $x_{vw} + x_{wv} \leq t$. A basic feasible solution of this smaller problem gives a bfs to (LP).

Section 2.2 specifies a collection of rounding rules that apply to a bfs to (LP). Each rule converts several variables into 0-1 values, preserving feasibility to (LP) and preserving at least 5/7 of the rounded variables' collective weight.

This leads to the following algorithm, a simple version of iterated rounding: Find a basic feasible solution to (LP). Apply rounding rules to this solution until no rule is applicable. Then define the residual orientation problem, by discarding the edges that have been rounded and reducing the degree constraints for edges that have been set to 1. If no edges remain in the residual graph, we are done. Otherwise apply the process recursively on the residual orientation problem.

We will now prove this algorithm is the desired approximation algorithm, assuming two properties of the rounding rules. We begin with some precise definitions (that illuminate some flexibility in the rounding rules). A *rounding rule* assigns values $x_{vw} = 1, x_{wv} = 0$ to some nonempty set R of directed edges vw. The *residual problem* has edge set $\{(v, w) : vw, wv \notin R\}$ and degree constraints $\text{out}(v) - |\{(v, w) : vw \in R\}|$, $\text{in}(v) - |\{(w, v) : wv \in R\}|$. Clearly a feasible solution to the residual problem can be combined with the rounded values to get a feasible solution to the original problem.

Fix a parameter $\alpha \geq 1$. In an *α-rounding* the residual problem has objective value at most $\alpha|R|$ below the original objective value. An *α-rounding rule* always produces an α-rounding. A set of rounding rules is *complete* if at least one rule is applicable to any basic feasible solution.

Given a complete set of α-rounding rules, our algorithm finds an orientation of size at least $1/\alpha$ times optimum, by an argument similar to [5]: Let O be the optimum objective value of (LP). Suppose we round a set of variables $R \neq \emptyset$, by completeness. By induction the algorithm finds an orientation for the residual problem containing $\geq (O - \alpha|R|)/\alpha = O/\alpha - |R|$ edges. Combining this with the $|R|$ rounded edges gives an orientation for the original problem with at least O/α edges. Since the optimum orientation has at most O edges, we have a $1/\alpha$ approximation.

Also it is easy to see the algorithm works in polynomial time: A trivial 1-rounding rule is to "round" an edge whose weight in the given solution is 1. So edges with large multiplicities present no problem, and the algorithm can recur at most m times, for m the number of distinct edges.

Section 2.2 will give a set of $7/5$-rounding rules. Section 2.3 will prove the set is complete. This gives our $5/7$ approximation algorithm. Section 3 is similar.

We take advantage of some flexibility in the definition of α-rounding: To establish an α-rounding rule we are allowed to modify values in the current bfs, without necessarily rounding them to integral values (since we are interested in the objective of the linear residual problem, not the integral residual problem). This differs slightly from rounding in [5], where the analysis does not change any bfs values except the rounded ones.

We will now derive an inequality that is the core of our argument for completeness. Fix a bfs to (LP). Assume no edge has weight 1, since such variables can be trivially rounded. We begin by classifying the variables and constraints of (LP).

We concentrate on the variables that are positive, so define $P = \{vw : x_{vw} > 0\}$. Say that an undirected edge $(v,w) \in E$ is *saturated* if $x_{vw} + x_{wv} = 1$. Note that both variables are positive, by the above assumption. S denotes the set of all saturated edges, so $S \subseteq E$. $E - S$ is the set of all undirected edges that are not saturated. However it is most natural to concentrate on the directed versions of these edges. Thus define a set of directed edges

$$U = \{vw : vw \in P,\ (v,w) \notin S\}.$$

(The second condition means $x_{vw} + x_{wv} < 1$.) U can contain one or both orientations of an undirected edge.

Next we classify the degree constraints. As usual a constraint of (LP) is *tight* if equality holds. A bfs is the unique solution to a set of linearly independent constraints that are tight. Fix such a set of *basic* constraints. It is important to keep in mind that not every tight constraint need be basic.

We partition the basic degree constraints into sets $ID_i, OD_i, i \geq 0$. Intuitively OD_i (ID_i) denotes the basic out-degree (in-degree) constraints involving exactly i edges of U. The precise definition of OD_i, $i \geq 0$ is as follows (ID_i is symmetric):

$$OD_i = \{v : \text{the constraint } \textstyle\sum_{(v,w)\in E} x_{vw} = \mathrm{out}(v) \text{ is basic, and } \delta_U(v) = i\}.$$

Note that $\delta_U(v)$ is the number of edges in U directed from v. As an example $v \in OD_0$ means that v's out-degree constraint is basic, and every directed edge vw with positive weight has (v,w) saturated.

Let OD_i^+ denote $\bigcup_{j\geq i} OD_j$, and similarly for ID_i^+. Let D_i denote the multiset $ID_i \cup OD_i$. (A vertex may occur twice in D_i, corresponding to its two constraints.) Let D_i^+ denote the multiset $ID_i^+ \cup OD_i^+$. For example D_0^+ denotes the set of all basic degree constraints.

Counting degrees gives the following inequality:

$$(1) \qquad \begin{aligned} 2|U| &= \sum_v \rho_U(v) + \delta_U(v) \\ &\geq \sum_{v\in ID_1^+} \rho_U(v) + \sum_{v\in OD_1^+} \delta_U(v) \\ &= \sum_i i|D_i| \\ &\geq 2|D_2^+| + |D_1|. \end{aligned}$$

We can assume each variable that equals 0 corresponds to a basic nonnegativity constraint $x_{vw} = 0$. Similarly we can assume that each saturated edge (v,w) corresponds to a basic orientation constraint $x_{vw} + x_{wv} = 1$. (In proof, this collection of nonnegativity and orientation constraints is linearly independent, since any variable occurs in at most one of the constraints. So these linearly independent constraints can be enlarged to a linearly independent set that uniquely determines the bfs.)

Since the basic constraints are linearly independent and uniquely determine all the variables, the number of basic constraints equals the number of variables. This implies

$$2|S| + |U| = |S| + |D_0^+|,$$

since the left-hand side counts the number of positive basic variables and the right-hand side counts the number of basic constraints excluding nonnegativity. This relation gives $|U| = |D_0^+| - |S| = |D_2^+| + |D_1| + |D_0| - |S|$. Substituting in (1) gives

$$|D_1| \geq 2(|S| - |D_0|). \tag{2}$$

To establish completeness we will show (in Section 2.3) that if no rounding rule applies to a bfs then

- (i) equality holds in (2), and
- (ii) every vertex on a saturated edge belongs to $D_0 \cup (ID_1^+ \cap OD_1^+)$.

We claim this implies that in reality a rounding rule *must* be applicable. In the proof, for convenience we use the term "constraint" to refer to the left-hand side of a constraint (so a constraint is a sum of variables x_{vw}). For instance a nonnegativity constraint is just a variable x_{vw} which is 0 in the bfs.

To prove the claim first note that (i) implies equality holds in each step of the derivation of (2), in particular in the second line of (1). Thus

(i') every edge of U occurs in two basic degree constraints, one for each end.

Remark. In Section 3 (i) can fail but we prove (i') always holds. The rest of this section does not use (i), only (i').

Consider the relation

$$\sum_v \delta_P(v) = \sum_v \rho_P(v). \tag{3}$$

Recall P is the set of positive directed edges, which is nonempty. We write each in- or out-degree as a sum of the corresponding variables x_{vw}.

Consider a term $\delta_P(v)$. If $v \in OD_0^+$ then $\delta_P(v)$ is a basic out-degree constraint (minus some nonnegativity constraints for weight 0 edges leaving v). If $v \notin OD_0^+$ then $\delta_U(v) = 0$, by (i'). We can assume v is on a saturated edge (else $\delta_P(v)$ has no terms). So (ii) shows $v \in ID_0$. In particular $\rho_U(v) = 0$ and every positive edge incident to v is saturated. Thus we can write $\delta_P(v)$ as

$$\sum_{(v,w)\in S} x_{vw} + x_{wv} \quad - \quad \sum_{(v,w)\in S} x_{wv}$$

which is a linear combination of orientation, in-degree and nonnegativity constraints, all basic.

Thus we have shown that the left-hand side of (3) can be written as a linear combination of basic constraints. Similarly for the right-hand side. The resulting relation is nontrivial, i.e., a nontrivial linear combination of basic constraints sums to 0. (It is nontrivial since one or more degree-constraints have coefficient 1 or 2.) This contradicts the definition of the basic constraints. In summary we have shown that conditions (i)–(ii) imply the rounding rules are complete.

We close this section by giving some other structural properties of our sets. The first property sets the approach for the proof of completeness.

Lemma 2.1 *The set of edges S is acyclic.*

Proof: Let C be a cycle of saturated edges. Let C_1 and C_{-1} be the two directed cycles that result from orienting C consistently. Choose any $\epsilon > 0$ and increase x_{vw} by $i\epsilon$ for every $vw \in C_i$, $i = -1, 1$.

Every edge of C remains saturated, and no in-degree or out-degree changes. Hence every orientation and degree constraint that is basic remains tight. Since $C_1 \cup C_2 \subseteq P$, none of these edges has a basic nonnegativity constraint, and all basic nonnegativity constraints remain tight. In summary all basic constraints remain tight, contradicting uniqueness of the bfs. □

The families ID and OD are symmetric. So all our results have a symmetric version obtained by interchanging O and I. Say the bfs (along with its basic constraints) is *normal* if

(i) every edge weight is < 1;

(ii) the basic constraints include the orientation constraint of every saturated edge and the nonnegativity constraint of every weight zero edge.

We have already indicated that normalcy can be assumed. Throughout the rest of the paper we always assume a normal bfs. For instance we do not explicitly state this assumption in the hypotheses of our lemmas.

Lemma 2.2 *(i) $ID_0 \cap OD_0 = \emptyset$.*

(ii) $ID_0 \cap OD_1 = \emptyset$.

(iii) Suppose $v \in ID_1 \cap OD_1$, and let the two edges of U incident to v be vw and uv. Then $x_{vw} + x_{uv} = 1$ and $u \neq w$.

Proof: Suppose $v \in ID_i \cap OD_j$ for arbitrary i, j. Thus $\sum_{(v,y)\in E} x_{yv} + x_{vy} = \text{in}(v) + \text{out}(v)$. Let $T \subseteq S$ be the set of saturated edges incident to v. Thus $\sum_{(v,y)\in T} x_{yv} + x_{vy} = |T|$. Subtracting, and dropping the terms x_{yv} and x_{vy} that are 0, gives

$$\sum_{yv\in U} x_{yv} + \sum_{vy\in U} x_{vy} = \text{in}(v) + \text{out}(v) - |T| \in \mathbf{Z}. \tag{4}$$

(i) Take $i = j = 0$. So the left-hand side of (4) has no terms. In other words the degree and orientation constraints at v, plus some nonnegativity constraints for weight zero edges, are linearly dependent. Normalcy implies all these constraints are basic, a contradiction.

(ii) Take $i = 0$ and $j = 1$. The left-hand side of (4) has exactly one term (it has the form x_{vy}) and it is positive. Now (4) contradicts normalcy.

(iii) Taking $i = j = 1$, (4) becomes the desired equation. (The right-hand side is 1 by normalcy.) If $u = w$ this equation makes (v, w) saturated, contradicting $vw \in U$. □

The next lemma provides the link to proving equality in (2), by giving a combinatorial interpretation of the right-hand side.

Lemma 2.3 *Let T be a tree of the forest S. Suppose T contains s edges, and its vertices have exactly d_0 constraints of D_0. Then T contains exactly $s - d_0 + 1$ vertices not in D_0.*

Proof: Part (i) of Lemma 2.2 shows T contains exactly d_0 vertices of D_0. Since T contains exactly $s+1$ vertices, the lemma follows. $\square$

2.2 Rounding Rules

Recall that R is the set of all directed edges vw that get rounded up, i.e., vw gets added to the algorithm's solution and we set $x_{vw} = 1$, $x_{wv} = 0$. We use Δ to denote the *cost* of a rounding operation, i.e., the amount by which the objective decreases when passing from the original problem to the residual problem for R. So an α-rounding has

$$\Delta \leq \alpha |R|.$$

Note that the cost always includes the sum $\sum_{vw \in R} x_{vw} + x_{wv}$. In general the cost includes other terms for variables decreased to ensure feasibility. It is also convenient to introduce the short-hand notation that, for any real value x, $\bar{x}$ denotes the value $1 - x$.

To illustrate these notions we show that any edge vw can be $(2 - x_{vw} + x_{wv})$-rounded. The round increases x_{vw} to 1 and decreases x_{wv} to 0. Since x_{vw} increases by $\bar{x}_{vw}$, two degree constraints can become violated by an amount $\leq \bar{x}_{vw}$. To satisfy the in-degree constraint of w we decrease values x_{zw}, $z \neq v$, by a total amount $\leq \bar{x}_{vw}$. (The decrease may be distributed over a number of edges zw.) Similarly we decrease values x_{vz}, $z \neq w$, by $\leq \bar{x}_{vw}$. So the cost is

$$\Delta \leq (x_{vw} + x_{wv}) + 2\bar{x}_{vw} = 2 + x_{wv} - x_{vw}.$$

Since we are rounding 1 edge, this gives the claimed rounding coefficient.

Call an edge $vw \in U$ a *singleton* if $wv \notin U$, i.e., $x_{wv} = 0$. vw is a *slack singleton* if in addition w's in-degree constraint (or v's out-degree constraint) has slack $\geq \bar{x}_{vw}$. We have this useful property:

Lemma 2.4 *Rounding a slack singleton costs* ≤ 1.

Proof: We pay x_{vw} to delete vw when forming the residual graph, The hypothesis implies there is no cost to preserve say the in-degree constraint at w. We pay $\leq \bar{x}_{vw}$ to preserve the out-degree constraint at v. $\square$

Rounding a directed path P means to round up P's edges, i.e., R is the set of edges of P. We always assume P is simple, but we allow P to be a simple cycle of length ≥ 3. (It is not meaningful to round a length 2 cycle.) The following lemma treats common cases; more complicated versions are possible.

Lemma 2.5 *Let* $P = v_0, v_1, \ldots, v_k$ *be a directed path or cycle and let* x_i *be the weight of edge* $v_{i-1}v_i$, $1 \leq i \leq k$. *Let every edge of* P *be saturated except possibly the first or last (or both).*

(i) If $x_1 \leq x_2 \leq \ldots \leq x_k$ *then the cost of rounding* P *is* $\leq k + 2\bar{x}_1$.

(ii) If $x_1 \geq x_2 \geq \ldots \geq x_k$ *then the cost of rounding* P *is* $\leq k + 2\bar{x}_k$.

(iii) In (i) the cost is $\leq k + 2\bar{x}_1 - \min\{\text{slack\_out}(v_0), \bar{x}_1\} - \min\{\text{slack\_in}(v_k), \bar{x}_k\}$, *where* $\text{slack\_out}(v_0)$ *is the slack in the constraint for* $\text{out}(v_0)$ *and similarly for* $\text{slack\_in}(v_k)$.

(iv) In (i) the cost is $\leq k + 2\bar{x}_1 - \bar{x}_k$ *if* $k > 1$ *and edge* v_0v_k *exists and has weight* $\geq \bar{x}_k$.

Proof: We note that $v_k = v_0$ when P is a simple cycle. We will show that for each $i \leq k$, rounding path $v_0, \ldots, v_i$ gives a residual problem having a solution that has cost $\leq i + 2\bar{x}_1$ and, when $i < k$,

(a) the constraint for $\text{out}(v_i)$ has slack $\geq \bar{x}_i$, and

(b) $v_{i+1}v_i$ has weight 0.

Clearly the case $i = k$ gives the lemma. The proof is by induction on i.

Suppose $i = 1$. The cost of discarding (v_0, v_1) and decreasing the weight of additional edges to achieve (i) is ≤ 1. To see this note that if $(v_0, v_1) \in S$ the discard costs 1 and it achieves (a). If $(v_0, v_1) \notin S$ it costs x_1 to discard v_0v_1, and a charge of $\leq \bar{x}_1$ accounts for voiding v_1v_0 and decreasing other edges to achieve (a).

In addition it costs $\leq \bar{x}_1$ to maintain the out-degree of v_0, and $\leq \bar{x}_1$ to maintain the in-degree of v_1. We can assume the latter voids edge v_2v_1, since $\bar{x}_1 \geq \bar{x}_2$ by hypothesis. So (b) holds.

Now assume the inductive assertion for $i-1$ and prove it for i as follows. (a)–(b) for $i-1$ along with $\bar{x}_{i-1} \geq \bar{x}_i$ show that $v_{i-1}v_i$ has become a slack singleton. Thus rounding it increases the cost by 1 as desired. Note this argument is unaffected if $i = k$ and $v_k = v_0$.

Now we show (a)–(b) if $i < k$: (a) follows since we have decreased the weight of v_iv_{i-1} from $\bar{x}_i$ to 0. (We are using $(v_{i-1}, v_i) \in S$ and the simplicity of P.) For (b) we paid $\leq \bar{x}_i$ to maintain the in-degree of v_i. We can assume this voids edge $v_{i+1}v_i$ since $\bar{x}_i \geq \bar{x}_{i+1}$ by hypothesis. (We are not assuming $(v_i, v_{i+1}) \in S$, which can fail for $i = k-1$.)

(ii) The argument is symmetric to (i). We show that for each $i < k$, rounding path $v_i, \ldots, v_k$ gives a residual problem having a solution that has cost $\leq k - i + 2\bar{x}_k - \min\{\sigma, \bar{x}_k\}$ and, when $i > 0$,

(a) the constraint for $\text{in}(v_i)$ has slack $\geq \bar{x}_{i+1}$, and

(b) v_iv_{i-1} has weight 0.

(iii) In (i) the charge of $\bar{x}_1$ to maintain the out-degree of v_0 can be decreased to $\bar{x}_1 - \min\{\text{slack\_out}(v_0), \bar{x}_1\}$. Similarly the charge of $\bar{x}_k$ to maintain the in-degree of v_k can be decreased to $\bar{x}_k - \min\{\text{slack\_in}(v_k), \bar{x}_k\}$.

(iv) Construct the solution to the residual problem in two steps. First, decrease the weight of v_0v_k by $\bar{x}_k$. This costs $\bar{x}_k$ and gives a feasible solution to the original problem with slack $\geq \bar{x}_k$ in the out-degree of v_0 and the in-degree of v_k. Then apply (iii). Since $\bar{x}_1 \geq \bar{x}_k$ the total cost is $\bar{x}_k + (k + 2\bar{x}_1 - 2\bar{x}_k) = k + 2\bar{x}_1 - \bar{x}_k$. □

One way we use the lemma is the following. Define the *co-weight* of edge vw as $\bar{x}_{vw}$. We'll see that the hypothesis on the first edge's co-weight gets satisfied because of appropriate singletons.

Corollary 2.6 *Let P be a path as in Lemma 2.5.*

(i) Rounding P is a 7/5-rounding (4/3-rounding) if P has length 3 (4), its edge weights are nondecreasing, and P starts with an edge of co-weight $\leq 3/5$ ($\leq 2/3$).

(ii) Rounding P is a 7/5-rounding (4/3-rounding) if P has length 3 (4), its edge weights are nonincreasing, and P ends with an edge of co-weight $\leq 3/5$ ($\leq 2/3$).

Proof: For both parts, the length 3 paths cost $\leq 3 + 6/5 = (7/5)3$. The length 4 paths cost $\leq 4 + 4/3 = (4/3)4$. □

We partition our set OD_i (ID_i) into sets OD_{ij} (ID_{ij}), $j \geq 0$. We just describe the O family, the I family is symmetric. Intuitively OD_{ij} denotes the basic out-degree constraints involving exactly i edges of U and j edges of S. In precise terms, for $i, j \geq 0$,

$$OD_{ij} = \{v : \text{the constraint } \textstyle\sum_{(v,w)\in E} x_{vw} = \text{out}(v) \text{ is basic, and } \delta_U(v) = i,\ d_S(v) = j\}.$$

Here $d_S(v)$ denotes the number of edges of S incident to v. We extend our other notational conventions to these sets, e.g., $D_{ij} = ID_{ij} \cup OD_{ij}$.

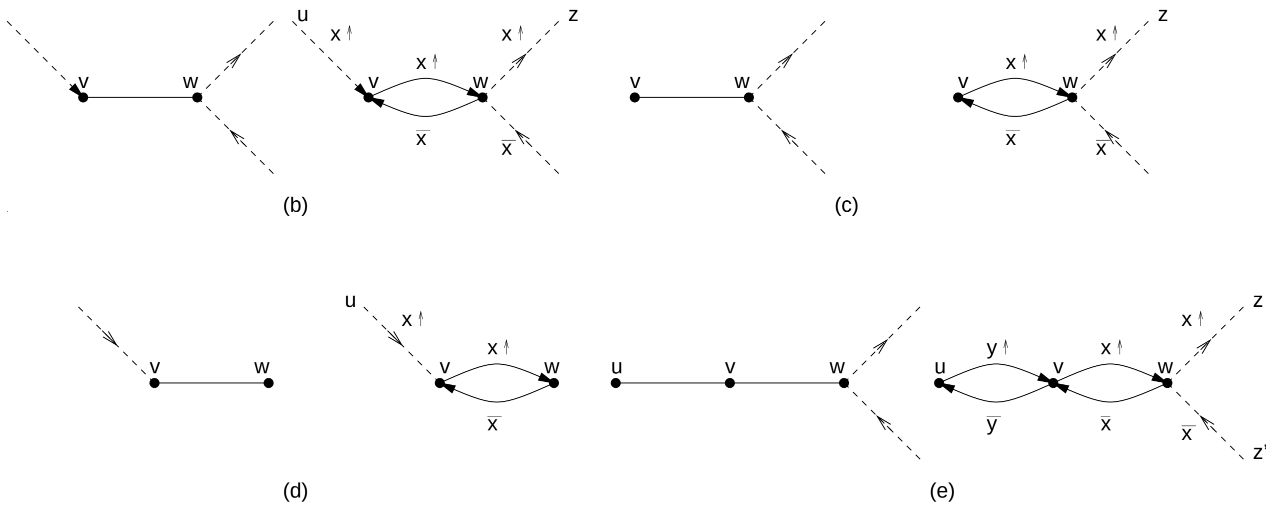

Figure 1: Rounding rules (b)–(e). Each rule is illustrated in two drawings. The first drawing illustrates the subgraph to which the rule applies. The second drawing shows the edge weights before the rounding. Weights that get rounded up to 1 are marked with an uparrow (in (e) just the 2nd case is illustrated). An edge is drawn solid or dashed, depending on whether it is saturated or unsaturated respectively. An arrowhead in the middle of an unsaturated edge indicates that the edge is a singleton.

Lemma 2.7 *(i) A vertex $v \in ID_1 \cap OD_1$ has two edges of U incident to v, say uv and vw. Both edges are singletons, and $x_{vw} = \bar{x}_{uv}$.*

(ii) A vertex $v \in OD_{11}$ has two positive edges directed from v, say vw and vu with $vw \in U$, $(v,u) \in S$. We have $x_{vw} = \bar{x}_{vu}$. Similarly $v \in ID_{11}$ has two positive edges directed to v, $wv \in U$, uv with $(u,v) \in S$, and $x_{wv} = \bar{x}_{uv}$.

(iii) A vertex $v \in ID_{11} \cap OD_{11}$ has four positive edges incident to v, say vw, wv with $(v,w) \in S$, and $uv, vz \in U$. The last two edges are singletons, and $x_{vz} = x_{wv} = \bar{x}_{vw}$, $x_{uv} = x_{vw}$.

Proof: (i) just restates Lemma 2.2(iii).

(ii) A vertex $v \in OD_{11}$ has a basic out-degree constraint which, if we drop all terms x_{vz} that equal 0 in our bfs, becomes $x_{vw} + x_{vu} = 1$ for edges $vw \in U$, $(v,u) \in S$. We know the right-hand side value is 1 by normalcy of the bfs. Thus $x_{vw} = \bar{x}_{vu}$.

(iii) follows by combining (i) and (ii). □

The following lemma presents our rounding rules. Part (a) is a consequence of the preceding calculation for rounding 1 edge. The remaining rounding rules are illustrated in Fig.1. Each rounding rule is a separate case (b)–(e) of the lemma. The rounding rule is illustrated twice in the corresponding part of Fig.1(b)–(e). The first drawing illustrates the definition of when the rule applies. The second drawing indicates the weights of the various directed edges and how they get rounded.

Lemma 2.8 *A 7/5-rounding can be done in each of the following cases:*

(a) Edge vw is a singleton with $x_{vw} \geq 3/5$, or edge $(v,w) \in S$ has $x_{vw} \geq 4/5$.

(b) Edge $(v,w) \in S$, $v \in ID_{11}$, $w \in ID_{11} \cap OD_{11}$.

(c) Edge $(v,w) \in S$, $w \in ID_{11} \cap OD_{11}$, and $\delta_P(v) = \{vw\}$.

(d) Edge $(v,w) \in S$, $v \in ID_{11}$, and $\delta_P(v) = \{vw\}$.
(e) Edges $(u,v),(v,w) \in S$, $u \neq w$ and $w \in ID_{11} \cap OD_{11}$.

Proof: (a) follows from the preceding calculation for rounding one edge. Note that if (v,w) is saturated, $x_{vw} \geq 4/5$ implies $x_{wv} \leq 1/5$.

Some cases of our lemma actually have a 4/3-rounding, used in Section 3. To prove them we use this version of (a):

(a′) A 4/3-rounding can be done if vw is a singleton with $x_{vw} \geq 2/3$, or vw has $x_{vw} - x_{wv} \geq 2/3$.

(b) Fig.1(b) gives the edge weights, shortening x_{vw} to x. (All the other figures are similar.) Note that edge uv need not be a singleton – this in fact governs our rounding strategy. The weights of the figure follow from the assumptions of case (b) and Lemma 2.7. It also follows that all 5 edges, plus vu if it exists, are distinct. It is possible that $u = z$.

We round path u, v, w, z. (This can be a cycle.) Since there is a singleton with weight $\bar{x}$ we can assume $\bar{x} < 3/5$, by (a). Now Corollary 2.6(i) shows we have a 7/5-rounding.

(c) The weights in Fig.1(c) follow as usual. All 4 edges of the figure are distinct.

Round path v, w, z. Lemma 2.5(iii) shows $\Delta \leq 2 + \bar{x}$, since vw is the only positive edge leaving v. By (a′), $\bar{x} < 2/3$. Thus $\Delta \leq 2 + 2/3 = (4/3) \times 2 = (4/3)|R|$.

(d) The weights in Fig.1(d) follow as usual. All 3 edges of the figure are distinct.

Round path u, v, w. We have

$$\Delta \leq (1 + \bar{x}) + 1.$$

The first term is for rounding vw; there is no charge for preserving the out-degree of v, since vw is the only positive edge leaving v. The second term is for rounding uv, which the first round makes a slack singleton.

Thus $\Delta \leq 2 + \bar{x}$. By (a), $\bar{x} < 4/5$. Thus $\Delta \leq 2 + 4/5 = (7/5) \times 2 = (7/5)|R|$.

(e) The weights in Fig.1(e) follow as usual. All edges in the figure are distinct, although it is possible that u coincides with z or z'. We consider two cases depending on the comparison between y and x.

Case 1. $y \leq x$ (and $\bar{y} \geq \bar{x}$). Round the right-to-left path (not marked in the figure) z', w, v, u. (a) implies $x < 3/5$ (singleton wz). So Corollary 2.6(i) shows this is a 7/5-rounding.

Case 2. $y \geq x$ (and $\bar{y} \leq \bar{x}$). Round the left-to-right path shown in the figure, u, v, w, z. As in Case 1, (a) implies $\bar{x} < 3/5$ (singleton $z'w$), and Corollary 2.6(ii) shows this is a 7/5-rounding. □

2.3 Completeness

We will prove the rounding rules given by Lemma 2.8*(a)–(e)* are complete. For succinctness we refer to the rules as (a)–(e).

We begin with some general results. Throughout this section T is a tree of the forest S.

Lemma 2.9 *A leaf of T known to belong to $D_0 \cup (ID_1 \cap OD_1)$ actually belongs to $ID_{11} \cap OD_{11}$.*

Proof: A vertex $v \in D_0$ has $d_S(v) > 1$, by normalcy. So the lemma's hypothesis puts $v \in ID_1 \cap OD_1$. Since leaf v is on 1 saturated edge, $v \in ID_{11} \cap OD_{11}$ as desired. $\square$

The proof of the next lemma roots the tree T. By convention the root is not considered to be a leaf of an oriented tree.

Lemma 2.10 *Rounding rule (e) can be applied if T has ≥ 2 edges and some vertex β has $V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta\}$.*

Proof: Choose β as the root of T. Lemma 2.9 shows that every leaf of T belongs to $ID_{11} \cap OD_{11}$. Hence rule (e) applies if T has a length 2 path ending at a leaf. This condition holds if T has height > 1 or the root has degree $d_S(\beta) > 1$. Since we assume T has ≥ 2 edges, one of the last two conditions holds. $\square$

As mentioned in Section 2.1 the proof of completeness is in two steps. We first show that if no rounding rule is applicable then (2) holds with equality. To accomplish this, we continue with the notation of Lemma 2.3: Suppose tree T of S contains $s \geq 0$ edges, with its vertices having exactly d_i constraints of D_i. We will prove

$$d_1 \leq 2(s - d_0). \tag{5}$$

Summing (5) over every tree of the forest gives $|D_1| \leq 2(|S| - |D_0|)$ (since we consider vertices with no S edges as trees). Since this is the opposite inequality from (2) we conclude that (2) holds with equality, as desired.

To prove (5) first suppose $s = 0$, i.e., we have a vertex v on no saturated edges. By definition $v \notin D_0$. Also $v \notin D_1$, since we assume the bfs is normal. Thus $d_1 = d_0 = s = 0$ and (5) holds.

Now suppose $s > 0$. A vertex of D_0 is not in D_1 (Lemma 2.2(ii)). Lemma 2.3 shows T contains exactly $s - d_0 + 1$ vertices not in D_0. Each of these can be in $ID_1 \cap OD_1$, so $d_1 \leq 2(s - d_0) + 2$. So (5) follows if we eliminate the possibility of exactly 0 or 1 "missing" ID_1/OD_1 constraints. More precisely it suffices show that under the assumption

(A) $\quad T$ contains a vertex $\beta \in ID_1 \cup OD_1$ with $V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta\}$,

some rounding rule is applicable.

Choose β as the root of T. Suppose T has ≥ 2 edges. Assumption (A) guarantees the hypothesis of Lemma 2.10. Hence a rounding rule can be applied.

The remaining case is that T consists of a single edge. From (A) and symmetry we can assume $\beta \in ID_1$. So actually $\beta \in ID_{11}$. The other vertex w of T belongs to $ID_{11} \cap OD_{11}$ (Lemma 2.9). Now rounding rule (b) applies to T.

This completes the first step of the proof. Before moving on let us summarize some consequences of this step. Section 2.1 shows that the first step implies

(i') every edge of U occurs in two basic degree constraints, one for each end.

In addition we have proved that each tree T of S satisfies (5) with equality. Recall that T has $s - d_0 + 1$ vertices not in D_0. So there are exactly 2 "missing" ID_1/OD_1 constraints. More precisely we have shown that for two vertices β, β' in $T - D_0$,

$$V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta, \beta'\}$$

and either

(a) $\beta' \in ID_1 \cap OD_1$ and $\beta \notin ID_1 \cup OD_1$, or

(b) $\beta, \beta' \in ID_1 \oplus OD_1$.

We turn to the second step of the proof. We must show that if no rounding rule is applicable,

(ii) $V(S) \subseteq D_0 \cup (ID_1^+ \cap OD_1^+)$.

We will prove the contrapositive: If (ii) fails then some rounding rule is applicable. A vertex failing (ii) is certainly not in $D_0 \cup (ID_1 \cap OD_1)$, so it must be β or β' in the above characterization of the vertices of T (for some tree T of S). So we can consider the two possibilities of the characterization, (a) and (b).

In possibility (a), $V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta\}$. Lemma 2.10 shows a rounding rule can be applied unless T has exactly 1 edge. So consider that case. The vertex failing (ii) is β, so $\beta \notin ID_1^+ \cap OD_1^+$. By symmetry assume $\beta \notin OD_1^+$. Property (i') implies no edge of U leaves β, i.e., exactly 1 edge of P leaves β. Now rounding rule (c) applies to the edge of T (Lemma 2.9 ensures the rest of the hypothesis).

Now we can assume possibility (b). Vertex β fails (ii), and β' may or may not fail. By I/O symmetry assume (b) gives $\beta \in ID_1$. Taking the failure of (ii) into account and also (i'), we get

$\beta \in ID_1 - OD_0^+$ and $\delta_U(\beta) = 0$, or β satisfies the I/O symmetric condition.

We show a rounding rule applies in possibility (b). Root T at β. If β has ≥ 2 children then T has a leaf $w \neq \beta'$, and there is a length 2 path ending at w. Lemma 2.9 shows $w \in ID_{11} \cap OD_{11}$. So rounding rule (e) applies to the path.

If β has exactly 1 child, then $\beta \in ID_{11}$ with $\delta_U(\beta) = 0$, or an I/O symmetric condition holds. Rounding rule (d) applies.

3 The Refined Algorithm

This section sketches how to refine our algorithm to achieve the performance bound of 4/3. We reduce the rounding coefficient by incorporating more context into the rounding rules. This comes at the expense of a large number of cases! Yet one rule seems to resist this approach: We could not achieve a 4/3-rounding for the entire case (b) of Lemma 2.8, 1-edge trees. Our failed efforts make us believe that no such rounding exists. Instead, we give a 4/3-rounding rule for a subclass of the case (b) 1-edge trees. To prove the rules are still complete necessitates modifying the high-level approach of Section 2.1. This section briefly covers the new rounding rules and the modified high-level approach.

We start by defining the case (b) trees that are not amenable to rounding. Let (v, w) be a case (b) edge, i.e., $(v, w) \in S$, $v \in ID_{11}$, $w \in ID_{11} \cap OD_{11}$. As in Fig.1(b) let uv be the unique edge of U directed to v. Call (v, w) a *special tree* of S if either

$\delta_U(v) > 2$, or

$\delta_U(v) = 2$ and $v \notin OD_2$, or

$x_{uv}, x_{vu} > 1/3$ and $u \notin V(S) \cup D_2$.

3.1 New Rounding Rules

Here is the complete set of rounding rules.

Lemma 3.1 *A 4/3-rounding can be done in each of the following cases:*

(a) Edge vw is a singleton with $x_{vw} \geq 2/3$, or edge vw has $x_{vw} - x_{wv} \geq 2/3$.

(b) Edge $(v, w) \in S$, $v \in ID_{11}$, $w \in ID_{11} \cap OD_{11}$, but (v, w) is not special.

(c) Edge $(v, w) \in S$, $w \in ID_{11} \cap OD_{11}$, and $\delta_P(v) = \{vw\}$.

(d) Edge $(v, w) \in S$, $v \in ID_{11}$, and $\delta_P(v) = \{vw\}$.

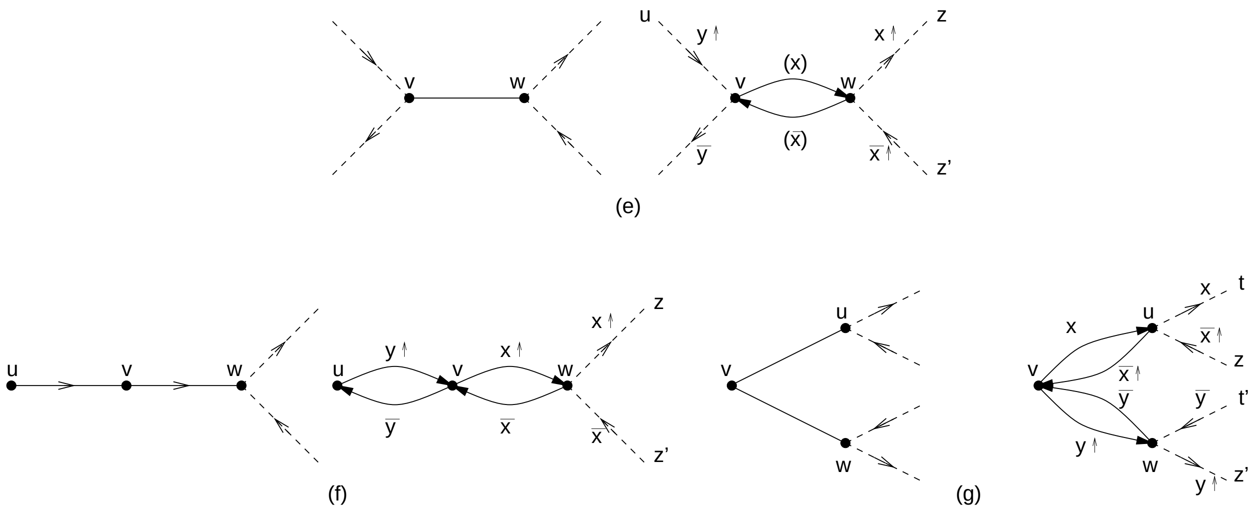

Figure 2: Rules (e)–(g). These additional conventions are used: Weights that get dropped are parenthesized. An arrowhead in the middle of a saturated edge gives the direction in which the edge is heavy.

(e) Edge $(v,w) \in S$, $v \in ID_1 \cap OD_1$, $w \in ID_{11} \cap OD_{11}$.

(f) Edges $(u,v),(v,w) \in S$, $u \neq w$, $w \in ID_{11} \cap OD_{11}$ *and* $x_{uv} \geq x_{vu}$, $x_{vw} \geq x_{wv}$, *or symmetrically,* $x_{vu} \geq x_{uv}$, $x_{wv} \geq x_{vw}$. *In particular one set of these inequalities holds when* $v \in D_{02}$.

(g) Edges $(u,v),(v,w) \in S$, $u \neq w$, $u,w \in ID_{11} \cap OD_{11}$.

(h) Edges $(u,v),(v,w) \in S$, $u \neq w$, $w \in ID_{11} \cap OD_{11}$, $v \in ID_{12}$, $\delta_P(v) = \{vw, vu\}$, *and either* $d_S(u) \geq 2$ *or* $u \in D_1$.

Comment. Single edges are treated in (a)–(e). (a) and (c) are (a′) and (c) of Lemma 2.8 respectively, where a 4/3-rounding was already achieved. (b) has been discussed. (d) is the same case as (d) of Lemma 2.8; the new proof is slightly more involved. (e) is a new case with simple proof.

The double edge case (e) of Lemma 2.8 expands into cases (f)–(h) (although these rules do not cover all the cases that (e) did). Here (f) and (g) have simple proofs. (h) is proved by combining a large number of rules.

The proofs use another rounding notion. To *drop* edge vw means to set $x_{vw} = 0$. This does not change the value x_{wv}. A drop is only used to construct a feasible solution to the residual problem in order to bound Δ.

We begin with 3 simple proofs for rules (e)–(g), illustrated in Fig.2.

Proof of (e): The weights of Fig.2(e) follow from the assumptions of case (e) and Lemma 2.7. All 6 edges are distinct. Some of the vertices can coincide, e.g.,u can be z or z'. In the following argument assume $y \geq x$ ($\bar{y} \leq \bar{x}$); the opposite case is symmetric.

Case 1. $u \neq z'$. Take $R = \{uv, z'w, wz\}$. To compute Δ, drop edges vw and wv. We get

$$\Delta \leq 1 + 1 + 1 + 1.$$

The first term is the cost of the drop. The next two terms are for rounding $z'w$ and wz, which the drop makes slack singletons. The last term is for rounding uv: The drop makes it a slack singleton since $\bar{x} \geq \bar{y}$. The assumption of Case 1 ensures uv is not affected by the round of $z'w$. We conclude $\Delta \leq 4 = (4/3)|R|$.

Case 2. $u = z'$. Round the path z', w, v. Lemma 2.5(iv) applies since edge $uv = z'v$ has weight $y \geq x$. Thus $\Delta \leq 2 + x$. Since a singleton has weight x (a) gives $\Delta \leq 2 + 2/3 = (4/3)|R|$. □

Proof of (f): The weights in Fig.2(f) follow from the assumptions of case (f) and Lemma 2.7. All 6 edges of the figure are distinct. Possibly $u = z$.

Case 1. $y \geq \bar{y}$, $x \geq \bar{x}$. Round the path u, v, w, z.

If $y \leq x$, Lemma 2.5(i) shows $\Delta \leq 3 + 2\bar{y}$. If $y \geq x$, Lemma 2.5(ii) shows $\Delta \leq 3 + 2\bar{x}$. Case 1 makes $\bar{y}, \bar{x} \leq 1/2$. Thus in both cases $\Delta \leq 3 + 1 = (4/3) \times 3 = (4/3)|R|$.

Case 2. $\bar{y} \geq y$, $\bar{x} \geq x$. This case is symmetric to Case 1. Round path z', w, v, u. If $\bar{x} \leq \bar{y}$, Lemma 2.5(i) shows $\Delta \leq 3 + 2x$. If $\bar{x} \geq \bar{y}$, Lemma 2.5(ii) shows $\Delta \leq 3 + 2y$.

Finally note that $v \in OD_{02}$ implies $\bar{y} + x = 1$, i.e., $y = x$. So one of the two pairs of inequalities in the hypothesis holds when $v \in OD_{02}$. □

Proof of (g): The weights in Fig.2(g) follow as usual. It is possible that two singletons are identical: we can have $uw = ut = t'w$ or $wu = zu = wz'$. But Lemma 2.2(iii) shows both alternatives cannot hold simultaneously. So *wlog* assume $z \neq w, z' \neq u$. We can still have $z = z'$.

Round path z, u, v, w, z'. Our assumption makes this a simple path or cycle. Since there are singletons of weight $x, \bar{y}$ we can assume $x, \bar{y} < 2/3$ by (a$'$). So we have a 4/3-rounding by Corollary 2.6 part (i) if $\bar{x} \leq y$, and part (ii) if $\bar{x} \geq y$. □

Proof of (d): As in the proof of Lemma 2.8(d) and Fig.1(d), rounding path u, v, w gives $\Delta \leq 2 + \bar{x}$. If $\bar{x} \leq 2/3$ we are done, so assume

$$\bar{x} > 2/3.$$

We can also assume there is an edge (w, z), $z \neq v$. For otherwise the cost of rounding u, v, w decreases by $\bar{x}$ to 2 and again we are done. We will show rounding path u, v, w, z gives $\Delta \leq 4$, thus achieving the desired 4/3-rounding. Let $x_{wz} = y$.

Suppose $y \geq x$ ($\bar{y} \leq \bar{x}$). The round of vw makes wz a slack singleton, since it voids wv and we can assume it voids zw. Thus $\Delta \leq 3 + \bar{x} < 4$.

Suppose $y \leq x$ ($\bar{y} \geq \bar{x}$). Let $x_{zw} = \alpha$. We can assume

$$\alpha - y < 2/3$$

by (a). Bound the cost of rounding u, v, w, z by increasing the cost of u, v, w $(2 + \bar{x})$ by the cost of rounding wz given by the 3 terms shown:

$$\Delta \leq (2 + \bar{x}) + (\bar{y} - \bar{x}) + 1 + \max\{\alpha - \bar{x}, 0\} = 4 - y + \max\{\alpha - \bar{x}, 0\}$$

The last expression is ≤ 4 since $-y + \alpha - \bar{x} < 2/3 - 2/3 = 0$. □

Proof of (b): Fig.3 illustrates the setting of the rule and the additional context that will get introduced. For now note that α denotes x_{vu}, where we know

$$0 \leq \alpha < \bar{x} \text{ and } \bar{\alpha} > x.$$

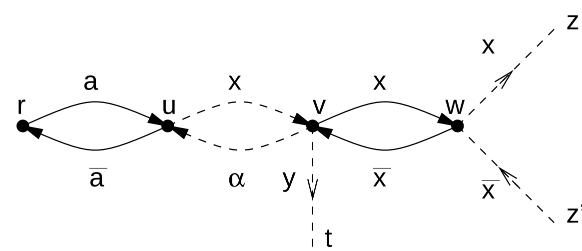

Figure 3: Rounding rule (b) for nonspecial trees.

Vertices r, u, v and w are distinct. ($r \neq v, w$ since the hypothesis of (b) makes (v, w) the only saturated edge incident to v or w. It is possible that other vertices coincide, e.g., u may coincide with z or z'.

The argument begins simlar to (e). We first eliminate the case $u = z'$. If $u = z'$ Lemma 2.5(iv) shows rounding path z', w, v costs $2 + x \leq 2 + 2/3$. Hence we can assume

$$u \neq z'.$$

We eliminate the case $u = z$ in a similar way: If $u = z$ Lemma 2.5(iv) shows rounding path w, u, v costs $2 + \bar{x} \leq 2 + 2/3$. Hence we can assume

$$u \neq z.$$

Next we eliminate the case $\alpha = 0$. Let $R = \{z'w, wz, uv\}$. Dropping vw and wv shows $\Delta \leq 1 + 3 = 4$, since (as in (e)) the 3 edges of R start as singletons and the drop makes them slack singletons, with $u \neq z'$ allowing us to round uv.

We now state the 3 main rounding strategies. The first is to round the path $P = u, v, w, z$. Lemma 2.5 shows $\Delta \leq 3 + 2\bar{x}$. Hence we can assume

$$x < 1/2 \text{ and } \bar{x} > 1/2.$$

The second strategy is to round the reverse path $\bar{P} = z', w, v, u$. It has $\Delta \leq (1 + x) + (1 + x) + (1 + \bar{\alpha} - x) = 3 + x + \bar{\alpha}$. Hence we can assume

$$\bar{\alpha} > \bar{x} \text{ and } \alpha < x.$$

In addition note that if vu is the only edge of U leaving v the cost of rounding vu decreases to 1, so $\Delta \leq 3 + 2x < 4$. Hence we can assume $\delta_U(v) = 2$ and $v \in OD_2$, as the tree is not special. As shown in Fig.3 vt is the new edge leaving v, a singleton. t may coincide with z or z'.

We eliminate the possibility $t = z$. Note that $v \in OD_2$, with $\alpha + x < 1$, implies $\alpha + x + y = 1$. Round path $u, v, w, z = t$. The cost is

$$\Delta \leq (1 + \alpha + \bar{x}) + (1 + y) + (1 - y).$$

The first term is for rounding uv. The round makes vw a singleton with slack α in the out-degree constraint of v. This justifies the second term for rounding vw, where we charge $\bar{x} - \alpha = y$ to maintain the out-degree of v. Do this by voiding edge vt. Then wz becomes a singleton with slack $\bar{x}$ at its tail w and slack y at its head $z = t$. Since $y < \bar{x}$ this gives the third term for rounding wz. We conclude $\Delta \leq 3 + \alpha + \bar{x} < 4$. Thus we can assume

$$t \neq z.$$

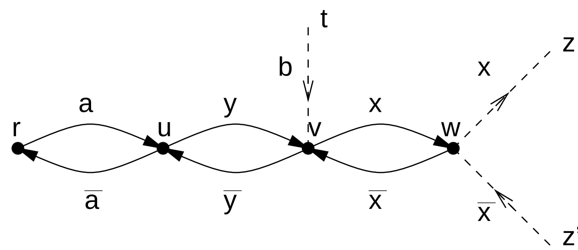

Figure 4: Rounding rule (h).

The third strategy extends the dropping strategy given above. We round the set $D = \{z'w, wz, uv, vt\}$. This gives $\Delta \le 5 + \alpha$. (Rounding uv costs $1 + \alpha$ and the other 3 rounds are slack singletons. Since $u \ne z'$ and $t \ne z$ all these rounds can be done together.) Hence we can assume

$$\alpha > 1/3 \text{ and } \bar{\alpha} < 2/3.$$

Since (v, w) is not special we must now have $u \in V(S) \cup D_2$. We consider these 2 sets in turn.

Case 1. $u \in V(S)$. This gives the saturated edge (r, u) of Fig.3. r may coincide with z, z' or t. We consider 2 subcases.

Case 1.1. $\alpha \ge \bar{a}$ $(\bar{\alpha} \le a)$: Round path $P \cup ru$. Since $z \ne u$ this path is simple. Using the cost for P given above we get $\Delta \le (3 + 2\bar{x}) + 1$. The second term is the cost to round ru, since rounding uv makes ru a slack singleton by $\bar{a} \le \alpha < x < \bar{x}$. Now $\Delta \le 4 + 2\bar{x} < 4 + 4/3$ since $\bar{x} < 2/3$ by (a).

Case 1.2. $\alpha \le \bar{a}$ $(\bar{\alpha} \ge a)$: Round $\bar{P} \cup ur$. Since $z' \ne u$ this path is simple. Using the cost for $\bar{P}$ given above we get $\Delta \le (3 + x + \bar{\alpha}) + (1 + \max\{a - x, 0\})$, since rounding vu makes ur a singleton $(\bar{\alpha} \ge a)$. Thus $\Delta \le 4 + \bar{\alpha} + \max\{a, x\} < 4 + 4/3$ since $x < 2/3$ and $a \le \bar{\alpha} < 2/3$.

Case 2. $u \in D_2 - V(S)$. If $u \in ID_2$ there is an edge ru of weight $\bar{\alpha}$. It is easy to see Case 1.1 applies (ru becomes a slack singleton since $\bar{x} > \alpha$). If $u \in OD_2$ there is an edge ur of weight $\bar{x}$. Case 1.2 applies (ur becomes a slack singleton since $\bar{\alpha} > x$). □

Proof of (h): Fig.4 illustrates the setting of the rule. (If $u \in D_1$ only 1 of the edges ur, ru need exist.) It is possible that t, z or z' coincide with other vertices. But note that $r \ne w$ if $(u, r) \in S$, by acyclicity. In contrast when $u \in D_1$ edge ru could be wz and ur could be $z'w$. Observe also that the hypothesis $u \in D_1$ of case (h) can be strengthened to $u \in D_{11}$. This follows since 2 saturated edges incident to u give the alternate hypothesis of (h), $d_S(u) \ge 2$.

The argument for case (h) is simlar to (b).

We first treat the case $y \ge x$ $(\bar{y} \le \bar{x})$. Rounding the rightwards path u, v, w, z costs

$$\Delta \le (1 + 2\bar{y}) + (1 - \bar{y}) + \bar{x} + 1 = 3 + \bar{y} + \bar{x}.$$

The first term $1 + 2\bar{y}$ is for rounding uv. This voids edge vu and decreases the weight of wv by $\bar{y}$. So the next 2 terms account for rounding vw. (There is no charge for the out-degree of v, since the hypothesis of (h) shows vw has become the only positive edge leaving v.) Finally the last term is for wz which has become a slack singleton.

Rounding the leftwards path z', w, v, u costs

$$\Delta \le (1+x) + (1+x) + (1-x) + y = 3 + x + y.$$

The first term $1+x$ is for rounding singleton $z'w$. The second term is for rounding wv, which has become a singleton. This decreases the weight of uv by x. So the last 2 terms account for rounding vu. (Again there is no charge for the out-degree of v.)

Since $(x+y) + (\bar{x}+\bar{y}) = 2$, one of the quantities $x+y$ or $\bar{x}+\bar{y}$ is ≤ 1. Thus one of our upper bounds on Δ is ≤ 4. This gives a 4/3-rounding. So we can now assume

$$y \le x \text{ and } \bar{y} \ge \bar{x}.$$

Next we eliminate several degenerate cases. First suppose $z = u$. Round path v, w, u. The hypothesis of Lemma 2.5(iv) is satisfied since $\bar{y} \ge \bar{x}$. Hence the cost is $\le 2 + \bar{x}$. Singleton $z'w$ makes this $\le 2 + 2/3$. Hence we can assume

$$z \ne u.$$

Next suppose $z' = u$. Rounding path v, u, w costs

$$\Delta \le (1+2y) + (1+x-2y).$$

The first term is for rounding vu. This voids edge uv, decreasing the out-degree of u by y, and also decreases vw by y ($y \le x$). Thus the singleton $uw = z'w$ has slack $y \le x$ in the relevant contraints at its head and tail. Hence its cost decreases from $1+x$ to $1+x-2y$. The total cost is $\Delta \le 2 + x < 2 + 2/3$ since singleton wz has weight $x < 2/3$. So we can now assume

$$z' \ne u.$$

Finally suppose $z' = t$. The hypothesis $v \in ID_{12}$ implies $b + y + \bar{x} \in \mathbf{Z}$. Each of the 3 quantities is $< 2/3$, by $y \le x$ and the singletons tv, wz and $z'w$. Hence $b + y + \bar{x} < 2$, giving $b + y + \bar{x} = 1$. In particular

$$\bar{b} \ge \bar{x}.$$

Round path t, v, w, z. We claim

$$\Delta \le (1+\bar{b}) + 1 + 1.$$

The first term is for rounding the singleton tv. This voids edges wv and $tw = z'w$, since $\bar{b} \ge \bar{x}$. Hence it makes vw a slack singleton, giving the 2nd term. The 3rd term is for the slack singleton wz. Thus $\Delta \le 3 + \bar{b} \le 4$. So we can now assume

$$z' \ne t.$$

Round the set $\{z'w, wz, tv, vu\}$. To estimate the cost drop vw and wv. We claim

$$\Delta \le 1 + 2 + (1+y) + 1.$$

The first term 1 is for the drop. This makes $z'w$ and wz slack singletons, so together they cost 2. The 3rd term is for rounding vu. There is no charge for the out-degree of v because of the drop and $x \ge y$. The last term is for tv. This edge starts as a singleton. It becomes a slack singleton because wv and uv have been voided, and these are the only other edges entering v (by $v \in ID_{12}$). Since $z' \neq t$ and $z \neq u$ all these rounds can be done together.

We conclude $\Delta \leq 5 + y$. If $y \leq 1/3$ then $\Delta \leq 5 + 1/3 = (4/3)4$. So we are left with the case

$$y > 1/3 \text{ and } \bar{y} < 2/3.$$

We complete the proof by examining length 4 paths.

Case 1. $(u, r) \in S$ and $a \leq y$ ($\bar{a} \geq \bar{y}$). The leftward path z', w, v, u, r is nondecreasing, since $\bar{x} \leq \bar{y} \leq \bar{a}$. The path is simple since $z' \neq u$. Its first edge has co-weight $x \leq 2/3$ by singleton wz. Hence Corollary 2.6(i) shows it gives a 4/3-rounding.

Case 2. $u \in OD_{11}$. This case gives edge $ur \in U$, with weight $\bar{y}$. Round the same leftward path as Case 1. Using $\bar{a} = \bar{y}$ shows we again get a 4/3-rounding, with one possible exception: A priori the path may not be simple because we can have $r = w$. But $ur = uw$, together with $z'w$ being the unique edge of U entering w, makes $u = z'$. This contradicts our previous assumption.

Case 3. $(u, r) \in S$ and $a \geq y$ ($\bar{a} \leq \bar{y}$). Round the rightward path r, u, v, w, z. The path is simple (or a cycle) since $z \neq u$ The nondecreasing subpath u, v, w, z costs $\leq 3 + 2\bar{y}$ (Lemma 2.5(i)). Rounding uv makes ru a slack singleton since $\bar{y} \geq \bar{a}$. So the total cost is $\Delta \leq 3 + 2\bar{y} + 1 = 4 + 2\bar{y} < 4 + 4/3$, using $\bar{y} \leq 2/3$.

Case 4. $u \in ID_{11}$. The argument is similar to Case 2: Case 4 gives edge $ru \in U$, with weight y. Round the same rightward path as Case 3. Using $a = y$ shows we get a 4/3-rounding, unless the path is not simple because $r = w$. But $ru = wu$ together with wz being the unique edge of U leaving w, makes $u = z$. This contradicts our previous assumption. □

3.2 Completeness

We begin by showing the trees that are not special have properties similar to Section 2.3. We have the following version of Lemma 2.10. It accounts for the fact that our double edge rules don't cover all the cases of our previous double edge rule. Throughout this section T is a tree of forest S. Note that Lemma 2.9 still holds.

Lemma 3.2 *A rounding rule can be applied if T has ≥ 2 edges and some vertex β has $V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta\}$.*

Proof: Choose β as the root of T. Lemma 2.9 shows that every leaf of T belongs to $ID_{11} \cap OD_{11}$.

If T has 2 leaves that are siblings, rounding rule (g) is applicable.

Otherwise, let w be a deepest leaf in T and let v be the parent of w. So w is the only child of v. Thus $v \neq \beta$ (since we assume T has ≥ 2 edges) and $d_S(v) = 2$. The hypothesis of the lemma gives two possibilities for v:

(a) $v \in ID_1 \cap OD_1$. Rounding rule (e) applies.

(b) $v \in D_0$. Thus $v \in D_{02}$. Rounding rule (f) applies. □

We have the following variation. The hypothesis assumes T has been rooted but nothing else (e.g., we don't assume the lemma's hypothesis).

Corollary 3.3 *Suppose the root of T has a child x that is not a leaf. A rounding rule can be applied if every descendant of x belongs to $D_0 \cup (ID_1 \cap OD_1)$.*

Proof: Apply the same argument as before to the subtree rooted at x. The hypothesis x not a leaf ensures v descends from x. □

We now show that any tree T that is not special satisfies inequality (5), which we restate for convenience as

$$d_1 \leq 2(s - d_0).$$

The argument is essentially the same as Section 2.3 so we just summarize it: Suppose $s > 0$ (the case $s = 0$ is unchanged). To prove (5) it suffices show that Assumption (A) implies some rounding rule is applicable. Root T at vertex β. If T has ≥ 2 edges Assumption (A) guarantees the hypothesis of Lemma 3.2. Hence a rounding rule can be applied. If T consists of a single edge, (A) and symmetry give $\beta \in ID_{11}$. The other vertex w of T belongs to $ID_{11} \cap OD_{11}$ (Lemma 2.9). Hence assuming T is not special rounding rule (b) applies.

Next we redo the high-level approach so that special trees can be handled. Define UX as the set of edges of U that occur in exactly 1 basic constraint. (Obviously every edge of U occurs in ≥ 1 basic constraint. So $vw \in UX$ means $v \in OD_0^+$ and $w \notin ID_0^+$, or the symmetric condition.) Thus $|UX| = \sum\{\delta_U(v) : v \notin OD_0^+\} + \sum\{\rho_U(v) : v \notin ID_0^+\}$. Our first goal is to prove $UX = \emptyset$ (this is (i') of Section 2.1).

Rewrite (1) as

$$2|U| = |UX| + \textstyle\sum_i i|D_i| \tag{1'}$$

and derive the analog of (2),

$$|D_1| = 2(|S| - |D_0|) + |UX| + \textstyle\sum_{i\geq 3}(i-2)|D_i|. \tag{2'}$$

We partition the special trees into "clusters" as follows. Let u, v, w be as in the definition of special tree. A special tree satisfying the condition on v ($\delta_U(v) > 2$, or $\delta_U(v) = 2$ and $v \notin OD_2$) forms a *v-cluster*. All special trees that are not v-clusters but share a common vertex u form a *u-cluster*.

We extend the parameters s and d_i to apply to a cluster, in the natural way. In a u-cluster (but not a v-cluster) the d_i values include the contribution from vertex u. There is no danger of double-counting since $u \notin V(S)$ implies u is not in any tree of S. Note that s equals the number of trees in the cluster. Also $d_1 = 3s$. In proof, case (b) gives 2 D_1 constraints at w and 1 at v. v cannot have 2 D_1 constraints, since $v \in ID_1 \cap OD_1$ with $uv, vu \in U$ contradicts Lemma 2.2(iii).

In addition we define μ to be a contribution of the cluster to $|UX|$: A v-cluster has $\mu = \delta_U(v)$ if $v \notin OD_0^+$, and $\mu = 0$ otherwise. A u-cluster has μ the sum of two terms, the first being $\delta_U(u)$ if $u \notin OD_0^+$ else 0, the second being $\rho_U(u)$ if $u \notin ID_0^+$ else 0. Clearly $|UX| \geq \sum \mu$ where we sum over all clusters. We show each cluster satisfies this analog of (5):

$$d_1 \leq 2(s - d_0) + \mu + \sum_{i\geq 3}(i-2)d_i, \tag{6}$$

where the inequality is strict if $\mu > 0$.

Note that this will achieve the goal of the first step of the proof: Summing all instances of (5) and (6) implies

$$|D_1| \leq 2(|S| - |D_0|) + \sum \mu + \sum_{i\geq 3}(i-2)|D_i|,$$

with strict inequality if some cluster has $\mu > 0$. Comparing this to (2′) gives $|UX| \le \sum \mu$ with strict inequality if the right-hand side is positive. Since strict inequality is impossible, we conclude both sides are 0. Thus $UX = \emptyset$ as desired. This argument also shows that every tree that is not special satisfies (5) with equality. This is the same as in Section 2.3, and is needed for the second step of the proof.

Now we prove (6). Substituting for d_1, s and d_0 gives this equivalent form: $s \le \mu + \sum_{i\ge 3}(i-2)d_i$. From now on (6) refers to this equivalent form. We prove it by considering these cases:

Case 1. v-clusters. Here $s = 1$. If $v \notin OD_0^+$ then $\mu = \delta_U(v) \ge 2$. If $v \in OD_0^+$ then it belongs to OD_3^+, and $d_i \ge 1$ for some $i \ge 3$.

The remaining cases refer to u-clusters, so fix vertex u. Note that u may be adjacent to v-clusters. u may also be on edges of U that are not associated with any cluster at all.

Case 2. $u \notin D_0^+$. Each tree of the cluster gives 2 edges of U not involved in any constraint at u. So $\mu \ge 2s > s$ as desired.

Now assume $u \in D_0^+$. Note that $u \in ID_i$ implies $i \ge 3$ (the definition of cluster rules out $i = 0$ or 2, and normalcy rules out $i = 1$).

Case 3. $u \in ID_0^+ - OD_0^+$. $\mu = s$ and $d_i \ge 1$ for some $i \ge 3$.

Case 4. $u \in ID_0^+ \cap OD_0^+$. Let $u \in ID_i \cap OD_j$. Consider the sum $\sum_{i\ge 3}(i-2)d_i$ of the right-hand side of (6). Since $i, j \ge 3$ the sum is $\ge 1 + 1$ so we are done if $s \le 2$. Assume $s \ge 3$. The edges incident to u from the trees of the cluster ensure $i, j \ge s$. So the sum is $\ge 2(s-2)$. We are done if $2(s-2) \ge s$, i.e., $s \ge 4$.

The only remaining case is $s = 3$. We claim $\max\{i, j\} \ge 4$. This makes the sum $\ge 2 + 1 = 3$ as desired. To prove the claim assume $i = j = 3$. Thus there are exactly 6 positive directed edges incident to u and they are the edges from 3 special trees. Let σ be the sum of the weights of these 6 edges. So $\sigma = \mathrm{in}(u) + \mathrm{out}(u) \in \mathbf{Z}$. Since the 6 directed edges come from 3 unsaturated edges, $\sigma < 3$. Since the 3 trees are special, each of the 6 weights is $> 1/3$ and $\sigma > 2$. This is a contradiction.

This completes the first step of the proof. The second step is to show $V(S) \subseteq D_0 \cup (ID_1^+ \cap OD_1^+)$ if no rounding rule is applicable. We begin by showing that special trees present no problem. If (v, w) is a special tree the definition obviously puts $w \in ID_1^+ \cap OD_1^+$ and $v \in ID_1$. The definition also implies $\delta_U(v) > 0$ in all 3 cases. Together with $UX = \emptyset$ this puts $v \in OD_1^+$. Hence the vertices of a special tree obey the inclusion of the second step.

The rest of the argument for the second step is essentially the same as Section 2.3 using the new rounding rules, so we just summarize it: Let T be a tree that is not special. We have proved that T satisfies (5) with equality. Exactly as in Section 2.3 we can consider the two possibilities (a) and (b) for β and β', and we must show that a rounding rule is applicable in both cases.

In possibility (a), $V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta\}$. Lemma 3.2 shows a rounding rule can be applied unless T has exactly 1 edge. The case of exactly 1 edge follows just as in Section 2.3: Rounding rule (c) applies to such a tree.

In possibility (b), Section 2.3 shows

$V(T) \subseteq D_0 \cup (ID_1 \cap OD_1) \cup \{\beta, \beta'\}$;

$\beta \in ID_1 - OD_0^+$ and $\delta_U(\beta) = 0$, or β satisfies the I/O symmetric condition;

$\beta' \in D_1$.

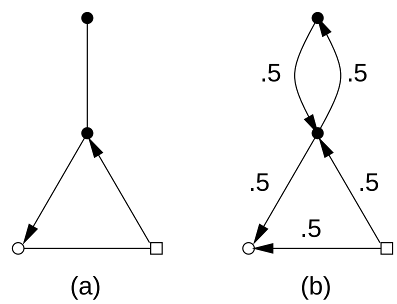

Figure 5: Example with integrality gap 5/4. (a) An optimum degree-constrained orientation, with 2 edges. (b) The optimum fractional orientation, with 2.5 edges.

Root T at β, and let β' descend from child γ of β. If β has a child $x \neq \gamma$ that is not a leaf, Corollary 3.3 shows a rounding rule applies. Hence assume every child $\neq \gamma$ is a leaf. Lemma 2.9 shows every such leaf belongs to $ID_{11} \cap OD_{11}$. Consider these 3 possibilities for the children of β:

(a) β has ≥ 2 children $\neq \gamma$. Rounding rule (g) applies.

(b) β has exactly 1 child $\neq \gamma$. Then $d_S(\beta) = 2$. Hence $\beta \in ID_{12}$ with $\delta_U(\beta) = 0$, or an I/O symmetric condition holds. In addition either γ has a child or is a leaf. In the first case γ is on a second edge of S. In the second case $\gamma = \beta' \in D_1$. Rounding rule (h) applies.

(c) γ is the only child of β. Then $\beta \in ID_{11}$ with $\delta_U(\beta) = 0$, or an I/O symmetric condition holds. Rounding rule (d) applies.

4 Lower Bounds

To describe instances of our orientation problem, say that a vertex has *type* (i, j) if it has in-degree bound i and out-degree bound j. Our examples use only vertices of type $(1, 1), (1, 0)$ and $(0, 1)$. The figures represent these vertices as solid circles, hollow circles, and hollow squares, respectively.

Integrality Gap

Fig.5 gives an orientation problem where (LP) has integrality gap 5/4. It is easy to see the largest (integral) orientation has 2 edges, by examining all possibilities for the type $(0, 1)$ vertex.

Theorem 4.1 *The integrality gap of (LP) is between* $5/4$ *and* $4/3$ *(inclusive).*

Hardness of Approximation

Fig.6 illustrates the reduction that shows our orientation problem is MAXSNP-hard. It is based on the gadget illustrated in Fig.6(a). In general the gadget is an undirected cycle of length $3k$ for some $k \geq 2$, where the pattern of degree constraints is one $(0, 1)$ vertex followed by two $(1, 1)$ vertices. It is easy to see that $\leq 2k$ edges can be oriented. Furthermore this bound can be achieved in exactly 2 ways, in the clockwise orientation illustrated in Fig.6(a) or a similar counterclockwise orientation.

We reduce MAXMUM INDEPENDENT SET to our problem, as illustrated in Fig.6(b)–(c). The gadget is used as vertex substitute. In addition the gadget has a $(1, 0)$ vertex on 1 undirected edge, which can be oriented in a clockwise gadget. Clockwise gadgets correspond to vertices in the independent set. Each edge of the given MIS problem has a $(1, 0)$ vertex on 2 undirected

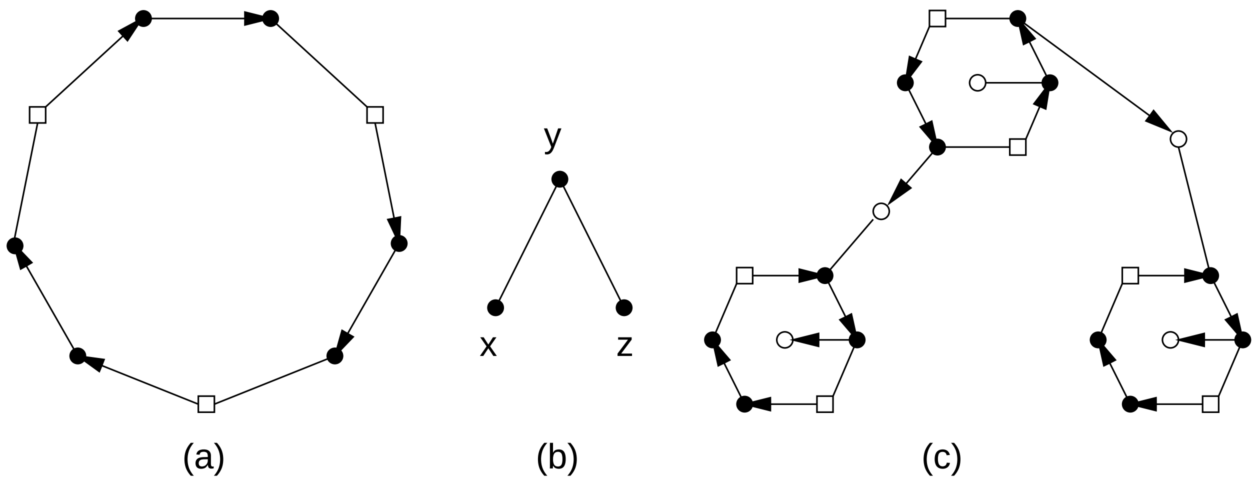

Figure 6: Reduction for MAXSNP-hardness. (a) Gadget, oriented clockwise. (b) Instance of independent set problem. (c) Orientation problem corresponding to (b). The orientation shown corresponds to the maximum independent set.

edges. These edges come from the corresponding vertex substitutes. If the vertex substitute is counterclockwise, its edge can be oriented. As a further example if we add edge (x, z) to Fig.6(b), we cannot orient either of the edges incident to the new edge substitute.

Theorem 4.2 *Finding a maximum upper-degree constrained orientation on a bipartite graph of maximum degree 3 is MAXSNP-complete.*

Acknowledgments

The orientation problem was formulated as part of Michael Duda's University of Colorado M.S. Thesis. The MAXSNP-hardness proof was obtained with Michael and appears in his thesis as well.

References

[1] J. Bang-Jensen and G. Gutin, *Digraphs: Theory, Algorithms and Applications*, Springer, NY, 2001.

[2] A. Frank and A. Gyárfás, How to orient the edges of a graph?, in *Combinatorics Volume I* (Proc. Fifth Hungarian Colloq. on Combinatorics, Keszthely, 1976, A. Hajnal, V.T. Sós, eds.), North-Holland, Amsterdam (1978), pp.353–364.

[3] S.L. Hakimi, On the degrees of the vertices of a directed graph, *J. Franklin Institute* 279 (1965), pp.290–308.

[4] C.A.J. Hurkens and A. Schrijver, On the size of systems of sets every t of which have an SDR, with an application to the worst-case ratio of heuristics for packing problems, *SIAM J. Disc. Math.* 2, 1 (1989), pp.68–72.

[5] K. Jain, A factor 2 approximation algorithm for the generalized Steiner network problem, *Combinatorica* 21 (2001), pp.39–60.

[6] N. Kakimura and S. Iwata, Computing the inertia from sign patterns, *Eleventh MPS Conf. on Integer Programming and Combinatorial Optimization*, 2005, M. Jünger and V. Kaibel eds., Springer-Verlag LNCS 3509, pp.236–248.

[7] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*, Springer, NY, 2003.