NODE-CAPACITATED RING ROUTING

András Frank* Bruce Shepherd† Vivek Tandon‡ and Zoltán Végh§

June 18, 2001

Abstract

We consider the node-capacitated routing problem in an undirected ring network along with its fractional relaxation, the node-capacitated multicommodity flow problem. For the feasibility problem, Farkas' lemma provides a characterization for general undirected graphs asserting roughly that there exists such a flow if and only if the so-called distance inequality holds for every choice of distance functions arising from non-negative node-weigths. For rings this (straightforward) result will be improved in two ways. We prove that, independent of the integrality of node-capacities, it suffices to require the distance inequality only for distances arising from (0-1-2)-valued node-weights, a requirement which will be called the double-cut condition. Moreover, for integer-valued node-capacities, the double-cut condition implies the existence of a half-integral multicommodity flow. In this case there is even an integer-valued multicommodity flow which violates each node-capacity by at most one.

Our approach gives rise to a combinatorial, strongly polynomial algorithm to compute either a violating double-cut or a node-capacitated multicommodity flow. A relation of the problem to its edge-capacitated counterpart will also be explained.

1 Introduction

Let $G = (V, E)$ be an undirected graph called a **supply** graph and $H = (V, F)$ a so-called **demand** graph on the same node set. Suppose that a nonnegative demand $h(f)$ is assigned to every demand edge f and nonnegative capacity $g(e)$ is assigned to every supply edge e. By a **path** P we mean an undirected graph $P = (U, A)$ where $U = \{u_1, u_2, \ldots, u_n\}$, $A = \{e_1, \ldots, e_{n-1}\}$, and $e_i = u_iu_{i+1}$, $i = 1, \ldots, n-1$. (We assume that the nodes u_i are distinct.) The edge-set A of a path P is denoted by $E(P)$ while its node set by $V(P)$. Nodes u_1 and u_n are called the **end-nodes** of P while the other nodes of P are called **internal nodes** and their set is denoted by $I(P)$. We say that a path P **connects** its end-nodes and that P **uses** an edge e if $e \in E(P)$.

For a demand edge $f \in F$, let $\mathcal{P}_f$ denote the set of paths of G connecting the end-nodes of f and let $\mathcal{P} := \cup(\mathcal{P}_f : f \in F)$. By a **path-packing** we mean a function $x : \mathcal{P} \to \boldsymbol{R}_+$. This is said to **fulfil** the demand if

$$\sum_{P \in \mathcal{P}_f} x(P) \geq h(f)$$

holds for every $f \in F$. The **occupancy** $o_x(e)$ and $o_x(v)$ (by x) of a supply edge $e \in E$ and of a node $v \in V$ is defined, respectively, by

$$o_x(e) := \sum_{\substack{P \in \mathcal{P} \\ e \in E(P)}} x(P) \text{ and } o_x(v) := \sum_{\substack{P \in \mathcal{P} \\ v \in I(P)}} x(P).$$

*Department of Operations Research, Eötvös University, Kecskeméti u. 10-12, Budapest, Hungary, H-1053 and Traffic Lab Ericsson Hungary, Laborc u. 1, Budapest, Hungary H-1037. e-mail: frank@cs.elte.hu. Supported by the Hungarian National Foundation for Scientific Research, OTKA T029772 and T030122.

†Bell Laboratories, Lucent Technologies, 600 Mountain Ave., Murray Hill, NJ, U.S.A. e-mail: bshep@research.bell-labs.com

‡Ciena Europe Ltd. High Holborn, London, UK

§Eötvös University, Kecskeméti u. 10-12, Budapest, Hungary, H-1053

A path-packing x is called **feasible** (with respect to edge-capacity vector g) if

$$o_x(e) \le g(e) \tag{1}$$

holds for every supply edge e.

Sometimes we are given a capacity function $c : V \to \boldsymbol{R}_+$ on the node set V rather than on E. A path-packing x is called **feasible** (with respect to node-capacity) if

$$o_x(v) \le c(v) \tag{2}$$

holds for every element $v \in V$. (1) and (2) are called, respectively the **edge-** and the **node-capacity constraints.**

The **edge-** or **node-capacitated multicommodity flow** problem, respectively, consists of finding an edge- or node-feasible path-packing fulfilling the demand. It is sometimes called the **fractional routing** problem. If x is required to be integer-valued, we speak of an **integer multicommodity flow** problem or a **routing** problem. If $2x$ is required to be integer-valued, we speak of a **half-integral multicommodity flow** problem or a **half-integral routing** problem. If each demand and each capacity is one, and x is also required to be (0–1)-valued, then we speak of an **edge-disjoint** or **node-disjoint** paths problem. That is, the edge-disjoint (node-disjoint) paths problem can be formulated as deciding if there is a path in G for each demand edge $f \in F$ connecting its end-nodes so that these $|F|$ paths are edge-disjoint (internally node-disjoint).

In the literature the term **through-capacity** is sometimes used in order to emphasize that the only paths which contribute to the occupancy of a node, are those which contain it as an internal node. Because we use the term node-capacity exclusively in this sense, the adjective *through-* will be dropped.

The node-capacitated versions are more general than their edge-capacitated counterparts since this latter can easily be reduced to the former one by subdividing each supply edge e by a new node v_e and assigning the capacity $g(e)$ to v_e. Therefore it is no wonder that there are many more results on edge-capacitated routing problems while our knowledge on node-disjoint or node-capacitated paths problem is much more limited. The operation of splitting nodes into two is a well-known elementary construction to transform node-disjoint paths problems into their edge-disjoint versions (for example, in this way, the directed node-disjoint version of Menger's theorem can be derived from its edge-disjoint counterpart, see [1]). This trick does not seem to work for undirected graphs, and even in directed graphs by splitting the nodes into two we may loose the speciality of the initial problem which ensured nice solvability for the edge-case. (For example, the directed edge-disjoint paths problem is tractable if $G + H$ is planar and G is acyclic, due to a theorem of Lucchesi and Younger [6] while nothing is known for the node-capacitated version of this special case.)

These problems are important from both the theoretical and practical viewpoints. For example, telecommunication network routing problems form a main source of practical demand of this type of problems. In particular, the present work was strongly motivated by engineering investigations in the area of so called passive optical networks, as is discussed briefly in the conclusions.

In this paper we consider the (fractional) routing problem for rings. By a **ring** (=cycle=circuit) we mean an undirected graph $G = (V, E)$ where $V = \{v_1, v_2, \ldots, v_n\}$, $E = \{e_1, \ldots, e_n\}$, and $e_i = v_iv_{i+1}$, $i = 1, \ldots, n$. [Notation: $v_{n+1} = v_1$.] We intuitively think that nodes of G are drawn in the plane in a clockwise cyclic order. The edge-capacitated half-integral version was solved earlier by Okamura and Seymour [8], while its integer-valued counter-part by the first-named author of the present work [4]. (Actually, both results concerned graphs more general than rings.)

The node-capacitated (fractional) routing problem for rings is the main concern of the present work. We characterize completely the solvability of the half-integral routing problem. This gives rise to a sufficient condition for the solvability of the routing problem; the problem of finding a necessary and sufficient condition remains open. Our approach also provides a combinatorial, strongly polynomial solution algorithm.

We conclude this introductory section with some further notation. In what follows, for any function $p : S \to \boldsymbol{R}$ and subset $X \subseteq S$, we use the notation $p(X) := \sum_{s \in X} p(s)$. For a number γ, let $\gamma^+ := \max(0, \gamma)$. For a function p, let p^+ be defined by $p^+(s) := p(s)^+$.

For a subset $X \subset V$, the set of edges of G with exactly one end-node in X is called a **cut** of G. The **capacity** of a cut (with respect to a capacity function g) is the sum of capacities of (supply) edges in the cut and is denoted by $d_g(X)$. The **load** of a cut (with respect to a demand function h) is the sum of demands of demand edges with exactly one end-node in X and is denoted by $d_h(X)$.

Any two edges e_i and e_j of a ring G determine a cut $\{e_i, e_j\}$ of G. A demand edge f is said to **cross** this cut if the two end-nodes of f belong to different components of of $G - \{e_i, e_j\}$. We denote the load of cut $\{e_i, e_j\}$ by $L_h(e_i, e_j)$, that is, $L_h(e_i, e_j) = d_h(\{v_{i+1}, \dots, v_j\})$. Two cuts $\{e_i, e_j\}$ and $\{e_k, e_l\}$ are **crossing** if the four edges are distinct and e_k and e_l belong to different components of $G - \{e_i, e_j\}$ which is equivalent to requiring that e_i and e_j belong to different components of $G - \{e_k, e_l\}$. Two pairs $\{u, u'\}$ and $\{v, v'\}$ of nodes of the ring are called **crossing** if they are distinct and u and u' belong to different components of $G - \{v, v'\}$ which is equivalent to requiring that v and v' belong to different components of $G - \{u, u'\}$. If the end-nodes of a demand edge are u and u' and the pairs $\{u, u'\}$ and $\{v, v'\}$ are crossing, then we say that f **crosses** the pair $\{v, v'\}$.

For two nodes $u, v \in V$, let P_{uv} denote the subpath of G identified by traversing the ring clockwise from u to v, i.e., for any edge e of the path, u, e, v occur clockwise in the ring. With this notation, the set $\mathcal{P}_f$ of paths connecting the end-nodes u, v of a demand edge f consists of the two paths P_{uv} and P_{vu}. Clearly, $d_h(I(P_{v_{i+1}v_j})) = d_h(I(P_{v_{j+1}v_i})) = L_h(e_i, e_j)$, $d_g(I(P_{v_{i+1}v_j})) = d_g(I(P_{v_{j+1}v_i})) = g(e_i) + g(e_j)$.

2 The edge-cut criterion for edge-capacitated routing

A simple necessary condition for solvability of the edge-disjoint paths problem is the so-called **edge-cut criterion** which requires that the load of every cut cannot exceed its capacity, that is, $d_G(X) \geq d_H(X)$ holds for every subset $X \subseteq V$ where $d_G(X)$ and $d_H(X)$ stand for the number of supply and demand edges, respectively, connecting X and $V - X$. For the routing or multicommodity flow problem the edge-cut criterion requires

$$d_g(X) \geq d_h(X) \tag{3}$$

for every subset $X \subseteq V$. (3) will be called the **edge-cut inequality**.

The following useful proposition is well-known and easy to prove anyway.

Proposition 2.1 *If the edge-cut inequality (3) holds for every subset X of V for which both X and $V - X$ induce a connected subgraph of G, then it holds for every subset $X \subseteq V$.*

An important special case when the edge-cut criterion is sufficient is due to H. Okamura and P. D. Seymour [8].

THEOREM 2.2 (Okamura and Seymour) *Suppose that (i) the supply graph G is planar and each edge has capacity one, (ii) each demand is one and is between two nodes on the outer face of G, and (iii) $G + H$ is Eulerian (the degree of each node is even). Then the edge-cut criterion is necessary and sufficient for the the solvability of the edge-disjoint paths problem. If only (i) and (ii) are assumed, then the cut-criterion implies the existence of a feasible half-integral routing.*

The second half is an immediate consequence of the first one. By Proposition 2.1 the theorem of Okamura and Seymour specializes to rings as follows. We formulate this result so as to apply to the (edge-)capacitated case.

Corollary 2.3 *In a ring G the multicommodity flow problem has a solution if and only if the edge-cut inequality*

$$g(e_i) + g(e_j) \geq L_h(e_i, e_j) \text{ holds whenever } 1 \leq i < j \leq n. \tag{4}$$

Moreover, if both h and g are integer-valued and (4) holds, then there is a half-integral routing.

In [4] a necessary and sufficient condition was given (much more complicated than the edge-cut criterion) for the solvability of the edge-disjoint paths problem when, in addition to (i) and (ii), we assume only a weakening of (iii), namely, (iii′) the degree in $G + H$ of every node not in the outer face is even. When G is a circuit with possible parallel edges, then (iii′) is automatically satisfied and in this case a simplified characterization follows for the routability. To formulate it, we call a cut $\{e_i, e_j\}$ of the ring **tight** if $g(e_i) + g(e_j) = L_h(e_i, e_j)$, that is, if the load and the capacity of the cut are equal. A subset $X \subseteq V$ and the cut of G defined by X are called **odd** (with respect to g and h) if $d_h(X) + d_g(X)$ is odd.

Corollary 2.4 ([4]) *Suppose that G is a ring and both h and g are integer-valued. The routing problem has a solution if and only if (4) holds and no odd component can arise by deleting the four edges of two crossing tight cuts. In particular, if (4) holds with strict inequality (that is, no tight cuts exist), then the routing problem always has a solution.*

Some remarks are in order. First, the original proof of Okamura and Seymour was algorithmic and this algorithmic proof can be made much simpler for rings as noted by Seymour, Schrijver and Winkler [9], a paper whose main achievement was an approximation algorithm for what is called an unsplittable flow problem in a ring network. (An *unsplittable flow* is one where, for every demand edge f, the flow realizing the demand $h(f)$ must use a single path: it cannot be split between the two possibilities.) Second, the proof in [4] is also algorithmic. And finally, Corollary 2.4 can be directly derived from Corollary 2.3, as well.

As far as the node-capacitated ring routing problem is concerned, it would be a natural attempt to mimic Corollary 2.3. The necessary condition in this case corresponding to (4) would require for every choice of pairs $\{v_i, v_j\}$ of nodes that the total demand of demand edges crossing $\{v_i, v_j\}$ should not exceed the capacity $c(v_i)+c(v_j)$ of this pair. While this condition is clearly necessary, it is not sufficient even for the existence of a fractional solution. To see this, let G be a ring of six nodes each having capacity one. There are two demand edges v_1v_4, v_2v_6 and their demands are $h(v_1v_4) = 1$, $h(v_2v_6) = 2$. Here, no pair of nodes crosses both demand edges and hence the necessary condition mentioned before is satisfied but easy case-checking shows that there is no multicommodity flow fulfilling the demands.

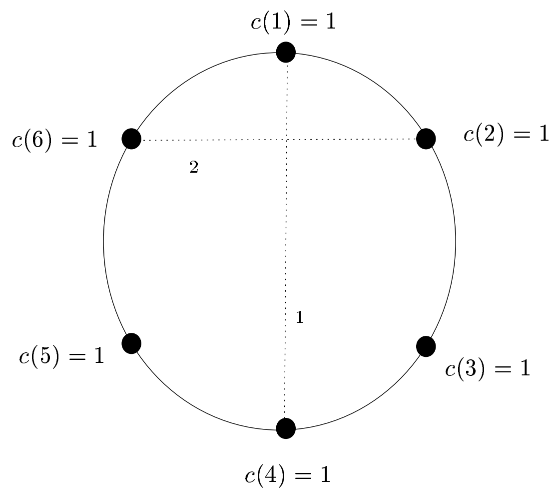

Figure 1: An instance with all cut inequalities satisfied but with a violated double cut.

Our plan therefore is turn to the multicommodity flow problem, both edge- and node-constrained. As these may be considered as plain linear programming problems, with the help of Farkas' lemma, a necessary and sufficient condition, called the distance criterion, will first be exhibited. After that we show for ring networks that a special form of the distance criterion is sufficient.

3 The distance criterion and the main result

We recall a well-known result, sometimes called the *Japanese* theorem [5],[7].

THEOREM 3.1 *Given a demand function h on the edge-set of a demand graph $H = (V, F)$ and a capacity function g on the edge-set of a supply graph $G = (V, E)$, there is a feasible solution to the edge-capacitated multicommodity flow problem fulfilling the demands if and only if the* **edge-distance criterion** *holds, that is, for every non-negative function $y : E \to \boldsymbol{R}_+$,*

$$\sum_{e \in E} y(e)g(e) \geq \sum_{f \in F} h(f)s'(f, y) \tag{5}$$

where $s'(f,y)$ denotes the minimum y-cost of a path connecting the two end-nodes of demand edge f.

That is, $s'(f,y)$ is the distance of the end-nodes of f in G with respect to cost-function y. (5) is called the **edge-distance inequality**. Note that it suffices to require (5) for integer-valued functions y. It is also useful to realize that if we choose y to be 1 on the edges of a cut of G and zero otherwise, then the distance inequality specializes to the edge-cut inequality.

Theorem 3.1 can be formulated so as to apply to the node-capacitated multicommodity flow problem. The counterpart (6) of (5) will be called the **node-distance inequality**.

THEOREM 3.2 *Given a demand function h on the edge-set of a demand graph $H=(V,F)$ and a capacity function c on the node-set V of a supply graph $G=(V,E)$, there is a feasible solution to the node-capacitated multicommodity flow problem fulfilling the demands if and only if the* **node-distance criterion** *holds, that is, for every non-negative function $y:V\to\boldsymbol{R}_+$,*

$$\sum_{v\in V} y(v)c(v) \geq \sum_{f\in F} h(f)s(f,y) \tag{6}$$

where $s(f,y) := \min\{\sum_{v\in I(P)} y(v) : P$ a path in G connecting the end-nodes of $f\}$, that is, $s(f,y)$ is the minimum total cost of internal nodes on path connecting the end-nodes of f.

Proof. Let $x\in\boldsymbol{R}_+^{\mathcal{P}}$ be a node-feasible multicommodity flow fulfilling the demands and let $y:V\to\boldsymbol{R}_+$. (Recall the notations $o_x(v)$ and $y(I(P))$ from the introduction.) We have

$$\sum_{v\in V} y(v)c(v) \geq \sum_{v\in V} y(v)o_x(v) = \sum_{v\in V} y(v) \sum_{\substack{P\in\mathcal{P}\\ v\in I(P)}} x(P) =$$

$$\sum_{f\in F}\sum_{P\in\mathcal{P}_f} x(P)y(I(P)) \geq \sum_{f\in F} h(f)s(f,y),$$

that is, (6) holds.

To see the sufficiency, suppose that the node-distance criterion holds and consider the polyhedron $Q=\{x\in\boldsymbol{R}_+^{\mathcal{P}} : o_x(v)\leq c(v)$ for every node $v\in V$, and $-\sum_{P\in\mathcal{P}_f)} x(P)\leq -h(f)$ for each $f\in F\}$. As the members of Q are exactly the node-feasible multicommodity flows fulfilling the demands, we want to show that Q is non-empty. If this were not true, then Farkas' lemma implies the existence of a non-negative vector y whose components correspond to the elements of V, and a non-negative vector z whose components correspond to the demand edges so that

$$y(I(P)) \geq z(f) \text{ for each path } P\in\mathcal{P}_f, f\in F \tag{7}$$

and

$$\sum_{v\in V} y(v)c(v) < \sum_{f\in F} h(f)z(f). \tag{8}$$

By (7), $s(f,u)\geq z(f)$ for $f\in F$ and hence (8) implies $\sum_{v\in V} y(v)c(v) < \sum_{f\in F} h(f)s(f,y)$ contradicting the node-distance inequality (6). •

In the special case when y is $(0-1)$-valued we say that y is a **cut** while if y is $(0-1-2)$-valued y is called a **double-cut** (so a cut is a special double-cut). Inequality (6) for a double-cut y is called the **double-cut condition** while the **double-cut criterion** is one requiring (6) for all double-cuts y, that is,

$$\sum_{f\in F}(h(f)s(f,y) \leq \sum_{v\in V} y(v)c(v) \text{ for every double-cut } y. \tag{9}$$

Now we are in a position to formulate the main result of the paper which is about the node-capacitated (fractional) routing problem when the supply graph is a ring. Since in a ring there are two paths between any pair $\{u,v\}$ of nodes,

$$s(f,y) = \min\{y(I(P_{uv})), y(I(P_{vu}))\}, \tag{10}$$

THEOREM 3.3 *Let us be given a ring G, a demand graph H endowed with a demand function h, and a node-capacity function c. There is a solution to the node-capacitated multicommodity flow problem (=fractional routing problem) if and only if the double-cut criterion (9) holds. If, in addition, both the demand function h and the node-capacity function c are integer-valued, then (9) is necessary and sufficient for the existence of a half-integral path-packing.*

Before proceeding to the proof, let us make some comments. At the end of the preceding section we indicated that a natural necessary condition involving two-element node-cuts is not sufficient. This condition may be interpreted as requiring the node-distance inequality (6) for cuts y with exactly two components of value 1. One may, however, hope that requiring (6) for every cut may perhaps be sufficient. However, the same example shows that this is not the case. By case-checking one can see that (6) is satisfied for every cut y. In this view, we may feel a bit fortunate in that (6) is not really required for all possible integer-valued y's but only for double-cuts. Incidentally, the example in Figure 1 does violate the double-cut criterion: take $y(v_1) = 2, y(v_2) = 0, y(v_3) = 1, y(v_4) = 0, y(v_5) = 1, y(v_6) = 0$. Then $\sum_{v \in V} y(v)c(v) = 2 + 1 + 1 \not\geq 2 \cdot 2 + 1 \cdot 1 = \sum_{f \in F} h(f)s(f, y)$.

In the view of the existing edge-capacitated version of the problem, it is tempting to believe that there must be a straightforward way to derive Theorem 3.3 from Corollary 2.3 via an elementary construction. A second look, however, may indicate some of the difficulties: how can one expect that the sufficiency of the pretty simple edge-cut criterion in the edge-capacitated case may be used to derive the sufficiency of the significantly more complex double-cut criterion? Though not very cheaply, such a derivation does turn out to exist, and this is the content of the next section.

4 The proof

We have seen already the necessity of the more general node-distance criterion. The proof of sufficiency, as indicated above, is based on a rather sophisticated reduction to the edge-capacitated case, that is, to Corollary 2.3.

We give a brief overview of how the proof proceeds. We first show that we may as well assume that each demand $h(f)$ is an even integer. We then show that if an edge-capacity function g satisfies some additional side constraints imposed by the node capacities, then every feasible fractional solution for the edge-capacitated problem will automatically satisfy the node constraints as well (Lemma 4.2). At this point we embark on finding such edge capacities. This is achieved by formulating a weighted directed graph in which we search for a potential function. If we succeed, then the desired edge capacities can be determined easily from the potential (Lemma 4.3). If we fail, then a negative circuit is found. It is then shown how, with some work, this circuit identifies a violated double-cut (Lemma 4.4).

We say that a demand function h is **Eulerian** if h is integer-valued and $d_h(v)$ is even for each node v. Obviously, for an Eulerian demand function the load of every cut is even.

Claim 4.1 *It suffices to prove the second half of the theorem for Eulerian demand functions h.*

Proof. Let v_i and v_{i+1} be two subsequent nodes in the ring. As a one-edge path connecting v_i and v_{i+1} has no inner node, increasing the demand between them affects neither the solvability of the problem nor the double-cut criterion.

If $d_h(v)$ is not even everywhere, then there are two nodes v_i and v_j so that both $d_h(v_i)$ and $d_h(v_j)$ are odd and so that $d_h(v_l)$ is even for each l with $i < l < j$. We can then increase by one the h-values on all demand edges $v_iv_{i+1},\ v_{i+1}v_{i+2}\ ,\ldots,\ v_{j-1}v_j$ and this way both $d_h(v_i)$ and $d_h(v_j)$ get even while all other $d_h(v_l)$ values remain even for $i < l < j$. •

Lemma 4.2 *Let $g : E \to \boldsymbol{R}$ be a function satisfying the edge-cut criterion and*

$$g(e_{i-1}) + g(e_i) \leq d_h(v_i) + 2c(v_i) \tag{11}$$

for every $v_i \in V$. Suppose that a path-packing $x : \mathcal{P} \to \boldsymbol{R}_+$ fulfils demand h and satisfies edge-capacity constraint (1) with respect to edge-capacity function g^+. Then x satisfies node-capacity constraint (2). Moreover, if g is integer-valued and there is an edge $e \in E$ with $g(e) \leq 0$, then x is integer-valued.

Proof. Let us consider an arbitrary node v_i. If $g(e_i) \leq 0$, then $g^+(e_i) = 0$ and hence $x(P) = 0$ for every path P using e_i. Therefore $x(P) = 0$ for every path P using v_i as an inner node, that is, in this case (2) is satisfied for v_i. The same argument works when $g(e_{i-1}) \leq 0$.

In the remaining case $g(e_i) > 0$ and $g(e_{i-1}) > 0$. For each i, let $\mathcal{P}(v_i)$ be the set of paths which have v_i as an endpoint. Let

$$\alpha := \sum_{\substack{P \in \mathcal{P}(v_i) \\ e_i \in E(P)}} x(P) \quad \text{and} \quad \beta := \sum_{\substack{P \in \mathcal{P}(v_i) \\ e_{i-1} \in E(P)}} x(P).$$

Then $\alpha + \beta \geq d_h(v_i)$. We have $o_x(v_i) + \alpha \leq g^+(e_i) = g(e_i)$ and $o_x(v_i) + \beta \leq g^+(e_{i-1}) = g(e_{i-1})$. By combining these with (11), we obtain $\alpha+\beta+2o_x(v_i) \leq g(e_{i-1})+g(e_i) \leq d_h(v_i)+2c(v_i) \leq \alpha+\beta+2c(v_i)$, from which $o_x(v_i) \leq c(v_i)$ follows, that is, x satisfies (2).

Let $g(e) \leq 0$ for some $e \in E$, that is, $g^+(e) = 0$. Let f be any demand edge and P_1, P_2 the two paths connecting its end-nodes. If P_1 contains e, then $x(P_1)$ must be zero and hence $x(P_2) = h(f)$, an integer. •

We define an auxiliary directed graph $D = (U, A)$ endowed with a cost function w on A as follows. Think of subdividing each edge e_i of G by two nodes s_i and t_i and let $U := \{s_i, t_i : i = 1, \ldots, n\}$ – see Figure 2. Imagine $\ldots t_{i-1}, v_i, s_i, t_i, v_{i+1}, s_{i+1} \ldots$ clockwise. We will denote a directed edge with head t and tail s by st (and hope this will not cause any ambiguity). Let A_1 consist of the following directed edges:

$$t_{i-1}s_i \text{ and } t_is_{i-1} \text{ for } i = 1, \ldots, n.$$

Let A_2 consist of the following directed edges:

$$s_it_j \text{ and } s_jt_i \text{ for } 1 \leq i < j \leq n.$$

(In particular, $s_it_{i-1} \in A_2$ and $t_{i-1}s_i \in A_1$). Let $A := A_1 \cup A_2$ and define a cost-function $w : A \to \boldsymbol{R}$ as follows. For A_1-edges let

$$w(t_{i-1}s_i) := w(t_is_{i-1}) := d_h(v_i) + 2c(v_i) \text{ for } i = 1, \ldots, n \tag{12}$$

and for A_2-edges let

$$w(s_it_j) := w(s_jt_i) := -L_h(e_i, e_j) \text{ for } 1 \leq i < j \leq n. \tag{13}$$

(Recall that $L_h(e_i, e_j)$ denotes the demand across cut $\{e_i, e_j\}$.)

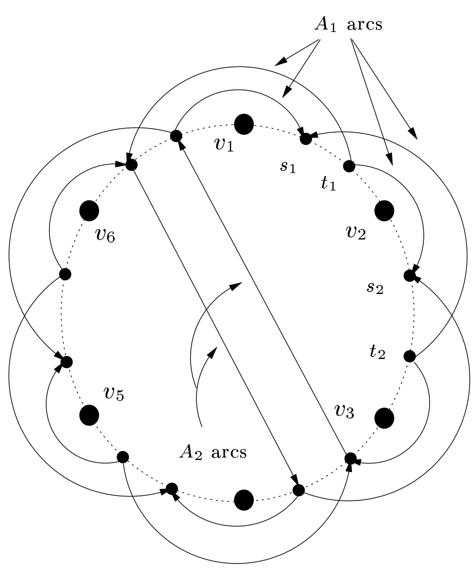

Figure 2: The auxiliary digraph.

Lemma 4.3 (A) *If there is no negative circuit in $D = (U, A)$ with respect to w, then there exists a function $g : E \to \boldsymbol{R}$ satisfying (11) and the edge-cut criterion.* (B) *If, in addition, c is integer-valued and h is Eulerian, then g can be chosen integer-valued.*

Proof. It is well-known that if one applies the Bellman-Ford algorithm (from any root node) in a digraph $D = (U, A)$, then either it detects a directed circuit of negative total weight or else it produces shortest path costs $\pi(s)$ which form a **feasible potential**, i.e., $\pi(t) - \pi(s) \leq w(st)$ for every $st \in A$. Moreover, π will be chosen integer-valued whenever w is integer-valued. This implies that if each component of w is an even integer, then π is even-integer-valued. Therefore the hypothesis of the lemma implies that there is a feasible potential π. For $i = 1, \ldots, n$ let us define

$$g(e_i) := (\pi(s_i) - \pi(t_i))/2. \tag{14}$$

Since in Case (B) each load is even, and so is the cost $w(a)$ of each directed edge $a \in A$, we have that π is even-integer-valued and g is integer-valued.

We claim that g satisfies (11). Indeed, for $t_{i-1}s_i \in A_1$ and $t_i s_{i-1} \in A_1$ we have

$$\pi(s_i) - \pi(t_{i-1}) \leq w(t_{i-1}s_i) = d_h(v_i) + 2c(v_i),$$

$$\pi(s_{i-1}) - \pi(t_i) \leq w(t_i s_{i-1}) = d_h(v_i) + 2c(v_i).$$

By adding these up and subdividing by two, we obtain $g(e_i) + g(e_{i-1}) \leq d_h(v_i) + 2c(v_i)$, as required for (11).

We claim that g satisfies the edge-cut criterion. Indeed, for $s_i t_j \in A_2$ and $s_j t_i \in A_2$ we have

$$\pi(t_j) - \pi(s_i) \leq w(s_i t_j) = -L_h(e_i, e_j),$$

$$\pi(t_i) - \pi(s_j) \leq w(s_j t_i) = -L_h(e_i, e_j).$$

By adding these up and subdividing by two, we obtain $-g(e_i) - g(e_j) \leq -L_h(e_i, e_j)$ which is just the edge-cut inequality for rings. •

Lemma 4.4 *If there is a negative directed circuit C in $D = (U, A)$ with respect to w, then there is a double-cut y violating the double-cut inequality.*

Proof. For $(i = 1, \ldots, n)$, let us define $y : V \to \{0, 1, 2\}$ as follows:

$$y(v_i) := |C \cap \{t_{i-1}s_i, t_i s_{i-1}\}|, \tag{15}$$

that is, $y(v_i)$ is the number of edges among $t_{i-1}s_i$ and $t_i s_{i-1}$ which belong to C. We are going to show that this double-cut y violates the double-cut inequality.

For $l = 1, 2$, let C_l denote the set of edges of C belonging to A_l. For an edge $a = s_i t_j \in A_2$, we will use the notation $L_h(a) := L_h(e_i, e_j)$. We will show, but defer to the end, that

$$\sum_{a \in C_2} L_h(a) \leq 2 \sum_{f \in F} h(f) s(f, y) + \sum_{v \in V} y(v) d_h(v). \tag{16}$$

Since C is a negative circuit, $0 > w(C) = w(C_1) + w(C_2)$, that is, $-w(C_2) > w(C_1)$. Hence, by using the definitions of y and w (in (15), (12), and (13)), we have

$$\sum_{a \in C_2} L_h(a) = -w(C_2) > w(C_1) = \sum_{a \in C_1} w(a) = \sum_{v \in V} y(v) \cdot (d_h(v) + 2c(v)). \tag{17}$$

By combining (16) and (17), we obtain

$$\sum_{v \in V} y(v) d_h(v) + 2 \sum_{v \in V} y(v) c(v) < \sum_{a \in C_2} L_h(a) \leq 2 \sum_{f \in F} h(f) s(f, y) + \sum_{v \in V} y(v) d_h(v),$$

that is,

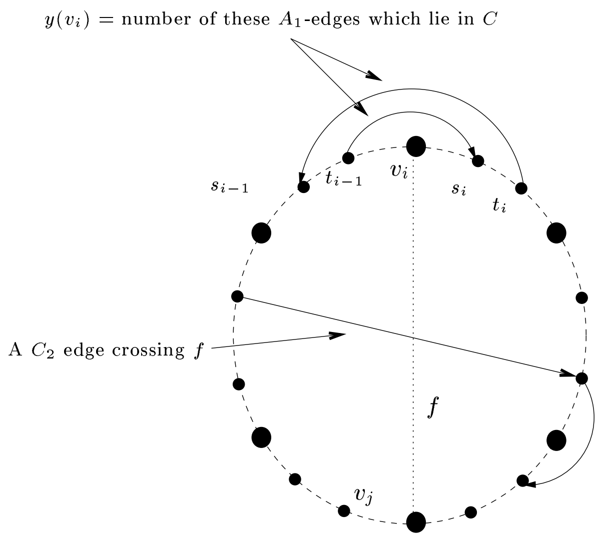

Figure 3:

$$\sum_{v \in V} y(v)c(v) < \sum_{f \in F} h(f)s(f,y)$$

showing that y indeed violates the double-cut inequality, completing the proof of the lemma.

It remains only to prove (16). Let $f = v_i v_j$ be a demand edge. For convenience, we assume that the ring is drawn in the plane in such a way that f is vertical and v_i is above v_j – see Figure 3. We say that a (directed) edge of digraph D **crosses** f if one of its end-nodes is on the left-hand side of f and the other on the right-hand side. Let $\alpha(f)$ denote the number of C_2-edges crossing f.

We claim first that

$$\alpha(f) \leq 2s(f,y) + y(v_i) + y(v_j). \tag{18}$$

To see this, by symmetry, we may suppose that $y(I(P_{v_i v_j})) \leq y(I(P_{v_j v_j}))$, that is, $s(f,y) = y(I(P_{v_i v_j}))$. Hence $s(f,y)$ is the number of C_1-edges whose end-nodes are both in the right hand side of f. For every C_2-edge a crossing f with head in the right-hand side of f we assign (or charge) to a the subsequent C_1-edge of C. For every C_2-edge a' crossing f with tail in the right-hand side of f we assign to a' the preceding C_1-edge of C.

If $t_{l-1}s_l$ is a C_1-edge entirely on the right-hand side of f, then it has been assigned to at most two C_2-edges. Analogously, a C_1-edge $s_l t_{l-1}$ in the right-hand side of f has been assigned to at most two C_2-edges. Note that the number of such C_1 edges is just $y(I(P_{v_i v_j})) = s(f,y)$.

Next, the tail of edge $t_{i-1}s_i$ is in the left-hand side of f, therefore $t_{i-1}s_i$ has been assigned to at most one C_2-edge. Similarly, the head of $t_i s_{i-1}$ is in the left-hand side of f, therefore $t_i s_{i-1}$ has been assigned to at most one C_2-edge crossing f. Analogous statements hold for index j.

By combining these, we obtain that $\alpha(f) \leq 2s(f,y) + y(v_i) + y(v_j)$, which is (18).

Finally we show that (16) holds. We first examine the quantity $\sum_{a \in C_2} L_h(a)$ which counts, for each edge $a \in C_2$, the total demand on edges which cross a. Alternatively, we may write this as the weighted sum which counts for each demand edge f, the number of C_2-edges which cross it. In other words, $\sum_{a \in C_2} L_h(a) = \sum_{f \in F} \alpha(f)h(f)$. By (18),

$$\sum_{f \in F} \alpha(f)h(f) \leq \sum_{f = v_i v_j \in F} h(f) \cdot [2s(f,y) + y(v_i) + y(v_j)] =$$

$$2 \sum_{f \in F} h(f)s(f,y) + \sum_{v \in V} y(v_i)d_h(v_i),$$

as required for (16). •

To return to the proof of sufficiency in Theorem 3.3, suppose now that the double-cut criterion holds. Then by Lemma 4.4 there is no negative circuit in the auxiliary digraph D. Hence, by Lemma

4.3, there is a function $g : E \to \boldsymbol{R}$ satisfying (11) and (3) which is integer-valued if g is integer-valued and h is Eulerian. By Corollary 2.3 there is a path-packing x obeying the edge-capacity constraints with respect to g and fulfilling the demand h, which is half-integer-valued if g is integer-valued. By Lemma 4.2 this x obeys the node-capacity constraints. • •

5 Routing in rings and algorithmic aspects

Theorem 3.3 provided a necessary and sufficient condition for the existence of a node-capacitated half-integral routing in a ring. The approach has the following consequence concerning routings (=integer-valued multicommodity flows) in rings.

THEOREM 5.1 *Suppose that in the ring-routing problem the capacity function c is integer-valued and assume that the double-cut criterion holds with respect to c. Let c' be a capacity function defined by $c'(v) := c(v)+1$ for $v \in V$. Then there is a feasible routing (with respect to c' fulfilling the demand).*

Proof. Lemma 4.3 guarantees the existence of integer-valued capacity function g satisfying the edge-cut criterion. Let us define g' by $g'(e) := g(e) + 1$ for $e \in E$. Then g' strictly satisfies all the edge-cut inequalities. Hence, by the second part of Corollary 2.4, there is an edge-feasible routing x fulfilling the demands. Furthermore, as g satisfies (11), we have the same inequality for g' and c', that is, $g'(e_{i-1}) + g'(e_i) \leq d_h(v_i) + 2c'(v_i)$. Lemma 4.2 now shows, when applied to g' and c', that x is node-feasible, as well. •

As far as algorithmic aspects are concerned, we mentioned earlier that there are simple algorithms for edge-capacitated ring-routing whose complexity is actually linear in $\max(|E|, |F|)$. The proof of the main theorem showed a reduction technique for the node-capacitated case. It required a subroutine to find a feasible potential or a negative circuit in a digraph D on $2n$ nodes. This can be achieved in polynomially bounded time, for instance, by applying the well-known Bellman-Ford algorithm which performs $O(n^3)$ basic operations. Thus in $O(n^3)$ time we can calculate a (half-integral) routing or find a double-cut whose corresponding inequality is violated.

6 Conclusion

This work grew out of investigations into the viability of so-called passive optical rings in access and metro-area networks. These networks are passive in the sense that if a certain node is transmitting on a given wavelength, then other nodes on the ring receive its signal (and so should not transmit on the same wavelength). This inefficient use of wavelengths led researchers at BT Labs to propose the use of absorbers at each node which can block specific wavelength signals from propagating in either the clockwise, or counterclockwise direction from that node. This allows a wavelength to be reused for several communication paths as long as the paths are node-disjoint. Thus the node-disjoint multicommodity problem is at the core of understanding routing algorithms for this proposed technology (although in practice there are several extra requirements not least of which is the need to assign wavelengths to the paths). In contrast, most theoretical routing problems arising from telecommunications are based on finding edge-disjoint paths. Even on the theoretical side, node-capacitated problems have been traditionally neglected; two prime exceptions are Menger's Theorem (and network flow theory), where there is a single commodity pair, and the work of Robertson and Seymour on graph minors, which applies to the case where capacities are all 1.

A general approach to handle integer-valued packing (and covering) problems is to associate first a polyhedron Q with the problem and then show the integrality of Q. A typical way to achieve this is to prove total dual integrality of an inequality system describing Q. The literature is quite rich in this type of result but important cases do not fit into this framework. Indeed, no general theory is known to derive half-integrality packing results as special cases. For example, the Lucchesi-Younger Theorem can be derived nicely from the theory submodular flows, whereas results such as the Okamura-Seymour Theorem are not known to have such a common source. In order to understand better the general behaviour of half-integrality results, it is highly desirable to look for new occurrences. This is another background motivation of our research.

The mathematical tool, shortest paths, we used is a common routine, with the advantage that long developed efficient algorithms are available for actual computation. On the other hand, the way in which shortest paths have been used for the routing problem is rather tricky, and these in combination may lend some merit to this work.

From a practical point of view, Theorem 5.1 seems to be quite satisfactory. It gives rise to an approximation algorithm in the sense that if there is a feasible node-capacitated ring routing fulfilling the demands, then the algorithm either finds one or finds an almost feasible routing, that is, one which fulfils the demand and violates every node-capacity constraint by at most one. From a theoretical point of view, it is certainly an interesting challenge to develop an algorithm which finds a routing if one exists, along with a necessary and sufficient condition for the existence of such a routing. A little encouragement in this direction is Corollary 2.4 showing that the edge-capacitated version of the ring routing problem is tractable.

Acknowledgements: We are very grateful to the referees whose suggestions greatly improved this article.

References

[1] Ford, L. R., Fulkerson, D. R. 1962. *Flows in Networks.* Princeton University Press, Princeton, NJ.

[2] A. Frank. 1990. Packing paths, circuits, and cuts - a survey. B. Korte, L. Lovász, H-J. Prömel, A. Schrijver, eds., *Paths, Flows and VLSI-Layouts*, Springer Verlag, 47-100.

[3] A. Frank. 1995. Connectivity and network flows. R. Graham, M. Grötschel and L. Lovász eds., *Handbook of Combinatorics*, Elsevier Science B.V., 111-177.

[4] A. Frank. 1985, Edge-disjoint paths in planar graphs. *J. of Combinatorial Theory Ser. B.* **39** 164-178.

[5] Iri, M. 1971. On an extension of the maximum-flow minimum cut theorem to multicommodity flows. *J. Operations Research Soc. of Japan.* **13** No. 3

[6] C. L. Lucchesi, D. H. Younger. 1978. A minimax relation for directed graphs. *J. London Math. Society.* (2) 17 369-374.

[7] K. Onaga, O. Kakusho. 1971. On feasibility conditions of multicommodity flows in Networks. *IEEE Trans. Circuit Theory.* **18** 425-429.

[8] H. Okamura, P. D. Seymour. 1981. Multicommodity flows in planar graphs. *J. Combinatorial Theory, Ser. B.* **31** 75-81.

[9] A. Schrijver, P. D. Seymour, P. Winkler. 1998. The ring loading problem. *SIAM J. Discrete Math.* **11** No 1 1-14.