Multiflows in Multihop Wireless Networks

[Extended Abstract]

Peng-Jun Wan*
Department of Computer Science
Illinois Institute of Technology
Chicago, IL 60616
wan@cs.iit.edu

ABSTRACT

This paper studies maximum multicommodity flow and maximum concurrent flow in multihop wireless networks subject to both bandwidth and interference constraints. The existing proof of the NP-hardness of both problems is too contrived to be applicable to meaningful multihop wireless networks. In addition, all known constant-approximation algorithms for both problems restricted to various network classes are super-exponential in running time. Some of them are simply incorrect. In this paper, we first provide a rigorous proof of the NP-hardness of both problems even in very simple settings. Then, we show that both problems restricted to a broad family of multihop wireless networks admit polynomial-time approximation scheme (PTAS). After that, we develop a unified framework for the design and analysis of polynomial approximation algorithms for both problems. Following such framework, we obtain polynomial constant-approximation algorithms for both problems restricted to a broad network family. The approximation ratios of these algorithms are also better than those known in the literature.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*wireless communication*; F.2.0 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity—*General*

General Terms

Algorithms, Theory

1. INTRODUCTION

In this paper, we study **Maximum Multiflow** (**MMF**) and **Maximum Concurrent Multiflow** (**MCMF**) in single-radio single-channel multihop wireless networks subject to both bandwidth and interference constraints. A single-radio single-channel multihop wireless network N is specified, in its most general format, by a triple (D, G, b), where D is a directed graph representing the communication topology of N, G is an undirected graph representing the conflict graph of the (communication) links in D, and b is the bandwidth function of the links in D. A set of links in D are said to be *conflict-free* if they are not adjacent pairwise in G. Consider a multihop wireless network $N = (D, G, b)$ with $D = (V, A)$. We use $\mathcal{I}$ to denote the collection of sets of conflict-free links in D. A *(fractional) link schedule* in N is a set

$$S = \{(I_j, \lambda_j) : 1 \leq j \leq k\}$$

with $I_j \in \mathcal{I}$, and $\lambda_j \in \mathbb{R}_+$ for each $1 \leq j \leq k$. The value $\sum_{j=1}^{k} \lambda_j$ is referred to as the *length* (or *latency*) of S, and $|S|$ is called the *size* of S. Any link schedule S in N of length at most one determines a link capacity function $c_S \in \mathbb{R}_+^A$ of D given by

$$c_S(e) = b(e) \sum_{1 \leq j \leq k} \lambda_j |I_j \cap \{e\}|$$

for each $e \in A$. Suppose that we are given with a set of commodities in D. For any link schedule S in N of length at most one, the maximum multiflow of these commodities in D subject to the capacity function c_S is referred to as the *maximum multiflow subject to* S. Suppose in addition that each commodity also has a demand associated with it. For any link schedule S in G of length at most one, the maximum concurrent multiflow of these commodities in D subject to the capacity function c_S is referred to as the *maximum concurrent multiflow subject to* S.

Now, we give a clean definition of the two optimization problems to be studied in this paper. Let $\mathcal{N}$ be a class of multihop wireless networks.

- **MMF** restricted to $\mathcal{N}$: Given a network $N = (D, G, b) \in \mathcal{N}$ and a set of commodities in D, find a link schedule S in N of length at most one such that the maximum multiflow subject to S is maximized.
- **MCMF** restricted to $\mathcal{N}$: Given a network $N = (D, G, b) \in \mathcal{N}$ and a set of commodities with demands in D, find a link schedule S in N of length at most one such that the maximum concurrent multiflow subject to S is maximized.

We will study both problems restricted to two network classes termed by *802.11 class* and *PIM class*, corresponding to the 802.11 interference model and the protocol interference model respectively. For both classes, an instance of a network N is specified by a finite planar set V of nodes together with a communication radius function $r \in \mathbb{R}_+^V$ and

*This work was partially supported by NSF under grant CNS-0831831.

MobiHoc'09, May 18–21, 2009, New Orleans, Louisiana, USA.

an interference radius function $\rho \in \mathbb{R}_+^V$. The communication (resp., interference) range of a node $v \in V$ is the disk centered at v of radius $r(v)$ and $\rho(v)$ respectively. The communication topology of N is the digraph $D = (V, A)$, in which there is an link from u to v if and only if v is within the communication range of u. If N belongs to the 802.11 class, two links in A are adjacent in the conflict graph G of N if and only if at least one link has an endpoint lying in the interference range of some endpoint of the other link. If N belongs to the PIM class, two links in A are adjacent in the conflict graph G of N if and only if the receiving endpoint of at least one link lies in the interference range of the transmitting endpoint of the other link.

1.1 Prior Works

Both **MMF** and **MCMF** restricted to various classes of multihop wireless networks have been studied in many recent works. Kodialam and Nandagopal [13] [14] studied both problems restricted to a network class in which the only interference constraint is that node may not transmit and receive simultaneously. Jain et al. [10] presented methods for computing upper and lower bounds on the maximum single-flow, but they didn't provide any polynomial algorithm for computing an approximation solution. In the same paper, the gave a very "unusual" proof for the NP-hardness of the maximum single-flow even in single-hop wireless networks under the protocol interference model. Even the authors themselves commented after their proof that "the above proof may come across as contrived since the wireless network we constructed is unlikely to arise in practice". Indeed, their proof of the NP-hardness is too contrived to be applicable to meaningful multihop wireless networks, and cannot even imply the NP-hardness of finding a maximum multiflow in wireless networks with a constant number of channels and a constant number of radios per node. We will discuss on this in Appendix 1 for the fairness. Nevertheless, such NP-hardness has been cited by many subsequent works for granted.

Kumar et al. [15] studied multiflow restricted to three classes of networks, different from the ones studied in this paper, corresponding to the following three interference models (in their term) respectively:

- Transmitter model with parameter $\rho \geq 1$: two links u_1v_1 and u_2v_2 are conflict-free if and only if $\|u_1u_2\| > \rho(r(u_1) + r(u_2))$.
- Protocol model with parameter $\rho \geq 1$: two links u_1v_1 and u_2v_2 are conflict-free if and only if $\|u_1v_2\| > \rho\|u_1v_1\|$ and $\|u_2v_1\| > \rho\|u_2v_2\|$.
- Transmitter-receiver model: two links u_1v_1 and u_2v_2 are conflict-free if and only if $\|u_1u_2\| > \max\{r(u_1), r(u_2)\}$ and $\|v_1v_2\| > \max\{r(v_1), r(v_2)\}$.

They developed constant-approximations by enforcing restrictive interference constraints on *links* to guarantee flow schedulability. However, we found that a key step in their algorithm for link scheduling (Section 3.2 in [15]) may have *super-exponential* running time, and we will provide a detailed explanation on this time complexity in Appendix 1 for the fairness.

Following the same approach (i.e. enforcing restrictive interference constraints on links to guarantee schedulability of the underlying flow), Alicherry et al. [1] and Wang et al. [21] obtained constant-approximations in other network classes of multihop wireless networks. Specifically, Alicherry et al. [1] considered the general multi-channel multi-radio multihop wireless network under 802.11 interference model in which all nodes have *uniform* communication radii, normalized to one, and *uniform* interference radius $\rho \geq 1$. For the single-channel single-radio configuration, they gave (in Lemma 1 in [1]) an approximation bound of $4, 8$ and 12 for $\rho = 1, 2, 2, 5$ resp., and claimed that in general the approximation bound is a constant growing with ρ. The bound of 4 for $\rho = 1$ is wrong and should be 8 (such false claim was also discovered by [3]). They also inherited the same link scheduling algorithm (VI.B in [1]) from [15] as part of their algorithm, which renders their algorithm to have *super-exponential* running time in the worst case.

Wang et al. [21] studied the multiflow restricted to either the 802.11 class or the PIM class. For the 802.11 class, their approximation bound is 120; for the subclass of the PIM class in which the interference radius of each node is at least c times its communication radius for some fixed constant $c > 1$, their approximation bound is $2\left\lceil 2\pi/\arcsin\frac{c-1}{2c}\right\rceil$. Once again, their constant-approximation algorithms have *super-exponential* running time in the worst case, as they used the same link scheduling algorithm (Section 6.1 in [21]) as in [1] and [15]. A serious mistake is that their latter result for the subclass of the PIM class is wrong. The restrictive interference constraints on links in their algorithm design and analysis *fails* to guarantee the flow schedulability. We will provide a counter-example in Appendix 1 to illustrate such failure.

Buragohain et al. [3] targeted at better approximation bounds for multiflows restricted to the subclass of the 802.11 class in which all nodes have uniform communication radii and uniform interference radii. They introduced less restrictive interference constraints on the *node* level to guarantee the flow schedulability, and were able to achieve an approximation bound of 3, which is an improvement on the approximation bound obtained earlier by Alicherry et al. [1]. But unfortunately, the interference constraints introduced by them are too less restrictive to guarantee the flow schedulability, and we will once again provide a counter-example for this failure in Appendix 1. Thus, their approximation algorithm is also wrong, and their approximation bound of 3 is invalid. Furthermore, they didn't provide any polynomial algorithm for link scheduling. If the same algorithm for link scheduling in [1] was adopted, then their algorithm wouldn't be polynomial either.

1.2 Our Contributions

The original purpose of this paper is to simply develop better approximation algorithms for **MMF** and **MCMF** restricted to either the 802.11 class or the PIM class. But a thorough literature review put us in a very awkward and uncommon situation. Almost all major prior results on the same subject have some technical bugs or even mistakes. It took us much effort to debug those bugs. For a purely scientific and fair treatment, we include in Appendix 1 detailed explanations on those bugs and mistakes identified in the previous subsection. We were then pressed to conduct a comprehensive study more than just better approximation algorithms. The following results are reported in this paper.

- NP-hardness: We provide a rigorous proof of the NP-hardness even restricted the subclass in which all nodes have uniform (and fixed) communication radii and uniform (and fixed) interference radii.
- Approximation hardness: We show that for the 802.11 class and two subclasses of the PIM class, both maximization problems admit a polynomial approximation scheme (PTAS). In other words, for any fixed $\varepsilon > 0$,

there is a polynomial-time (depending on ε) $(1+\varepsilon)$-approximation algorithm for each of them. However, these PTAS's are of only theoretical interest and are quite infeasible practically.

- Faster and better polynomial approximation algorithms: We first develop a unified framework for both design and analysis of polynomial approximation algorithms. Following such general framework, we obtain improved approximations in the following classes of networks:

  1. For the 802.11 class, we obtain a 23-approximation algorithm. This is a significant cut-down from the approximation bound of 120 derived by Wang et al. in [21].
  2. For the subclass of the 802.11 class in which all nodes have uniform communication radii, normalized to one, and uniform interference radii $\rho \geq 1$, a 7-approximation is obtained regardless of ρ.
  3. For the subclass of the PIM class in which the interference radius of each node is at least c times its communication radius, we obtain a $2\left(\lceil \pi/\arcsin\frac{c-1}{2c}\rceil - 1\right)$-approximation algorithm.

At the end of this section, we introduce some standard notations and terms used throughout this paper. The disk centered at a node v of radius r is denoted by $B(v,r)$. For a finite planar point set V and a number $r>0$, the r-disk graph on V is a simple geometric graph on V in which there is an edge between two nodes if and only if their distance is at most r. In particular, a 1-disk graph is referred to as unit-disk graph. Let Π be a finite subset. For any real function $f \in \mathbb{R}^{\Pi}$ and any subset $\Pi' \subseteq \Pi$, $f(\Pi')$ denotes $\sum_{e\in\Pi'} f(e)$. For two functions $f,g \in \mathbb{R}^{\Pi}$ with $g(e) \neq 0$ for each $e \in \Pi$, we use f/g to denote the function in $\mathbb{R}^{\Pi}$ defined by $(f/g)(e) = f(e)/g(e)$.

2. PRELIMINARIES

Let $G=(V,E)$ be an undirected graph. A subset I of V is an *independent set* (IS) of G if no two nodes in I are adjacent. If I is an IS of G but no proper superset of I is an IS of G, then I is called a *maximal IS* of G. Any node ordering $\langle v_1, v_2, \cdots, v_n\rangle$ of V induces a maximal IS I in the following first-fit manner: Initially, $I=\{v_1\}$. For $i=2$ up to n, add v_i to I if v_i is not adjacent to any node in I. An IS of the largest size is called a *maximum IS*. Let $\mathcal{I}$ be the collection of all independent sets of G. The *independence number* of G, denoted by $\alpha(G)$, is defined to be $\max_{I\in\mathcal{I}}|I|$. For any $d \in \mathbb{R}_+^V$, the *(weighted) independence number* of (G,d), denoted by $\alpha(G,d)$, is defined to be $\max_{I\in\mathcal{I}} d(I)$. For any $d\in\mathbb{R}_+^V$, a *fractional coloring* of (G,d) is a set of k pairs (I_j,λ_j) with each $I_j\in\mathcal{I}$ and $\lambda_j\in\mathbb{R}_+$ for $1\le j\le k$ satisfying that for each $v\in V$,

$$\sum_{1\le j\le k, v\in I_j,} \lambda_j = d(v).$$

The two values k and $\sum_{j=1}^{k}\lambda_j$ are referred to as the number and total weight of the coloring respectively. The *fractional chromatic number* $\chi_f(G,d)$ of (G,d) is defined as the minimum weight of all fractional colorings of (G,d). It's obvious that

$$\chi_f(G,d) \geq \frac{d(V)}{\alpha(G)}.$$

The *independence polytope* P of G is the convex hull of the incidence vectors of independent sets in G. Equivalently, it consists of all $d\in\mathbb{R}_+^V$ with $\chi_f(G,d)\le 1$. A polytope Q is said to be a μ-approximation of P for some $\mu>1$ if $Q\subseteq P\subseteq \mu Q$.

Now, we describe several optimization problems. Let $\mathcal{G}$ be a class of graphs.

- **Maximum Independent Set** (**MIS**) restricted to $\mathcal{G}$: Given any $G\in\mathcal{G}$, find an IS I of G with $|I|=\alpha(G,d)$.
- **Maximum Weighted Independent Set** (**MWIS**) restricted to $\mathcal{G}$: Given any $G\in\mathcal{G}$ and $d\in\mathbb{R}_+^{V(G)}$, find an IS I of G with $d(I)=\alpha(G,d)$.
- **Minimum Fractional Weighted Coloring** (**MFWC**) restricted to $\mathcal{G}$: Given any $G\in\mathcal{G}$ and $d\in\mathbb{R}_+^{V(G)}$, find a fractional coloring of (G,d) with total color weight equal to $\chi_f(G,d)$.

The next theorem follows from general theorems on separation and optimization given by Grötschel et al. [6] in the exact case and by Jansen [11] in the approximation case.

THEOREM 2.1. *For any class $\mathcal{G}$ of graphs,*

1. *there is a polynomial algorithm for **MFWC** restricted to $\mathcal{G}$ if and only there is a polynomial algorithm for **MWIS** restricted to $\mathcal{G}$;*
2. *if there is a polynomial μ-approximation algorithm for **MWIS** restricted to $\mathcal{G}$, then there is a polynomial μ-approximation algorithm for **MFWC** restricted to $\mathcal{G}$.*

Now, suppose that $\mathcal{G}$ is the class of conflict graphs of networks in a class $\mathcal{N}$ of multihop wireless networks. The problem **MIS** (resp., **MWIS**, **MFWC**) restricted to $\mathcal{G}$ is referred to as **Maximum Conflict-Free Links** (**MCFL**) (resp., **Maximum Weighted Conflict-Free Links** (**MWCFL**), **Minimum Fractional Weighted Link Schedule** (**MFWLS**)) restricted to $\mathcal{N}$.

3. NP-HARDNESS

In this section, we will establish the following four hardness results.

THEOREM 3.1. *Even restricted to the subclass of the 802.11 class or the PIM class in which all nodes have uniform (and fixed) communication radii and uniform (and fixed) interference radii and the positions of all nodes are available, all of the following problems are NP-hard: (1) **MWCFL**, (2) **MFWLS**, (3) **MMF**, and (4) **MCMF**.*

Our reductions make use of the NP-hardness of three optimization problems restricted to UDGs.

LEMMA 3.2. *Restricted to UDGs in which the positions of all nodes are available, all of the following three problems are NP-hard (1) **MIS**, (2) **MWIS**, and (3) **MFWC**.*

The NP-hardness of the first problem in the above lemma was proved in [20], which then implies the NP-hardness of the second problem, which further implies the NP-hardness of the third problem by Theorem 2.1. We first present a construction of a (connected) multihop wireless network from a connected UDG $G=(V,E)$ and a fixed constant $\rho\ge 1$.

Let L be the distance between the closest pair of nodes in V that are not adjacent in G. Set $\rho' = \frac{L+2}{3}$ and $r' = \frac{\rho'}{\rho}$.

Then $1 < \rho' < L$. We first construct a set W of at most $\rho(|V|-1)$ points such that the r'-disk graph on $V \cup W$ is connected. Compute an Euclidean minimum spanning tree T of G. Since G is connected, all edges of T have length at most one. Initially, W is empty. We subdivide each edge uv in T with $\|uv\| > r'$ into $\lceil \|uv\| / r' \rceil$ segments of equal length and adding those $\lceil \|uv\| / r' \rceil - 1$ endpoints of these segments other than u and v to W. Since $\rho' > 1$, we have

$$\left\lceil \frac{\|uv\|}{r'} \right\rceil - 1 \leq \left\lceil \frac{1}{r'} \right\rceil - 1 = \left\lceil \frac{\rho}{\rho'} \right\rceil - 1 \leq \left\lfloor \frac{\rho}{\rho'} \right\rfloor < \rho.$$

Thus, $|W| \leq \rho(|V|-1)$. In addition, the r'-disk graph on $V \cup W$ is connected.

Now, we construct a copy V' of V as follows. Let l be the distance between the closest pair of nodes in $V \cup W$ and set

$$\sigma = \min\left\{ r', \frac{L-1}{3}, \frac{l}{2} \right\}.$$

Then, $1+\sigma \leq \rho' < L-\sigma$. For each $v \in V$, we make a copy v' satisfying that v' is straightly below v and $\|vv'\| = \sigma$. Then, $v' \notin V \cup W$. Let V' denote the set of copies v' constructed in this way. Clearly, $|V'| = |V|$.

Then, the multihop wireless network consists of all nodes $V \cup W \cup V'$ in which each node has communication radius r' and interference radius $\rho' = \rho \cdot r'$. Let G' be the r'-disk graph on $V \cup W \cup V'$. Then, G' is connected, and the communication topology $D = (V, A)$ of the network is the directed version of G'. All links in D have the same link capacity normalized to one unit. Additional geometric properties of the network are given in the next two claims.

Claim 3.3. *Let u and v be any two distinct nodes in V. If $\|uv\| \leq 1$, then*

$$\max\{\|uv\|, \|u'v'\|, \|uv'\|, \|u'v\|\} \leq \rho';$$

otherwise,

$$\min\{\|uv\|, \|u'v'\|, \|uv'\|, \|u'v\|\} > \rho'.$$

Claim 3.4. *Let v be a node in V and x be any node in $V \cup W \cup V'$. If $\|vx\| \leq r'$, then $\|v'x\| \leq \rho'$.*

The above two claims can be easily verified and their proofs are omitted due to the space limitation.

For each subset U of V, we denote

$$A_U = \{(v', v) : v \in U\}.$$

Claim 3.3 implies that U is an independent set of G if and only if all links in A_U are conflict-free under either 802.11 model or the protocol interference model. Claim 3.4 implies that at any time when v' is communicating with some node other than v, then v must be idle under either 802.11 model or protocol model. These two implications are essential to the correctness of our following reductions.

Consider an arbitrary $d \in \mathbb{R}_+^V$. By necessary scaling, we assume that for each $v \in V$, either $d(v) = 0$ or $d(v) \geq 1$ (where 1 represents the normalized unit link capacity). Such scaling does not change the computational complexity. Let $d' \in \mathbb{R}_+^A$ be such that $d'(e) = d(v)$ for any $e = (v', v) \in A_V$ and $d'(e) = 0$ for any other $e \in A \setminus A_V$. We make the following three claims:

1. By treating d' as the link-weight function of D, the maximum weight of conflict-free links in D is $\alpha(G, d)$.
2. By treating d' as the link-demands, the minimum length of the fractional weighted schedule for (D, d') is $\chi_f(G, d)$.
3. By treating d' as the commodity traffic demands, the maximum concurrency of these demands in D is $1/\chi_f(G, d)$.

These relations together with Lemma 3.2 (2) and (3) imply the NP-hardness of the three problems in Theorem 3.1 except the third one. The proofs of the first two relations are easy and so are skipped. In the sequel, we give the proof of the third one, which is more complicated because the routing is involved. We denote by ϕ the maximum concurrency.

First we show that $\phi \geq 1/\chi_f(G, d)$. Consider a minimum fractional coloring of (G, d) given by k pairs (I_j, λ_j) for each $1 \leq j \leq k$. Then,

$$\sum_{j=1}^{k} \lambda_j = \chi_f(G, d).$$

We construct the following routing and schedule for the commodities in A_V. Each commodity in A_V takes a single-hop route. For each $1 \leq j \leq k$, we schedule all the links in A_{I_j} for $\lambda_j / \chi_f(G, d)$ amount of time. Then, $d(v) / \chi_f(G, d)$ units of commodity from v' to v is transported in a unit of time. Hence, the concurrency of such schedule is exactly $1/\chi_f(G, d)$, and hence $\phi \geq 1/\chi_f(G, d)$.

Next, we show that $\phi \leq 1/\chi_f(G, d)$. The key observation is that we can convert any routing and schedule with concurrency ϕ to a "canonical" one in which the traffic of each commodity is transported directly in one hop from the source to the sink without sacrificing the concurrency. Indeed, suppose that a portion of traffic from some node v' to v is routed along a path p rather than the direct link $v'v$. By Claim 3.4, v must be idle whenever v' is transmitting. We replace the path p by the direct link $v'v$.The same amount of traffic carried along p would be routed along $v'v$ but without changing the time schedule. This new routing and schedule does not affect any other traffic transported concurrently. By repeatedly taking such switching operation, we obtain a routing and schedule of the same concurrence ϕ in which all commodities are transported along the direct link. Now, we decompose the schedule into a set of pairs (A_j, λ_j) for $1 \leq j \leq k$ where a period of duration λ_j is dedicated to links in A_j which concurrently. Then,

$$\sum_{1 \leq j \leq k} \lambda_j = 1,$$

and each A_j is a set of conflict-free links. Let I_j denote the set of receiving nodes of the links in A_j. Then, each I_j is an independent set of G. In addition, for each $v \in V$ we have

$$\sum_{v \in I_j, 1 \leq j \leq k} \lambda_j = \sum_{(v', v) \in A_j, 1 \leq j \leq k} \lambda_j = \phi \cdot d(v),$$

which implies that

$$\sum_{v \in I_j, 1 \leq j \leq k} \frac{\lambda_j}{\phi} = d(v).$$

This means that the set of k pairs $(I_j, \lambda_j / \phi)$ for $1 \leq j \leq k$ is a fractional weighted coloring of (G, d). Thus,

$$\chi_f(G, d) \leq \sum_{j=1}^{k} \frac{\lambda_j}{\phi} = \frac{1}{\phi} \sum_{j=1}^{k} \lambda_j = \frac{1}{\phi},$$

So, $\phi \leq 1/\chi_f(G, d)$.

Therefore, $\phi = 1/\chi_f(G, d)$ and hence the third claimed relation is true.

Finally, we claim that by treating A_V as the $|V|$ commodities, the maximum multiflow of these commodities is $\alpha(G)$. The proof of this claim is almost the same as the proof of the third relation above by using the concept of "canonical" routing and schedule. Therefore, the proof is omitted. This relation together with Lemma 3.2(1) implies the NP-hardness of the third problem in Theorem 3.1. This completes the proof of Theorem 3.1.

We conclude this section by remarking that our reduction can actually be used to show the NP-hardness of **MWCFL** even restricted to $\{0,1\}$-weight.

4. POLYNOMIAL-TIME APPROXIMATION SCHEMES

While the NP-hardness established in the previous section brought us negative news, the approximation hardness of the same problems should bring us good news. We start with the following general results on approximality.

THEOREM 4.1. *Suppose that $\mathcal{N}$ is a network class satisfying that there is a polynomial (resp., a polynomial μ-approximation) algorithm for **MWCFL** restricted to $\mathcal{N}$. Then, there is a polynomial (resp., a polynomial μ-approximation) algorithm for each of all of following three problem restricted to $\mathcal{N}$: (1) **MFWLS**, (2) **MMF**, and (3) **MCMF**.*

PROOF. For **MFWLS** restricted to $\mathcal{N}$, the theorem is implied by Theorem 2.1. So, we move on to the other two problems. For simplicity of presentation, we treat a polynomial algorithm which produces an optimal solution as a polynomial 1-approximation algorithm. Let $\mathcal{A}$ be a μ-approximation algorithm for **MWCFL** restricted to $\mathcal{N}$. The proof leverages an ellipsoid method for exponential-sized linear program (LP) with an (approximate) separation oracle. Given an approximate separation oracle for the dual LP of a primal LP, both the primal LP and the dual LP can be solved with the ellipsoid method within the same approximation factor as the approximate separation oracle. This powerful technique was first used by Karmarkar and Karp for the bin packing problem [12] and has been successfully applied to produce approximation algorithms for a number of other optimization problems (see, e.g., [4], [9], [11], [18].)

Let $\mathcal{A}$ be a μ-approximation algorithm for **MWCFL** restricted to $\mathcal{N}$. Consider a network $N = (D, G, b) \in \mathcal{N}$. Suppose that $D = (V, A)$. We use $\mathcal{I}$ to denote the collection of sets of conflict-free links in A. In addition, we are given with k commodities with s_j, t_j being the source and sink, respectively, for commodity j for $1 \leq j \leq k$. Let $\mathcal{P}_j$ be the set of (s_j, t_j)-paths in $D = (V, A)$ for all $1 \leq j \leq k$, and define $\mathcal{P}$ to be the union of $\mathcal{P}_1, \cdots, \mathcal{P}_k$. Also, let $\mathcal{P}_e$ be the set of paths in $\mathcal{P}$ that use link e for all $e \in A$.

Now, we give the proof for **MMF**. The path-flow LP formulation for **MMF** has a variable $x(p)$ for the flow sent along each path $p \in \mathcal{P}$ and a variable $y(I)$ each independent set $I \in \mathcal{I}$:

$$
\begin{array}{lll}
(P_{mmf}) & \max & \sum_{p\in\mathcal{P}} x(p) \\
 & s.t. & \sum_{p\in\mathcal{P}_e} x(p) \leq b(e) \sum_{e\in I} y(I), \forall e \in A \\
 & & \sum_{I\in\mathcal{I}} y(I) \leq 1 \\
 & & x, y \geq 0
\end{array}
$$

The dual to this LP associates a length $l(e)$ for each link $e \in A$ and another variable ω:

$$
\begin{array}{lll}
(D_{mmf}) & \min & \omega \\
 & s.t. & \sum_{e\in p} l(e) \geq 1, \forall p \in \mathcal{P} \\
 & & \sum_{e\in I} b(e) l(e) \leq \omega, \forall I \in \mathcal{I} \\
 & & l, \omega \geq 0
\end{array}
$$

The dual LP can be interpreted the as follows. Let $dist_j(l)$ be the length of the shortest (s_j, t_j)-path in D with respect to length function $l \in \mathbb{R}_+^A$ for $1 \leq j \leq k$. Also let

$$
\alpha(l) = \min_{1\leq j\leq k} dist_j(l),
$$
$$
D(l) = \max_{I\in\mathcal{I}} \sum_{e\in I} b(e) l(e) = \max_{I\in\mathcal{I}} l(I),
$$

Then, $\alpha(l)$ is the minimum length of the shortest paths between all pairs of terminals of the commodities, and $D(l)$ is the maximum weight of conflict-free links in the conflict graph of A in which each link $e \in A$ has weight $b(e) l(e)$. Thus, (D_{mmf}) is equivalent to finding a length function $l \in \mathbb{R}_+^A$ such that $D(l)$ is minimized subject to $\alpha(l) \geq 1$.

We run the ellipsoid algorithm on the dual LP using $\mathcal{A}$ as the approximate separation oracle. More precisely, we use binary search to find the smallest value of ω for which the dual linear program is feasible. The separation oracle acts as follows: First, we compute $dist_j(l)$ for each $1 \leq j \leq k$, and then compute $\alpha(l)$. We consider two cases.

Case 1: $\alpha(l) < 1$. Then l is not feasible. Let j be such that $\alpha(l) = dist_j(l)$ and p be the shortest (s_j, t_j)-path with respect to l. Then, the constraint corresponding to p is a separating hyperplane.

Case 2: $\alpha(l) \geq 1$. We apply $\mathcal{A}$ to compute an $I \in \mathcal{I}$ with respect to l. If

$$
\sum_{e\in I} b(e) l(e) > \omega,
$$

then l is not feasible and the constraint corresponding to I is a separating hyperplane; otherwise, we accept l as a feasible solution and therefore the ellipsoid algorithm decides that the LP is feasible.

Of course, since $\mathcal{A}$ is just an approximation algorithm, the above conclusion might be incorrect, and the dual LP might actually be infeasible. However, since the approximation factor of $\mathcal{A}$ is at most μ, we know that in this case, l and $\mu\omega$ constitute a feasible solution of the dual LP. Therefore, if ω^* is the minimum value of ω for which the algorithm decides that the dual LP is feasible, then we know that the dual LP is infeasible for $\omega^* - \epsilon$ (where ϵ depends on the precision of the algorithm), and is feasible for $\mu\omega^*$. Therefore, the value of the dual LP, and hence the value of the primal LP as well, is between ω^* and $\mu\omega^*$.

The above algorithm computes the approximate value of the primal LP. In order to compute the actual solution, we use the technique used in [4] [9]. The total number of separating hyperplanes found by the above separation oracle while running the ellipsoid algorithm for $\omega^* - \epsilon$ is bounded by a polynomial. These separation oracles are enough to show that the value of the dual LP is at least ω^*. Therefore, if we consider the set of primal variables that correspond to these separating hyperplanes, we get a set of polynomial many primal variables. By LP-duality, if we fix the values of the other variables to 0, the resulting LP still has solution at least ω^*. However, after fixing the values of other variables to 0 we obtain a polynomial size LP, which we can solve in polynomial time, and find the optimum solution. By the above argument this optimum solution has value at least ω^*, and thus is a μ-approximation.

The proof for **MCMF** uses the similar method. Due to the space limitation, we omit the proof in this paper. □

The next theorem gives three network classes restricted to which the problem **MWCFL** has a PTAS.

THEOREM 4.2. *The problem **MWCFL** has a PTAS when restricted to any of the following three network classes:*

1. *the 802.11 class;*
2. *the subclass of the PIM class in which the interference radius of each node is at least c times its communication radius for some constant $c > 1$;*
3. *the subclass of the PIM class in which every k-hop neighborhood in the conflict-graph contains at most $O(k^c)$ conflict-free links for some constant $c > 0$.*

The proof of the above theorem is omitted mainly for the following reason. The PTAS for the first two classes is an almost verbatim repetition of the PTAS for **MIS** restricted to disk graphs proposed in [5] or a faster one proposed in [19], both of which utilize the shifting strategy [7] [8] [16] combined with dynamic programming. The PTAS for the third class is also an almost verbatim repetition of the PTAS for **MIS** restricted to disk graphs proposed in [17]. We also would like to acknowledge that a PTAS for the *unweighted* variant of **MWCFL** restricted to the 802.11 class was given in [2] in the same vein.

From Theorem 4.1 and Theorem 4.2, we obtain the following corollary.

Corollary 4.3. *Restricted to any of the three network classes given in Theorem 4.2, all of the following three problems have a PTAS: (1)* ***MFWLS****, (2)* ***MMF****, and (3)* ***MCMF****.*

To conclude this section, we remark that the PTAS's presented in this section are of theoretical interest only, but are practically quite infeasible. In the coming sections, we will develop practically feasible constant-approximation algorithms for those problems.

5. FRACTIONAL WEIGHTED COLORING AND INDEPENDENCE POLYTOPES

In this section, we temporally divert from multihop wireless networks to developing a general graph-theoretic algorithm for fractional weighted coloring and constructing polytopes with linear explicit representations approximating the independence polytope within guaranteed factors.

Let $G = (V, E)$ be an undirected graph. Consider a vertex ordering $\langle v_1, v_2, \cdots, v_n \rangle$ of V. For any $U \subseteq V$,the *first-fit MIS* of U is the MIS of $G[U]$ selected in the first-fit manner in the ordering $\langle v_1, v_2, \cdots, v_n \rangle$. Our algorithm, referred to as **First-Fit Fractional Weighted Coloring (F^{3}WC)**, takes as input a graph $G = (V, E)$ together with an ordering $\langle v_1, v_2, \cdots, v_n \rangle$ of V and a node-demand (or weight) function $d \in \mathbb{R}_+^V$. The description of this algorithm is given in Table 1. The idea of this algorithm is simple. It runs in iterations. In each iteration, a first-fit MIS of remaining nodes with positive residue demand is selected, and then assign to this MIS a color with a weight to "saturate" at least one node. This ensures that at least one node gets its demand satisfied and stops moving onto the subsequent iteration. As a result, the number of iterations, or equivalently the number of colors, is bounded by n, and the running time is $O(n^2)$. The next theorem gives a bound on the total color weight of the output coloring.

Theorem 5.1. *The coloring output by* ***F^{3}WC*** *uses at most n colors of total weight at most* $\max_{1 \le i \le n} d(V_i)$*, where* V_i *consists of* v_i *and all its neighbors in* $\{v_1, v_2, \cdots, v_{i-1}\}$ *for each* $1 \le i \le n$.

First-Fit Fractional Weighted Coloring (F^{3}WC)
Input: a graph $G = (V, E)$, $d \in \mathbb{R}_+^V$,
and an ordering $\langle v_1, v_2, \cdots, v_n \rangle$ of V
Output: a fractional weighted coloring Π of (G, d).
Begin
$\Pi \leftarrow \emptyset$;
$U \leftarrow \{v \in V : d(v) > 0\}$;
while $U \neq \emptyset$ do
$I \leftarrow$ the first-fit MIS of U;
$\lambda \leftarrow \min_{v \in I} d(v)$;
add (I, λ) to Π;
for each $v \in U$,
$d(v) \leftarrow d(v) - \lambda$;
if $d(v) = 0$, remove v from U;
output Π;
End

Table 1: The first-fit algorithm for fractional weighted coloring.

Proof. Suppose that the algorithm runs in k iterations. For each $1 \le j \le k$, let U_j be the subset U (of nodes with residue demands) at the beginning of the j-th iteration, and (I_j, λ_j) be the pairs of independent set and color weight selected in the j-th iteration. Since at least one node gets satisfied in each iteration, the k subsets $U_1, U_2, \cdots, U_k$ are strictly decreasing. Hence, $k \le |U_1| \le n$. Now, consider an arbitrary node $v_i \in U_k$. For each $1 \le j \le k$, let $V_{i,j} = V_i \cap U_j$. Then, $I_j \cap V_{i,j} \neq \emptyset$ for each $1 \le j \le k$ by the first-fit criteria for selection the MIS. Hence,

$$d(V_i) = \sum_{j=1}^{k} \lambda_j |I_j \cap V_{i,j}| \ge \sum_{j=1}^{k} \lambda_j .$$

Thus, the theorem follows. □

Theorem 5.1 has a profound application in approximating the independence polytope, denoted by P, of G. Consider an arbitrary ordering $\langle v_1, v_2, \cdots, v_n \rangle$ of V, and define V_i as in Theorem 5.1. The *inductive independence polytope* of G by the ordering $\langle v_1, v_2, \cdots, v_n \rangle$ is defined to be

$$Q = \left\{ d \in \mathbb{R}_+^V : \max_{1 \le i \le n} d(V_i) \le 1 \right\},$$

and the *inductive independence number* of G by the ordering $\langle v_1, v_2, \cdots, v_n \rangle$ is defined to be the maximum size, denoted by α^*, of any IS of G contained in some V_i for $1 \le i \le n$.

Corollary 5.2. $Q \subseteq P \subseteq \alpha^* Q$*. In addition, the weight of the coloring output by* ***F^{3}WC*** *is at most* $\alpha^* \chi_f(G, d)$.

Proof. By Theorem 5.1, for any $d \in Q$, $\chi_f(G, d) \le 1$ and hence $d \in P$. So, $Q \subseteq P$. To prove that $P \subseteq \alpha^* Q$, it's sufficient to show that for any IS I of G, its incidence vector belongs to $\alpha^* Q$. Let d be the incidence vector of I. Then, for each $1 \le i \le n$, $d(V_i) = |I \cap V_i| \le \alpha^*$. Hence, $d \in \alpha^* Q$. Thus, $P \subseteq \alpha^* Q$. Finally, consider an arbitrary $d \in \mathbb{R}_+^V$. Then,

$$d \in \chi_f(G, d) P \subseteq \alpha^* \chi_f(G, d) Q.$$

By Theorem 5.1, the coloring of (G, d) output by **F^{3}WC** has weight at most $\alpha^* \chi_f(G, d)$. □

In the next, we present another approximate independence polytope based on graph orientations. An *orientation* of an undirected graph G is a digraph obtained from G by imposing an orientation on each edge of G. Suppose that the digraph $D = (V, A)$ is an orientation of $G = (V, E)$. For each $u \in V$, let $N^{in}(u)$ (resp., $N^{out}(u)$) denote the set of in-neighbors (resp., out-neighbors) of u in D, and let $N^{in}[u]$ (resp., $N^{out}[u]$) denote the union of $N^{in}(u)$ (resp., $N^{out}(u)$) and $\{u\}$. For any $d \in \mathbb{R}_+^V$, a node $u \in V$ is said to be a *surplus* node of (D, d) if $d(N^{in}(u)) \geq d(N^{out}(u))$. It's easy to show that for any $d \in \mathbb{R}_+^V$, (D, d) contains at least one surplus node. For any $d \in \mathbb{R}_+^V$, we construct an ordering of V, depending on d, iteratively as follows. Initialize H to D. For $i = n$ down to 1, let v_i be a vertex of the largest surplus in (H, d) and delete v_i from H. Then the ordering $\langle v_1, v_2, \cdots, v_n \rangle$ is called a *largest surplus last ordering* of (D, d).

LEMMA 5.3. *The coloring of (G, d) output by $\boldsymbol{F^3 WC}$ in the largest surplus last ordering of (D, d) uses at most n colors with total weight at most $2 \max_{u \in V} d(N^{in}[u])$.*

PROOF. Suppose that $\langle v_1, v_2, \cdots, v_n \rangle$ is a the largest surplus last ordering of (D, d). For each $1 \leq i \leq n$, let V_i^{in} (resp., V_i^{out}) denote the set consisting of the in-neighbors (resp., out-neighbors) of v_i in $\{v_1, v_2, \cdots, v_{i-1}\}$ in D, and let

$$V_i = V_i^{in} \cup V_i^{out} \cup \{v_i\}.$$

By the construction of the largest surplus last ordering, we have $d(V_i^{in}) \geq d(V_i^{out})$ and hence

$$\begin{aligned} d(V_i) &= d(v_i) + d\left(V_i^{in}\right) + d\left(V_i^{out}\right) \leq d(v_i) + 2d\left(V_i^{in}\right) \\ &\leq 2d\left(N^{in}[v_i]\right) \leq 2 \max_{u \in V} d\left(N^{in}[u]\right). \end{aligned}$$

By Theorem 5.1, the lemma follows. □

The *independence polytope* of D is defined to be

$$Q' = \left\{ d \in \mathbb{R}_+^V : \max_{u \in V} d\left(N^{in}[u]\right) \leq 1/2 \right\}.$$

It is also called as a *directional independence polytope* of G. The *local independence number* of D is defined to be the maximum size of any IS of G contained in some $N^{in}[u]$ for some $u \in V$, and is denoted by β^*.

COROLLARY 5.4. *$Q' \subseteq P \subseteq 2\beta^* Q'$. In addition, the coloring of (G, d) output by $\boldsymbol{F^3 WC}$ in the largest surplus last ordering of (D, d) has weight at most $2\beta^* \chi_f(G, d)$.*

The proof of the Corollary 5.4 is almost the same as that of Corollary 5.2, and is thus omitted.

6. RESTRICTED MULTIFLOW

Consider a wireless network $N = (D, G, b)$ and suppose that $D = (V, A)$. For each node $v \in V$, we use $\delta^{in}(v)$ (resp., $\delta^{out}(v)$) to denote the set of links in D entering (resp., leaving) v. Consider two distinct nodes $s, t \in V$. A vector $f \in \mathbb{R}_+^A$ is called a flow from s to t, or simply a $s - t$ flow, if for each for each $v \in V \setminus \{s, t\}$,

$$f\left(\delta^{out}(v)\right) = f\left(\delta^{in}(v)\right)$$

This condition is called the *flow conservation law*: the amount of flow entering a vertex $v \neq s, t$ should be equal to the amount of flow leaving v. The value of a flow f from s to t is, by definition:

$$val(f) = f\left(\delta^{out}(s)\right) - f\left(\delta^{in}(s)\right).$$

So the value is the net amount of flow leaving s, which is also equal to the net amount of flow entering t.

Suppose that we are given with k commodities with s_i, t_i being the source and sink, respectively, for commodity i. We use $\mathcal{F}_i$ to denote the set of s_i–t_i flows. A k-flow is s sequence of flows $\langle f_1, f_2, \cdots, f_k \rangle$ with $f_i \in \mathcal{F}_i$ for each $1 \leq i \leq k$. Let P be the independence polytope of G (also known as the capacity region), and Q be a polytope contained in P. A k-flow $\langle f_1, f_2, \cdots, f_k \rangle$ is said to be *Q-restricted* if $\left(\sum_{j=1}^k f_j\right)/b \in Q$. The maximum Q-restricted multiflow is defined by the following LP:

$$\begin{array}{lll} (MMF_Q) & \max & \sum_{j=1}^k val(f_j) \\ & s.t. & f_j \in \mathcal{F}_j, \forall 1 \leq j \leq k \\ & & \left(\sum_{j=1}^k f_j\right)/b \in Q \end{array}$$

Similarly, the maximum concurrent Q-restricted multiflow with demands $d(j)$ for $1 \leq j \leq k$ is defined by the following LP:

$$\begin{array}{lll} (MCMF_Q) & \max & \phi \\ & s.t. & f_j \in \mathcal{F}_j, \forall 1 \leq j \leq k \\ & & val(f_j) \geq \phi d(j), \forall 1 \leq j \leq k \\ & & \left(\sum_{j=1}^k f_j\right)/b \in Q \end{array}$$

Of course, the two LPs may not be optimal as Q is a subset of P. However, if Q is a μ-approximation of P for some $\mu \geq 1$, i.e., $Q \subseteq P \subseteq \mu Q$, then both of them are μ-approximations of the respective optimum. Furthermore, if Q has an explicit polynomial representation, both LPs are of polynomial size and can be solved in polynomial time. But we still cannot claim a polynomial-time μ-approximate solution yet, even if Q is a μ-approximation of P and has an explicit polynomial representation. We have to make sure that a fractional schedule of length at most one for the k-flows output by the two LPs can be found in polynomial time. If in addition there is a polynomial algorithm which produces a fractional schedule of length at most one for (G, d) for any $d \in Q$, then we can safely make the claim. Indeed, the μ-approximation algorithm runs in two steps: We first solve the Q-restricted LP to obtain a k-flow $\langle f_1, f_2, \cdots, f_k \rangle$, and then compute a fractional link schedule of length at most one for $\left(G, \left(\sum_{j=1}^k f_j\right)/b\right)$.

Now, the picture is clear. Both design and analysis of a polynomial-time μ-approximation algorithm for either **MMF** or **MCMF** boil down to find a polytope Q satisfying the three conditions: (1) Q is a μ-approximation of P, (2) Q has an explicit polynomial representation, and (3) there is a polynomial algorithm which produces a fractional schedule of length at most one for (G, d) for any $d \in Q$. The inductive independence polytopes and directional independence polytopes in Section 5 are perfect candidates for such Q. Based on the above discussions and the theory developed in Section 5, we have reached the following two master theorems.

THEOREM 6.1. *Suppose that $\mathcal{N}$ is network class satisfying that there is a polynomial algorithm to find for any network in $\mathcal{N}$ an ordering of the communication links by which the inductive independence number of the conflict graph is at most μ. Then, when restricted to $\mathcal{N}$, both **MMF** and **MCMF** have a polynomial μ-approximation algorithm.*

THEOREM 6.2. *Suppose that $\mathcal{N}$ is a network class satisfying that there is a polynomial algorithm to find for any network in $\mathcal{N}$ an orientation of conflict graph whose local independence number in the conflict graph is at most μ. Then,* *when restricted to $\mathcal{N}$, both* ***MMF*** *and* ***MCMF*** *have a polynomial 2μ-approximation algorithm.*

In the remaining of this section, we focus on the 802.11 class and the PIM class. For any network in the 802.11 class, we consider the following two orderings of the communication links:

- Interference radius decreasing ordering: Define the inference radius of a link to be larger one of the interference radii of its endpoints, and sort all links in descending order of the interference radius.
- Lexicographic ordering: Sort all links in the lexicographic order of their right endpoints.

These two orderings have the following two properties.

LEMMA 6.3. *For any network in the 802.11 class, the inductive independence number of the conflict graph by the interference radius decreasing ordering is at most* 23.

LEMMA 6.4. *For any network in the subclass of the 802.11 class in which all nodes have uniform communication radii and uniform interference radii, the inductive independence number of the conflict graph by the lexicographic ordering is at most* 7.

The proofs of the above two lemmas are purely geometric and yet quite involved, and so are relegated to Appendix 2 and Appendix 3 respectively. So, for the 802.1 class, we have the following results.

THEOREM 6.5. *Restricted to the 802.11 class, both* ***MMF*** *and* ***MCMF*** *have a polynomial* 23*-approximation algorithm.*

THEOREM 6.6. *Restricted to the subclass of the 802.11 class in which all nodes have uniform communication radii and uniform interference radii, both* ***MMF*** *and* ***MCMF*** *have a polynomial* 7*-approximation algorithm.*

For any network in the PIM class, we consider the following orientation of its conflict graph. For any pair of conflicting links $a_1 = u_1v_1$ and $a_2 = u_2v_2$, if v_1 is within the interference range of u_2 and v_2 is within the interference range of u_1, take an arbitrary orientation; otherwise, if v_1 is within the interference range of u_2, take the orientation from a_2 to a_1; otherwise, take the orientation from a_1 to a_2. Then, the orientation of the conflict graph has the following important property: for any in-neighbor $a' = u'v'$ of $a = uv$, v is within the interference range of u'. This property leads to the following theorem.

LEMMA 6.7. *For any network in the subclass of the PIM class in which the interference radius of each node is at least c times its communication radius for some $c > 1$, the local independence number of the above orientation of its conflict graph is at most $\lceil \pi / \arcsin \frac{c-1}{2c} \rceil - 1$.*

The proof of this lemma is given in Appendix 4. We remark that the bound in the above lemma is at most half of the bound $\lceil 2\pi / \arcsin \frac{c-1}{2c} \rceil$ derived in [21]. So, for the PIM class we have the following result.

THEOREM 6.8. *Restricted to the subclass of the PIM class in which the interference radius of each node is at least c times its communication radius for some $c > 1$, both* ***MMF*** *and* ***MCMF*** *have a polynomial $2\left(\lceil \pi / \arcsin \frac{c-1}{2c} \rceil - 1\right)$-approximation algorithm.*

7. CONCLUSION

In this paper, we have conducted comprehensive studies on both **MMF** and **MCMF** restricted to either the 802.11 class or the PIM class. Not only have we provided the full characterization of their NP-hardness and approximation hardness, we have also developed polynomial algorithms with better approximation bounds. In addition, the two relevant problems **MWCFL** and **MFWLS** restricted to either the 802.11 class or the PIM class have also been studied. While the 802.11 class and the PIM class are the two classes of multihop wireless networks focused on in this paper, most of our results are given for an arbitrary network class. In particular, we have developed a unified framework for both the design and the analysis of polynomial approximation algorithms.

8. REFERENCES

[1] M. Alicherry, R. Bhatia, and L.E. Li, Joint Channel Assignment and Routing for Throughput Optimization in Multiradio Wireless Mesh Networks, *IEEE JSAC* 24(11):1960–1971 (2006). Also appeared in *Proc. of ACM MobiCom* 2005.

[2] H. Balakrishnan, C.L. Barrett, V.S.A. Kumar, M.V. Marathe, and S. Thite, The distance-2 matching problem and its relationship to the MAC-Layer capacity of ad hoc wireless networks, *IEEE JSAC* 22(6): 1069-1079 (2004).

[3] C. Buragohain, S. Suri, C. D. Toth, and Y. Zhou, Improved Throughput Bounds for Interference-Aware Routing in Wireless Networks, *Proc. COCOON* 2007, LNCS 4598, Springer, pp. 210-221.

[4] B. Carr and S. Vempala, Randomized meta-rounding, *Proc. of ACM STOC* 2000.

[5] T. Erlebach, K. Jansen, E. Seidel, Polynomial-time approximation schemes for geometric graphs, *SIAM Journal on Computing* 34 (6) (2005) 1302–1323. Also appeared in *Proc. of ACM-SIAM SODA* 2001, pp. 671–679.

[6] M. Grötschel, L. Lovász, and A. Schrijver, The ellipsoid method and its consequences in combinatorial optimization, *Combinatorica* 1 (1981) 169–197.

[7] D.S. Hochbaum, and W. Maass, Approximation schemes for covering and packing problems in image processing and VLSI, *Journal of the ACM* 32 (1) (1985) 130–136.

[8] H.B. Hunt III, M.V. Marathe, V. Radhakrishnan, S.S. Ravi, D.J. Rosenkrantz, and R.E. Stearns, NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs, *Journal of Algorithms* 26 (2) (1998) 238–274.

[9] K. Jain, M. Mahdian, and M.R. Salavatipour, Packing Steiner trees, *Proceedings of ACM-SIAM SODA* 2003, pp.266–274.

[10] K. Jain, J. Padhye, V.N. Padmanabhan, and L. Qiu, Impact of interference on multi-hop wireless network performance, *ACM/Springer Wireless Networks* 11:471–487, 2005. Also appeared in *Proc. of ACM MobiCom* 2003.

[11] K. Jansen, Approximate strong separation with application in fractional graph coloring and preemptive scheduling, *Theoretical Computer Science* 302 (1–3) (2003) 239–256.

[12] N. Karmarkar, R.M. Karp, An efficient approximation scheme for the one-dimensional bin-packing problem, *Proc. of IEEE FOCS* 1982, pp. 206–213.

[13] M. Kodialam and T. Nandagopal, Characterizing achievable rates in multi-hop wireless networks: the joint routing and scheduling problem, *Proc. of ACM MobiCom* 2003.

[14] M. Kodialam and T. Nandagopal. Characterizing the capacity region in multi-radio multi-channel wireless mesh networks, *Proc. of ACM MobiCom* 2005.

[15] V.S.A. Kumar, M.V. Marathe, S. Parthasarathy, and A. Srinivasan, Algorithmic aspects of capacity in wireless networks, *SIGMETRICS Perform. Eval. Rev.* 33(1):133–144, 2005.

[16] T. Matsui, Approximation algorithms for maximum independent set problems and fractional coloring problems on unit disk graphs, *JCDCG* 1998, LNCS 1763, Springer, pp. 194–200.

[17] T. Nieberg, J. Hurink, and W. Kern, A Robust PTAS for Maximum Weight Independent Sets in Unit Disk Graphs, *Proc. WG 2004*, LNCS 3353, Springer, pp. 214–221.

[18] Z. Nutov, I. Beniaminy, and R. Yuster, A (1-1/e)-approximation algorithm for the generalized assignment problem. *Oper. Res. Lett.* 34(3): 283-288 (2006).

[19] E.J. van Leeuwen, Better Approximation Schemes for Disk Graphs, *Proc. SWAT* 2006, LNCS 4059, Springer, pp. 316–327.

[20] D. Wang and Y.-S. Kuo, A study on two geometric location problems, *Information Processing Letters*, 28:281–286, 1988.
[21] Y. Wang, W. Wang, X.-Y. Li, and W.-Z. Song, Interference-Aware Joint Routing and TDMA Link Scheduling for Static Wireless Networks, *IEEE TPDS* 19(12): 1709-1726, Dec. 2008. An early version appeared in *Proc. of ACM MobiCom* 2006.

Appendix 1: Debugging

Proof of NP-hardness in [10]: Theorem 1 in [10] claimed the NP-hardness of **MMF** restricted to the PIM class using a reduction from **MIS** in general graphs. Let $G = (V, E)$ be an arbitrary graph. Jain et al. [10] constructed a wireless network $N = (D, G, b)$ as follows. D consists of only two nodes s and t. For each $v \in V$, they added to D a link e_v from s to t and set $b(e_v) = 1$. In the conflict graph G', two links e_u and e_v conflict with each other if and only if u and v are adjacent in G. Then, the maximum s-t flow in N is $\alpha(G)$. They argued that such network N "may arise, for instance, if nodes s and r are each equipped with multiple radios set either to the same channel or to separate channels", but later on admitted that "the above proof may come across as contrived since the wireless network we constructed is unlikely to arise in practice". Such proof is indeed contrived as all wireless networks in their paper were assumed implicitly to be single-channel and single-radio. Multi-channel and multi-radio appeared nowhere else in their paper except this proof. Even with multi-channel and multi-radio wireless networks, their reduction is incomplete as they do not provide any mapping from the links to the channels and radios. A more important truth is that their reduction does not work for wireless networks with constant number of channels and constant number of radios per node.

Link scheduling in [15] [1] [21]: Let G be the interference graph, A be the set of its communication links, and $d \in \mathbb{R}_+^A$ be the cumulative flow. All the three papers [15] [1] [21] computed a schedule of (G, d) as follows. They first chose an integer w such that for all $e \in A$, $w \cdot d(e)$ is integral. Let us assume that d is rational such that such integer w exists. They set $d' = w \cdot d$ and run the conventional first-fit coloring on (G, d') in some ordering of A: each $e \in A$ is assigned with the first $d'(e)$ available colors which have not been used by any link preceding and conflicting with e. Finally, such first-fit schedule of (G, d) is scaled down to obtain a schedule of (G, d). This algorithm is not polynomial as it has super-exponential worst-case running time. Indeed, consider a rational $d \in \mathbb{R}_+^A$ satisfying that for each $e \in A$, each $d(e)$ is a positive reduced fraction whose denominator is a distinct (sufficiently large) prime number $p(e)$. The least common denominator of all $d(e)$ for $e \in A$ is $\prod_{e \in A} p(e)$, which is a lower bound on w. So, for each $e \in A$, $d'(e)$ is super-exponential. Thus, when run on (G, d'), the conventional first-fit coloring in any ordering would take super-exponential running time.

Flow schedulability in [21]: Consider a network in the PIM subclass. Let A be the set of its communication links. For each link $e = uv$, define $I^{in}(e)$ to be set of links $e' = u'v'$ in A satisfying that v is within the interference range of u'. For each cumulative flow (or demand) $d \in \mathbb{R}_+^A$, Wang et al. [21] introduced the following constraint to guarantee the schedulability of d (i.e., d belongs to the capacity region):

$$\max_{e \in A} \left(d(e) + \sum_{e' \in I^{in}(e)} d(e') \right) \leq 1.$$

But such constraint cannot guarantee the schedulability of d, even restricted to the subclass of the PIM class in which all nodes have unit communication radius and a uniform interference radius equal to some constant $\rho \geq 1$. We construct a counter-example below. Consider a set U of 10 nodes u_0 through u_9 on a circle C (see Figure 1) in the counterclockwise order satisfying that for each $0 \leq i \leq 4$, $\|u_{2i}u_{2i+1}\| = 1$ and $\|u_{2i-1}u_{2i}\| = \rho$ (u_{-1} is treated as u_9 in the cyclic order). The rest nodes $V \setminus U$ are located on the fives minor arcs between u_{2i-1} and u_{2i} for $0 \leq i \leq 4$ respectively such that the unit-disk graph on V is connected. Let e_i be the link $u_{2i}u_{2i+1}$ for $0 \leq i \leq 4$. Consider the demand $\in \mathbb{R}_+^A$ given by $d(e_i) = 1/2$ for each $0 \leq i \leq 4$ and $d(e) = 0$ for all other links $e \in A$. It's easy to verify that

$$\max_{e \in A} \left(d(e) + \sum_{e' \in I^{in}(e)} d(e') \right) = 1.$$

We claim that the shortest schedule for d has length 1.25. Indeed, let G' be the conflict graph of the five links e_i for $0 \leq i \leq 4$. Then, the shortest schedule for d is equivalent to the minimal fractional coloring of (G', d') in which $d'(e_i) = 1/2$ for each $0 \leq i \leq 4$. Since G' is a 5-cycle, its independence number is two, and hence $\chi_f(G', d') \geq 2.5/2 = 1.25$. On the other hand, the fractional coloring which assigns a weight of 1/4 to each of the 5 maximum independent sets has weight $5/4 = 1.25$. Hence $\chi_f(G', d') = 1.25$.

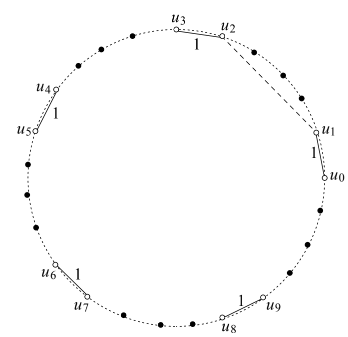

Figure 1: A counter-example to the flow schedulability in [21] and [3].

Flow schedulability in [3]: Consider a network in the subclass of the 802.11 class in which all nodes have unit communication radius and a uniform interference radius equal to some constant $\rho \geq 1$. Let V be the networking nodes and A be the set of its communication links. Sort V in the lexicographic order. For each $v \in V$, let Γ_v denote the set consisting of v itself and all nodes preceding v in the lexicographic order and lying in the interference range of v, and A_v denote the set of links in A with both endpoints in Γ_v. For each cumulative flow $d \in \mathbb{R}_+^A$, τ_v denotes the total flow through the links in A incident to v for any $v \in V$, and Buragohain et al. [3] introduced the following constraint to guarantee the schedulability of d:

$$\max_{v \in V} \left(\sum_{u \in \Gamma_v} \tau_u - \sum_{e \in A_v} d(e) \right) \leq 1.$$

Again this constraint fails to guarantee the schedulability of d. Indeed, consider the same instance of nodes given in the previous paragraph. It's also easy to verify that

$$\max_{v \in V} \left(\sum_{u \in \Gamma_v} \tau_u - \sum_{e \in A_v} d(e) \right) = 1.$$

But under the 802.11 interference model, the shortest schedule for d is also 1.25 using the same argument as above.

Appendix 2: Proof of Lemma 6.3

For each node v, we write ρ_v for $\rho(v)$, B_v for the disk of radius ρ_v centered at v, and $\Gamma(v)$ for the set of nodes w satisfying that $\rho_w \geq \rho_v$ and $v \in B_w$. Consider an arbitrary link e and let I be a set of conflict-free links in which each link conflict with e and has interference radii at least that of e. We will prove that $|I| \leq 23$, from which Lemma 6.3 follows. If $e \in I$, then $|I| = 1$. So we assume that $e \notin I$ in the sequel. As any pair of links in I have no common endpoint and the directions of the links have no impact on the interference, we will ignore the directions of the links in I and call the links in I as edges for simplicity. Let u and v be the two endpoints of e with $\rho_u \geq \rho_v$. We partition I into four subsets I_1, I_2, I_3 and I_4 as follows. I_1 consists of edges in I with at least one endpoint in $\Gamma(u)$, I_2 consists of edges in $I \setminus I_1$ with at least one endpoint in B_u, I_3 consists of edges in $I \setminus (I_1 \cup I_2)$ with at least one endpoint in $\Gamma(v)$, and I_4 consists of rest edges in I. We define the representatives of the edges in I. The representative of each edge in I_1 (resp., I_3) is one of its endpoints belonging to $\Gamma(u)$ (resp., $\Gamma(v)$), the representative of each edge in $I_2 \cup I_4$ is its endpoint with larger interference radius.

LEMMA 8.1. *The following statements are true:*

1. *Suppose that w_1 and w_2 are representatives of two edges in $I_1 \cup I_2$. Then, $\widehat{w_1 u w_2} > 2\arcsin\frac{1}{4}$.*
2. *Suppose that w_1 and w_2 are representatives of two edges in $I_3 \cup I_4$. Then, $\widehat{w_1 v w_2} > 2\arcsin\frac{1}{4}$.*
3. *Suppose that w is the representative of an edge in $I_3 \cup I_4$. Then $\widehat{uvw} > 30^o$.*

The proof of the above lemma is lengthy, and is omitted due to the limitation on space. By Lemma 8.1(1),

$$|I_1 \cup I_2| \leq \left\lceil \frac{2\pi}{2\arcsin\frac{1}{4}} \right\rceil - 1 = 12.$$

By Lemma 8.1(2) and (3),

$$|I_3 \cup I_4| \leq \left\lceil \frac{2\pi - \frac{\pi}{3}}{2\arcsin\frac{1}{4}} \right\rceil = 11.$$

Therefore, $|I| = |I_1 \cup I_2| + |I_3 \cup I_4| \leq 23$.

Appendix 3: Proof of Lemma 6.4

By proper scaling, we assume the communication of each node is one and the interference of each node is $\rho \geq 1$. For each node v, we use B_v to denote the disk of radius ρ centered at v. Consider an arbitrary link e and let I be a set of conflict-free links in which each link conflict with e and its endpoints both lie to the left side of the right endpoint of e. We will prove that $|I| \leq 7$, from which Lemma 6.4 follows. If $e \in I$,then $|I| = 1$. So we assume that $e \notin I$ in the sequel. As any pair of links in I have no common endpoint and the directions of the links have no impact on the interference, we will ignore the directions of the links in I and call the links in I as edges for simplicity. Let u and v be the two endpoints of e with u be the right endpoint of e. For each edge in I, we pick one of its endpoints belonging to $B_u \cup B_v$ as its representative. We partition I into two subsets I_1 and I_2 as follows. Let I_1 be the set of edges in I whose representative lie in B_u. As the representatives of the edges in I_1 all lie in a half disk of B_u and have pairwise distances greater than the radius ρ of B_u, we have $|I_1| \leq 3$. As the representatives of the edges in I_2 all lie in the "moon" $B_v \setminus B_u$ and have pairwise distances greater than ρ, we have $|I \setminus I_1| \leq 4$. Therefore, $|I| \leq |I_1| + |I \setminus I_1| \leq 7$.

Appendix 4: Proof of Lemma 6.7

The following geometric lemma is key to the proof for Lemma 6.7.

LEMMA 8.2. *Consider a triangle $\triangle puv$ with $\|pu\| = \rho$ and $\|pv\| = 1$ (see Figure 8.2). Let q be a point on the same side of pu satisfying that $\|uq\| = \|uv\|$ and $\|pq\| = \rho$. Then,*

$$\widehat{quv} \geq 2\arcsin\frac{\rho - 1}{2\rho}.$$

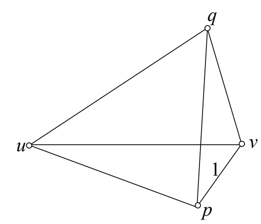

Figure 2: Figure for Lemma 8.2.

The proof of the above lemma is also lengthy and is omitted due to the limitation on space. Now we proceed to prove Lemma 6.7. For each node v, we write ρ_v for $\rho(v)$, Let $e_1 = u_1 v_1$ be an arbitrary link, and I be a set of conflict-free links in which each link is either e_1 itself or an in-neighbor of e_1 in the orientation of the conflict graph. We will prove that $|I| \leq \lceil \pi / \arcsin\frac{c-1}{2c} \rceil - 1$, from which Lemma 6.7 follows. If $e \in I$ then $|I| = 1$. So we assume that $e \notin I$ in the sequel. Consider any pair of links $e_2 = u_2 v_2$ and $e_3 = u_3 v_3$ in I. Then, $\|u_2 v_3\| > \rho_{u_2}$. We claim that

$$\widehat{v_2 v_1 v_3} > 2\arcsin\frac{c-1}{2c}.$$

By symmetry, we assume that v_2 is farther away from v_1 than v_3. Then, u_2 lies in the intersection of $B(v_1, \rho_{u_2})$ and $B(v_2, r_{u_2})$. Let p be the intersection point of $\partial B(v_1, \rho_{v_1})$ and $\partial B(v_2, r_{u_2})$ which lies on the different side of $v_1 v_2$ from v_3, and q be the point on the same side of $v_1 p$ satisfying that $\|v_1 q\| = \|v_1 v_2\|$ and $\|pq\| = \rho_{u_2}$. By Lemma 8.2,

$$\widehat{v_2 v_1 q} \geq 2\arcsin\frac{\rho_{u_2} - r_{u_2}}{2\rho_{u_2}} \geq 2\arcsin\frac{c-1}{2c}.$$

For any point w in the sector $v_2 v_1 q$ (centered at v_1 with radius $\|v_1 v_2\|$),

$$\|u_2 w\| \leq \|pw\| \leq \max\{\|pv_1\|, \|pv_2\|, \|pq\|\} = \rho_{u_2}.$$

Thus, v_3 cannot be in the sector $v_2 v_1 q$ and consequently,

$$\widehat{v_2 v_1 v_3} > \widehat{v_2 v_1 q} \geq 2\arcsin\frac{c-1}{2c}.$$

Therefore, the claim holds. Hence,

$$|I| \leq \left\lceil \frac{\pi}{\arcsin\frac{c-1}{2c}} \right\rceil - 1.$$

So, the lemma follows.