

Universal Framework for Wireless Scheduling Problems^{*†}

Eyjólfur I. Ásgeirsson¹, Magnús M. Halldórsson², and Tigran Tonoyan³

- 1 School of Science and Engineering, Reykjavik University, Reykjavik, Iceland
eyjo@ru.is
- 2 ICE-TCS, School of Computer Science, Reykjavik University, Reykjavik, Iceland
mmh@ru.is
- 3 ICE-TCS, School of Computer Science, Reykjavik University, Reykjavik, Iceland
tigran@ru.is

Abstract

An overarching issue in resource management of wireless networks is assessing their *capacity*: *How much communication can be achieved in a network, utilizing all the tools available: power control, scheduling, routing, channel assignment and rate adjustment?* We propose the first framework for approximation algorithms in the physical model that addresses these questions in full, including rate control. The approximations obtained are doubly logarithmic in the link length and rate diversity. Where previous bounds are known, this gives an exponential improvement.

A key contribution is showing that the complex interference relationship of the physical model can be simplified into a novel type of amenable conflict graphs, at a small cost. We also show that the approximation obtained is provably the best possible for any conflict graph formulation.

1998 ACM Subject Classification C.2.1 Network Architecture and Design

Keywords and phrases Wireless, Scheduling, Physical Model, Approximation framework

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.129

1 Introduction

The effective use of wireless networks revolves around utilizing fully all available diversity. This can include power control, scheduling, routing, channel assignment and *transmission rate control* on the links, the latter being an issue of key interest for us. The long-studied topic of *network capacity* deals with how much communication can be achieved in a network when its resources are utilized to the fullest. This can be formalized in different ways, involving a range of problems. The communication ability of packet networks is characterized by the capacity region, i.e. the set of traffic rates that can be supported by any scheduling policy. In order to achieve *low delays and optimal throughput*, the classic result of Tassiulas and Ephremides [28] and followup work in the area (e.g. [25]) point out a core optimization problem that lies at the heart of such questions – the *maximum weight independent set of links* (MWISL) problem: from a given set of communication links with associated weights/utilities, find an *independent (conflict-free, subject to the interference model) subset of maximum total weight*.

* A full version of the paper is available at <https://arxiv.org/abs/1705.10104>.

† This work was supported by grants 152679-05 and 174484-05 from the Icelandic Research Fund.



© Eyjólfur I. Ásgeirsson, Magnús M. Halldórsson, and Tigran Tonoyan;
licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 129; pp. 129:1–129:15



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



This reduction applies to very general settings involving single-hop and multi-hop, as well as fixed and controlled transmission rate networks. Moreover, approximating MWISL within any factor implies achieving the corresponding fraction of the capacity region. This makes MWISL a central problem in the area. Unfortunately, solving this problem in its full generality is notoriously hard, since it is well known that MWISL is effectively inapproximable (under standard complexity theory) e.g. in models described by general conflict relations or general graphs. Moreover, in general, even approximating the capacity region in polynomial time within a non-trivial bound, while keeping the delays in reasonable bounds, is hard under standard assumptions [27].

We tackle this question in the *physical model* of communication. Towards this end, we develop a general approximation framework that not only helps us to approximate MWISL, but can also be used for tackling various other scheduling problems, such as TDMA scheduling, joint routing and scheduling and others. The problems handled can additionally involve path or flow selection, multiple channels and radios, and packet scheduling. We obtain *doubly-logarithmic* (in link and rate diversity) approximation for these problems, exponentially improving the previously known logarithmic approximations, and, importantly, extending them to incorporate *different fixed rates and rate control*. The crucial feature of the approach (which makes it so general) is that it involves transforming the complex physical model into an unweighted/undirected conflict graph and solving the problems simply on these graphs. Perhaps surprisingly, we find that our schema attains the best possible performance of *any* conflict graph representation. Numerical simulations show that the conflict graph framework is a good approximation for the physical model on randomly placed network instances as well. Our approach also finesses the task of selecting optimum power settings by using *oblivious* power assignment, one that depends only on the properties of the link itself and not on other links. The performance bounds are however in comparison with the optimum solution that can use arbitrary power settings.

Technically, our approach generalizes our earlier framework [14]. Our extensions required substantial changes throughout the whole body of arguments. That formulation works only for uniform constant rates, and the generalization requires substantial new ideas. One indicator of the challenges overcome is that we could prove that our doubly-logarithmic approximation is best possible in the presence of different rates, while better approximations are known to hold in the case of uniform rates [14].

We make some undemanding assumptions about the settings. We assume that the networks are *interference-constrained*, in that interference rather than the ambient noise is the determining factor of proper reception. This assumption is common and is particularly natural in settings with rate control, since the impact of noise can always be made negligible by avoiding the highest rates, losing only a small factor in performance. We also assume that nodes are (arbitrarily) located in a doubling metric, which generalizes Euclidean space, allowing the modeling of some of non-geometric effects seen in practice.

Our Results. Our results can be summarized as follows:

- We establish a general framework for tackling wireless scheduling and related problems,
- Our approximations hold for nearly all such problems, including variable rates settings,
- We obtain exponential improvement over previously known approximations,
- The approximations are obtained via simple conflict graphs, as opposed to the complicated physical model, and by using *oblivious* power assignments,
- We establish tight bounds indicating the limitations of our method.

Related work. Gupta and Kumar introduced the physical model of interference/communication with log-path fading in their influential paper [10], and it has remained the default in

analytic studies. Moscibroda and Wattenhofer [26] initiated worst-case analysis of scheduling problems in networks of arbitrary topology, which is also the setting of interest in this paper. There has been significant progress in understanding scheduling problems with fixed uniform rates. NP-completeness results have been given for different variants [8, 21, 24]. Early work on approximation algorithms involve (directly or indirectly) partitioning links into length groups, which results in performance guarantees that are at least logarithmic in Δ , the link length diversity: TDMA scheduling and uniform weights MWISL [8, 5, 11], non-preemptive scheduling [7], joint power control, scheduling and routing [4], and joint power control, routing and throughput scheduling in multiple channels [2], to name a few. Constant-factor approximations are known for uniform weight MWISL (in restricted power regimes [12] and (general) power control [22]). Standard approaches translate the constant-factor approximations for the uniform weight MWISL into $O(\log n)$ approximations for TDMA scheduling and general MWISL. Many problems become easier, including MWISL and TDMA scheduling, in the regime of linear power assignments [6, 33, 13, 29]. Recently, a $O(\log^* \Delta)$ -approximation algorithm was given for TDMA scheduling and MWISL [14], by transforming the physical model into a conflict graph. We build on this approach, and extend it into a general framework that covers other problems and incorporates support for rate control.

Very few results are known for problems involving rate control. The constant-factor approximation for MWISL with uniform weights and arbitrary but fixed rates proposed by Kesselheim [23] can be used to obtain $O(\log n)$ -approximations for TDMA scheduling and MWISL with rate control, where n is the number of links. Another recent work [9] handles the TDMA scheduling problem (with fixed but different rates), obtaining an approximation independent of the number of links n , but the ratio is polynomial in Δ . There have been numerous algorithms that try to approximate or replace MWISL in the context of packet scheduling. Several examples include Longest-Queue-First Scheduling (LQF) [20], Maximal Scheduling [34], Carrier Sense Multiple Access (CSMA) [19]. The approximations obtained usually depend on some parameter of the conflict graph, such as the *interference degree*. In the case of CSMA (and other similar protocols), it is known that the algorithms are throughput-optimal, but in general they take exponential time to stabilize, or otherwise require constant degree conflict graphs [18]. It is also well known that many scheduling problems such as vertex coloring and MWISL are easy to approximate in bounded *inductive independence* graphs, such as geometric intersection graphs or protocol model. However, fidelity to the cumulative nature of interference and the question of modeling rate control are among the significant issues faced by such graph models.

Paper Organization. The fundamental ideas of our approximation framework are described in Section 2. After introducing the model and definition in Section 3, we derive the core technical part, the approximation of the physical model by the conflict graphs, in Section 4, and the optimality of approximation. The framework is applied to obtain approximations for fixed rate scheduling problems in Section 5 and for problems with rate control in Section 6 (the latter two can be read separately from Section 4). Due to space constraints, several technical proofs are deferred to the full version of this paper.

2 Approximation Method

Before defining the details, let us describe the main idea behind the approximation technique. In essence, we define a notion of approximation of an independence system¹ $\mathcal{I}_{\mathcal{P}} = (L, \mathcal{E}_{\mathcal{P}})$ by

¹ An independence system \mathcal{I} over a set of vertices V is a pair $\mathcal{I} = (V, \mathcal{E})$, where $\mathcal{E} \subseteq 2^V$ is a collection of subsets of vertices that is closed under subsetting: if $S \in \mathcal{E}$ and $S' \subset S$, then $S' \in \mathcal{E}$.

a graph $\mathcal{G} = (L, E)$ over the set L of links. The system $\mathcal{I}_{\mathcal{P}}$ corresponds to the cumulative interference in the physical model, while \mathcal{G} is a conflict graph describing pairwise conflicts between links. We will refer to independent sets in $\mathcal{I}_{\mathcal{P}}$ as *feasible sets*, and to independent sets in \mathcal{G} as *independent sets*, to avoid confusion.

The approximation is described by several key properties.

Refinement (Feasibility of Independent Sets). Every independent set S in \mathcal{G} must be feasible, i.e. $S \in \mathcal{E}_{\mathcal{P}}$. Thus, finding an independent set in \mathcal{G} gives also a feasible set in $\mathcal{I}_{\mathcal{P}}$.

Tightness (of refinement). There is a small number k such that every feasible set $S \in \mathcal{I}_{\mathcal{P}}$ is a union of at most k independent sets in \mathcal{G} . The smallest such k is called the *tightness* of refinement. This property guarantees that even an optimal (for a problem in question) feasible set can be covered with a few independent sets.

The two properties above establish a tight connection between the two models. That allows us to take nearly *every* scheduling problem in the physical model and reduce it to the corresponding problem in conflict graphs (in a way formalized in Section 5), by paying only an approximation factor depending on the tightness k . However, in order for this scheme to work, it should be easier to solve such problems in \mathcal{G} , which leads to the third key property.

Computability. There are efficient (approximation) algorithms for scheduling-related problems such as *vertex coloring* and *maximum weight independent set* in \mathcal{G} .

A graph \mathcal{G} satisfying the properties above is said to be a *refinement* of $\mathcal{I}_{\mathcal{P}}$. The main effort in the following two sections is to define an appropriate conflict graph refinement for the physical model and prove these key properties. We find such a family that approximates the physical model with nearly constant tightness, i.e. double-logarithmic in length and rate diversity and show that this is best possible for any conflict graph, up to constant factors. This approximation allows us to bring to bear the large body of theory of graph algorithms, greatly simplifying both the exposition and the analysis.

3 Model

In scheduling problems, the basic object of consideration is a set L of n communication links, numbered from 1 to n , where each link $i \in L$ represents a single-hop communication request between two wireless nodes located in a metric space – a sender node s_i and receiver node r_i .

We assume the nodes are located in a metric space with distance function d . We denote $d_{ij} = d(s_i, r_j)$ and $l_i = d(s_i, r_i)$. The latter is called the *length* of link i . Let $d(i, j)$ denote the minimum distance between the nodes of links i and j .

The nodes have adjustable transmission power levels. A *power assignment* for the set L is a function $P : L \rightarrow \mathbb{R}_+$. For each link i , $P(i)$ defines the power level used by the sender node s_i . In the physical model of communication, when using a power assignment P , a transmission of a link i is successful if and only if

$$SIR(S, i) = \frac{P(i)/l_i^\alpha}{\sum_{j \in S \setminus \{i\}} P(j)/d_{ji}^\alpha} \geq \beta_i, \quad (1)$$

where $\beta_i > 1$ denotes the minimum signal to noise ratio required for link i , $\alpha \in (2, 6)$ is the path loss exponent and S is the set of links transmitting concurrently with link i . Note that we omit the noise term in the formula above, since we focus on interference limited networks. This can be justified by the fact that one can simply slightly decrease the transmission rates

to make the effect of the noise negligible, then restore the rates by paying only constant factors in approximation.

A set S of links is called P -feasible if the condition (1) holds for each link $i \in S$ when using power assignment P . We say S is feasible if there exists a power assignment P for which S is P -feasible.

Effective Length. Let us denote $l_i = \beta_i^{1/\alpha} l_i$ and call it the *effective length* of link i . Let $\Delta(L) = \max_{i,j \in L} \{l_i/l_j\}$ denote the (*effective*) *length diversity* of instance L . We call a set S of links *equilength* if for every two links $i, j \in S$, $l_i \leq 2l_j$, i.e., $\Delta(S) \leq 2$. Note that with the introduction of effective length, the feasibility constraint (1) becomes: $\frac{P(i)}{l_i^\alpha} \geq \sum_{j \in S \setminus \{i\}} \frac{P(j)}{d_{ji}^\alpha}$. This looks like the same formula but with uniform rates $\beta'_i = 1$. However, there is an essential difference between the two: the quantities l_i are not related to the metric space in the same way as lengths l_i , as l_i can be arbitrarily larger than l_i .

Metrics. The *doubling dimension* of a metric space is the infimum of all numbers $\delta > 0$ such that for every ϵ , $0 < \epsilon \leq 1$, every ball of radius $r > 0$ has at most $C\epsilon^{-\delta}$ points of mutual distance at least ϵr where $C \geq 1$ is an absolute constant. For example, the m -dimensional Euclidean space has doubling dimension m [16]. We let m denote the doubling dimension of the space containing the links. We will assume $\alpha > m$, which is the standard assumption $\alpha > 2$ in the Euclidean plane.

4 Conflict Graph Approximation of Physical Model

In this section we present the $O(\log \log \Delta)$ -tight refinement of the physical model by conflict graphs. The first part introduces our conflict graph \mathcal{G}_f that generalizes the conflict graph definition of [14] and extends it to general thresholds/rates. The three subsequent parts give the proofs of the three key properties: refinement, tightness and computability. The last part argues the asymptotic optimality of $O(\log \log \Delta)$ -tightness for any conflict graph, which contrasts the $O(\log^* \Delta)$ bound known in the uniform thresholds setting.

► **Theorem 1.** *There is an $O(\log \log \Delta)$ -tight refinement of the physical model by a conflict graph family $\mathcal{G}(L)$.*

Conflict Graphs. We define the conflict graph family as follows.

► **Definition 2.** Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a positive non-decreasing function. Links i, j are said to be f -independent if $d_{ij}d_{ji} > l_i l_j f(l_{max}/l_{min})$, where $l_{min} = \min\{l_i, l_j\}$, $l_{max} = \max\{l_i, l_j\}$, and otherwise f -adjacent. A set of links is f -independent (f -adjacent) if they are pairwise f -independent (f -adjacent).

The conflict graph $\mathcal{G}_f(L)$ of a set L of links is the graph with vertex set L , where two vertices are adjacent if and only if they are f -adjacent.

This definition extends the conflict graphs introduced in [14], where the independence criterion was $d(i, j)/l_{min} > f(l_{max}/l_{min})$ (l_{max}, l_{min} are the length of the longer and shorter links, resp.). When all threshold values β_i are constant, the latter essentially follows from the definition above by “canceling” l_{max} with the larger value of d_{ij}, d_{ji} (modulo constant factors). In general, however, the effective lengths can be very different from the actual link lengths, and feasibility requires more separation than given by graphs involving distances only. A technical difficulty introduced by the new definition is that we have to keep track of two distances d_{ij} and d_{ji} instead of the single distance $d(i, j)$, but this appears to be necessary.

We will be particularly interested in *sub-linear* functions $f(x) = O(x)$. A function f is *strongly sub-linear* if for each constant $c \geq 1$, there is a constant c' such that $cf(x)/x \leq f(y)/y$ for all $x, y \geq 1$ with $x \geq c'y$. Note that if f is strongly sub-linear then $f(x) = o(x)$. For example, the functions $f(x) = x^\delta$ ($\delta < 1$) and $f(x) = \log x$ are strongly sub-linear.

Refinement: Feasibility of Independent Sets. Our goal now is to find a function f such that each independent set in \mathcal{G}_f is feasible. It is clear that this can be achieved by letting f grow sufficiently fast. But we should not let it grow too fast, so as to not affect tightness. We also need to indicate which power assignment makes the independent sets in \mathcal{G}_f feasible. Our approach is to preselect a family of *oblivious power assignments*, that are local to each link and do not depend on others, and then find an appropriate function f . Consider the family of power assignments P_τ parameterized by $\tau \in (0, 1)$, where $P_\tau(i) \sim l_i^\tau$ for each link i . In order to obtain P_τ -feasibility, we take conflict graphs \mathcal{G}_f with $f(x) = \gamma x^\delta$ for $\delta \in (0, 1)$ and $\gamma \geq 1$. Such graphs are denoted as $\mathcal{G}_\gamma^\delta$. We show that every independent set in $\mathcal{G}_\gamma^\delta$ for appropriate γ and δ is P_τ -feasible for some τ .

► **Theorem 3.** *Let $\delta_0 = \frac{\alpha-m+1}{2(\alpha-m)+1}$. If $\delta \in (\delta_0, 1)$ and the constant $\gamma > 1$ is large enough, there is a value $\tau \in (0, 1)$ such that each independent set in $\mathcal{G}_\gamma^\delta$ is P_τ -feasible.*

The proof is an adaptation of the ideas used in the proof of [15, Cor. 6] to our definition of conflict graphs and effective lengths. Given an independent set S in $\mathcal{G}_\gamma^\delta$ and a link i , we bound the interference of S on i by first splitting S into equilength subsets, bounding the contribution of each subset separately, then combining the bounds into one. The core of the proof is a careful application of a common packing argument in doubling metric spaces.

Tightness of Refinement. Now, let us bound the number of f -independent sets that are necessary to cover a feasible set. We show that this number is $O(f^*(\Delta(S)))$ for any feasible set S , where f^* is defined for every strongly sub-linear function, as follows. For each integer $c \geq 1$, the function $f^{(c)}(x)$ is defined inductively by: $f^{(1)}(x) = f(x)$ and $f^{(c)}(x) = f(f^{(c-1)}(x))$. Let $x_0 = \inf\{x \geq 1, f(x) < x\} + 1$; such a point exists for every $f(x) = o(x)$. The function $f^*(x)$, is defined by: $f^*(x) = \arg \min_c \{f^{(c)}(x) \leq x_0\}$ for arguments $x > x_0$, and $f^*(x) = 1$ for the rest. Note that for a function $f(x) = \gamma x^\delta$ with constants $\gamma > 0$ and $\delta \in (0, 1)$, $f^*(\Delta) = \Theta(\log \log \Delta)$, which is the tightness bound we are aiming for.

► **Theorem 4.** *Consider a non-decreasing strongly sub-linear function f . Every feasible set S can be split into $O(f^*(\Delta(S)))$ subsets, each independent in $\mathcal{G}_f(S)$.*

Let us fix a function f with properties as in the theorem. We establish the partition in Thm. 4 in two steps. The first step is to show that feasible set S can be partitioned into a constant number of independent sets in $\mathcal{G}_\rho^0(S)$ for any constant ρ , i.e., subsets S' such that for every pair of links $i, j \in S'$, $d_{ij}d_{ji} > \rho l_i l_j$. Such subsets are called ρ -independent for short. The second step is to show that for an appropriate constant ρ , each ρ -independent set can be partitioned into at most $O(f^*(\Delta))$ of f -independent subsets.

The first step is easy. Each feasible set can be partitioned into at most $2\rho^{\alpha/2}$ subsets, each of them feasible with updated thresholds $\{\rho^{\alpha/2}\beta_i\}$. This is a direct application of Corollary 2 of [3]. Let S' be such a subset and let $i, j \in S'$. The feasibility constraint for i and j implies:

$$P(i)/l_i^\alpha \geq \rho^{\alpha/2}\beta_i P(j)/d_{ji}^\alpha \text{ and } P(j)/l_j^\alpha \geq \rho^{\alpha/2}\beta_j P(i)/d_{ij}^\alpha.$$

By multiplying together the inequalities above, canceling $P(i)$ and $P(j)$ and raising to the power of $1/\alpha$, we obtain: $d_{ij}d_{ji} \geq \rho l_i l_j$, as required.

The proof of the second step requires the following lemmas, which constitute the most significant technical difference from the proof of the corresponding theorem in [14], as they encapsulate the technicalities of dealing with our definition of conflict graphs: It is not sufficient to bound only one of the distances between links (such as $d(i, j)$ in [14]); we need a bound on the product of two distances.

► **Lemma 5.** *Let i, j, k be such that $l_i \leq l_j \leq l_k$ and i is f -adjacent with both j and k , where f is a non-decreasing sublinear function. Then*

$$d_{jk}d_{kj} < 18l_i l_k f(l_k/l_i) + 13l_j l_k + 2l_j \sqrt{l_i l_k f(l_k/l_i)} + l_k \sqrt{l_i l_j f(l_j/l_i)}.$$

► **Lemma 6.** *Let i be a link and $\rho > 1$. If E is a ρ -independent set of links where each $j \in E$ is f -adjacent with i and satisfies $l_i \leq l_j \leq cl_i$ for a constant c , then $|E| = O(1)$.*

Proof of Theorem 4. By the discussion above, it is sufficient to show that each ρ -independent set S , for appropriate constant $\rho > 1$, can be partitioned into a small number of f -independent sets. We choose $\rho = 3c_f + 31$, where c_f is such that $f(x) \leq c_f x$ for all $x \geq 1$ (recall that f is sub-linear). Partitioning is done by the following *inductive* coloring procedure: 1. Consider the links in a non-increasing order by effective length, 2. Assign each link the first natural number that has not been assigned to an f -adjacent link yet. Clearly, such a procedure defines a partitioning of S into f -independent subsets.

Fix a link $i \in S$. Let T denote the set of links in $j \in S$ that have greater effective length than i and are f -adjacent with i . In order to complete the proof, it is enough to show that $|T| = O(f^*(\Delta))$, as $|T|$ is an upper bound on the number assigned to link i .

Since $f(x)$ is strongly sub-linear, there exists $x_0 = \inf\{x \geq 1, f(x) < x\} + 1$. Let us split T into two subsets A and B , where A contains the links $j \in T$ such that $l_j \leq x_0 l_i$ and $B = T \setminus A$. By Lemma 6, we have that $|A| = O(1)$, so we concentrate on B .

Let j, k be arbitrary links in B such that $l_j \leq l_k$. By applying Lemma 5 and using the definition of c_f , we obtain: $d_{jk}d_{kj} < 18l_i l_k f(l_k/l_i) + (3c_f + 13)l_j l_k$. Recall that j and k are $(\rho = 3c_f + 31)$ -independent, so $d_{jk}d_{kj} > (3c_f + 31)l_j l_k$, which gives us $l_j/l_i < f(l_k/l_i)$. Let $1, 2, \dots, t = |B|$ be an arrangement of the links in B in a non-decreasing order by effective length and let $\lambda_j = l_j/l_i$ for $j = 1, 2, \dots, t$. We have just shown that

$$x_0 \leq \lambda_1 < f(\lambda_2) \leq f(f(\lambda_3)) \leq \dots \leq f^{(t-1)}(\lambda_t),$$

namely, $t - 1 \leq f^*(\lambda_t) = O(f^*(\Delta))$. Thus, $|T| = |A| + |B| = O(1) + O(f^*(\Delta))$. ◀

Computability. Computability of our conflict graph construction is demonstrated through the notion of *inductive independence*. An n -vertex graph G is k -inductive independent if there is an ordering v_1, v_2, \dots, v_n of vertices such that for each v_i , the subgraph of G induced by the set $N_G(v_i) \cap \{v_i, v_{i+1}, \dots, v_n\}$ has no independent set larger than k , where $N_G(v)$ denotes the neighborhood of vertex v . It is well known, e.g. [1, 35], that vertex coloring and MWISL problems are k -approximable in k -inductive independent graphs.

► **Theorem 7.** *Let f be a non-decreasing strongly sub-linear function with $f(x) \geq 40$ for all $x \geq 1$. For every set L , the graph $\mathcal{G}_f(L)$ is constant inductive independent.*

The proof is somewhat similar to that of Thm. 4. The inductive independence ordering non-decreasing order of links by length. With this in mind, the proof of Thm. 4 can be applied, with the following core difference: while in Thm. 4 the goal was, for a link i , to bound the number of ρ -independent links that have greater effective length and are f -adjacent with i , here we need to bound the number of f -independent links that have greater effective length and are f -adjacent with i .

Optimality of $O(\log \log \Delta)$ -tightness. Here we show that the obtained tightness is essentially best possible, by demonstrating that every reasonable conflict graph formulation must incur an $O(\log \log \Delta)$ factor. We depart from some basic assumptions on conflict graphs. First, since the feasibility of a set of links is precisely determined by the values l_i and d_{ij} , we assume that in a conflict graph, the adjacency of two links i, j is a predicate of variables l_i, l_j, d_{ij}, d_{ji} . Another basic observation is that the feasibility formula is scale-free w.r.t. those values; hence, we assume that so is a conflict graph formulation. This allows us to reduce the number of variables in the adjacency predicate: $\frac{l_{max}}{l_{min}}, \frac{d_{ij}}{l_{min}}, \frac{d_{ji}}{l_{min}}$, where l_{min} and l_{max} are the smaller and larger values of l_i, l_j , respectively. Our construction will consist of only *unit-length* links (i.e. $l_i = 1$) of mutual distance at least 3. In this case, we can further reduce the number of variables by noticing that in such instances, $d_{ij} = \Theta(d_{ji}) = \Theta(d(i, j))$. Thus, the conflict relation is essentially determined by two variables: $\frac{d(i, j)}{l_{min}}$ and $\frac{l_{max}}{l_{min}}$. By separating the variables, the conflict predicate boils down to a relation $\frac{d(i, j)}{l_{min}} > f(\frac{l_{max}}{l_{min}})$ for a function f . Note that this is similar to the conflict graph definition of [14], except that the lengths are replaced with effective lengths.

Let us show that the refinement property requires that $f(x) = \Omega(\sqrt{x})$ in such a graph. Let us fix a function $f : [1, \infty) \rightarrow [1, \infty)$. Let i, j be unit-length links with $\beta_j = 1$ and $\beta_i = X^\alpha > 1$, where X is a parameter. Assume further that the links i, j are placed on the plane so that $d(i, j) = 3f(X) = 3f(l_i/l_j)$, which means the links are f -independent. Thus, i, j must form a feasible set: $\frac{P(i)}{l_i^\alpha} > \frac{P(j)}{d_{ji}^\alpha}$ and $\frac{P(j)}{l_j^\alpha} > \frac{P(i)}{d_{ij}^\alpha}$. Multiplying these inequalities together and canceling $P(i)$ and $P(j)$ out, gives: $d_{ij}d_{ji} > l_i l_j = X$. This implies that we must have $d(i, j) = \Theta(\sqrt{d_{ij}d_{ji}}) = \Omega(\sqrt{X})$, which in turn implies that $f(X) = d(i, j)/3 = \Omega(\sqrt{X})$.

Now, a simple modification of the construction in [14, Thm. 9] gives a set S of unit-length links arranged on the line and with appropriately chosen thresholds β_i and distances $d(i, j)$, such that every two links are f -adjacent, but the whole set S is feasible. Such a construction can be done with the number of links $n = \Omega(f^*(\Delta))$, i.e. there is a *feasible* set of links that cannot be split in less than $\Omega(f^*(\Delta))$ f -independent subsets. Since $f(x) = \Omega(\sqrt{x})$, we have $f^*(x) = \Omega(\log \log x)$, which proves that the tightness must be at least $\Omega(\log \log \Delta)$.

5 Approximating Fixed-Rate Scheduling

We detail now the more classical problems that can be handled with our framework, starting with those involving fixed data rates. Intuitively, our framework can handle a problem if there is a correspondence between solutions in the physical model instance and solutions in the refinement graph. The refinement property ensures that the graph solutions map directly to feasible solutions in the physical model — we need to ensure a (approximate) correspondence in the other direction. We will argue that an optimal solution in the physical model has a counterpart in the graph instance, whose quality decreases only by the tightness factor k .

General Approximation Framework. Common scheduling-related optimization problems can be classified as *covering* or *packing*.

In covering problems, a feasible solution σ contains a (ordered) covering of the set L of links with feasible sets $\pi = \langle S_1, S_2, \dots, S_t \rangle$ (i.e., $\cup_t S_t = L$), which we call *time slots*, and the objective is to minimize a function $f_\sigma(\pi)$ of the covering, which may also depend on other problem constraints.

In packing problems, a feasible solution σ contains a *fixed* number c of feasible sets (packing), $\eta = \langle S_1, S_2, \dots, S_c \rangle$, not necessarily covering L , which we call *channels*, and the objective is to maximize a function $g_\sigma(\eta)$ of the packing.

Given a refinement \mathcal{G} and a cover $\pi = \langle S_1, S_2, \dots, S_t \rangle$ of L by feasible sets, we call another cover $\pi' = \langle S_1^1, \dots, S_1^{h_1}, S_2^1, \dots, S_2^{h_2}, \dots, S_t^1, \dots, S_t^{h_t} \rangle$, a refinement of π if $\langle S_i^1, \dots, S_i^{h_i} \rangle$ is a cover of S_i by *independent* sets in \mathcal{G} . Similarly, given a packing $\eta = \langle S_1, S_2, \dots, S_c \rangle$, a refinement of η is another packing $\eta' = \langle S'_1, S'_2, \dots, S'_c \rangle$, where $S'_i \subseteq S_i$ is an independent set in \mathcal{G} .

Formally, a covering problem is *refinable* if for every k -tight refinement \mathcal{G} and a solution σ with cover π , there is a feasible solution σ' containing a refinement π' of π , and such that $f_\sigma(\pi) \geq \frac{f_{\sigma'}(\pi')}{k}$. A packing problem is *refinable* if for every k -tight refinement \mathcal{G} and a solution σ with a packing $\eta = \langle S_1, S_2, \dots, S_c \rangle$, there is a feasible solution σ' containing a refinement η' of η , and such that $g_\sigma(\eta) \leq k \cdot g_{\sigma'}(\eta')$.

► **Theorem 8.** *Let \mathcal{G} be a k -tight refinement of the physical model. For every refinable problem, a ρ -approximation algorithm in \mathcal{G} gives $k \cdot \rho$ -approximation in the physical model.*

Thus, in order to obtain an approximation for a specific problem, it is sufficient to show that the problem is refinable: then the solution in a k -tight refinement gives a solution with an additional approximation factor k . Refinability requires the objective function of the problem to have certain linearity property. Examples of refinable covering problems include the ones where the objective function is the number of time slots or the sum of completion times (i.e. indices of time slots). Perhaps the simplest example of a refinable packing problem is the *maximal independent set of links* problem, where the objective is the size of the feasible set (i.e., there is only a single channel). Below, we apply the refinement framework to some important scheduling problems, which leads to $O(\log \log \Delta)$ -approximation for all of them.

MWISL with Fixed Weights. Consider the MWISL problem, where the weights ω_i of links are fixed positive numbers. It is easy to see that this is a refinable packing problem, as the objective function – the sum of weights – is linear with respect to partition. Thus, since there is a constant factor approximation to MWISL in $\mathcal{G}(L)$ (by computability), it gives an $O(\log \log \Delta)$ -approximation in the physical model (by Thm. 8).

Multi-Channel Selection. Given a natural number c – the number of channels – the goal is to select a maximum number of links that can be partitioned into c feasible subsets (a subset for each channel). Again, this is easily seen to be a refinable packing problem, as the objective function – the total number of links across all channels, is linear w.r.t. partitioning. A simple greedy algorithm gives constant factor approximation to multi-channel selection in constant-inductive independent graphs, which translates to an $O(\log \log \Delta)$ -approximation in the physical model.

TDMA Scheduling. The goal is to partition the set L of links into the minimum number of feasible subsets. This is a covering problem, and the objective function is the number of slots, which is linear w.r.t. partitioning. A simple *first-fit* style greedy algorithm gives constant factor approximation to vertex coloring in constant inductive independent graphs, which gives an $O(\log \log \Delta)$ -approximation to TDMA scheduling in the physical model.

Fractional Scheduling. This is a fractional variant of TDMA scheduling with an additional constraint of link demands. A *fractional schedule* for a set L of links is a collection of feasible sets with rational values $\mathcal{S} = \{(I_k, t_k) : k = 1, 2, \dots, q\} \subseteq \mathcal{E}_{\mathcal{P}} \times \mathbb{R}_+$, where $\mathcal{E}_{\mathcal{P}}$ is the set of all feasible subsets of L . The sum $\sum_{k=1}^q t_k$ is the *length* of the schedule \mathcal{S} . The *link capacity vector* $c_{\mathcal{S}} : L \rightarrow \mathbb{R}_+$ associated with the schedule \mathcal{S} is given by $c_{\mathcal{S}}(i) = \sum_{(I,t) \in \mathcal{S}: I \ni i} t$.

Essentially, the link capacity shows how much scheduling time each link gets. Finally, a *link demand vector* $d : L \rightarrow \mathbb{R}_+$ indicates how much scheduling time each link needs.

The *fractional scheduling problem* is a covering type problem, where given a demand vector d , the goal is to compute a minimum length schedule that serves the demands d , namely, for each link $i \in L$, $c_{\mathcal{S}}(i) \geq d(i)$. Since the cost function $\sum_{k=1}^q t_k$ is again linear w.r.t. partitioning of a schedule, it is readily checked that the fractional scheduling problem is also refinable. A simple greedy algorithm presented in [31] achieves constant factor approximation for fractional scheduling in constant inductive independent graphs. This gives an $O(\log \log \Delta)$ -approximation in the physical model.

Joint Routing and Scheduling. Consider an ordered set of p source-destination node pairs (multihop communication requests) (u_i, v_i) , $i = 1, 2, \dots, p$, with associated weights/utilities $\omega_i > 0$, in a multihop network given by a directed graph G , where the *edges* of the graph are the transmission links. Let \mathcal{P}_i denote the set of directed (u_i, v_i) paths in G and let $\mathcal{P} = \cup_i \mathcal{P}_i$. Then a *path flow* for the given set of requests is a set $\mathcal{F} = \{(P_k, \delta_k) : k = 1, 2, \dots\} \subseteq \mathcal{P} \times \mathbb{R}_+$. The *link flow vector* $f_{\mathcal{F}}$ corresponding to path flow \mathcal{F} , with $f_{\mathcal{F}}(i) = \sum_{(P, \delta) \in \mathcal{F} : P \ni i} \delta$ for each link i , shows the flow along each link.

The *multiflow routing and scheduling problem* is a covering problem, where given source-destination pairs with associated utilities, the goal is to find a path flow \mathcal{F} together with a fractional link schedule \mathcal{S} of length 1, such that² for each link i , the link flow is at most the link capacity provided by the schedule, $f_{\mathcal{F}}(i) \leq c_{\mathcal{S}}(i)$, and the *flow value*

$$W = \sum_{i=1}^p \omega_i \cdot \sum_{(P_k, \delta_k) \in \mathcal{F}, P_k \in \mathcal{P}_i} \delta_k$$

is maximized. Let us verify that this problem is also refinable. Consider a feasible solution in (the physical model) that consists of a path flow $\mathcal{F} = \{(P_k, \delta_k) : k = 1, 2, \dots\}$ and a schedule $\mathcal{S} = \{(I_k, t_k) : k = 1, 2, \dots\}$ of length $\sum_{k \geq 1} t_k = 1$, such that $f_{\mathcal{F}}(i) \leq c_{\mathcal{S}}(i)$. As observed in the previous section, the schedule \mathcal{S} can be refined into a schedule $\mathcal{S}' = \{(I_k^s, t_k^s)\}_{k,s}$ in $\mathcal{G}(L)$, where \mathcal{S}' serves the same demand vector as \mathcal{S} does, and \mathcal{S}' has length at most $K = O(\log \log \Delta)$ times more than the length of \mathcal{S} . Now we normalize the refined schedule to have length 1. Then, the following modified path flow $\mathcal{F}' = \{(P_k, \delta_k/K) : k = 1, 2, \dots\}$ together with the new schedule will be feasible in $\mathcal{G}(L)$, as all link demands will be served. Moreover, the value of \mathcal{F}' is at most K times that of \mathcal{F} . Hence, the problem is refinable.

Thus, applying the constant factor approximation algorithm of [32] for constant inductive independent conflict graphs (the result holds with unit utilities) gives an $O(\log \log \Delta)$ -approximation for multiflow routing and scheduling problem in the physical model. It should also be noted that the fractional scheduling and routing and scheduling problems can be reduced to the MWISL problem using linear programming techniques (described e.g. in [17]), as it was shown in [30]. We will further discuss this in Section 6.

Extensions to Multi-Channel Multi-Antenna Settings. All problems above may be naturally generalized to the case when there are several channels (e.g. frequency bands) available and moreover, wireless nodes are equipped with multiple antennas and can work in different channels simultaneously. We denote the setting with multiple antennas/channels as *MC-MA*.

² Essentially, the schedule here gives a probability distribution over the feasible sets of links.

It is easy to show that our refinement framework can be extended to MC-MA with very little change. Assume that each node u is equipped with $a(u)$ antennas numbered from 1 to $a(u)$ and can use a set $\mathcal{C}(u)$ of channels. Consider a link i that corresponds to the pair of nodes s_i and r_i . There are $a(s_i)a(r_i)|\mathcal{C}(s_i) \cap \mathcal{C}(r_i)|$ *virtual* links corresponding to each selection of an antenna of the sender node s_i , an antenna of receiver node r_i and a channel $c \in \mathcal{C}(s_i) \cap \mathcal{C}(r_i)$ available to both nodes. Thus a virtual link is described by the tuple (i, a_s, a_r, c) , where a_s (a_r) denotes the antenna index at s_i (r_i , respectively), and c denotes the channel. We call link i *the original* of its virtual links. Note that the formulation above can easily be generalized to the case where certain antennas don't work in certain channels, e.g., due to multi-path fading.

A set of (virtual) links S is feasible in MC-MA if and only if no two links in S share an antenna (i.e., they do not use the same antenna of the same node), and for each channel c , the set of originals of links in S using channel c is feasible in the physical model. Then, an $O(\log \log \Delta)$ -tight refinement for the MC-MA physical model by a conflict graph can be found by a simple extension of the existing refinement for the single channel case to the virtual links. This implies, in particular, that all scheduling problems considered in the previous sections can also be approximated in the MC-MA setting within an approximation factor $O(\log \log \Delta)$, as the corresponding approximations for the conflict graph hold with MC-MA [32].

6 Rate Control and Scheduling

The most important application of efficient approximation algorithms for scheduling problems with different thresholds is the application to scheduling with rate control. This is achieved first by obtaining a double-logarithmic approximation to MWISL with rate control. This will then lead to similar approximations for fractional scheduling and joint routing and scheduling problems.

MWISL with Rate Control. By Shannon's theorem, given a set S of links simultaneously transmitting in the same channel, the transmission rate $r(S, i)$ of a link i is a function of $SIR(S, i)$. Thus, we consider the MWISL problem where each link i has an associated *utility function* $u^i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, and the weight of link i is the value of u^i at $SIR(S, i)$ if link i is selected in the set, and 0 otherwise. As before, the goal is, given the links with utility functions, to find a subset S that maximizes the total weight $\sum_{i \in S} u^i(r(S, i))$. We assume that $u^i(SIR(S, i)) = 0$ if $SIR(S, i) < 1$.

An $O(\log n)$ -approximation for this variant of MWISL has been obtained in [23]. We show that this can be replaced with $O(\log \log \Delta')$, where $\Delta'(L) = \max_{i,j \in L} \frac{u_{max}^i l_i}{u_{min}^j l_j}$ and u_{min}^i, u_{max}^i are the minimum and maximum possible utility values for the given instance and link. This is achieved by reducing the problem to MWISL in an extended instance.

Let us fix a utility function u . First, assume that the possible set of weights for each link is a discrete set $u_{min} = u_1 < u_2 < \dots < u_t = u_{max}$. Then, we can replace each link i with t copies i_1, i_2, \dots, i_t with different thresholds and fixed weights, where $\omega_{i_k} = u_k$ and $\beta_{i_k} = \min\{x : u^{i_k}(x) \geq u_k\}$ if $\beta_{i_k} \geq 1$ and $\omega_{i_k} = 0$ otherwise. Now, the problem becomes a MWISL problem for the modified instance L' with link replicas and fixed weights. Observe that no feasible set in L' contains more than a single copy of the same link, as the copies occupy the same geometric place, implying that each feasible set of the extended instance corresponds to a feasible set of the original instance, with an obvious transformation. The effective length diversity of the extended instance is $\Delta(L') = \Delta'(L)$. Thus, using the approximation

algorithm for the fixed rate MWISL problem, we obtain an $O(\log \log \Delta'(L))$ -approximation for MWISL with rate control.

For the case when the number of possible utility values is too large or the set is continuous, a standard trick can be applied. Let u_{max}^i, u_{min}^i be as before. The extended instance L' is constructed by replacing each link i with $O(\log u_{max}^i/u_{min}^i)$ copies i_1, i_2, \dots of itself and assigning each replica i_k weight $\omega_k = 2^{k-1}$ and threshold $\beta_k = \min\{x : 2^{k-1} \leq u^i(x) \leq 2^k\}$ if $\beta_k \geq 1$ and let $\omega_k = 0$ otherwise. It is easy to see that the optimum value of MWISL with fixed rates in L' is again an $O(\log \log \Delta'(L))$ -approximation to MWISL with rate control.

If the value $\log u_{max}^i/u_{min}^i$ is still too large, it may be inefficient to have $O(\log u_{max}^i/u_{min}^i)$ copies for each link. It is another standard observation that only the last $O(\log n)$ copies of each link really matter, as restricting to only those links degrades approximation by a factor of at most 2.

Fractional Scheduling with Rate Control. In this formulation, we redefine a fractional schedule to be a set $\mathcal{S} = \{(I_k, t_k) : k = 1, 2, \dots, q\} \subseteq 2^L \times \mathbb{R}_+$, namely, I_k are arbitrary subsets, rather than independent ones. We redefine the link capacity vector $\hat{c}_{\mathcal{S}}$ to incorporate the rates as follows:

$$\hat{c}_{\mathcal{S}}(i) = \sum_{(I,t) \in \mathcal{S}: I \ni i} t \cdot r(i, I). \quad (2)$$

The *fractional scheduling with rate control* problem is to find a minimum length schedule \mathcal{S} that serves a given demand vector d , namely, such that for each link $i \in L$, $\hat{c}_{\mathcal{S}}(i) \geq d(i)$.

The problem can be formulated as an exponential size linear program LP_1 , as follows.

$$\begin{aligned} \min \sum_{I \subseteq L} t_I \text{ subject to } & \sum_{I \subseteq L: I \ni i} t_I \cdot r(i, I) \geq d(i) & \forall i \in L \\ & t_I \geq 0 & \forall I \subseteq L \end{aligned}$$

The dual program LP_2 looks as follows:

$$\begin{aligned} \max \sum_{i \in L} d(i) y_i \text{ subject to } & \sum_{i \in I} y_i \cdot r(i, I) \geq 1 & \forall I \subseteq L \\ & y_i \geq 0 & \forall i \in L \end{aligned}$$

As [17, Thm. 5.1] states, if there is an approximation algorithm that finds a set \hat{I} such that $\sum_{i \in \hat{I}} y_i r(i, \hat{I}) \geq \frac{1}{a} \max_{I \subseteq L} \sum_{i \in I} y_i r(i, I)$, then there is an a -approximation algorithm for LP_1 , where the former algorithm acts as an approximate separation oracle for LP_1 . But this auxiliary problem is simply a special case of the MWISL with rate control, which we can approximate within a double-logarithmic factor. Thus, there is an approximation preserving reduction from the fractional scheduling with rate control to MWISL with rate control. By the obtained approximation for MWISL, we obtain an $O(\log \log \Delta')$ -approximation for fractional scheduling with rate control.

Routing, Scheduling and Rate Control. The rate-control variant of the routing and scheduling problem is formulated in the same way as for the fixed rate setting, with the only modified constraint being the capacity constraints, which, instead of the link capacity vector $c_{\mathcal{S}}$, now use the modified variant $\hat{c}_{\mathcal{S}}$ that incorporates the link rates (see the definition in (2)).

This problem can also be reduced to MWISL with rate control, using similar methods as for the fractional scheduling problem. The reduction is nearly identical to the reduction of fixed rate versions of these problems to MWISL, presented in [30, Thm. 4.1].

Thus, we can conclude that there is an $O(\log \log \Delta')$ -approximation algorithm for joint routing, scheduling and rate control that uses MWISL with rate control as a subroutine.

References

- 1 Karhan Akcoglu, James Aspnes, Bhaskar DasGupta, and Ming-Yang Kao. Opportunity cost algorithms for combinatorial auctions. *CoRR*, cs.CE/0010031, 2000.
- 2 Mahmoud Al-Ayyoub and Himanshu Gupta. Joint routing, channel assignment, and scheduling for throughput maximization in general interference models. *IEEE Trans. Mob. Comput.*, 9(4):553–565, 2010. doi:10.1109/TMC.2009.144.
- 3 Jørgen Bang-Jensen and Magnús M. Halldórsson. Vertex coloring edge-weighted digraphs. *Inf. Process. Lett.*, 115(10):791–796, 2015. doi:10.1016/j.ipl.2015.05.007.
- 4 Deepti Chafekar, V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Cross-layer latency minimization in wireless networks with SINR constraints. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2007, Montreal, Quebec, Canada, September 9-14, 2007*, pages 110–119, 2007. doi:10.1145/1288107.1288123.
- 5 Michael Dinitz. Distributed algorithms for approximating wireless network capacity. In *INFOCOM 2010. 29th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 15-19 March 2010, San Diego, CA, USA*, pages 1397–1405, 2010. doi:10.1109/INFCOM.2010.5461905.
- 6 Alexander Fanghänel, Thomas Kesselheim, and Berthold Vöcking. Improved algorithms for latency minimization in wireless networks. *Theor. Comput. Sci.*, 412(24):2657–2667, 2011. doi:10.1016/j.tcs.2010.05.004.
- 7 Liqun Fu, Soung Chang Liew, and Jianwei Huang. Power controlled scheduling with consecutive transmission constraints: Complexity analysis and algorithm design. In *INFOCOM 2009. 28th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 19-25 April 2009, Rio de Janeiro, Brazil*, pages 1530–1538, 2009. doi:10.1109/INFCOM.2009.5062070.
- 8 Olga Goussevskaia, Yvonne Anne Oswald, and Roger Wattenhofer. Complexity in geometric SINR. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2007, Montreal, Quebec, Canada, September 9-14, 2007*, pages 100–109, 2007. doi:10.1145/1288107.1288122.
- 9 Olga Goussevskaia, Luiz Filipe M. Vieira, and Marcos Augusto M. Vieira. Wireless scheduling with multiple data rates: From physical interference to disk graphs. *Computer Networks*, 106:64–76, 2016. doi:10.1016/j.comnet.2016.06.016.
- 10 Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000. doi:10.1109/18.825799.
- 11 Magnús M. Halldórsson. Wireless scheduling with power control. *ACM Trans. Algorithms*, 9(1):7:1–7:20, 2012. doi:10.1145/2390176.2390183.
- 12 Magnús M. Halldórsson and Pradipta Mitra. Wireless capacity with oblivious power in general metrics. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1538–1548, 2011. doi:10.1137/1.9781611973082.119.
- 13 Magnús M. Halldórsson and Pradipta Mitra. Wireless capacity and admission control in cognitive radio. In *Proceedings of the IEEE INFOCOM 2012, Orlando, FL, USA, March 25-30, 2012*, pages 855–863, 2012. doi:10.1109/INFCOM.2012.6195834.
- 14 Magnús M. Halldórsson and Tigran Tonoyan. How well can graphs represent wireless interference? In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 635–644, 2015. doi:10.1145/2746539.2746585.

- 15 Magnús M. Halldórsson and Tigran Tonoyan. The price of local power control in wireless scheduling. In *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2015, December 16-18, 2015, Bangalore, India*, pages 529–542, 2015. doi:10.4230/LIPIcs.FSTTCS.2015.529.
- 16 Juha Heinonen. *Lectures on Analysis on Metric Spaces*. Springer, 1. edition, 2000.
- 17 Klaus Jansen. Approximate strong separation with application in fractional graph coloring and preemptive scheduling. *Theor. Comput. Sci.*, 302(1-3):239–256, 2003. doi:10.1016/S0304-3975(02)00829-0.
- 18 Libin Jiang, Mathieu Leconte, Jian Ni, R. Srikant, and Jean C. Walrand. Fast mixing of parallel glauber dynamics and low-delay CSMA scheduling. *IEEE Trans. Information Theory*, 58(10):6541–6555, 2012. doi:10.1109/TIT.2012.2204032.
- 19 Libin Jiang and Jean C. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.*, 18(3):960–972, 2010. doi:10.1109/TNET.2009.2035046.
- 20 Changhee Joo, Xiaojun Lin, and Ness B. Shroff. Understanding the capacity region of the greedy maximal scheduling algorithm in multihop wireless networks. *IEEE/ACM Trans. Netw.*, 17(4):1132–1145, 2009. doi:10.1145/1618562.1618572.
- 21 Bastian Katz, Markus Völker, and Dorothea Wagner. Energy efficient scheduling with power control for wireless networks. In *8th International Symposium on Modeling and Optimization in Mobile, Ad-Hoc and Wireless Networks (WiOpt 2010), May 31 - June 4, 2010, University of Avignon, Avignon, France*, pages 160–169, 2010.
- 22 Thomas Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1549–1559, 2011. doi:10.1137/1.9781611973082.120.
- 23 Thomas Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *Algorithms - ESA 2012 - 20th Annual European Symposium, Ljubljana, Slovenia, September 10-12, 2012. Proceedings*, pages 659–670, 2012. doi:10.1007/978-3-642-33090-2_57.
- 24 Henry Lin and Frans Schalekamp. On the complexity of the minimum latency scheduling problem on the euclidean plane. *CoRR*, abs/1203.2725, 2012.
- 25 Xiaojun Lin and N. B. Shroff. Joint rate control and scheduling in multihop wireless networks. In *2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601)*, volume 2, pages 1484–1489 Vol.2, Dec 2004. doi:10.1109/CDC.2004.1430253.
- 26 Thomas Moscibroda and Roger Wattenhofer. The complexity of connectivity in wireless networks. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 23-29 April 2006, Barcelona, Catalunya, Spain, 2006*. doi:10.1109/INFOCOM.2006.23.
- 27 Devavrat Shah, David N. C. Tse, and John N. Tsitsiklis. Hardness of low delay network scheduling. *IEEE Trans. Information Theory*, 57(12):7810–7817, 2011. doi:10.1109/TIT.2011.2168897.
- 28 L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, Dec 1992. doi:10.1109/9.182479.
- 29 Tigran Tonoyan. On some bounds on the optimum schedule length in the SINR model. In *Algorithms for Sensor Systems, 8th International Symposium on Algorithms for Sensor Systems, Wireless Ad Hoc Networks and Autonomous Mobile Entities, ALGOSENSORS 2012, Ljubljana, Slovenia, September 13-14, 2012. Revised Selected Papers*, pages 120–131, 2012. doi:10.1007/978-3-642-36092-3_14.

- 30 Peng-Jun Wan. Multiflows in multihop wireless networks. In *Proceedings of the 10th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2009, New Orleans, LA, USA, May 18-21, 2009*, pages 85–94, 2009. doi:10.1145/1530748.1530761.
- 31 Peng-Jun Wan, Xiaohua Jia, Guojun Dai, Hongwei Du, Zhiguo Wan, and Ophir Frieder. Scalable algorithms for wireless link schedulings in multi-channel multi-radio wireless networks. In *Proceedings of the IEEE INFOCOM 2013, Turin, Italy, April 14-19, 2013*, pages 2121–2129, 2013. doi:10.1109/INFCOM.2013.6567014.
- 32 Peng-Jun Wan, Zhu Wang, Lei Wang, Zhiguo Wan, and Sai Ji. From least interference-cost paths to maximum (concurrent) multiflow in MC-MR wireless networks. In *2014 IEEE Conference on Computer Communications, INFOCOM 2014, Toronto, Canada, April 27 - May 2, 2014*, pages 334–342, 2014. doi:10.1109/INFCOM.2014.6847955.
- 33 Lixin Wang, C. P. Abubucker, William F. Lawless, and Anthony J. Baker. A constant-approximation for maximum weight independent set of links under the SINR model. In *Seventh International Conference on Mobile Ad-hoc and Sensor Networks, MSN 2011, Beijing, China, December 16-18, 2011*, pages 9–14, 2011. doi:10.1109/MSN.2011.1.
- 34 Xinzhou Wu, R. Srikant, and James R. Perkins. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. *IEEE Trans. Mob. Comput.*, 6(6):595–605, 2007. doi:10.1109/TMC.2007.1061.
- 35 Yuli Ye and Allan Borodin. Elimination graphs. *ACM Trans. Algorithms*, 8(2):14:1–14:23, 2012. doi:10.1145/2151171.2151177.