How Well Can Graphs Represent Wireless Interference?*

[Extended Abstract]†

Magnús M. Halldórsson
ICE-TCS, Reykjavik University
RIMS, Kyoto University
mmh@ru.is

Tigran Tonoyan
ICE-TCS, Reykjavik University
ttonoyan@gmail.com

ABSTRACT

Efficient use of a wireless network requires that transmissions be grouped into feasible sets, where feasibility means that each transmission can be successfully decoded in spite of the interference caused by simultaneous transmissions. Feasibility is most closely modeled by a signal-to-interference-plus-noise (SINR) formula, which unfortunately is conceptually complicated, being an asymmetric, cumulative, many-to-one relationship.

We re-examine how well graphs can capture wireless receptions as encoded in SINR relationships, placing them in a framework in order to understand the limits of such modelling. We seek for each wireless instance a pair of graphs that provide upper and lower bounds on the feasibility relation, while aiming to minimize the gap between the two graphs. The cost of a graph formulation is the worst gap over all instances, and the price of (graph) abstraction is the smallest cost of a graph formulation.

We propose a family of conflict graphs that is parameterized by a non-decreasing sub-linear function, and show that with a judicious choice of functions, the graphs can capture feasibility with a cost of $O(\log^* \Delta)$, where Δ is the ratio between the longest and the shortest link length. This holds on the plane and more generally in doubling metrics. We use this to give greatly improved $O(\log^* \Delta)$-approximation for fundamental link scheduling problems with arbitrary power control. We also explore the limits of graph representations and find that our upper bound is tight: the price of graph abstraction is $\Omega(\log^* \Delta)$. In addition, we give strong impossibility results for general metrics, and for approximations in terms of the number of links.

*The work is supported by grant-of-excellence 120032011 and grant 152679-051 from the Icelandic Research Fund.

†Full version is available at http://arxiv.org/abs/1411.1263

STOC'15, June 14–17, 2015, Portland, Oregon, USA.

http://dx.doi.org/10.1145/2746539.2746585.

Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless communication*; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Sequencing and scheduling*

General Terms

Theory, Algorithms

Keywords

Wireless Networks, SINR, Conflict Graphs, Scheduling

1. INTRODUCTION

Wireless scheduling. At the heart of any wireless network is a mechanism for managing *interference* between simultaneous transmissions. The medium access (MAC) layer manages access to the shared resource, the wireless spectrum, balancing the aim of maximizing simultaneous use with the impact of the resulting interference. We can represent a transmission as a communication *link*, a sender-receiver pair of nodes in a metric space. Wireless scheduling mechanisms assign the links to different "slots", involving different frequencies, phases and/or time steps.

The model of communication that most closely captures actual conditions, nicknamed the *physical model*, uses a formula based on the ratio of the (intended) signal strength to the received interference strength (SINR) to determine if decoding is successful. A subset X of links is *feasible* if there exists a power assignment to the senders such that each link i satisfies the SINR formula $\frac{S_i}{\sum_{j\in X} I_{ji}+N} \geq \beta$ within its subset, where S_i is the received signal strength on link i, β and N are fixed constants (dependent on technology and environment), and I_{ji} is the interference strength of link j on link i [15] (see Sec. 5 for full definitions). We can avoid dealing directly with power assignments using a condition for feasibility due to Kesselheim [28] (and shown here to be necessary). The point remains, however, that the feasibility predicate is an asymmetric, cumulative, many-to-one relationship.

Capturing interference with graphs. The aim of this work is to capture the complex feasibility relationship of the physical model with *graphs*, a much more amenable and better studied model. Preferably, a feasible set of links should correspond to an *independent set* in the graph on the links, and vice versa, but since exact capture is impossible, we seek instead approximate representations. Specifically, we want a formulation that constructs a *pair* of graphs that bound feasibility both from below and from above. The two graphs should also be "close" in some sense; specifically we want an independent set in the lower bound graph to induce a low chromaticity subgraph in the upper bound graph. That is, given a set L of links, form graphs $G_l(L)$ and $G_u(L)$ on L such that:

- If S is feasible subset of links, then it is an independent set in $G_l(L)$,
- If S is an independent set in $G_u(L)$, then S is feasible set of links, and
- Chromatic numbers of the subgraphs induced by a subset S are close, or $\max_{S\subseteq L} \frac{\chi(G_u[S])}{\chi(G_l[S])} \leq \rho$.

We may dub the worst ratio ρ over all instances as the *cost* of that graph formulation, the price we pay for using that simpler pairwise and binary graph representation. The least such cost over all graph formulation can then be called the *price of (graph) abstraction*.

Our results. We propose a family of conflict graph representations, parameterized by a sub-linear, non-decreasing function f. It generalizes known families, e.g., disc graphs correspond to linear functions f and pairwise feasibility is captured by constant functions. The graphs in the family have a structural property that allows for effective approximability. We argue that this family captures all meaningful conflict graph representations, modulo constant factors. Our main positive result is that for the right choice of f, our conflict graph representation has a cost of $O(\log^* \Delta)$, where Δ is the ratio between longest and the shortest link length. We also show that all meaningful representations must pay this $\log^* \Delta$ factor. Thus, the price of abstraction, for the SINR model with arbitrary power control, is $\Theta(\log^* \Delta)$. The upper bounds hold for planar instances, and more generally in doubling metrics, while the lower bounds are on the line. We also find that no such results are possible in general metrics nor in terms of the parameter n, the number of links.

We apply our formulations to obtain greatly improved bounds for fundamental wireless scheduling problems. In the link Scheduling problem, we want to partition a given set of links into fewest possible feasible sets. In the *weighted capacity* problem, WCapacity, the links have associated positive weight, and we want to find the maximum weight feasible subset. In both cases, our $O(\log^* \Delta)$-approximations are the first sub-logarithmic approximations known.

Related work. Gupta and Kumar [15] proposed the geometric version of the SINR model, where signal decays as a fixed polynomial of distance; it has since been the default in analytic and simulations studies. They also initiated the average-case analysis of network capacity, giving rise to a large body on "scaling laws". Moscibroda and Wattenhofer [37] initiated worst-case analysis in the SINR model.

Graph-based models of wireless communication have been very common in the past. Most common are geometric graphs involving circular ranges: *unit-disc graphs* (UDG) have all ranges of the same radius, while in *disc graphs*, the radius can vary with the power assigned, and in the *protocol* model [15] the communication and interference ranges are different. Various attempts have been made to add realism, such as with 2-hop model, or *quasi-unit disc graphs* [2, 33], and the recently studied model of *dual graphs* [32] captures arbitrary unreliability in networks. None of these known models offers though any guarantees on fidelity for representing SINR relationships, see e.g. [38]. Graph formulations for modeling SINR relationships were given previously in [17] followed by [46], but the cost factor was either $O(\log\log\Delta \cdot \log n)$ or $O(\log \Delta)$.

Early work on the Scheduling problem includes [5, 7, 4, 13]. NP-completeness results have been given for different variants [13, 27, 34], but as of yet no APX-hardness or stronger lower bounds are known for any related problem in geometric settings, perhaps indicating the difficulty of dealing with the SINR constraints. The related Capacity problem, where we seek to find a maximum feasible subset of links, admits constant-factor approximation [28]. This immediately implies a $O(\log n)$-approximation for Scheduling, where n is the number of links. Another approach is to solve links of similar lengths in groups, which results in a $O(\log \Delta)$-approximation [13, 11, 17]. The question of improved approximation for Scheduling has been frequently cited as perhaps the most fundamental open problem in the field [35, 25, 10, 12, 16]. The WCapacity problem has applications in several extensions of unit-demand link scheduling problems, such as *stochastic packet scheduling* [44, 41, 30], *general demand vectors, multi-path flow* etc. [47]. A more general variant of this problem has been considered in the framework of *combinatorial auctions* [24, 23]. Scheduling and WCapacity have also been considered with *fixed oblivious power assignments* [21, 8, 17, 19, 9], but the only known sub-logarithmic approximations are known in the case of the *linear power scheme* [19, 45].

Issues of models. A wide range of areas in computer science and mathematics deal with finding simpler abstractions of complex phenomena. Some examples include: a) discrepancy theory, b) dimensionality reductions and embeddings, c) graph augmentations and sandwiching properties, d) graph sparsification, e) curve fitting (including least squares, finite methods, and regression), f) approximation theory (in math), including generalized Fourier series and Chebyshev polynomials, and g) PAC learning.

There are tradeoffs between the accuracy of a model and its complexity of detail. There are legitimate concerns that models and problem formulations are sometimes overly detailed and "brittle", possibly exceeding reasonable levels of precision (see, e.g., discussion in [39]). The benefits of a coarser model tend to include simpler algorithms, easier analysis, but also less sensitivity to incidental details that may or may not be modelable.

A case in point is the SINR model, which has its issues. Whereas the additivity of interference and the near-threshold nature of signal reception has been borne out in experiments, the geometric decay assumption is far off in essentially all actual environments [43, 36, 42, 14]. One practical alternative is to use facts-on-the-ground in the form of signal strength measurements, instead of the prescriptive distance-based formula [14, 3]. To model that formally, the pessimistic reaction would be to replace the distance assumption with an arbitrary signal-quality matrix, but that runs into the computational intractability monster, since such a formulation can encode the coloring problem in general graphs [12]. A more moderate approach is to relax the Euclidean assumption to more general metric spaces [9]. The determinacy of the model is another issue. To capture the probabilistic factors observed in the capture of transmissions, one approach is to extend the basic SINR model accordingly, such as with Rayleigh fading; in that case, it has been shown that applying algorithms based on the deterministic formula results in nearly equally good results in that probabilistic setting [6].

Our results suggest that a reassessment of the role of graphs as wireless models might be in order. By paying a small factor (recalling, as well, that $\log^*(x) \leq 5$ in this universe), we can work at higher levels of abstraction, with all the algorithmic and analytical benefits that it accrues. At the same time, hopes for fully constant-factor approximation algorithms for core scheduling problems have receded. It remains to be seen what abstractions are possible for other related settings, especially the case of uniform power.

Roadmap to the rest of the paper. Following the basic definitions in Sec. 2, we derive from first principles what properties link conflict graphs must satisfy (Sec. 3). This can be read independently of the rest of the paper. We next derive (in Sec. 4) two key properties of the family of graphs: how their chromatic numbers relate and how their colorings can be approximated. We then introduce the definitions of the SINR model before starting the technical core of the paper. In Sec. 6, we show that feasibility is captured by two members of our conflict graph family. The implications are discussed in Sec. 7: our main upper bound result, i.e. $O(\log^* \Delta)$-approximation for Scheduling and WCapacity; a necessary and sufficient condition for feasibility; and an explicit polynomial-time computable measure of interference. Limitation results are given in Sec. 8, in particular that no better bounds are possible via conflict graphs. For space reasons, only sketches of the proofs are given in this extended abstract.

2. DEFINITIONS: METRICS, FUNCTIONS, GRAPHS

Communication Links. Consider a set L of n *links*, numbered from 1 to n. Each link i represents a unit-demand communication request from a sender s_i to a receiver r_i — point-size wireless transmitter/receivers (nodes) located in a metric space with distance function d. We denote $d_{ij} = d(s_i, r_j)$ and refer to $l_i = d(s_i, r_i)$ as the *length* of link i. We let $\Delta(L)$ denote the ratio between the longest and the shortest link lengths in L, and drop L when clear from context. We shall assume in the rest of the paper that all link lengths are distinct, which can be achieved by arbitrarily (but consistently) breaking ties as needed. For sets S_1, S_2 of links, we let $d(S_1, S_2)$ denote the minimum distance between *a node* in S_1 and a node in S_2. In particular, we will extensively use $d(i,j) = \min(d_{ij}, d_{ji}, d(s_i, s_j), d(r_i, r_j))$, the minimum distance between nodes on links i, j.

Let $S_i^+ = \{j \in S : l_j > l_i\}$ denote the subset of links in a set S that are longer than link i, and similarly $S_i^- = \{j \in S : l_j < l_i\}$ the subset of links shorter than i.

Doubling 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, and $0 < \epsilon \leq 1$. Metrics with finite doubling dimensions are said to be *doubling*. For example, the m-dimensional Euclidean space has doubling dimension m [22]. We let m denote the doubling dimension of the space containing the links.

Functions. A function f is *sub-linear* if $f(x) = O(x)$. 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 $f(x) = o(x)$.

Examples. The functions $f(x) = x^{1-\epsilon}$ for any constant $\epsilon > 0$ and $f(x) = \log x$ are strongly sub-linear[1], while $f(x) = x/\log x$ is not strongly sub-linear even though $f(x) = o(x)$.

Let f be a strongly sub-linear function. 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 any $f(x) = o(x)$. The function *iterated* f, denoted $f^*(x)$, is defined by:

$$f^*(x) = \begin{cases} \min_c\{f^{(c)}(x) \leq x_0\}, & \text{if } x > x_0, \\ 1, & \text{otherwise.} \end{cases}$$

Examples. For $f(x) = \log x$, $f^*(x) = \log^* x$ is the well known *iterated logarithm*. It is also easy to check that for $f(x) = \log^{(c)} x$, $f^*(x) = \lceil \frac{\log^* x}{c} \rceil$ and for $f(x) = \sqrt{x}$, $f^*(x) = \lceil \log\log x \rceil$. Note also that $g^*(x) = \Theta(f^*(x))$ if $g(x) = cf(x)$ for a constant $c > 0$.

Graphs. For a graph G and a vertex v, $N_G(v)$ denotes the neighborhood of v in G, i.e., the set of vertices adjacent with v. $\chi(G)$ denotes the chromatic number of graph G, the minimum number of colors needed for a proper (vertex) coloring of G.

A *d-inductive order* of a graph is an arrangement of the vertices from left to right such that each vertex has at most d *post-neighbors*, or neighbors appearing to its right. A *k-simplicial elimination order* is one where the post-neighbors of each vertex can be covered with k cliques. A graph is d-inductive (k-simplicial) if it has a d-inductive (k-simplicial elimination) order. It is well known that a d-inductive graphs are $d+1$-colorable, while the coloring and weighted maximum independent set problems are k-approximable in k-simplicial graphs [1, 26, 48]. The only inductive or simplicial order we consider for conflict graphs is the *increasing order* of links by length.

3. CONFLICT GRAPHS FORMULATIONS

What kind of graphs are conflict graphs? We propose here that essentially the only reasonable forms of conflict graphs, at least for capturing SINR properties, are the graphs $\mathcal{G}_f(L)$ defined below.

Definition 1. Let $f : \mathbb{R}_+ \to \mathbb{R}_+$ be a non-decreasing sub-linear function. Two links i, j are said to be *f-independent* if

$$\frac{d(i,j)}{l_{min}} > f\left(\frac{l_{max}}{l_{min}}\right),$$

where $l_{min} = \min\{l_i, l_j\}, l_{max} = \max\{l_i, l_j\}$, and otherwise they are *f-adjacent*. A set of links is f-independent if they are pairwise f-independent. The conflict graph $\mathcal{G}_f(L)$, for a set L of links, is the graph with vertex set L, where two vertices $i, j \in L$ are adjacent if and only if they are f-adjacent.

Remark. For the constant function $f(x) \equiv \gamma$ for a number γ, we use the notation $\mathcal{G}_\gamma(L)$ for the corresponding conflict graph.

[1] When not otherwise identified, logarithms are base 2.

First of all, a conflict graph formulation should specify a *deterministic rule* for forming a graph on top of a set of links. It should be defined only *in terms of the pairwise relationship* between links, and thus the only relevant properties of concern are the $\binom{4}{2} = 6$ distances between the nodes on the links.

Further, the formulation should be *independent of positions and scale*, which is a distinguishing feature of SINR relationships. Thus, we can factor out the length of the shorter of the two links under consideration. A useful feature of the feasibility characterization that we derive in Sec. 7.3 (based on a sufficient condition of Kesselheim, given here as Thm. 2) is that conflicts should be *symmetric* with respect to senders and receivers.

The formulation should also *respect pairwise incompatibility*: if two links cannot coexist in a feasible solution, they should be adjacent in the graph. In the SINR model, this means that the distance between the links must be at least a fixed constant factor of the length of the shorter link (see Thm. 4). This implies that distances involving one endpoint of the shorter link are within a constant factor of distances involving the other endpoint. Further, the distances from the short link to the further endpoint of the longer link are within a constant factor of the length of the longer link. Thus, we are left with only two parameters, up to constant factors: length l_{max} of the longer link, and the (minimum) distance $d(i,j)$ between the links, both scaled by the length of the shorter link l_{min}. We prove in the full version of the paper that constant factors only affect the chromatic numbers by constant factors. Hence, all conflict graphs are essentially represented by a two-variable predicate.

Finally, a conflict formulation should be *monotonic* with increasing distance. The reasoning is that a conflict formulation should represent the degree of conflict between pairs of links, or their relative "nearness". Specifically, if two links conflict and their separation (i.e., one of the distances between endpoints on distinct links) decreases while the links stay of the same length, then the links still conflict. A monotonic two variable predicate can be represented as a comparison of one variable with some monotone predicate of the other variable. Therefore, modulo constant factors, we are left with the conflict graph formulations $\mathcal{G}_f$. Finally, the monotonicity property also requires that the function f be sub-linear (as it could otherwise hold that $l_{max} + \epsilon < f(l_{max} + \epsilon)$ even though $l_{max} < f(l_{max})$).

4. PROPERTIES OF CONFLICT GRAPHS

We explore two types of properties of conflict graphs. The first type is concerned with gaps between the chromatic numbers of conflict graphs, or the relative difference of the chromatic numbers of graphs $\mathcal{G}_f$ and $\mathcal{G}_{f'}$. We show that the introduction of f increases the chromatic number of $\mathcal{G}_\gamma$ by a rather small factor depending on f. This is a key result that will be used to derive the approximation factor in the main result of this paper.

In the second part, we consider algorithmic properties of graphs $\mathcal{G}_f$. We prove that graphs $\mathcal{G}_f$ are constant-simplicial. Thus, constant factor approximation algorithms for vertex coloring, weighted maximum independent set and several other $\mathcal{NP}$-hard problems follow. This allows us to algorithmically approximate feasibility with graphs.

4.1 Gaps Between Chromatic Numbers of Conflict Graphs

We start by showing that the difference between the chromatic numbers of $\mathcal{G}_\gamma$ and $\mathcal{G}_{\gamma f}$ is a factor of at most $O(f^*(\Delta))$, where f^* is the iterated f function. This result is obtained by proving that for any independent set S in $\mathcal{G}_\gamma(L)$, the graph $\mathcal{G}_{\gamma f}(S)$ is $O(f^*(\Delta))$-inductive. To this end, we want to show that, for any given link i in S, any set T of mutually γ-independent links in S_i^+ that are γf-neighbors of i is small, or $O(f^*(\Delta))$. We do so by showing that the progression of lengths of the links in T must be fast growing, or inversely with f; the number of links must therefore be bounded by the iterated f function.

Theorem 1. *For any set of links L, a constant $\gamma \geq 1$ and a non-decreasing strongly sub-linear function f, $\chi(\mathcal{G}_{\gamma f}(L)) = O(f^*(\Delta)) \cdot \chi(\mathcal{G}_\gamma(L))$.*

Proof. Let S be a γ-independent set in L. Consider any link $i \in S$ and let T denote the set of links in S_i^+ that are f-adjacent with i. We will show that $|T| = O(f^*(\Delta))$. Note that for each $j \in T$, $d(i,j) \leq \gamma l_i f(l_j/l_i)$. Let p_j denote the endpoint of link $j \in T$ closest to link i. We split T into two subsets T_1 and T_2, where

$$T_1 = \{j \in T : d(p_j, r_i) \leq \gamma l_i f(l_j/l_i)\},$$
$$T_2 = T \setminus T_1 \subseteq \{j \in T : d(p_j, s_i) \leq \gamma l_i f(l_j/l_i)\}.$$

Let us first consider T_1. Let $j,k \in T_1$ be two links with $l_k < l_j$. Then,

$$d(p_j, r_i) \leq \gamma l_i f(l_j/l_i), \qquad \text{(because } j,k \in T_1) \qquad (1)$$
$$d(p_k, r_i) \leq \gamma l_i f(l_k/l_i), \qquad (2)$$
$$d(p_j, p_k) \geq d(j,k) > \gamma l_k, \quad (j,k \text{ are } \gamma\text{-independent}) \qquad (3)$$

By plugging the inequalities above into the triangle inequality $d(p_j,p_k) \leq d(p_j,r_i) + d(r_i,p_k)$, we obtain

$$\gamma l_k < \gamma l_i f(l_j/l_i) + \gamma l_i f(l_k/l_i) \leq 2\gamma l_i f(l_j/l_i),$$

where the last inequality follows from the assumption that $l_k < l_j$ and that f is a non-decreasing function. Thus,

$$l_k/l_i < 2f(l_j/l_i). \qquad (4)$$

Denote $g(x) \equiv 2f(x)$. Note that $g(x)$ is strongly sub-linear; hence, there exists $x_0 = \inf\{x \geq 1, g(x) < x\} + 1$. Let $1, 2, \ldots, t = |T_1|$ be the arrangement of the links in T_1 in increasing order by length and let $\lambda_j = \frac{l_j}{l_i}$ for $j = 1, 2, \ldots, t$. Let h be the link with the smallest index such that $\lambda_h \geq x_0$. We will bound the number of links in $A = \{1, 2, \ldots, h-1\}$ and $B = \{h, h+1, \ldots, t\}$ separately. $|A|$ can be bounded by a simple application of the doubling property of the space. Note that for all $j \in A$, $f(l_j/l_i) \leq f(x_0) = O(1)$ because $l_j/l_i < x_0$. Thus, the system of inequalities (1-3) implies that the mutual distance between different points p_j with $j \in A$ is at least γl_i, while their distance from r_i is at most $\gamma f(x_0) l_i$; hence, we have that $|A| = O(f(x_0)^m) = O(1)$.

Now let us bound $|B|$, using (4). We have that

$$x_0 \leq \lambda_h < g(\lambda_{h+1}) \leq g(g(\lambda_{h+2})) \leq \cdots \leq g^{(t-h)}(\lambda_t),$$

which implies that $t - h \leq g^*(\lambda_t) = O(f^*(\Delta))$. Recall that $h = |A| = O(1)$; hence, $t = O(f^*(\Delta))$.

This completes the proof that $|T_1| = O(f^*(\Delta))$. The set T_2 is handled similarly. In this case, for any pair of links $j,k \in T_2$ with $l_j > l_k$, the system of inequalities (1-3) holds with r_i replaced with s_i; the rest of the argument is identical. These results imply the theorem. □

4.2 Algorithmic Properties of Conflict Graphs

We prove that every conflict graph $\mathcal{G}_f$ with strongly sub-linear function f is constant-simplicial. This guarantees, among other properties, that the vertex coloring and maximum weighted independent set problems in these graphs can be efficiently approximated within constant factors [1, 26, 48].

We give a simple argument that holds in the plane, where it holds for essentially all sub-linear functions. It is based on splitting the plane into 60° sectors emanating from a given node of a link, and arguing that all adjacent longer links within a sector must form a clique. With a more detailed argument, the result can be extended to general doubling metrics, but requires then strong sub-linearity.

PROPOSITION 1. *Let f be a non-decreasing function such that $f(x)/x$ is non-increasing and let L be a set of links in the plane. Then $\mathcal{G}_f(L)$ is 12-simplicial.*

5. DEFINITIONS: SINR, FEASIBILITY

SINR Model and Feasibility. A *power assignment* for a set L of links is a function $P : L \to \mathbb{R}_+$. For each link i, $P(i)$ defines the power level used by the sender node s_i. In the *physical model (or SINR model)* of communication [40], when using a power assignment P, a transmission of a link i is successful if and only if

$$\frac{P(i)}{l_i^\alpha} \geq \beta \cdot \left(\sum_{j \in S \setminus \{i\}} \frac{P(j)}{d_{ji}^\alpha} + N \right), \qquad (5)$$

where N is a constant denoting the ambient noise, β denotes the minimum SINR (Signal to Interference and Noise Ratio) required for a message to be successfully received, $\alpha \in (2, 6)$ is the path loss constant and S is the set of links transmitting concurrently with link i. Here the left side of the inequality is interpreted as the received signal power of link i and the sum on the right side is interpreted as the interference on link i caused by concurrently transmitting links.

A set S of links is called *P-feasible* if the condition (5) holds for each link $i \in S$ when using power P. We say S is *feasible* if there exists a power assignment P for which S is P-feasible. Similarly, a collection of sets is P-feasible/feasible if each set in the collection is. Note that we do not assume limits on the available power, which means that the noise term can be ignored. The case of a maximum power limit requires primarily that the links that are close to maximum length be handled separately using the maximum power available [29], something that remains to be studied.

The Influence Operator and a Sufficient Condition for Feasibility. The *influence operator* I is defined as follows. For links i, j, let $I(i,j) = \frac{l_i^\alpha}{d(i,j)^\alpha}$ and define $I(i,i) = 0$ for simplicity of notation. The operator I is additively expanded: for a set S of links and a link i, let $I(S,i) = \sum_{j \in S} I(j,i)$ and $I(i,S) = \sum_{j \in S} I(i,j)$. We will use the notation $I(L) = \max_{i \in L} I(L_i^-, i)$.

In order to identify feasible sets, we will use the following sufficient condition for feasibility.

THEOREM 2. *[28] For any set of links L in a metric space, if $I(L) < \frac{1}{2 \cdot 3^\alpha (4\beta + 2)}$, then L is feasible.*

Sensitivity of Feasible Sets. A set of links is called p-P-feasible if it is P-feasible with the parameter β replaced with number p. The following sensitivity argument has proved useful. It shows, in particular, that constant factor changes to the threshold parameter β do not affect asymptotic results by more than a constant factor.

THEOREM 3. *[18] Let p, p' be positive values, P be a power assignment, and L be a p-P-feasible set. Then L can be partitioned into $\lceil 2p'/p \rceil$ sets each of which is p'-P-feasible.*

Fading Metrics. *Fading metrics* are doubling metrics with doubling dimension $m < \alpha$. We shall assume, without stating so explicitly, that the links are located in a fading metric.

6. CAPTURING FEASIBILITY WITH CONFLICT GRAPHS

We show that for appropriate constant $\gamma > 0$ and function f, SINR-feasibility is "trapped" between graph representations $\mathcal{G}_\gamma$ and $\mathcal{G}_f$; namely, each feasible set is an independent set in $\mathcal{G}_\gamma(L)$ and each independent set in $\mathcal{G}_f(L)$ is feasible. In particular, this holds for $f(x) = \gamma' \widehat{\log}(x)$ for an appropriate constant $\gamma' > 0$, where the function $\widehat{\log}(x)$ is defined for $x \geq 1$ by $\widehat{\log}(x) = \max(\log^{2/(\alpha-m)}(x), 1)$. The gap between these approximations is quantified using our results in Sec. 4.1, ultimately leading to $O(\log^* \Delta)$ approximation for scheduling problems.

6.1 Independence of Feasible Sets

The theorem below is based on the simple observation that two links in the same "highly feasible" set must be spatially separated by at least a multiple of the length of the shorter link, implying γ-independence for some $\gamma > 0$. The constant γ may then be adapted using Thm. 3, i.e. a feasible set can be split into a constant number of γ'-independent sets for any constant $\gamma' > 0$.

THEOREM 4. *For any constant $\gamma > 0$, a $(\gamma+1)^\alpha$-feasible set is γ-independent. In particular, if $\beta > 1$ then each feasible set is $(\beta^{1/\alpha} - 1)$-independent.*

PROOF. It suffices to show that two links in the same $(\gamma+1)^\alpha$-feasible set must be γ-independent. Let i, j be such links. Since i, j are in the same $(\gamma+1)^\alpha$-feasible set, the SINR condition implies that there is a power assignment P such that:

$P(i)/l_i^\alpha > (\gamma+1)^\alpha P(j)/d_{ji}^\alpha$ and $P(j)/l_j^\alpha > (\gamma+1)^\alpha P(i)/d_{ij}^\alpha$.

By multiplying together the inequalities above, canceling $P(i)$ and $P(j)$ and raising to the power of $1/\alpha$, we obtain:

$$d_{ij} d_{ji} > (\gamma+1)^2 l_i l_j. \qquad (6)$$

Let us show first that $\min\{d_{ij}, d_{ji}\} > \gamma \min\{l_i, l_j\}$. Indeed, if the opposite was true, e.g. if $d_{ij} \leq \gamma \min\{l_i, l_j\}$, the the triangle inequality would imply that $d_{ji} \leq d_{ij} + l_i + l_j \leq (\gamma+2)\max\{l_i, l_j\}$, which would contradict to (6): $d_{ij} d_{ji} \leq \gamma(\gamma+2) l_i l_j \leq (\gamma+1)^2 l_i l_j$.

Now consider $d(s_i, s_j)$. Let us assume, for contradiction, that e.g. $d(s_i, s_j) \leq \gamma l_i \leq \gamma l_j$. Then the triangle inequality would imply $d_{ji} \leq d(s_i, s_j) + l_i \leq (\gamma+1) l_i$ and $d_{ij} \leq d(s_i, s_j) + l_j \leq (\gamma+1) l_j$, which would again yield a contradiction to (6). We prove in the same manner that $d(r_i, r_j) > \gamma \min\{l_i, l_j\}$ and conclude that $d(i,j) > \gamma \min\{l_i, l_j\}$, i.e., i and j are γ-independent. □

6.2 Feasibility of Independent Sets

Here we show that for a large enough constant $\gamma > 0$, $\gamma\widehat{\log}$-independence implies feasibility. In particular, we show that if a set S is $\gamma\widehat{\log}$-independent then $I(S) = O(\gamma^{m-\alpha})$. Since we assumed that $m < \alpha$, an appropriate choice of γ yields feasibility via Thm. 2.

The argument consists of the following stages. For any given link $i \in S$, we first split S_i^- into length classes, or *equilength subsets*, where each equilength subset contains links differing by at most a factor of 2 in length. We bound the influence on link i for each of those subsets separately, and then combine those bounds using the additivity of the influence operator I.

For each equilength subset S the following common technique is applied: partition the plane into concentric annuli around the link i, count the number of links in each annulus and bound $I(S_i^-, i)$ based on these numbers and the fact that the links within the same annulus have almost the same influence on link i (because they are at roughly the same distance from i and have roughly similar lengths). The number of links in each annulus can be bounded using the doubling property of the space and independence of the links. The influence bound obtained for each subset S is $O((\gamma\widehat{\log}(l_i/\ell))^{m-\alpha})$, where ℓ is the longest link length in S. The function $\widehat{\log}$ is chosen so that combining those bounds in a sum results in an upper bound of $I(S_i^-, i) = O(\gamma^{m-\alpha})$.

The following lemma bounds the influence of an equilength 1-independent set S on a longer link i that is f-independent from the set S. This will be the main building block to be used for showing that $\gamma\widehat{\log}$-independent sets are feasible. The proof uses the annuli argument mentioned above.

LEMMA 1. *Let f be a non-decreasing function, s.t. $f(x) \geq 1$ whenever $x \geq 1$. Let S be an equilength 1-independent set of links, and let i be a link s.t. for each $j \in S$, $l_i \geq l_j$ and i and j are f-independent. Then $I(S, i) = O\left((f(l_i/\ell))^{m-\alpha}\right)$, where ℓ denotes the longest link length in S.*

Having a bound for the influence of each equilength set, we can now split the whole set into equilength subsets (length classes), bound the influence of each equilength subset using Lemma 1 and combine them into a series that converges when we choose $f(x) = \gamma\widehat{\log}(x)$.

THEOREM 5. *Let L be a $\gamma\widehat{\log}$-independent set with $\gamma \geq 1$. Then $I(L) = O\left(\gamma^{m-\alpha}\right)$.*

PROOF. Let us fix an arbitrary link $i \in L$. We have for each $j \in L_i^-$, $d(i, j) > \gamma l_j \widehat{\log}(l_i/l_j)$ because of $\gamma\widehat{\log}$-independence and that $l_i \geq l_j$. Let ℓ_0 denote the minimum link length in L_i^-. We partition L_i^- into at most $\lceil \log l_i/\ell_0 \rceil$ equilength subsets $L_1, L_2, \dots$ as follows:

$$L_t = \{j \in L_i^- : 2^{t-1}\ell_0 \leq l_j < 2^t \ell_0\},$$

for $t = 1, 2, \dots$. Let ℓ_t be the longest link length in L_t. The conditions of Lemma 1 hold for each L_t: it is an equilength 1-independent set ($\gamma\widehat{\log}$-independence implies 1-independence for $\gamma \geq 1$) and is f-independent from link i, with $f = \gamma\widehat{\log}$. Note also that $f(x) \geq 1$ when $x \geq 1$. Applying the lemma, we obtain

$$I(L_t, i) = O\left((\gamma\widehat{\log}(l_i/\ell_t))^{m-\alpha}\right).$$

Let d denote the largest index t for which L_t is not empty. By the definition of function $\widehat{\log}$, we have that $\widehat{\log}(l_i/\ell_d) \geq 1$ and for each $t < d$, $\widehat{\log}\frac{l_i}{\ell_t} = \log^{2/(\alpha-m)}\frac{l_i}{\ell_t} \geq (d-t)^{2/(\alpha-m)}$. Thus,

$$\begin{aligned} I(L_i^-, i) &= \sum_{t=1}^{d} I(L_t, i) \\ &\leq c\gamma^{m-\alpha}\left(1 + \sum_{t=1}^{d}\left((d-t)^{2/(\alpha-m)}\right)^{m-\alpha}\right) \\ &= O(\gamma^{m-\alpha}), \end{aligned}$$

where c is a constant. Since this holds for arbitrary $i \in L$, we have that $I(L) = O(\gamma^{m-\alpha})$. $\square$

Since the theorem above holds for any $\gamma \geq 1$, we obtain the desired result.

COROLLARY 1. *There is a constant $\gamma \geq 1$ such that each $\gamma\widehat{\log}$-independent set is feasible.*

7. IMPLICATIONS

7.1 Approximation Algorithms

Using our method of capturing feasibility with graphs, we approximate Scheduling and WCapacity problems within a factor of $O(\log^* \Delta)$. Let us first formally define the problems and related terms.

A *schedule* for a set L of links is a partition of L into feasible subsets (or *slots*). The *length* of the schedule is its number of slots. The Scheduling *problem* is to find a minimum length schedule for a given set L. The length of an optimal schedule for L is denoted $OPTS(L)$.

The WCapacity problem is the generalized dual of Scheduling, where given a set L of links with weights $\omega : L \to \mathbb{R}^+$, the goal is to find a feasible subset $S \subseteq L$ of maximum weight $\sum_{i \in S} \omega(i)$.

THEOREM 6. *There are polynomial $O(\log^* \Delta)$-approximation algorithms for Scheduling and WCapacity. The approximation is obtained by coloring the graph $\mathcal{G}_{\gamma\widehat{\log}}$ (for an appropriate constant $\gamma \geq 1$) in the case of Scheduling and by approximating its maximum weighted independent set in the case of WCapacity.*

PROOF. We present the proof for Scheduling approximation. The WCapacity problem is handled in a similar manner.

Let L be an input to Scheduling. We construct and color the graph $\mathcal{G}_{\gamma\widehat{\log}}(L)$ with constant γ chosen as in Corollary 1. By Corollary 1, such a coloring corresponds to a feasible schedule.

To derive the approximation factor, observe on one hand that in view of Thms. 3 and 4, any schedule of L can be refined into a coloring of $\mathcal{G}_\gamma(L)$ with only constant factor increase in the number of slots. Thus,

$$OPTS(L) = \Omega(\chi(\mathcal{G}_\gamma(L))).$$

On the other hand, by Thm. 1,

$$\chi(\mathcal{G}_{\gamma\widehat{\log}}(L)) = O(\widehat{\log}^*(\Delta)) \cdot \chi(\mathcal{G}_\gamma(L)) = O(\log^* \Delta) \cdot OPTS(L).$$

It is readily verified that the function $\gamma\widehat{\log}$ is strongly sublinear, implying, via , that $\mathcal{G}_{\gamma\widehat{\log}}(L)$ is constant-simplicial and thus colorable within constant approximation factor. $\square$

7.2 Measure of Interference

While approximation algorithms give bounds relative to an optimal value, it is often advantageous to have bounds in terms of some intrinsic parameters or more easily computable properties. Thus the interest in bounding chromatic numbers of graphs in terms of clique numbers, broadcast algorithms in terms of network diameter, and routing time in terms of "congestion + dilation". Our results also imply bounds for the optimum schedule length that can be efficiently computed from the network topology. Previous such results involved logarithmic factors in n and/or Δ [9, 31].

Let G be a k-simplicial graph and let $v_1, v_2, \ldots, v_n$ be a k-simplicial elimination order of vertices, which for our conflict graphs is by increasing link length. A k-approximate coloring of G is obtained by coloring the vertices greedily in reverse order. The number of colors used is at most the maximum post-degree plus 1, or $\max_i\{|N(v_i)\cap\{v_{i+1},\ldots,v_n\}|\}+1 \leq k\cdot\chi(G)+1$. We therefore define

$$B_f(L) = \max_{i\in L} |\{j \in L : l_j \geq l_i, d(i,j) \leq l_i f(l_j/l_i)\}|,$$

for a function f, and observe that $\chi(\mathcal{G}_f(L)) = \Theta(B_f(L))$. The results of Sec. 6 and 8 then imply the following theorem.

THEOREM 7. *There are constants $a, b > 0$ and $\gamma \geq 1$, such that for any set L, $a\cdot B_\gamma(L) \leq OPTS(L) \leq b\cdot B_{\widehat{\log}}(L)$, where $\frac{B_{\widehat{\log}}(L)}{B_\gamma(L)} = O(\log^* \Delta(L))$. Moreover, there are infinitely many instances L' and L'' s.t. $\frac{OPTS(L')}{B_\gamma(L')} = \Omega(\log^*\Delta(L'))$ and $\frac{B_{\widehat{\log}}(L'')}{OPTS(L'')} = \Omega(\log^*\Delta(L''))$.*

7.3 A Necessary and Sufficient Condition for Feasibility

Another interesting implication of Thm. 5 is the following result that shows that the sufficient condition for feasibility stated in Thm. 2 is essentially necessary in doubling metric spaces. This result is of independent interest, as it may prove useful for improved analysis of various problems. It should be noted that this theorem does not hold in general metric spaces.

The proof consists of two parts, bounding the influence on a link i by faraway links (i.e., links that are highly independent from link i) on one hand using Thm. 5, and by near links (the rest) on the other hand, using simple manipulations of the SINR condition.

THEOREM 8. *Let L be a 3^α-feasible set of links. Then, $I(L) = O(1)$.*

PROOF. Let us fix a link $i \in L$ and denote $S = L_i^-$. We split S into two subsets S_1 and S_2, where for each link $j \in S_1$, j and i are f-independent with $f(x) = 2x$, and $S_2 = S \setminus S_1$.

Note that 3^α-feasibility implies 2-independence of S_1, by Thm. 4. The bound $I(S_1, i) = O(1)$ then follows by applying an analogue of Thm. 5 with $\gamma = 1$ and with f-independence instead of $\widehat{\log}$-independence, which can be done because $\widehat{\log}(x) = O(f(x))$.

It remains to show that $I(S_2, i) = O(1)$. Let P be a power assignment for which L is P-feasible. Then, the SINR condition gives us the following inequalities:

$$\frac{P(i)}{l_i^\alpha} > 3^\alpha \sum_{j\in S_2} \frac{P(j)}{d_{ji}^\alpha}, \text{ and } \frac{P(j)}{l_j^\alpha} > 3^\alpha \frac{P(i)}{d_{ij}^\alpha} \text{ for all } j \in S_2.$$

By replacing $P(j)$ with $3^\alpha \frac{P(i)l_j^\alpha}{d_{ij}^\alpha}$ in the first inequality and simplifying the expression, we get:

$$\sum_{j\in S_2} \frac{l_i^\alpha l_j^\alpha}{d_{ij}^\alpha d_{ji}^\alpha} \leq 9^{-\alpha}. \tag{7}$$

In order to extract a bound on $I(S_2, i)$ from (7), we will show that one of the values $d_{ij}^\alpha, d_{ji}^\alpha$ in the denominator can be canceled out with l_i^α in the numerator and the other one can be replaced with $d(i,j)^\alpha$ by only introducing additional constant factors in the expression. Such a modification will transform the left side of (7) into $I(S_2, i)$.

Let us assume w.l.o.g. that $d_{ij} \geq d_{ji}$. Recall that for each $j \in S_2$, $d(i,j) \leq 2l_i$ by definition of S_2. Using the triangle inequality, we obtain $d_{ij} \leq d(i,j) + l_i + l_j \leq 4l_i$. On the other hand, as it was mentioned above, the set S_2 is 2-independent, which implies that $d_{ji} \geq d(i,j) > 2l_j$. Using the triangle inequality again, we obtain:

$$d(s_i, s_j) \geq d_{ij} - l_j \geq d_{ji} - l_j > d_{ji}/2,$$
$$d(r_i, r_j) \geq d_{ji} - l_j > d_{ji}/2.$$

Thus, $d(i,j) > d_{ji}/2$. By replacing d_{ij} with $4l_i$ and d_{ji} with $2d(i,j)$ in the left-hand part of (7), we obtain the desired bound: $I(S_2, i) \leq (8/9)^\alpha$. Since this holds for an arbitrary $i \in L$, we get that $I(L) = O(1)$. □

Remark. Note that Thm. 3 implies that any feasible set can be refined into a constant number of 3^α-feasible subsets. Thus, the influence function fully captures feasibility in fading metrics, modulo constant factors.

8. LIMITATIONS OF THE GRAPH-BASED APPROACH

We have found that conflict graphs can achieve a remarkably good, yet super-constant, approximation for scheduling problems in doubling metrics. We examine in this section how far this approach can be pushed, obtaining essentially tight bounds. We treat these issues in terms of the Scheduling problem.

In the first part of the section, we expose the limitations of the graph method in Euclidean spaces. We show, in particular, that conflict graphs do not yield any non-trivial approximation to the Scheduling problem in terms of the number of links n. In particular, they cannot lead to constant factor approximation. We also consider approximation limits in terms of the parameter Δ, and show that for all reasonable functions f, the approximation factor is at least $\Omega(\log^*\Delta)$. Thus, the approximation factor we obtained cannot be improved within a conflict-graph framework. Note that the instances we construct are embedded on the real line, i.e., in one dimensional space.

In the second part, we find that the graph method cannot provide any non-trivial approximation guarantees in general metric spaces, neither in terms of n nor Δ.

8.1 Euclidean Spaces

In the following theorem, we construct, for any function $f = \omega(1)$, a feasible set of f-adjacent links. The construction is based on the following observations. On the one hand, it follows from Thm. 2 that any set of exponentially growing links arranged sequentially by the order of length on the real line is (almost) feasible. On the other hand, given such a set

S of links on the line, a new link j can be formed so that j is f-adjacent to all the links in S while the set $S \cup j$ stays feasible; the only requirement is that j be long enough. Our construction then builds recursively on these ideas.

THEOREM 9. *Let $f(x) = \omega(1)$. For any integer $n > 0$, there is a* feasible *set L of n links arranged on the real line, such that $\mathcal{G}_f(L)$ is a clique, i.e., $\chi(\mathcal{G}_f(L)) = n$. Moreover, if $f(x) \geq g(x)$ $(x \geq 1)$ for a strongly sub-linear increasing function $g(x)$ with $g(x) = \omega(1)$, then $n = \Omega(g^*(\Delta))$.*

PROOF. Consider a set of cn links $\{1, 2, \ldots, cn\}$ arranged sequentially from left to right on the real line, where $c > 0$ is a constant to be chosen later. Each link i is directed from left to right and for each $i = 1, 2, 3, \ldots, n-1$, the nodes s_{i+1} and r_i share the same location on the line, i.e., $r_i = s_i + l_i = s_{i+1}$. See Figure 1. The lengths of links are defined inductively, as follows. We set $l_1 = 1$, and for $i \geq 1$, we choose l_{i+1} to be the minimum value satisfying:

$$l_{i+1} \geq 2l_i \tag{8}$$

$$2d(i+1, j) = 2d_{i+1,j} \leq l_j f(l_{i+1}/l_j) \text{ for all } j \leq i. \tag{9}$$

Such a value of l_{i+1} can be chosen as follows. By the inductive hypothesis, we have $l_j \geq 2l_{j-1}$ for $j = 2, 3, \ldots, i$. This implies that $l_i \geq \sum_{j=1}^{i-1} l_j$. Then, we have that $d_{i+1,j} = \sum_{t=j+1}^{i} l_t \leq 2l_i$ for $j = 1, 2, \ldots, i$. Thus, it is enough to choose l_{i+1} so that $l_{i+1} \geq 2l_i$ and $4l_i \leq l_j f(l_{i+1}/l_j)$, which can be done using $f = \omega(1)$ and the fact that the values of l_j for $j = 1, 2, \ldots, i$ are already fixed at this point. This completes the construction. First note that (9) implies that $\mathcal{G}_f(L)$ is a clique. It remains to argue feasibility. Consider the odd numbered links $S = \{1, 3, \ldots, \ldots, 2t+1\}$. Let us fix a link $2k+1 \in S$. Note that for each $j \in S^-_{2k+1}$, $d(j, 2k+1) \geq l_{2k}$. We have that

$$\begin{aligned} I(S^-_{2k+1}, 2k+1) &= \sum_{j \in S^-_{2k+1}} \frac{l_j^\alpha}{d(j, 2k+1)^\alpha} \\ &\leq \sum_{j \in S^-_{2k+1}} \left(\frac{l_j}{l_{2k}}\right)^\alpha \\ &\leq \sum_{j \in S^-_{2k+1}} \frac{l_j}{l_{2k}} \leq 1, \end{aligned}$$

where the second inequality holds because $l_j/l_{2k} \leq 1$ and the last inequality follows from (8). Thus, we can extract a constant fraction S' of S that is feasible, using Thm. 3. With the right choice of the constant c in the beginning of the proof we have that $|S'| = n$. This proves the first part of the theorem.

Now let us assume that $f(x) \geq g(x)$ for a strongly sublinear function $g(x)$ with $g(x) = \omega(1)$. Then, there is a constant x_0 such that $g(x) < x$ for all $x \geq x_0$ (because $g(x) = o(x)$) and there is a constant c such that $2g(x)/x \leq g(y)/y$ whenever $x \geq cy$ (strong sub-linearity). In this case we repeat the construction above with slight modifications.

We set $l_1 = 1$ and set $l_{i+1} > \max\{c, x_0\}$ be the minimum value s.t. $g(l_{i+1}) \geq 2l_i$, for $i = 1, 2, \ldots$ (such a value exists because $g(x) = \omega(1)$). Let us show that the conditions (8-9) hold with these lengths.

Since $l_{i+1} \geq x_0$, we have that $l_{i+1} > g(l_{i+1}) \geq 2l_i$, which implies (8). This in turn implies, as observed in the first part of the proof, that $d(i+1, j) < 2l_i$ for all $2 \leq j \leq i$. Let us denote $x = l_{i+1}/l_1 = l_{i+1}$ and $y = l_{i+1}/l_j$. Note that $x/y = l_j \geq c$, so we have, by strong sub-linearity of g, that $g(y)/y \geq 2g(x)/x$, or equivalently, that $l_j \cdot g(l_{i+1}/l_j) \geq 2 \cdot g(l_{i+1})$; hence $l_j \cdot g(l_{i+1}/l_j) \geq 4l_i > 2d(i+1, j)$ for all $2 \leq j \leq i$, which means that (9) also holds.

It remains to prove the lower bound for n. Recall that the value of l_{i+1} is the minimum satisfying $g(l_{i+1}) \geq 2l_i$ for $i = 1, 2, \ldots, n-1$. Then, we have $g(l_{i+1}/2) < 2l_i$ or, equivalently, $h(l_{i+1}/2) < l_i/2$, where $h(x) = g(x)/4$. Thus,

$$\begin{aligned} 1/2 = l_1/2 > h(l_2/2) > h(h(l_3/2)) > \ldots > h^{(n-1)}(l_n/2) \\ = h^{(n-1)}(\Delta/2), \end{aligned}$$

which implies that $n = \Omega(h^*(\Delta/2)) = \Omega(g^*(\Delta))$. □

COROLLARY 2. *In terms of the number of links n, the approximation factor for Scheduling when using $\mathcal{G}_f$ with any $f = \omega(1)$ is no better than n.*

By choosing $g(x) = \gamma\widehat{\log}(x)$ in Thm. 9, we obtain that the approximation factor of $O(\log^* \Delta)$ cannot be improved for $\mathcal{G}_{\gamma\widehat{\log}}$.

COROLLARY 3. *Let $f(x) = \Omega(\log^{(c)} x)$ for a constant c. Then, for each $\Delta > 0$, there is a* feasible *set of links L with $\Delta(L) = \Omega(\Delta)$, such that $\mathcal{G}_f(L)$ is a clique of size $\Theta(\log^* \Delta(L))$.*

While the theorem above shows that graphs $\mathcal{G}_f$ with $f = \Omega(\log^{(c)} x)$ for some constant c require too much separation, the theorem below shows that graphs $\mathcal{G}_f$ with $f = O(\log^{1/\alpha} x)$ provide insufficient separation, leading, perhaps surprisingly, to a similar sized gap of $\log^* \Delta$. Namely, we have $\chi(\mathcal{G}_f(L')) = \frac{OPTS(L')}{\Omega(\log^* \Delta)}$ for certain instances L'. The construction follows the general structure of Thm. 7 in [20] of a lower bound for scheduling the edges of a minimum spanning tree of a set of points in the plane. There are two technical challenges to overcome, in order to implement this structure in our setting. First, the construction of [20] is not f-independent. Second, even when ignoring the f-independence requirement, the lower bound for the scheduling number obtained in [20] is only $\Omega(\log\log^* \Delta)$.

THEOREM 10. *Let $f(x) = O(\log^{1/\alpha} x)$. For each $\Delta > 0$, there is an f-independent set of links L on the real line with $\Delta(L) = \Omega(\Delta)$ that cannot be scheduled in fewer than $\Theta(\log^* \Delta(L))$ slots.*

We describe the idea of the construction informally. The construction is inductive, starting from a trivial instance L_1 containing a single link. For $t \geq 1$, assume there is an instance L_t having the desired properties, i.e., L_t is f-independent and with $OPTS(L) \geq t$. In order to construct the instance L_{t+1}, consider a single link j that is longer than the links in L_t and place it at distance d from L_t so that all the links in L_t are f-independent from j. Let I_0 denote the minimum influence of a link from L_t on link j. Now, take k identical copies of L_t and place them at a distance d from j. This will of course violate the independence between different instances, which we will address shortly, but they will still be independent from link j. The idea is that if the number of copies k is large enough, then for any set S containing at least one link from each copy, we will have

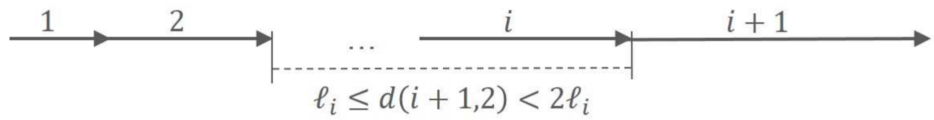

Figure 1: The construction in Thm. 9.

$I(S, j) = kI_0 > c_0$, where c_0 is a constant large enough to ensure that $S \cup \{j\}$ is infeasible (based on Thm. 8). This will mean that any schedule of the link j and the k copies must place at least *one whole copy* of L_t in slots separate from j. Since it takes at least t slots to schedule one copy of L_t, it takes at least $t + 1$ slots to schedule all the copies together with link j.

It remains to address the issue of f-independence between different copies. Note that because of the scale-invariance of the influence operator, we can scale a copy of L_t by a factor s and place it further than before, at a distance $s \cdot d$ from link j and still have the minimum influence of I_0 on j. However, in order for this influence to be taken into the account, the link j must still be longer than the links in the scaled instance. Using this trick, we can scale different copies by different factors and guarantee their mutual independence, while preserving the properties we had in the case of identical copies. Since the link lengths must grow exponentially at each step t, the number t of slots required will be small compared to the number of links and the parameter Δ, but will still be $\Omega(\log^* \Delta)$.

8.2 General Metric Spaces

The following theorem shows that conflict graphs can be arbitrarily far from schedules in general metric spaces. Given a function f, the construction consists of an f-independent set of *unit length* links. Since all links have length 1, f-independence is equivalent to $f(1)$-independence. The separation between the links is just enough to ensure $f(1)$-independence. However, since all the links are equally ($f(1)$-) separated from any given link, their interference accumulates and only a constant number of links can be scheduled in the same slot. This leads to schedules of length $\Theta(n)$.

PROPOSITION 2. *For each function f and any $n \geq 1$, there is an f-independent set of n unit length links (hence, $\Delta = 1$) that cannot be scheduled into less than $\Theta(n)$ slots.*

PROOF. Let $L = \{1, 2, \ldots, n\}$ be the set of links. We define the lengths and the distances between the links so as to ensure the metric constraints hold. For each link i we define $l_i = 1$. The distances are defined as follows:

$$d(s_i, s_j) = f(1) \cdot (l_i + l_j) = 2f(1),$$
$$d(s_i, r_j) = d(s_i, s_j) + l_j = 2f(1) + 1,$$
$$d(r_i, r_j) = d(s_i, s_j) + l_i + l_j = 2f(1) + 2.$$

It is straightforward to check that such distances define a metric. Moreover, the set L is f-independent, since $d(i, j) > f(1) \cdot l_i = l_i f(l_j / l_i)$. Let us consider any P-feasible subset S of k links for a power assignment P. Let us fix a link $i \in S$. The SINR condition implies: $P(i) > \beta \sum_{j \in S \setminus \{i\}} \frac{P(j) l_i^\alpha}{d_{ji}^\alpha}$ and $P(j) > \beta \frac{P(i) l_j^\alpha}{d_{ij}^\alpha}$ for all $j \in S \setminus \{i\}$. By replacing $P(j)$ with $\frac{\beta P(i) l_j^\alpha}{d_{ij}^\alpha}$ in the first inequality and canceling the term $P(i)$, we obtain:

$$1 > \beta^2 \sum_{j \in S \setminus \{i\}} \frac{l_i^\alpha l_j^\alpha}{d_{ij}^\alpha d_{ji}^\alpha} = \beta^2 \sum_{j \in S \setminus \{i\}} \frac{1}{(2f(1)+1)^2} = \frac{\beta^2(|S| - 1)}{(2f(1)+1)^2},$$

which implies that $|S| < \left(\frac{2f(1)+1}{\beta}\right)^2 + 1 = O(1)$. Since S was an arbitrary feasible subset of L, we conclude that L cannot be split into less than $\Theta(n)$ feasible subsets. $\square$

Acknowledgments

We thank Christian Konrad and Jaikumar Radhakrishnan for helpful suggestions.

9. REFERENCES

[1] K. Akcoglu, J. Aspnes, B. DasGupta, and M.-Y. Kao. Opportunity-cost algorithms for combinatorial auctions. In E. J. Kontoghiorghes, B. Rustem, and S. Siokos, editors, *Applied Optimization 74: Computational Methods in Decision-Making, Economics and Finance*, pages 455–479. Kluwer Academic Publishers, 2002.

[2] L. Barriére, P. Fraigniaud, L. Narayanan, and J. Opatrny. Robust position-based routing in wireless ad hoc networks with irregular transmission ranges. *Wireless Communications and Mobile Computing*, 3(2):141–153, 2003.

[3] M. Bodlaender and M. M. Halldórsson. Beyond geometry: Towards fully realistic wireless models. In *PODC*, 2014.

[4] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer latency minimization for wireless networks using SINR constraints. In *Mobihoc*, 2007.

[5] R. L. Cruz and A. Santhanam. Optimal Routing, Link Scheduling, and Power Control in Multi-hop Wireless Networks. In *INFOCOM*, 2003.

[6] J. Dams, M. Hoefer, and T. Kesselheim. Scheduling in wireless networks with Rayleigh-fading interference. In *SPAA*, pages 327–335, 2012.

[7] T. ElBatt and A. Ephremides. Joint Scheduling and Power Control for Wireless Ad-hoc Networks. In *INFOCOM*, 2002.

[8] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *PODC*, pages 220–229, 2009.

[9] A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. *Theoretical Computer Science*, 412(24):2657–2667, 2011.

[10] A. Fanghänel and B. Vöcking. Scheduling and power assignments in the physical model. In S. Nikoletseas and J. D. Rolim, editors, *Theoretical Aspects of Distributed Computing in Sensor Networks*,

Monographs in Theoretical Computer Science. An EATCS Series, pages 31–57. Springer, 2011.

[11] L. Fu, S. C. Liew, and J. Huang. Power controlled scheduling with consecutive transmission constraints: complexity analysis and algorithm design. In *INFOCOM*, pages 1530–1538. IEEE, 2009.

[12] O. Goussevskaia, M. M. Halldórsson, and R. Wattenhofer. Algorithms for wireless capacity. *IEEE/ACM Trans. Netw.*, 22(3):745–755, 2014.

[13] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *MobiHoc*, pages 100–109, 2007.

[14] H. Gudmundsdottir, E. I. Ásgeirsson, M. Bodlaender, J. T. Foley, M. M. Halldórsson, and Y. Vigfusson. Measurement based interference models for wireless scheduling algorithms. In *MSWiM*, 2014. arXiv:1401.1723.

[15] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.

[16] M. Halldórsson and C. Newport. Making wireless algorithm theory more useful: Five ideas from the 2013 Workshop on Realistic Models for Algorithms in Wireless Networks. *ACM SIGACT News*, 45(3):72–74, 2014.

[17] M. M. Halldórsson. Wireless scheduling with power control. *ACM Transactions on Algorithms*, 9(1):7, December 2012.

[18] M. M. Halldórsson and J. Bang-Jensen. A note on vertex coloring edge-weighted digraphs. Technical Report 29, Institute Mittag-Leffler, Preprints Graphs, Hypergraphs, and Computing, 2014.

[19] M. M. Halldórsson and P. Mitra. Wireless capacity and admission control in cognitive radio. In *INFOCOM*, pages 855–863, 2012.

[20] M. M. Halldórsson and P. Mitra. Wireless Connectivity and Capacity. In *SODA*, 2012.

[21] M. M. Halldórsson and R. Wattenhofer. Wireless communication is in APX. In *ICALP*, pages 525–536, 2009.

[22] J. Heinonen. *Lectures on Analysis on Metric Spaces.* Springer, 1 edition, 2000.

[23] M. Hoefer and T. Kesselheim. Secondary spectrum auctions for symmetric and submodular bidders. In *EC*, pages 657–671, 2012.

[24] M. Hoefer, T. Kesselheim, and B. Vöcking. Approximation algorithms for secondary spectrum auctions. In *SPAA*, pages 177–186, 2011.

[25] K. Jansen, C. Mathieu, H. Shachnai, and N. E. Young. Packing and scheduling algorithms for information and communication services. Report from Dagstuhl Seminar 11091, March 2011.

[26] F. Kammer and T. Tholey. Approximation algorithms for intersection graphs. *Algorithmica*, 68(2):312–336, 2014.

[27] B. Katz, M. Volker, and D. Wagner. Energy efficient scheduling with power control for wireless networks. In *WiOpt*, pages 160–169. IEEE, 2010.

[28] T. Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *SODA*, pages 1549–1559, 2011.

[29] T. Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *ESA*, pages 659–670, 2012.

[30] T. Kesselheim. Dynamic packet scheduling in wireless networks. In *PODC*, pages 281–290, 2012.

[31] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *DISC*, pages 163–178. Springer, 2010.

[32] F. Kuhn, N. Lynch, C. Newport, R. Oshman, and A. Richa. Broadcasting in unreliable radio networks. In *PODC*, pages 336–345. ACM, 2010.

[33] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-hoc networks beyond unit disk graphs. In *FOMC*, pages 69–78. ACM, 2003.

[34] H. Lin and F. Schalekamp. On the complexity of the minimum latency scheduling problem on the Euclidean plane. *arXiv preprint 1203.2725*, 2012.

[35] T. Locher, P. von Rickenbach, and R. Wattenhofer. Sensor networks continue to puzzle: Selected open problems. In *ICDCN*, pages 25–38. Springer, 2008.

[36] R. Maheshwari, S. Jain, and S. R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *SenSys*, pages 141–154, 2008.

[37] T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *INFOCOM*, pages 1–13, 2006.

[38] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *HotNets*, 2006.

[39] C. H. Papadimitriou. NP-completeness: A retrospective. In *ICALP*, pages 2–6. Springer, 1997.

[40] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 2 edition, 2002.

[41] J. Ryu, C. Joo, T. T. Kwon, N. B. Shroff, and Y. Choi. DSS: Distributed SINR-based scheduling algorithm for multihop wireless networks. *IEEE Trans. Mobile Computing*, 12(6):1120–1132, 2013.

[42] V. Sevani and B. Raman. SIR based interference modeling for wireless mesh networks: A detailed measurement study. In *COMSNETS*. IEEE, 2012.

[43] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *SenSys*, pages 237–250. ACM, 2006.

[44] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Contr.*, 37(12):1936–1948, 1992.

[45] T. Tonoyan. On some bounds on the optimum schedule length in the SINR model. In *ALGOSENSORS*, pages 120–131, 2012.

[46] T. Tonoyan. Conflict graphs and the capacity of the mean power scheme. *In ALGOSENSORS*, 2013.

[47] P. Wan, O. Frieder, X. Jia, F. F. Yao, X. Xu, and S. Tang. Wireless link scheduling under physical interference model. In *INFOCOM*, pages 838–845, 2011.

[48] Y. Ye and A. Borodin. Elimination graphs. *ACM Transactions on Algorithms*, 8(2):14:1–14:23, 2012.