Algorithmica
https://doi.org/10.1007/s00453-021-00866-z

Network Design under General Wireless Interference

Magnús M. Halldórsson[1] · Guy Kortsarz[2] · Pradipta Mitra[3] · Tigran Tonoyan[4]

Received: 13 June 2020 / Accepted: 3 August 2021

Abstract

We introduce the problem of finding a spanning tree along with a partition of the tree edges into the fewest number of feasible sets, where constraints on the edges define feasibility. The motivation comes from wireless networking, where we seek to model the irregularities seen in actual wireless environments. Not all node pairs may be able to communicate, even if geographically close—thus, the available pairs are specified with a link graph $\mathcal{G} = (V, E)$. Also, signal attenuation need not follow a nice geometric formula—hence, interference is modeled by a conflict (hyper)graph $\mathcal{C} = (E, F)$ on the links. The objective is to maximize the efficiency of the communication, or equivalently, to minimize the length of a schedule of the tree edges in the form of a coloring. We find that in spite of all this generality, the problem can be approximated linearly in terms of a versatile parameter, the inductive independence of the conflict graph. Specifically, we give a simple algorithm that attains a $O(\rho \log n)$-approximation, where n is the number of nodes and ρ is the inductive independence. For an extension to Steiner trees, modeling multicasting, we obtain a $O(\rho \log^2 n)$-approximation. We also consider a natural geometric setting when only links longer than a threshold can be unavailable, and analyze the performance of a geometric minimum spanning tree.

✉ Tigran Tonoyan
ttonoyan@gmail.com

Magnús M. Halldórsson
mmh@ru.is

Guy Kortsarz
guyk@camden.rutgers.edu

Pradipta Mitra
ppmitra@gmail.com

1 ICE-TCS, Department of Computer Science, Reykjavik University, Reykjavík, Iceland

2 Department of Computer Science, Rutgers University, Camden, NJ, USA

3 Google Research, New York, NY, USA

4 Department of Computer Science, Technion Israel Institute of Technology, Haifa, Israel

Published online: 10 August 2021

Keywords Network Design · Wireless interference · Connectivity · Scheduling

1 Introduction

We introduce the problem of finding a spanning tree along with a partition of the tree edges into the fewest number of feasible sets, which are independent sets in a given conflict (hyper)graph. The motivation comes from wireless networking, where we seek a basic communication structure while capturing the irregularities seen in actual wireless environments.

A spanning tree is the minimal structure for connecting the given set of nodes into a mutually communicable network. The *cost* of a communication spanning tree is the *time* required to *schedule* all the tree edges—the transmission *links*—while obeying the *interference* caused by simultaneous transmissions. The *scheduling complexity* of the tree represents its throughput capacity: how much communication can be sustained in the long run. The task might be to *aggregate* the data measured at the sensor nodes, or to *broadcast* using one-to-one communication to all nodes of the network.

Technically, the former premise means that the set of usable or *available or reliable links* is now given as a *link graph* $\mathcal{G} = (V, L)$ over the set V of nodes. We place no restrictions on the structure of this graph. The second premise implies another (hyper)graph, the *conflict graph* $\mathcal{C} = (L, W)$, this time on top of the *links*. In the Connectivity Scheduling problem, we seek a spanning tree T of $\mathcal{G}$ and a coloring of the links of T minimizing the number of colors used, where the conflict graph $\mathcal{C}$ specifies whether a given set of links in L can coexist in the same color class.

These formulations naturally raise a number of questions: Can arbitrary sets of available/usable links actually be handled effectively? Can we disconnect the conflicts/interference from the geometry? Since the ugly specter of intractability is bound to raise its head somewhere, what are minimal restrictions that keep these problems well-approximable?

Formal definitions The set of *available links* is given as the edges of a *link graph* $\mathcal{G} = (V, L)$ over the set V of nodes. We place no restrictions on the structure of this graph. The interference conflicts are captured by a *fractional conflict graph* $\mathcal{C} = (L, W)$ on top of the set L of communication links. Here $W : L \times L \rightarrow \mathbb{R}^+$ is a function on ordered pairs of links, where $W(e, f)$ represents (or approximates) the degree to which a transmission on link e interferes with a transmission on link f. For convenience, let $W(e, e) = 0$. Following Kesselheim [29], we define the asymmetric function $\hat{W}(e,f)$ on a pair of links by $\hat{W}(e,f) = W(e,f) + W(f, e)$, if $e \prec f$, and $\hat{W}(e,f) = 0$, otherwise. We shall write $W(S,f) = \sum_{e \in S} W(e,f)$ ($\hat{W}(S,f) = \sum_{e \in S} \hat{W}(e,f)$) and $W(e, S) = \sum_{f \in S} W(e,f)$ ($\hat{W}(e, S) = \sum_{f \in S} \hat{W}(e,f)$). Also, $\hat{W}(S, S') = \sum_{e \in S} \hat{W}(e, S')$.

A set S of links is an *independent* or a *feasible set* if $W(S, e) < 1$, for all $e \in S$. A *coloring* of $\mathcal{C} = (L, W)$ is a partition of L into independent sets. Observe that when W is a 0–1 function, we have the usual independent sets and colorings of graphs. Also, the formulation with fractional conflicts corresponds to hypergraphs that contain a hyperedge for each minimal set S' where $W(S', e) \geq 1$ holds for some $e \in S'$. Let $\mathcal{C}[Y] = (Y, W\restriction_Y)$ denote the subgraph induced by a given subset $Y \subseteq L$ of links.

We can now state our Connectivity Scheduling problem formally:

> Given a link graph $\mathcal{G} = (V, L)$ and a fractional conflict graph $\mathcal{C} = (L, W)$, we seek a spanning tree T of $\mathcal{G}$ and a coloring of $\mathcal{C}[T]$, using the fewest number of colors.

A fractional conflict graph $\mathcal{C} = (L, W)$ is *ρ-inductive independent*, w.r.t. an ordering $\prec$ of the links, if, for every link e and every feasible set I (not necessarily containing e), $\hat{W}(e, I) \leq \rho$. Here, "inductive" refers to how the interference is measured only towards later links, and "independence" that it is towards independent sets. We write $e \prec I$ to mean that e precedes each link in I. In geometric settings (including range-based and SINR models), $\prec$ corresponds to a non-decreasing ordering by link length.

Our results A key contribution of this work is the formulation of the Connectivity Scheduling problem that captures in a highly general way the irregularity and unreliability of wireless setting, while avoiding the inapproximability monster.

We show that Connectivity Scheduling can be approximated within a $O(\rho \log n)$ -factor, where ρ is the *inductive independence* of the (fractional) conflict graph. This is particularly relevant since ρ is known to be constant in both of the predominant interference models: the physical (or SINR) model, and the protocol model. This is attained by a simple greedy algorithm that can be viewed as a combination of Kruskal's MST algorithm and a link scheduling algorithm for the physical model. Interestingly, the result implies that the approximability of Connectivity Scheduling is not significantly affected by restricting the set of allowable links. We also generalize the problem to Steiner trees and obtain a $O(\rho \log^2 n)$-approximation by a novel reduction to the multi-dimensional Steiner tree problem.

In contrast, we find that in the geometric setting, where the links are embedded in the plane, the (perhaps more natural) approach of selecting and coloring an MST fails badly. However, we show that MST gives improved results in the natural special case where all short links are available.

Our results extend the state of the art for Connectivity Scheduling in the geometric SINR model in three ways: our algorithms give good approximations in the presence of unavailable links (in contrast with the known results), they work in general metrics (first such results, even with all links available), and we obtain similarly good approximations in the Steiner extension.

Related work The connectivity problem in the geometric SINR model was first considered by Moscibroda and Wattenhofer [37]. It was, in fact, the first work on worst-case analysis in the SINR model. They show that unlike in random networks, the worst-case connectivity depends crucially on the use of power control, and with optimal power control, $O(\log^4 n)$ colors suffice to connect the nodes. They soon improved this to $O(\log^2 n)$ [36, 38]. Currently, the best upper bounds known are $O(\log n)$ [17] and $O(\log^* \Lambda)$ [22], where Λ is the ratio between the longest to the shortest length of a link in a minimum spanning tree (MST), a structural parameter that is independent of n. Both of these results hold for the MST of the pointset; there are pointsets where $\Omega(\log^* \Lambda)$ colors are necessary for coloring an MST [22]. For a

Table 1 Results in the Euclidean SINR model

Setting	Approximation
No unreliable links	$O(\log n)$ [17], $O(\log^* \Lambda)$ [22]
With unreliable links	$O(\log n)$ [This work]
With unreliable links, Steiner variant	$O(\log n \log^* \Lambda)$ [This work]
Only long unreliable links	$O(\Pi + \min(\log^* \Lambda', \log n))$ [This work]

Λ (Λ') denotes the ratio of the longest and shortest link lengths in an MST (distances between the nodes, respectively). Π denotes the ratio of the longest reliable and shortest unreliable link lengths

comparison, we present our results for the Euclidean SINR setting together with the corresponding state of the art in Table 1 (see Sect. 5 for details).

The scheduling complexity of connectivity relates closely to the efficiency of *aggregation*, a key primitive for wireless sensor networks. We refer the reader to Incel and Ghosh [27] for bibliography on aggregation/collection problems.

Inductive independence was first defined by Akcoglu et al. [1] and studied by Ye and Borodin [40] in the graph setting, while the weighted version was introduced by Hoefer et al. [26]. It has been used as a performance measure for various throughput problems related to wireless networks, including admission control [11], dynamic packet scheduling [23, 31], and spectrum auctions [23, 25, 26]. On the other hand, it has not previously been applied to connected structures.

Outline of the paper We first discuss the modeling of irregularity and variability in Sect. 2. Our main technical results are given in Sect. 3. We give in Sect. 3.1 a greedy algorithm for Connectivity Scheduling achieving $O(\rho \log n)$-approximation, where ρ is the inductive independence number of the conflict graph. We also obtain a similar approximation of a *Steiner or multicast* version of the problem in Sect. 3.2. We examine, in Sect. 4, how the standard approach—finding a minimum spanning tree—fares for our problem in the geometric setting where the links are embedded in the plane. Implications of our results to the SINR (or physical) model are given in Sect. 5. The rest of the paper can safely be read without any background in that model. We then close with open problems. A brief primer on SINR concepts is given in Sect. 6 for completeness.

2 Modeling Irregularity

Algorithmic studies of wireless connectivity to date have generally involved strong "niceness" assumptions. One core assumption is that points are located in the Euclidean plane and *all (close enough) pairs of nodes are available as links* for use in the spanning tree. Interference modeling has become progressively more realistic, starting with range-based graph models to the fractional *SINR* model of interference, but the common thread is that *interference is a direct function of the geometry*.

The objective of this work is to embrace the irregularity in connectivity problems. We replace the previous assumptions by the opposite premises: *a link may not be usable even if it should be* and *Interference need not follow the underlying* *geometry*. The primary objective at the outset of this research was to extend results that hold in the SINR model to somehow capture irregularity. The SINR model is based on three axioms: (a) Signal decreases in strength as a fixed polynomial of the distance traveled, (b) interference is additive, and (c) message is received if the strength of the signal is sufficiently stronger than the total interference. We embrace the last two axioms, which have been verified experimentally (e.g., [39]), while aiming to modify the first one.

One approach would be to allow for variability of some restricted form. This could take the form of, say, bounded variation in pathloss (the decrease of the signal strength with distance); or, two different path loss constants (i.e., the exponent of the polynomial decay); or, the addition of specific forms of obstructions. Such modifications would however neither be very general nor particularly easy to reason about. One could easily grow a whole subtopic of "reasonable" extensions to SINR.

The other extreme is to treat the most general version, sometimes called *abstract SINR*, where signal strength and position are completely disassociated. However, this has the serious downside of intractability: computing a good solution to most basic problems becomes impossibly hard, or $n^{1-\epsilon}$-hard to approximate, for any fixed $\epsilon > 0$ [13].

We believe our formulation combines the best of both worlds: generality and computational tractability. This is a key contribution of this work. Our formulation, e.g., captures the following natural variations non-stochastically:

- Arbitrary forms of obstructions or loss of line-of-sight.
- Heterogeneous networks sharing the same channel, e.g., using differing coding or technology.
- Arbitrary variation in signal propagation and pathloss.
- Existence of unavailable links, either with or without a corresponding increase in pathloss.

The difficulty of the environment is then captured by a simple parameter, ρ, the inductive independence, which interpolates nicely between perfect environments and the worst-case instances that are hard to approximate.

Notable instantiations Connectivity Scheduling has a number of special cases of independent interest, both graph-based and geometric:

- A well-studied setting is where two links conflict if they are incident on a common link, i.e., when $\mathcal{C}$ is the square of the line graph of the link graph $\mathcal{G}$. This case corresponds to bidirectional version of the classic *radio network* model. The directed version of Connectivity Scheduling was treated in [9] as the *radio aggregation scheduling problem*.
- In *range-based or disk models*, nodes are embedded in the plane and two links are adjacent if the distance between (the closest points on) them is less than K times the length of the longer link, where K is some fixed constant. In a variant, the condition is on distances between particular nodes on the links. Also, in the the related *protocol* model, adjacency occurs if the distance is less than K_1 times

the length of the longer link plus K_2 times the length of the shorter link, for some constants K_1, K_2.

- The original driving motivation is when nodes and links are embedded in a metric space and the fractional conflicts follow the *geometric SINR model* of interference in terms of the lengths and distances between links. The implications for this setting are treated in Sect. 5. Before this work, only the case when $\mathcal{G}$ is the complete graph over a set of points in a Euclidean (or doubling) metric was considered.
- A natural special case occurs when link availability is restricted by link length, so that only reasonably long links are unavailable or attenuated, but short links follow the normal SINR laws (short links are available). This is treated in Sect. 4.
- Finally, when the conflict graph $\mathcal{C}$ is the *line graph* of the link graph $\mathcal{G}$, i.e., $\mathcal{C} = L(\mathcal{G})$, we obtain the well-known *minimum degree spanning tree* (MDST) problem, where given a graph $\mathcal{G}$, the goal is to find a spanning tree of smallest maximum degree. By König's theorem, the chromatic number of the line graph of a tree (in fact, of any bipartite graph) is equal to the maximum degree of the tree. This problem has more structure that allows for better solution: while it is NP-hard, it can be approximated within an additive one [8]. In particular, $L(\mathcal{G})$ is claw-free (does not contain an induced star graph $K_{1,3}$), which is stronger than being 2-inductive independent), and is intimately related to $\mathcal{G}$.

Related work on modeling Wireless networking in the real world behaves quite differently from the theoretical models [10, 33, 41] and typically displays a high degree of irregularity. This manifests in how the strength of signals (and the corresponding interference) often varies greatly within the same region, and is often poorly correlated with distance [2]. This behavior holds even in simple outdoor environments, but is magnified inside buildings. It is also evidenced by fluctuations, sensitivity to environmental changes (even levels of humidity), and hard-to-explain unreliability.

There are many approaches that have been proposed to model irregularity in wireless networks. We first examine static cases, or the modeling of non-geometric behavior. The basic SINR model allows the pathloss constant α to be adjusted [14], giving a first-order approximation of the signal gain. In the engineering community, it is most common to assume that the deviations are drawn from a particular stochastic distribution, typically assuming independence of events. There are, however, issues with such assumptions, including how they can be validated and how to deal with actual instances (rather than probabilistic distributions). On the computer science side, the prevailing approach is to view the variations as conforming the plane into a non-Euclidean metric space [6, 16], while retaining some tractable characteristics. This can also entail identifying appropriate parameters [4].

For frequent temporal changes, the standard engineering assumption is Rayleigh fading. Dams et al. [5] (see also [21]) showed that link scheduling algorithms are not significantly affected by such variation, assuming independence across time.

For unpredictably changing behavior, there is much research on adapting to new conditions, particularly with exponential backoff. A theoretic model proposed to specifically capture unreliability is the *dual graph model* [35], which extends the radio network model to a pair of graphs, the reliable and the unreliable links, where the latter are under adversarial control. The focus there is on distributed algorithms for one-shot problems, like global and local broadcast problems, where the nodes do not know which links are reliable. As far as we know, it has not been considered in settings involving a long-term communication structure.

3 Approximations in Terms of Inductive Independence

3.1 Greedy Algorithm

A natural greedy approach is to find as many edges as possible that form a feasible set. This set can be assigned the first color and the process repeated, ensuring that edges added always connect different trees of the forest being grown. The key step is obtaining a constant-approximation for a maximum feasible subset.

We assume in this section that $\mathcal{G}$ can have parallel edges but no loops. We assume that the conflict graph $\mathcal{C}$ is ρ-inductive independent, for a number $\rho > 0$, and that the corresponding conflict function W and ordering of edges $\prec$ are given. In typical applications, such as wireless scheduling, the ordering is known. For bounded ρ and ordinary graphs $\mathcal{C}$, it can be computed in polynomial time. In general, the problem of finding a (near-)optimal inductive independence ordering is hard. Note also that the results of this section can be recovered even if no particular ordering is known, if we assume that $\mathcal{C}$ is an ordinary graph, and replace ρ by the simpliciality of $\mathcal{C}$ [24].

In the *maximum feasible forest problem*, the goal is to find a maximum cardinality subset of edges of the link graph $\mathcal{G}$ that is both independent in the conflict graph $\mathcal{C}$ and acyclic in $\mathcal{G}$.

Algorithm 1 CapKruskal($\mathcal{G}, \mathcal{C}$)

MakeSet(v), for each $v \in V(\mathcal{G})$
$S \leftarrow \emptyset$
for $e = (u, v)$ in L in $\prec$ order **do**
    if $W(S, e) + W(e, S) < 1/2$ **and not** Connected(u, v) **then**
        $S \leftarrow S \cup \{e\}$
        Union(u, v)
    end if
end for
return $S' = \{e \in S : W(S, e) < 1\}$

The algorithm, given as Algorighthm 1, is a greedy algorithm that mixes the edge selection criteria of wireless capacity algorithms [16, 29] with the classic MST algorithm of Kruskal [34], thus the name CapKruskal. It processes the edges in order of precedence $\prec$ and adds an edge to the forest S if: a) the interference on that edge from previously selected edges S is small, and b) the edge does not induce a cycle with S (as per Kruskal). We state it in terms of the classic union-find operations of MakeSet, Connected, and Union.

Theorem 1 *Let F be a maximum feasible forest of $\mathcal{G}$. Then the output of* CapKruskal$(\mathcal{G}, \mathcal{C})$ *is a feasible forest containing $\Omega(|F|/\rho)$ edges, where ρ is the inductive independence of $\mathcal{C}$.*

Proof Let S and S' be the final sets computed in CapKruskal$(\mathcal{G}, \mathcal{C})$. By definition, S' is feasible. To argue that S' is large, we examine an arbitrary feasible forest, break it into three parts, and show that none of the parts can be too large compared to S'. This will hold, in particular, for the optimal feasible forest.

First, observe that the selection condition of the algorithm implies that $\hat{W}(S, e) < 1/2$, for each $e \in S$. It then follows that

$$\sum_{e \in S} W(S, e) = \sum_{e \in S} \hat{W}(S, e) < \frac{|S|}{2} ,$$

so for at least half of the links $e \in S$ it holds that $W(S, e) < 1$, i.e., $|S'| \geq |S|/2$.

Let I be an arbitrary feasible forest. Let I_R be those edges e in I that failed the degree condition ($\hat{W}(S, e) \geq 1/2$), and I_T those edges $e = (u, v)$ in I that failed the connectivity condition (Connected(u, v)). The rest, $I_S = I \setminus (I_R \cup I_T)$ are contained in S. We bound these sets in terms of S.

Since I_T contains only edges inside components that S also connects (recalling that I induces a forest), $|I_T| \leq |S|$. Also, clearly $I_S \subseteq I \cap S \subseteq S$, so $|I_S| \leq |S|$. To bound the size of I_R, observe first that by the definition of ρ-inductive independence,

$$\hat{W}(S, I_R) = \sum_{f \in S} \hat{W}(f, I_R) \leq \sum_{f \in S} \rho = \rho \cdot |S| .$$

On the other hand, by the selection criteria,

$$\hat{W}(S, I_R) = \sum_{e \in I_R} \hat{W}(S, e) \geq \sum_{e \in I_R} \frac{1}{2} = \frac{|I_R|}{2} .$$

Thus, $|I_R| \leq 2\rho \cdot |S|$ and $|I| \leq (2\rho + 2)|S| \leq 4(\rho + 1)|S'|$. □

Coloring algorithm Algorithm Connect repeatedly calls CapKruskal to obtain a large independent set of links and assigns it a new color class. These links are then contracted in the graph and the process repeated until we have obtained a spanning tree.

A *simple* contraction of an edge $e = (u, v)$ in a graph $G = (V, E)$ results in the graph $G/e = (V', E')$, where $V' = V \setminus \{u, v\} \cup \{uv\}$ and $E' = E \setminus \{(w, u), (z, v) \in E\} \cup \{(w, uv) : (w, u) \in E \wedge (w, v) \in E\}$. Note that *contraction leaves the conflict graph $\mathcal{C}$ intact*. The operation *Contract*$(\mathcal{G}, S)$ contracts all edges in S of a link graph $\mathcal{G}$ and outputs the resulting graph. Observe that contraction preserves upper bounds on inductive independence (but can improve it).

Algorithm 2 CONNECT($\mathcal{G}, \mathcal{C}$)

```
i ← 0
G_0 ← G
while G_i has an edge do
    S_i ← CapKruskal(G_i, C[G_i])
    G_{i+1} ← Contract(G_i, S_i)
    i ← i + 1
end while
return S_0, S_1, ..., S_{i-1}
```

The pseudocode of the algorithm is given in Algorithm 2. The proof of the following theorem follows the classic set cover argument [28].

Theorem 2 *CONNECT terminates after $O(\rho \log n) \cdot \chi$ iterations, where χ is the minimum number of colors of a spanning tree of $\mathcal{G}$, and ρ is the inductive independence of $\mathcal{C}$.*

Proof Let $S_0, S_1, \ldots, S_{i-1}$ be the sequence of edge-sets returned by CONNECT. For each k, denote $n_k = |V(\mathcal{G}_k)|$ and x_k the cardinality of the optimum independent (in $\mathcal{C}[\mathcal{G}_k]$) forest in $\mathcal{G}_k$. By averaging, $x_k \geq n_k/\chi$, and by Theorem 1, $|S_k| \geq x_k/(4\rho + 4)$. Thus,

$$n_k = n_{k-1} - |S_{k-1}| \leq n_{k-1} - \frac{x_{k-1}}{4\rho + 4} \leq n_{k-1}\left(1 - \frac{1}{\chi(4\rho + 4)}\right)$$

and by induction, for any $k_0 > (4\rho + 4)\chi \ln n$,

$$n_{k_0} \leq n_0\left(1 - \frac{1}{\chi(4\rho + 4)}\right)^{k_0} \leq n_0 \cdot e^{-k_0/((4\rho+4)\chi)} < n_0 \cdot e^{-\ln n} = 1\ ,$$

so no nodes remain in $\mathcal{G}_{k_0}$. □

A near-linear dependence on ρ, the inductive independence, is unavoidable for polynomial-time approximation algorithms. For instance, when the link graph $\mathcal{G}$ is already a spanning tree, CONNECTIVITY SCHEDULING becomes precisely the classical graph coloring problem (of $\mathcal{C}$), which is hard to approximate within an $n^{1-\epsilon}$-factor [7].

3.2 Algorithm for a Steiner Tree Extension

A natural generalization of CONNECTIVITY SCHEDULING is to allow for relay nodes that can be optionally used in the tree construction but need not be spanned. Formally, the node set V contains a subset X of terminals and we seek a Steiner (or multicast) tree that spans all the terminals. As before, we ask for the minimum number of colors to color the tree links under the conflict graph $\mathcal{C}$. We refer to this as the Steiner Connectivity Scheduling problem.

It is not hard to construct examples for which optimal multicast trees are arbitrarily better than trees that use only the terminals, even in a geometric setting. One such instance is given in Sect. 4.3.

We give a $O(\rho \log^2 n)$-approximation algorithm for Steiner Connectivity Scheduling with ρ-inductive independent conflict graph $\mathcal{C}$. The approximation factor reduces to $O(\rho \log n)$ when $\mathcal{C}$ is a *ordinary* graph (i.e., when edges have 0/1-weights). As before, the weight function W and the ordering $\prec$ are known to the algorithm.

Our algorithm is a reduction to a multi-dimensional version of the Steiner tree (MMST) problem, recently treated by Bilò et al. [3]. In MMST, each edge e of the input graph has an associated d-dimensional weight vector $\bar{w}_e \in \mathbb{R}^d_+$, where the weight of edge e along dimension i, $\bar{w}_e[i]$, indicates how much of the i-th *resource* is required by e. The objective is to find a tree that minimizes the ℓ_p-norm of its load vector, where the load vector $\bar{w}_T$ of a Steiner tree T is the sum of the weight vectors of its edges. We use the ℓ_∞-norm, as we want to minimize the maximum use of a resource. For that case, Bilò et al. [3] gives a greedy $O(\log d)$-approximation algorithm.

Given an instance of Steiner Connectivity Scheduling with link graph $\mathcal{G}$ and conflict graph $\mathcal{C}$, our reduction is as follows. Each link e in $\mathcal{G}$ is itself (or corresponds to) a resource, so there are $n = |L|$ resources. The weight of link e along dimension f is $w_e[f] = \hat{W}(f, e)$.

Our algorithm simply applies the MMST algorithm of Bilò et al. [3] to this reduced instance and then colors the tree using the algorithm InductivenessColoring below. Thus, we find a tree T that (approximately) maximizes the quantity $\max_{f\in L} \sum_{e\in E(T)} w_e[f]$. We refer to this quantity as the *inductiveness* of the tree, which is formally defined as follows:

The inductiveness of a subset S of links in $\mathcal{G}$ (w.r.t. ordering $\prec$) is $\mathcal{I}^{\prec}(S) = \max_{f\in L} \hat{W}(f, S)$.

Algorithm 3 InductivenessColoring$(S, \mathcal{C})$

```
1:  S_i ← ∅, for i ≥ 1
2:  for links e ∈ S in opposite order of ≺ do
3:      i ← min{j : Ŵ(e, S_j) < 1/2}
4:      S_i ← S_i ∪ {e}
5:  end for
6:  for i ≥ 1 do
7:      k ← 1
8:      while S_i ≠ ∅ do
9:          S_i^k ← {e ∈ S_i : W(S_i, e) < 1}
10:         S_i ← S_i \ S_i^k
11:     end while
12: end for
13: return {S_i^k}_{i,k≥1}
```

It remains to relate inductiveness to the chromatic number (Lemma 1).

Lemma 1 *Let $S \subseteq L$ be a subset of links, and denote by χ the minimum number of colors needed to color S in $\mathcal{C}$. Then $\chi = O(\mathcal{I}^{\prec}(S)\log n)$, and if additionally $\mathcal{C}$ is an ordinary graph, then $\chi = O(\mathcal{I}^{\prec}(S))$. Also, $\chi \geq \mathcal{I}^{\prec}(S)/\rho$, when $\mathcal{C}$ is ρ-inductive independent (and possibly fractional).*

Proof The first bound, $\chi = O(\mathcal{I}^{\prec}(S)\log n)$, is achieved by Algorithm 3. The algorithm takes the links in S in *opposite* order of $\prec$ and adds each link to the first set S_i where $\hat{W}(e, S_i) < 1/2$. Let i be the index of the last non-empty set S_i and let e be a link in S_i. The fact that e was not added to the sets $S_1, S_2, \dots, S_{i-1}$ implies that $\hat{W}(e, S_t) \geq 1/2$, for all $t < i$, both when e was added and later. Thus, $\hat{W}(e, S) = \sum_t \hat{W}(e, S_t) \geq (i-1)/2$, which in turn implies the total number (i) of colors used is at most $2\mathcal{I}^{\prec}(S) + 1$.

The second observation is that in the *while* loop, the index k grows to at most $\lceil \log n \rceil$. Indeed, by an averaging argument similar to the one in Theorem 1, we know that $|S_i^1| \geq |S_i|/2$. Similarly, in each iteration of the loop, $|S_i|$ is halved, which implies that the number of iterations is at most $\lceil \log n \rceil$.

Thus, the total number of colors used by the algorithm is at most $(2\mathcal{I}^{\prec}(S) + 1)\lceil \log n \rceil = O(\mathcal{I}^{\prec}(S)\log n)$. The claim for ordinary graphs $\mathcal{C}$ is proved similarly, but here the second loop is unnecessary, since every subset S_i is an independent set.

Finally, we show that $\chi \geq \mathcal{I}^{\prec}(S)/\rho$. Consider an optimal coloring of S, consisting of color classes $T_1, T_2, \dots, T_\chi$. Then,

$$\mathcal{I}^{\prec}(S) = \max_{f \in L} \hat{W}(f, S) = \max_{f \in L} \sum_{t=1}^{\chi} \hat{W}(f, T_t) \leq \sum_{t=1}^{\chi} \max_{f \in L} \hat{W}(f, T_t) \leq \rho\chi \ ,$$

as required, where in the last line we used the definition of inductive independence and the assumption that the T_t are independent. □

Theorem 3 *There is an $O(\rho \log^2 n)$-approximation algorithm for* Steiner Connectivity Scheduling *with a ρ-inductive independent conflict graph $\mathcal{C}$. If $\mathcal{C}$ is an ordinary graph, the approximation ratio is $O(\rho \log n)$.*

Proof Directly follows from Lemma 1, since the MMST algorithm of Bilò et al. [3], applied to our instances, gives a $O(\log n)$-approximation to the inductiveness of the optimal tree. □

4 All Short Links Available on the Plane: How Good is an MST?

In the geometric setting of Connectivity Scheduling, nodes are located in the plane (or in a doubling metric), and the interference between two links is a function of the lengths of links (distance between the two end-nodes), and the distance between the (endpoints of) links. For instance, in the SINR model, the interference between two links is a function that is decreasing in their distance and increasing in the length of the interfered link.

In this setting, a Euclidean minimum spanning tree (MST) over the set of nodes is a natural candidate for connectivity, since it favors short links and has low degree (or, more generally, contains few links in the vicinity of any node). Indeed, geometric properties of the conflict graphs can be used to show that an MST of n nodes is, e.g., $O(\log n)$-colorable under the Euclidean SINR model (i.e., with all links available) [17]. In this section, we examine the chromatic number of an MST in the presence of unavailable links.

We begin with introducing some basic geometric properties of the conflict graphs that can be extracted from all major interference models defined in the plane. The remainder is then split into two parts. In the first part, we focus on the case when all links shorter than some threshold are available, and prove a bound on the chromatic number of an MST, depending on the mentioned threshold. In the second part, we show that the obtained result cannot be improved significantly, and that in general an MST can have extremely large chromatic number compared with the optimum.

4.1 The Conflict Graph on the Plane

We make limited assumptions about the conflict graph $\mathcal{C}$. We first define some notions. By a t-square we mean a square of side t in the plane. A square *hits* a link if an endpoint of the link is within the square. A set of links of length *at most* ℓ is said to be s-sparse if every ℓ-square hits at most s links, and a set of links of length *at least* ℓ is d-dense if some ℓ-square hits at least d links.

Geometric model assumptions We assume that in $\mathcal{C}$, every s-sparse set of links is $O(s)$-colorable, while a d-dense set requires $\Omega(d)$ colors.

These assumptions are satisfied by all major interference models defined in the plane (or in doubling metrics); we argue this for the SINR model in Sect. 6.

4.2 MST When All Short Links are Available

We consider here the case when all short links are available. This is motivated by experimental results which indicate on one hand that signal strength is poorly correlated with distance, but also that short links are nevertheless almost always strong and reliable [41], with most of the variability in the links of intermediate range. This is probably the most natural relaxation of the problem involving geometry.

The setting is as follows. All pairs of nodes $u, v \in V$ within unit distance (after normalization) form an available link in L, i.e., are connected by an edge in $\mathcal{G}$. Also, there is a maximum distance Π so that pairs of nodes of distance more than Π are not connected by an edge. Node pairs $u, v \in V$ of distance in the range 1 to Π may or may not form an edge in $\mathcal{G}$. We call the links of length at most 1 *short* links.

We examine the approximability of Connectivity Scheduling w.r.t. Π.

Theorem 4 *Every MST of $\mathcal{G}$ can be colored using $\zeta + O(\Pi\sqrt{\chi})$ colors under the conflict graph $\mathcal{C}$, where χ is the optimum number of colors of a spanning tree and ζ is the number of colors required to schedule an MST of the complete graph over V.*

In many settings, ζ is a negligible term (see, e.g., Corollary 2), in which case we obtain a $O(\Pi/\sqrt{\chi})$-approximation.

Before proceeding to the proof, we state several technical lemmas. We assume below a fixed MST T of the link graph $\mathcal{G}$. Denote $a = \Pi\sqrt{\chi}$. We split the non-short links into *medium* links, of length from 1 to $\sqrt{a}$, and *long*, of length at least $\sqrt{a}$. We refer to the *maximal* connected subgraphs of T containing only short (non-long) links as *clusters* (resp. *blocks*). A t-square *hits* a cluster (or a block) if it contains a vertex of that cluster (block).

The plan is to show that both medium and long links of T form a $O(a)$-sparse subset, and hence can be colored using $O(a)$ colors. To that end, we show that such links are only used to connect different clusters and blocks, which cannot be too close to each other, since otherwise the MST property would be violated. As for the short links, they are part of a MST of the complete graph over V, and hence can be colored using ζ colors.

Lemma 2 *There is no short link in $\mathcal{G}$ connecting two clusters. Similarly, there is no non-long link connecting two blocks.*

Proof Suppose there is a short link e connecting two clusters C_1 and C_2. By the maximality of the clusters, $e \notin T$. Then there is a non-short link f in T connecting C_1 and C_2. Replacing f with e in T results in a smaller spanning tree, contradicting that T is an MST. □

Lemma 3 *Every $\sqrt{a}$-square S hits $O(a)$ clusters.*

Proof A $1/\sqrt{2}$-square can hit at most one cluster, as otherwise the respective vertices contained in the $1/\sqrt{2}$-square would be within unit distance and could be connected by a short link, contradicting Lemma 2. Thus, a given $\sqrt{a}$-square S hits at most $(\sqrt{2a}+1)^2 = O(a)$ clusters, since it can be covered with that many $1/\sqrt{2}$ -squares. □

Lemma 4 *Every Π-square S hits $O(a)$ blocks.*

Proof First, observe that if there are at least 2 blocks, every block B must have a vertex incident to a long link in T, because T has to connect B to some other block, and by Lemma 2, it has to use a long link for that. We partition the set of blocks hit by S into Class 1 and 2, where the former consists of blocks containing a vertex inside S that is incident to a long link in T, and Class 2, the remaining blocks. We bound the two classes separately.

Let t denote the number of long links of T that are hit by S. Since S can be covered with $O(\Pi/\sqrt{a}+1)^2 = O(\Pi/\sqrt{\chi})$ of $\sqrt{a}$-squares, at least one of them hits $\Omega(t/(\Pi/\sqrt{\chi})) = \Omega(t\sqrt{\chi}/\Pi)$ long links. Thus, the long links of T are $(t\sqrt{\chi}/\Pi)$ -dense, and by our assumption on $\mathcal{C}$, $\chi = \Omega(t\sqrt{\chi}/\Pi)$. Rearranging, we have that $t = O(a)$. We conclude that the number of Class 1 blocks hit by S is in $O(a)$.

Next, we consider Class 2 blocks. Each such block must have a vertex incident on a long link e in T, with both its endpoints outside of S. Thus, a Class 2 block must have vertices both inside and outside of S, and T uses only short or medium links to connect vertices from the two sides. For each Class 2 block B, identify a single link used in T to connect the vertices of B inside S to those outside S, and refer to it as B's *linker*. Note that each $1/\sqrt{2}$-square hits at most one linker, as otherwise the corresponding blocks could be connected with a short edge, contradicting Lemma 2. A *short linker* is one with length at most $\sqrt{\chi}$; those must have an endpoint in S within distance $\sqrt{\chi}$ from the border of S, as they must cross the border. Thus, the total area in S that can contain an endpoint of a short linker is at most $2\Pi\sqrt{\chi}$, and by covering it with $1/\sqrt{2}$-squares, we see that there are $O(\Pi\sqrt{\chi}) = O(a)$ short linkers.

We partition the non-short linkers into *i-linkers*, which have length between $Q_i := 2^i \cdot \sqrt{\chi}$ and $2Q_i = 2^{i+1} \cdot \sqrt{\chi}$, for $i = 0, 1, \ldots$. Let q_i be the number of Class 2 blocks with i-linkers. Observe that an i-linker has an endpoint in S within distance $2Q_i$ from the border of S. Thus, the total area in S that can contain an i-linker is at most $8\Pi Q_i$, and it can be covered with $O(\Pi/Q_i)$ different Q_i-squares. Thus, some Q_i-square hits $q_i/O(\Pi/Q_i) = O(q_i Q_i/\Pi)$ of i-linkers (each in T, and of length at least Q_i), so T is $\Omega(q_i Q_i/\Pi)$-dense. Hence, by our assumption on $\mathcal{C}$, $\chi = \Omega(q_i Q_i/\Pi)$, and by rearranging, $q_i = O(\sqrt{\chi}\Pi/2^i) = O(a/2^i)$. The number of Class 2 blocks is then bounded by

$$O(\Pi\sqrt{\chi}) + \sum_{i=0} q_i = O(a) + \sum_{i=0} O(a/2^i) = O(a) \sum_{i=0} 2^{-i} = O(a) .$$

□

Proof of Theorem 4 The short links are contained in an MST of the complete graph on the pointset, and hence they can be colored using ζ colors.

Next, we show that medium links form a $O(a)$-sparse set, and can be colored using $O(a)$ colors, by our assumptions on $\mathcal{C}$. Let S be a $\sqrt{a}$-square and let S' be the $3\sqrt{a}$-square with S in the center. Note that the medium links with an endpoint in S have the other endpoint within S' and are used to connect clusters that hit S'. Since there is no cycle of cluster connections, the number of medium links touching S must be fewer than the number of clusters hitting S', or $O(a)$, by Lemma 3.

Concerning the long links, we can apply a nearly identical argument as above, but considering a Π-square instead of a $\sqrt{a}$-square, and blocks instead of clusters. This shows that long links form a $O(a)$-sparse set, and can be colored using $O(a)$ colors as well. □

4.3 Limitations of MST

The bound of Theorem 4 is quite good if the optimum solution of Connectivity Scheduling has large chromatic number, compared to the maximum usable link length Π. Can we improve the bound when this is not the case? The answer is negative. The MST can actually fail quite badly.

Theorem 5 *For any integer $n > 0$, there is an instance $\mathcal{G} = (V, L)$ of links over n nodes embedded in the plane, such that $\mathcal{G}$ contains a spanning tree that is $O(1)$-colorable while every MST requires $\Omega(n^{1/3})$ colors. Moreover, this is a tight instance for Theorem 4, up to constant factors.*

Proof Let $k \geq 1$ be an integer and $K = 2k^2$. Let $V = \{o\} \cup \{v_{i,j} : i = 0, 1, \dots, k-1, j = 0, 1, \dots, K-1\}$ denote the set of $n = kK + 1 = 2k^3 + 1$ nodes. We position the nodes in the plane using polar coordinates, with the node o as the origin. For node $v_{i,j}$, angular coordinate $r_{i,j}$ is $2\pi \cdot i/k$, while its radial coordinate is $k + j$.

The available links are given by $L = O \cup S \cup Y$, where

$$O = \{(o, v_{i,1}) : i = 0, \dots, k-1\},$$
$$S = \{(v_{i,j}, v_{i,j+1}) : i = 0, \dots, k-1, j = 0, \dots, K-2\},$$
$$Y = \{(v_{i,K-1}, v_{i+1 \bmod k, K-1}) : i = 0, \dots, k-1\},$$

or the *ordinary*, the *short* and the *yuge* links. That is, the link graph is in the form of a wheel, centered at origin, with k spokes, and K nodes on each spoke (see Fig. 1). Ordinary links are incident with the origin, while the yuge links form the tire of the wheel.

We observe that $d(v_{i,K-1}, v_{i+1 \bmod k, K-1}) > k = d(o, v_{i',1})$, for any i, i'. Thus, the MST consists of the short and the ordinary links, $S \cup O$. Since all the ordinary links have an endpoint in the origin, they form a $\Theta(k)$-dense set, and they (and the MST) must be colored with $\Theta(k) = \Theta(n^{1/3})$ colors in $\mathcal{C}$.

On the other hand, a more efficient solution is to use the short links, the yuge links, and a single (arbitrary) ordinary link. As a union of three $O(1)$-sparse subsets, it can be colored with $O(1)$ colors.

Let us now view the constructed instance in the context of Theorem 4. Observe that the threshold under which all links are available is 1 (and all short links are present in $\mathcal{G}$), while $\Pi < 2k$. We have shown that the MST must use $\Theta(k) = \Omega(\Pi)$ colors, while there is a spanning tree that can be colored using $O(1)$ colors. That is, the number of colors needed for the MST is $\Omega(\Pi) = \Omega(\Pi\sqrt{\chi})$, which is what Theorem 4 claims (note that in this instance, $\zeta = O(1)$). □

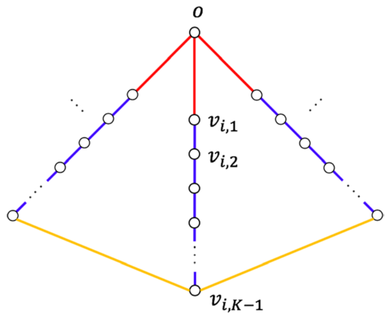

Fig. 1 The construction from Theorem 5. Yuge edges are colored orange, ordinary edges red, and short edges blue (Color figure online)

This same example shows why the known results for Euclidean SINR do not carry over to general metric spaces (even without missing links). Namely, one could simply form a metric space on the n nodes by shortest-path distances in the link graph. This example also shows how optimal multicast trees can be much more efficient than spanning trees, when we restrict the terminals to the nodes incident on yuge links along with the origin.

One way to try to overcome the hard example above would be to consider *bounded degree minimum spanning trees*. However, the example above can be modified so that the maximum degree of the resulting link graph $\mathcal{G}$ is at most 3, but the result is similar. To this end, one can replace the top vertex o in the construction with a chain of k equally spaced nodes connected into a simple path (which is a sparse subset), where each node is incident with one ordinary link. The ordinary links would still make a $\Theta(k)$-dense set.

5 Implications to Signal Strength Models

We consider in this section the implementation and implications of our results to signal strength models, most importantly the metric SINR model.

SINR-feasibility, besides the underlying metric, also depends on the *transmission power control* regime. Different regimes give different notions of feasibility. Nevertheless, it is known that for most interesting cases, SINR-feasibility has the constant-inductive independence property. In particular, power control is usually considered in two modes: *oblivious power schemes*, where links use only local information, such as the link length, to define the power level, and *global power control*, where all power levels are controlled simultaneously to give larger independent sets. The former includes the uniform power mode, where all links use equal power. Another technical issue is *directionality* of links, which is not explicitly addressed by our general results, but will be addressed below.

Let us start the discussion from Euclidean metrics (or more generally, doubling metrics). For the global power control mode, Kesselheim [29] introduced a weight function W that gives a constant-inductive independent conflict graph for any set of links (see [29, Theorem 1]), so our results apply here directly (except for directionality issues, addressed below). Similarly, for oblivious power schemes (excluding uniform power), Halldórsson et al. [23] showed that in order to get constant-inductive independence, one may take the natural weight function, *affectance* (a.k.a. relative or normalized interference) [23, Theorem 3.3]. In all cases, any non-decreasing order of links by length can be taken as the ordering $\prec$.

For general metric spaces, a slightly more technical definition of inductive independence is used. A fractional conflict graph $\mathcal{C} = (L, W)$ is (ρ, γ) − inductive *independent*, w.r.t. an ordering $\prec$ of the links, if, for every link e and every feasible set $I \in \mathcal{F}$, there is a subset $I' \subseteq I$ of size $|I'| \geq |I|/\gamma$, such that $\hat{W}(e, I) \leq \rho$. The old definition corresponds to the setting $\gamma = 1$. It is easily verified that Theorems 1 and 2 extend to cover this new definition, with approximation ratios multiplied by a factor of γ. Now, the counterparts of the results from the previous paragraph in general metrics can be found in [18, Lemmas 2,4] and [30, Theorem 1, Lemma 3], where it is shown that with appropriate weight functions, feasibility for any oblivious power scheme (including uniform power), as well as feasibility with global power control, can be expressed by an $(O(1), O(1))$-inductive independent fractional conflict graph.

5.1 Link directions

In general, weight functions derived from the SINR model depend on link directions, i.e., each orientation of links gives us a different weight function. In Connectivity Scheduling, however, we would like to have a schedule for a spanning tree which includes every link in both directions.

This issue is not present for the global power control mode, where the weight function of Kesselheim [29] is independent of directions. In particular, it gives a coloring, such that whatever direction is assigned to the links, one can find a power assignment that makes it work (the power assignment could be different for different orientations of links).

For oblivious power schemes, the following trick applies. It is known that for a set of links with some direction and an oblivious power scheme, and with the weight function W defined in terms of the affectances, if $W(e, S) < 1/2$, for all $e \in S$ (call this *dual-feasibility*), then there is another oblivious power assignment (called the *dual* of the original one) that makes S feasible with the reversed directions of links [32]. Thus, we would like to have a coloring where each color class S is also dual-feasible. To this end, it is enough to modify CapKruskal, so that the threshold 1/2 in the acceptance condition is replaced with 1/4, and the output set S' is given by $S' = \{e \in S : (W(S, e) < 1) \wedge (W(e, S) < 1/2)\}$. Similar methods then show that this again gives an $O(\rho)$-approximation to the maximum feasible forest problem. The rest of the analysis is left intact, so that we obtain an $O(\log n)$-approximation as before, but with color classes that are both feasible and dual-feasible. Then we can replace each color class with its two copies and reverse the directions of links in one of the copies. Each link thus gets a color for both directions, while the number of colors used increases by a factor of two.

We summarize the observations above in the following theorem.

Theorem 6 *There is an $O(\log n)$-approximation algorithm for Connectivity Scheduling in the SINR model in arbitrary metric spaces. This holds both in the case of oblivious power schemes and for arbitrary power control. It holds even when only a subset of the node-pairs are available as links (but interferences follow the metric SINR definitions).*

These are the first results on SINR connectivity that hold in general metrics. They are necessarily relative approximations, since in general metric spaces, there is no good upper bound on the connectivity number, even for complete graphs. A simple example is the unit metric over n nodes with pairwise unit distances between nodes and no unavailable links: The size of a maximum independent set of links is bounded by a constant, so every spanning tree has to be colored with (trivial) $\Omega(n)$ number of colors.

For the case of points in the plane (i.e., a complete link graph with conflicts induced by distances), connectivity can be achieved using $O(\log n)$ colors (with general power control) [17]. Since it is not known if $O(1)$ colors always suffice, this result is not directly implied by Theorem 2. However, it was also shown in [17] that the MST contains a feasible forest of $\Omega(n)$ edges. The rest of our analysis (using constant-inductive independence) then implies a result matching [17].

Corollary 1 *Let P be a set of points in the plane. Then, Connect finds and colors a spanning tree of P with $O(\log n)$ colors.*

5.2 Steiner trees

Similarly to Theorem 6, we can summarize our results for the Steiner variant of the problem that follow from Theorem 3.

Theorem 7 *There is an $O(\log^2 n)$-approximation algorithm for Steiner Connectivity Scheduling in the SINR model in arbitrary metric spaces. This holds both in the case of oblivious power schemes and for arbitrary power control. It holds even when only a subset of the node-pairs are available as links (but interferences follow the metric SINR definitions).*

Using global power control, we can do considerably better. The main result of Halldórsson and Tonoyan [19] shows that, for any set L of links, there is an *ordinary* (non-fractional) conflict graph $\mathcal{C}(L)$, such that every independent set in $\mathcal{C}$ is feasible under the geometric SINR model, and the chromatic number of $\mathcal{C}$ is at most $O(\log^* \Lambda)$ factor away from the chromatic number of L under SINR (using global power control). Recall that Λ is the ratio between the largest and smallest link length. Moreover, $\mathcal{C}$ is constant-inductive independent [19, Prop. 1]. This directly gives us the following Corollary of Theorem 3.

Theorem 8 *There is a $O(\log n \log^* \Lambda)$-approximation algorithm for Steiner Connectivity Scheduling in geometric SINR with global power control.*

A similar result with $O(\log \log \Lambda)$-factor instead of $O(\log^* \Lambda)$ holds also for certain oblivious power schemes (but not, for instance, uniform power) [20].

5.3 All short links available

Recall that the parameter ζ in Theorem 4 was defined as the number of colors required to color an MST in the complete graph setting, i.e., when $\mathcal{G}$ is the complete graph. For Euclidean SINR with general power control, $\zeta = O(\min(\log n, \log^* \Lambda'))$, where Λ' is the ratio of the longest and the shortest distances between the nodes [17, 22].

Corollary 2 For a graph $\mathcal{G}$ containing all short links, an MST gives a $O(\Pi + \min(\log n, \log^* \Lambda'))$-approximation in geometric SINR with power control.

6 SINR Definitions

For completeness, we include here various definitions and facts regarding the SINR model.

The *abstract SINR* model has two key properties: (i) signal decays as it travels from a sender to a receiver, and (ii) interference—signals from other than the intended transmitter—accumulates. Transmission succeeds if and only if the interference is below a given threshold. The *Metric SINR* model additionally assumes *geometric path-loss*: that signal decays proportional to a fixed polynomial of the distance, where the *path-loss constant* α is assumed to be an arbitrary but fixed constant between 1 and 6. This assumption is valid with $\alpha = 2$ in free space and perfect vacuum [12, Sect. 3.1]. In the *Euclidean SINR* model, the distances are planar.

Formally, a *link* $l_v = (s_v, r_v)$ is given by a pair of nodes, sender s_v and a receiver r_v, which are located in a metric space. Let $d(x, y)$ denote the distance between points x and y in the metric, and use the shorthand $d_{vw} = d(s_v, r_w)$. The strength of a signal transmitted from point x as received at point y is $d(x, y)^\alpha$. The *interference* I_{uv} of sender s_u (of link l_u) on the receiver r_v (of link l_v) is P_u / d_{uv}^α, where P_v is the power used by s_v. When $u = v$, we refer to I_{vv} as the *signal strength* of link l_v. If a set S of links transmits simultaneously, then the *signal to noise and interference ratio* (SINR) at l_v is

$$\mathrm{SINR}_v := \frac{I_{vv}}{N + \sum_{u \in S} I_{uv}} = \frac{P_v / d_{vv}^\alpha}{N + \sum_{u \in S} P_v / d_{uv}^\alpha}, \tag{1}$$

where N is the ambient noise. The transmission of l_v is *successful* iff $\mathrm{SINR}_v \geq \beta$, where $\beta \geq 1$ is a hardware-dependent constant.

6.1 Additional definitions: power, affectance, separability

We will work with a total order $\prec$ on the links, where $l_v \prec l_w$ implies that $d_{vv} \leq d_{ww}$. A power assignment $\mathcal{P}$ is *oblivious* if both $P_v \leq P_w$ and $\frac{P_w}{d_{ww}^\alpha} \leq \frac{P_v}{d_{vv}^\alpha}$ hold whenever $l_v \prec l_w$. This captures the main power strategies, including uniform and linear power.

The affectance $a_w^{\mathcal{P}}(v)$ [13, 32] of link l_w on link l_v under power assignment $\mathcal{P}$ is the interference of l_w on l_v normalized to the signal strength (power received) of l_v, or

$$a_w(v) = \min\left(1, c_v \frac{P_w}{P_v} \frac{d_{vv}^\alpha}{d_{wv}^\alpha}\right),$$

where $c_v = \frac{\beta}{1 - \beta N / (P_v / d_{vv}^\alpha)} > \beta$ is a factor depending only on universal constants and the signal strength P/d_{vv}^α of l_v, indicating the extent to which the ambient noise affects the transmission. We drop $\mathcal{P}$ when clear from context. Furthermore let $a_v(v) = 0$. For a set S of links and link l_v, let $a_v(S) = \sum_{l_w \in S} a_v(w)$ be the *out-affectance* of v on S and $a_S(v) = \sum_{l_w \in S} a_w(v)$ be the *in-affectance*. Assuming S contains at least two links we can rewrite Eq. 1 as $a_S(v) \leq 1$ and this is the form we will use. A set S of links is *feasible* if $a_S(v) \leq 1$ and more generally *K-feasible* if $a_v(S) \leq 1/K$.

The following theorem shows that the interference model assumptions of Sects. 4.3 and 4 hold for geometric SINR. This fact is widely known, see e.g., [17]. We outline a proof for completeness.

Theorem 9 [17] *If a link set is s-sparse, then it can be colored using $O(s)$ colors under geometric SINR, and if it is d-dense, then it requires $\Omega(d)$ colors.*

Proof The former claim essentially follows from the results of Halldórsson [15]. Here is a crude sketch of a proof. Let L be a s-sparse set of links of length at most ℓ. Partition the plane into squares of side ℓ. Assign each link to a square where it has an endpoint, ties broken arbitrarily. It is easy to color the squares using constant number of colors, such that for each color class $\mathcal{C}$, the distances between the squares in $\mathcal{C}$ are greater than $c\ell$, where c is a constant of our choice. Let $\mathcal{C}$ be any color class. Using sparsity, partition the set of links assigned to the squares in $\mathcal{C}$ into at most s subsets $S_1, S_2, \dots, S_k$, such the intersection of each S_i and each square in $\mathcal{C}$ is at most a single link. Then a standard area argument (see, e.g. [15]) shows that if the constant c is sufficiently large, S_i are feasible sets (e.g. under uniform power assignment). Note that it is important here that all links have length at most ℓ, so they are "attached" to their corresponding squares.

Now consider a subset $S \subseteq L$ that is $s(L)$-dense, and let ℓ be the minimum link length in S, and let X be a ℓ-by-ℓ square with $s(L)$ endpoints from S. Let $T \subseteq S$ be the subset of links with endpoints in X, and note that $|T| \geq s(L)/2$. The distance between any two points within X is at most $\sqrt{2}\ell$. It follows that no pair of links in T can coexist in a $\sqrt{2}^{\alpha}$-feasible set. That is, T, and therefore also L, requires $|T| \geq s(L)/2$ colors when $\beta \geq \sqrt{2}^{\alpha}$. By signal strengthening, the exact value of β changes the chromatic number of the set only by a constant factor. □

7 Conclusion and Open Questions

We introduced a new formulation of Connectivity Scheduling that captures unreliable links in wireless networks, thus extending the "vanilla" SINR setting. This was done in a way that incorporates certain abstract properties of the geometric SINR model, without overly relying on geometry, on the one hand, and without overly generalizing the model (which would imply intractability), on the other hand. Our new algorithms, besides working in the extended model, also produce results of similar quality for a more general variant, namely the Steiner variant of the problem.

Many related problems are left addressing; we mention the most prominent ones. In the *latency minimization problem*, the goal is bounding the time it takes for a packet to filter through the tree from a leaf to the root (and back). This requires optimizing both the height of the tree as well as the ordering of the links in the coloring (and the number of colors). Another challenging variation is the *directed case*, where the structure we are looking for is an arborescence, rather than a tree. This requires new techniques, as our arguments crucially depend on the graph being undirected. Finally, in various practical scenarios, it may be necessary for wireless nodes to self-organize, and, in particular, compute schedules and connect to each other via *distributed algorithms*. In the context of our work, a key challenge is handling the sequential nature of inductive independence in a distributed setting, or finding an alternative measure.

Funding Funding was provided by Icelandic Centre for Research (Grant No. 174484-051).

References

1. Akcoglu, K., Aspnes, J., DasGupta, B., Kao, M.Y.: Opportunity cost algorithms for combinatorial auctions. In: Computational Methods in Decision-Making, Economics and Finance, pp. 455–479. Springer (2002)
2. Baccour, N., Koubaa, A., Mottola, L., Zuniga, M.A., Youssef, H., Boano, C.A., Alves, M.: Radio link quality estimation in wireless sensor networks: a survey. ACM Trans. Sens. Netw. **8**(4), 34 (2012)
3. Bilò V, Caragiannis, I., Fanelli, A., Flammini, M., Monaco, G.: Simple greedy algorithms for fundamental multidimensional graph problems. In: ICALP, pp 125:1–125:13 (2017)
4. Bodlaender, M.H.L., Halldórsson, M.M.: Beyond geometry: towards fully realistic wireless models. In: PODC, pp 347–356 (2014)
5. Dams, J., Hoefer, M., Kesselheim, T.: Scheduling in wireless networks with Rayleigh-fading interference. IEEE Trans. Mob. Comput. **14**(7), 1503–1514 (2015)
6. Fanghänel, A., Kesselheim, T., Räcke, H., Vöcking, B.: Oblivious interference scheduling. In: PODC, pp 220–229 (2009)
7. Feige, U., Kilian, J.: Zero knowledge and the chromatic number. J. Comput. Syst. Sci. **57**, 187–199 (1998)
8. Fürer, M., Raghavachari, B.: Approximating the minimum-degree Steiner tree to within one of optimal. J. Algorithms **17**(3), 409–423 (1994)
9. Gandhi, R., Halldórsson, M.M., Konrad, C., Kortsarz, G., Oh, H.: Radio aggregation scheduling. In: ALGOSENSORS, pp 169–182 (2015)
10. Ganesan, D., Krishnamachari, B., Woo, A., Culler, D., Estrin, D., Wicker, S.: Complex behavior at scale: an experimental study of low-power wireless sensor networks, p. 02. Tech. rep, UCLA/CSD-TR (2002)
11. Göbel, O., Hoefer, M., Kesselheim, T., Schleiden, T., Vöcking, B.: In: Online independent set beyond the worst-case: secretaries, prophets, and periods. In: ICALP, pp. 508–519 (2014)
12. Goldsmith, A.: Wireless Communications. Cambridge University Press (2005)
13. Goussevskaia, O., Halldórsson, M.M., Wattenhofer, R.: Algorithms for wireless capacity. IEEE/ACM Trans. Netw. **22**(3), 745–755 (2014)
14. Gupta, P., Kumar, P.R.: The capacity of wireless networks. IEEE Trans. Inf. Theory **46**(2), 388–404 (2000)
15. Halldórsson, M.M.: Wireless scheduling with power control. ACM Trans. Algorithms **9**(1), 7 (2012)
16. Halldórsson, M.M., Mitra, P.: Wireless capacity with oblivious power in general metrics. In: SODA, pp. 1538–1548 (2011)
17. Halldórsson, M.M., Mitra, P.: Wireless connectivity and capacity. In: SODA, pp. 516–526 (2012)
18. Halldórsson, M.M., Mitra, P.: Nearly optimal bounds for distributed wireless scheduling in the SINR model. Distrib. Comput. **29**(2), 77–88 (2016)
19. Halldórsson, M.M., Tonoyan, T.: How well can graphs represent wireless interference? In: STOC, pp. 635–644 (2015)

20. Halldórsson, M.M, Tonoyan, T.: The price of local power control in wireless scheduling. In: FSTTCS, pp. 529–542 (2015)
21. Halldórsson, M.M., Tonoyan, T.: Wireless link capacity under shadowing and fading. In: MobiHoc, pp. 27:1–27:10 (2017)
22. Halldórsson, M.M., Tonoyan, T.: Wireless aggregation at nearly constant rate. In: ICDCS, pp. 753–763 (2018)
23. Halldórsson, M.M., Holzer, S., Mitra, P., Wattenhofer, R.: The power of oblivious wireless power. SIAM J. Comput. **46**(3), 1062–1086 (2017)
24. Halldórsson, M.M., Tonoyan, T.: Computing inductive vertex orderings. Inf. Process. Lett. **172**, 106159 (2021). https://doi.org/10.1016/j.ipl.2021.106159
25. Hoefer, M., Kesselheim, T.: Secondary spectrum auctions for symmetric and submodular bidders. ACM Trans. Econ. Comput. **3**(2), 9 (2015)
26. Hoefer, M., Kesselheim, T., Vöcking, B.: Approximation algorithms for secondary spectrum auctions. ACM Trans. Internet Technol. **14**(2–3), 16 (2014)
27. Incel, Ö.D., Ghosh, A.A., Krishnamachari, B.: Scheduling algorithms for tree-based data collection in wireless sensor networks. In: Nikoletseas, S.E., Rolim, J.D.P. (eds.) Theoretical Aspects of Distributed Computing in Sensor Networks, Monographs in Theoretical Computer Science, An EATCS Series, pp. 407–445. Springer (2011)
28. Johnson, D.S.: Approximation algorithms for combinatorial problems. J. Comput. Syst. Sci. **9**, 256–278 (1974)
29. Kesselheim, T.: A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In: SODA, pp. 1549–1559 (2011)
30. Kesselheim, T.: Approximation algorithms for wireless link scheduling with flexible data rates. In: ESA, pp. 659–670 (2012)
31. Kesselheim, T.: Dynamic packet scheduling in wireless networks. In: PODC, pp. 281–290 (2012)
32. Kesselheim, T., Vöcking, B.: Distributed contention resolution in wireless networks. In: DISC, pp. 163–178 (2010)
33. Kotz, D., Newport, C., Gray, R.S., Liu, J., Yuan, Y., Elliott, C.: Experimental evaluation of wireless simulation assumptions. In: MSWiM, pp. 78–82 (2004)
34. Kruskal, J.B.: On the shortest spanning subtree of a graph and the traveling salesman problem. Proc. Am. Math. Soc. **7**(1), 48–50 (1956)
35. Kuhn, F., Lynch, N., Newport, C., Oshman, R., Richa, A.: Broadcasting in unreliable radio networks. In: PODC, pp. 336–345 (2010)
36. Moscibroda, T.: The worst-case capacity of wireless sensor networks. In: IPSN, pp. 1–10 (2007)
37. Moscibroda, T., Wattenhofer, R.: The complexity of connectivity in wireless networks. In: INFOCOM (2006)
38. Moscibroda, T., Wattenhofer, R., Zollinger, A.: Topology control meets SINR: the scheduling complexity of arbitrary topologies. In: MOBIHOC, pp. 310–321 (2006)
39. Son, D., Krishnamachari, B., Heidemann, J.: Experimental study of concurrent transmission in wireless sensor networks. In: SenSys, ACM, pp. 237–250 (2006)
40. Ye, Y., Borodin, A.: Elimination graphs. ACM Trans. Algorithms **8(2)**, 14:1–14:23 (2012)
41. Zamalloa, M.Z., Krishnamachari, B.: An analysis of unreliability and asymmetry in low-power wireless links. ACM Trans. Sens. Netw. **3**(2), 7 (2007)