

Scheduling Multicast Transmissions Under SINR Constraints

Thomas Erlebach and Tom Grant

Dept. of Computer Science, University of Leicester,
{te17,tg53}@mcs.le.ac.uk

Abstract. The problem of scheduling wireless transmissions under signal to interference-plus-noise ratio (SINR) constraints has received increasing attention recently. While previous work has considered the unicast case where each transmission has one sender and one receiver, we consider the setting of multicast requests where each transmission has one sender and a set of receivers. A set of multicast transmissions can be scheduled in the same round if the SINR at all receivers is above a certain threshold. The goal is to minimise the number of rounds. Building on the relationship between SINR scheduling and unit disk graph colouring established by Halldórsson (ESA 2009), we present an $\mathcal{O}(\log \Gamma)$ -approximation algorithm for multicast scheduling in the SINR model, where Γ is the ratio of the longest to the shortest link length, considering only the longest link of each multicast request. The algorithm uses uniform power assignment and can be implemented online. We also compare the model of atomic multicasts (where all receivers of a multicast must receive the transmission in the same round) to the model of splittable multicasts (where a multicast sender can transmit in several rounds, each time serving a subset of its receivers). Furthermore, we consider the throughput maximisation problem and obtain an $\mathcal{O}(\log \Gamma)$ -competitive randomised online algorithm and show that every deterministic algorithm, even for unicast links and using arbitrary power assignments, has competitive ratio $\Omega(\log \Gamma)$.

1 Introduction

Wireless ad-hoc networks allow a set of wireless nodes to form a network without any pre-existing infrastructure. Such networks are very versatile, but there are limitations. To receive a transmission successfully, the reception strength of the signal must be greater than the sum of other transmissions in the network with the addition of any background noise. In other words, the signal to interference-plus-noise ratio (SINR) must be above a certain threshold. This is an important feature of wireless networks, especially in situations where nodes have close proximity to each other or where the network has a high level of traffic. Such scenarios can reduce the capacity and impair the performance of wireless networks. Therefore, it is essential that scheduling schemes take into account interference constraints, otherwise they may suffer a considerable amount of transmission errors such as corrupted or lost packets.

In earlier theoretical studies of wireless networks, interference was typically modelled in a simplistic way, e.g., by assuming that a transmission can be received by any node that is at distance at most r_1 from the sender, and the transmission creates interference at any node that is at distance at most $r_2 > r_1$ from the sender. One of the shortcomings of such models is that the cumulative effects of interference are neglected. To address this and other shortcomings, it has been proposed that the SINR model, also called the physical model, be used for theoretical worst-case analysis as well [14], and that model has now been widely adopted, see, e.g., [11, 4, 1, 6, 2, 10].

Given a set of transmission requests, each requesting a direct transmission from a sender to a receiver, a fundamental problem is to compute a shortest schedule for completing the requests. A schedule proceeds in rounds, and each round consists of a set of transmissions that can take place simultaneously without violating SINR constraints. The length of a schedule is the number of rounds. Another natural optimisation problem, called the throughput problem, is to maximise the number of transmissions that can be scheduled simultaneously in one round.

An important aspect of a wireless transmission schedule is the power that is assigned to the nodes for their transmissions. The simplest approach is to use a *uniform* power assignment where all nodes use the same transmission power. Uniform power assignments may be preferred due to the simpler implementation and must be used in homogeneous networks where power control is not available due to hardware constraints. Other power assignments that have been considered in the literature are *oblivious* power assignments, where the power assigned to the sender of a transmission is a function of the distance to the receiver, and *arbitrary* power assignments, where the power assigned to a sender can be set arbitrarily (for example, depending on the interference caused by simultaneous transmissions by other senders).

Previous work on transmission scheduling under SINR constraints has considered unicast transmissions, i.e., the case where each transmission has one sender and one receiver. However, a fundamental property of wireless transmissions is that a single transmission can be received by several nodes that are within the sender's transmission range. Furthermore, there are many scenarios where nodes of the network may want to transmit the same message to a set of other nodes, e.g., in the exchange of routing information between neighbours in a virtual topology maintained on top of the physical network or in the flooding of information across the network. Therefore, in this paper we consider the wireless transmission scheduling problem for multicast requests. Our aim is to investigate how existing methods developed for the unicast setting can be adapted to the multicast setting. We consider algorithms that use uniform power assignment. Although the lack of power control has been shown to be sub-optimal by a factor logarithmic in the maximum power used [14, 2], we argue that it is meaningful to study uniform power schedules. The predominant reason for this is that solutions based on uniform power are simpler for implementation, both in hardware and software. Thus, such methods are more likely to be adopted and

utilised by practitioners. A similar justification has also been given in [2, 10]. Furthermore, we measure the approximation ratio of our algorithms compared to optimal solutions with arbitrary power assignment.

The unicast case of both the scheduling and the throughput problem was proved \mathcal{NP} -complete by Goussevskaia et al. [8] for uniform power. As the unicast case is a special case of the multicast case, the multicast version is also \mathcal{NP} -complete for scheduling and throughput with uniform power. Andrews and Dinitz [1] prove that the throughput problem is \mathcal{NP} -hard for arbitrary power. These results provide motivation for studying approximation algorithms for these problems.

1.1 Related Work

Most of the related work that we are aware of and that we discuss in this section has been for unicast transmissions. A starting point for the theoretical analysis of the capacity of wireless networks was the seminal work by Gupta and Kumar [9] that studied the throughput in a setting where the nodes are distributed uniformly at random. More recently, there has been an extensive study of offline scheduling and throughput problems in the SINR model with respect to arbitrary networks and with a focus on worst-case analysis and approximation algorithms [14, 13, 5, 11, 4, 1, 6, 2, 10]. The uniform power variant has also been under intense scrutiny especially due to its appeal to both theoreticians and applied researchers, leading to constant-factor approximation algorithms for scheduling and throughput (compared to the optimal solution with uniform power) [6, 11]. The difference between uniform and non-uniform power assignments was first investigated by Moscibroda and Wattenhofer [14].

Halldórsson [10] considers the wireless scheduling problem in comparison with the optimal solution that uses an arbitrary power assignment. He shows that the scheduling problem can be related to the colouring of unit disk graphs at a constant-factor loss of approximation ratio in the case of links of similar length. This yields simple online algorithms with uniform power that achieve ratio $\mathcal{O}(1)$ for links of similar lengths and ratio $\mathcal{O}(\log \Gamma)$ for arbitrary links, where Γ is the ratio between the maximum and minimum link lengths. Our results for the multicast setting build on these results. Furthermore, in [10] Halldórsson also presents an $\mathcal{O}(\log n \cdot \log \log \Gamma)$ -approximation algorithm using an oblivious power assignment, the so-called *square-root assignment* [5]. This is the first approximation algorithm with ratio polylogarithmic in the size of the input. He complements this result by a lower bound construction showing that any oblivious power assignment cannot achieve a better ratio than $\Omega(\log \log \Gamma)$ in the worst case. Furthermore, he mentions that the results also hold for the throughput problem.

We note that broadcast transmissions, which can be considered as an extreme case of multicast transmissions, were studied in [7] and [15]. Additionally the online version of the wireless throughput problem has also been studied very recently by Fanghänel et al. [3]. They assume that requests can have a duration

in the interval $[1, T]$. They show that no deterministic algorithm with oblivious power assignment can achieve competitive ratio better than $\Omega(T \cdot \Gamma^{d/2})$. Furthermore, they present an $\mathcal{O}(T \cdot \Gamma^{(d/2)+\varepsilon})$ -competitive deterministic online algorithm and a randomised $\mathcal{O}(\log \Gamma \cdot \log T)$ -competitive online algorithm. They also consider a generalisation of the problem where requests have to be assigned to one of k channels.

1.2 Our Results

We consider the scheduling problem for multicast requests under SINR constraints, assuming that the nodes are points in the two-dimensional Euclidean plane and the received signal strength is proportional to the power of the transmitted signal divided by the distance to the power of α , where α is the path loss exponent and assumed to be a constant greater than 2 (in practice, α is typically assumed to be between 2 and 6). We consider multicast transmissions to be *atomic*, which is to say that in a given multicast group, all receivers must successfully receive the transmission in the same round. Building on Halldórsson's [10] results for the unicast case, we present an $\mathcal{O}(\log \Gamma)$ -approximation algorithm for multicast scheduling under SINR constraints. In the multicast case, Γ denotes the ratio of the longest to the shortest link length, considering only the longest link of each multicast request. Note that Γ can be much smaller than the ratio of the longest to the shortest link among all unicast links that are part of a multicast request. As our algorithm is based on partitioning the requests into length classes and using unit disk graph colouring to schedule each length class, it can also be used as a simple online algorithm.

We also discuss the relationship between schedules for such atomic multicasts and schedules for splittable multicasts, where the sender of a multicast can transmit in several rounds, serving a subset of its receivers in each round.

As with previous work, we show that our approach is also applicable to the case of throughput maximisation. Moreover, we present a lower bound showing that every deterministic online algorithm has competitive ratio $\Omega(\log \Gamma)$ even for the unicast version of the throughput problem. This complements previous lower bounds where the algorithm was restricted to oblivious power assignments.

2 Preliminaries

We assume that senders and receivers are points in two-dimensional Euclidean space, and we use $\delta(p, q)$ to denote the Euclidean distance between two points p and q . For a given undirected graph $G = (V, E)$, $\Delta(G)$ denotes the maximum degree of any vertex in V . Note that any graph G can be coloured with $\mathcal{O}(\Delta(G))$ colours using any greedy colouring algorithm. A graph $G = (V, E)$ is a *unit disk graph* (for disks with radius r) if each vertex $v \in V$ can be associated with a disk of radius r centred at a point p_v in the plane such that two vertices $u, v \in V$ are adjacent if and only if the corresponding disks intersect (or, equivalently, if the distance between p_u and p_v is at most $2r$). We note the following simple property of unit disk graphs.

Lemma 1. *If a unit disk graph with disks of radius r has maximum degree $k \geq 1$, increasing the radius of the disks from r to cr for some $c > 1$ increases the maximum degree to at most $\mathcal{O}(kc^2)$.*

In the following, we first discuss notation for unicast requests and multicast requests and then define the problems we are considering. As we are presenting adaptations of previous work, in particular [10], we strive to keep similar notation for clarity.

2.1 Unicast Requests

A unicast request ℓ_v is a transmission request from a single sender s_v to a single receiver r_v . We also interchangeably refer to requests as *links*, as is common practice in the literature.

We override notation when the context is clear and denote the *length* of a link ℓ_v by $\ell_v = \delta(s_v, r_v)$. Additionally, we use the shorthand δ_{uv} to refer to the distance $\delta(s_u, r_v)$ between the sender s_u of a link ℓ_u and the receiver r_v of another link ℓ_v . A set L of links is *nearly equi-length* if the lengths of the links in L are within a factor of 2 of each other. When discussing a set of nearly equi-length links, we let D be a value such that the lengths of all links lie in the interval $[D, 2D]$.

For a set L of nearly equi-length unicast links with lengths in $[D, 2D]$, as in [10] $G'_q(L)$ denotes the unit disk graph formed by disks with radius $qD/2$ and with the unicast receivers as centres.

For a given set L of unicast links, we denote by Γ the ratio $\max_v \ell_v / \min_u \ell_u$ of the maximum link length to the minimum link length.

We denote the power assigned to a sender s_v as \mathcal{P}_v . For all our algorithms we assume that all senders are assigned the same power and that there is no constraint on the maximum power level.

To model the degradation of a transmitted signal over distance, we let α be the path loss exponent and adopt the common assumption that $2 < \alpha \leq 6$. The received signal strength of a transmission on link ℓ_v from s_v to r_v is $\mathcal{P}_v / \ell_v^\alpha$.

Let \mathcal{S}_v be the set of senders that are transmitting concurrently with a sender s_v and let N represent the background noise in the network. We use the SINR model of interference, which is interchangeably termed the *physical model*. A unicast transmission ℓ_u from s_u to r_u is successfully received if the following constraint is satisfied:

$$\frac{\mathcal{P}_u / \ell_u^\alpha}{N + \sum_{s_v \in \mathcal{S}_u \setminus \{s_u\}} \mathcal{P}_v / \delta_{vu}^\alpha} > \beta \quad (1)$$

Here, $\beta \geq 1$ denotes the minimum SINR required for a successful reception. As in previous work, we assume throughout the paper without loss of generality that $N = 0$. This assumption can be justified by noting that the effect of noise can be made arbitrarily small by scaling up the power of all senders. Furthermore, one sometimes considers the case $\beta = 1$ for simplicity, as by Lemma 6 any constant $\beta > 1$ can be achieved losing only a constant factor in the approximation ratio.

2.2 Multicast Requests

A *multicast request* or *multicast group* is a set of unicast links with a common sender. A multicast group m_v is represented as a pair (s_v, R_v) where s_v is the sender and $R_v = \{r_{v_1}, r_{v_2} \dots r_{v_{k_v}}\}$ is a set of $k_v \geq 1$ receivers. Intuitively, a multicast request (s_v, R_v) asks for a single transmission by the sender s_v that is successfully received by all receivers in R_v simultaneously. For $1 \leq i \leq k_v$, we use ℓ_{v_i} to refer to the link with sender s_v and receiver r_{v_i} . Again, we override notation and also use ℓ_{v_i} to refer to the length of that link. Without loss of generality, we assume that the receiver with index 1 is a receiver that is furthest from the sender of m_v , i.e., $\ell_{v_1} = \max_{1 \leq i \leq k_v} \ell_{v_i}$. The distance between the sender of multicast group m_u to a given receiver r_{v_i} in multicast group m_v is denoted by $\delta_{uv_i} = \delta(s_u, r_{v_i})$.

For a given set M of multicast groups, we denote by Γ the ratio of the maximum link length to the minimum link length amongst all longest links of the multicast groups in M , i.e., the value $\max_v \ell_{v_1} / \min_u \ell_{u_1}$.

A set M of multicast links is *nearly equi-length* if the lengths of the longest links in each group are within a factor of 2 of each other, i.e., there is a D such that $\ell_{v_1} \in [D, 2D]$ for all $m_v \in M$.

A multicast transmission from s_u to R_u is successfully received if for all receivers $r_{u_i} \in R_u$ inequality (1) holds.

A *schedule* for a set M of multicast links is a partition of M into subsets, called *rounds* or *slots*, and an assignment of powers to all senders of requests in M , such that the multicast transmissions assigned to the same slot are successfully received by all their receivers.

Halldórsson [10] defines the affectance on the receiver r_v of a unicast link ℓ_v to be the ratio of the interference received from concurrent transmissions by other senders to the received signal strength at r_v from s_v . In the multicast context we define the affectance on a receiver r_{v_i} in a multicast group m_v to be the ratio of the interference received from concurrent transmissions by other senders to the received signal strength at r_{v_i} from s_v . Note that a receiver r_{v_i} successfully receives a transmission from s_v if and only if the affectance of concurrent transmissions on r_{v_i} is at most $1/\beta$.

For $p \geq 1$, a *p-signal* schedule is a schedule for which the affectance on any receiver of any multicast request is at most $1/p$. (Equivalently, the SINR at every multicast receiver is at least p .) A *p-signal slot* or *p-signal set* is one round of a *p-signal* schedule.

2.3 Problem Definitions

We are concerned with the following two optimisation problems. The multicast scheduling problem, denoted by M-SCHEDULING, is to compute, for a given set M of multicast requests, a *p-signal* schedule with a minimum number of rounds. The multicast throughput problem, denoted by M-THROUGHPUT, is to compute, for a given set M of multicast requests, a largest subset of M that forms a *p-signal*

slot. The corresponding problems for unicast requests are called SCHEDULING and THROUGHPUT.

An algorithm for one of these problems is called a ρ -approximation algorithm, and has approximation ratio ρ , if it runs in polynomial time and always outputs a solution that is at most a factor of ρ away from the optimum.

In the online versions of these problems, the requests are presented to the algorithm one by one, and the algorithm must process each request without knowledge of future requests. In the scheduling problems, this means that the algorithm must assign a round and a power to the request, and in the throughput problem, the algorithm must accept or reject the request and, if accepted, assign a power to the sender. In both problems, decisions of the algorithm are irrevocable, and the solution must be feasible at all times. In the online version, we compare the quality of the solution produced by an algorithm with the quality of an optimal offline solution for the same input. An online algorithm is ρ -competitive, or has *competitive ratio* ρ , if it always produces a solution that is within a factor ρ of the optimum.

Even though our algorithms use uniform power assignments, we compare their solutions with an optimal solution that can use arbitrary power assignments. This strengthens the approximation results. For a given set R of unicast or multicast links, we denote an optimal solution by $\text{OPT}_p(R)$. For convenience, when discussing the scheduling problem, we may also use $\text{OPT}_p(R)$ to refer to the length of the optimal schedule. If $p = 1$, we also write OPT for OPT_p .

So far we have assumed that multicast requests are *atomic*, i.e., the sender of a multicast request can transmit only once and all its receivers must successfully receive the transmission in the same round. One can also consider a variant of the scheduling problem with *splittable* multicast requests. In that variant, the sender of a multicast request can transmit in several rounds, and each of its receivers must successfully receive the transmission in at least one of those rounds. It is clear that the optimal splittable schedule cannot be longer than the optimal atomic schedule.

3 Algorithm for Multicast Scheduling

We present an $\mathcal{O}(\log \Gamma)$ -approximation algorithm for M-SCHEDULING. We follow the approach of Halldórsson [10] and show that a constant-factor approximation for nearly equilength multicast groups can be achieved by a greedy colouring of a suitably defined unit disk graph. The difficulty is that a set of nearly equilength multicast groups may contain unicast links that are much shorter than the longest links of the groups, and hence it is not enough to argue about nearly equilength unicast links. We recapitulate some results by Halldórsson [10] which we require.

As in [10], two unicast links ℓ_u and ℓ_v are called *q-independent* if they satisfy

$$\delta_{uv} \cdot \delta_{vu} \geq q^2 \cdot \ell_v \ell_u. \quad (2)$$

A set of links is *q-independent* if any pair of links in the set is *q-independent*.

For a given set L of unicast requests, the *link graph* $G_q(L)$ is a graph with a vertex for each request in L and an edge between two vertices if the corresponding requests are not q -independent.

Recall that $G'_q(L)$ denotes the unit disk graph with a disk of radius $qD/2$ centred at each receiver of a link in L .

Halldórsson shows that there is a close relationship between link graphs $G_q(L)$ and unit disk graphs $G'_{q'}(L)$ for the same set of unicast links.

Lemma 2. [10] *For any $q \geq 1$ and any set L of nearly equilength unicast links, $G'_q(L) \subseteq G_{q+1}(L)$ and $G_q(L) \subseteq G'_{2(q+1)}(L)$.*

Halldórsson shows that unicast links that belong to the same q^α -signal slot are q -independent, and he establishes a slightly weaker version of the converse statement as stated in the following lemma.

Lemma 3. [10] *Let S be a z -independent set of nearly equilength unicast links. Then, with uniform power assignment, S is an $\Omega(z^\alpha)$ -signal set.*

Furthermore, Halldórsson gives the following lower bound on OPT for nearly equilength links.

Lemma 4. [10] *Let L be a set of nearly equilength unicast links and q be any constant. Then $\text{OPT}(L) = \Omega(\Delta(G_q(L)))$.*

Halldórsson combines the above statements to show that a constant factor approximation for scheduling a set L of nearly equilength unicast links can be obtained by greedily colouring either $G_q(L)$ or $G'_{q'}(L)$, for suitably chosen constants q and q' . The results listed above can be combined to derive the statement of Theorem 1.

Theorem 1. [10] *Let $p \geq 1$ be an arbitrary constant. Then there is a $q = q(p) \geq 1$ such that for any set L of nearly equilength unicast links, any colouring of $G'_q(L)$ with $\mathcal{O}(\Delta(G'_q(L)))$ colours gives a p -signal schedule with uniform power that is within a constant factor of the optimal p -signal schedule for L .*

Proof. By Lemma 3, for every p there is a z such that any z -independent set of nearly equilength links is a p -signal set (with uniform power). Let $q = 2(z + 1)$. By Lemma 2, $G_z(L) \subseteq G'_q(L)$, and hence any independent set in $G'_q(L)$ is z -independent and therefore a p -signal slot. Thus, any colouring of $G'_q(L)$ constitutes a p -signal schedule. By Lemma 4, the optimal 1-signal schedule for L , and therefore also the optimal p -signal schedule for L , has length $\Omega(\Delta(G_{q+1}(L)))$ and thus, by Lemma 2, length $\Omega(\Delta(G'_q(L)))$. \square

Consider Algorithm 1 for scheduling a set of nearly equilength multicast groups. It creates a unit disk graph H that has one disk for every multicast group, centred at the receiver of the longest link in that multicast group. The radius is set to $(q/2 + 4)D$, where q is chosen according to Theorem 1. Then a greedy colouring of that unit disk graph is returned as the schedule, i.e., we

Data: a set M of multicast requests with longest link lengths in $[D, 2D]$

Result: a p -signal schedule S_M for M with uniform power

- 1 let $q = q(p)$ be the value from Theorem 1;
- 2 construct the unit disk graph H with disks of radius $(q/2 + 4)D$ centred at the receivers r_{v_1} of all $m_v \in M$;
- 3 greedily colour H with $\mathcal{O}(\Delta(H))$ colours;
- 4 return the colouring of H as a p -signal schedule S_M with uniform power;

Algorithm 1: Algorithm for nearly equilength multicast requests

schedule each multicast group in the round given by the colour assigned to the corresponding disk, and assign uniform power to all senders.

For a set M of multicast groups, let L_M be the set of unicast links obtained by taking the longest link ℓ_{v_1} from every multicast group $m_v \in M$.

Lemma 5. *Let M be a set of nearly equilength multicast groups. For any constant $p \geq 1$, any greedy colouring of the unit disk graph H constructed by Algorithm 1 gives a p -signal schedule for M whose length is at most a constant factor longer than the optimal p -signal schedule for L_M .*

Proof. Let D be such that the length of the longest link in each multicast group lies in the interval $[D, 2D]$.

The algorithm chooses $q = q(p) \geq 1$ according to Theorem 1. Hence we have that any colouring of $G'_q(L_M)$ with $\mathcal{O}(\Delta(G'_q(L_M)))$ colours gives a p -signal schedule with uniform power for L_M that is within a constant factor of $\text{OPT}_p(L_M)$.

The unit disk graph H constructed by the algorithm is the unit disk graph obtained from $G'_q(L_M)$ by increasing the radius of the disks from $qD/2$ to $(q/2 + 4)D$. As $q \geq 1$, this increases the radius of the disks by at most a factor of 9, so by Lemma 1 the maximum degree of H is within a constant factor of the maximum degree of $G'_q(L_M)$. Hence, a greedy colouring of H with $\mathcal{O}(\Delta(H))$ colours uses only $\mathcal{O}(\Delta(G'_q(L_M)))$ colours. Note that any greedy colouring of H is also a colouring of $G'_q(L_M)$ with $\mathcal{O}(\Delta(G'_q(L_M)))$ colours and thus, by Theorem 1, constitutes a p -signal schedule S_{L_M} for L_M that is a constant-factor approximation of $\text{OPT}_p(L_M)$.

Observe that for any multicast group m_v , the disk in H with centre r_{v_1} contains the disks with radius $qD/2$ centred at any receiver r_{v_j} of the multicast group m_v , as the distance between two receivers of the same multicast group is at most $4D$ by the triangle inequality (using that each unicast link in m_v has length at most $2D$).

We claim that a greedy colouring of H gives a p -signal schedule S_M for M with uniform power that is within a constant factor of $\text{OPT}_p(L_M)$. Since the maximum degree of H is within a constant factor of the maximum degree of $G'_q(L_M)$, the number of colours in any greedy colouring of H is a constant-factor approximation of $\text{OPT}_p(L_M)$. It remains to show that the schedule S_M derived from any greedy colouring of H constitutes a p -signal schedule. Consider an arbitrary link ℓ_{v_i} of a multicast group m_v . Let U be the set of all multicast

groups $m_u \neq m_v$ that are scheduled in the same round as m_v in S_M . If $i = 1$, i.e., if ℓ_{v_i} is the longest link of m_v , we can simply argue as follows: The received signal strength from s_v at r_{v_1} and the total strength of interfering signals received at r_{v_1} are the same in S_M and in S_{L_M} , so the affectance at r_{v_1} is at most $1/p$ in S_M . If $i \neq 1$, we need a more elaborate argument. Let s'_v be an arbitrary point in the plane that has distance $2D$ from r_{v_i} , and let ℓ'_v be the unicast link with sender s'_v and receiver r_{v_i} . Consider the unit disks of radius $qD/2$ centred at r_{v_i} and at all r_{u_1} for $m_u \in U$. Observe that these unit disks are the disks that constitute $G'_q(L')$, where L' is the set of unicast links containing the link ℓ'_v and the links ℓ_{u_1} for all $m_u \in U$. Furthermore, these unit disks are disjoint since they are contained in the respective disks of radius $(q/2 + 4)D$ that have received the same colour in H . By Theorem 1, the links in L' constitute a p -signal set $S_{L'}$. In the round in which m_v is scheduled in S_M , the received signal strength from s_v at r_{v_i} is at least as large as the received signal strength from s'_v at r_{v_i} in $S_{L'}$, and the total strength of interfering signals received at r_{v_i} is the same in S_M and in $S_{L'}$. Therefore, the affectance at r_{v_i} is at most $1/p$ in S_M . \square

As any p -signal schedule for M is also a p -signal schedule for L_M , it is clear that $\text{OPT}_p(L_M)$ cannot be larger than $\text{OPT}_p(M)$. Thus, we obtain the following corollary.

Corollary 1. *For any constant $p \geq 1$, Algorithm 1 is a constant-factor approximation algorithm for M-SCHEDULING with nearly equilength multicast groups.*

We can now tackle arbitrary sets M of multicast requests using the standard approach of partitioning the requests into a logarithmic number of length classes. Assume without loss of generality that the lengths of the longest links of all multicast groups lie in the interval $[1, \Gamma]$. Partition the set M of multicast groups into $\lceil \log \Gamma \rceil$ classes M_i where M_i consists of all multicast groups with a longest link of length in $[2^i, 2^{i+1})$. Apply Algorithm 1 to each class M_i separately and obtain a schedule for M by concatenating the schedules for the classes M_i . As the schedule for each class M_i is a constant-factor approximation of $\text{OPT}_p(M_i)$ and therefore also of $\text{OPT}_p(M)$, we obtain a p -signal schedule that is an $\mathcal{O}(\log \Gamma)$ -approximation of $\text{OPT}_p(M)$.

Theorem 2. *For every constant $p \geq 1$, there is an $\mathcal{O}(\log \Gamma)$ -approximation algorithm for the problem of computing a shortest p -signal schedule for a given set of multicast requests, i.e., for the M-SCHEDULING problem.*

Since the partition of multicast requests into length classes and the greedy colouring of the unit disk graphs for each length class can be performed online, the same approach gives an $\mathcal{O}(\log \Gamma)$ -competitive online algorithm for multicast scheduling. Furthermore, as pointed out by Halldórsson [10], approaches based on colouring of unit disk graphs are amenable to a distributed implementation.

3.1 Signal Strengthening

It is also interesting to relate the length of the optimal p -signal schedule to the optimal 1-signal schedule. The following result from [11] shows that in the

unicast case, a larger SINR for all receivers can be achieved at a constant-factor loss in the schedule length.

Lemma 6. [11] *There is a polynomial-time algorithm that takes a p -signal schedule for a set of unicast links and refines it into a p' -signal schedule, for $p' > p$, increasing the number of slots by a factor of at most $\lceil 2p'/p \rceil^2$.*

Adapting this lemma to the multicast setting does not seem straightforward, but we are able to establish an analogous result at least for nearly equilength multicast links.

Lemma 7. *If there is a p -signal schedule for a set M of nearly equilength multicast groups that has length k , then there is a p' -signal schedule for M , for any constant $p' > p$, of length $\mathcal{O}(k)$. Moreover, such a schedule can be computed in polynomial time.*

Proof. Let \mathcal{A} be a p -signal schedule for a set M of multicast groups. Let L_M be the set of unicast links obtained by taking the longest link ℓ_{v_1} from every multicast group m_v in M . It is clear that \mathcal{A} can also be viewed as a p -signal schedule for L_M . As L_M is a set of unicast links, we can transform \mathcal{A} into a p' -signal schedule \mathcal{A}' for L_M by Lemma 6, such that the length of the schedule increases by only a constant factor. As the optimal p' -signal schedule for L_M cannot be longer than \mathcal{A}' , we have that the length of $\text{OPT}_{p'}(L_M)$, the optimal p' -signal schedule for L_M , is within a constant factor of the length of \mathcal{A} .

Consider the p' -signal schedule S_M computed for M by Algorithm 1 in polynomial time. By Lemma 5, the length of S_M is within a constant factor of $\text{OPT}_{p'}(L_M)$, and therefore also within a constant factor of the length of \mathcal{A} . Thus, S_M is a p' -signal schedule for M that is within a constant factor of the length of \mathcal{A} . \square

Applying Lemma 7 to an optimal 1-signal schedule for M , we obtain the following corollary.

Corollary 2. *For any constant $p \geq 1$ and any set M of nearly equilength multicast groups, $\text{OPT}_p(M)$ is at most a constant factor longer than $\text{OPT}(M)$.*

This result shows that if we require strengthening of the SINR in a schedule, we will lose only a constant factor in the schedule length, at least for nearly equilength multicast requests.

4 Splittable versus Atomic Multicast

In this section we discuss the relationship between splittable and atomic multicast requests. For ease of presentation, we only consider 1-signal schedules in this section. For a given set M of multicast requests, denote by $\text{OPT}^s(M)$ the length of an optimal 1-signal schedule that is allowed to split a multicast request. As before, $\text{OPT}(M)$ denotes the length of an optimal 1-signal schedule with atomic multicasts.

Since every atomic schedule is also a splittable schedule, it is clear that $\text{OPT}^s(M) \leq \text{OPT}(M)$ for any set M of multicast requests. Furthermore, we have the following lemmas.

Lemma 8. *For any set M of nearly equilength multicast requests, $\text{OPT}(M) = \mathcal{O}(\text{OPT}^s(M))$.*

Proof. Let L_M be defined as in Section 3. We have $\text{OPT}(L_M) \leq \text{OPT}^s(M)$ since even a splittable schedule must schedule the longest link of each multicast group in some round. Furthermore, by Lemma 5 there is an atomic 1-signal schedule for M that is within a constant factor of $\text{OPT}(L_M)$. \square

Lemma 9. *For any set M of multicast requests, $\text{OPT}(M) = \mathcal{O}(\log \Gamma) \cdot \text{OPT}^s(M)$.*

Proof. Consider a partition of M into $\mathcal{O}(\log \Gamma)$ length classes M_i , such that each length class is nearly equilength. The maximum of $\text{OPT}(L_{M_i})$ over all i is a lower bound on $\text{OPT}^s(M)$. Furthermore, the algorithm of Section 3 computes an atomic schedule for M that is within a constant factor of the sum of the values $\text{OPT}(L_{M_i})$ over all i . \square

5 Throughput Maximisation

5.1 Algorithms

In this section, we discuss how the approach described in Section 3 can be adapted to the problem M-THROUGHPUT. We consider only atomic multicasts. For a given set M of nearly equilength multicast groups, one can construct the unit disk graph H as in Algorithm 1 and then compute a maximal independent set I in H , for example using a greedy algorithm. The set I forms a 1-signal set. It is known that in unit disk graphs, the size of any maximal independent set is within a factor of 5 of the maximum independent set [12].

Let $I^* \subseteq M$ be a 1-signal set of largest size. By Lemma 5, the unit disks in H corresponding to the multicast groups in I^* can be coloured with a constant number of colours. Hence, I^* contains a set of multicast groups that corresponds to an independent set in H of size $\Omega(|I^*|)$. The set I computed by the algorithm then also has size $\Omega(|I^*|)$ and therefore constitutes a constant-factor approximation of the largest 1-signal set.

For general sets of multicast requests, we can partition the multicast requests into $\mathcal{O}(\log \Gamma)$ length classes, compute a 1-signal set for each length class as described above, and output the largest of these 1-signal sets. This gives an $\mathcal{O}(\log \Gamma)$ -approximation for M-THROUGHPUT. The same approach can be used to obtain a randomised $\mathcal{O}(\log \Gamma)$ -competitive online algorithm. The randomisation is only needed to select one of the length classes at the beginning of the algorithm.

5.2 Online Lower Bound

In this section we present a lower bound showing that no deterministic online algorithm for the throughput problem can achieve competitive ratio better than $\Omega(\log \Gamma)$ even in the case of unicast requests. We make use of the following fact that was stated by Avin et al. [2].

Lemma 10. [2] *Two senders s_1 and s_2 cannot transmit successfully at the same time if their respective receiver is closer to the other sender i.e., if $\delta(s_1, r_1) > \delta(s_1, r_2)$ and $\delta(s_2, r_2) > \delta(s_2, r_1)$.*

Now utilising Lemma 10 we provide a construction bounding the performance of any arbitrary deterministic algorithm for the online variant of THROUGHPUT.

Theorem 3. *The competitive ratio of every deterministic online algorithm, even with arbitrary power assignments, is $\Omega(\log \Gamma)$ for THROUGHPUT.*

Proof. Consider the following construction. Let n be an arbitrary positive integer, and let $\Gamma = 2b^n$ for a sufficiently large constant $b > 1$. All senders and receivers are located on the x-axis, so we can identify a node with its x-coordinate.

Let A be an arbitrary deterministic online algorithm. The adversary first presents a request ℓ_0 with sender $s_0 = 2b^n$ and receiver $r_0 = 0$. The algorithm must accept ℓ_0 as otherwise its competitive ratio is unbounded if the request sequence stops here. Next, the adversary presents requests ℓ_1, \dots, ℓ_n where for each $1 \leq i \leq n$ the sender and receiver of ℓ_i are $s_i = -b^i$ and $r_i = b^i$, respectively.

By Lemma 10, the algorithm cannot accept any of the requests ℓ_1, \dots, ℓ_n as none of them can transmit at the same time as ℓ_0 . It remains to show that an optimal solution can reject ℓ_0 and accept all other requests.

Let ℓ_1, \dots, ℓ_n transmit simultaneously using the square root power assignment, i.e., assign power $\sqrt{b^i}$ to s_i for all i .

The strength of the signal received at r_i from s_i is

$$\frac{\sqrt{b^i}}{(2b^i)^\alpha} = \frac{b^{(0.5-\alpha)i}}{2^\alpha}.$$

The total interference received at r_i is

$$\sum_{j < i} \frac{\sqrt{b^j}}{(b^j + b^i)^\alpha} + \sum_{j > i} \frac{\sqrt{b^j}}{(b^j + b^i)^\alpha}. \quad (3)$$

We can bound the first sum in (3) as follows:

$$\sum_{j < i} \frac{\sqrt{b^j}}{(b^j + b^i)^\alpha} \leq \sum_{j < i} \frac{\sqrt{b^j}}{(b^i)^\alpha} \leq \frac{\sqrt{b^i}}{(\sqrt{b} - 1) \cdot (b^i)^\alpha} = \frac{b^{(0.5-\alpha)i}}{\sqrt{b} - 1}$$

The second sum in (3) can be bounded as follows:

$$\sum_{j > i} \frac{\sqrt{b^j}}{(b^j + b^i)^\alpha} \leq \sum_{j > i} \frac{\sqrt{b^j}}{(b^j)^\alpha} = \sum_{j > i} b^{(0.5-\alpha)j} \leq 2 \cdot b^{(0.5-\alpha)(i+1)}$$

where the last inequality holds for sufficiently large b . The SINR at r_i is

$$\frac{b^{(0.5-\alpha)i} 2^{-\alpha}}{b^{(0.5-\alpha)i} \left(\frac{1}{\sqrt{b-1}} + 2 \cdot b^{0.5-\alpha} \right)} = \frac{2^{-\alpha}}{\frac{1}{\sqrt{b-1}} + 2 \cdot b^{0.5-\alpha}}$$

For b a sufficiently large constant (chosen depending on α), the SINR is larger than 1 (or any other desired constant SINR threshold). \square

We remark that the lower bound of Theorem 3 applies even in the case where the algorithm can change the power assigned to previously accepted requests upon the acceptance of a new request.

6 Conclusion

In this paper we have studied wireless scheduling and throughput problems for multicast requests in the SINR model. We have presented $\mathcal{O}(\log \Gamma)$ -competitive algorithms for both problems by exploiting and extending the relationship between SINR scheduling and unit disk graph colouring that has been established by Halldórsson for unicast links [10]. This shows that the approach of reducing SINR scheduling problems for nearly equilength links to unit disk graph colouring extends to multicast requests provided that the longest links in the multicast groups are nearly equilength. We have also given an $\Omega(\log \Gamma)$ lower bound on the competitive ratio of any deterministic online algorithm for throughput even in the case of unicast links and arbitrary power assignments, and discussed relationships between scheduling with atomic and splittable multicast requests.

Several questions are left open. First, it would be interesting to find out whether offline approximation algorithms with ratio better than $\mathcal{O}(\log \Gamma)$ exist for the multicast scheduling and throughput problems. Halldórsson's $\mathcal{O}(\log n \cdot \log \log \Gamma)$ -approximation algorithm for unicast links seems difficult to extend to the multicast setting. Furthermore, it would also be useful to determine whether signal strengthening can be done for arbitrary multicast requests while losing only a constant factor. This is true for arbitrary unicast requests (Lemma 6), but our current proof of Lemma 7 works only for nearly equilength multicast requests. Finally, it would be interesting to know whether the factor $\mathcal{O}(\log \Gamma)$ bounding the difference in schedule length between atomic and splittable schedules in Lemma 9 can be reduced.

References

1. M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *INFOCOM*, pages 1332–1340, April 2009.
2. C. Avin, Z. Lotker, and Y.-A. Pigolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, LNCS 5757, pages 373–384. Springer, 2009.

3. A. Fanghänel, S. Geulen, M. Hofer, and B. Vöcking. Online capacity maximization in wireless networks. In *Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 92–99. ACM, 2010.
4. A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 220–229. ACM, 2009.
5. A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 5556, pages 447–458. Springer, 2009.
6. O. Goussevskaia, M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of arbitrary wireless networks. In *INFOCOM*, pages 1872–1880, April 2009.
7. O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local Broadcasting in the Physical Interference Model. In *ACM SIGACT-SIGOPT International Workshop on Foundations of Mobile Computing (DialM-POMC), Toronto, Canada*, August 2008.
8. O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 100–109. ACM, 2007.
9. P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.
10. M. Halldórsson. Wireless Scheduling with Power Control. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, LNCS 5757, pages 368–380. Springer, 2009. Preliminary full version available at: <http://www.hr.is/faculty/mmh/papers/ESAFull.pdf>.
11. M. Halldórsson and R. Wattenhofer. Wireless Communication is in APX. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 5555, pages 525–536. Springer, 2009.
12. M. V. Marathe, H. Breu, H. B. Hunt III, S. S. Ravi, and D. J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25:59–68, 1995.
13. T. Moscibroda. The worst-case capacity of wireless sensor networks. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–10, New York, NY, USA, 2007. ACM.
14. T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *INFOCOM*, 2006.
15. G. Resta and P. Santi. Latency and capacity optimal broadcasting in wireless multihop networks. In *ICC*, pages 1–6. IEEE, 2009.