

Approximation Algorithms for Wireless Link Scheduling with SINR-based Interference

Douglas M. Blough, *Senior Member, IEEE*, G. Resta, and P. Santi *Member, IEEE*,

Abstract—In this paper, we consider the classical problem of link scheduling in wireless networks under an accurate interference model, in which correct packet reception at a receiver node depends on the signal to interference plus noise ratio (SINR). While most previous work on wireless networks has addressed the scheduling problem using simplistic graph-based or distance-based interference models, a few recent papers have investigated scheduling with SINR-based interference models. However, these papers have either used approximations to the SINR model or have ignored important aspects of the problem. We study the problem of wireless link scheduling under the exact SINR model, and present the first known true approximation algorithms for transmission scheduling under the exact model. We also introduce an algorithm with a proven approximation bound with respect to the length of the optimal schedule under primary interference. As an aside, our study identifies a class of “difficult to schedule” links, which hinder the derivation of tighter approximation bounds. Furthermore, we characterize conditions under which scheduling under SINR-based interference is within a constant factor from optimal under primary interference, which implies that secondary interference only degrades performance by a constant factor in these situations.

Index Terms—Wireless link scheduling, spatial-TDMA, SINR-based interference model, approximation algorithms.

I. INTRODUCTION

The classical problem of scheduling transmissions in multihop wireless networks, first studied in [12], has recently gained renewed interest from the networking research community, mainly because of its potential application in wireless mesh networks [2], [3], [15], [17], [21], [22], where the tight time synchronization between nodes needed to schedule wireless transmissions is deemed technically feasible. In fact, the 802.16 [1] standard for mesh networks is considering TDMA-based MAC implementation. Study of the transmission scheduling problem is motivated by the likelihood that good scheduling techniques can optimize mesh performance by maximizing throughput and improving fairness properties.

There are several different versions of the scheduling problem for multihop wireless networks. In one version [6], flows are given, flow demands are elastic, and the problem is to construct a set of routes, a bandwidth assignment, and a link schedule, over a sufficiently long scheduling interval, that maximize the total throughput (i.e. the sum of individual flow rates). In a second version of the problem [7], typically referred to as *link scheduling*, link demands are given and the problem is to construct a schedule of minimum length

that satisfies all demands. If new demands are placed on the network as soon as one set of demands is satisfied, minimizing schedule length will also maximize throughput. While our algorithms are formulated for the link scheduling problem, we discuss related work on SINR-based scheduling for the first version of the problem as well. Yet another version, which is less common, specifies end-to-end flow demands [6] instead of link demands. For this version, as for the first version mentioned above, routing and scheduling must be jointly considered. Given a routing algorithm, this version can be trivially converted to the link demand version. However, note that this version still differs from the first in that demands are fixed rather than elastic.

A crucial point to address when tackling any of the scheduling problems is to determine whether a certain set of wireless transmissions can occur in parallel without corrupting each other. This, in turn, requires using appropriate models for radio signal propagation and interference. In fact, in contrast to wired networks, wireless communications share the same (radio) communication channel, and interference between simultaneously transmitting links must be considered. The interference model has been shown to have a major impact on the complexity of optimal wireless link scheduling. In fact, while the first problem version is solvable in polynomial time under the primary interference model¹ [9], both problem versions are known to be NP-hard under models that account for secondary interference (e.g., graph-based models [20] and SINR-based models [7]). Furthermore, while algorithms with constant-factor approximation bounds have been developed for geometric graphs under some of the simpler secondary interference models, such as k-hop interference [20] and protocol interference [21], no approximation bounds have yet been proven for the most accurate SINR model.

Using accurate radio signal propagation and interference models is fundamental to ensure that the schedules computed by a certain algorithm do not lead to collisions in a practical scenario. This explains the efforts in the research community to derive near-optimal scheduling algorithms using increasingly accurate models. In particular, several interference models have been considered in the literature, including, most recently, the physical (or SINR-based) interference model, used for the first time in [8] to investigate asymptotical wireless network capacity. The basic assumption of the SINR-based interference model, i.e., that it is *cumulative interference power*, and not the *number of interfering signals*, which

D.M. Blough is with School of ECE, Georgia Institute of Technology, Atlanta, US.

G. Resta and P. Santi are with the Istituto di Informatica e Telematica del CNR, Pisa, ITALY.

¹In the primary interference model, two links conflict if and only if they share a common endpoint.

determines whether a packet is correctly received at the intended receiver, has been recently experimentally validated in a low-power wireless network [13]. Furthermore, in the same paper, the authors show that graph-based and distance-based interference models are inaccurate, and lead to throughput degradation with respect to SINR-based models. Thus, usage of the physical interference model for determining feasible transmission schedules is motivated by its practical relevance.

According to the physical interference model, a set of concurrent transmissions along links l_1, \dots, l_k do not corrupt each other if and only if the SINR (Signal to Interference + Noise Ratio) at each receiver is at least a certain threshold β , when all nodes at the transmitter end of links l_1, \dots, l_k are concurrently transmitting.

Designing and analyzing algorithms under this SINR-based model is especially difficult, since, contrary to what happens with simpler models such as graph-based [20] and protocol interference [8], interference under this model is not localized (i.e., even far-away interferers can potentially corrupt a transmission), and does not induce binary “conflict” relationships between links (e.g., transmission along a link l_1 might not corrupt a transmission along link l_2 by itself, but it might corrupt l_2 when a third transmission along link l_3 is also occurring). These intrinsic properties of the SINR model have hindered derivations of computationally efficient wireless link scheduling algorithms under this model. Only very recently have these problems been approached in the literature [3], [4], [6], [7], [14]. However, these prior works are either based on unrealistic assumptions (unbounded transmit power) [14], or use approximate SINR models in which the effect of noise [7] or of far-away interferers [3], [4] is ignored, or ignore certain “difficult to schedule” links [6]. Thus, to the best of our knowledge, the problem of designing computationally efficient wireless link scheduling algorithms with *proven, deterministic* approximation bounds under the *exact* SINR model remains open for both versions of the wireless scheduling problem.

Summary of contributions. In this paper, we present computationally efficient scheduling algorithms for the schedule-length-minimization version of wireless link scheduling under the assumption that nodes use the same, fixed, transmission power P . Through analysis of these algorithms, we prove, *for the first time under the exact SINR model*, deterministic approximation bounds on schedule length. By combining our main algorithm with algorithms for scheduling under the primary interference model in a novel way, we are also able to prove an approximation bound that is relative to the optimal schedule computed *under primary interference*. Since the optimal schedule under primary interference is, in general, shorter than the optimal schedule under SINR-based interference, this approximation bound is even stronger than approximation bounds relative to the optimal SINR-based solution. Based on near-optimal primary interference scheduling algorithms and, under certain conditions, constant factor approximation for our SINR scheduling algorithm, the combined algorithm produces schedules that are within a constant factor of the optimal schedule for primary interference under these conditions. This result shows that, under certain conditions, the impact of secondary interference on schedule length is minimal and that

efficient algorithms exist to derive schedules that are provably close to optimal primary interference schedules.

As a second major contribution of this work, we present insights into which factors affect the complexity of optimally scheduling links under the SINR model. In particular, we identify a paradox that we call the “*black-gray link paradox*”, which leads to the identification of a class of links that are ‘difficult to schedule’ (gray links). We show that gray links are directly responsible for problems that prior work had in producing approximation results for this problem. Gray links have an impact on our results, as well. As our analysis shows, the higher the fraction of gray links in the set of links to schedule, the looser the approximation bounds we are able to prove for our algorithms. In case no gray link occurs in the network, our algorithms are proven to be within a *constant factor from optimal* (Corollary 1). The same constant approximation bound can be obtained if nodes are deployed in a region of constant diameter, independently of the number of black/gray links to schedule (Theorem 4). Furthermore, if density of nodes per unit area is assumed to be constant, one of our algorithms is proven to be within a *constant factor from optimal under primary interference* (Corollary 2). To the best of our knowledge, ours is the first characterization of conditions under which wireless link scheduling under the SINR model asymptotically achieves the same performance as that achievable under primary interference. This result indicates that under some (reasonable) assumptions, the impact of secondary interference on network throughput is asymptotically marginal. The approximation bounds of the various algorithms considered in this paper, as well as of existing algorithms, are summarized later in Table I.

The rest of this paper is organized as follows. In Section II, we survey recent works related to our study. In Section III, we introduce notation and define the problem considered in the rest of the paper. In Section IV, we describe a phenomenon which we call the “black-gray link paradox”, which hinders the derivation of tight lower bounds to wireless link scheduling complexity under the exact SINR-based model. In Section V we present our approximation algorithms based on the exact SINR-based model, prove their properties in terms of performance bounds and time complexity, and discuss guidelines for distributed implementation. Section VI concludes.

II. RELATED WORK

The problem of transmission scheduling in multihop wireless networks has been deeply investigated in the literature following the seminal work by Nelson and Kleinrock [12]. It has been investigated either in isolation, known as link scheduling [7], or as part of the more general problem of investigating network capacity limits, which encompasses also finding optimal routes between source destination pairs, transmission power assignment, and so on. Herein, we focus primarily on the link scheduling version of the problem.

Given the shared nature of the communication medium in a wireless network, how to model interference between transmissions occurring on different, spatially separated links is a fundamental component of the network model, which considerably affects computational and algorithmic complexity. A

first fundamental distinction is between models that consider only *primary interference*, and those considering also effects of *secondary interference*. Primary interference, according to which two links interfere with each other if and only if they share a common endpoint, is a mandatory constraint to model single radio per node networks. On the other hand, secondary interference models a peculiar feature of wireless communications, i.e., usage of a shared communication medium, implying that two links can interfere with each other even if they do not share endpoints. Several models accounting for both primary and secondary interference have been considered in the literature, which can be roughly classified into graph-based (e.g., hop-based [20]), distance-based (e.g., the homogeneous protocol model of [8]), and SINR-based models (e.g., the physical model of [8]).

It is important to observe that there is a significant difference in the complexity of computing a minimum-length schedule, depending on whether secondary interference is considered or ignored: while this problem can be solved in polynomial time with primary interference only (which is equivalent to 1-hop, graph-based interference) [9], it becomes NP-hard when considering models that include secondary interference, e.g. SINR-based models [7].

A second important distinguishing feature of interference models is whether only local, pair-wise interference is considered, or all possible simultaneous communications occurring in the network are taken into account. Models that can be represented by a *conflict graph*, e.g. hop-based interference and protocol interference, belong to the former category, whereas SINR-based interference models belong to the latter. Recently, constant-factor approximation bounds have emerged for geometric graphs under several graph-based interference models, e.g. k-hop interference [20] and protocol interference [21]. However, the global nature of the SINR-based interference models challenges the design of algorithms with proven approximation bounds with respect to optimal, since under these models it is not possible to spatially divide the deployment area into smaller regions, with the property that link scheduling in different regions can be done independently. Furthermore, the global nature of SINR-based interference models hinders the design of localized, distributed scheduling algorithms based on these models.

Recently, a few papers have addressed some of the challenges related to SINR-based interference models.

In [14], Moscibroda and Wattenhofer derive upper bounds on the length of schedules built according to the exact SINR model, under the assumption that nodes can use arbitrarily high transmission power, which is unrealistic in practical scenarios. The works that are more closely related to ours are [3], [7], in which the authors present computationally efficient scheduling algorithms with proven approximation bounds under the assumption that nodes use the same, constant, transmission power. However, the bounds proven in [3] are very loose, hold only in a probabilistic sense under a random node distribution assumption, and are obtained for an SINR-based interference model in which interference from far-away transmitters is neglected. The interference model used in [7] is also an approximation of the SINR model, in which the

effect of noise is neglected. As we will thoroughly discuss in the following, the interference model considered in [7] is significantly different from the exact SINR model, and the approximation bound given in [7] does not hold in the exact model.

Another work which is closely related to ours is [6], in which the authors consider the throughput maximization version of the link scheduling problem under the exact SINR-based interference model. The authors present polynomial time algorithms that achieve a throughput within certain, non-trivial bounds from optimal. Interestingly, the approximation bounds are valid only under the assumption that the transmit power of nodes is slightly decreased (by a constant, multiplicative amount $(1 + \epsilon)$, where $\epsilon > 0$ is an arbitrary constant) with respect to the transmit power used in the original problem instance. A consequence of this is that certain links, including the so-called black and gray links that we have identified as critical, are ignored by the algorithm. Thus, the approximation bounds hold only when no such links are present in the network.

It is also worth mentioning two recent papers in which localized, distributed algorithms working under SINR-based interference models have been designed. In [4], Brar et al. present a localized, distributed implementation of the GreedyPhysical scheduling algorithm presented in [3], which thus achieves the same approximation bounds (under the approximate SINR interference model described above). In [19], Scheideler et al. present a localized, distributed algorithm for building a dominating set under the exact SINR interference model. Both algorithms make extensive use of physical carrier sensing to achieve local computation at the nodes, while guaranteeing algorithm operations under the global SINR interference model.

III. BACKGROUND

A. Problem Formulation

Given is a weighted communication graph $G = (V, E)$ where edge $e = (s, r) \in E$ represents a directed wireless link between a sender s and a receiver r in the network. Graph G is *geometric*, i.e., nodes in V can be thought of as points in the two-dimensional plane, and edge $e = (s, r) \in E$ implies that $d(s, r) \leq r_t$, where $d(s, r)$ is the Euclidean distance between nodes s and r . In other words, a directed edge between nodes s and r can exist only if r is within distance r_t from s , where r_t is the maximum transmission range (see below for a formal definition). Directed edges in E represent directed wireless communication links used to carry traffic between network nodes. For this reason, in the following we will use terms “edge” and “link” interchangeably.

Note that E in general is a *subset* of the set of all possible links between nodes in V . This is to model situations in which a routing algorithm has selected a subset of all possible communication links to actually carry traffic (e.g., a set of trees routed at the gateways in case of wireless mesh networks). Again, we stress that the focus of this paper is on optimal link scheduling, hence we assume that a routing algorithm has already been executed, and the set of links to schedule with relative demands (see below) is given.

For each edge $e \in E$, integer weight d_e represents the (current) traffic demand on link e . In case of time-varying traffic demands, we can think of periodically re-executing the presented algorithms in order to adapt to the new traffic distribution. In some places, we assume all edges have a weight of one. This is referred to as the *unit demand* case, as has been done in some previous work [7], [14]. The problem is then to construct a schedule, S_1, S_2, \dots, S_k , where edge e is assigned to exactly d_e of the S_i slots, the set of edges E_i assigned to slot S_i is a *feasible transmission set* with respect to the interference model under consideration, and k is as small as possible. A feasible transmission set is a set \bar{S} of wireless links for which, if all of nodes at the transmitter end of edges in \bar{S} are transmitting concurrently, all of the receivers will correctly receive their intended packets, according to some model of the interference caused by one communication on reception of another communication. We discuss interference models later in this section. A schedule S_1, S_2, \dots, S_k is *feasible* if and only if all the S_i 's are feasible transmission sets (w.r.t. the considered interference model). We are interested in constructing minimum-length feasible schedules, so as to maximize the overall throughput of the network. This is the problem considered, e.g., in [3], [7], [9], [14], [20]. Note that the dimension of our problem is the number n of links to schedule, i.e., $|E|$, and not the number of network nodes.

Summarizing, the problem considered in this paper is the following:

Input: Parameters α, β, N , and P . A set V of wireless nodes and for each $v \in V$, its location $l(v) \in \mathbb{R}^2$. A set E of links, where each $e = (u, v) \in E$ connects two nodes $u, v \in V$, where the Euclidean distance between u and v is no greater than the transmission range r_t (given by Equation (3) in the following). For each $e \in E$, its demand $d_e \in \{0, 1, \dots, m(E)\}$, where $m(E)$ is the maximum link demand.

Problem: Find a schedule $\mathcal{S} = (S_1, S_2, \dots, S_k)$ of k slots satisfying:

- 1) for $1 \leq i \leq k$, $S_i \subseteq E$,
- 2) every link $e \in E$ appears in d_e slots of \mathcal{S} , and
- 3) for $1 \leq i \leq k$ and for every $(u, v) \in S_i$, Inequality (1) (the SINR inequality) is satisfied at v when all u such that $(u, v) \in S_i$ are transmitting,

and having the minimum possible k .

B. Radio signal propagation

We adopt the classical model for radio signal propagation in wireless networks, which is referred to as the *log-distance path loss model*. In this model, the radio signal strength (power) at a distance d from the transmitter is given by P/d^α , where P is the transmission power and $\alpha > 2$ is the path loss coefficient [16] (the actual value of the constant α depends on the environment – e.g., indoor or outdoor)². Up to technical details, our results can be extended to more general radio propagation models that accounts for irregular

²Note that $\alpha > 2$ is a standard assumption in literature, justified by the fact that area (and, approximately, the number of interferers) grows as a power of 2, hence the only way of having bounded aggregate interference is to have radio signal decay larger than 2.

radio coverage area, such as the cost-based model proposed in [19], which is shown to closely approximate log-normal shadowing propagation. In the following, we assume all nodes use the same transmit power, an arbitrary constant P .

C. Interference Models

The simplest interference model for single radio, half-duplex wireless networks (i.e., networks in which each node is equipped with a single, half-duplex radio) is the primary interference model, according to which two links cannot transmit simultaneously if and only if they share an endpoint.

More accurate interference models consider also *secondary interference*, which accounts for the fact that all nodes in a wireless network share the same radio communication channel. In particular, in this paper we use the SINR interference model [8] (a.k.a. physical interference model), according to which the successful reception of a packet sent by node s and destined to node r depends on the Signal-to-Noise-plus-Interference-Ratio (SINR) at r . To be specific, denoting by $P_r(x)$ the received power at node r of the signal transmitted by node x , a packet along link (s, r) is correctly received if and only if:

$$\frac{P_r(s)}{N + \sum_{w \in V' - \{s\}} P_r(w)} \geq \beta, \quad (1)$$

where constant N is the background noise, V' is the subset of nodes in V that are transmitting simultaneously, and $\beta > 0$ is a constant threshold (the *SINR threshold*)³. In the SINR model, *every* concurrent transmission in the network (including transmissions between very distant nodes) must be explicitly considered when evaluating whether any single given transmission is successful.

Combining equation (1) with the formula for radio signal propagation, we have that a packet sent along link (s, r) is correctly received if and only if:

$$\frac{\frac{P}{d(s,r)^\alpha}}{N + \sum_{w \in V' - \{s\}} \frac{P}{d(w,r)^\alpha}} \geq \beta. \quad (2)$$

We stress that the SINR interference model with log-distance path loss, under the assumption that $\beta > 0$, implicitly accounts for primary interference. In fact, it is easy to see that the equation (2) cannot be satisfied in each of the following situations leading to primary interference: *i*) node u transmits simultaneously to two receiver nodes r_1 and r_2 ; *ii*) node u receives simultaneously from two sender nodes s_1 and s_2 ; *iii*) node u receives from node s while transmitting to node r . In case *i*), at both r_1 and r_2 we would have a SINR below 0dB (at least one interferer – node u itself – is at the same distance as the intended transmitter), which is not sufficient to correctly receive a packet. In case *ii*), at least the packet sent by the farther transmitter cannot be correctly received due to negative SINR value. Similarly, in case *iii*) at least the packet transmitted to node u cannot be correctly received due to negative SINR value – an interferer co-located

³In practice, the exact value of β depends on the desired data rate on the wireless channel, the modulation scheme, etc.

with the receiver. Note that the log-distance path loss model implies the use of omni-directional antennas and, hence, we do not consider the use of specific technologies such as directional antennas and MIMO, wherein primary interference could possibly occur without violating SINR constraints.

Equation (2) leads to the notion of *maximum transmission range*, denoted r_t , which is defined as the maximum distance up to which a packet can be correctly received *in absence of interference*, i.e., when set $V' - \{s\}$ is empty. Formally,

$$r_t = \sqrt[\alpha]{\frac{P}{\beta N}}. \quad (3)$$

In [7], the SINR model is considered without noise, i.e., it is assumed that $N = 0$ in Inequality (1). We refer to this as the SIR model. Observe that in the SIR model, a link has a positive budget up to *infinite* distances in absence of interference, i.e., the transmission range is infinite. This is a major difference w.r.t. to the SINR model in which the link budget is positive up to a *finite* distance even in absence of interference. As we shall see, this difference between the two models has a significant impact on the complexity of deriving efficient approximation algorithms for the scheduling problem. More specifically, the fact that the transmission range is finite in the SINR model gives rise to what we call the “black-gray link paradox”, which we carefully describe in the next section.

IV. THE BLACK-GRAY LINK PARADOX

Consider a transmission from a node s to a node r and assume that r is *exactly* at the border of s 's maximum transmission range (i.e., the SINR at node r in absence of interference – the SNR – is exactly β). In this situation, independently of the size of the deployment region, no other transmission concurrent with that along link (s, r) is possible, since even a very tiny contribution to the interference is sufficient to drive the SINR at r below β . Hence, if we call *black* those links such that the sender-receiver distance is exactly equal to r_t , then sequentially scheduling black links is the best that even the optimal algorithm can do (i.e., black links are “easy to schedule”). However, let us now consider a link (w, z) (call it a *gray* link) such that the SINR at z in absence of interference is $\beta + \epsilon$, for some arbitrarily small constant $\epsilon > 0$. If the deployment region is unbounded, then we can put a concurrent transmitter t_1 at a large enough distance from z in such a way that the SINR at z when t_1 is transmitting is some value $\beta + \epsilon'$, for some $0 < \epsilon' < \epsilon$. Since we can repeat this argument over and over, if the deployment region is unbounded, we can have an *infinite* number of transmissions going on in parallel to the transmission along link (w, z) without impairing correct message reception at z . We thus have the following apparent paradox, which we call the *black-gray link paradox*: if l is a black link, it must be scheduled sequentially (i.e., only one simultaneous transmission is possible) even if the deployment region is unbounded; on the other hand, if l is a gray link, then an infinite number of transmissions can potentially be scheduled in parallel with transmission along l (i.e., $\Omega(n)$ simultaneous transmissions are, in principle, possible). This

huge difference in potential concurrency comes despite an arbitrarily small difference in SNR values for black and gray links.

As we shall see, the “black-gray link paradox” hinders the derivation of tight lower bounds to the number of time slots needed to optimally schedule gray links. As a consequence, the quality of our derived approximation bounds for the wireless link scheduling problem under the SINR model is heavily affected by the black-gray link paradox. In particular, gray links are ‘difficult to schedule’: if few or no gray links are present, we can prove an $O(1)$ approximation bound for our scheduling algorithm; however, if relatively more gray links are present in the network, the approximation bound becomes looser; in the extreme case in which all the links to schedule are gray, we can prove only the trivial $O(n)$ approximation bound.

A possible way to limit the extent of the “black-gray link paradox” is to make some assumptions on the *size* of the deployment region. In particular, if the deployment region is assumed to have bounded diameter, we can prove an $O(1)$ approximation bound for the problem of optimally scheduling transmissions under the SINR model regardless of the number of gray links to schedule.

We close this subsection with a brief analysis of how the black-gray link paradox manifests itself also in the work presented in [6], which, similarly to our work, is based on the exact SINR interference model. In fact, the approximation bounds reported therein (which are of the form $O(g(L))$ for the case of homogeneous transmit power assignment, where $g(L)$, similarly to [7], is defined as the number of link length classes) are valid up to the following technical trick. Let $\mathcal{I} = (V, E, \mathcal{D}, P)$ be a problem instance, where V is the set of nodes, E is the set of possible links, \mathcal{D} is the vector of source/destination pairs, and P is the nodes transmit power⁴. Furthermore, let $\mathcal{I}'(\epsilon)$ be the ϵ power-reduced instance of \mathcal{I} (ϵ is an arbitrary constant > 0), i.e., the problem instance in which all nodes use transmit power $P/(1 + \epsilon)$, and the set $E' \subseteq E$ of all possible links is reduced accordingly. The authors of [6] show that the solution S of a certain linear program built on \mathcal{I}' is feasible also for the original problem instance \mathcal{I} , and that the throughput achieved by S is within a factor $O(g(L))$ from the optimal solution of \mathcal{I}' , not of the original instance \mathcal{I} . Technically speaking, the one presented in [6] is not an approximation bound, since no relationship between the optimal throughput for instance \mathcal{I} and that for instance \mathcal{I}' is proven. Note that the introduction of the $(1 + \epsilon)$ margin to the nodes transmission power is needed to keep possible communication links (the set E') sufficiently above the required threshold β for correct message reception, i.e., to *avoid occurrence of gray links*.

V. APPROXIMATION ALGORITHMS

A. Analysis of Algorithm GOW

Our first approximation algorithm, referred to as GOW*, is presented in the next subsection. GOW* builds on the

⁴We recall that the authors of [6] consider the more general problem of approximating network capacity, hence routes between source/destination pairs are not specified, but are part of the computed solution.




Fig. 1. 4-coloring used in step 6. of algorithm GOW*.

algorithm of [7], which is referred to as GOW in the rest of this paper. In order to understand GOW*, we first review the operation of Algorithm GOW, and we also evaluate GOW's performance under the true SINR model in light of the black-gray link paradox.

Algorithm GOW is designed for the case of unit link demands, i.e., $d_e = 1$ for each $e \in E$, and is based on the idea of partitioning links into length classes C_0, C_1, \dots , with the property that $\forall l_i, l_j \in C_h, \frac{1}{2} \leq \frac{\text{length}(l_i)}{\text{length}(l_j)} < 2$. For each such class C_i , a proper square cell partitioning of the (possibly infinite) deployment region (i.e., of the portion of the plane where nodes in V lie) is defined. The analysis of GOW in [7] ignores the effect of noise and thus, given that $r_t = \infty$ in the SIR model, the possible number of length classes is infinite as well. Cells are then 4-colored in such a way that no two adjacent cells have the same color (see Figure 1). A fundamental property of algorithm GOW is that, if the side of the square cells is chosen properly, links whose receivers are in different cells with the same color can be scheduled in the same slot without corrupting each other's transmissions. A crucial point in the derivation of the approximation factor reported in [7] is that the above property holds if the cell side for length class C_x is set to a *constant* multiple of the minimum length of links in C_x . The value of this constant is denoted μ . For each of the cell partitionings corresponding to the various length classes, the authors then prove an upper bound on the maximum number of feasible transmissions whose receivers are in the same cell, which is shown to be a certain *constant* $q(\mu)$ depending on μ . This upper bound is used to prove the approximation bound: the basic argument here is that algorithm GOW schedules all links whose receivers are in the same length class and same cell sequentially, while the optimal algorithm can schedule at most $q(\mu)$ of them in a single slot, resulting in an overall $O(g(L))$ approximation bound, where $g(L)$ is the number of length classes (which can be as high as $O(n)$ if link lengths grow as a geometric series).

To factor noise into the analysis, we first observe that links have a maximum possible length, which equals r_t , as defined in Equation 3. Another effect of considering noise is that the value of μ used to determine the cell partitioning is *no longer a constant*, but is instead a function $\mu(x)$ of the considered length class C_x . In particular, $\mu(x)$ is *unbounded* for the last

length class (the one including links whose length is within a factor $\frac{1}{2}$ from the transmission range). Since $\mu(x)$ for the last length class is unbounded, also the number $q(\mu(x))$ of feasible concurrent transmissions in the same cell becomes unbounded for this length class. On the other hand, algorithm GOW schedules *all* links in the last length class sequentially (because $\mu(x)$ is unbounded), implying a trivial $O(n)$ approximation bound (even if $g(L)$ is a constant) for the performance of the scheduling algorithm. Note that this last length class, which causes the problem with the approximation bound, contains all gray links.

In summary, due to the difficulty of scheduling gray links, the result of directly applying Algorithm GOW in the true SINR model is a trivial approximation bound. In the next subsection, we show how Algorithm GOW can be extended to handle gray links and provide non-trivial approximation bounds even in the true SINR model.

B. Algorithm GOW*

Similarly to [7], we assume unit link demands throughout this section. Algorithm GOW* is reported in Figure 2. There are two main differences between Algorithms GOW and GOW*: 1) links are grouped according to a SNR-based, instead of distance-based, criterion; and 2) links with the smallest SNR values (black and gray links according to the terminology introduced above) are treated separately.

Grouping links according to SNR instead of distance reflects the fact that what is relevant when scheduling a link is its strength (expressed in terms of the SNR), rather than the length. Furthermore, by combining usage of SNR-based link grouping with the exact SINR interference model we have that, independently of the actual link lengths, there are a *constant* number of classes to consider. This is because: *i*) the minimum possible SNR value for a feasible link is β , by the very same definition of SINR model; and *ii*) the maximum possible SNR value is P/N , corresponding to the situation in which the receiver is within a distance known as *near-field* [16] from the transmitter, where the received signal power is the same as the transmitted power⁵. It is worth noting that situations in which the minimum link length converges to 0 (e.g., when an increasing number of nodes is distributed uniformly at random in a region of unit area – see, e.g., the analyses of [3], [8]) are also handled by our model, since the maximum possible SNR value of P/N is assigned to arbitrarily short links.

Note that, under our working assumption of log-distance radio propagation with path loss exponent $\alpha > 2$, links in the k -th SNR class have length

$$D_{k+1} = \left(\frac{P}{(1+\epsilon)^{k+1}\beta N} \right)^{\frac{1}{\alpha}} < L_k \leq \left(\frac{P}{(1+\epsilon)^k\beta N} \right)^{\frac{1}{\alpha}} = D_k,$$

and that black and gray links are included in class $k = 0$ (this amounts to classifying as gray those links whose SNR is below $(1+\epsilon)\beta$ and greater than β , for a given value of $\epsilon \geq 1/7$).

⁵By fundamental laws of physics, the received signal power can be at most as large as the transmitted power.

Algorithm GOW*:

- Input:* A set E of n edges (links) with unit demand
 Parameter $\epsilon \geq 1/7$ that defines link classes
Output: A feasible schedule $S_1^0, \dots, S_{k_0}^0, S_1, \dots, S_t$
 under SINR model
1. $t = 1$
 2. Let $C = \{C_0, \dots, C_{\lfloor \log_{1+\epsilon}(P/\beta N) \rfloor} = \bar{k}\}$ be such that C_k is the set of links l_i such that $(1+\epsilon)^k \beta \leq \text{SNR}(l_i) < (1+\epsilon)^{k+1} \beta$ in E
 3. sequentially schedule all the links in class C_0 in slots $S_1^0, \dots, S_{k_0}^0$
 4. **for each** $C_k \neq \emptyset$, with $1 \leq k \leq \bar{k}$
 5. Partition network deployment region into squares of width $\mu \cdot D_{k+1}$
 6. 4-color the squares such that no two adjacent squares have the same color
 7. **for** $j = 1, \dots, 4$
 8. Select color j
 9. **repeat**
 10. For each square A of color j , choose a link $l_i \in C_k$ with receiver in A ; $L_j^k = L_j^k \cup \{l_i\}$
 11. $t = t + 1$; $S_t = L_j^k$
 12. **until** all links of C_k in selected squares are scheduled
 13. **return** $S_1^0, \dots, S_{k_0}^0, S_1, \dots, S_t$

Fig. 2. The GOW* Algorithm.

When considering links in class $k \geq 1$, the deployment region is divided into square cells of side μD_{k+1} , where constant μ is defined as follows⁶:

$$\mu = 2 \left(\frac{64(1+\epsilon)\beta(\alpha-1)}{\alpha-2} \right)^{\frac{1}{\alpha}}.$$

Cells in the same class are then 4-colored in such a way that no two adjacent cells have the same color (recall Figure 1). Then, at steps 7–12 links are greedily scheduled in successive slots, with the property that only links with the same color whose receivers are in different cells are assigned with the same slot.

We first prove that the schedule computed by Algorithm GOW* is feasible under the SINR model. We recall that usage of the SINR interference model ensures that links scheduled in the same slot are not subject to primary interference. Then, we show an upper bound to the length of the schedule constructed by our algorithm w.r.t. optimal.

Theorem 1: Assume that $\epsilon \geq \frac{1}{7}$. Then, the schedule computed by Algorithm GOW* is feasible under the SINR model.

Proof: We first observe that links in class C_0 are scheduled sequentially, thus the resulting slots are obviously feasible under the SINR model. Let us now consider a slots containing links in class C_k , for some $1 \leq k \leq \lfloor \log_{1+\epsilon}(P/\beta N) \rfloor = \bar{k}$. We now upper bound the interference experienced by a receiver r in a certain cell S in the partitioning obtained for class C_k . Once we focus on a receiver r in specific cell S , the cells containing receivers of the interfering links can be arranged in circumcentric square frames around S . The inner frame contains $3^2 - 1^2 = 8$ cells, the second frame contains

$5^2 - 3^2 = 16$ cells, and in general the h -th frame will contain $(2h+1)^2 - (2h-1)^2 = 8 \cdot h$ cells. The generic receiver contained in the h -th frame will be at least $(2h-1)\mu D_{k+1}$ apart from r . Considering that in class k all links have a length smaller than D_k , the minimum distance between r and a sender relative to frame h is $(2h-1)\mu D_{k+1} - D_k = (2h-1)\mu D_k / (1+\epsilon)^{1/\alpha} - D_k = D_k((2h-1)\mu(1+\epsilon)^{-1/\alpha} - 1)$. Hence, the total interference I_r experienced by r can be upper bounded by

$$I_r < \sum_{h=1}^{\infty} \frac{8h \cdot P}{D_k^\alpha \cdot ((2h-1)\mu(1+\epsilon)^{-1/\alpha} - 1)^\alpha} \quad (4)$$

$$\leq \frac{8P}{D_k^\alpha} \sum_{h=1}^{\infty} \frac{h}{(\frac{1}{2}(2h-1)\mu(1+\epsilon)^{-1/\alpha})^\alpha} \quad (5)$$

$$= \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu^\alpha D_k^\alpha} \sum_{h=1}^{\infty} \frac{h}{(2h-1)^\alpha} \quad (6)$$

$$\leq \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu^\alpha D_k^\alpha} \sum_{h=1}^{\infty} \frac{h}{(2h-h)^\alpha} \quad (7)$$

$$= \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu^\alpha D_k^\alpha} \sum_{h=1}^{\infty} \frac{1}{h^{\alpha-1}} \quad (8)$$

$$\leq \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu^\alpha D_k^\alpha} \cdot \frac{\alpha-1}{\alpha-2} \quad (9)$$

where (5) follows because $x-1 > x/2$ for $x > 2$ and indeed $(2h-1)\mu(1+\epsilon)^{-1/\alpha}$ is always greater than 2, and (9) follows from a known bound on Riemann's zeta function.

The SINR for the receiver r can thus be bounded by

$$\text{SINR}(r) \geq \frac{P}{I_r + N} = \frac{8 \cdot (1+\epsilon)^k}{(1+\epsilon)^k + 8} \cdot \beta \geq \beta$$

since $k \geq 1$ and $\epsilon \geq \frac{1}{7}$. ■

Definition 1: Given a problem instance $G = (V, E)$, the *maximal SNR density* Δ_{max} is the maximal number of receivers in a cell of class C_k , for some $1 \leq k \leq \lfloor \log_{1+\epsilon}(P/\beta N) \rfloor = \bar{k}$. Constant \bar{k} is called the *SNR diversity* of instance G .

We now prove an upper bound on the length of the schedule computed by Algorithm GOW*.

Theorem 2: The schedule computed by Algorithm GOW* has $O(|C_0| + \Delta_{max})$ length, where $|C_0|$ is the number of links in class C_0 .

Proof: Links in class C_0 are scheduled sequentially, which brings the $|C_0|$ term in the big-O notation. Links in class C_k , for $1 \leq k \leq \bar{k}$, whose receivers are in a cell of color, say, j , are scheduled in parallel if they are in different cells; hence, the number of slots needed to accommodate all links in class C_k is the number of receivers in the maximally occupied cell, which we denote Δ_{max}^k , times the number of colors. We then have that the total schedule length is upper bounded by $|C_0| + \sum_{k=1}^{\bar{k}} 4 \cdot \Delta_{max}^k \leq |C_0| + 4 \cdot \bar{k} \cdot \Delta_{max} \in O(|C_0| + \Delta_{max})$. ■

We are now ready to prove the approximation bound for Algorithm GOW*.

Theorem 3: Algorithm GOW* computes a schedule whose length is within a factor:

⁶We recall that we assume $\alpha > 2$, so constant μ is well defined.

- i) $O(|C_0|) = O(n)$ from optimal when $\Delta_{max} = 0$;
- ii) $O(\frac{|C_0|}{\Delta_{max}})$ from optimal when $\Delta_{max} > 0$ and $\Delta_{max} \in o(|C_0|)$;
- iii) $O(1)$ from optimal when $\Delta_{max} > 0$ and $\Delta_{max} \in \Omega(|C_0|)$.

Proof: We first consider case i). In this case, we can use the trivial lower bound 1 on optimal schedule length and Theorem 2 to conclude that Algorithm GOW* is within a factor $O(|C_0|)$ from optimal, which is $O(n)$ since $\Delta_{max} = 0$ implies that all links are in class C_0 .

Assume now that $\Delta_{max} > 0$. We start by proving an upper bound on the number of feasible transmissions with receivers belonging to the same cell S relative to length class C_k , for $1 \leq k \leq \bar{k}$. In particular, we prove that no more than

$$q = \frac{((1 + \epsilon)^{1/\alpha} + \sqrt{2}\mu)^\alpha}{\beta}$$

links with receiver belonging to the same cell can be scheduled together, where link class C_k as above is considered. The value of q is obtained by solving the following inequality

$$\frac{P/D_{k+1}^\alpha}{N + x \cdot P/(\sqrt{2}\mu D_{k+1} + D_k)^\alpha} < \beta \quad (10)$$

which leads to

$$x < (1 - \frac{1}{(1 + \epsilon)^{k+1}}) \frac{((1 + \epsilon)^{1/\alpha} + \sqrt{2}\mu)^\alpha}{\beta}$$

from which the above value of q is obtained. Inequality (10) comes from assuming the largest possible received power at the numerator, and the minimum possible contribution to interference from links whose receiver end is in cell S .

From the above fact, it follows that links belonging to a cell for which link occupancy is Δ_{max} are scheduled in at least $\frac{\Delta_{max}}{q}$ slots in the optimal schedule. This leads to an $\Omega(\Delta_{max})$ lower bound on the length of the optimal schedule. By Theorem 2, we have an approximation ratio with respect to optimal of $O(\frac{|C_0| + \Delta_{max}}{\Delta_{max}})$. In case ii), we have that $O(|C_0| + \Delta_{max}) = O(|C_0|)$, and the approximation bound becomes $O(\frac{|C_0|}{\Delta_{max}})$. In case iii), we have $O(|C_0| + \Delta_{max}) = O(\Delta_{max})$, and the approximation bound becomes $O(1)$. This concludes the proof of the theorem. ■

Corollary 1: If $|C_0| \in O(1)$, then the approximation bound of algorithm GOW* becomes $O(1)$.

As observed above, the GOW* approximation bound depends on the number of long links (black and gray links) in the network: the higher this number, the worse the approximation bound.

The relationship between number of black/gray links and quality of approximation bound lies in the fact that, when SNR is arbitrarily close to β , it is not possible to define a bounded grid size which is able to spatially separate links. In other words, when scheduling links whose SNR is arbitrarily close to β , the deployment region can be considered to be composed of a single cell comprising the whole region, and links have to be scheduled sequentially according to the GOW/GOW* approach. A possible way of improving the approximation bounds in presence of relatively many black/gray links would be to define more accurate scheduling policies within a single

cell, thus possibly allowing spatial reuse within the same cell. However, this would essentially entail re-designing the entire scheduling approach. On the other hand, we stress that in most practical situations the approximation bound is indeed accurate, since: i) the approximation bound becomes $O(1)$ independently of the number of black/gray links when the deployment region is bounded (see below); and ii) the number of black/gray links to schedule can be reduced by appropriate routing decisions and/or transmit power control (see below).

Given the relationship between number of black/gray links and approximations bounds, the network designer might want to select links with a relatively high SNR value for routing network traffic (i.e., to include in link set E), so that the performance of the scheduling algorithm becomes close to optimal. Note that wireless links are typically subject to link quality fluctuations, due to the time varying nature of radio signal propagation. Hence, selecting links with relatively high SNR values is very reasonable in practice, since they provide good transmission quality also in presence of radio signal fluctuations⁷. For instance, if we consider typical settings of $P = 100mW = 20dBm$, $N = -90dBm$, $\beta = 10dB$, and $\alpha = 3$, we have a nominal transmission range of $2.154Km$; if we set $\epsilon = 1/7$ when defining the SNR classes, we get an $O(1)$ approximation bound for GOW-PI by excluding links of length larger than $2.060Km$ from the set of links used to route traffic.

Another possible way of limiting the number of black/gray links is to use transmit power control. In principle, it is possible to slightly increase the transmit power used at some transmitters in such a way that the SNR at receivers is sufficiently above β , so to reduce/avoid the occurrence of black/gray links. Note that, up to properly re-defining SNR link classes and constant μ , all the results presented in the paper remain valid also when transmitters use different transmit powers, as long as the ratio between the highest and the lowest transmit power used by the transmitters is $O(1)$ (i.e., an arbitrary constant).

We now prove a stronger approximation bound for the case where network nodes V are deployed in a region R with bounded diameter.

Theorem 4: Assume nodes in V are deployed in a region R of constant, bounded diameter $diam(R)$. Then, Algorithm GOW* computes a schedule whose length is within a factor $O(1)$ from optimal, regardless of the number of links in class C_0 .

Proof: To prove the theorem it is sufficient to show that, under the assumption that $diam(R)$ is a constant, at most a constant number q' of links in class C_0 can be scheduled concurrently by the optimal algorithm. From this observation, we have a lower bound of $\Omega(|C_0| + \Delta_{max})$ on the optimal schedule length, and the theorem follows by Theorem 2.

The above fact can be easily shown by solving inequality

$$\frac{P/D_1^\alpha}{N + q' \cdot P/diam(R)^\alpha} < \beta$$

⁷Time varying radio signal propagation is not accounted for in the log-distance path loss model, which can be thought of as representing the time averaged intensity of the radio signal at a certain distance from transmitter.

where $diam(R)$ is the maximum distance between a pair of nodes in R , which results in $q' > diam(R)^\alpha \cdot N/P$, which is a constant under the theorem assumptions. ■

To the best of our knowledge, the bounds proved in Theorems 3 and 4 are the first non-trivial, deterministic approximation bounds for the problem of wireless link scheduling under the SINR model.

Before ending this section, we formally prove that GOW* has polynomial time complexity.

Theorem 5: Algorithm GOW* has $O(n \cdot \bar{k} \cdot \#C \cdot \Delta_{max})$, where $\#C$ is the maximum number of cells in a partitioning of the deployment region computed at step 5. of the algorithm.

Proof: The outer **for** cycle (steps 4.–12.) is executed \bar{k} times. At each iteration, for each of the 4 colors considered, all the squares of current color are scanned ($O(\#C)$ operations), and possibly a link is selected in each square ($O(\Delta_{max})$ operations). This **repeat-until** cycle is repeated $O(|E| = n)$ times. Hence, GOW* time complexity is $O(n \cdot \bar{k} \cdot \#C \cdot \Delta_{max})$. ■

C. Approximation Bounds with respect to Primary Interference Model

In this subsection, we present a two-phase approach to prove approximation bounds with respect to optimal scheduling under the primary interference model. We recall that a transmission set $S = \{l_1, \dots, l_h\}$ is feasible under the primary interference model if and only if no two links share a common endpoint. Furthermore, throughout this subsection we allow arbitrary integer link demands.⁸

Proving approximation bounds with respect to optimal under primary interference is important to isolate the effects of radio interference, which, we recall, is responsible for secondary interference, on the wireless link scheduling problem. In particular, in this section we will prove that, under some (reasonable) conditions, accounting for secondary interference when building the schedule incurs no schedule length increase with respect to the primary interference model (in an asymptotic sense). Thus, our results indicate that, in some conditions, radio interference has no (asymptotic) effect on schedule length, and the wireless link scheduling problem can be accurately characterized in asymptotic terms even using the simple primary interference model. It is also worth observing that, since the optimal schedule under primary interference cannot be longer than that under SINR interference, the approximation bounds proved in this section can be considered as relatively stronger than those reported in Section V-B and in existing literature on the SINR-based interference model.

The reader might wonder why we do not directly use algorithm GOW*, which is able to produce schedules free of primary interference, to derive performance bounds with respect to optimal scheduling under primary interference. The difficulty here lies in the fact that, by directly applying algorithm GOW*, it is not easy to relate the length of the optimal

⁸We only assume that the sum of the link demands remains polynomial in the number of links, so that when expanding the demand graph to a multigraph with unit demands, the problem size remains polynomial in the number of links.

Algorithm PRIMARYSINR:

Input: A set $G = (V, E)$ of links to schedule
 A weight $d_i > 0$ associated with each link $l_i \in E$
 A scheduling algorithm \mathcal{A}_p for primary interference
 A scheduling algorithm \mathcal{A}_s for SINR interference
Output: A feasible schedule \mathcal{S}_{SINR} under SINR model
 Phase1: Use \mathcal{A}_p to compute schedule under primary interference
 Let $\mathcal{S}_{Pr} = \{S_1, \dots, S_l\}$ be the resulting schedule
 Phase2:
 1. $\mathcal{S}_{SINR} = \emptyset$
 2. **for** $h = 1, \dots, l$
 3. Execute algorithm \mathcal{A}_s on links in S_h
 4. Let S_h^1, \dots, S_h^k be the transmission slots computed by \mathcal{A}_s
 5. $\mathcal{S}_{SINR} = \mathcal{S}_{SINR} \cup S_h^1 \cup \dots \cup S_h^k$
 6. **return** \mathcal{S}_{SINR}

Fig. 3. Algorithm PRIMARYSINR

schedule under primary interference to the upper bound on the length of the schedules produced by GOW* (which accounts for both primary and secondary interference).

To get around this difficulty, we present a two-phase approach, which we call PRIMARYSINR. The key idea is to separate primary and secondary interference contributions when building the schedule: in the first phase, links are scheduled according to some algorithm \mathcal{A}_p considering only primary interference as the feasibility criterion; then, transmission sets corresponding to each slot computed by \mathcal{A}_p are considered, and are possibly split into several slots in order to satisfy also requirements of the SINR interference model. For this purpose, a second scheduling algorithm \mathcal{A}_s is considered. This two-phase approach allows us to easily derive bounds with respect to optimal schedule length under primary interference, if approximation bounds for \mathcal{A}_p and upper bounds on the length of the schedule built by \mathcal{A}_s are known. The PRIMARYSINR scheduling approach is summarized in Figure 3.

The following Proposition shows how to turn approximation bounds for \mathcal{A}_p and upper bounds on schedule length for \mathcal{A}_s into an approximation bound for PRIMARYSINR with respect to the optimal schedule under primary interference.

Proposition 1: Let \mathcal{A}_p be a scheduling algorithm for primary interference with approximation bound $O(f(n))$ with respect to optimal, and let \mathcal{A}_s be a scheduling algorithm for SINR interference that builds schedules of length at most $O(g(n))$, for some functions $f(n), g(n)$. Algorithm PRIMARYSINR builds a schedule whose length with respect to optimal under primary interference is upper bounded by $O(f(n) \cdot g(n))$.

Proof: It is sufficient to observe that the number of slots computed by \mathcal{A}_p is at most $f(n) \cdot S_{opt}$, where S_{opt} is the length of the optimal schedule under primary interference, and that each of these slots is divided into at most $O(g(n))$ slots by algorithm \mathcal{A}_s in the second phase of the algorithm. Hence, the total length of schedule computed by PRIMARYSINR is within a $O(f(n) \cdot g(n))$ factor from the optimal schedule under primary interference. ■

An immediate consequence of Proposition 1 is that, when-

ever \mathcal{A}_p is an optimal algorithm for primary interference scheduling and \mathcal{A}_s produces a schedule of $O(1)$ length, the PRIMARYSINR approach builds a feasible schedule under SINR interference whose length is within a constant factor from the length of the optimal schedule under primary interference. In particular, if both \mathcal{A}_p and \mathcal{A}_s are polynomial time algorithms, the schedule can be computed efficiently and used, e.g., for benchmarking more practical SINR interference scheduling approaches such as GreedyPhysical [3]. This brings us to the question of whether such algorithms for primary and SINR interference scheduling exist.

We first observe that building an optimal schedule under primary interference consists in finding a minimum edge coloring of the multi-graph obtained from G by replacing each edge $e \in E$ with demand d_e with d_e multi-edges between the same pair of nodes. Thus, an edge-coloring algorithm for multi-graphs can be used in phase 1 of PRIMARYSINR. The minimum edge coloring problem is known to be NP-hard even for simple graphs [11]. However, the optimal solution can be computed in polynomial time if the graph is bipartite⁹. If graph G is not bipartite, a $4/3$ approximation of the optimal schedule can be computed using, for instance, the algorithm of [10], yielding the same asymptotical approximation bounds as in the case where graph G is bipartite.

On the other hand, Theorem 2 states that the length of the schedules computed by SINR interference scheduling algorithm GOW* is $O(|C_0| + \Delta_{max})$, which implies that GOW* builds schedules of $O(1)$ length whenever the number of black/gray links to schedule is $O(1)$, and we have a constant node density per unit area. Thus, we can conclude this section with the following corollary:

Corollary 2: Assume G is an arbitrary graph with integer link demands, \mathcal{A}_p is a polynomial time edge-coloring algorithm for multi-graphs that is at most a constant factor from optimal, and \mathcal{A}_s is the GOW* algorithm. Furthermore, assume the number of black/gray links in the network is $O(1)$, and the node density per unit area is constant. Under these assumptions, algorithm PRIMARYSINR builds in polynomial time a schedule which is feasible under SINR interference, and whose length is within a factor $O(1)$ from the optimal schedule under primary interference.

Proof: The proof is an immediate implication of Proposition 1, Theorem 2, and of the assumptions on link length and node density. ■

To the best of our knowledge, this is the first characterization of conditions under which wireless scheduling under SINR interference achieves a performance comparable to that under primary interference presented in the literature. This result is very interesting, since it shows that, under certain conditions, secondary interference reduces network throughput by at most a constant factor with respect to primary interference. It is also worth observing that the conditions for the above to happen are indeed satisfied in most realistic application scenarios.

⁹Note that the multi-graph to edge color is bipartite if and only if the original graph G is bipartite.

D. Summary of approximation bounds

In this subsection, we summarize the approximation bounds of the various scheduling algorithms introduced in the previous sections. We first observe that algorithms for unit link demands (i.e., GOW and GOW*) can be extended to the case of arbitrary integer link demands by replacing a link with demand d_e into d_e copies of the same link with unit demand. The resulting approximation bounds are $O(g(L))$ for algorithm GOW and $O(\frac{|C_0^d| + \Delta_{max}^d}{\max\{1, \Delta_{max}^d\}})$ for algorithm GOW*, where $|C_0^d|$ is the sum of the demands on links in class C_0 , and Δ_{max}^d is the maximal sum of demands of links whose receivers are in the same cell.

The approximation bounds of the various algorithms considered in this paper are summarized in Table I. The bounds for the PRIMARYSINR approach refer to the case in which a constant factor to optimal approximation algorithm for primary interference scheduling is used, and GOW* is used as the SINR interference scheduling algorithm. In the table, ‘interference model’ refers to the model used to build the schedule, while ‘reference model’ refers to the interference model used to build the optimal schedule to which performance of the considered algorithm is compared.

In the table, we have reported a generalized version of the approximation bound for PRIMARYSINR stated in Corollary 2 that can be easily derived from Theorem 2. This bound holds for networks not necessarily satisfying the link length and node density conditions stated in the corollary, and is defined as $O(\Psi)$, where Ψ is defined as $\max_{h=1, \dots, d(E)} \{|C_0^h| + \Delta_{max}^h\}$, $|C_0^h|$ is the number of links in slot S_h belonging to class C_0 , and Δ_{max}^h is the maximum number of links in slot S_h whose receivers are in the same cell (for some link class C_x , with $x \geq 1$).

It is worth observing that ours are the only bounds using the exact SINR model to build the schedule, and possibly using primary interference as a reference model. While bounds of the various algorithms depend on specific quantities defined therein (e.g., number $g(L)$ of link length classes for GOW, or maximal SNR density for GOW* and PRIMARYSINR), all the considered bounds (excluding those for GREEDYPHYSICAL) have the property that they can be as tight as $O(1)$ in some conditions, but they can be as loose as the trivial $O(n)$ bound in worst-case conditions. The design of scheduling algorithms based on SINR interference models with provable, non-trivial worst-case bounds is, to the best of our knowledge, still open.

E. Distributed Implementation of Scheduling Algorithms

For some types of wireless multihop networks, centralized operation of scheduling algorithms is quite realistic. For example, in small-to-medium sized mesh networks, gateway nodes serve as natural centralization points for collecting network information, computing schedules, and disseminating the schedules to the other network nodes. However, large mesh networks and other types of networks might require distributed scheduling algorithms. Here, we briefly discuss the suitability of Algorithms GOW* and PRIMARYSINR for distributed execution.

Algorithm	Approx. factor	Interference model	Reference model	Networks	Link demands
GREEDYPHYSICAL [3]	$O(n^{1-c(\alpha)}(\log n)^{c(\alpha)})$ w.h.p., $0 < c(\alpha) < 1$	bounded SINR	bounded SINR	random	integer
GOW [7]	$O(g(L))$	SIR	SIR	arbitrary	unit/integer
GOW*	$O(\frac{ C_0 + \Delta_{max}}{\max\{1, \Delta_{max}\}})$	SINR	SINR	arbitrary	unit
	$O(1)$ – bounded diameter	SINR	SINR	arbitrary	unit
	$O(\frac{ C_0^d + \Delta_{max}^d}{\max\{1, \Delta_{max}^d\}})$	SINR	SINR	arbitrary	integer
PRIMARYSINR	$O(\Psi)$	SINR	primary int.	arbitrary	integer
	$O(1)$ – no black/gray links, constant density	SINR	primary int.	arbitrary	integer

TABLE I
COMPARISON OF THE SCHEDULING ALGORITHMS.

For Algorithm GOW*, many of the parameters can be determined prior to, or at the beginning of, deployment. The length classes for the links, along with the associated cell sizes, are dependent on the transmission power, P , the SINR threshold, β , the noise floor, N , and the path loss exponent, α . Parameters P and β are functions of the devices and technology used in the network, while N and α can be easily measured in the deployment environment [5]. Prior to network operation, these parameters can be distributed to all nodes along with the location of a fixed reference point so that every node can calculate the grids used for each length class. We also assume that nodes know their approximate locations, so that they know in what cell they reside for every length class grid. Location could be determined from built-in GPS, from GPS measurements done during deployment, or through estimation techniques such as triangulation from known points. The grid colorings can also be done in a deterministic way so that every node can reproduce them.

The remaining values that need to be known prior to schedule construction are the link demands. These can be estimated periodically by sampling queue lengths at transmitters and, for distributed execution need only to be communicated from the transmitter of every link to the associated receiver, since scheduling is done at the receiver side. For every cell within every length class, all nodes that serve as a receiver for at least one link in the class must coordinate to determine a schedule. This can be done in several ways. One possibility is to elect a cell leader, send all of the link demands for the cell to the leader, and then have the leader calculate and distribute the cell schedule to all nodes in the cell. The nodes who are receivers for scheduled links would then simply communicate the schedule slots to their associated transmitters. It is also possible to have every receiver broadcast its demand to all nodes in the cell and then have every node independently calculate in a deterministic fashion the same cell schedule. The important points to emphasize are that: (1) coordination is limited to within individual cells, so there is no global network-wide communication performed, and (2) the number of length classes is constant, so the number of grids for which these distributed scheduling computations must be performed does not increase with the size of the network. As a result of these two points, we believe that distributed implementation of Algorithm GOW* is realistic, even for quite large networks.

VI. DISCUSSION AND CONCLUDING REMARKS

In this paper, we have investigated the problem of wireless link scheduling under the exact SINR interference model.

First, we have shown that even tiny contributions to the SINR at the receivers (e.g., the noise, or the interference from far-way transmitters) cannot be ignored when building the schedule, nor when evaluating the approximation bounds. In view of this observation, we have provided the first known algorithm for wireless link scheduling with proven, deterministic approximation bounds under the SINR model, and identified a class of links that are “difficult to schedule” and hinder the derivation of tighter approximation bounds. We have also introduced an algorithm with proven approximation bounds with respect to primary interference, and identified conditions under which scheduling under SINR interference is within a constant factor from optimal under primary interference. These conditions might serve as a guideline in the design of wireless networks (e.g., avoid routing along links whose SINR value is close to the threshold β for correct message reception).

The study reported in this paper leaves several avenues open for further research on this intriguing problem. In particular, the problem of finding better lower bounds on the length of optimal schedules for the class of “difficult to schedule” links, which might lead to the derivation of non-trivial *worst-case* approximation bounds, remains open. In view of this, it is interesting to observe that the “black-gray link paradox” is a consequence of using a thresholded interference model, according to which the packet reception rate on a link is 100% if the SINR value at the receiver is $\geq \beta$, and it is 0 if the SINR is even slightly below β . Indeed, such a sharp SINR threshold for correct packet reception is unlikely to occur in practical settings, where transition between near 0 packet reception rates and near 100% rates spans a few dBs (see, e.g., [13]): hence, there exists a gray SINR area in which packet reception is still possible, although with a rate significantly below 100% (and above 0). A promising direction for future work is studying the wireless link scheduling problem using a non-thresholded SINR-based interference model, and proving approximation bounds under this model. Usage of non-thresholded SINR-based interference model has the potential to improve (expected) throughput (e.g., 30% throughput improvement with respect to thresholded SINR model has been observed in [13] in an experimental testbed), while at the same time (possibly) countering the “black-gray link paradox” described in this paper. Initial steps in this direction can be found in [18].

REFERENCES

- [1] <http://iee802.org/16/>
- [2] M. Alicherry, R. Bhatia, L. Li, “Joint Channel Assignment and Routing for Throughput Optimization in Multi-Radio Wireless Mesh Networks”, *Proc. ACM Mobicom*, pp. 58–72, 2005.

- [3] G. Brar, D. Blough, and P. Santi, "Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks," *Proc. ACM MobiCom*, pp. 2–13, 2006.
- [4] G. Brar, D. Blough, and P. Santi, "The SCREAM Approach for Efficient Distributed Scheduling with Physical Interference in Wireless Mesh Networks", *Proc. IEEE ICDCS*, pp. 214–224, 2008.
- [5] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network," *Proceedings of the 4th ACM International Conference on Mobile Systems, Applications and Services* 2006.
- [6] D. Chafekar, V.S. Anil Kumar, M.V. Marathe, S. Parthasarathy, A. Srinivasan, "Approximation Algorithms for Computing Capacity of Wireless Networks with SINR Constraints", *Proc. IEEE Infocom*, pp. 1166–1174, 2008.
- [7] O. Goussevskaia, Y.V. Oswald, R. Wattenhofer, "Complexity in Geometric SINR", *Proc. ACM MobiHoc*, pp. 100–109, 2007.
- [8] P. Gupta and P.R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, Vol. 46, No. 2, pp. 388–404, 2000.
- [9] B. Hajek, G. Sasaki, "Link Scheduling in Polynomial Time", *IEEE Trans. on Information Theory*, vol. 34, n. 5, Sept. 1988.
- [10] D.S. Hochbaum, T. Nishizeki, D.B. Shmoys, "A Better than "Best Possible" Algorithm to Edge Color Multigraphs", *Journal of Algorithms*, Vol. 7, pp. 79–104, 1986.
- [11] I. Holyer, "The NP-completeness of Edge Coloring", *SIAM Journal of Computing*, Vol. 10, n. 4, pp. 718–720, 1981.
- [12] R. Nelson and L. Kleinrock, "Spatial-TDMA: A Collision-free Multihop Channel Access Protocol," *IEEE Transactions on Communication*, Vol. 33, pp. 934–944, Sept. 1985.
- [13] R. Maheshwari, S. Jain, S.R. Das, "A Measurement Study of Interference Modeling and Scheduling in Low-Power Wireless Networks", *Proc. 6th ACM Conference on Embedded Networked Sensor Systems (ACM SenSys 2008)*, pp. 141–154, 2008.
- [14] T. Moscibroda, R. Wattenhofer, A. Zollinger, "Topology Control Meets SINR: The Scheduling Complexity of Arbitrary Topologies", *Proc. ACM MobiHoc*, pp. 310–321, 2006.
- [15] B. Raman, K. Chebrolu, "Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks" *Proc. ACM Mobicom*, pp. 156–169, 2005.
- [16] T.S. Rappaport, *Wireless Communications*, Prentice Hall, Upper Saddle River, NJ, 2002.
- [17] T. Salonidis, L. Tassioulas, "Distributed Dynamic Scheduling for End-to-End Rate Guarantees in Wireless Ad Hoc Networks", *Proc. ACM MobiHoc*, pp. 145–156, 2005.
- [18] P. Santi, R. Maheshwari, G. Resta, S. Das, D. Blough, "Wireless Link Scheduling under a Graded SINR Interference Model, *Proc. ACM Int. Workshop on Foundations of Wireless Ad Hoc and Sensor Networking and Computing (FOWANC)*, pp. 3–12, 2009.
- [19] C. Scheideler, A. Richa, P. Santi, "An $O(\log n)$ Dominating Set Protocol for Wireless Ad Hoc Networks under the Physical Interference Model", *Proc. ACM MobiHoc*, pp. 91–100, 2008.
- [20] G. Sharma, R. Mazumdar, N. Shroff, "On the Complexity of Scheduling in Wireless Networks", *Proc. ACM Mobicom*, pp. 227–238, 2006.
- [21] W. Wang, Y. Wang, X.-Y. Li, W.-Z. Song, O. Frieder, "Efficient Interference-Aware TDMA Link Scheduling for Static Wireless Networks", *Proc. ACM Mobicom*, pp. 262–273, 2006.
- [22] Z. Wu and D. Raychaudhuri, "D-LSMA: Distributed Link Scheduling Multiple Access Protocol for QoS in Ad-hoc Networks," *Proc. IEEE Globecom*, pp. 1670–1675, 2004.

PLACE
PHOTO
HERE

Douglas M. Blough received the B.S. degree in electrical engineering and the M.S. and Ph.D. degrees in computer science from The Johns Hopkins University, Baltimore, MD, in 1984, 1986, and 1988, respectively. Since Fall 1999, he has been Professor of Electrical and Computer Engineering at the Georgia Institute of Technology, where he also holds a joint appointment in the College of Computing. From 1988 to 1999, he was on the faculty of Electrical and Computer Engineering at the University of California, Irvine. Dr. Blough was Program Co-Chair for the 2000 International Conference on Dependable Systems and Networks (DSN) and the 1995 Pacific Rim International Symposium on Fault-Tolerant Systems. He has been on the Program Committees of numerous other conferences, was Associate Editor for IEEE Transactions on Computers from 1995 through 2000, and was Associate Editor for IEEE Transactions on Parallel and Distributed Systems from 2001 through 2005. His research interests include distributed systems, dependability and security, and wireless multihop networks.

PLACE
PHOTO
HERE

G. Resta received the MS degree in computer science from the University of Pisa, Italy, in 1988. In 1996, he became a researcher at the Istituto di Matematica Computazionale of the Italian National Research Council (CNR), Pisa. He is now a senior researcher at the Istituto di Informatica e Telematica (CNR) in Pisa. His research interests include computational complexity (especially in relation to linear algebra problems), parallel and distributed computing, and the study of structural properties of wireless ad hoc networks.

PLACE
PHOTO
HERE

P. Santi received the Laura Degree and Ph.D. degree in computer science from the University of Pisa in 1994 and 2000, respectively. He has been researcher at the Istituto di Informatica e Telematica del CNR in Pisa, Italy, since 2001, and he has now Senior Researcher at the same Institute. During his career, he visited Georgia Institute of Technology in 2001, and Carnegie Mellon University in 2003. His research interests include fault-tolerant computing in multi-processor systems (during PhD studies), and, more recently, the investigation of fundamental properties of wireless multihop networks such as connectivity, lifetime, capacity, mobility modeling, and cooperation issues. He has contributed more than 40 papers and a book in the field of wireless ad hoc and sensor networking, he is Associate Editor of IEEE Trans. on Mobile Computing, he has been General Co-Chair of ACM VANET 2007 and 2008, and he is involved in the organizational and technical program committee of several conferences in the field. He is a member of IEEE CS and a senior member of ACM and SIGMOBILE.