Scheduling Wireless Links in the Physical Interference Model by Fractional Edge Coloring

Guilherme I. Ricardo, José F. de Rezende, and Valmir C. Barbosa

***Abstract*—We consider the link scheduling problem in wireless mesh networks for capacity maximization. Unlike all previous approaches, ours views capacity in terms of covering a graph's edges by matchings that are feasible in the sense of the physical interference model, thus leading to a novel fractional edge-chromatic indicator of the graph. Network capacity is then defined as a function of this indicator, whose exact computation can be formulated as a linear programming problem. Solving this problem for an extensive collection of random geometric graphs that instantiate networks in the physical interference model reveals that, depending on node density, often a network admits increased capacity by virtue of simply scheduling links based on fractional, rather than integer, edge coloring. This settles a heretofore unaddressed issue in the physical interference model.**

***Index Terms*—Wireless mesh networks, physical interference model, link scheduling, capacity maximization, fractional edge coloring.**

I. INTRODUCTION

WE CONSIDER a set of nodes operating in a wireless mesh network under the constraints imposed by the physical interference model of wireless communication [1]. These nodes are interconnected by a set L of links, each link $e \in L$ being characterized by a sender node s_e and a receiver node r_e. Any node may in principle act either as sender or as receiver, depending on the links in which it participates. When all links in a set $S \subseteq L$ are concomitantly active, the ability of receiver r_e to decode what it receives from sender s_e for any given link $e \in S$ is constrained by the signal-to-interference-plus-noise ratio (SINR) that results from the combined activity of the group, given by

$$\mathrm{SINR}(e, S) = \frac{P/d^{\alpha}_{s_e r_e}}{\gamma + \sum_{f \in S \setminus \{e\}} P/d^{\alpha}_{s_f r_e}}. \tag{1}$$

In this expression, P is a node's transmission power (assumed the same for all nodes), γ is the noise floor, d_{ab} is the Euclidean distance between nodes a and b, and $\alpha > 2$ is used to determine how power decays away from the transmitter with the distance to it. The SINR constraint operates by affecting the so-called feasibility of set S. We say that a nonempty $S \subseteq L$ is *feasible* if no two of its links share a node and $\mathrm{SINR}(e, S) \geq \beta$ for all $e \in S$, where $\beta > 1$ is a parameter related to a receiver's decoding capabilities, assumed the same for all receivers. Every singleton $\{e\} \subseteq L$ is assumed feasible.

The problem of maximizing network capacity, broadly understood as the rate of effective communication among nodes, is closely related to that of scheduling the links in L for operation. This, in turn, is often posed under the so-called physical interference model (in which the SINR constraint is fully taken into account) but sometimes assumes only the constraints imposed by the so-called protocol-based interference model (which depend essentially on graph-based distances). Solving the link-scheduling problem has given rise to numerous proposals, some approaching the scheduling problem by itself [2]–[9], some in conjunction with others (such as power control and end-to-end routing, for example) [10]–[12]. Many of these proposals are formulated within the framework of spatial time-division multiple access (STDMA), which divides time into slots and reduces the scheduling problem to the selection of which links to activate simultaneously in each one. The type of selection strategy that is by far the most adopted asks that a sequence $\mathcal{S} = \langle S_1, S_2, \ldots, S_T \rangle$ be determined for some $T > 1$. In this sequence, each S_i is a subset of the link set L complying with the constraints imposed by the interference model in use and moreover ensuring that each $e \in L$ appears in exactly one of the T subsets. Any approach to maximize network capacity by solving the link-scheduling problem (as in [1], [2], [8], [9]) will therefore seek to minimize T (maximize $1/T$) without violating any of these constraints. Once a solution is available, repeating the sequence $\mathcal{S}$ guarantees interference-free communication for as long as needed.

II. OPEN ISSUES AND CONTRIBUTIONS

One common abstraction for reasoning about such proposals is that of graph coloring or related notions. In terms of the formulation outlined above, clearly the links in each set S_i, to be scheduled for operation in the same time slot, can be regarded as being assigned the same color (the ith of the T available colors) if we interpret the conditions for membership in S_i in the context of graph coloring. Whether it is vertex coloring or edge coloring that is being considered depends on how the graph in question is set up to represent how the various links relate to each other given the interference model at hand. In either case, once the least possible value of T is found (or approximated), each vertex or edge ends up having exactly one color.

Manuscript received November 2, 2019; revised December 11, 2019; accepted December 17, 2019. Date of publication December 20, 2019; date of current version April 9, 2020. This work was supported in part by Conselho Nacional de Desenvolvimento Científico e Tecnológico, in part by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior, and in part by the BBP grant from Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro. The associate editor coordinating the review of this article and approving it for publication was K. Yamamoto. *(Corresponding author: Valmir C. Barbosa.)*

The authors are with the Systems Engineering and Computer Science Program, Centro de Tecnologia, Federal University of Rio de Janeiro, Rio de Janeiro 21941-914, Brazil (e-mail: valmir@cos.ufrj.br).

Digital Object Identifier 10.1109/LWC.2019.2961361

To the best of our knowledge, the decades-long effort to come up with capacity-optimizing strategies for link scheduling has almost completely failed to recognize that such a single-color abstraction is inherently limited and may in many cases fall short of leading to as much network capacity as possible. The exceptions we know of are only three and separated by many years. The earliest one is based on the coloring of a graph's edges and adopts what would pass for the protocol-based interference model [13]. The other two are much more recent and both based on the coloring of vertices, the earliest one given for the protocol-based interference model as well [12], the latest for the physical interference model [9]. (Further exceptions exist related to models of physical interference other than the one we assume, as for example the form of node coloring used in index coding [14].)

These three proposals use the fractional variety of graph coloring. In terms of STDMA, they let each link appear in exactly $q \geq 1$ of the T sets in $\mathcal{S}$, the same value of q for all links, instead of in one single set. Moreover, instead of minimizing T alone, they seek to minimize the ratio T/q while treating q as a variable. If T^1 is the least number of slots to accommodate all $|L|$ link activations, one per link, in the single-color case, and if the pair (T^*, q^*) provides the least possible T/q ratio while accommodating all $q|L|$ link activations, q per link, in the fractional-coloring case, then conceivably it may happen that $T^*/q^* < T^1$. If this does happen, then clearly we have $T^* < T^1 q^*$, so the T^*-slot sequence is shorter than q^* repetitions of the T^1-slot sequence and therefore the former is preferable to the latter, since in the two cases we have the same total number of link activations, viz., $q^*|L|$. Readily, in this case the sequence to be repeated in order for interference-free communication to be provided while needed is the one comprising T^* sets. Thus, so far as we seek to maximize network capacity via link scheduling, what needs to be done is maximize the ratio q/T. This ratio is how we define network capacity henceforth.

Important though these three contributions have been, they have each left relevant problems open as well. The most relevant one is the search for approaches for the exact determination of optimal fractional colorings in the physical interference model. This has been attempted neither by the proposals in [12], [13] (both of which target the protocol-based interference model, though the latter is exact while the former is a heuristic) nor by the one in [9] (which is a heuristic, even if one for the physical interference model).

In this letter we describe such an exact method and use it to chart the landscape of a class of random networks regarding the possibility of fractional-coloring-based link scheduling that is superior to its single-color counterpart. We have been able to do this for a reasonably wide range of the parameters involved in network generation, which has led us to conclude that, given the uncertainties afforded by the confidence intervals obtained, networks with the potential to benefit from a fractional-coloring approach occur in a non-negligible proportion. These empirical findings constitute one of our contributions. Our main contribution, though, is that by formulating capacity optimization in the framework of fractional coloring, we automatically provide for the fallback solution in which it is integer (single-color) coloring, rather than nontrivial fractional (more-than-one-color) coloring, that yields optimal capacity. Integer coloring, after all, is simply the particular case of fractional coloring in which it is better to use one, rather than more than one, color per link.

III. Mathematical Formulation

Given the set L of links to be scheduled, and letting N be the set of all nodes acting as sender or receiver in at least one link in L, we consider the undirected graph $G = (N,L)$, that is, the graph having N for set of vertices and L for set of undirected edges. We begin with the presentation of a linear programming (LP) problem for the determination of maximum network capacity as defined in Section II. That is, we aim to formulate the problem of finding the integers T and q that minimize the ratio T/q while allowing every link in L to be active in exactly q of T time slots and respecting the constraints imposed by the physical interference model. By the definition of a feasible set of links and also the definition of graph G above, clearly every feasible set of links corresponds to a matching in G, though the converse may not be true. Henceforth we refer as a *feasible matching* to any matching whose edges constitute a feasible set of links.

Let $\mathcal{M}$ be the set of all feasible matchings of G. For each $M \in \mathcal{M}$, let x_M be a real variable and consider the following LP problem.

$$\text{minimize} \quad w = \sum_{M \in \mathcal{M}} x_M \tag{2}$$

$$\text{subject to} \quad x_M \geq 0, \quad \forall M \in \mathcal{M} \tag{3}$$

$$\sum_{M \in \mathcal{M} | e \in M} x_M = 1, \quad \forall e \in L \tag{4}$$

This problem asks that the sum of all x_M's (the objective function w in Eq. (2)) be minimized while respecting the constraints that none of them be allowed to become negative (Eq. (3)) and that, for each edge $e \in L$, those x_M's for which $e \in M$ add up to 1 (Eq. (4)). Because the coefficients of the x_M's in Eqs. (2) and (4) are all equal to 1, hence rational numbers, at least one solution exists minimizing w with every x_M a rational number as well. Let $\mathcal{P}$ be the subset of $\mathcal{M}$ such that $M \in \mathcal{P}$ if and only if $x_M > 0$ in this solution.

For $M \in \mathcal{P}$, let p_M/q_M be such positive rational value of x_M minimizing w. If q^* denotes the least common multiple of all q_M's over $M \in \mathcal{P}$, then the desired minimum value of w, call it w^*, can be written as

$$w^* = \frac{\sum_{M \in \mathcal{P}} T_M}{q^*}, \tag{5}$$

where $T_M = q^* p_M / q_M$ is necessarily a positive integer.

Now consider any edge $e \in L$ and let $\mathcal{P}_e$ be the subset of $\mathcal{P}$ such that $M \in \mathcal{P}_e$ if and only if $e \in M$. That is, $\mathcal{P}_e$ is the set of all feasible matchings M of G that contribute to the minimum value of w with a positive x_M and moreover include edge e. Set $\mathcal{P}_e$ is necessarily nonempty, since the matching containing e and no other edge is by assumption feasible. By the constraint in Eq. (4), we have

$$\sum_{M \in \mathcal{P}_e} T_M = q^* \sum_{M \in \mathcal{P}_e} \frac{p_M}{q_M} = q^*. \tag{6}$$

If we view each $T_M \geq 1$ as a sort of multiplicity of matching *M*, then this equation is saying that the added multiplicities of all matchings in $\mathcal{P}_e$ equals q^*. In the context of scheduling the links in *L*, this means that, if we let all links in *M* be concomitantly active for T_M time slots and do this for all $M \in \mathcal{P}_e$, then after all $\sum_{M \in \mathcal{P}_e} T_M$ time slots link *e* will have appeared q^* times, regardless of the particular link *e* under consideration. Thus, ensuring that this happens for every $e \in L$ requires $\sum_{M \in \cup_{e \in L} \mathcal{P}_e} T_M = \sum_{M \in \mathcal{P}} T_M$ time slots. We denote this overall number of time slots by T^* and, by Eq. (5), conclude that $w^* = T^*/q^*$ is the desired minimum value of the ratio *T*/*q*.

The preceding development admits the graph-theoretic interpretation in which the edge set of *G* has to be covered by *T* feasible matchings while mandatorily including every edge in exactly *q* of them and minimizing *T*/*q*. A new fractional edge-chromatic indicator of *G*, denoted by $\chi'^*_{\rm phys}(G)$, can then be defined to be

$$\chi'^*_{\rm phys}(G) = w^* = \min_{k \geq 1} \frac{\chi'^k_{\rm phys}(G)}{k}, \tag{7}$$

where $\chi'^k_{\rm phys}(G)$ is the minimum number of feasible matchings that cover all edges in such a way that each one appears in exactly *k* matchings. Clearly, network capacity as defined in Section II is given by $1/\chi'^*_{\rm phys}(G)$. Moreover, by Eq. (7) we have

$$\chi'^*_{\rm phys}(G) \leq \chi'^1_{\rm phys}(G). \tag{8}$$

One of the core elements of our study in this letter is the determination, for some given *G*, of whether coloring its edges fractionally provides more capacity than coloring them with one single color per edge. Put differently, for each *G* we must be able to determine whether the inequality in Eq. (8) is strict, which clearly is true if and only if the minimum in Eq. (7) is achieved for $k > 1$. Solving the LP problem in Eqs. (2)–(4) already gives us the value of $\chi'^*_{\rm phys}(G)$ along with the corresponding x_M's that are nonzero. It would then seem that checking whether all of these x_M's equal 1 suffices, since if they do we can immediately conclude that $\chi'^*_{\rm phys}(G) = \chi'^1_{\rm phys}(G)$. However, that LP problem may admit several optimal solutions, including some that involve non-unit x_M's even when another equally optimal solution involves unit x_M's only. For this reason, testing whether every nonzero x_M equals 1 in the optimal solution returned by the LP solver is meaningful only in the affirmative case. In the negative case the test is meaningless, since it does not necessarily follow that $\chi'^*_{\rm phys}(G) < \chi'^1_{\rm phys}(G)$. Given this difficulty, we address the direct calculation of $\chi'^1_{\rm phys}(G)$ as well, by modifying the LP program of Eqs. (2)–(4) so that each x_M must be an integer equal to 0 or 1. The result is the following integer linear programming (ILP) problem.

$$\text{minimize} \quad w_{\rm int} = \sum_{M \in \mathcal{M}} x_M \tag{9}$$

$$\text{subject to} \quad x_M \in \{0, 1\}, \quad \forall M \in \mathcal{M} \tag{10}$$

$$\sum_{M \in \mathcal{M} | e \in M} x_M = 1, \quad \forall e \in L \tag{11}$$

Clearly, any valuation of the x_M's satisfying the constraints in Eqs. (10) and (11) characterizes a partition of the link set *L* into feasible matchings (specifically, a matching $M \in \mathcal{M}$ is in the partition if and only if $x_M = 1$). The objective function in Eq. (9) counts the corresponding number of matchings and therefore its optimal value, call it $w^*_{\rm int}$, is such that $\chi'^1_{\rm phys}(G) = w^*_{\rm int}$.

In summary, the following is how we find out whether $\chi'^*_{\rm phys}(G) < \chi'^1_{\rm phys}(G)$. Step 1: Find $\chi'^*_{\rm phys}(G)$ by solving the LP problem in Eqs. (2)–(4); Step 2: If every nonzero x_M in the solution equals 1, then conclude that $\chi'^*_{\rm phys}(G) = \chi'^1_{\rm phys}(G)$ and stop; Step 3: Find $\chi'^1_{\rm phys}(G)$ by solving the ILP problem in Eqs. (9)–(11); Step 4: Test whether $\chi'^*_{\rm phys}(G) < \chi'^1_{\rm phys}(G)$. Clearly, Steps 1–4 amount to solving the ILP problem only in those cases in which the solution to the LP problem is inconclusive as far as comparing $\chi'^*_{\rm phys}(G)$ and $\chi'^1_{\rm phys}(G)$ is concerned.

IV. Experimental Setup

Given a fixed graph *G* = (*N*, *L*), of vertex set *N* and edge set *L*, the computational core of our experiments is carrying out Steps 1 and 3 given at the end of Section III, which solve an LP problem and an ILP problem on *G*, respectively. In our experiments, graph *G* is an instance of the following random geometric graph. Given a $d \times d$ region in two-dimensional Euclidean space, each of the $|N|$ vertices is placed in it uniformly at random. For vertices $a, b \in N$, the unordered pair (*a, b*) is an edge of *L* if and only if $d_{ab} \leq (P/\beta\gamma)^{1/\alpha}$ (equivalently, if and only if $\mathrm{SINR}((a, b), \{(a, b)\}) \geq \beta$, where the role taken up by *a* as sender or receiver relative to *b* is immaterial). Put differently, (*a, b*) is an edge in *G* if and only if the singleton {(*a, b*)} is feasible. For each $e = (a, b) \in L$, sender s_e is either *a* or *b* uniformly at random, with receiver r_e set correspondingly. We use $\alpha = 4$, $\beta = 316.23$ (25 dB), and $\gamma = 8 \times 10^{-11}$ mW (-100.97 dBm), as well as $P = 300$ mW (24.78 dBm) throughout all experiments.

For fixed $|N|$ and *d*, we generated 1 000 graph instances and tested each one for suitability to Steps 1–4. Failing instances were dropped, so all our results are expressed as averages over the passing instances. An instance can fail for at least one of three reasons: edge set *L* is empty; edge set *L* has more than 128 edges, in which case we lack the computational resources to enumerate all feasible matchings that go in set $\mathcal{M}$; the number of feasible matchings in set $\mathcal{M}$ is greater than 50×10^6, which is as far as we can go given 128 GB of RAM and given that we use the Gurobi suite (www.gurobi.com) for solving both the LP and ILP problems, always with the presolver disabled and Simplex as the core linear programming solver. Choosing $\beta = 25$ dB has therefore aimed at minimizing the number of failing instances by avoiding too many networks with more than 128 edges. Even so, many of the instances for the higher values of $|N|$ reached the mark of several million feasible sets. Moreover, $\beta = 25$ dB lies within the range for WiFi, a key component of wireless mesh networks [15].

Given a $|N|, d$ pair, enumerating all the feasible matchings in $\mathcal{M}$ for a passing graph instance *G* has taken on average up to about 597 seconds to complete. Running Step 1 to find $\chi'^*_{\rm phys}(G)$, or Step 3 to find $\chi'^1_{\rm phys}(G)$ whenever reaching that step, has required on average up to about 199 and

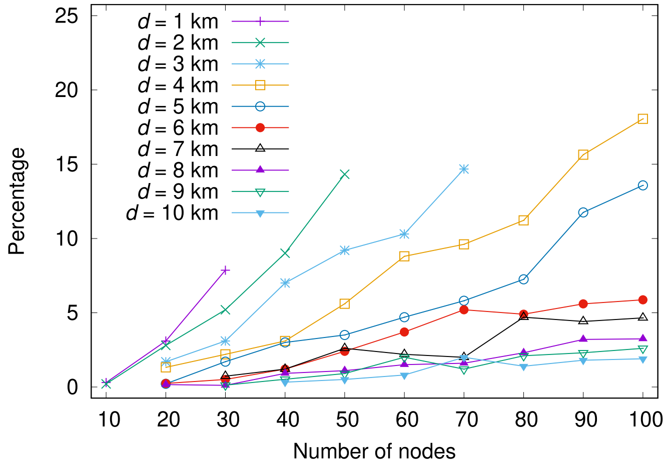

Fig. 1. Percentage of graph instances G for which $\chi'^{*}_{\text{phys}}(G) < \chi'^{1}_{\text{phys}}(G)$. Missing points indicate either the absence of passing instances or the absence of passing instances for which $\chi'^{*}_{\text{phys}}(G) < \chi'^{1}_{\text{phys}}(G)$. The latter is the case of all points missing for $|N| = 10$, 20, or 30, as well as for $|N| = 60$ with $d = 2$ km.

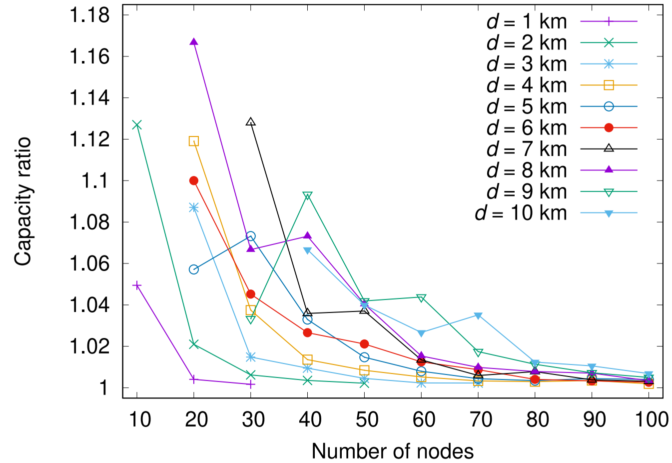

Fig. 2. Average value of the capacity ratio $\chi'^{1}_{\text{phys}}(G)/\chi'^{*}_{\text{phys}}(G)$ over the graph instances G accounted for in Figure 1. Confidence intervals (not shown for the sake of clarity) at the 95% level lie below 7% of the means for the lowest values of $|N|$ and below 0.5% for the highest values.

3 575 seconds, respectively. These figures refer to an Intel Xeon E5-1650 v4 running at 3.6 GHz on 128 GB of RAM. Such maximum averages were all observed for $|N| = 100$ and $d = 4$ km. Taken in conjunction with the limitations discussed above regarding the value of $|L|$ and a graph instance's number of feasible matchings, these running times make it unlikely that problem instances can be scaled up significantly while still being amenable to solution by our exact method. We return to this issue in the Appendix.

V. Results and Discussion

We give results in Figures 1 and 2, with Figure 1 showing the percentage of graph instances G for which $\chi'^{*}_{\text{phys}}(G) < \chi'^{1}_{\text{phys}}(G)$. Such graph instances are cases of network-capacity improvement when substituting a schedule based on fractional edge coloring (of capacity $1/\chi'^{*}_{\text{phys}}(G)$) for one based on edge coloring that employs one single color per edge (of capacity $1/\chi'^{1}_{\text{phys}}(G)$). The ratios of capacity improvement, given by $\chi'^{1}_{\text{phys}}(G)/\chi'^{*}_{\text{phys}}(G)$, are shown in Figure 2 as averages over the pertinent graph instances G. These figures are best analyzed in terms of a geometric graph's node density (its number of nodes divided by the area of deployment), since this automatically provides an indirect basis for analysis in other terms as well (higher densities imply higher numbers of edges and in turn higher numbers of feasible matchings).

Clearly, a tendency is shown in Figure 1 of higher percentages for higher node densities. This is easily seen as we fix d while $|N|$ is increased, but holds across values of d as well: note that node density is of the order of 10^{-5} to 10^{-4} for $d = 1$ km, 10^{-6} to 10^{-5} for $d = 2$ km, and so on, which in general correlates well with lower-d plots being located higher up in the figure. Of course, owing to the total absence of passing graph instances for the highest values of $|N|$ and lowest values of d, at this point we can only speculate as to what would happen if such instances' numbers of edges and of feasible matchings could be handled, but the trend seems clear nonetheless. Indeed, increasing a geometric graph G's node density tends to lead to a higher number of edges, and therefore a pressure exists for the value of $\chi'^{1}_{\text{phys}}(G)$ to increase as well. Intuitively, this presents an opportunity for some $k > 1$ to prevail in Eq. (7) and for $\chi'^{*}_{\text{phys}}(G) < \chi'^{1}_{\text{phys}}(G)$ to occur.

Another observable of interest is the capacity ratio, or gain, given by $\chi'^{1}_{\text{phys}}(G)/\chi'^{*}_{\text{phys}}(G)$, for those instances G for which $\chi'^{*}_{\text{phys}}(G) < \chi'^{1}_{\text{phys}}(G)$ is obtained. This is shown in Figure 2 as averages over the instances accounted for in Figure 1. A relationship is seen to exist also between the graphs' node densities and their capacity gains, but now the trend is for the lower-node-density graphs to afford higher marks. This can be seen as we fix d and increase $|N|$ (plots in the figure are generally decreasing toward 1) and, to a limited extent, across values of d as well. At this point we must not state the latter more firmly because the plots in Figure 2 deviate from what they would look like ideally (plots nicely nested one above the other with increasing d). This may have to do with the higher confidence intervals occurring precisely where deviations from the said ideal are most striking (confidence intervals up to nearly 7% of the corresponding means in some cases), but only further experimentation will clarify the issue.

VI. Concluding Remarks

Considering the set of feasible matchings of G in the physical interference model has led to the definition of a new fractional edge-chromatic indicator for G, $\chi'^{*}_{\text{phys}}(G)$, and correspondingly a new definition of network capacity, $1/\chi'^{*}_{\text{phys}}(G)$. The single-color-per-edge definition, $1/\chi'^{1}_{\text{phys}}(G)$, can be optimal whenever $\chi'^{*}_{\text{phys}}(G) = \chi'^{1}_{\text{phys}}(G)$, which as we have found can happen relatively often. This notwithstanding, denser networks can also yield $\chi'^{*}_{\text{phys}}(G) < \chi'^{1}_{\text{phys}}(G)$, which sometimes accounts for non-negligible capacity improvements (i.e., increases in the ratio $\chi'^{1}_{\text{phys}}(G)/\chi'^{*}_{\text{phys}}(G)$). Conveniently, adopting the fractional-coloring framework allows optimization to be carried out exclusively by solving the corresponding LP problem, without any need whatsoever to call upon the accompanying ILP problem, which is substantially more time-consuming than the former.

APPENDIX
REGARDING SCALABILITY

It is clear from our discussion in Section IV that the networks we experimented with were limited by the need to enumerate all feasible matchings in $\mathcal{M}$ in order to exactly solve the LP problem given by Eqs. (2)–(4) and the ILP problem given by Eqs. (9)–(11). The number of such matchings eventually exhausts all computational resources available, thus making it impossible for $\mathcal{M}$ to be enumerated and for either problem to be solved. Ultimately, however, it is the LP problem that matters most, and for this one a clear path exists for the search for scalability.

To see that a substantially more efficient alternative may be available, consider the particular case in which $\mathcal{M}$ is the set of all matchings of G (i.e., not necessarily feasible in the sense of the physical interference model). In this case, the LP problem given in Eqs. (2)–(4) can be taken as defining the fractional chromatic index of G [16] and the better alternative is to consider it in its dual formulation, given as follows.

$$\text{maximize} \quad z = \sum_{e \in L} y_e \tag{12}$$

$$\text{subject to} \quad \sum_{e \in M} y_e \leq 1, \quad \forall M \in \mathcal{M} \tag{13}$$

In this formulation, for each $e \in L$ we have a real variable y_e (which can be negative, zero, or positive, by virtue of the equality constraint in Eq. (4)), and for each matching $M \in \mathcal{M}$ we have a constraint forbidding the y_e's for $e \in M$ to add up to more than 1 (Eq. (13)). The goal is to maximize the sum of all y_e's (the objective function z in Eq. (12)). By LP duality, if the LP problem in Eqs. (2)–(4) defines the fractional chromatic index of G, then so does the one in Eqs. (12) and (13).

It would seem that the new formulation suffers from the same problem as the previous one, the only difference being that now the size of $\mathcal{M}$ is reflected in the number of constraints, not the number of variables. While the latter is clearly true, it is in principle possible to solve the problem without listing all constraints explicitly. We start by maximizing z subject to only a minimal set of constraints (one for each singleton matching $\{e\} \subseteq L$, which ensures a finite maximum value for z). Then we iterate, each time expanding the set of explicitly listed constraints with the addition of an unlisted one that is currently violated. We do this until no violated unlisted constraints remain.

In order to succeed with this approach, we must ensure that both the time required to identify a violated unlisted constraint and the overall number of iterations are polynomially bounded. The first of these goals is achieved by resorting to the problem of finding a maximum-weight matching in G, which is known to be solvable in polynomial time by a variety of methods (see [17] and references therein). To see how this problem can be of use, let y_e^* be the value of y_e for each $e \in L$ after one of the iterations, then find a maximum-weight matching of G with the y_e^*'s as weights. Let M^* be the matching obtained, of weight $W^* = \sum_{e \in M^*} y_e^*$. If $W^* > 1$, then clearly the constraint in Eq. (13) for $M = M^*$ is being violated and should therefore be listed explicitly for the next iteration. If $W^* \leq 1$, then clearly no further violated constraints exist (since M^* has maximum weight) and no further iterations are needed. As for the second goal, that of iterating for only a polynomially-bounded number of times, the ellipsoid method for linear programming, though impractical for direct use, provides the necessary theoretical guarantee [18]. An essentially equivalent path is followed in [13].

The case in which the matchings in $\mathcal{M}$ are all feasible in the sense of the physical interference model is substantially more complex, but at least we have the results of [18] to rely on for guidance. Specifically, what we must do is discover a polynomial-time algorithm to find a maximum-weight feasible matching of G. Such an algorithm will depend on all the intricacies underlying the definition of SINR in Eq. (1), and whether one exists is for now an open problem. Should it not exist, or should it prove too elusive to find, a more costly algorithm will also do: though requiring more computational effort to determine the required feasible matching, the expected savings derived from not having to list a huge number of constraints explicitly are bound to be worth the additional resources expended.

REFERENCES

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.

[2] G. Brar, D. M. Blough, and P. Santi, "Computationally efficient scheduling with the physical interference model for throughput improvement in wireless mesh networks," in *Proc. 12th Annu. Int. Conf. Mobile Comput. Netw. (MobiCom)*, Los Angeles, CA, USA, 2006, pp. 2–13.

[3] D. M. Blough, G. Resta, and P. Santi, "Approximation algorithms for wireless link scheduling with SINR-based interference," *IEEE/ACM Trans. Netw.*, vol. 18, no. 6, pp. 1701–1712, Dec. 2010.

[4] D. Chafekar, V. S. A. Kumart, M. V. Marathe, S. Parthasarathy, and A. Srinivasan, "Approximation algorithms for computing capacity of wireless networks with SINR constraints," in *Proc. 27th Conf. Comput. Commun. (INFOCOM)*, 2008, pp. 1166–1174.

[5] C. Boyaci, B. Li, and Y. Xia, "An investigation on the nature of wireless scheduling," in *Proc. IEEE INFOCOM*, San Diego, CA, USA, 2010, pp. 1–9.

[6] M. Leconte, J. Ni, and R. Srikant, "Improved bounds on the throughput efficiency of greedy maximal scheduling in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 709–720, Jun. 2011.

[7] Y. Shi, Y. T. Hou, S. Kompella, and H. D. Sherali, "Maximizing capacity in multihop cognitive radio networks under the SINR model," *IEEE Trans. Mobile Comput.*, vol. 10, no. 7, pp. 954–967, Jul. 2011.

[8] O. Goussevskaia, M. M. Halldórsson, and R. Wattenhofer, "Algorithms for wireless capacity," *IEEE/ACM Trans. Netw.*, vol. 22, no. 3, pp. 745–755, Jun. 2014.

[9] F. R. J. Vieira, J. F. de Rezende, and V. C. Barbosa, "Scheduling wireless links by vertex multicoloring in the physical interference model," *Comput. Netw.*, vol. 99, pp. 125–133, Apr. 2016.

[10] A. D. Gore and A. Karandikar, "Link scheduling algorithms for wireless mesh networks," *IEEE Commun. Surveys Tuts.*, vol. 13, no. 2, pp. 258–273, 2nd Quart., 2011.

[11] I. Rubin, C.-C. Tan, and R. Cohen, "Joint scheduling and power control for multicasting in cellular wireless networks," *EURASIP J. Wireless Commun.*, vol. 2012, p. 250, Aug. 2012.

[12] F. R. J. Vieira, J. F. de Rezende, V. C. Barbosa, and S. Fdida, "Scheduling links for heavy traffic on interfering routes in wireless mesh networks," *Comput. Netw.*, vol. 56, no. 5, pp. 1584–1598, 2012.

[13] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," *IEEE Trans. Inf. Theory*, vol. IT-34, no. 5, pp. 910–917, Sep. 1988.

[14] F. Arbabjolfaei and Y.-H. Kim, *Fundamentals of Index Coding*. Hanover, MA, USA: Now, 2018.

[15] *Receiver Sensitivity Tables for MIMO-OFDM 802.11n*. Accessed: Dec. 10, 2019. [Online]. Available: https://mentor.ieee.org/802.11/dcn/03/11–03-0845–00-000n-11–03-0845–00-000n-receiver-sensitivity-tables-mimo-ofdm.ppt

[16] M. Stiebitz, D. Scheide, B. Toft, and L. M. Favrholdt, *Graph Edge Coloring*. Hoboken, NJ, USA: Wiley, 2012.

[17] C.-C. Huang and T. Kavitha, "New algorithms for maximum weight matching and a decomposition theorem," *Math. Oper. Res.*, vol. 42, no. 2, pp. 411–426, 2017.

[18] M. Grötschel, L. Lovász, and A. Schrijver, "The ellipsoid method and its consequences in combinatorial optimization," *Combinatorica*, vol. 1, no. 2, pp. 169–197, 1981.