Coordinating Concurrent Transmissions: A Constant-Factor Approximation of Maximum-Weight Independent Set in Local Conflict Graphs

Petteri Kaski[1], Aleksi Penttinen[2], and Jukka Suomela[1]

[1] Helsinki Institute for Information Technology HIIT
Department of Computer Science, University of Helsinki
P.O. Box 68, FI-00014 University of Helsinki, Finland
`firstname.lastname@cs.helsinki.fi`
[2] Networking Laboratory, Helsinki University of Technology
P.O. Box 3000, FI-02015 TKK, Finland
`firstname.lastname@tkk.fi`

Abstract. We study the algorithmic problem of coordinating transmissions in a wireless network where radio interference constrains concurrent transmissions on wireless links. We focus on pairwise conflicts between the links; these can be described as a conflict graph. Associated with the conflict graph are two fundamental network coordination tasks: (a) finding a nonconflicting set of links with the maximum total weight, and (b) finding a link schedule with the minimum total length. Our work shows that two assumptions on the geometric structure of conflict graphs suffice to achieve polynomial-time constant-factor approximations: (i) bounded density of devices, and (ii) bounded range of interference. We also show that these assumptions are not sufficient to obtain a polynomial-time approximation scheme for either coordination task.

Key words: geometric graphs, maximum-weight independent set, radio interference.

1 Introduction

A fundamental challenge in wireless networking is the shared transmission medium, which in many cases prevents concurrent transmissions due to radio interference. This brings forth the algorithmic problem of coordinating the transmissions so that performance loss due to interference does not occur. In this work, we investigate the polynomial-time approximability of network coordination within the following framework.

Interference in Wireless Networks. A wireless network consists of devices which communicate with each other by radio transmissions. We study unicast networks where each radio transmission has one designated recipient, see Fig. 1 for an illustration.

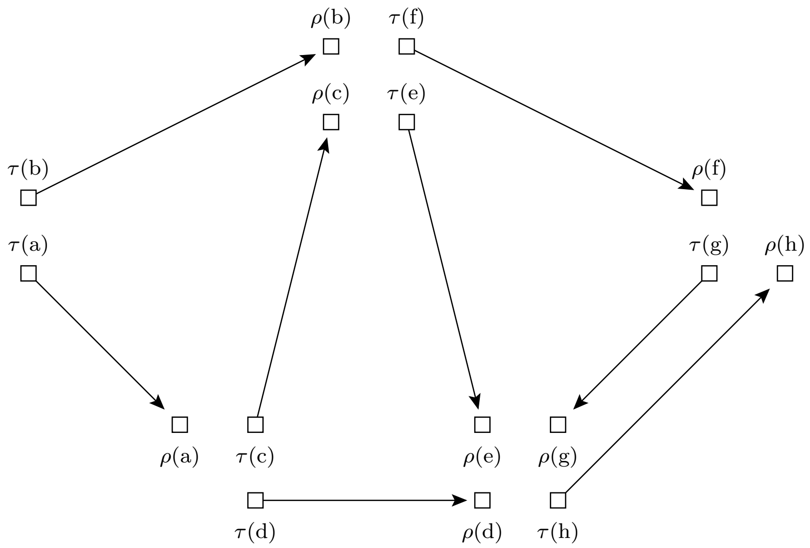

Fig. 1. A wireless network. Wireless communication links are marked with arrows. Devices are marked with boxes, τ denotes a transmitter and ρ denotes a receiver. For clarity, each device in this illustration takes part in only one transmission.

A radio transmission may interfere with other transmissions. We focus on systems where the radio interference is dominated by the near-far effect: radio reception from a distant transmitter may be blocked by other transmitters which are much closer to the receiver.

Figure 2 illustrates transmitter-receiver pairs where the near-far effect might occur in our example. In the illustration, the link from $\tau(\mathrm{a})$ to $\rho(\mathrm{a})$ and the link from $\tau(\mathrm{d})$ to $\rho(\mathrm{d})$ cannot be active simultaneously: the transmitter of the latter blocks the receiver of the former.

We focus on pairwise conflicts between the links. The pairwise conflicts can be described as a *conflict graph* [1]. A conflict graph $G = (V, E)$ is an undirected graph where each vertex $v \in V$ corresponds to a communication link and an edge $\{u, v\} \in E$ describes that the links u and v are mutually conflicting. Figure 3 illustrates a conflict graph; here the vertex $\mathrm{a} \in V$ corresponds to the communication link from $\tau(\mathrm{a})$ to $\rho(\mathrm{a})$ and the vertex $\mathrm{d} \in V$ corresponds to the communication link from $\tau(\mathrm{d})$ to $\rho(\mathrm{d})$. As these two links cannot transmit simultaneously, there is an edge $\{\mathrm{a}, \mathrm{d}\} \in E$ in the conflict graph.

Algorithmic Problems and Earlier Work. Associated with a conflict graph are the following network coordination tasks.

(1) Given some weights (such as priorities or utilities) on each vertex, find an *independent set* of the maximum total weight; in other words, find a nonconflicting set of links of the maximum total weight. See Fig. 3 for an example.

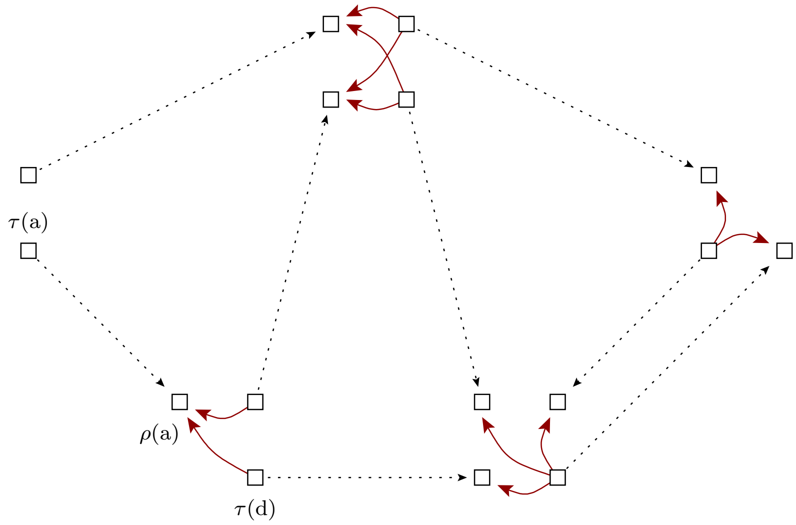

Fig. 2. The near-far effect in a wireless network. Solid arrows point from an interfering transmitter to the interfered receiver. For example, if the transmitter $\tau(\mathrm{d})$ is active, the device $\rho(\mathrm{a})$ cannot receive the transmission from $\tau(\mathrm{a})$. The signal power received from $\tau(\mathrm{a})$ is too low in comparison with the interfering power received from $\tau(\mathrm{d})$.

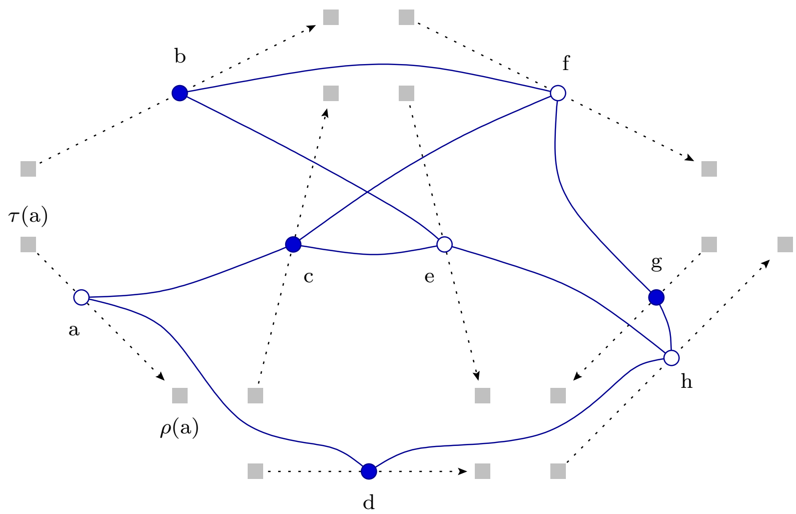

Fig. 3. The conflict graph for the example in Fig. 2. There is one vertex for each communication link and one edge for each pair of conflicting links; vertices are marked with circles and edges are marked with solid lines. Note that even though interference in Fig. 2 was highly localised, such locality is no longer immediately visible in the conflict graph. An independent set $\{\mathrm{b}, \mathrm{c}, \mathrm{d}, \mathrm{g}\}$ is highlighted; these communication links can be active simultaneously.

(2) Given some data transmission needs on each link, find a *link schedule* of the minimum length such that at each point in time, the set of active links is nonconflicting, and each link is active for a time that suffices to cover its data transmission needs. See Sect. 6 for a precise linear programming formulation.

An approximation algorithm for maximum-weight independent set also implies an approximation algorithm for the link scheduling problem in the same class of graphs [2, 3]. Unfortunately, both problems are prohibitively hard to approximate in general graphs [4–6]. Jain et al. [1] have put forth the question of whether there is a family of conflict graphs which arises in realistic network deployments and which makes the problem of finding an independent set easier.

Contribution. Our work shows that two assumptions on the structure of conflict graphs suffice in order to achieve a polynomial-time constant-factor approximation of maximum-weight independent set and link scheduling:

(i) Bounded density of the devices. Radios are points in a low-dimensional space and they are not located in an arbitrarily dense manner.
(ii) Bounded range of interference. Conflicts are caused by the near-far effect; if there is a conflict, the interfering transmitter is close to the interfered receiver.

Note that we do not need to assume that there has to be interference in certain situations, say, between nodes close to each other; indeed, such assumptions are often not valid in practice [7, 8]. We can make measurements in the deployed physical system to determine whether a pair of links is mutually conflicting; there is no need to use a simplifying model of radio propagation and interference.

We also show that the assumptions (i) and (ii) alone are not sufficient in order to achieve an arbitrarily small approximation ratio; further assumptions such as a bounded range of the wireless link are required.

2 Statement of Results

Let N be a constant which controls the relative density of the nodes.

Definition 1. *An N-*local conflict graph *is a tuple (G, τ, ρ) where $G = (V, E)$ is an undirected graph and τ, ρ are functions $V \to \mathbb{R}^2$ such that*

(i) *the function $v \mapsto (\tau(v), \rho(v))$ is an injection, and no unit disk in $\mathbb{R}^2$ contains more than N points in $\tau(V) \cup \rho(V)$*
(ii) *for all $\{v, u\} \in E$ it holds that $d(\tau(v), \rho(u)) < 1$ or $d(\tau(u), \rho(v)) < 1$ where $d(\cdot, \cdot)$ is the Euclidean distance.*

We call $\tau(v)$ the *transmitter* and $\rho(v)$ the *receiver*, the intuition being that a pair $(\tau(v), \rho(v))$ corresponds to a data transmission link. Note that $d(\tau(v), \rho(v))$ is unrestricted and that some receivers and transmitters may coincide; however, the pair $(\tau(v), \rho(v))$ must be unique for every vertex.

If we required $\tau(v) = \rho(v)$ for each $v \in V$, we would obtain what we call $(2, N)$-*local graphs* [9]; these are similar to *civilised graphs* [10, §8.5]. Thus, N-local conflict graphs can be interpreted as a natural generalisation of the families of local graphs and civilised graphs.

In Sect. 3, we derive some basic properties of N-local conflict graphs. We show that N-local conflict graphs are not contained in families such as bounded-degree, bounded-density, bipartite, planar, or disk graphs.

In Sect. 4, we prove our main result; *MWIS* refers to the problem of finding a maximum-weight independent set:

Theorem 1. *MWIS for N-local conflict graphs admits a polynomial-time* $(5+\epsilon)$*-approximation algorithm for any constants* $\epsilon > 0$ *and N.*

While the time complexity of the algorithm depends on the parameter N, we emphasise that the approximation ratio does not depend on N. This is unlike families such as bounded-degree graphs where achievable approximation ratios typically depend on the parameters of the family [11]; for example, MWIS in graphs of maximum degree Δ can be approximated within a factor of $O(\Delta \log\log \Delta / \log \Delta)$ [12].

In Sect. 5, we show that approximating beyond a certain constant factor remains hard:

Theorem 2. *MWIS for N-local conflict graphs admits no polynomial-time approximation scheme (PTAS) for any N unless P = NP.*

This is unlike families such as disk graphs; for example, MWIS in disk graphs admits a PTAS [13].

In Sect. 6, we consider the problem of fractional covering by independent sets in local conflict graphs, obtaining analogous approximability and inapproximability results for the covering LP, which captures the link scheduling problem.

3 Representability

It is not immediate from Definition 1 which graphs admit representation as a local conflict graph. The purpose of this section is to shed some light on this question. We begin by showing that local conflict graphs are not contained in families of graphs such as planar graphs, bounded-degree graphs, or disk graphs.

A first observation is that the family of N-local conflict graphs is closed under deletion of edges and vertices. Furthermore, an N_1-local conflict graph is an N_2-local conflict graph for any $N_2 \geq N_1$.

Theorem 3. *Any bipartite graph can be represented as a* 1*-local conflict graph.*

Proof. Consider a bipartite graph $G = (V, E)$. The set V can be partitioned into $A = \{a_1, a_2, \ldots, a_m\}$ and $B = \{b_1, b_2, \ldots, b_n\}$ for some m, n such that all edges are between A and B. Let $\tau(a_i) = (3i, -3)$ and $\rho(a_i) = (0, 0)$ for all i; let $\tau(b_j) = (0, 0)$ and $\rho(b_j) = (3j, 3)$ for all j.

It follows that the maximum degree of a vertex in a 1-local conflict graph can be as high as $|V| - 1$ (consider the complete bipartite graph $K_{1,n}$), the average degree and the minimum degree can be as high as $|V|/2$ (consider $K_{n,n}$), and there are 1-local conflict graphs that are not planar and not disk graphs (consider $K_{3,3}$).

A local conflict graph need not be bipartite. To illustrate the rich substructure that can occur in a local conflict graph, we show that a local conflict graph may contain relatively large but not arbitrarily large cliques.

Theorem 4. *A complete graph on N^2 vertices is representable as an N-local conflict graph.*

Proof. Consider a complete graph with vertices $V = \{v_{i,j} \mid i, j \in \{1, 2, \ldots, N\}\}$. Let $\tau(v_{i,j}) = (0, i/N)$ and $\rho(v_{i,j}) = (0, j/N)$ for all i and j.

Lemma 1. *For every tournament (complete oriented graph) $G = (V, A)$ on n vertices, there are $s \in V$, $t \in V$ and $X \subseteq V$ with $\{s\} \times X \subseteq A$, $X \times \{t\} \subseteq A$, and $|X| \geq (n - 2)/6$.*

Proof. Let $v, u \in V$, $v \neq u$. Let

$$Q = \{x \mid (v, x) \in A, (x, u) \in A\}, \qquad R = \{x \mid (u, x) \in A, (x, v) \in A\},$$
$$S = \{x \mid (x, v) \in A, (x, u) \in A\}, \qquad T = \{x \mid (v, x) \in A, (u, x) \in A\}.$$

If $|Q| \geq (n - 2)/6$ or $|R| \geq (n - 2)/6$, we are done. Otherwise, $|S| + |T| \geq 2(n + 1)/3$. If $|S| \geq (n + 1)/3$, the subgraph induced by S contains a vertex a with outdegree at least $(n - 2)/6$; let $s = a$, $t = u$, and let $X \subseteq S$ consist of the successors of a. The case $|T| \geq (n + 1)/3$ is analogous.

Theorem 5. *A complete graph on $6N^2 + 8$ vertices cannot be represented as an N-local conflict graph.*

Proof. To reach a contradiction, assume that there is a complete graph on $6N^2 + 8$ vertices that is an N-local conflict graph. We say that $P(v, u)$ holds if $d(\tau(v), \rho(u)) < 1$. Orient the graph as follows: if $P(v, u)$ and not $P(u, v)$, assign the direction (v, u) on $\{v, u\} \in E$; if $P(u, v)$ and not $P(v, u)$, assign the opposite direction; otherwise both $P(v, u)$ and $P(u, v)$ hold, in which case assign an arbitrary direction.

Choose s, t, and X as in Lemma 1; $|X| \geq N^2 + 1$. Now, $P(s, x)$ and $P(x, t)$ hold for all $x \in X$. A unit disk centred at $\tau(s)$ contains all points $\rho(X)$. There can be at most N distinct points; thus, there is a set $X' \subseteq X$ and a point $\rho' \in \mathbb{R}^2$ such that $\rho(X') = \{\rho'\}$ and $|X'| \geq N + 1$.

A unit disk centred at $\rho(t)$ contains all points $\tau(X')$. Again, there can be at most N distinct points; thus, there are two distinct vertices $v, u \in X'$ and a point $\tau' \in \mathbb{R}^2$ with $\tau(v) = \tau(u) = \tau'$. This is a contradiction because $v \mapsto (\tau(v), \rho(v))$ is an injection by Definition 1.

It follows immediately that the family of N-local conflict graphs is not closed under taking minors (form a bipartite graph by splitting each edge of a large complete graph in two).

4 Proof of Theorem 1

The input of the algorithm consists of a graph $G = (V, E)$, points $\tau(v)$ and $\rho(v)$ for each $v \in V$, and a weight $w(v)$ for each $v \in V$.

We present the algorithm in a somewhat more general setting than required by the MWIS problem. Write $C(v, I)$ for the *contribution* of the vertex $v \in V$ in the proposed solution $I \subseteq V$. Let $W(v, I) = w(v)C(v, I)$ for each $I \subseteq V$ and $v \in V$, and let $W(A, I) = \sum_{v \in A} W(v, I)$ for any $A \subseteq V$. The objective is to find a solution $I \subseteq V$ that maximises $W(V, I)$, that is, maximises weighted contributions.

In the case of MWIS, we define $C(v, I) = 1$ if $v \in I$ and there is no $u \in I$ with $\{v, u\} \in E$ and $d(\tau(u), \rho(v)) < 1$, otherwise $C(v, I) = 0$. The set I that maximises $W(V, I)$ is (after removing vertices with zero contribution) a maximum-weight independent set in G.

Define the set of possibly interfering vertices

$$U(v) = \{v\} \cup \{u \in V \mid d(\tau(u), \rho(v)) < 1\}.$$

The algorithm makes use of the following assumptions on C; these are immediate in the case of MWIS for local conflict graphs: $C(v, I)$ can be evaluated in polynomial time; $C(v, I) = 0$ for all $v \notin I$; $C(v, I_1) \geq C(v, I_2)$ for all $I_1 \subseteq I_2$ with $v \in I_1$ (contributions are nonincreasing); and $C(v, I) = C(v, I \cap U(v))$ for all I and $v \in V$ (locality).

In the algorithm, the full problem is divided into *subproblems*. Each subproblem is defined by a subset $A \subseteq V$, and the associated task is to find a set $I \subseteq A$ that maximises $W(A, I)$.

Let $\hat{W}(v) = W(v, \{v\})$ for each $v \in V$, and let $\hat{W}(A) = \sum_{v \in A} \hat{W}(v)$. We make use of the following properties. As $C(v, I) = 0$ for $v \notin I$, we have $W(A, I) = W(A \cap I, I)$ for all $A \subseteq V$. By nonincreasingness, $W(v, I_1) \geq W(v, I_2)$ for all $I_1 \subseteq I_2 \subseteq V$ and $v \in I_1$; in particular, $\hat{W}(v) \geq W(v, I)$ for all $v \in I$ and $I \subseteq V$. As $C(v, I) = 0$ for $v \notin I$, we have $\hat{W}(v) \geq W(v, I)$ for all $v \in V \setminus I$ and $I \subseteq V$. In summary, $\hat{W}(A) \geq W(A, I)$ for all $A, I \subseteq V$, and $\hat{W}(I) \geq W(I, I) = W(V, I)$ for all $I \subseteq V$.

To create the subproblems, we apply a shifting strategy [14, 15]. We make the following initial assignments. Choose an integer $k \geq 3$ such that $5k^2/(k-2)^2 < 5 + \epsilon$. Let

$$\begin{aligned}
A'_i &= \{(x, y) \in \mathbb{R}^2 \mid i \leq x < i + 1\}, && i \in \mathbb{Z}, \\
B'_j &= \{(x, y) \in \mathbb{R}^2 \mid j \leq y < j + 1\}, && j \in \mathbb{Z}, \\
A_i &= \bigcup \{A'_t \mid t \in \mathbb{Z},\, t \equiv i \pmod{k}\}, && i = 0, 1, \ldots, k-1, \\
B_j &= \bigcup \{B'_t \mid t \in \mathbb{Z},\, t \equiv j \pmod{k}\}, && j = 0, 1, \ldots, k-1, \\
D_{ij} &= \mathbb{R}^2 \setminus A_i \setminus B_j, && i, j = 0, 1, \ldots, k-1.
\end{aligned}$$

Each D_{ij} consists of squares $k-1$ units wide and high. We write $D_{ij1}, D_{ij2}, \ldots$ for the nonempty squares of D_{ij}. Let $Z_{ij} \subseteq V$ be the set of vertices v with both

$\tau(v)$ and $\rho(v)$ in D_{ij}, and let $X_{ij\beta} \subseteq Z_{ij}$ be the set of vertices v with both $\tau(v)$ and $\rho(v)$ in $D_{ij\beta}$. Form the set of "short links" $X_{ij} = \bigcup_\beta X_{ij\beta}$ and the set of "long links" $Y_{ij} = Z_{ij} \setminus X_{ij}$.

Now we can use the following procedure to find an approximately optimal solution. In the first part we solve small subproblems by exhaustive search; in the second part we apply the standard greedy algorithm for finding a large cut.

1. Initialise $\mathcal{N}$ to an empty set.
2. For all $i = 0, 1, \ldots, k-1$ and $j = 0, 1, \ldots, k-1$:
   (a) For each nonempty square $D_{ij\beta}$, find a subset $I \subseteq X_{ij\beta}$ which maximises $W(X_{ij\beta}, I)$. Call this set $S_{ij\beta}$.
   (b) Let $S_{ij} = \bigcup_\beta S_{ij\beta}$. Insert S_{ij} into $\mathcal{N}$.
3. For all $i = 0, 1, \ldots, k-1$ and $j = 0, 1, \ldots, k-1$:
   (a) Initialise $\varGamma$ and $\varLambda$ to empty sets.
   (b) For each nonempty square $D_{ij\beta}$, in an arbitrary order: Write $[\beta, \varGamma]$ for the set of vertices $v \in Y_{ij}$ such that one of the points $\tau(v), \rho(v)$ is located in $D_{ij\beta}$ and the other point is located in $D_{ij\varGamma} = \bigcup_{\gamma \in \varGamma} D_{ij\gamma}$; define the set $[\beta, \varLambda]$ similarly. Add β to $\varLambda$ if $\hat{W}([\beta, \varGamma]) > \hat{W}([\beta, \varLambda])$; otherwise add β to $\varGamma$.
   (c) Let $T_1 = \{v \in Y_{ij} \mid \tau(v) \in D_{ij\varGamma},\ \rho(v) \in D_{ij\varLambda}\}$ and let $T_2 = \{v \in Y_{ij} \mid \tau(v) \in D_{ij\varLambda},\ \rho(v) \in D_{ij\varGamma}\}$. Let $T \in \{T_1, T_2\}$ be the set that maximises $\hat{W}(T)$. Call this set R_{ij}, and insert R_{ij} into $\mathcal{N}$.
4. Return $\tilde{I} = \arg\max_{I \in \mathcal{N}} W(V, I)$.

The time complexity of the algorithm is polynomial in the size of the input since the number of nonempty squares $D_{ij\beta}$ is bounded by $2|V|$, the number of distinct transmitters or receivers in each square is bounded by a constant, and a pair $(\tau(v), \rho(v))$ uniquely determines v for all $v \in V$. The following three lemmata establish the correctness of the algorithm. We denote by $I^*(A)$ an optimal solution to the subproblem A.

Lemma 2. *Each S_{ij} is an optimal solution of the subproblem X_{ij}.*

Proof. To reach a contradiction, assume that S_{ij} is not optimal, i.e., there exists an $I \subseteq X_{ij}$ with $W(X_{ij}, I) > W(X_{ij}, S_{ij})$. Then there is a β with $W(X_{ij\beta}, I) > W(X_{ij\beta}, S_{ij})$. As the squares $D_{ij\delta}$ are separated by stripes of width one, $U(v) \cap X_{ij} \subseteq X_{ij\beta}$ for all $v \in X_{ij\beta}$. Thus, $W(X_{ij\beta}, I) = W(X_{ij\beta}, I \cap X_{ij\beta})$ and $W(X_{ij\beta}, S_{ij}) = W(X_{ij\beta}, S_{ij} \cap X_{ij\beta}) = W(X_{ij\beta}, S_{ij\beta})$. Thus, $W(X_{ij\beta}, I \cap X_{ij\beta}) > W(X_{ij\beta}, S_{ij\beta})$, contradicting the choice of $S_{ij\beta}$.

Lemma 3. *Each R_{ij} is a 4-approximate solution of the subproblem Y_{ij}.*

Proof. The greedy algorithm in parts (3a) and (3b) finds a cut (in the directed graph with an arc of weight $\hat{W}(v)$ from $\tau(v)$ to $\rho(v)$ for each $v \in Y_{ij}$) of a total weight at least $\hat{W}(Y_{ij})/2$, implying by (3c) that $\hat{W}(Y_{ij})/4 \leq \hat{W}(R_{ij})$. All transmitters of R_{ij} are in $D_{ij\varGamma}$ and all receivers of R_{ij} are in $D_{ij\varLambda}$ or vice versa. The distance between any receiver and transmitter is larger than 1. Thus,

$U(v) \cap R_{ij} = \{v\}$ for each $v \in R_{ij}$, implying that $W(v, R_{ij}) = W(v, \{v\}) = \hat{W}(v)$. Therefore, $\hat{W}(R_{ij}) = W(Y_{ij}, R_{ij})$. In summary, it holds that $W(Y_{ij}, I^*(Y_{ij})) \leq \hat{W}(Y_{ij}) \leq 4W(Y_{ij}, R_{ij})$.

Lemma 4. *The set $\tilde{I}$ is a $\big(5k^2/(k-2)^2\big)$-approximate solution.*

Proof. There is exactly one set A_i and exactly one set B_j that contains any given point in $\tau(V) \cup \rho(V)$. For each vertex $v \in V$, there are at most two sets A_i and at most two sets B_j that contain $\tau(v)$ or $\rho(v)$. Thus, there are at least $(k-2)^2$ pairs (i,j) such that D_{ij} contains both $\tau(v)$ and $\rho(v)$. In notation, $|\{(i,j) \mid v \in Z_{ij}\}| \geq (k-2)^2$ for all $v \in V$.

Let $I^*_{ij} = I^*(V) \cap Z_{ij}$. As the contributions are nonincreasing, we have $W(v, I^*(V)) \leq W(v, I^*_{ij})$ for all $v \in I^*_{ij}$ and $W(Z_{ij}, I^*(V)) \leq W(Z_{ij}, I^*_{ij})$ for all Z_{ij}, implying that $(k-2)^2 W(V, I^*(V)) \leq \sum_{i,j} W(Z_{ij}, I^*(V)) \leq \sum_{i,j} W(V, I^*_{ij})$. Thus, there is a pair (i,j) satisfying $(k-2)^2 W(V, I^*(V)) \leq k^2 W(V, I^*_{ij})$. As the sets X_{ij} and Y_{ij} partition Z_{ij}, we have $W(V, I^*_{ij}) = W(X_{ij}, I^*_{ij}) + W(Y_{ij}, I^*_{ij}) \leq W(X_{ij}, I^*(X_{ij})) + W(Y_{ij}, I^*(Y_{ij}))$. By Lemmata 2 and 3, we obtain

$$\begin{aligned} W(X_{ij}, I^*(X_{ij})) + W(Y_{ij}, I^*(Y_{ij})) &\leq W(V, S_{ij}) + 4W(V, R_{ij}) \\ &\leq (1+4) \max\{W(V, S_{ij}), W(V, R_{ij})\} \leq 5W(V, \tilde{I}). \end{aligned}$$

Remarks. As the above divide-and-conquer approach gives a constant-factor approximation, it is of interest whether a similar result can be obtained by considerably simpler greedy algorithms. However, this does not seem to be the case. First, maximal independent sets may be arbitrarily small: $K_{1,n}$ is a 1-local conflict graph with a maximal independent set of size 1 and an independent set of size n. Second, vertices which are part of a large independent set may also be relatively rare: take a complete graph on N^2 vertices and remove the edges of a clique, possibly with some extra edges.

The next section proves that MWIS in local conflict graphs does not admit a PTAS unless P = NP. Note that if we restricted to the subfamily of N-local conflict graphs where there is no vertex v with $d(\tau(v), \rho(v)) > 1$, a simplification of our algorithm would give a PTAS for MWIS.

5 Proof of Theorem 2

The maximum-weight directed cut problem is defined as follows: given a directed graph $G = (V, A)$ and a nonnegative weight $w(a)$ for each arc $a \in A$, find a subset $S \subseteq V$ such that the total weight of the resulting cut $\delta^+(S) = \{(u,v) \in A \mid u \in S, v \notin S\}$ is maximised. The problem is APX-complete already in the unweighted case [16, 17].

We show that for any N, approximating MWIS in N-local conflict graphs within factor α implies approximating maximum-weight directed cut in an arbitrary directed graph within factor α. The reduction is similar to the one used by

Chvátal and Ebenegger [18]; applied here, the reduction actually shows that the underlying undirected graph of the directed line graph of an arbitrary directed graph is a 1-local conflict graph.

Consider an instance of maximum-weight directed cut $G = (V, A)$ with the arc weights $w(a)$. Without loss of generality, we may arbitrarily relabel the vertices so that $V = \{(0, 3), (0, 6), \ldots, (0, 3|V|)\} \subseteq \mathbb{R}^2$. Construct an N-local conflict graph $G' = (V', E')$ with weights w' as follows. Let $V' = A$ and $E' = \{\{(t, u), (u, v)\} \mid (t, u) \in A, (u, v) \in A\}$. For each $(u, v) \in A$, let $\tau((u, v)) = u$, $\rho((u, v)) = v$, and $w'((u, v)) = w((u, v))$. The construction is a valid N-local conflict graph for any $N \geq 1$.

Let $S \subseteq V$ be a directed cut of G. Now, $\delta^+(S)$ is an independent set of the same weight in G' because there is no pair of arcs $(t, u), (u, v)$ in $\delta^+(S)$. Conversely, let $I' \subseteq V'$ be an independent set in G'. Let S consist of all transmitters of the vertices in I'. Note that S contains no receiver of a vertex in I'. Thus, $I' \subseteq \delta^+(S)$, i.e., S defines a directed cut with weight at least that of I'.

It follows that if W^* is the maximum weight of a directed cut in G, there is an independent set with weight W^* in G'. An α-approximation algorithm for MWIS finds an independent set of weight at least W^*/α in G', which transforms to a directed cut with weight at least W^*/α in G. □

6 Link Scheduling

We conclude this paper by considering fractional covering by independent sets in local conflict graphs. For an independent set I, we let $I(v) = 1$ if $v \in I$ and $I(v) = 0$ if $v \notin I$.

Definition 2. *In the* link scheduling problem*, the input consists of an N-local conflict graph (G, τ, ρ) and a* requirement $r(v) \geq 0$ *for each vertex v in G. The task is to*

$$\begin{array}{lll} \textit{minimise} & \sum_I x(I) & \\ \textit{subject to} & \sum_I I(v)x(I) \geq r(v) & \textit{for all } v, \\ & \qquad x(I) \geq 0 & \textit{for all } I, \end{array}$$

where v ranges over all vertices in G and I ranges over all independent sets in G. The value $L = \sum_I x(I)$ is called the length of the schedule*.*

If G is a conflict graph and $r(v)$ is the amount of data that has to be transmitted on the link from $\tau(v)$ to $\rho(v)$, the link scheduling problem corresponds to finding an optimal schedule of data transmissions in a wireless communication network. In this setting, the vector x is interpreted as a schedule that assigns to independent set I the time slice $x(I)$.

The special case $r(v) = 1$ for all v corresponds to fractional colouring; the minimum schedule length is the fractional chromatic number of G.

Example 1. Consider the example in Fig. 3, and let $r(v) = 1$ for each v. Let $I_1 = \{\mathrm{b, c, d, g}\}$, $I_2 = \{\mathrm{b, c, h}\}$, $I_3 = \{\mathrm{a, f, h}\}$, $I_4 = \{\mathrm{a, e, f}\}$, and $I_5 = \{\mathrm{d, e, g}\}$; each of these sets is an independent set. Choose $x(I_1) = x(I_2) = x(I_3) = x(I_4) = x(I_5) = 1/2$. Now x is a solution to the link scheduling problem, and the length of the schedule equals $5/2$.

The number of variables in the above LP may be exponential in the size of the input. However, it is not necessary to construct the LP explicitly: the link scheduling problem can be solved $(\alpha + \epsilon)$-approximately in polynomial time for any $\epsilon > 0$ as long as there is a polynomial-time α-approximation algorithm for finding the maximum-weight independent set of G for arbitrary weights [2, 3]. Thus, Theorem 1 has the following corollary.

Corollary 1. *The link scheduling problem for N-local conflict graphs admits a polynomial-time* $(5 + \epsilon)$*-approximation algorithm for any constants* $\epsilon > 0$ *and N.*

The main result of this section shows that approximating beyond a certain constant factor remains hard.

Theorem 6. *The link scheduling problem for N-local conflict graphs admits no PTAS for any N unless P = NP.*

We begin with the following lemma.

Lemma 5. *There are constants* k, $c > 1$, *and* Δ *such that the following problem is NP-hard: Given a graph with maximum vertex degree at most* Δ, *decide whether its fractional chromatic number is at most* k *or at least* ck.

Proof. Khot [5] established that it is NP-hard to colour a k-colourable graph with $k^{\log(k)/25}$ colours for all sufficiently large constants k, even for graphs of bounded degree. In fact, Khot's proof shows that distinguishing between the following two cases is NP-hard for graphs of bounded degree: (i) There is a k-colouring. Thus, also the *fractional* chromatic number is at most k. (ii) The ratio of the number of vertices to the maximum size of an independent set is at least $k^{\log(k)/25}$. Thus, the *fractional* chromatic number is at least $k^{\log(k)/25}$.

Proof of Theorem 6. We show that a PTAS for link scheduling in N-local conflict graphs can be used to solve the NP-hard problem in Lemma 5. Let $G = (V, E)$ be an arbitrary graph with maximum vertex degree at most Δ. Label the vertices so that $V \subseteq \mathbb{Z}$. We associate with each $v \in V$ three points, v_1, v_2, and v_3, by setting $v_j = (3v, 3j)$. No unit disk contains more than one point.

Construct an instance of the link scheduling problem: an N-local conflict graph $G' = (V', E')$ and the corresponding requirements $r(v')$ for each $v' \in V'$. We refer to the elements $v' \in V'$ as *links*, the intuition being that they correspond to a pair of transmitter and receiver; we reserve the word *vertex* for the elements $v \in V$.

For each vertex $v \in V$, introduce two links $(v_1, v_2) \in V'$ and $(v_2, v_3) \in V'$ with the requirements $r((v_1, v_2)) = k - 1$ and $r((v_2, v_3)) = 1$. For each edge $\{u, v\} \in$

E, introduce two links $(u_2, v_2) \in V'$ and $(v_2, u_2) \in V'$ with the requirements $r((u_2, v_2)) = r((v_2, u_2)) = 1$. Finally, let $\tau((x, y)) = x$, $\rho((x, y)) = y$, and $E' = \{\{(x, y), (y, z)\} \mid (x, y) \in V', (y, z) \in V'\}$.

Select a positive $\epsilon' < \min\{(\sqrt{c}-1)/2, (1-1/\sqrt{c})/(2k\Delta)\}$, and use the PTAS to solve the constructed link scheduling instance within factor $(1+\epsilon')$. Let the length of the schedule be L.

Let the fractional chromatic number of G be χ_f. If $\chi_f \leq k$, we may use a fractional colouring of size k to construct a feasible schedule of length k for the link scheduling instance as follows. Interpret the fractional colouring as a schedule in which the vertices may be active or inactive. Without loss of generality we may assume that the schedule is exact; that is, each vertex is active for exactly 1 time unit. Whenever $v \in V$ is active, transmit data on links (v_2, v_3) and (v_2, u_2) for each vertex u adjacent to v in G; this is possible since the vertices adjacent to v cannot be active and thus are not transmitting at the same time. Whenever $v \in V$ is inactive, transmit data on (v_1, v_2). Note that each vertex is active for 1 time unit and inactive for $k-1$ time units, implying that the requirements $r(v')$ are met. Observe that if $\chi_f \leq k$, we have $L \leq (1+\epsilon')k$ since the length of the schedule cannot be more than k and we solved the link scheduling instance within factor $(1+\epsilon')$.

On the other hand, if $L \leq (1+2\epsilon')k$, we may use a schedule of length $(1+2\epsilon')k$ to construct a fractional colouring. Let $T(v')$ be the union of all time intervals when the link $v' \in V'$ is active in the schedule, and let $|T(v')|$ be the total length of these time intervals. Consider an arbitrary vertex $v \in V$. Let $V'(v)$ be the set of all links $(v_2, u_2) \in V'$ with $u \in V$. We have $T((v_1, v_2)) \cap T(v') = \emptyset$ for all $v' \in V'(v)$. Furthermore, $|T(v_1, v_2)| \geq k-1$, and the total schedule length equals $k + 2\epsilon' k$. Thus, $|\bigcup_{v' \in V'(v)} T(v')| \leq 1 + 2\epsilon' k$. On the other hand, $|T(v')| \geq 1$ for each $v' \in V'(v)$. Observe that $|\bigcap_{v' \in V'(v)} T(v')| \geq 1 - 2\epsilon' k |V'(v)|$ since each new edge introduced to the intersection shortens the intersection by at most $2\epsilon' k$ units. Also observe that $|V'(v)| \leq \Delta$.

A non-isolated vertex v in the fractional colouring problem may be active in time intervals $\bigcap_{v' \in V'(v)} T(v')$ because none of its neighbours may be active at the same time; for an isolated vertex (the case of an empty $V'(v)$), colouring is trivial. Thus, we have a partial fractional covering of length $(1+2\epsilon')k$ that covers each vertex for at least $1 - 2\epsilon' k \Delta$ units of time. Multiply all time assignments by $1/(1-2\epsilon' k\Delta)$ to obtain a fractional colouring of size $(1+2\epsilon')k/(1-2\epsilon' k\Delta) < (1+(\sqrt{c}-1))k/(1-(1-1/\sqrt{c})) = ck$ that covers each vertex for at least 1 unit of time. Thus, $\chi_f < ck$.

To summarise, $\chi_f \leq k$ implies $L \leq (1+\epsilon')k$ and $\chi_f \geq ck$ implies $L > (1+2\epsilon')k$. This shows that we can use a PTAS to distinguish in polynomial time between the two cases in Lemma 5.

Remark. It is not known whether Theorem 6 could be obtained as a simple corollary of Theorem 2. For example, the conversion method by Erlebach and Jansen [19] cannot be applied directly as it requires that the family of graphs is closed not only under the deletion of vertices but also under the duplication of vertices.

Acknowledgements. We thank Harri Haanpää, Patrik Floréen, Pekka Orponen, and Marja Hassinen for useful discussions and comments.

This work was supported in part by the Academy of Finland, Grants 116547, 117499, 202203, 202204, and 202205, by Helsinki Graduate School in Computer Science and Engineering (Hecse), and by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778.

References

1. Jain, K., Padhye, J., Padmanabhan, V.N., Qiu, L.: Impact of interference on multi-hop wireless network performance. Wireless Networks **11**(4) (2005) 471–487
2. Jansen, K.: Approximate strong separation with application in fractional graph coloring and preemptive scheduling. Theoretical Computer Science **302**(1–3) (2003) 239–256
3. Young, N.E.: Sequential and parallel algorithms for mixed packing and covering. In: Proc. 42nd Annual Symposium on Foundations of Computer Science (FOCS, Las Vegas, NV, USA, October 2001), Washington, DC, USA, IEEE Computer Society (2001) 538–546
4. Håstad, J.: Clique is hard to approximate within $n^{1-\epsilon}$. Acta Mathematica **182** (1999) 105–142
5. Khot, S.: Improved inapproximability results for MaxClique, chromatic number and approximate graph coloring. In: Proc. 42nd Annual Symposium on Foundations of Computer Science (FOCS, Las Vegas, NV, USA, October 2001), Washington, DC, USA, IEEE Computer Society (2001) 600–609
6. Lund, C., Yannakakis, M.: On the hardness of approximating minimization problems. Journal of the ACM **41**(5) (1994) 960–981
7. Goldsmith, A.: Wireless Communications. Cambridge University Press (2005)
8. Krishnamachari, B.: Networking Wireless Sensors. Cambridge University Press, Cambridge, UK (2005)
9. Suomela, J.: Approximability of identifying codes and locating-dominating codes. Information Processing Letters **103**(1) (2007) 28–33
10. Doyle, P.G., Snell, J.L.: Random Walks and Electric Networks. The Mathematical Association of America, Washington, DC, USA (1984)
11. Halldórsson, M.M.: Approximations of independent sets in graphs. In: Proc. Workshop on Approximation Algorithms for Combinatiorial Optimization Problems (APPROX, Aalborg, Denmark, July 1998). Volume 1444 of Lecture Notes in Computer Science. Berlin, Germany, Springer-Verlag (1998) 1–13
12. Halldórsson, M.M.: Approximations of weighted independent set and hereditary subset problems. Journal of Graph Algorithms and Applications **4**(1) (2000) 1–16
13. Erlebach, T., Jansen, K., Seidel, E.: Polynomial-time approximation schemes for geometric intersection graphs. SIAM Journal on Computing **34**(6) (2005) 1302–1323
14. Hochbaum, D.S., Maass, W.: Approximation schemes for covering and packing problems in image processing and VLSI. Journal of the ACM **32**(1) (1985) 130–136

15. Hunt III, H.B., Marathe, M.V., Radhakrishnan, V., Ravi, S.S., Rosenkrantz, D.J., Stearns, R.E.: NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. Journal of Algorithms **26**(2) (1998) 238–274
16. Khanna, S., Motwani, R., Sudan, M., Vazirani, U.: On syntactic versus computational views of approximability. SIAM Journal on Computing **28**(1) (1999) 164–191
17. Papadimitriou, C.H., Yannakakis, M.: Optimization, approximation, and complexity classes. Journal of Computer and System Sciences **43**(3) (1991) 425–440
18. Chvátal, V., Ebenegger, C.: A note on line digraphs and the directed max-cut problem. Discrete Applied Mathematics **29**(2–3) (1990) 165–170
19. Erlebach, T., Jansen, K.: Conversion of coloring algorithms into maximum weight independent set algorithms. Discrete Applied Mathematics **148**(1) (2005) 107–125