Maximum Weight Independent Sets in an Infinite Plane

Jarno Nousiainen, Jorma Virtamo, Pasi Lassila
jarno.nousiainen@tkk.fi, jorma.virtamo@tkk.fi, pasi.lassila@tkk.fi
Department of Communications and Networking
Helsinki University of Technology
P.O.Box 3000, FI-02015 TKK, Finland

***Abstract*—We study the maximum weight independent sets of links between nodes distributed as a spatial Poisson process in an infinite plane. Three different definitions of the weight of a link are considered, leading to slight variations of what is essentially a spatial reuse problem in wireless multihop networks. A simple Boolean interference model is assumed with the interference radius equaling the transmission radius. We study both the case where the transmission radius is fixed and the case where it can be reduced (by power control) so as to just reach the destination to minimize the interference. For the case of a fixed transmission radius, we give asymptotic results for the low density regime and present a rudimentary analysis for the high density asymptotics. The main contribution of this paper is in the numerical results for the maximum weight for the considered infinite networks and in thus establishing some previously unknown parameters of stochastic geometry. For instance, we find that in the unweighted case, just counting the number of independent links, the maximum possible packing is 0.322 links per node attained with the mean neighborhood size of 2.73. The results are obtained by the Moving Window Algorithm that is able to find the maximum weight independent set in a strip of limited height but unlimited length. By studying the results as the function of the height of the strip, we are able to extrapolate to the infinite plane.**

I. Introduction

We study the following fundamental problem: given a network in an infinite plane with nodes placed according to a planar Poisson process and connected to each other via wireless links with given weights and a given transmission range, what is the total weight (per unit area or per node) of the maximum weight set of non-interfering links under the Boolean interference model, the interference range being assumed to equal the transmission range? In the terminology of graph theory, the set of non-interfering links maps to what is called an independent set; we use these terms interchangeably. The studied problem is purely one of stochastic geometry, but it has a close connection to modeling the capacity of large-scale wireless multihop networks, and this paper is written with that viewpoint in mind.

The question arises, e.g., in the context of analyzing the local forwarding capacity of massively dense wireless networks, see [2], [6], motivated by the future applications of large scale sensor networks. From a single node's perspective the surrounding network appears as an infinite network of randomly placed wireless nodes. The solution to the maximum weight problem gives the maximum *instantaneous* forwarding capacity in the neighborhood of the considered node, which in turn sets an upper bound for the local sustainable mean forwarding capacity, i.e., the average rate at which traffic can be "moved" in a given direction, see [1], [7]. (Note that the maximum weight independent set cannot be used repeatedly for forwarding traffic because it consists of independent, isolated links that do not form a connected network.) Results like these yield useful information about the achievable gains from utilizing optimal global coordination in multihop communications, and thus complement the well-known scaling results for the capacity of multihop networks, see [5] and the more recent results in [4] on random networks.

In this paper, we consider the above problem with three different kinds of weights: A) unweighted, i.e., each link has the weight one, B) weighted by the length of the projection of the link in a given direction, e.g., on the x-axis, and C) weighted by the length of the link. In case A, only the number of links in an independent set, i.e., the size of the set is counted. The problem is then just the maximum independent set problem. Case B corresponds to the above discussion of forwarding capacity, as the length of the projection tells how much the traffic is moved in the x-direction or the amount of x-progress. Case C is similar but without sense of direction.

Each of these three cases defines a challenging problem in stochastic geometry. Finding the maximum independent set for a given finite graph is known to be an NP-complete problem, and in our case the graph is even infinite. However, in the present problem the graphs are not arbitrary but the interference is localized; links separated far enough do not interfere with each other. In our approach, we will efficiently exploit this special structure of the graph.

We study the problem also with two different assumptions on the transmission range. In the basic case the transmission range is assumed to be fixed. Additionally, we examine the case where the transmission range can be adjusted (by power control) up to a given maximum radius. The idea is that not every sending node uses the maximum radius but a radius just large enough to reach the receiving node, thereby minimizing the interference.

Our analytical contribution comprises of asymptotic analyses of the different cases A, B, and C with a fixed transmission radius. The asymptotics are analyzed both when the mean number of neighbors tends to zero and to infinity (loosely speaking, the low and high density asymptotics). In the previous case, the analysis is simple and the results are exact. In the high density limit, the problem is more intricate, and we present only a rudimentary analysis, which however, we believe, captures the essential dependency. The asymptotic behaviors at both ends yield insight on the behavior of a curve also in the intermediate range.

The main contribution of this paper is in the numerical results representing the total weight of the maximum independent sets of the studied systems as a function of the mean neighborhood size. In the case of a fixed transmission radius, these curves have a maximum, which is of special interest in defining the best that can be obtained as well as the optimal value of the transmission radius in relation to the mean distance between the nodes. In the case of an adjustable transmission range (up to a given maximum) no maximum exists, since all the independent sets that are feasible with a given maximum range are feasible when the maximum range is made larger. Similarly, we see that the results with an adjustable range are always as good as or better than when only the maximum range can be used since anything that can be done with a fixed range can be done with an adjustable range.

The results are obtained by applying the Moving Window Algorithm (MWA), originally developed in a different context in [8], but fine-tuned for the present application. The algorithm is able to find the maximum weight independent set (and its weight) of any network realization in a strip of limited height but unlimited length. The algorithm works by moving a window along the strip, say from left to right, starting from an initial position. The window is high enough to cover the strip in the vertical direction and wide enough so that no two links on different sides of the window interfere with each other. The algorithm maintains a tree representing all possible conflict-free sets within the window, with each leaf of the tree being assigned a value equal to the cumulative weight of the maximum weight independent set so far, conditioned on the choice of the conflict-free set in the window corresponding to the leaf in question. The tree and the assigned values are updated as the window moves. By studying the results as a function of the height of the strip we are able to extrapolate to the infinite plane.

We have applied the MWA algorithm for case B (weighted with x-progress) in our earlier paper [7]. In the present paper, we generalize the algorithm and extend the results for the cases A and C, and additionally consider the impact of power control. The algorithm is computationally demanding, especially for large windows and high mean neighborhood sizes, where, besides long running times, the memory requirements of the tree become a bottleneck. The results are presented as far as it has been possible to proceed with a computer with a 4 GB memory. For cases A and B the explored range covers the points where the curves with fixed transmission range reach their maxima; for case C this is still unattainable. Notably, for the unweighted case A without power control, we are able to establish as a new result that the maximum size of an independent set is 0.322 links per node attained at mean neighborhood size of 2.73.

The rest of this paper is organized as follows. In Section II, we introduce the notation and present scaling considerations to reduce the unknowns to the minimum. The asymptotic results are derived in Section III. In Section IV, we describe the MWA algorithm. The numerical results are presented in Section V, and we conclude in Section VI.

II. Notation and scaling considerations

As discussed in the Introduction, we consider a system where the locations of nodes are assumed to obey a planar Poisson process. The intensity of this process is denoted by λ. In the basic case, the transmission radius is fixed and denoted by R. In the case of power control, the transmission range can be chosen up to a maximum R_{max}. Thus, there is a feasible link between any two nodes within R (or R_{max}) from each other. Interference is modeled using the Boolean interference model with the interference radius equaling the (used) transmission radius. A transmission interferes with all the receptions inside its range, implying that a link is only possible if the receiver hears exactly one transmission (excluding simultaneous transmission and reception).

The beauty of the system with the Boolean interference model is that there are only two parameters in the model: λ and R (or R_{max}). And there is only one (independent) dimensionless parameter that can be formed from these two system parameters. We use the most natural one, viz. the mean number of neighbors within the range, denoted by ν,

$$\nu(\lambda, R) = \pi \lambda R^2,$$

with R replaced with R_{max} in the case of an adjustable transmission range. With proper scaling considerations, as detailed below, all unknown functions of the two system parameters can be reduced to functions of this single variable.

Our goal is to find the maximum weight (or size) of an independent set of links per unit area in the following cases[1]:

A. Unweighted,
B. Weighted by the x-progress of the links,
C. Weighted by the length of the links.

In the first case (A), we aim to calculate the maximum number of links per unit area, denoted by $U(\lambda, R)$. By dimensional analysis, we can write

$$U(\lambda, R) = \lambda\, u(\nu(\lambda, R)), \tag{1}$$

where $u(\nu)$ is a dimensionless function of a single variable to be determined. In fact, $u(\nu(\lambda, R))$ represents the number of links per node.

[1]The graph theoretic correspondence of our problem to the well known maximum weight independent set problem results from the following mapping between the graph model of the wireless network and the so-called interference graph. An independent set of links in a wireless network is an independent set (of vertices) in the network's interference graph, where each link corresponds to a vertex, and two vertices are adjacent if the links interfere with each other. Determining the maximum weight of an independent set of links is equivalent with finding the maximum weight independent set in the interference graph.

The second case (B) differs from the previous as we are interested in the maximum density of progress, $U_x(\lambda, R)$, that is the maximum x-progress (the progress of the link in a fixed direction) per unit area, and by dimensional analysis we get

$$U_x(\lambda, R) = \sqrt{\lambda}\, u_x(\nu(\lambda, R)), \tag{2}$$

where $u_x(\nu)$ is another unknown dimensionless function.

In the last case (C) we are interested in the total length of the links per unit area, U_l, that can be expressed with the help of a yet another dimensionless function u_l, exactly as in (2),

$$U_l(\lambda, R) = \sqrt{\lambda}\, u_l(\nu(\lambda, R)). \tag{3}$$

In short, our task is to find the dimensionless functions $u(\nu)$, $u_x(\nu)$, and $u_l(\nu)$. In the sequel, speaking generally we use $u_*(\nu)$ to represent any of these functions. The notation $U_*(\lambda, R)$ is used similarly.

III. Low and high density asymptotics

In this section we consider the asymptotic behavior of the dimensionless functions $u_*(\nu)$ when the mean node degree approaches zero or infinity. The analysis when ν approaches infinity is rudimentary but believed to capture the essential dependency. The obtained theoretical limits provide useful insight and will be compared to simulations in Section V.

A. Asymptotics in the limit $\nu \to 0$

In the unweighted case (A), a general upper bound for the function U is given by

$$U(\lambda, R) \le \frac{1}{2}\, \lambda(1 - e^{-\nu}).$$

The reasoning with the above inequality is that there are on the average λ nodes per unit area, and that one obviously gets an upper bound for $U(\lambda, R)$ if each node can freely choose the neighbor to form a link with, without any restrictions imposed by other links. The factor $\frac{1}{2}$ accounts for the fact that it takes two nodes to form a link. The parenthetical expression is the probability that a node has a neighbor. Written in terms of $u(\nu)$ the upper bound takes the form

$$u(\nu) \le \frac{1}{2}(1 - e^{-\nu}).$$

It is also obvious that asymptotically when $\nu \to 0$ the upper bound becomes tight, since in the rare cases when a node has a neighbor within its transmission radius, they can indeed form a link with a high probability without any other link interfering. In this asymptotic regime the probability $(1 - \exp\{-\nu\}) \approx \nu$, and we have

$$u(\nu) \to \frac{1}{2}\, \nu. \tag{4}$$

For the case weighted by the x-progress (B), the general upper bound for the function U_x becomes

$$U_x(\lambda, R) \le \frac{1}{2}\, \lambda R\, X(\nu),$$

where $X(\nu)$ is the mean distance (in units of R) from a randomly chosen node to its most distant neighbor node in the x-direction, i.e., absolute value of the x-distance (if there is none, the distance is taken to be zero). For $u(\nu)$, we have

$$u_x(\nu) \le \frac{1}{2}\sqrt{\frac{\nu}{\pi}}\, X(\nu).$$

When ν is small, $X(\nu) \to 4\nu/(3\pi)$, where $4/(3\pi)$ is the mean x-distance to a neighbor, and ν is the approximate probability of having a neighbor. Thus, we have

$$u_x(\nu) \to \frac{2}{3}\left(\frac{\nu}{\pi}\right)^{3/2}. \tag{5}$$

The third case (C) is similar to the second case, but we have to replace the x-distance between the nodes by the actual distance, $L(\nu)$. Hence, the limits become

$$L(\nu) \to \frac{2}{3}\, \nu, \quad \text{and} \quad u_l(\nu) \to \frac{1}{3\sqrt{\pi}}\, \nu^{3/2}, \tag{6}$$

when $\nu \to 0$.

B. Asymptotics in the limit $\nu \to \infty$

We now turn our attention to how $U(\lambda, R)$ behaves for large λ when R is considered to be fixed and present a plausible reasoning for the asymptotics. The starting observation is that if the end points of a link can be arbitrarily placed on a continuous plane, then the most efficient way of packing links is to form vertical columns. The claim is most credible in the case with x-progress (B). There has to be a distance larger than R between two consecutive links, as illustrated in Figure 1, but the vertical distance between the links can be small. In fact, the Boolean interference model (unrealistically) sets no limit on how densely the links can be vertically packed: two parallel links of maximal length R, however close, never interfere with each other. This suggests that for a very high λ, when there are nodes almost everywhere, a good strategy is to try to form vertical columns.

In cases A and C, that are undirected, the packing can be done even more efficiently by changing the direction of every other column. This way a small distance ε is enough between the columns as the endpoints near each other are all either transmitters or receivers.

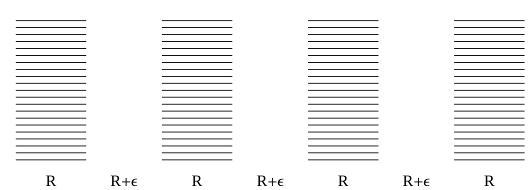

Fig. 1. On a continuous plane links can be efficiently stacked in vertical columns (case B). In cases A and C, links can be packed even tighter since ε margin is enough between columns transmitting in alternating directions.

The next step is to estimate the expected vertical distance between the links. Based on the above observation we consider a naive model where, starting from a vertical link of length R, the end points of the next link above are determined independently by proceeding in the vertical direction in the

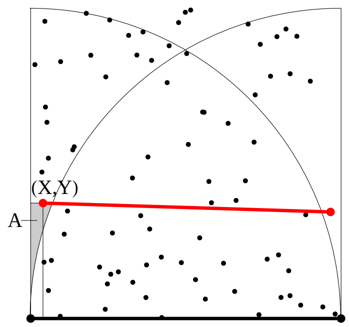

Fig. 2. The simplified model for estimating the vertical distance between the stacked links.

shown areas of Figure 2 until next node (from the Poisson process) is found.

The width x of the area between the vertical line and the circle is for small heights y approximately parabolic, $x \approx y^2/2R$. Denote the coordinates (random variables) of the node by (X, Y). Since $A \sim \text{Exp}(\lambda)$ and $A \approx Y^3/(6R)$, we have the complementary distribution function of Y,

$$\text{P}\{Y > y\} = \text{P}\{A > \frac{y^3}{6R}\} = e^{-\frac{\lambda y^3}{6R}},$$

from which

$$\begin{aligned}\text{E}[Y^n] &= \int_0^\infty e^{-\frac{\lambda y^{3/n}}{6R}} dy = \left(\frac{6R}{\lambda}\right)^{n/3} \int_0^\infty e^{-y^{3/n}} dy \\ &= \left(\frac{6R}{\lambda}\right)^{n/3} \Gamma\Big(1 + \frac{n}{3}\Big).\end{aligned}$$

In particular we have

$$\text{E}[Y] = \Gamma\Big(\frac{4}{3}\Big)\Big(\frac{6R}{\lambda}\Big)^{1/3}, \quad \text{E}[Y^n] = \frac{\Gamma(1+\frac{n}{3})}{\Gamma(\frac{4}{3})^n}\,\text{E}[Y]^n,$$

whence the variance is

$$\text{V}[Y] = \text{E}[Y^2] - \text{E}[Y]^2 = \left(\frac{\Gamma(\frac{5}{3})}{\Gamma(\frac{4}{3})^2} - 1\right)\text{E}[Y]^2 \approx 0.132\text{E}[Y]^2.$$

The distribution of X is determined by that of Y, $X \sim \text{Uniform}(0, Y^2/(2R))$, from which

$$\text{E}[X] = \text{E}[\text{E}\,[X\,|\,Y]] = \text{E}[\frac{Y^2}{4R}] = \frac{\Gamma(\frac{5}{3})}{4R\Gamma(\frac{4}{3})^2}\,\text{E}[Y]^2$$

and

$$\begin{aligned}\text{V}[X] &= \text{E}[\text{E}\left[X^2\,|\,Y\right]] - \text{E}[X]^2 = \text{E}[\frac{1}{3}(\frac{Y^2}{2R})] - \text{E}[\frac{Y^2}{4R}]^2 \\ &= \frac{\frac{1}{3}\Gamma(\frac{7}{3}) - \frac{1}{4}\Gamma(\frac{5}{3})^2}{4\Gamma(\frac{4}{3})^4 R^2}\,\text{E}[Y]^4 \approx \frac{0.0759}{R^2}\,\text{E}[Y]^4.\end{aligned}$$

Now, consider the random walk $\mathbf{X}_n = \sum_{i=1}^n (X_i, Y_i)$, $n = 1, 2, \ldots$. When $\lambda \to \infty$ this random walk tends to a deterministic motion along the vertical line with constant rate. This is because both $\text{E}[X]$ and $\text{V}[Y]$ go to zero quadratically in $\text{E}[Y]$. Thus over a finite interval y, which takes on the average $n = y/\text{E}[Y]$ steps, the expected total displacement in the x-direction is $n\text{E}[X] \sim y\text{E}[Y]$, which goes to zero with $\text{E}[Y]$ as $\lambda \to \infty$. Similarly the total variance of the displacement in the y-direction after n steps is $n\text{V}[Y] \sim y\text{E}[Y]$ and goes to zero as $\lambda \to \infty$ (the total variance of the x-displacement goes to zero even faster as the one step variance $\text{V}[X] \sim \text{E}[Y]^4$).

The fact that the independent random walks of both the end points tend to constant deterministic motion along the vertical lines, in the hindsight justifies considering each step starting from a vertical link of maximal length R; the wiggle and contraction of the added links tend to zero.

Finally, we are able to calculate the asymptotic behavior in the three cases starting from the unweighted one (A). From the above it follows that $\text{E}[Y]$ defines the vertical packing distance. As there is one vertical link in every rectangle of height $\text{E}[Y]$ and width $(1+\varepsilon)R$, cf. Figure 1, the reward per unit area is asymptotically $U(\lambda, R) \approx 1/(R\,\text{E}[Y])$,

$$U(\lambda, R) \approx \frac{1}{\Gamma(\frac{4}{3})}\left(\frac{\lambda}{6R^4}\right)^{1/3}, \quad u(\nu) = \frac{1}{\Gamma(\frac{4}{3})}\left(\frac{\sqrt{6}}{\pi}\,\nu\right)^{-2/3}.$$

Similarly in the weighted case (B), as there is one vertical link of length R in every rectangle of height $\text{E}[Y]$ and width $(2+\varepsilon)R$, the weight per unit area is asymptotically $U_x(\lambda, R) \approx 1/(2\text{E}[Y])$,

$$U_x(\lambda, R) \approx \frac{1}{2\,\Gamma(\frac{4}{3})}\left(\frac{\lambda}{6R}\right)^{1/3}, \quad u_x(\nu) = \frac{1}{2\,\Gamma(\frac{4}{3})}\left(\frac{36}{\pi}\,\nu\right)^{-1/6}.$$

Finally in the third case (C), $U_l(\lambda, R) = 2\,U_x(\lambda, R)$, and $u_l(\nu) = 2\,u_x(\nu)$ as the number of links compared to the packing in Figure 1 can be doubled. Hence,

$$U_l(\lambda, R) \approx \frac{1}{\Gamma(\frac{4}{3})}\left(\frac{\lambda}{6R}\right)^{1/3}, \quad u_l(\nu) = \frac{1}{\Gamma(\frac{4}{3})}\left(\frac{36}{\pi}\,\nu\right)^{-1/6}.$$

This asymptotic behavior of $u_*(\nu)$ presumably gives everywhere an upper bound to the true curve. We return to the comparison with the numerical values later in Section V. For more realistic interference models, one can conjecture that the asymptotic tail of $u_*(\nu)$ comes down more rapidly than for the considered Boolean interference model due to the fact that this model unrealistically allows multiple transmissions just outside the interference range of a receiving node.

Note also that the $\nu \to \infty$ asymptotics do not apply when adjusting the transmission radius, R, is allowed. Though the number of potential neighbors increases with a greater R_{max}, it is always possible to use the previous link configuration unless a better one becomes available. Thus in these cases, $u_*(\nu)$ approaches some limit.

IV. Moving Window Algorithm

In this section we derive an algorithm similar to Retrospective optimization introduced in a study of reservation systems [8].

The algorithm considers a small portion of the network at a time, a rectangular window that moves, and regarding the strip that the moving window covers during a simulation, the result is exact. The algorithm uses a binary tree to enumerate all the possible link combinations in the window area to find the maximum size or weight of an independent set of links per unit area so far conditioned on the choice of the combination of conflict-free links. The length of the simulation is not limited, and the covered strip can be of any desired length. We repeat the simulation for windows of different height to extrapolate the value of the maximum weight per area for an infinitely large network.

Because the height of the window in practice is limited, the top and the bottom of the strip can be connected to diminish the border effect and represent an infinite dimension, see Figure 3. The perimeter of the formed cylinder needs to be large enough for the results to be meaningful. The other direction can be handled by moving the window along the cylindrical network. The width of the window (i.e., the length of the cylindrical window) needs to be large enough for the window to contain all the links that can possibly interfere with the links that are going to enter the window in the future (that is, $3R$ which is the maximum length of two links and a $R+\varepsilon$ margin). The possible combinations of these links are maintained in the binary tree. The links that have already left the window do not affect the possible on/off-state of the links entering the window and can thus be removed by a procedure explained next.

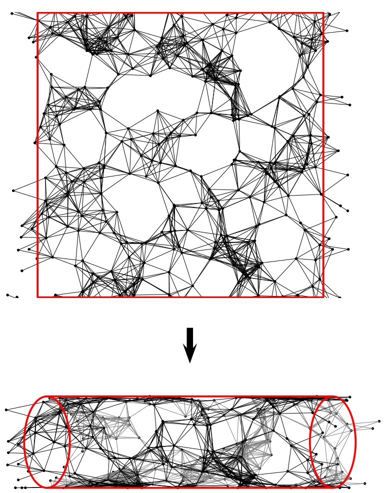

Fig. 3. The top and the bottom of the window are connected together to diminish the border effect. The formed cylinder is moved in the direction of its axis.

A rooted binary tree represents all the possible link combinations in the window area. Every edge of the tree describes whether the link corresponding to that level is active or not, and the value assigned to each leaf shows the maximum size or weight of the independent set thus far (starting from the initial position of the window) conditioned on the combination of active links in the window represented by the leaf. This is illustrated in Figure 4. The figure represents an example of a situation where the simulation of the unweighted case (A) (the values in the vertices represent the size of the independent set of links) has just started, and the first four nodes have entered the window making it possible to form six links. The maximum size of an independent set of links is 2 corresponding to activation set $\{uv, xw\}$ or $\{vu, wx\}$.

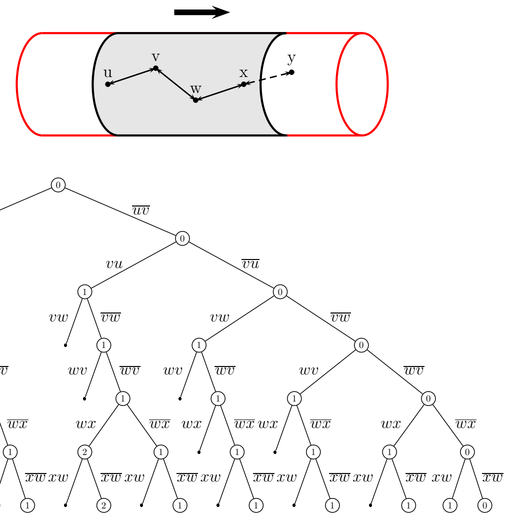

Fig. 4. A window containing 4 links and the corresponding binary tree with 6 levels representing the links (in alphabetical order) in the window.

When the window of Figure 4 is being moved to the right, the first event is the node u leaving the window. Since the entering and exiting links are independent, we can combine the on- and off-branches corresponding to a link whose endpoint has been dropped out of the window and choose the greater values for the new tree. That is, we compare leaves[2] that only differ in the dropped link and choose the maximum of those to be the value of the same node in the new tree where the level corresponding to the dropped link has been eliminated in this way. For example, when the first link to leave the window, uv, is being eliminated from the tree, the leaf with value 2 corresponding to the activation set $\{uv, xw\}$ (first from the left in Figure 4) is compared to the leaf with value 1 corresponding to the activation set $\{xw\}$ (second from the right), and the value of leaf $\{xw\}$ in the new tree in Figure 5 (second from the right) is thus 2. Also the link vu has to be removed from the tree when the node u leaves the window.

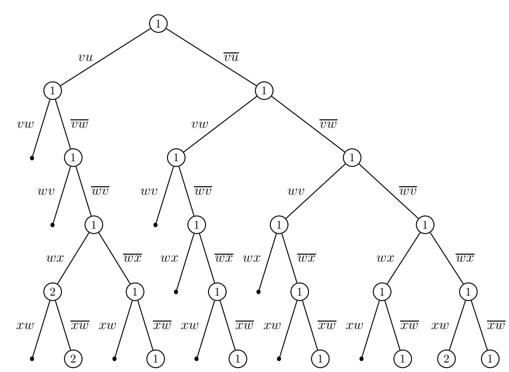

Fig. 5. The binary tree of Figure 4 after the first link to leave the window uv has been removed. In addition to the leaves, also the other vertices have been updated.

[2]Besides the leaves also the other vertices can be updated, but they hold no significance to the final result since the leaves cover all the possible link combinations.

The next event, when moving the window, happens when node y enters and makes two new links possible. These new links are then added to the binary tree, after which the shape of the tree is the same as in Figure 4, but the value assigned to each leaf, except for those with vw or wv, is one higher since in these cases it is possible to use either the link uv or vu that have already exited the window. At this point, it is not explicitly visible which dropped links can be activated. Thus, the maximum value in the tree is the maximum size of the independent set of links so far given the set of active links in the window area. In this way we can generate the network realizations on the fly and progressively find the maximum size or weight of the independent set of links.

We do not maintain information about the links belonging to the maximum weight set, although, this information could be extracted from the algorithm with the cost of used memory. To further minimize the memory requirements the links are removed from the window as soon as they stop interfering with links that are going to enter the window in the future. When a link does not interfere with future links anymore, the information whether the link belongs to the maximum weight independent set is no longer required in the calculations, but the link can be removed from the binary tree maintaining the on/off status of the relevant links. This way the size of the binary tree, that is the bottleneck limiting the usefulness of the algorithm, can be kept as small as possible.

The algorithm limits in no way the length of the simulation in the direction in which the window moves, and when the execution is continued, the result converges without bias towards the true value. When the simulation is ended, the maximum size or weight of the independent set of links is the maximum of the values assigned to the leaves of the binary tree. The simulation is repeated to produce confidence intervals for the value. In the other direction, we have to rely on extrapolation and estimate the maximum value for an infinitely wide cylinder, as discussed in the next section.

A. Extrapolation

This section concentrates on extrapolating the maximum weight per area for the infinite plane from the measurements considering only strips of the network with limited height.

The simulations with the moving window algorithm produce values $u_*(\nu, p)$, where p (in units of R) is the perimeter of the cylinder. For a given ν, several values of p are needed to extrapolate $u_*(\nu)$ with an infinitely wide cylinder. Figure 6 represents $u(p)$ of case A for different values of ν. As seen from the figure, the narrowest cylinders do not give a reliable estimate for larger values of ν. The exact number of values of p required for the extrapolation depends heavily on the case studied as discussed later related to the x-progress case.

The second case (B) with x-progress differs from the others as it is the only directed case. The working principle of the algorithm does not depend on the direction of the traffic, i.e., the direction in which the progress of the maximal independent set is calculated, but it has to be fixed. We have two extremes: the direction is parallel with the direction in which the cylinder moves (along the cylinder) or the progress is calculated perpendicular to the movement of the cylinder (around the cylinder). In the latter case $u(p)$ depends heavily on the number of link columns that we are able to fit around the cylinder. The maxima appear when the perimeter is approximately a multiple of $2R$, meaning that we are able to fit full-length links and the margins $R + \varepsilon$ between them. When the direction of the progress is turned by a right angle, we get more stable results as the vertical distance between the links in a column is more stochastic. This effect is illustrated in Figure 7. The observation supports the assumption made in Section III-B about the most efficient way of packing the links, that is, to form vertical columns.

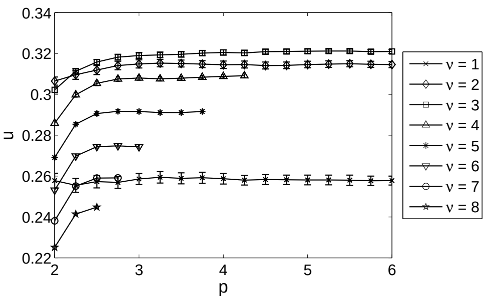

Fig. 6. Function $u(p)$ for different values of ν, and the 95 % confidence intervals.

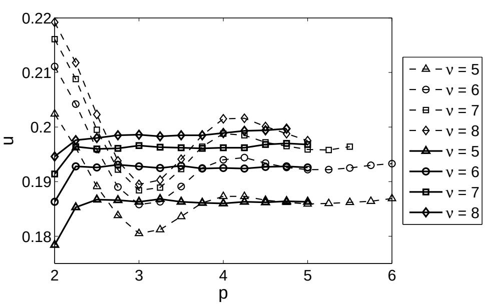

Fig. 7. Function $u_x(p)$ for different values of ν when direction of progress is around the cylinder (dashed lines) and along the cylinder (solid lines).

B. Half-space considerations

Here we study the maximum weight problem by considering the Poisson process in a half-space configuration. We show that the maximum weight over the whole plane can be related to a local additive contribution from a single node added on the border of the half space. The result also allows an alternative way of justifying the MWA algorithm.

Consider the problem of the maximum weight of an independent set of links per area in a half-space configuration, see Figure 8. The border introduces a boundary effect but far from the boundary, inside the body of the Poisson process, the expected total weight per unit area is given by $U_*(\lambda, R)$.

Now, consider moving the boundary incrementally to the right, so that a new area dA is covered. The increase in the total weight of the maximal independent set can be evaluated in two different ways: a) one can think that the slice dA has

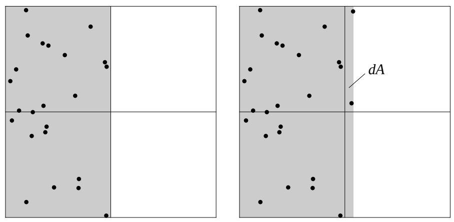

Fig. 8. Poisson process in the half space and an incremental shift of the boundary.

been added in the body, pushing the boundary to the right; then the added weight is $U_*(\lambda, R)\, dA$, b) one can think that the slice has been added to the right boundary introducing new nodes, as shown in Figure 8; In the limit $dA \to 0$, the added nodes are far apart and the increase is the number of nodes, $\lambda\, dA$, times the contribution from a single added node at the boundary, as illustrated in Figure 8. Equating these two yields the average weight per node

$$\frac{1}{\lambda}\, U_*(\lambda, R) = C,$$

where C is the expected total increase of weight due to a single node added at the border. This is in itself an interesting result as it relates mean value over the whole infinite plane to a quantity that has a local character.

To study C, we have to compare two cases. In the first case, there is a node on the border of the network, and in the second there is not. Now, the window starts from a point where one side (right) of the window corresponds to the border of the network and moves away (left) from the border. The top and the bottom of the window are again connected to diminish the border effect and form a cylinder. The effect of the node is the difference in the maximum value in the binary tree in these two simulations. The problem of this formulation is that one simulation produces a single sample instead of some kind of mean, and this causes a large deviation.

From the algorithm point of view, it makes no difference whether the additional node is the first node of the simulation or the last. We can draw multiple samples from a single simulation by assuming that every node entering the window is the node on the border of the network. Not even the differences need to be calculated separately in this case – the sum of the differences is simply the total weight of the independent set (and is given by the maximum leaf value in the tree). Hence, we have returned to the original algorithm.

V. Numerical results

In this section we present the numerical results obtained by the Moving Window Algorithm of Section IV for the three cases: unweighted (A), weighted by x-progress (B), and weighted by length (C). In addition to fixed transmission radius R, we consider transmission radii limited by a maximum value $R \le R_{max}$. In this case the parameter ν is defined to correspond to the mean number of neighbors within the maximum range $\nu(\lambda, R_{max})$.

Figure 9 shows $u(\nu)$ for the unweighted case (A) with both fixed and adjustable transmission radius. With a fixed transmission radius the maximum occurs at $\nu^* = 2.73$ and equals 0.322. The curve with power control is an increasing one as all configurations that are feasible with a given R_{max} are also possible with a greater R_{max}, and being upper bounded by the theoretical maximum of $\frac{1}{2}$, i.e., one link per two nodes, it tends to a limit when $\nu \to \infty$. As can be seen, the limit is relatively close to the theoretical maximum, implying that the maximum gain from a freely adjustable transmission radius is approximately 50 % (30 % at ν^*).

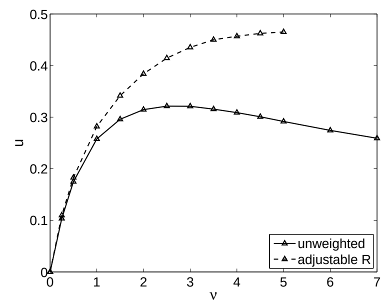

Fig. 9. Function $u(\nu)$ for unweighted case (A) with and without power control.

Figure 10 represents $u_*(\nu)$ for the weighted cases (B and C) with fixed transmission radius and with the possibility to reduce the transmission power to the minimum required. Even though the computational complexity grows with the number of links, it is possible to simulate x-progress with fixed transmission radius up to the optimal size of the neighborhood. The maximum occurs at $\nu^* \approx 10$ and equals 0.20. In the case of a length-weighted set (C), the number of links in the window is doubled compared to the second case, since we have to consider both directions separately. Thus we are not able to find the optimal neighborhood size. As with the unweighted case, the curves corresponding to cases with power control do not have a maximum but are increasing functions of ν tending to a limit when $\nu \to \infty$. Again, the maximum gain from an adjustable transmission radius is close to 50 % in case B.

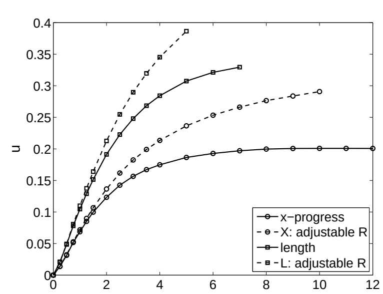

Fig. 10. Functions $u_*(\nu)$ for weighted cases, x-progress (B) and length (C), and for their versions with adjustable transmission radius.

As mentioned, the size of the binary tree places limitations on the feasible simulation parameters (ν and p). Since the process is stochastic, the number of links in the window

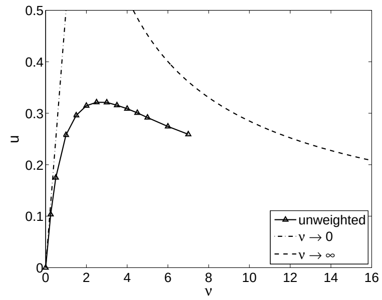

Fig. 11. Numerically evaluated curve for the function $u(\nu)$ along with the low- and high-ν asymptotic curves.

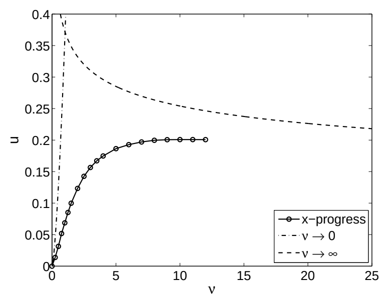

Fig. 12. Numerically evaluated curve for the function $u_x(\nu)$ along with the low- and high-ν asymptotic curves.

may temporarily grow very large, and the size of the tree may exceed the available memory (4 GB). Thus, we are only able to simulate cylinders wide enough until a certain value of ν in each case. An adjustable transmission radius is always computationally more complex than a fixed one since it increases the number of conflict-free link combinations and the size of the tree.

The time needed for the calculations depends on the size of the tree. In the case weighted by the x-progress (B), we do not have to consider links with negative weight, and thus, we are able to do simulations with relative high values of ν. Hence, the number of links in the window stays continuously on a high level, and the simulations are slow. For example, the last points of case B (20 repeats with simulation length of $1000R$) took over 24 h of computer time.

Finally, we compare the numerical results with the asymptotic results of Section III. Figure 11 represents these for $u(\nu)$ and Figure 12 for $u_x(\nu)$. As can be seen from these figures, even the rudimentary analysis seems to yield a plausible asymptotic behavior for large ν. However, the figures also show that asymptotes alone do not characterize the curves accurately in the most interesting parameter area.

VI. Conclusions

We studied the maximum weight independent sets in an infinite plane. This is a problem of stochastic geometry that relates to the question of the largest possible number of simultaneous successful transmissions, i.e., the spatial reuse in wireless multihop networks.

We illustrated the working principles of the Moving Window Algorithm that allowed us to study the problem numerically. The algorithm produces exact results (weight per area) for any network realization in an arbitrarily long strip or, to reduce boundary effects, a cylinder obtained from the strip by joining its upper and lower ends. The network realization can be generated on-the-fly as the window moves, thus enabling unlimited simulations and accurate unbiased estimates. The height of the strip (perimeter of the cylinder) is, however, limited and to obtain results for an infinite plane an extrapolation technique was used.

Three different cases were studied. The first considered the number of simultaneous transmissions per unit area, the second the number of transmissions weighted by the progress of the transmissions in a given direction, and the third the number of transmissions weighted by their lengths. In addition, the effect of power control was studied.

As expected, the problem turned out to be computationally demanding. However, we were able to produce previously unknown numerical results for all the cases, and in two cases to cover the optimal operating region. Only in the third case, with links weighted by their lengths, were we unable to reach the most interesting parameter region, leaving room for more computational science oriented work in the future.

We presented also asymptotic analyses of the systems. The high-density asymptotics are a challenging problem, and our analysis is to be considered as a first attempt. Though capturing the essence of the problem, the analysis can be made more refined and improved in rigor.

Another direction for future research would be to experiment with other possible approaches provided by stochastic optimization in its various incarnations, like simulated annealing or the "packing approach" proposed in [3]. We believe that the MWA algorithm's ability to give exact results in a stripe or cylinder, however, gives it an advantage that is hard to beat.

Acknowledgment

This work was done in the ABI project funded by TEKES, Ericsson, and Nokia Siemens Networks.

References

[1] F. Baccelli, B. Blaszczyszyn, and P. Muhlethaler. An aloha protocol for multihop mobile wireless networks. *IEEE Transactions on Information Theory*, 52(2):421–436, 2006.

[2] R. Catanuto, S. Toumpis, and G. Morabito. Opti{c,m}al: Optical/optimal routing in massively dense wireless networks. *Proceedings of IEEE INFOCOM*, pages 1010–1018, 2007.

[3] M. Durvy and P. Thiran. A packing approach to compare slotted and non-slotted medium access control. In *Proceedings of IEEE INFOCOM*, pages 1–12, April 2006.

[4] M. Franceschetti, O Dousse, D.N.C. Tse, and P. Thiran. Closing the gap in the capacity of wireless networks via percolation theory. *IEEE Transactions on Information Theory*, 53(3):1009–1018, 2007.

[5] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[6] E. Hyytiä and J. Virtamo. On traffic load distribution and load balancing in dense wireless multihop networks. *EURASIP Journal on Wireless Communications and Networking*, 2007:Article ID 16932, 15 pages, 2007.

[7] J. Nousiainen, J. Virtamo, and P. Lassila. Forwarding capacity of an infinite wireless network. In *Proceedings of ACM MSWiM*, pages 177–184, Vancouver, CA, Oct 2008.

[8] J.T. Virtamo. A model of reservation systems. *IEEE Transactions on Communications*, 40(1):109–118, 1992.