Publication VII

Jarno Nousiainen, Pasi Lassila and Jorma Virtamo. Instantaneous Forwarding Capacity under the SINR Threshold Interference Model. In *Proceedings of the 11th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net 2012)*, pages 67–74, Ayia Napa, Cyprus, June 2012.

Instantaneous Forwarding Capacity Under the SINR Threshold Interference Model

Jarno Nousiainen, Pasi Lassila, Jorma Virtamo
jarno.nousiainen@tkk.fi, pasi.lassila@aalto.fi, jorma.virtamo@tkk.fi
Aalto University, Department of Communications and Networking
P.O.Box 13000, FI-00076 Aalto, Finland

***Abstract*—Spatial reuse is a key aspect of wireless network design, and the choice of an appropriate interference model is important for capturing the intrinsic characteristics of such a network. We address the problem of finding optimal transmission modes, and their capacities, in a network consisting of nodes distributed as a spatial Poisson process in an infinite plane, i.e., such combinations of transmitting links that maximize the instantaneous forwarding capacity of the network. A stochastic optimization method called simulated annealing is used to obtain the results. The approach is applied to a SINR threshold interference model that treats the sum of all the other simultaneously transmitted signals as noise. We find out how the maximum capacity behaves for different network densities, signal attenuation coefficients, and thresholds for the required SINR. These numerical results shed light on the spatial reuse problem in wireless multihop networks. We further characterize the asymptotic behavior of the sum capacity of the optimal combination of transmitting links and the fraction of transmitting nodes in the low and high interference regimes.**

I. Introduction

Characterizing the exact capacity of a large wireless network is notoriously hard due to the complexity of analyzing the impact of the interference from simultaneously transmitting nodes. So far, the analysis has mostly focused on the scaling of the network capacity [5], [6]. To fully understand the formation of the capacity of networks comprising of a plethora of nodes, it is also important to know the magnitude of the capacity, or in other words, the constant in front of the scaling part. When possibly hundreds of nodes transmit over the same channel, a realistic modeling of the impact of wireless interference requires taking into account its additive nature. Hence, we study an SINR-based interference model where the interaction between the links can no longer be described using a simple interference graph.

We formulate the problem as follows. A network in an infinite plane is given with node locations obeying a planar Poisson process. The nodes communicate with each other through the wireless medium. Simultaneous transmissions interfere with each other according to the rules of a given interference model. In the SINR threshold model, if the *signal-to-interference-and-noise-ratio* (SINR) of a link is above a given threshold, a transmission is possible with a fixed rate. We assume that the origin nodes transmit with constant power and that the transmitted signal attenuates according to a power law as a function of distance from its transmitter. A set of links that transmit at the same time is called a transmission mode. The target is to find the transmission mode that maximizes the value of an objective function per area.

As an objective function we use the instantaneous forwarding capacity. With the term forwarding capacity, we refer to the capability of the network to relay information. The capacity is measured in bits per second or bit-meters per second depending on the used link weight. The weight of the link is either one (unweighted) or the length of the projection of the link in a given direction, to represent the progress of information in that direction. We use the word *instantaneous* to emphasize that we are considering only a single transmission mode (individual time slot) whose capacity is a natural, although rather loose, upper bound for continuous multihop flows.

The results are obtained by simulated annealing (SA), a probabilistic method for solving difficult optimization problems. It is based on the work of Metropolis *et al.* in 1953 [15], and was later formulated as a more general optimization technique by Kirkpatrick *et al.* in 1983 [13]. The suitability of SA for this kind of a problem stems from the fact that the method is able to process a wide variety of objective functions and constraints. It is statistically guaranteed to find an optimal solution though, on a negative side, finding the optimum may be time consuming. The main merit of the method is that it permits the transition from graph based interference models to a more realistic modeling of the underlying wireless medium using SINR-based interference models. The reader should keep in mind that though simulated annealing is usually described as a minimization algorithm, our text is written from the maximization point of view.

In this paper, we aim to broaden the knowledge on the theoretical maximum capacity of the SINR threshold model. In the analysis, we establish a so-called neighborhood-size parameter, that captures all the dependencies of the model, and characterize the asymptotic behavior of the system as a function of the parameter. A simple theoretical analysis on the capacity is also conducted to serve as a point of comparison. The simulated annealing algorithm and its parameter selection process are presented, and we illustrate the suitability of different cooling schedules for the problem. The approach is used to determine the weight of the optimal transmission mode as a function of the neighborhood-size parameter for different combinations of the attenuation coefficient of the power law and the threshold required for the SINR. Corresponding

978-1-4673-2039-9/12/$31.00 ©2012 IEEE

values from a model where the data rate of an active link is determined using Shannon's formula (from [20]) are given as a point of comparison. The results are the most accurate ones available for the SINR threshold model thus far. The methodology extends the knowledge of the properties of the theoretical maximum capacity of large-scale wireless networks and supports the known asymptotic scaling characteristics when the number of nodes approaches infinity. It can be used to compare the performance of more practical implementations to the optimal transmission configurations that offer the upper bound of the capacity.

The remainder of this paper is organized as follows. In Section II, we present the related work. Section III describes the network model and the associated objective functions whose behavior is analyzed in Section IV. In Section V, we outline the main principle of the SA method, while Section VI is dedicated to some implementation aspects and to the effects of the choice of simulation parameters. The numerical results are presented in Section VII. We conclude with some closing remarks in Section VIII.

II. Related work

The question of optimal transmission modes appears in wireless networking, e.g., in the context of large scale sensor networks of the future. An extensive network is often modeled as a *massively dense* network [4], [8], [10], [12] that from a single node's perspective appears as an infinite network of randomly placed wireless nodes. The solution of a maximum weight transmission mode 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 information can be "moved" in a given direction [17], [19]. 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. The concept is similar to density of progress, see [1], [21]. Results like these yield useful information about the achievable gains from utilizing optimal global coordination in multihop communications, and thus they complement the well-known scaling results for the capacity of multihop networks, see [5], [6]. The maximum weight independent set problem is also related to the challenging global optimization phase of the original maximum weight scheduling algorithm [22] and its distributed variants [11] and [16].

We consider the maximum weight independent set problem assuming that interference between the links is represented by the SINR threshold model. The SINR-based interference models are of interest as an inapt interference model can lead to unjustified structural decision as there are qualitative differences in the predictions of different interference models [9]. Furthermore, an experimental comparison study [14] shows the accuracy of SINR-based interference in real-life applications. The significance of the interference model selection has been discussed widely in the literature (see [3] for a survey).

In this paper, we extend the work of [20] and further develop stochastic optimization algorithms based on simulated annealing for tackling the maximum weight independent set problem, which essentially corresponds to determining the maximal spatial reuse. Our objective is to determine the total weight of the maximum weight transmission mode per unit area (or node) in an infinitely large wireless network as a function of an appropriately defined neighborhood-size parameter using the much used and in practice relevant SINR threshold model. More specifically, comparing with [20], which focused on a model where the link weights were given by the Shannon capacity of the link, (i) we derive for the threshold model a new analytical model of the capacity, (ii) we analyze the impact of the cooling strategies and show the robustness of the used linear strategy and (iii) we produce numerical results on the forwarding capacity as a function of the various parameters. These numerical results are unknown in the literature, as far as we know, and also enable a comparison between the SINR threshold model and our previous results on the Shannon capacity-based model.

III. Network model

The network consists of nodes, $v \in \mathcal{V}$, distributed randomly over a plane according to a spatial Poisson point process with density n. Each node has one transceiver and can thus participate in only one transmission at any time. A pair of nodes, a transmitter and a receiver, may according to the rules of the interference model form a directed link $l \in \mathcal{L}$. Link l has a weight, w_l, that is either

- one (unweighted) or
- the x-progress of the link (length of the projection of the link onto the x-axis)

Note that x-axis represents an arbitrary direction, the direction in which the information is moved. Links that can successfully transmit at the same time form a transmission mode $m \in \mathcal{M}$.

A. SINR threshold model

We assume that all the nodes transmit with the same constant power P_0. The attenuation factor between the transmitter of link l and the receiver of link l' is assumed to follow power law

$$g(l, l') = \left(\frac{|\boldsymbol{t}_l - \boldsymbol{r}_{l'}|}{\rho_0}\right)^{-\alpha}, \quad (1)$$

where ρ_0 is a reference distance, and α is a given attenuation coefficient.

Under the SINR interference model, there exists a link $l = (u, v) \in \mathcal{L},\ u, v \in \mathcal{V}$ if the signal to interference and noise ratio (SINR) is greater than or equal to a given threshold, θ. The SINR at $\boldsymbol{r}_l$ for link l in transmission mode m is

$$\text{SINR}(\boldsymbol{r}_l, m) = \frac{P_0 \cdot g(l, l)}{\sum_{l' \in m \setminus \{l\}} P_0 \cdot g(l', l) + \sigma^2}, \quad (2)$$

where σ^2 is the thermal noise power. The spectral efficiency of links with $\text{SINR}(\boldsymbol{r}_l, m) \geq \theta$ is assumed constant irrespective

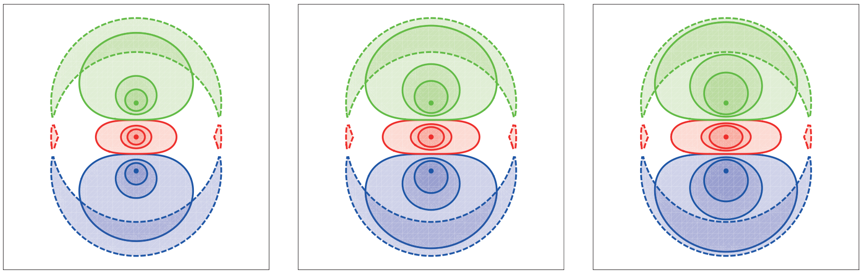

Fig. 1: The areas where a reception is possible for three transmitters (dots) under the Boolean interference model (dashed lines) and the SINR threshold model with threshold 1 (outermost solid lines). The inner solid lines are contours for SINR values 7 and 31.

of the SINR,

$$R_l(m) = R(\theta) = \log_2(1+\theta), \tag{3}$$

i.e., the spectral efficiency at the threshold according to Shannon's formula.

As the attenuation coefficient α and the threshold θ (and possibly also the thermal noise power, σ^2) are given constants, we can define a length unit

$$\rho(P_0) = \rho_0 \sqrt[\alpha]{P_0/\theta\sigma^2}, \tag{4}$$

that is, the maximum distance at which a reception is possible if there are no competing transmissions (zero interference).

The difference between the SINR threshold model and the Boolean interference model in an example scenario with three active transmitters is illustrated in Figure 1. The areas where a reception is possible from each of the transmitters are drawn in the figure. Under the simple Boolean interference model, a node is able to receive a transmission if it is inside the (fixed) transmission radius of only one active node. The main difference between the SINR interference model and the Boolean interference model is that, although the border of the reception area approaches the one of the Boolean interference model when α grows, the SINR interference model is more realistic in always allowing a reception near the transmitter.

B. Objectives

Our goal is to find the instantaneous forwarding capacity, i.e, the maximum capacity of a transmission mode per unit area, $\hat{I}$, as a function of the system parameters. The capacity of a transmission mode is measured either in bits per second (link weight one) or bit-meters per second (weighted by x-progress). Hence, the unit of $\hat{I}$ is either $1/s/m^2$ or $1/s/m$ respectively. The link weights one (unweighted) and the x-progress of the link (length of the projection of the link onto the x-axis) are denoted by subindices 1 and x in the context of $\hat{I}$ (and its dimensionless counterpart u introduced next), while subindex $*$ refers to either of the weights.

According to the Buckingham π theorem [2], any physically meaningful equation is equivalent to another equation involving all[1] the independent dimensionless parameters that can be constructed from the original variables. In our case, in addition to $\hat{I}$, only one independent parameter can be formed, and hence $\hat{I}$ can be expressed in terms of it. The most natural choice for this dimensionless parameter is

$$\nu(n,\rho) = \pi n \rho^2, \tag{5}$$

which has the interpretation as the mean neighborhood size. All unknown functions of the system parameters can be reduced to functions of this single variable.

Depending on the definition of the link weight, $\hat{I}^{\alpha,\theta}(\nu)$ can be expressed as

$$\hat{I}_1^{\alpha,\theta}(\nu) = C_0\, n\, u_1^{\alpha,\theta}(\nu(n,\rho)), \text{ or} \tag{6}$$

$$\hat{I}_x^{\alpha,\theta}(\nu) = C_0 \sqrt{n}\, u_x^{\alpha,\theta}(\nu(n,\rho)), \tag{7}$$

where C_0 denotes the bandwidth [Hz], and $u_1^{\alpha,\theta}(\nu)$ and $u_x^{\alpha,\theta}(\nu)$ are dimensionless functions of the independent dimensionless parameter ν. Note that functions $u_*^{\alpha,\theta}(\nu)$ are different for different values of α and θ.

IV. Analysis

In order to gain a better understanding of how the dimensionless functions $u_*^{\alpha,\theta}(\nu)$ behave, it is useful to consider them analytically. The asymptotic analysis studies the performance in small densities, $\nu \ll 1$, and in high densities, $\nu \gg 1$. See [18], [20] for other interference models. With a simple analytical model, we study if or when it is possible to reproduce the results of the well-chosen transmission modes with a naive model.

A. Asymptotic characterization

Let us rewrite (2) as follows,

$$\text{SINR}(\boldsymbol{r}_l, m) = \frac{\bar{g}(l,l)\,\nu^{\alpha/2}}{\sum_{l' \in m\setminus\{l\}} \bar{g}(l',l)\,\nu^{\alpha/2} + 1}, \tag{8}$$

[1] If there are a physical variables that are expressed in terms of b independent physical units, the number of dimensionless parameters is $a - b$.

where $\bar{g}$ is the dimensionless function

$$\bar{g}(l,l') = \left(\sqrt{\pi n}\, |\boldsymbol{t}_l - \boldsymbol{r}_{l'}|\right)^{-\alpha}. \quad (9)$$

An important observation is that, because of insertion of the factor $\sqrt{\pi n}$, for any realization of the spatial Poisson process the function $\bar{g}(l,l)$ is independent of the scale. That is, if all the distances are stretched or contracted by some factor, the value of $\bar{g}(l,l')$ remains unchanged for any pair of links $\{l,\, l'\}$. Thus, the dependence on the density is fully incorporated in the factors $\nu^{\alpha/2}$ in (8). Since this factor controls the ratio of the two terms in the denominator, the low and high density limits may equivalently be called the *noise-limited* and *interference-limited* cases, respectively.

Curves $u_*^{\alpha,\theta}(\nu)$ defined in (6) and (7) are increasing functions of ν. This stems from the fact that $\mathrm{SINR}(\boldsymbol{r}_l, m)$ of (8) is an increasing function of ν. The total capacity of any transmission mode m is constant, but the maximizing mode, and thus the maximum capacity, may change as ν increases and makes more transmission modes feasible.

Let us now consider the interference-limited case $\nu \gg 1$. In this case, the one in the denominator of (8) may be neglected, whence the factors $\nu^{\alpha/2}$ cancel out. Therefore, provided that ν is large enough, the problem becomes completely scale-free, independent of ν. No matter how the scale is stretched or contracted, it is always the same mode of active links that realizes the optimum.

Next we turn our attention to the noise-limited case $\nu \ll 1$. Now, the interference term in the denominator of (8) may be neglected, whence the $\mathrm{SINR}(\boldsymbol{r}_l, m)$ reduces to the numerator of the expression on the right hand side. Without interference, reception is possible anywhere in the transmission region, and a higher SINR threshold θ directly leads to better spectral efficiency (3) as the system is studied as a function of ν, given by (5).

Under a reasonable assumption that $\theta > 1$, a reception is only possible from the closest transmitting node. Now, strong attenuation (large α) increases the capacity independent of ν. As α tends to infinity, the interference is dominated by the interfering transmitter closest to the receiver of the link (or the noise if there are no interfering transmitters within the distance of ρ) that is still farther than the transmitter of the link. Hence, the SINR (8) tends to infinity as $\alpha \to \infty$. The effect is stronger when ν is large, and there are more competing transmissions.

B. *Simple model for the unweighted case*

Here we consider the unweighted case and derive a rough approximation for the dimensionless capacity $u_1^{\alpha,\theta}(\nu)$. Without loss of generality, we consider a Poisson process with intensity $n = 1$, $\rho_0 = 1$ and $\sigma^2 = 1$, and study the dimensionless capacity as a function of the power P_0, which corresponds to a specific value of ν. Two modeling assumptions are made. First, we assume that every node is able to form a link with its nearest neighbor without conflicts, and that links are indeed formed starting from the nearest pair of nodes up to a node pair at the maximum distance $\hat{r}$ from each other. Then, the fraction $F(\hat{r}) = 1 - e^{-\pi\hat{r}^2}$ of nodes is used to form links, and

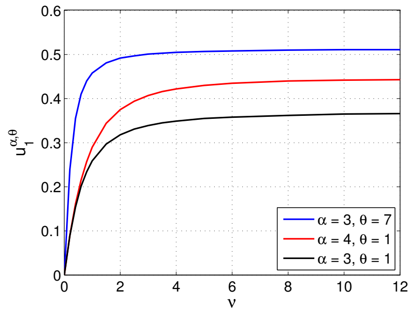

Fig. 2: The dimensionless capacity in the unweighted case as a function of ν for (α,θ)-pairs (3,1), (4,1) and (3,7).

the density of transmitters is $F(\hat{r})/2$. Second, the interference power to a receiver with its nearest neighbor at distance r, coming completely from transmitters farther than r from the receiver, is represented by its expectation,

$$I(r) = \frac{F(\hat{r})}{2}\int_r^\infty P_0 s^{-\alpha} 2\pi s\, ds = \frac{P_0 \pi r^{2-\alpha}}{\alpha - 2}(1 - e^{-\pi\hat{r}^2}).$$

It is easy to reason that it is advantageous to increase $\hat{r}$ up to the point where the receiver of a link of this maximum length experiences a SINR equaling the threshold value θ. Thus, $\hat{r}$ is determined from the equation

$$P_0\, \hat{r}^{-\alpha} = \theta\,(I(\hat{r}) + 1),$$

and the dimensionless capacity is $(F(\hat{r})/2)\log_2(1+\theta)$. In Figure 2, we illustrate $u_1^{\alpha,\theta}(\nu)$ as a function of ν for (α,θ)-pairs (3,1), (4,1) and (3,7). As will be seen later, in spite of the heuristic nature of the model, these estimates are surprisingly accurate for the cases with $\theta = 1$. However, as soon as θ becomes larger the model underestimates the capacity, as the example with $\theta = 7$ shows. In order to obtain more precise results, we next develop a technique based on simulated annealing that correctly captures the spatial dependencies of the active links.

V. Simulated Annealing

In order to determine the instantaneous forwarding capacity of a large network (that represents the entire plane), we want to find the transmission mode that maximizes the capacity per unit area. This is done using simulated annealing (SA).

The idea of the method comes from the physical process of annealing, where a material cooled slowly enough approaches the ground state of the system, i.e., the state with minimum energy (maximum of the negative energy). As the name indicates, SA tries to simulate this kind of process. In this method, the current solution is randomly moved to a “neighboring” solution with a probability that depends on the height of the ascent/descent and a parameter called *temperature*. By allowing the algorithm to move to a worse solution, it is possible to avoid being stuck at local optima. When the temperature parameter is properly modified during the optimization, the algorithm also eventually reaches the optimal solution.

The SA algorithm uses the following elements in its operation (physical analogues in parentheses):

1) A finite set of possible states.
2) A real-valued target function (energy) that defines the set of optimal states.
3) A set of neighboring states for each state and the rule for randomly choosing the next state from the neighbors of the current state.
4) A cooling schedule (temperature) that "freezes" the probability distribution of the states to the set of optimal states over time.

Also required are the initial state and the termination condition. In our case, the set of transmission modes forms the state space of the system. The order of the states is determined by the weight of the transmission mode $w(m)$, and neighboring states are those transmission modes that differ only by one link.

Let us consider a Markov chain with the state space composed of the different transmission modes and with the steady-state distribution

$$\pi(m) \sim \exp\{w(m)/T\}, \tag{10}$$

where T is a constant. When T is small, the probability mass is concentrated to the maximum we are interested in,

$$m^* = \arg\max_{m \in \mathcal{M}} w(m). \tag{11}$$

The time needed for the Markov chain to reach the steady-state with small T can, depending on the heights of the local maxima, be inordinate. The idea of simulated annealing is to avoid this problem by slowly decreasing the temperature T. Even then, the cooling schedule has to be slow enough for the system to avoid being quenched in a local extremum [7].

A. Algorithm

A Markov chain M_τ with state space $\mathcal{M}$ is formed, and the steady-state probability of a transmission mode m is chosen to be

$$\pi(m) = \frac{e^{w(m)/T}}{\sum_{m' \in \mathcal{M}} e^{w(m')/T}}, \tag{12}$$

where T, i.e, the temperature, is a positive parameter and $w(m)$ is the weight of transmission mode m equal to

$$w(m) = \sum_{l \in m} w_l R_l(m), \tag{13}$$

where w_l is the weight of link l, and the spectral efficiency, $R_l(m)$, under the Boolean interference model is one. As mentioned, we study two different link weights

$$w_l = \begin{cases} 1, & \text{unweighted} \\ |[\boldsymbol{t}_l]_1 - [\boldsymbol{r}_l]_1|/\rho_0, & x\text{-progress.} \end{cases} \tag{14}$$

When T is small, the probability of the transmission mode with the highest capacity is close to one.

To assure the required steady-state distribution, the transition probabilities, $p(m', m)$, are chosen so that π shows detailed balance, i.e., $p(m', m)\pi(m') = p(m, m')\pi(m)$. This is achieved using a proposal distribution $q(m', m)$ along with an acceptance/rejection procedure. The proposal distribution $q(m', m)$ gives the probability that transmission mode m is the candidate to be selected as the next transmission mode when the current transmission mode is m'. When a proposal m with a higher capacity is accepted with probability 1, a transition to a lower capacity, $w(m) < w(m')$, is accepted with probability, r, that can be solved from the detailed balance equation,

$$\begin{aligned} p(m', m)\pi(m') &= p(m, m')\pi(m) \\ \Leftrightarrow r \cdot q(m', m)\pi(m') &= 1 \cdot q(m, m')\pi(m) \\ \Leftrightarrow r &= \frac{q(m, m')\pi(m)}{q(m', m)\pi(m')}, \end{aligned}$$

and is rejected otherwise. Note that r is always defined since a transition from m' to m can only occur if both $q(m', m)$ and $\pi(m')$ are nonzero.

When the proposal m is obtained by randomly choosing a link, $l \in \mathcal{L}$, and adding it to m' if l is feasible and does not belong to m' and removing it from m' if it does, we have $q(m', m) = q(m, m') = 1/|\mathcal{L}|$, and

$$r = e^{-(w(m') - w(m))/T}. \tag{15}$$

This equation is referred to as Metropolis (acceptance) criterion.

Markov chain M_τ with known steady-state distribution is simulated to find transmission modes with near optimal capacity. As the temperature T is decreased, the samples come from a process that more and more heavily favors modes with a large capacity.

VI. Implementation Aspects

In this section, we discuss implementation aspects of the simulated annealing algorithm and some considerations that need to be taken into account in setting up the simulations.

A. General parameters

A key parameter of the simulated annealing algorithm is the temperature of the system. When the temperature parameter is properly modified during the optimization, the algorithm also eventually reaches the optimal solution. One way to properly modify the temperature is to use logarithmic cooling, $T(t) = c/\log(1+t)$, where t is the time and c is a constant that needs to be large enough [7].

In order to avoid being stuck at local optima, the initial temperature, that depends on c, needs to be large enough. On the other hand, in order to find the true optimum, the simulation time has to be long. That is, the temperature at the end of the simulation needs to be very small. Using the logarithmic cooling schedule this might take a long time (depending on c).

A linear cooling schedule, $T(t) = c/t$, is faster than the logarithmic cooling schedule and is not guaranteed to "freeze" the process to the optimal state. However, by using a linear cooling schedule, it is easier to choose an initial temperature that is large enough for the process to explore the search space and still have a low end temperature that "freezes" the process to a good pseudo-optimal solution. This is a desired property especially when the simulation time is limited, and the discovery of the optimum cannot be guaranteed (even using logarithmic cooling).

Figure 3 illustrates the performance of the two cooling schedules, logarithmic and linear, as a function of the final temperature when the simulation time is fixed to one million steps. From the figure one can see that, though both schedules give nearly the same performance with correctly chosen c, that fixes the end temperature, the linear schedule is much less sensitive to the selection of the parameter. The larger error bars of the logarithmic schedule indicate that some simulation runs have been unable to escape local maxima and have thus resulted in much worse performance.

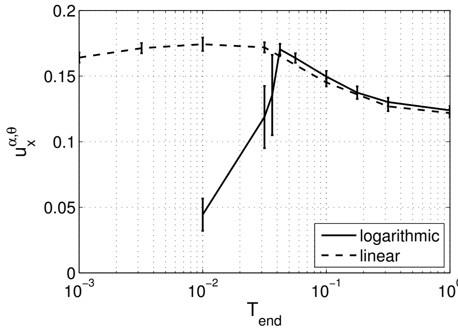

Fig. 3: The curves $u_x^{\alpha,\theta}(T_{\text{end}})$ for the logarithmic and linear cooling schedules in a one million step simulation with parameters $\alpha = 3$, $\theta = 7$, $\nu = 6$.

It is to be noted that when the simulation time is limited, the simulated annealing algorithm cannot be guaranteed to find the optimal solution. Since both the initial and the end temperature are significant to the end result, the parameter selection is easier with a linear cooling schedule. Hence, we use a fixed simulation length and linear cooling, and the temperatures of the system form a harmonic sequence.

For the results to be generalizable to a plane, the simulated network needs to be large enough. To eliminate border effects, we identify the top and the bottom of the square and the circular edges of the formed cylinder to form a torus. Additionally, we allow an interfering signal to travel around the torus for a given number of rounds as explained in the next subsection. Figure 4 depicts the capacity $u_x^{\alpha,\theta}$ as a function of the expected network size for $\nu = 6$. It shows that the difference in the capacity becomes negligible for network sizes greater than 600.

B. Handling residual interference

To estimate the amount of interference that is not covered by considering a network with a finite number of nodes, we do the

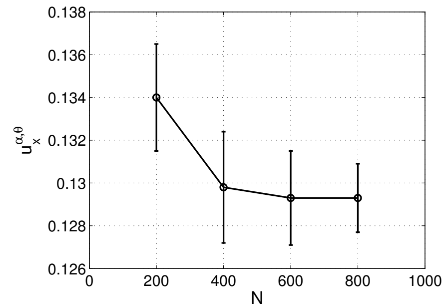

Fig. 4: The effect of the expected network size (number of nodes) on the network capacity in the simulations.

following calculation. Let $B = [0, \sqrt{A}] \times [0, \sqrt{A}]$ be a square of a plane with nodes placed according to a Poisson point process of density n. Now, the fraction of interference coming from outside a circle that can be fitted into B compared to the interference coming from outside a circle with a radius that equals the mean distance of a node to its nearest neighbor, $1/2\sqrt{n}$, is

$$\frac{\displaystyle\int_{\sqrt{A}/2}^{\infty} z\, n\, r^{-\alpha}\, 2\pi r\, dr}{\displaystyle\int_{1/2\sqrt{n}}^{\infty} z\, n\, r^{-\alpha}\, 2\pi r\, dr} = \frac{1}{\sqrt{nA}} = \frac{1}{\sqrt{N}}, \qquad (16)$$

where N is the total number of nodes in the simulation area and z is the fraction of transmitting nodes.

For example, if $N = 600$, less than 4 % of the interference is caused by nodes outside the square. We take such a residual interference approximately into account by letting an interfering signal to travel around the torus for a given number of rounds. Essentially, this is the same as adding copies of B around it in a network that has not been wrapped up to a torus (our implementation has 201×201 squares). This way, we are able to accurately take into account most of the interference coming from outside B.

VII. Numerical results

This section presents the results for different combinations of α and θ that have been obtained using simulated annealing. The SA algorithm produces as a result the maximum weight of a transmission mode $w(m)$ (13) for a finite network realization. When the reference distance is chosen suitably, $\rho_0 = 1/\sqrt{n}$, the dimensionless functions can be calculated simply as $u^*(\nu) = w(m)/N$, where N is the expected number of nodes in the network.

First, we consider the unweighted case. Figure 5 shows the curve $u_1^{\alpha,\theta}(\nu)$ for different combinations of α and θ. The results are averages over 10 network realizations with the mean network size of 600 nodes. The 95 % confidence intervals are shown as error bars. The end value of the temperature parameter T was chosen to be 0.1 after the 10-million-step simulation. As predicted in Section IV-A, the curves are increasing functions of ν. The figure also matches with the deductions for the dependence on α and θ.

When parameter α grows, the interference attenuates faster, and more links can be activated. The threshold parameter has an opposite effect, but higher θ also increases the link capacity. When the threshold grows, a higher SINR value is required for a successful transmission, and a smaller fraction of the links can be activated. However, the spectral efficiency (3) of the links is higher. When the model of Section IV-B is set against these values, the results for $\theta = 1$ are similar, but the difference when θ is greater is evident, see also Figure 2.

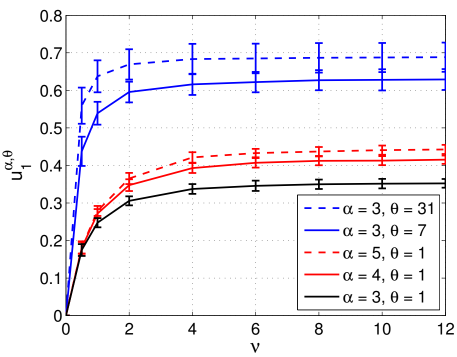

Fig. 5: The results for the unweighted SINR threshold model.

Next, we consider the case weighted by the x-progress. Figure 6 shows the curve $u_x^{\alpha,\theta}(\nu)$ for different combinations of α and θ. The results are averages over 5 network realizations with the mean network size of 800 nodes. The 95 % confidence intervals are shown as error bars. The end value of the temperature parameter T was chosen to be 0.01 after the 20-million-step simulation. These curves also match the predictions of Section IV-A.

Threshold θ has a notable significance when ν is small while α has almost no importance. When ν is small, interference has very little role in constituting the capacity as the problem revolves around being able to form the links. Thus, the attenuation coefficient α is less important than θ that directly affects the spectral efficiency of the links that can be formed.

When the mean number of neighbors, ν, grows, the effect of α increases as the interfering signals and their attenuation becomes more important. In this case, higher α naturally leads to higher capacities. The effect of θ with larger ν is not straight forward, but the SINR threshold value that maximizes the capacity depends on the neighborhood size. Figure 7 depicts the curve $u_x^{\alpha,\theta}(\theta)$ for different values of ν when α is equal to three. From the figure, it can be seen that when ν is small a higher threshold leads to better performance, but the optimal θ becomes smaller as ν grows.

The percentage of transmitting nodes in the optimal transmission mode as a function of ν is illustrated Figure 8. It is comparable with Figure 5 after the effect of spectral efficiency has been eliminated (divide the values of Fig. 5 by $\log_2(1+\theta)$) and shows the same behavior with respect to α and θ. Naturally, the tendency towards using longer links means that a smaller fraction of nodes can be activated. The curves start leveling off when ν grows indicating that the asymptotically optimal transmission mode starts to be established already with relatively small values of ν.

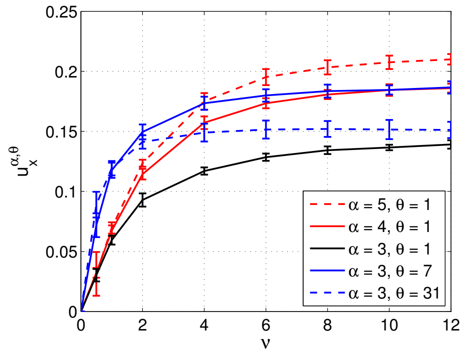

Fig. 6: The results for the SINR threshold model weighted by the x-progress.

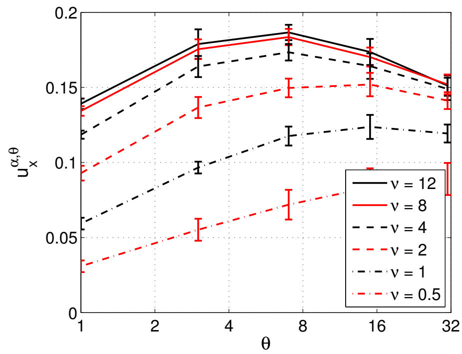

Fig. 7: The curve $u_x^{\alpha,\theta}(\theta)$ for the SINR threshold model weighted by the x-progress when $\alpha = 3$ and ν varies.

The previous results are upper bounded by results from a model where the spectral efficiency of a link is calculated according to Shannon's formula. That is, instead of (3), the spectral efficiency at the threshold, the actual SINR is used to calculate the spectral efficiency $\log_2(1 + \mathrm{SINR}(\boldsymbol{r}_l, m))$. The Shannon model represents what is achievable with perfect adaptive coding and modulation and corresponds to a (graded) model with infinitely many thresholds. Table I represents comparable values for $\nu = 10$ with 95 % confidence intervals.

VIII. Conclusions

Finding and using efficient combinations of transmitting links is crucial to the performance of a wireless network. In

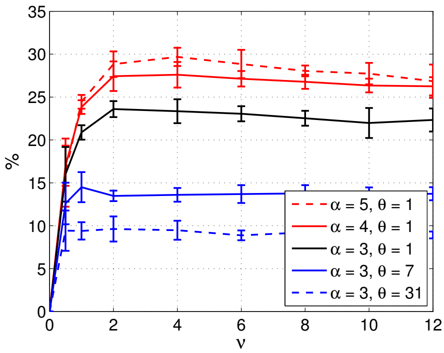

Fig. 8: The fraction of transmitting nodes under the SINR threshold model weighted by the x-progress.

TABLE I: Results for the Shannon model with $\nu = 10$

α	u_1^α	u_x^α
3	1.15 $\pm$ 0.02	0.252 $\pm$ 0.002
4	1.83 $\pm$ 0.03	0.431 $\pm$ 0.003
5	2.46 $\pm$ 0.04	0.598 $\pm$ 0.005

this paper, we studied producing (near) optimal transmission modes with stochastic optimization technique called simulated annealing. The approach is suitable for determining the behavior of the theoretical maximum capacity of a large-scale wireless network. The obtained results illustrate the formation of the overall capacity of a wireless network and support the known asymptotic laws that tell how the capacity scales as the network size grows to infinity by providing an accurate estimate of the capacity. The results also serve as a point of comparison for practical medium access control protocols.

The simulated annealing algorithm and its parameter selection process were presented, and a linear cooling schedule was selected for its robustness compared with a logarithmic one in finite simulation time. We studied extensively how the maximum capacity behaves for different network densities, attenuation coefficients, and thresholds. The numerical results show that the capacity exhibits the deduced asymptotic behavior with respect to the model parameters; the threshold, the attenuation coefficient, and the mean neighborhood size. The results were compared with the results from a model where the data rate of an active link is determined using Shannon's formula.

The aforementioned Shannon model offers an upper bound for the threshold model with the same attenuation coefficient, and illustrates the gain that can be achieved with perfect adaptive coding and modulation. Hence, it also gives an upper bound for graded models that allow the use of multiple thresholds to better utilize the potential of the links. The optimal selection of multiple thresholds remains as a field of future work, as well as the study of other interference models that fall between the two extremes of the simple Boolean interference model and the complex SINR-based additive interference.

Acknowledgements

The work was done in the HEWINETS project funded by TEKES, Ericsson, and Cassidian Systems.

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] E. Buckingham. On physically similar systems; illustrations of the use of dimensional equations. *Phys. Rev.*, 4(4):345–376, 1914.

[3] P. Cardieri. Modeling interference in wireless ad hoc networks. *IEEE Communications Surveys & Tutorials*, 12(4):551–572, 2010.

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

[5] 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.

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

[7] B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 1988.

[8] 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.

[9] A. Iyer, C. Rosenberg, and A. Karnik. What is the right model for wireless channel interference? *IEEE Transactions on Wireless Communications*, 8(5):2662–2671, 2009.

[10] P. Jacquet. Geometry of information propagation in massively dense ad hoc networks. In *ACM MobiHoc 04*, pages 157–162, 2004.

[11] K. Jung and D. Shah. Low delay scheduling in wireless network. In *IEEE ISIT*, pages 1396–1400, 2007.

[12] M. Kalantari and M. Shayman. Energy efficient routing in wireless sensor networks. In *Conference on Information Sciences and Systems*, 2004.

[13] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, 1983.

[14] R. Maheshwari, S. Jain, and S.R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *SenSys '08*, pages 141–154, 2008.

[15] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 1953.

[16] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless networks via gossiping. *SIGMETRICS Perform. Eval. Rev.*, 34(1):27–38, 2006.

[17] J. Nousiainen and P. Lassila. Approximating maximum directed flow in a large wireless network. In *IEEE ICC*, Dresden, Germany, 2009.

[18] J. Nousiainen, J. Virtamo, and P. Lassila. Maximum weight independent sets in an infinite plane with uni- and bidirectional interference models. *Annals of Telecommunications*, 66(1–2):119–132, 2010.

[19] J. Nousiainen, J. Virtamo, and P. Lassila. On the achievable forwarding capacity of an infinite wireless network. In *ACM MSWiM*, Bodrum, Turkey, 2010.

[20] J. Nousiainen, J. Virtamo, and P. Lassila. Optimal transmission modes by simulated annealing. In *Proceedings of PM^{2}HW^{2}N'11*, Miami, Florida, USA, Oct 2011.

[21] H. Takagi and L. Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *IEEE Transactions on Communications*, 32(3):246–257, 1984.

[22] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1949, 1992.