

Multidirectional Forwarding Capacity in a Massively Dense Wireless Network

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

Abstract—The problem of the capacity of a massively dense wireless multihop network can be broken down into separate problems at macroscopic and microscopic levels. At the microscopic level, from the local perspective, the network appears like an infinite plane with traffic that is uniform but flowing in many directions. Previous studies have assumed that it is sufficient to find the maximum sustainable density of packet flow in a single direction and use time sharing to serve flows in different directions. We show that this time-sharing limit can be exceeded by scheduling that truly interleaves the traffic flows in different directions. Determining the forwarding capacity for multidirectional traffic defines a new problem that has not been studied earlier. For finding numerical values, we adopt a constructive approach by simulating a finite but large network using greedy maximum weight scheduling. Bi- and four-directional balanced traffic patterns are studied. For the latter, an improved greedy algorithm is developed, using insights from our earlier work. Isotropic traffic plausibly yields the highest benefits for multidirectional forwarding, and our results show that a significant gain compared with single-directional forwarding can be achieved.

I. INTRODUCTION

Efficient utilization of radio resources in wireless multihop networks, to achieve maximum capacity, is of great importance especially when the network in question is large. The seminal papers [5], [6] give important insight in the asymptotic scaling of this maximum capacity. However, these studies leave the question of the magnitude of the capacity unresolved. We aim at providing numerical results that shed light on this fundamental capacity problem.

In the limit of a massively dense network, a natural separation of spatial scales emerges, and the network capacity problem can be separated into two independent problems at the macroscopic and microscopic levels, respectively, see [7], [8]. At the macroscopic level, the underlying network is considered a continuous medium, where the routes are smooth geometric curves [2], [9], [12], [18], and the problem is to define the most efficient routing given the traffic matrix and the constraints set by the microscopic level. At the microscopic level, representing the network from a single nodes point of view, a massively dense network appears to be infinite, and the task is to forward traffic as efficiently as possible.

Here we focus on the microscopic level problem. In general, at the microscopic level there is traffic flowing in many directions, as characterized by the directional distribution of the traffic. A key assumption made in the above macroscopic

level studies is that the constraint set by the microscopic level reduces the microscopic level problem to one of determining the maximum density of packet flow in a given direction and this single-directional forwarding capacity is shared in time between the different directions. In our previous papers [15], [16], we have provided increasingly accurate estimates on the single-directional forwarding capacity.

However, an important observation is that the earlier form of the constraint only represents a sufficient condition for the capacity problem in dense networks, i.e., time sharing between directions is not optimal. Significantly higher capacities can be achieved by exploiting the fact that traffic flowing in different directions can be interleaved in the microscopic level scheduling. The necessary condition characterizing the microscopic level problem states that the sustainable local traffic load is bounded by a limit, which is a functional of the directional distribution of the traffic, and which we call the multidirectional forwarding capacity. To our knowledge the multidirectional forwarding capacity problem has not been studied earlier. Our objective is to evaluate the gain from the multidirectional forwarding compared with only single-directional forwarding in certain special cases, namely balanced bidirectional and balanced four-directional cases. Throughout this paper we use the word balanced to describe a traffic pattern where the traffic streams are equal in the different directions.

At the microscopic level, the network consists of nodes distributed according to a spatial Poisson process in an infinite plane, and the traffic to be carried is uniform but non-isotropic. The interference caused by simultaneous transmissions is modeled using the Boolean interference model. To obtain an estimate of the multidirectional forwarding capacity, we simulate a large but finite network in a square area carrying only relay traffic. In the simulations, we apply maximum weight scheduling where the decision about the resource allocation in each time slot is based on the current queue-length-based weight of each link. In this case, the different directions of the traffic correspond to classes. It is well known that maximum weight scheduling is throughput optimal [20]. However, considering the size of the network realizations that we use, finding the true maximum weight independent set in each time slot is infeasible. Therefore, we need to resort to a greedy method.

While being suboptimal, greedy scheduling [3], [14], [21] represents a more practical approach from the network simu-

lation point of view for determining the maximal forwarding capacity than other throughput optimal approaches, such as randomized maximum weight scheduling [4], [19] or distributed CSMA approaches [11]. Also, approximate linear programming formulations using only a subset of the necessary schedulability constraints [10], [13] are computationally inefficient for the network sizes we consider, or the bounds become loose, see, e.g., [15].

In the examples of the directional distribution that we consider (balanced bi- and four-directional), the naive basic greedy method works well for the bidirectional case. However, it fails when cross traffic is mixed with the bidirectional stream as happens in the four-directional case. Based on our insights at high network densities, we develop an improved greedy algorithm that solves the problem. It represents a parameterized variant of the basic greedy; in one extreme it works as the basic greedy and in the other extreme the four directions are treated as two bidirectional ones.

Our results demonstrate the significant gains from the multidirectional forwarding. We argue that the isotropic traffic case provides the highest gains in comparison with traffic in a single direction and that already the results with four directions provide a rough order of magnitude estimate of the gain.

The remainder of this paper is structured as follows. In Section 2, we formulate the problem along with the used notation and the network model. Section 3 presents the greedy maximum weight scheduling algorithms, while section 4 discusses the implementation issues. The numerical results are given in Section 5. Finally, Section 6 concludes the study.

II. PROBLEM FORMULATION

In this section, we present the network model and the microscopic level multidirectional forwarding problem with the related definitions.

A. 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 . We assume that time is slotted and that the length of a time slot matches the duration of a packet transmission with the nominal link capacity C . The network is further modeled as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{L})$, where there exists a link $(u, v) \in \mathcal{L}$, $u, v \in \mathcal{V}$ if the distance between the nodes, $|\mathbf{u} - \mathbf{v}|$, is less than the transmission range ρ . Two links l and l' interfere with each other if

$$|\mathbf{t}_l - \mathbf{r}_{l'}| \leq \rho \vee |\mathbf{r}_l - \mathbf{t}_{l'}| \leq \rho,$$

where \mathbf{t}_l is the location of the transmitting node, \mathbf{t}_l , of link $l \in \mathcal{L}$, \mathbf{r}_l is the location of the receiving node, \mathbf{r}_l . This interference model, referred to as the Boolean interference model, says that a node is only capable of receiving a transmission if it is inside the transmission radius of only one active node.

Because not all links can be active simultaneously, we have to establish a schedule α which tells us how the links are used. All the links that are active simultaneously have to belong to the same independent set of links to avoid collisions.

The independent sets that are used for transmitting are also referred to as transmission modes, and we denote the set of transmission modes with $\mathcal{M} = \{m_1, \dots, m_M\}$. The schedule $\alpha = \{\alpha_1, \dots, \alpha_M\}$ assigns each transmission mode m_i with the proportion of time α_i that it is used. Now the effective capacity of link l is

$$c(l) = C \sum_{i=1}^M \alpha_i \mathbf{1}_{\{l \in m_i\}},$$

that is, the nominal capacity times the time share the link is active.

B. Multidirectional forwarding capacity problem

A massively dense network corresponds to a dense network in a closed domain \mathcal{A} with the nodes having an infinitesimal transmission range, and the paths being smooth geometric curves allowing a continuous representation of the network at the global scale. From the local perspective, a single node sees the network as an infinite network, and locally the carried traffic consists essentially only of relay traffic. According to [8], the network capacity problem separates into two scales:

- 1) Macroscopic level routing tries to find routes enabling to carry as much traffic as possible through the network without exceeding the microscopic level capacity constraint.
- 2) Microscopic level forwarding aims at coordinating the transmissions so that the packets are relayed hop-by-hop as efficiently as possible.

In general, the problem at the microscopic level relates to the following macroscopic level quantities, the directional distribution of the traffic $f(\theta)$ and the local traffic load Φ .¹ The directional distribution of traffic $f(\theta)$ represents the fraction of traffic in a given direction θ . The local traffic load Φ gives the total offered traffic intensity in [pkts/s/m] summed over all the angles. Previous work on the capacity problem, based on the above separation of microscopic and macroscopic scale problems [7], [8], [18], has assumed that the microscopic scale sets the limit $\Phi \leq I_1^*$ for the local traffic load, where I_1^* is the (single-directional) forwarding capacity defined as the maximum sustainable density of directed packet flow [pkts/m/s]. In fact, this constraint expresses a sufficient condition for the sustainable traffic. Namely, if the traffic flow of intensity I_1^* can be sustained in a single direction, then a traffic load Φ with arbitrary directional distribution $f(\theta)$ satisfying $\Phi \leq I_1^*$ can be handled by a simple time-sharing, by allocating the traffic in the direction increment $(\theta, \theta + d\theta)$ the time share $f(\theta)d\theta$.

While the stated constraint is sufficient, it is not a necessary condition. Namely, when the traffic consists of a mixture

¹Angular flux of packets [8] in direction θ , denoted by $\varphi(\theta)$ is equal to the rate [pkts/s/m/rad] at which packets flow in the angle interval $(\theta, \theta + d\theta)$ across a small line segment of the length ds perpendicular to direction θ divided by $ds \cdot d\theta$ in the limit when $ds \rightarrow 0$ and $d\theta \rightarrow 0$. We write $\varphi(\theta) = \Phi \cdot f(\theta)$, where Φ is the scalar flux $\Phi = \int_0^{2\pi} \varphi(\theta) d\theta$ and $f(\theta)$ is the directional distribution $\int_0^{2\pi} f(\theta) d\theta = 1$.

of flows in different directions, it is, in general, possible to carry more traffic by properly interleaving the use of links for flows in different directions in the same time slot. The actual limit, the multidirectional forwarding capacity, depends on the directional distribution, $f(\theta)$, of the traffic, and we denote it by $I^*[f(\theta)]$, using square brackets to emphasize the functional dependence on $f(\theta)$. The multidirectional forwarding capacity can be interpreted to represent the maximum sustainable mean density of progress, i.e., the density of packets times their mean velocity in their respective directions [pkts/m/s].

In earlier papers [15], [16], we have found bounds or approximations for the constant I_1^* (previously denoted by I^*). But finding the functional $I^*[f(\theta)]$, to our knowledge, has not been addressed before. Virtually nothing is known about $I^*[f(\theta)]$, but on general grounds one can state that for any $f(\theta)$

$$I_1^* \leq I^*[f(\theta)] \leq I_\infty^*,$$

where I_∞^* is the limit for the case where the directional distribution is uniform, $f(\theta) = 1/2\pi$, and I_1^* is the above limit with traffic in a single direction with $f(\theta) = \delta(\theta)$, i.e., the Dirac delta function. The first inequality is the above sufficiency condition, and the second is rather obvious², i.e., the interleaving advantage is greatest for isotropic traffic. It should be noted that all the quantities depend also on the network parameters, suppressed here for clarity (see Sec. II-D).

In the simulations, we apply a greedy implementation of the maximum weight scheduling algorithm to obtain estimates of the multidirectional forwarding capacity for certain special cases of the directional distribution. The covered cases include two opposite directions with equal flows, $f(\theta) = (\delta(\theta) + \delta(\theta - \pi))/2$, and the four cardinal directions with balanced flows, $f(\theta) = (\delta(\theta) + \delta(\theta - \pi/2) + \delta(\theta - \pi) + \delta(\theta - 3\pi/2))/4$. We will find that the corresponding limits I_2^* and I_4^* indeed can be considerably greater than I_1^* depending on the network parameters.

C. The macroscopic level problem

Note that the fact that the actual capacity limit on the microscopic scale, $I^*[f(\theta)]$, depends functionally on $f(\theta)$ does not destroy the separation into two subproblems on the microscopic and macroscopic scales: the microscopic limit still depends on the local characterization of the traffic only, on its directional distribution, but not on the global characterization, i.e., where the actual sources of a stream in any given direction are located. It is certainly true that having the functional $I^*[f(\theta)]$ instead of a constant I_1^* as the limit renders the routing problem on the macroscopic scale more difficult.

As mentioned, on the macroscopic scale the problem is the following: given a network area \mathcal{A} and the traffic matrix, find a routing system \mathcal{P} , i.e., a set of paths (smooth curves), such that at every point \mathbf{x} , the local microscopic scale capacity

constraint is satisfied. With the multidirectional forwarding capacity it reads,

$$\Phi(\mathbf{x}; \mathcal{P}) \leq I^*[f(\theta, \mathbf{x}; \mathcal{P})] \quad \forall \mathbf{x} \in \mathcal{A},$$

where the scalar flux Φ and the directional distribution $f(\theta)$ are now functions of \mathbf{x} as determined by the routing system \mathcal{P} . In particular, the network capacity problem is to find a routing system \mathcal{P} such that the above condition is satisfied with the maximal possible scalar multiplier of a given form of the traffic matrix. To be explicit, this leads to the following modified load balancing problem

$$\max_{\mathcal{P}} \min_{\mathbf{x} \in \mathcal{A}} I^*[f(\theta, \mathbf{x}; \mathcal{P})]/\Phi(\mathbf{x}; \mathcal{P}),$$

where the numerator, instead of being constant I_1^* as was assumed in the earlier work [8], [18], now depends on \mathcal{P} . When the above maxmin problem is solved with a unit traffic matrix with the total traffic of 1 pkts/s, then the maxmin value gives the network capacity. While solving the maxmin problem is outside the scope of this paper, we will return to the question of the impact of multidirectional forwarding on the macroscopic level problem later on in Section V-D.

D. Dimensional analysis

In the microscopic level problem, the number of parameters needed to describe the problem can be reduced by dimensional analysis [1]. The multidirectional forwarding capacity depends on the physical parameters at hand: density of nodes n [1/m²], transmission range ρ [m], and nominal capacity of a link C [1/s]. For a given directional distribution $f(\theta)$, the multidirectional forwarding capacity, $I^*[f(\theta)]$, can be expressed as any combination of the parameters having the dimension 1/m/s times a function of all the independent dimensionless parameters that can be formed. A combination of parameters of dimension 1/m/s is provided by $C\sqrt{n}$, and there is only one dimensionless parameter, namely the mean number of neighbors (or the mean degree in undirected graph) of a node $\nu = n\pi\rho^2$ (the constant π is unimportant as it can be absorbed in the definition). Thus,

$$I^*[f(\theta)](C, n, \rho) = C\sqrt{n} u^*(\nu; f), \quad (1)$$

where u^* is an unknown dimensionless function to be determined.

III. GREEDY MAXIMUM WEIGHT SCHEDULING

In this section, we describe a method for resolving the problems related to finding a feasible way to schedule the transmissions efficiently in order to simulate the microscopic level problem. The (centralized) greedy maximum weight scheduling is used here to obtain an estimate of $I^*[f(\theta)]$, i.e., in practice, it cannot be used for scheduling when the packets are routed along the actual macroscopic level paths. We also modify the basic greedy maximum weight scheduling algorithm to be more suitable for traffic flowing in the four cardinal directions.

²The multidirectional forwarding capacity for isotropic traffic is at least as great as that for any other traffic pattern, as uniform distribution can be constructed as a linear superposition of rotated (by $d\theta$) copies of any other distribution and correspondingly served in time-sharing manner.

A. Maximum weight scheduling

In their paper [20], the authors present a maximum throughput policy that stabilizes the network for all arrival rates of multiclass traffic for which it is stabilizable. In our setting, the customer classes of the algorithm correspond to the traffic flows in different directions. In the simulations, the classes are discrete. Generally, the traffic in the direction increment $(\theta, \theta + d\theta)$ is equivalent to a class. The original algorithm also allows multiple link capacities, but here all the links are assumed to have the same nominal capacity.

The algorithm has three stages. In time slot t , the first stage is to calculate a weight w_l^t for each link $l \in \mathcal{L}$ as follows,

$$w_l^t = \max_j (q_j^{t-1}(t_l) - q_j^{t-1}(r_l)), \quad (2)$$

where j is the class index and $q_j(t_l)$ ($q_j(r_l)$) is the queue length of class j at node t_l (r_l). In the second stage, a maximum weight transmission mode is selected

$$m^*(t) = \arg \max_{m \in \mathcal{M}} \sum_{l \in m} w_l^t. \quad (3)$$

Finally, in the third and last stage, if we index the links with $i = 1, \dots, |\mathcal{L}|$, and denote by j^* the direction for which $w_i^t = q_{j^*}^{t-1}(t_l) - q_{j^*}^{t-1}(r_l)$, we get the activation vector $E(t)$ at time slot t as follows,

$$E_{ij}(t) = \begin{cases} 1, & \text{if } l_i \in m^*(t), j = j^*, \text{ and } q_j^{t-1}(t_{l_i}) > 0, \\ 0, & \text{otherwise.} \end{cases}$$

In each timeslot the policy finds the transmission mode that is of the maximum weight. Since all the links have the same nominal capacity, the weight of each link is simply the maximum (over the classes) difference in the queue lengths between the transmitting and receiving end of the link. A link in the maximum weight transmission mode is always activated if the transmitter has a packet to send.

B. Basic greedy algorithm

If it were possible to run the maximum weight algorithm properly, we could find out the true forwarding capacity. Unfortunately, finding the maximum weight transmission mode (3) is NP-complete. Because of the limitations in the computing capacity and the large network size, it is necessary to simplify the problem. Therefore, we choose the links of the transmission mode greedily.

In the basic greedy algorithm, the used transmission mode is chosen in each time slot in the following fashion. The links are sorted in the descending order according to their weight (2). The heaviest link is added to the transmission mode, and all the links interfering with it are removed from the list of possible links. The next heaviest link (that does not interfere with first one) is then added to the transmission mode. Again, the links interfering with the added link are removed from the list. The procedure is continued until there are no links left to add to the transmission mode, i.e., the mode is maximal.

The basic greedy algorithm was found to work well in all the cases when ν is small. When ν is very large, however, better results were obtained in the case of traffic in four cardinal

directions by making time sharing, with equal shares, between the two pairs of opposite streams and applying the greedy algorithm for one pair in turn.

This behavior can be understood as follows. According to [17], if links can be arbitrarily placed on a continuous plane, then the most efficient way of packing them for carrying traffic in the horizontal direction is to form vertical columns of horizontal links. The Boolean interference model sets no limit on how close two parallel links of maximal length, ρ , can be. When the direction of the columns alternate the endpoints close to each other are all either transmitters or receivers and a small distance ε is enough between the columns.

The above suggests that for a high ν , when there are nodes almost everywhere, a good strategy is to try to form vertical columns. Since the transport network is very dense, the same kind of transmission mode can be used in every time slot. The configuration is just rotated and/or moved by ε resulting in a flow where the same packet is moved once in every ρ/ε time slots allocated for this pair of opposite directions. Thus, packets traversing in two opposite directions are transmitted in a single timeslot, and different directions are handled via time sharing.

In order to have a single algorithm that covers all values of ν in the case of equal traffic streams in the four cardinal directions, we developed an improved algorithm described in detail in the next subsection.

C. Improved greedy algorithm

In the basic greedy algorithm, as described, the links are ordered in a list according to their weights. The links are chosen from this list starting from the link with greatest weight and skipping links that are in conflict with some of the already chosen links.

In the improved algorithm we introduce two parameters, one discrete, $k = 1, 2$, and one continuous, $\beta \in [0, 1]$. For given values of k and β , the algorithm works as follows. First, the greatest weight of all the links, w_{\max} , is found. Then, priority is given to the pair k of opposite streams, $k = 1$ corresponds to the left-right pair, and $k = 2$ corresponds to the up-down pair. Weights (2) are calculated taking into account only the two customer classes in the preferred pair of directions. Links are ordered according to these weights and conflict-free links are chosen as in the basic greedy algorithm from this list, starting from the link with the greatest weight, as long as the weight exceeds the value $\beta \cdot w_{\max}$. After this point, the weights (2) are calculated for all the remaining links taking into account all four customer classes and links are chosen as in the basic greedy algorithm.

For each setting of the parameters (k, β) we get a different algorithm yielding a different set of selected links with associated total weight. In principle, we could make an optimization to find the parameters that give the highest total weight in each time slot, to become as close as possible to the max weight scheduling. However, making such an optimization in every time slot for the continuous parameter β would be too time consuming. Therefore, in the algorithm we use a fixed value of

β , the same in all time slots. This value is, however, optimized externally to give maximal flow for a given ν . Optimization over the two values of k is done explicitly in each time slot, i.e., both values are tried, and the one that gives the higher total weight is selected.

Note that when $\beta = 1$ the algorithm reduces to the basic greedy algorithm. One can also see that when $\beta = 0$ the result cannot be worse than time sharing between the two pairs of opposite directions. Hence, with the optimized k and β , we always get results that are at least as good as those of the two simpler algorithms, as substantiated in Figure 3 of Section V.

IV. IMPLEMENTATION ISSUES

This section covers the simulation setup and the factors affecting the accuracy of the results.

A. Simulation setup

Since in a massively dense network the nodes communicating with each other are much further apart than two neighboring ones, a route between a source and a destination consists of a large number of hops. Therefore, the relay traffic dominates the amount of traffic in a specific area of the network. In order to simulate this setting, we consider a finite transport network with a set of artificial sources and sinks on different sides of the network. The sources generate the relay traffic that actually consists of the packets of multiple OD-pairs with the same direction, as determined by the macroscopic level routing \mathcal{P} , at this point of the network. The problem is studied as a function of the network size to determine a size that properly represents an infinite network.

The transport network with an average of N nodes resides in a unit square that has a strip of sources added to each side of the network with a large fixed queue length of q_0 . Each source doubles as a sink for the sources of the opposite strip. The queue length at sinks is always set to be zero. The width of the strips is ρ .

The top and the bottom of the unit square are connected together to form a cylinder and the ends of the 'unit-square' cylinder are connected together to further form a torus. The packets moving in the horizontal direction (originating from the sources on the left or right) can move over the seam of the cylinder. The packets moving in the vertical direction can move over the other seam of the torus and also see the network as a cylinder. This is done to reduce the border effects.

The above construction allows us study multidirectional forwarding capacities when the directional distribution $f(\theta)$ contains traffic in two and four directions, in addition to the single-directional forwarding capacity. We are still limited to the cases with right angles between the flows, but these cases include arguably the most interesting ones, e.g., two flows in opposite directions. If all of the four directions are not used the corresponding sources are turned off.

As a result, the method gives the average progress of packets in their direction per slot, \bar{x} , when the network has been stabilized. Alternatively, it is possible to monitor the number of packets per slot arriving to the sinks or the number of

packets leaving the sources. When these three quantities are approximately the same, the simulation can be considered to have reached steady state. As we have a unit-area network, we can calculate the dimensionless component of the achieved forwarding capacity as $u = \bar{x}/\sqrt{N}$.

Now, there are three sources of error affecting the result. Firstly, the true maximum weight transmission mode is not found but a suboptimal greedy approach is used. Secondly, the fixed queue length of the sources, q_0 , affects the set of possible link weights and thus the accuracy of the scheduling. The queue length initialization of the other nodes is also important because of the length of the initial transient. Thirdly, the size of the network used in the simulations should be as large as possible to properly represent the infinite network. The effects of these approximations are discussed in the following.




Fig. 1. The dimensionless function u for four directions and $\nu = 10$ as a function of N for different values of q_0 .

B. Greedy scheduling

The main factor limiting the applicability of the original method is our inability to solve the maximum weight independent set problem exactly since this is a computationally complex problem. Although the improved algorithm changes the order of the links compared to their original weights, both the basic and the improved version of the algorithm still select the activated links from an ordered list. Because of this greedy scheduling, we get a lower bound for the capacity with the particular size of the simulated network.

C. Queue length initialization

When the number of packets in the network is small, the link weights are also small, and the probability that two links have equal weights is notable. Since the choice between two transmission modes that are of the maximum weight is arbitrary, the operation of the scheduler is, at this point, more random. This leads to suboptimal performance that results in increasing queue lengths. When the queues grow, some of the ties are resolved, and eventually the policy stabilizes the queues for all arrival rates for which it is possible. This might however require very long queues. When the queue length at sources is fixed, the number of possible link weights is also fixed. This means that to assure the best possible performance, the queue length at sources, q_0 , should be as large as possible.

In practice, the choice of the queue length at sources is a compromise since a large q_0 means a long initial transient. Because of the finite queue length at the sources, we get a capacity lower than the true capacity. Figure 1 shows the effect in different size networks when the queue length at sources varies from 100 to 1000 packets.

D. Network size

The size of the simulated network needs to be large enough in order for the results to be meaningful for a network with infinite size. A small network gives too high values for the achievable forwarding capacity since it is easier to establish a flow through shorter paths.

Because of the finite network size, we get a capacity higher than the true capacity. By eliminating the effect of the limited network size, it would be possible to study capacities that are actually achievable with greedy maximum weight scheduling. Figure 1 represents u as a function of the network size N .

V. RESULTS AND DISCUSSION

This section contains the numerical results and the discussion on their significance on more general directional distributions as well as on the macroscopic level problem.

A. Improved greedy algorithm

The performance of the improved greedy algorithm as a function of the priority threshold parameter β with $q_0 = 100$ is presented in Figure 2. The figure illustrates how the optimal value of β becomes smaller as the transport network density, ν , increases. The data points are averages over 10 network realizations and the error bars show the 95 % confidence intervals. The maximum values of each curve, corresponding to the optimization over β , form the capacity curve of the improved algorithm. This curve is represented in Figure 3 along with the results from the basic greedy algorithm with two and four directions. The figure shows clearly how the improved algorithm is able to outperform the basic greedy algorithm with four directions when ν is large, achieving the level of the two-directional one, but is still able to utilize the multidirectional gain from all of the four directions when ν is smaller.




Fig. 2. The results from the improved greedy algorithm as a function of the threshold β .




Fig. 3. The improved algorithm for traffic in four directions compared with the basic greedy algorithm with two and four directions.

B. Comparison with the single-directional capacity

Figure 4 presents the dimensionless components, u^* in (1), of I_1^* , I_2^* , and I_4^* obtained using $N = 1000$ and $q_0 = 100$ from the basic greedy maximum weight scheduling algorithm with one and two directions and the improved greedy maximum weight scheduling algorithm with four directions. The results are averages over 10 network realizations and the error bars show the 95 % confidence intervals. The values of q_0 and N are a practical compromise, based on Figure 1, between the accuracy and the necessity to keep the simulation times reasonable.




Fig. 4. Multidirectional forwarding capacity for traffic in one, two and four directions.

As can be seen from the figure, the multidirectional gain is, indeed, considerable ranging from a factor of 1.5 in a dense transport network to over 2 in a sparse transport network. Traffic in two opposite directions is enough to generate the gain in a dense transport network, as discussed in Section III-B, but having the four-directional distribution is beneficial in a sparse transport network. In general, the possibility to use more directions can only improve the result as it is always possible to regress back to using only a subset of the directions in a time-shared manner. It is safe to presume that the four cardinal directions with equal flows should already be a good approximation for the isotropic traffic.

The forwarding capacity with two opposite directions is close to 1.5 times the single-directional when the network

is sparse. This can be understood as follows. When the transport network is operated close to the percolation threshold ($\nu \approx 4.5$), the number of paths connecting different sides of the network is small. The small number of paths also implies that they are far from each other and the interference between two paths is negligible. When the schedule is chosen in the way that the links of the path only interfere with the previous and following link of the path, it is possible to use every third link of the path simultaneously with a single direction. This idea has been utilized in [16] to construct a lower bound for the forwarding capacity also when the mean neighborhood size, ν , is larger. In the multidirectional case, it is possible to activate more than every third link. With two opposite directions, every other link can be active when the directions of the transmissions alternate. A single-directional schedule consisting of three transmission modes leads to capacity $C/3$ for a single path while a two-directional schedule consisting of four transmission modes leads to capacity $C/2$ for a path. Hence, the two-directional case gives 1.5 times the single-directional capacity when the transport network is sparse.

C. On the accuracy of the results

For a very small ν , the error in the simulations is mainly due to the network size, N . The $C/3$ flow that the one-directional scheduling is able to achieve along a single path is maximal. Also the two-directional greedy scheme can schedule close to the maximum flow of one half along a path. Thus, the results do not suffer from the greedy heuristic with the smallest neighborhoods. The used network size of 1000 nodes leads to a slightly too high capacity as there is more likely to be better connectivity through a small network.

When ν is large, the effect of N is less relevant since the network is more homogeneous from the scheduling point of view. That is to say, there are always multiple possible links to choose from, and no clear bottlenecks appear as the network is made larger. The capacities are achievable since the greedy scheduling moves the result downwards. Greedy scheduling is not able to coordinate the transmissions efficiently enough when the number of interfering links is large, and, in addition to the weight, one should also consider how the links interfere with each other.

This is manifested in Figure 3 in that the greedy scheduling performs worse with four directions than with two when the network is dense. As concluded in Section III-B, the optimal configuration for a dense network uses only two opposite directions at a time. The two-directional greedy scheduling automatically selects transmission modes that have this characteristic, and it is thus able to outperform the four-directional greedy method that cannot achieve the same spatial reuse. The improved greedy algorithm is able to rectify this for four-directional traffic, but it is still unable to coordinate the transmissions further, e.g., to form columns.

D. Discussion

Next we consider what can be said about $I^*[f(\theta)]$ with a general directional distribution $f(\theta)$. Assume that the multi-

directional forwarding capacities are known for some set of directional distributions, $\{g_1(\theta), g_2(\theta), \dots\}$, and let these capacities be $\{J_1^*, J_2^*, \dots\}$, respectively. We can always express the distribution $f(\theta)$ in the form

$$f(\theta) = \sum_i a_i g_i(\theta) + b h(\theta), \quad (4)$$

where the a_i and b are non-negative constants, and the remainder term $h(\theta) \geq 0$ for all θ . Note that also $h(\theta)$ represents a distribution (with the integral over the angle equaling one).

Given $f(\theta)$, we try to determine the maximal scalar flux Φ , i.e., the constant multiplier in front of $f(\theta)$, such that the traffic can be sustained. Each of the components of the sum in (4) can be handled in $\Phi \cdot a_i / J_i^*$ fraction of time. The remainder requires a fraction smaller than or equal to $\Phi \cdot b / I_1^*$, where I_1^* is the single-directional forwarding capacity. The total traffic can be sustained using time sharing between the components if the sum of the time shares is at most 1. Then, for this $f(\theta)$ we have the following lower bound for the multidirectional forwarding capacity

$$I^*[f(\theta)] \geq \left(\sum_i a_i / J_i^* + b / I_1^* \right)^{-1}.$$

This is a sure lower bound since the handling of the remainder term is upper bounded (by only using time-shared single-directional forwarding) and also since time-sharing among the components i is not necessarily optimal.

In our case, the directional distributions for which the multidirectional forwarding capacity is known are the single-, bi- and four-directional balanced traffic patterns. These can be utilized to forward non-balanced four-directional traffic by first separating the four-directional balanced traffic pattern, in which case the remaining traffic equals zero in at least one direction. In the other orthogonal direction, the balanced bidirectional traffic can again be extracted. This only leaves two single-directional orthogonal flows that can be handled using time sharing with single-directional forwarding. This yields the lower bound. By rotating this pattern over all angles $(0, \pi/2)$, a lower bound is obtained for $I^*[f(\theta)]$ for any directional distribution. Explicitly, we have

$$I^*[f(\theta)] \geq \left(\frac{K_1 - K_2}{I_1^*} + \frac{K_2 - K_4}{I_2^*} + \frac{K_4}{I_4^*} \right)^{-1},$$

where

$$\begin{cases} K_1 = \int_0^{2\pi} f(\theta) d\theta = 1, \\ K_2 = 2 \int_0^\pi \min\{f(\theta), f(\theta+\pi)\} d\theta, \\ K_4 = 4 \int_0^{\pi/2} \min\{f(\theta), f(\theta+\frac{\pi}{2}), f(\theta+\pi), f(\theta+\frac{3\pi}{2})\} d\theta. \end{cases}$$

Finally, we make a remark on the impact of multidirectional forwarding in the macroscopic level routing problem with a uniform traffic matrix. Under the assumption that only single directional forwarding is used at the microscopic level, the macroscopic-level routing problem is to determine the routes so that the maximum local load is minimized, i.e., a problem

of load balancing, see [8], [18]. This causes the routes to be long so that traffic is pushed away from the center towards the edges, in order to avoid congesting the center of the area. Now, the multidirectional forwarding capacity increases the capacity compared with single directional forwarding, especially when the traffic is nearly isotropic. Because the traffic is naturally more isotropic at the center of the area (e.g., a disk) than at the border, there is less need to push the traffic away from the central area, and the optimal paths under multidirectional forwarding will be more straight. It might even be that shortest path routing (direct lines) is close to optimum.

VI. CONCLUSIONS

In this paper, we studied the forwarding capacity of a homogeneous infinite network of Poisson distributed nodes for multidirectional, non-isotropic but uniform traffic under the Boolean interference model. This problem is relevant for determining the network capacity of a massively dense network, where a natural spatial separation of scales occurs, and the multidirectional forwarding problem constitutes the microscopic level part of the overall problem.

By using a constructive, simulation-based approach, we were able to obtain numerical results for the hitherto unknown multidirectional forwarding capacity, which depends both on the directional distribution of the traffic and on the network parameters through the mean node degree. More precisely, we applied the multiclass maximum weight scheduling with a greedy heuristic for finding a good candidate for the maximum weight transmission mode in each time slot. This provided a practical approach for our simulations, allowing us to consider reasonably large networks and to produce what we believe are the most accurate numerical results that are presently available.

Specifically, the multidirectional forwarding capacity was determined as a function of the mean number of neighbors of a node for different directional traffic patterns: bi- and four-directional balanced traffic. The results confirm the intuition that as the traffic becomes more isotropic the total forwarding capacity increases. Indeed, a notable gain can be achieved over the single-directional forwarding capacity.

An interesting observation was that in the regime of high network density, i.e., high node degree, bidirectional balanced traffic defines the maximum capacity. This is due to the fact that in a dense network the Boolean model allows very efficient packing of links by stacking them into columns of parallel links, pointing to one traffic direction in one column and to the opposite direction in the adjacent column. The best strategy for any balanced traffic pattern is to make time sharing between opposite pairs of traffic in different directions. This insight was also used for designing an improved greedy algorithm for four-directional traffic, that is able to produce good results both at low- and high-density regimes. A general, time-sharing lower bound for any directional distribution of the traffic was derived based on our results for the three basic traffic patterns.

The possibility for the very tight packing of the transmitting links in dense networks is a peculiarity of the Boolean

interference model. If the footprint of an active link is larger, as it is when using, e.g., RTS/CTS-type two-way handshake or additive interference, the gain is likely to be smaller. The study of other interference models and a more detailed study of the macroscopic level problem are left as future research.

ACKNOWLEDGEMENTS

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

REFERENCES

- [1] E. Buckingham. On physically similar systems: illustrations of the use of dimensional equations. *Phys. Rev.*, 4(4):345–376, 1914.
- [2] 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.
- [3] P. Chaporkar, K. Kar, and S. Sarkar. Throughput guarantees through maximal scheduling in wireless networks. In *43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005.
- [4] A. Eryilmaz, O. Asuman, and E. Modiano. Polynomial complexity algorithms for full utilization of multi-hop wireless networks. In *IEEE INFOCOM*, pages 499–507, 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] Esa Hyttiä and Jorma Virtamo. Near-optimal load balancing in dense wireless multi-hop networks. In *NGI 2008, 4th Conference on Next Generation Internet Networks*, pages 181–188, 2008.
- [8] Esa Hyttiä and Jorma Virtamo. On the optimality of field-line routing in massively dense wireless multi-hop networks. *Performance Evaluation*, 66(3–5):158–172, 2009.
- [9] P. Jacquet. Geometry of information propagation in massively dense ad hoc networks. In *ACM MobiHoc*, pages 157–162, 2004.
- [10] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *ACM MobiCom*, pages 66–80, 2003.
- [11] L. Jiang and J. Walrand. A distributed csma algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transactions on Networking*, 18(3):960–972, 2010.
- [12] M. Kalantari and M. Shayman. Energy efficient routing in wireless sensor networks. In *Conference on Information Sciences and Systems*, 2004.
- [13] M. Kodialam and T. Nandagopal. Characterizing achievable rates in multihop wireless mesh networks with orthogonal channels. *IEEE Transactions on Networking*, 13(4):868–880, 2005.
- [14] X. Lin and N.B. Shroff. The impact of imperfect scheduling on cross-layer rate control in wireless networks. In *IEEE INFOCOM*, 2005.
- [15] J. Nousiainen and P. Lassila. Approximating maximum directed flow in a large wireless network. In *IEEE ICC*, 2009.
- [16] J. Nousiainen, J. Virtamo, and P. Lassila. On the achievable forwarding capacity of an infinite wireless network. In *ACM MSWiM*, pages 151–159, 2010.
- [17] 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, 2011.
- [18] L. Popa, A. Rostamizadeh, R.M. Karp, C. Papadimitriou, and I. Stoica. Balancing traffic load in wireless networks with curveball routing. In *ACM MobiHoc*, 2007.
- [19] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *IEEE INFOCOM*, pages 533–539, 1998.
- [20] 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.
- [21] X. Wu, R. Srikant, and J.R. Perkins. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. *IEEE Transactions on Mobile Computing*, 6(6):595–605, 2007.