Telescopic beamforming for large wireless networks

Alla Merzakreeva, Olivier Lévêque
Swiss Federal Institute of Technology - Lausanne, Switzerland
{alla.merzakreeva, olivier.leveque}@epfl.ch

Ayfer Özgür
Stanford University, Stanford CA
aozgur@stanford.edu

***Abstract*—We consider a wireless network with n users distributed over a square area $A = n$. Under line-of-sight propagation, this network has $\Theta(\sqrt{n})$ degrees of freedom. At high SNR, these degrees of freedom can be readily achieved by multi-hop relaying between nodes. At low SNR however, the performance is determined by the power transfer in the network. We show that none of the existing architectures can achieve optimal capacity scaling. We develop a beamforming architecture where signals are relayed by coherent combining over multiple clusters. The key ingredient is an analysis of the beamforming gain achievable between two clusters of nodes under the line-of-sight propagation model. This result reveals a new regime for large two-dimensional wireless networks, where beamforming techniques are needed to achieve capacity.**

I. Introduction

Multi-hop is the traditional communication architecture for wireless adhoc networks. Information is routed from source nodes to destinations via multiple point-to-point transmissions between intermediate nodes acting as relays. Is multi-hop fundamentally good for wireless adhoc networks or can we design new cooperation architectures that significantly outperform multi-hop? This question has been extensively studied in the large network regime, following the *scaling law* approach first introduced in [2] and followed in [3], [6], [7], [8], [10].

The current progress shows that the answer depends on the operating regime of the network. The operating regime of a large wireless network is determined by the average SNR between neighboring nodes, the spatial degrees of freedom of the network, dictated by the area and the carrier wavelength, and by the power path loss exponent of the environment [7], [9]. When there are sufficient spatial degrees of freedom in the network (this is, for example, the case when the pairwise channels are subject to i.i.d. fading), a hierarchical cooperation architecture based on distributed MIMO transmission can exploit these degrees of freedom [8]. This provides significant capacity gain over multi-hop both in the high and the low SNR regimes. When the spatial degrees of freedom in a wireless network are limited by physical constraints, however, these few degrees of freedom can be readily achieved by multi-hop. Multi-hop becomes scaling optimal for such networks in the high SNR regime, where the capacity is degrees-of-freedom -limited [1].

The situation is more delicate for networks at low SNR and with limited spatial degrees of freedom. Multi-hop is not anymore optimal, as the performance is limited by the power transfer in the network, and not the degrees of freedom achieved. Earlier in [4], we considered the extremal case when there is only a single degree of freedom for communication in the network: a one-dimensional wireless network. We showed that a hierarchical beamforming architecture can significantly outperform existing strategies in this case, providing an n-fold capacity increase over multi-hop in a network with n users. Here, nodes first broadcast their information to a small cluster around them. In a second step, nodes in this cluster beamform this information to a larger cluster. Continuing in a hierarchical fashion, the information of each source node is broadcasted to the whole network, including the destination node.

It is non-trivial to extend this architecture to two-dimensional networks. A two-dimensional network with n users distributed uniformly over an area A has at least $\sqrt{n}$ degrees of freedom [9]. Therefore, transmissions can not coherently combine simultaneously at many destination nodes providing large beamforming gains as in the one-dimensional case. We analyze the beamforming gain achievable between two clusters of nodes under the line-of-sight propagation model, and show that it critically depends on the shape, orientation and distance between the two-dimensional clusters. We develop a communication architecture that uses bursty amplify-and-forward between successive clusters, with cluster sizes and shapes chosen carefully to ensure maximum beamforming gain at each step. The architecture allows to achieve a throughput significantly larger than that achieved by multi-hop cooperation ($n^{6/7}\,\mathrm{SNR}_s$ versus $\sqrt{n}\,\mathrm{SNR}_s$ for multi-hop, where SNR_s is the typical nearest neighbor SNR in the network and $\mathrm{SNR}_s \ll 0$ dB in the low-SNR regime). It has a small gap to the information theoretic upper bound $n\,\mathrm{SNR}_s$. Achieving the exact capacity remains work in progress.

II. Model

There are n nodes uniformly and independently distributed in a square of area $A = n$ (so the node density $\rho = A/n = 1$). Every node is both a source and a destination, and the sources and destinations are randomly paired up one-to-one. All source nodes want to communicate to their destination node at the same rate $R(n)$. The aggregate throughput of the network is defined as $T(n) = nR(n)$.

We assume that communication takes place over a flat channel with bandwidth W and that the received signal $Y_k[m]$ by node k at time m is given by

$$Y_j[m] = \sum_{k\in\mathcal{K}} h_{jk} X_k[m] + Z_j[m]$$

where $\mathcal{K}$ is the set of transmitting nodes, $X_k[m]$ is the signal sent at time m by node k and $Z_j[m]$ is additive white circularly symmetric Gaussian noise (AWGN) of power spectral density $N_0/2$ Watts/Hz. In a line-of-sight environment, the complex baseband-equivalent channel gain h_{jk} between transmit node k and receive node j is given by

$$h_{jk} = \sqrt{G}\, \frac{\exp(2\pi i r_{jk}/\lambda)}{r_{jk}} \tag{1}$$

where G is Friis' constant, λ is the carrier wavelength and r_{jk} is the distance between node k and node j. This line-of-sight model clearly departs from the classical i.i.d. phase assumption: it implies in particular that the channel matrix between two clusters of nodes is not necessarily full-rank. Finally, we assume full channel state information at both the transmitters and receivers (which is a reasonable assumption in a static line-of-sight environment), as well as a common average power budget per node P.

III. Main result

Let us denote by SNR_s the signal-to-noise ratio over the typical nearest neighbor distance in the network. In our setup, the typical nearest neighbor distance is 1, therefore, the short-distance SNR is

$$\mathrm{SNR}_s = \frac{GP}{N_0 W}.$$

In this paper, we are interested in the low SNR regime, where $\mathrm{SNR}_s = n^{-\gamma}$ for some $\gamma > 1$. We will use the notation $\mathrm{SNR}_s \ll 0$ dB to denote this condition[1]. This models the scenario when the pairwise channels between nearest neighbors are in the low SNR regime. Note that since SNR_s is jointly determined by system parameters P and W, this can be the case when the available power per node is small or when the bandwidth is large.

In the low SNR regime and under the line-of-sight model described above, the multi-hop scheme proposed in [2] achieves an aggregate throughput of order

$$T(n) = \Omega\Big(\mathrm{SNR}_s \sqrt{n/\log n}\Big)$$

with high probability as n gets large. While multi-hop can be shown to be order optimal at high SNR (see [1]), the best known information theoretic upper bound on the throughput scaling at low SNR is of order

$$T(n) = O\Big(\mathrm{SNR}_s\, n \log n\Big)$$

This is obtained by showing that every source node in the network cannot transmit more than $O(\mathrm{SNR}_s \log n)$ bits to the rest of network, therefore the aggregate throughput can not exceed n times this number. The question therefore remains whether a more sophisticated strategy would allow to achieve higher throughput scaling than multi-hop at low SNR. We answer this question by the affirmative in the following theorem.

[1]The extension of our results to the intermediate SNR regime $0 \le \gamma \le 1$ is not trivial. Here the capacity is both power and degrees-of-freedom limited and the beamforming techniques considered in this paper become suboptimal.

Theorem 1. *Let us assume that* $\mathrm{SNR}_s \ll 0$ *dB (i.e.* $\mathrm{SNR}_s = n^{-\gamma}$ *for some* $\gamma > 1$*). Then there exists a communication scheme (referred to as "telescopic beamforming" in the sequel) that achieves the following aggregate throughput scaling with high probability as* n *gets large:*

$$T(n) = \Omega\left(\mathrm{SNR}_s\, n^{6/7-\varepsilon}\right) \tag{2}$$

for every $\varepsilon > 0$.

At low SNR, the telescopic beamforming scheme achieves therefore a throughput scaling close to the best known upper bound and clearly outperforms classical multi-hop. Before giving a description of this scheme, we present in the next section a simple scheme, whose throughput scaling is lower than that of the telescopic beamforming scheme, but already higher than that of the multi-hop scheme.

IV. Two-cluster scheme

In this section, we first describe a simple three-stage relaying scheme that allows to enhance communication rates at low SNR in the network, and then proceed to its performance analysis.

In the low SNR regime, the performance is not degrees-of-freedom but power-limited, so we can handle the communication between the source-destination pairs one at a time. Let us therefore consider a single source-destination pair. Here is the basic structure of the scheme: the source node s communicates its signal to the destination node d with the help of two clusters of relay nodes, each of size M, one surrounding s and the other surrounding d, as illustrated on Fig. 1. The first cluster amplifies and forwards the signals received from the source to the second cluster, where the signals are again amplified and forwarded to the destination.

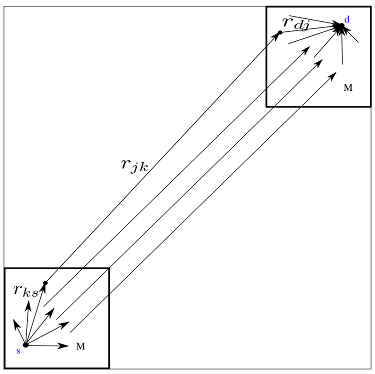

Fig. 1: Two-cluster scheme.

Notice that in the low SNR regime, the signal power is much smaller than the noise power; a pure amplify-and-forward strategy would therefore mostly amplify noise. This can be compensated by using a *bursty* amplify-and-forward scheme, following [5], where the nodes operate only a small fraction of time with increased power and stay silent the rest of the time. A second key ingredient of the scheme is the *beamforming gain* achieved by letting the relay nodes coherently combine their signals at each retransmission. The performance of the scheme is stated in the theorem below.

Theorem 2. *In the low SNR regime, the following throughput scaling is achievable:*

$$T(n) = \Omega\Big(\mathrm{SNR}_s \, n^{2/3}\Big)$$

with high probability as n gets large. The scheme achieving this throughput scaling is a bursty amplify-and-forward scheme employing two clusters of M relay nodes with M chosen such that

$$n^{1/3} \leq M \ll n^{1/2}. \tag{3}$$

Before going into the details of the proof, let us briefly explain the trade-off behind condition (3). On the one hand, it is desirable to have many nodes participating to the transmission (and therefore a large cluster size M) in order to combat noise amplification via coherent combining of the signals. On the other hand, we will see that under the line-of-sight propagation model (1), coherent combination of the signals can only take place when the cluster size M is smaller than the inter-cluster distance $\sqrt{n}$. We elaborate on this fact in Section V, as it plays an important role for the extensions of the scheme described in later sections.

Proof sketch. In order to simplify notation, we assume in the following, without loss of generality, a choice of units such that $\mathrm{SNR}_s = P$.

We consider that all nodes operate only a fraction of time $1/a$, with a being some constant greater than 1. This compensates for low power by providing a power gain factor a at each transmission. We also assume that s-d pairs operate one at a time in the network, so that we may focus on a single s-d pair.

Equivalent end-to-end channel model. The source node s communicates a signal X to the destination node d in 3 steps:

- It first broadcasts the signal X to a square cluster of M neighboring nodes. The signal received by node k in the cluster is given by

$$Y_k^{(1)} = \frac{e^{2\pi i r_{ks}/\lambda}}{r_{ks}} \, X + Z_k^{(1)},$$

  where r_{ks} is the distance between the source node and node k, and $Z_k^{(1)}$ is additive white Gaussian noise.
- The received signals are then amplified and forwarded to the square cluster around the destination, containing also M nodes and located at a distance of order $\sqrt{n}$. The signal received by node j in this cluster is given by

$$Y_j^{(2)} = \sum_{k=1}^{M} \frac{e^{2\pi i r_{jk}/\lambda}}{r_{jk}} \, X_k^{(1)} + Z_j^{(2)},$$

  where r_{jk} is the distance between transmit node k and receive node j, $Z_j^{(2)}$ is additive white Gaussian noise and $X_k^{(1)} = C_k \, Y_k^{(1)}$ is the signal emitted by node k, amplified to meet the power constraint.
- The nodes in the cluster around the destination finally retransmit, after amplification, the received signals to the destination node d. The signal received by the destination node is given by

$$Y = \sum_{j=1}^{M} \frac{e^{2\pi i r_{dj}/\lambda}}{r_{dj}} \, X_j^{(2)} + Z,$$

  where r_{dj} is the distance between node j in the cluster and the destination node d, Z is additive white Gaussian noise and $X_j^{(2)} = D_j \, Y_j^{(2)}$ is the signal emitted by node j, amplified to meet the power constraint.

At the end of the three steps, the signal received at the destination node d is therefore given by

$$Y = F \, X + \widetilde{Z}, \tag{4}$$

where

$$F = \sum_{j,k=1}^{M} \frac{e^{2\pi i r_{dj}/\lambda}}{r_{dj}} \, D_j \, \frac{e^{2\pi i r_{jk}/\lambda}}{r_{jk}} \, C_k \, \frac{e^{2\pi i r_{ks}/\lambda}}{r_{ks}} \tag{5}$$

and $\widetilde{Z}$ is the noise accumulated at the destination:

$$\begin{aligned} \widetilde{Z} \;&=\; Z + \sum_{j=1}^{M} \frac{e^{2\pi i r_{dj}/\lambda}}{r_{dj}} \, D_j \, Z_j^{(2)} \\ &\quad + \sum_{j,k=1}^{M} \frac{e^{2\pi i r_{dj}/\lambda}}{r_{dj}} \, D_j \, \frac{e^{2\pi i r_{jk}/\lambda}}{r_{jk}} \, C_k \, Z_k^{(1)}. \end{aligned}$$

Power constraints and amplifying coefficients C_k and D_j. Since s-d pairs operate one at a time in the network and all of them operate only a fraction $1/a$ of the time, the power constraint at the source node is $E(|X|^2) \leq anP$.

The relay nodes are busier, as they need to relay the transmissions for M different s-d pairs, so the power constraint at each relay node is $E(|X_k^{(1)}|^2) \leq anP/M$ in the first cluster and $E(|X_j^{(2)}|^2) \leq anP/M$ in the second cluster. The amplifying coefficients C_k and D_j should be chosen so as to satisfy these power constraints. Assuming that the amplifying factor a is chosen such that the incoming signal power does not exceed the noise power at each stage (which will indeed turn out to be an optimal choice at low SNR), we obtain

$$|C_k|, |D_j| \simeq \sqrt{anP/M}. \tag{6}$$

The phases of C_k and D_j should also be chosen so as to compensate for the phase shifts arising from the line-of-sight propagation in equation (5), in order to maximize the beamforming gain. Let therefore the phases ϕ_k of C_k be

$$\phi_k = -2\pi i (r_{ks} + x_k)/\lambda$$

where x_k denotes the horizontal coordinate of the transmit node k (i.e. x_k is the coordinate along the main direction of transmission to the other cluster: see Fig. 2). The first term $-2\pi i r_{ks}/\lambda$ cancels exactly the phase shift arising from the first transmission, while the second term $-2\pi i x_k/\lambda$ aims to compensate for the phase shifts $2\pi i r_{jk}/\lambda$ arising in the second transmission.

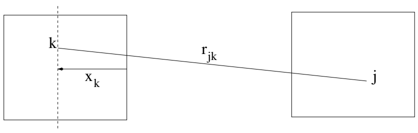

Fig. 2: Compensating factor x_k for phase shits.

As this second transmission is intended for multiple receive nodes j, it is a priori not guaranteed that compensating simultaneously for all phase shifts is doable. Using a second order Taylor expansion of the square root appearing in the expression for the inter-node distances r_{jk}, it can actually be shown that when the cluster size M is much smaller than the inter-cluster distance $\sqrt{n}$ (strictly speaking, when $M \le n^{1/2-\varepsilon}$ for some $\varepsilon > 0$) , the following holds:

$$E\left(\sum_{k=1}^{M} e^{2\pi i (r_{jk}-x_k)/\lambda}\right) = \Theta(M), \tag{7}$$

where the expectation is taken over the node positions in the receive cluster. An optimal beamforming gain is therefore ensured in this case. Similarly, it can be shown that under the same condition $M \ll \sqrt{n}$, the phases ψ_j of the coefficients D_j can be chosen so as to ensure a full beamforming gain for the last transmission towards the destination.

Back-of-the-envelop computation of the end-to-end SNR. Since all nodes operate only a fraction of time $1/a$, where $a > 1$, the aggregate throughput of the above two-cluster scheme is given by[2]

$$T(n) = \frac{1}{a}\,\log(1+\mathrm{SNR}_{\mathrm{tot}}),$$

where the end-to-end SNR is the one resulting from equation (4):

$$\mathrm{SNR}_{\mathrm{tot}} = \frac{|F|^2\,E(|X|^2)}{E(|\widetilde{Z}|^2)}. \tag{8}$$

In order to operate the system the most efficiently, we should tune the amplifying factor a and the cluster size M so as to obtain an end-to-end SNR of order 1. As seen above, the condition $M \ll \sqrt{n}$ is needed in order to allow for an (approximate) compensation of all phase shifts. A back-of-the-envelop computation of the numerator in (8) then gives

$$\begin{aligned} |F|^2\,E(|X|^2) &\simeq \frac{1}{M^2 n}\left(\sum_{j,k=1}^{M} |C_k|\,|D_j|\right)^2 anP \\ &\simeq \frac{M^4}{M^2 n}\,(anP/M)^2\,(anP) = \frac{(anP)^3}{n}, \end{aligned} \tag{9}$$

where we have used successively the approximations $r_{sk}, r_{dj} \simeq \sqrt{M}$, $r_{jk} \simeq \sqrt{n}$ and the approximation (6). Similarly, the noise $\widetilde{Z}$ accumulated at the destination can be approximated as

$$\widetilde{Z} \simeq Z + \frac{\sqrt{anP/M}}{\sqrt{M}}\sum_{j=1}^{M} Z_j^{(2)} + \frac{anP/M}{\sqrt{Mn}}\,M\sum_{k=1}^{M} Z_k^{(1)},$$

[2]Notice indeed that as s-d pairs operate one at a time, $T(n)$ represents the total number of bits per second that travel in the network.

so by the independence of the random variables $Z_j^{(2)}$ and $Z_k^{(1)}$, the denominator in (8) is of order

$$E(|\widetilde{Z}|^2) \simeq 1 + \frac{anP}{M} + \frac{(anP)^2}{n}. \tag{10}$$

From equations (8), (9) and (10), we see that in order to ensure an end-to-end SNR of order 1, the optimal choice for a and M is

$$a = \frac{1}{n^{2/3}\,P} \quad \text{and} \quad M \ge anP = n^{1/3},$$

which luckily does not contradict the condition $M \ll \sqrt{n}$ found previously. This finally implies that the two-cluster scheme achieves an aggregate throughput scaling

$$T(n) \simeq \frac{1}{a} = n^{2/3}\,P$$

and concludes the proof sketch of Theorem 2. □

V. Beamforming gain

In the two-cluster scheme presented in the previous section, it was stated in equation (7) that the maximal beamforming gain between two square clusters of size M separated by distance $\sqrt{n}$ can be achieved when $M \ll \sqrt{n}$, by using a proper compensation of the phase shifts at the transmit side. This claim can be generalized to the case of two rectangular clusters separated by arbitrary distance d.

Claim 1. Consider two clusters placed on the same horizontal line, of sizes $M_1 = d_1 \times e_1$ and $M_2 = d_2 \times e_2$, respectively, and separated by distance d, as illustrated on Fig. 3. Provided the following relationship holds:

$$e_1\,e_2 \ll \max(d, d_1, d_2) \tag{11}$$

the maximum beamforming gain between cluster M_1 and cluster M_2 can be achieved by using a proper compensation of the phase shifts at the transmit side, that is,

$$E\left(\sum_{k=1}^{M_1} e^{2\pi i (r_{jk}-x_k)/\lambda}\right) = \Theta(M_1).$$

where x_k denotes the horizontal position of node k and the expectation is taken over the node positions in the receive cluster.

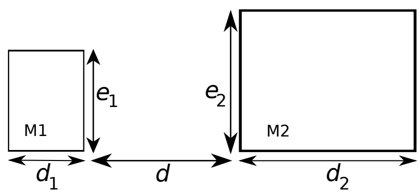

Fig. 3: Two rectangular clusters separated by distance d.

Remark 1. It is interesting to notice that condition (11) is also the condition under which the number of spatial degrees of freedom between the two clusters is minimal, that is, of order 1 (see [9]). At the other extreme, when the channel coefficients are i.i.d., for example, the spatial degrees of freedom are many, but no significant beamforming gain can be achieved by aligning the phases at the relays. We see here an interesting

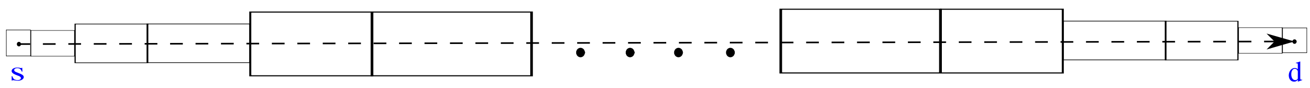

Fig. 4: Telescopic beamforming.

duality between the maximum achievable beamforming gain and the number of spatial degrees of freedom between the two clusters.

VI. Four-cluster scheme

We would like now to iterate the idea of the two-cluster scheme, so as to further enhance the throughput scaling. A natural idea seems to increase the number of relay stages, so as to:

1) reach larger and larger clusters of nodes and therefore increase the beamforming gain of transmissions;

2) take advantage at the same time of the burstiness of transmissions, and increase the transmit power with the number of stages.

However, given the restriction imposed by the claim made in the previous section, it is not trivial how to proceed. The difficulty lies in the fact that in order to remain optimal (within this type of scheme), transmissions should only occur between clusters that satisfy condition (11). It turns out in this case that taking clusters of rectangular shape allows to achieve higher beamforming gain and therefore higher end-to-end SNR.

This leads us to propose the following four-cluster scheme depicted in Fig. 5:

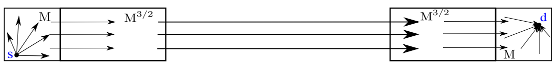

Fig. 5: Four-cluster scheme.

We assume here that the source node s communicates to the destination node d in five steps:

- It first broadcasts its signal to a square cluster of $\sqrt{M} \times \sqrt{M} = M$ neighboring nodes around itself.
- The received signals are then amplified and retransmitted to an adjacent rectangular cluster of size $\sqrt{M} \times M = M^{3/2}$ nodes; this ensures that the full beamforming gain can be achieved during the transmission.
- Similarly, the nodes in the rectangular cluster simultaneously rescale and retransmit their signals to a rectangular cluster of the same size at distance of order $\sqrt{n}$. During this transmission, the full beamforming gain can again be attained, as long as $M \ll \sqrt{n}$, according to condition (11). Notice that the condition is the same as in the two-cluster case, but the number of nodes in each cluster is now $M^{3/2}$.
- Finally, the reverse of steps 1 and 2 takes place in the last two steps, as illustrated on Fig. 5.

By computations similar to those made in the two-cluster scheme, one can show that an aggregate throughput scaling of order

$$T(n) = \Omega\left(\mathrm{SNR}_s\, n^{3/4}\right)$$

is achieved in this case by taking $M = n^{1/4}$.

VII. Telescopic Beamforming

The idea of a telescopic beamforming strategy is coming from a further iteration of the four-cluster scheme. The scheme is illustrated on Fig. 4. In this scheme, multiple retransmissions of the source signal are performed via many clusters increasing in size (notice however that the number of clusters remains fixed compared to the number of nodes n in the network). In order for this scheme to work, the following two conditions should be satisfied:

1) The first ingredient is that for every transmission from one cluster to the next, condition (11) should be satisfied.

2) The second ingredient is that the overall noise amplification should be kept at its minimum level all the way, so that the end-to-end SNR remains of order 1.

Imposing these two conditions, the optimal cluster sizes can be computed via a MATLAB program for a given number of clusters on each side. From this, we deduce that the best throughput scaling achieved with this strategy can be made as large as

$$T(n) = \Omega\left(\mathrm{SNR}_s\, n^{6/7-\varepsilon}\right).$$

for every $\varepsilon > 0$.

References

[1] M. Franceschetti, M.D. Migliore, P. Minero, *The Capacity of Wireless Networks: Information-Theoretic and Physical Limits*, IEEE Trans. on Information Theory 55 (8), 2009, 3413–3424.

[2] P. Gupta, P. R. Kumar, *The Capacity of Wireless Networks*, IEEE Trans. on Information Theory 42 (2), 2000, 2313–2328.

[3] A. Jovicic, P. Viswanath and S. R. Kulkarni, *Upper Bounds to Transport Capacity of Wireless Networks*, IEEE Trans. on Information Theory 50 (11), 2004, 2555–2565.

[4] A. Merzakreeva, O. Lévêque, A. Özgür, *Hierarchical Beamforming for Large One-Dimensional Wireless Networks*, Proceedings of the IEEE International Symposium on Information Theory, Cambridge, MA, July 1-6, 2012, 1533–1537.

[5] U. Niesen, S. Diggavi, *The Approximate Capacity of the Gaussian N-Relay Diamond Network*, IEEE Trans. on Information Theory 59 (2), 2013, 845–859.

[6] U. Niesen, P. Gupta, D. Shah, *On Capacity Scaling in Arbitrary Wireless Networks*, IEEE Trans. on Information Theory 55 (9), 3959–3982, September 2009.

[7] A. Özgür, R. Johari, D. Tse, O. Lévêque, *Information-Theoretic Operating Regimes of Large Wireless Networks*, IEEE Trans. on Information Theory 56 (1), 2010, 427–437.

[8] A. Özgür, O. Lévêque, D. Tse, *Hierarchical Cooperation Achieves Optimal Capacity Scaling in Ad-Hoc Networks*, IEEE Trans. on Information Theory 53 (10), 2007, 3549–3572.

[9] A. Özgür, O. Lévêque, D. Tse, *Spatial Degrees of Freedom of Large Distributed MIMO Systems and Wireless Ad hoc Networks*, to appear in the IEEE Journal on Selected Areas in Communications, February 2013.

[10] L.-L. Xie and P. R. Kumar, *A Network Information Theory for Wireless Communications: Scaling Laws and Optimal Operation*, IEEE Trans. on Information Theory 50 (5), 2004, 748–767.