

Hierarchical Cooperation Achieves Linear Capacity Scaling in Ad Hoc Networks

Ayfer Özgür, Olivier Lévêque

Faculté Informatique et Communications
Ecole Polytechnique Fédérale de Lausanne
1015 Lausanne, Switzerland
{ayfer.ozgur,olivier.leveque}@epfl.ch

David Tse

Wireless Foundations
University of California at Berkeley
Berkeley, CA 94720, USA
dtse@eecs.berkeley.edu

Abstract— n source and destination pairs randomly located in a fixed area want to communicate with each other. It is well known that classical multihop architectures that decode and forward packets can deliver at most a \sqrt{n} -scaling of the aggregate throughput. The performance is limited by the mutual interference between communicating nodes. We show however that a *linear* scaling of the capacity with n can in fact be achieved by more intelligent node cooperation and distributed MIMO communication. The key ingredient is a *hierarchical and digital* architecture for nodal exchange of information for realizing the cooperation.

I. INTRODUCTION

The seminal paper by Gupta and Kumar [1] initiated the study of scaling laws in large ad-hoc wireless networks. Their by-now-familiar model considers n nodes randomly located in the unit disk, each of which wants to communicate to a random destination node at a rate $R(n)$ bits/second. They ask what is the maximally achievable scaling of the total throughput $T(n) = nR(n)$ with the system size n . They showed that classical multihop architectures with conventional single-user decoding and forwarding of packets cannot achieve a scaling of better than $O(\sqrt{n})$, and that a scheme that uses only nearest-neighbor communication can achieve a throughput that scales as $\Theta(\sqrt{n/\log n})$. This gap was later closed by Franceschetti et al [2], who showed using percolation theory that the $\Theta(\sqrt{n})$ scaling is indeed achievable.

Gupta-Kumar model makes certain assumptions on the physical-layer communication technology. In particular, it assumes that the signals received from nodes other than one particular transmitter are interference to be regarded as noise degrading the communication link. Given this assumption, long-range communication between nodes is not preferable, as the interference generated would preclude most of the other nodes from communicating. Instead, the optimal strategy is to confine to nearest neighbor communication and maximize the number of simultaneous transmissions (spatial reuse). However, this means that each packet has to be retransmitted many times before getting to the final destination, leading to a sub-linear scaling of system throughput. Thus, fundamentally, the Gupta-Kumar result is an *interference-limited* result.

A natural question is whether this result is a consequence of the physical-layer assumptions or whether one can do better using more sophisticated physical-layer processing. In a recent

work [3], Aeron and Saligrama have showed that the answer is the latter: they exhibited a scheme which yields a throughput scaling of $\Theta(n^{2/3})$ bits/second. However, it is not clear if one can do even better. The main result in this paper is that one can in fact achieve arbitrarily close to *linear* scaling: for any $\epsilon > 0$, we present a scheme that achieves an aggregate rate of $\Theta(n^{1-\epsilon})$. This is a surprising result: a linear scaling means the rate for *each* source-destination pair does not degrade significantly even as one puts more and more users in the network. It is easy to show, using an information theoretic argument, that one cannot get a better capacity scaling than $O(n \log n)$, so up to logarithmic terms, our scheme is optimal.

To achieve linear scaling, one must be able to perform *many* simultaneous long-range communications. A physical-layer technique which achieves this is MIMO (multi-input multi-output): the use of multiple transmit and receive antennas to multiplex several streams of data and transmit them simultaneously. MIMO was originally developed in the point-to-point setting, where the transmit antennas are co-located at a single transmit node, each transmitting one data stream, and the receive antennas are co-located at a single receive node, jointly processing the vector of received observations at the antennas. A natural approach to apply this concept to the network setting is to have both source nodes and destination nodes cooperate in *clusters* to form distributed transmit and receive antenna arrays respectively. In this way, mutually interfering signals can be turned into useful ones that can be jointly decoded at the receive cluster and spatial multiplexing gain can be realized. In fact, if *all* the nodes in the network could cooperate for free, then a classical MIMO result [4], [5] says that a sum rate scaling proportional to n could be achieved. However, this may be over-optimistic : communication between nodes is required to set up the cooperation and this may drastically reduce the useful throughput. The Aeron-Saligrama scheme is MIMO-based and its performance is precisely limited by the cooperation overhead between receive nodes. Our main contribution is a *multi-scale, hierarchical* cooperation architecture without significant overhead. Cooperation first takes place between nodes within very small local clusters to facilitate MIMO communication over a larger spatial scale. This can then be used as a communication infrastructure for cooperation within larger clusters at the next level of the hierarchy. Continuing on

this fashion, cooperation can be achieved at an almost global scale.

Since the publication of [1], there have been several works dealing with information theoretic scaling laws of wireless adhoc networks [6], [7], [8], [9], [10]. All of them deal with *extended* networks, which scale to cover an increasing geographical extent with the density of nodes fixed and the source-destination distances increasing large. The performance in such a regime is, however, *power-limited* rather than *interference-limited*: the role of relays is to ensure that power can be transferred over large distances rather than to minimize the mutual interference between communicating nodes. Nevertheless, it has recently been showed that the hierarchical scheme here can be adapted to achieve optimal power transfer and optimal capacity scaling in extended networks as well [11].

The dense scaling is relevant whenever one wants to design networks to serve many nodes, all within communication range of each other (within a campus, an urban block, etc.). This scaling is also a reasonable model to study problems such as *spectrum sharing*, where many users in a geographical area are sharing a wide band of spectrum. Consider the scenario where we segregate the total bandwidth into many orthogonal bands, one for each separate network supporting a *fixed* number of users. As we increase the number of users, the number of such segregated networks increases but the *spectral efficiency*, in bits/s/Hz, does not scale with the *total* number of users. In contrast, if we build one large ad hoc network for all the users on the entire bandwidth, then our result says that the spectral efficiency actually increases *linearly* with the number of users. The gain is coming from a *network* effect via cooperation between the many nodes in the system.

The rest of the paper is summarized as follows. In Section II, we present the model. Section III contains the main result and an outline of the proposed architecture together with a back-of-the-envelope analysis of its performance. The details of its performance analysis are given in Section IV. Section V contains our conclusions.

II. MODEL

There are n nodes uniformly and independently distributed in a square of unit area. Every node is both a source and a destination. The sources and destinations are paired up one-to-one in an arbitrary way. Each source has the same traffic rate $R(n)$ to send to its destination node and a common average transmit power budget of P Watts. The total throughput of the system is $T(n) = nR(n)$.¹

We assume that communication takes place over a flat channel of bandwidth W Hz around a carrier frequency of f_c , $f_c \gg W$. The complex baseband-equivalent channel gain between node i and node k at time m is given by:

$$h_{ik}[m] = \sqrt{G} r_{ik}^{-\alpha/2} \exp(j\theta_{ik}[m]) \quad (1)$$

¹In the sequel, whenever we say a total throughput $T(n)$ is achievable, we implicitly mean that that a rate of $T(n)/n$ is achievable for every source-destination pair.

where r_{ik} is the distance between the nodes, $\theta_{ik}[m]$ is the random phase at time m , uniformly distributed in $[0, 2\pi)$ and $\{\theta_{ik}[m]\}$ are i.i.d. random processes across all i and k . The $\theta_{ik}[m]$'s and the r_{ik} 's are also assumed to be independent. The parameters G and $\alpha \geq 2$ are assumed to be constants; α is called the path loss exponent. For example, under free-space line-of-sight propagation, Friis' formula applies and

$$|h_{ik}[m]|^2 = \frac{G_{Tx} \cdot G_{Rx}}{(4\pi r_{ik}/\lambda_c)^2} \quad (2)$$

so that

$$G = \frac{G_{Tx} \cdot G_{Rx} \cdot \lambda_c^2}{16\pi^2}, \quad \alpha = 2.$$

where G_{Tx} and G_{Rx} are the transmitter and receiver antenna gains respectively and λ_c is the carrier wavelength.

Note that the channel is random, depending on the location of the users and the phases. The locations are assumed to be fixed over the duration of the communication. The phases are assumed to vary in a stationary ergodic manner (fast fading).² We assume that the channel gains are known at all the nodes. The received signal is a sum of the received signals plus white circular symmetric Gaussian noise of variance N_0 per symbol.

Several comments about the model are in order:

- The path loss model is based on a *far-field* assumption: the distance r_{ik} is assumed to be much larger than the carrier wavelength. When the distance is of the order or shorter than the carrier wavelength, the simple path loss model obviously does not hold anymore as path loss can potentially become path "gain". The reason is that near-field electromagnetics now come into play.
- The phase $\theta_{ik}[m]$ depends on the distance between the nodes modulo the carrier wavelength [12]. The random phase model is thus also based on a far-field assumption: we are assuming that the nodes separation is at a much larger spatial scale compared to the carrier wavelength, so that the phases can be modelled as completely random and independent of the actual positions.
- It is realistic to assume the variation of the phases since they vary significantly when users move a distance of the order of the carrier wavelength (fractions of a meter). The positions determine the path losses and they on the other hand vary over a much larger spatial scale. So the positions are assumed to be fixed.
- We essentially assume a line-of-sight type environment and ignore multipath effects. The randomness in phases is sufficient for the long range MIMO transmissions needed in our scheme. With multipaths, there is a further randomness due to random constructive and destructive interference of these paths. It can be seen that our result easily extends to the multipath case.

Theoretically, as the number of nodes increases, the far-field assumption eventually becomes invalid as nodes become closer. In reality, the typical separation between nodes is so

²With more technical efforts, we believe our results can be extended to the slow fading setting where the phases are fixed as well.

much larger than the carrier wavelength that the number of nodes when the far-field assumption fails is humongous, i.e. there is a clear separation between the large and the small spatial scales. Consider the following numerical example. Suppose the area of interest is 1 sq. km, well within the communication range of many radio devices. With a carrier frequency of 3 GHz, the carrier wavelength is 0.1m. Even with a very large system size of $n = 10000$ nodes, the typical separation between nearest neighbors is 10 m, very much in the far-field. Under free-space propagation and assuming unit transmit and receive antenna gains, the attenuation given by Friis' formula (2) is about 10^{-6} , much smaller than unity. To have a nearest-neighbor distance of $0.1m$ (the carrier wavelength), 10^8 nodes would be needed in the area! Hence, there is a wide range of system parameters for which simultaneously the number of nodes is large and the far-field assumption holds.

In the following discussions, we will simplify the notation by suppressing the dependency of the channel gains on the time index m .

III. MAIN RESULT

We first give an information-theoretic upper bound on the achievable scaling law for the aggregate throughput in the network. Before starting to look for good communication strategies, Theorem 3.1 establishes the best we can hope for.

Theorem 3.1: The aggregate throughput in a network with n nodes is bounded above by

$$T(n) \leq K'n \log n$$

with high probability³ for some constant $K' > 0$ independent of n .

Proof: Consider a source-destination pair (s, d) in the network. The transmission rate $R(n)$ from source node s to destination node d is upperbounded by the capacity of the single-input multiple-output (SIMO) channel between source node s and the rest of the network. Using a standard formula for this channel (see eg. [12]), we get:

$$R(n) \leq \log \left(1 + \frac{P}{N_0} \sum_{\substack{i=1 \\ i \neq s}}^n |h_{is}|^2 \right) = \log \left(1 + \frac{P}{N_0} \sum_{\substack{i=1 \\ i \neq s}}^n \frac{G}{r_{is}^\alpha} \right).$$

It is a well known fact that in a random network with n nodes uniformly distributed on a fixed two-dimensional area, the minimum distance between any two nodes in the network is larger than $\frac{1}{n^{1+\delta}}$ with high probability, for any $\delta > 0$. Using this fact, we obtain

$$R(n) \leq \log \left(1 + \frac{GP}{N_0} n^{\alpha(1+\delta)+1} \right) \leq K' \log n$$

for some constant $K' > 0$ independent of n , for all source-destination pairs in the network with high probability. The theorem follows. \square

³i.e. probability going to 1 as system size grows.

In the view of what is ultimately possible, established by Theorem 3.1, we are now ready to state the main result of this paper.

Theorem 3.2: Let $\alpha \geq 2$. For any $\epsilon > 0$, with high probability an aggregate throughput

$$T(n) \geq Kn^{1-\epsilon}$$

is achievable in the network for all possible pairings between sources and destinations. $K > 0$ is a constant independent of n and the source-destination pairing.

Theorem 3.2 states that it is actually possible to perform arbitrarily close to the bound given in Theorem 3.1. The two theorems together establish the capacity scaling for the network up to logarithmic terms. Note how dramatically different is this new linear capacity scaling law from the well-known throughput scaling of $\Theta(\sqrt{n})$ implied by [1], [2] for the same model. Note also that the upper bound in Theorem 3.1 assumes a genie-aided removal of interference between simultaneous transmissions from different sources. By proving Theorem 3.2, we will show that it is possible to mitigate such interference without a genie but with cooperation between the nodes.

The proof of Theorem 3.2 relies on the construction of an explicit scheme that realizes the promised scaling law. The construction is based on recursively using the following key lemma, which addresses the case when $\alpha > 2$.

Lemma 3.1: Consider $\alpha > 2$ and a network with n nodes subjected to interference from external sources. Let the interference signals received by different nodes in the network be uncorrelated and the interference power received by each node be upperbounded by

$$P_I \leq K_I$$

for some $K_I > 0$ independent of n . Let us assume there exists a scheme such that for each n , with probability at least $1 - e^{-n^{c_1}}$, it achieves an aggregate throughput

$$T(n) \geq K_1 n^b$$

for every possible source-destination pairing in a network of n nodes. K_1 and c_1 are positive constants independent of n and the source-destination pairing, and $0 \leq b < 1$. Let us also assume that the per node average power budget required to realize this scheme is:

$$P \leq \frac{K_p}{n} \tag{3}$$

for some $K_p > 0$ independent of n .

Then one can construct another scheme that achieves a *higher* aggregate throughput

$$T(n) \geq K_2 n^{\frac{1}{2-b}}$$

for every source-destination pairing in a network of n nodes under the same interference conditions, where $K_2 > 0$ is another constant independent of n and the pairing. Moreover, the failure rate for the new scheme is upper bounded by $e^{-n^{c_2}}$ for another positive constant c_2 , while the per node average

power needed to realize the scheme is also bounded above by (3).

Lemma 3.1 is the key step to build a hierarchical architecture. Since $\frac{1}{2-b} > b$ for $0 \leq b < 1$, the new scheme is always better than the old one. We will now give a rough description of how the new scheme can be constructed given the old scheme, as well as a back-of-the-envelope analysis of the scaling law it achieves. Next section is devoted to its precise description and performance analysis.

The constructed scheme is based on clustering and long-range MIMO transmissions between clusters. We divide the network into clusters of M nodes. Let us focus for now on a particular source node s and its destination node d . s will send M bits to d in 3 steps:

- 1) Node s will distribute its M bits among the M nodes in its cluster, one for each node;
- 2) These nodes together can then form a distributed transmit antenna array, sending the M bits *simultaneously* to the destination cluster where d lies;
- 3) Each node in the destination cluster obtained one observation from the MIMO transmission, and it quantizes and ships the observation back to d , which can then do joint MIMO processing of all the observations and decode the M transmitted bits.

From the network point of view, all source-destination pairs have to eventually accomplish these three steps. Step 2 is long-range communication and only one source-destination pair can operate at the same time. Steps 1 and 3 involve local communication and can be parallelized across source-destination pairs. Combining all this leads to three phases in the operation of the network:

Phase 1: Setting Up Transmit Cooperation Clusters work in parallel. Within a cluster, each source node has to distribute M bits to the other nodes, 1 bit for each node, such that at the end of the phase, each node has 1 bit from each of the source nodes in the same cluster. Since there are M source nodes in each cluster, this gives a traffic demand of exchanging M^2 bits. The key observation is that this is similar to the original problem of communicating between n source and destination pairs, but on a network of size M . More specifically, this traffic demand of exchanging M^2 bits is handled by setting up M sub-phases, and assigning M source-destination pairs for each sub-phase. Since our channel model is scale invariant, note that the scheme given in the hypothesis of the lemma can be used in each sub-phase by simply scaling down the power with cluster area. Having aggregate throughput M^b , each sub-phase is completed in M^{1-b} time slots, while the whole phase takes M^{2-b} time slots. See Figure 1.

Phase 2: MIMO Transmissions We perform successive long-distance MIMO transmissions between source-destination pairs, one at a time. In each one of the MIMO transmissions, say the one between s and d , the M bits of s are simultaneously transmitted by the M nodes in its cluster to the M nodes in the cluster of d . Each of the long-distance MIMO transmissions are repeated for each source node in the

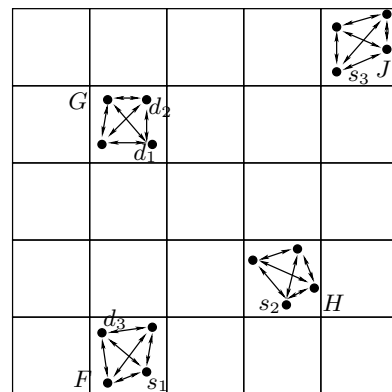


Fig. 1. Nodes inside clusters F , G , H and J are illustrated while exchanging bits in Phases 1 and 3. Note that in Phase 1 the exchanged bits are the source bits whereas in Phase 3 they are the quantized MIMO observations. Clusters work in parallel. In this and the following figure Fig. 2, we highlight three source-destination pairs $s_1 - d_1$, $s_2 - d_2$ and $s_3 - d_3$, such that nodes s_1 and d_3 are located in F , nodes s_2 and s_3 are located in H and J respectively, and nodes d_1 and d_2 are located in G .

network, hence we need n time slots to complete the phase. See Figure 2.

Phase 3: Cooperate to Decode Clusters work in parallel. Since there are M destination nodes inside the clusters, each cluster received M MIMO transmissions in phase 2, one intended for each of the destination nodes in the cluster. Thus, each node in the cluster has M received observations, one from each of the MIMO transmissions, and each observation is to be conveyed to a different destination node in its cluster. Nodes quantize each observation into fixed Q bits, so there are now a total of QM^2 bits to exchange inside each cluster. Using exactly the same scheme as in Phase 1, we conclude the phase in QM^{2-b} time slots. See Figure 1.

Assuming that each destination node is able to decode the transmitted bits from its source node from the M quantized signals it gathers by the end of Phase 3, we can calculate the rate of the scheme as follows: each source node is able to transmit M bits to its destination node, hence nM bits in total are delivered to their destinations in $M^{2-b} + n + QM^{2-b}$ time slots, yielding an aggregate throughput of

$$\frac{nM}{M^{2-b} + n + QM^{2-b}}$$

bits per time slot. Maximizing this throughput by choosing $M = n^{\frac{1}{2-b}}$ yields $T(n) = \frac{1}{2+Q} n^{\frac{1}{2-b}}$ for the aggregate throughput, which is the result in Lemma 3.1.

Clusters can work in parallel in phases 1 and 3 because for $\alpha > 2$, the aggregate interference at a particular cluster caused by other active nodes is bounded. For $\alpha = 2$ the aggregate interference scales like $\log n$, leading to a slightly different version of the lemma.

Lemma 3.2: Consider $\alpha = 2$ and a network with n nodes subjected to interference from external sources. Let the interference signals received by different nodes in the network be uncorrelated and the interference power received by each node




Fig. 2. Successive MIMO transmissions are performed between clusters. The first figure depicts MIMO transmission from cluster F to G , where bits originally belonging to s_1 are simultaneously transmitted by all nodes in F to all nodes in G . The second MIMO transmission is from H to G , while now bits of source node s_2 are transmitted from nodes in H to nodes in G . The third picture illustrates MIMO transmission from cluster J to F corresponding to the source destination pair $s_3 - d_3$.

be upperbounded by

$$P_I \leq K_I \log n$$

for some $K_I \geq 0$ independent of n . Let us assume there exists a scheme which for each n , with failure probability at most $e^{-n^{c_1}}$, achieves an aggregate throughput

$$T(n) \geq K_1 \frac{n^b}{\log n}$$

for every source-destination pairing in a network with n nodes. K_1 and c_1 are positive constants independent of n and the source-destination pairing, and $0 \leq b < 1$. Let us also assume that the per node average power budget required to realize this scheme is:

$$P \leq \frac{K_p}{n} \quad (4)$$

for some $K_p > 0$ independent of n .

Then one can construct another scheme that achieves a *higher* aggregate throughput

$$T(n) \geq K_2 \frac{n^{\frac{1}{2-b}}}{(\log n)^2}$$

for every source-destination pairing in a network of n nodes under the same interference conditions, where $K_2 > 0$ is another constant independent of n and the pairing. Moreover, the failure rate for the new scheme is upper bounded by $e^{-n^{c_2}}$ for another positive constant c_2 , while the per node average power needed to realize the scheme is also bounded above by (4).

We can now use Lemma 3.1 and 3.2 to prove Theorem 3.2.

Proof of Theorem 3.2: We only focus on the case $\alpha > 2$. The case $\alpha = 2$ proceeds similarly, differing only with a reduction of a factor of $\log n$ in the throughputs.

We start by observing that the simple scheme of transmitting directly between the source-destination pairs one at a time (TDMA) satisfies the requirements of the lemma. The aggregate throughput is $\Theta(1)$, so $b = 0$. The failure probability is 0. Since each source is only transmitting $\frac{1}{n}$ th of the time and the distance between the source and its destination is bounded, the average power consumed per node is of the order of $\frac{1}{n}$.

As soon as we have a scheme to start with, Lemma 3.1 can be applied recursively, yielding a scheme that achieves higher throughput at each step of the recursion. More precisely, starting with a TDMA scheme with $b = 0$ and applying Lemma 3.1 recursively h times, one gets a scheme achieving $O(n^{\frac{h}{h+1}})$ aggregate throughput. Given any $\epsilon > 0$, we can now choose h such that $\frac{h}{h+1} \geq 1 - \epsilon$ and we get a scheme that achieves $O(n^{1-\epsilon})$ aggregate throughput scaling with high probability. This concludes the proof of Theorem 3.2. \square

Putting everything together, we have built a hierarchical scheme to achieve the desired throughput. At the lowest level of the hierarchy, we use the simple TDMA scheme to exchange bits for cooperation among small clusters. Combining this with longer range MIMO transmissions, we get a higher throughput scheme for cooperation among nodes in larger clusters at the next level of the hierarchy. Finally, at the top level of the hierarchy, the cooperation clusters are almost the size of the network and the MIMO transmissions are over the global scale to meet the desired traffic demands. Figure 3 shows the resulting hierarchical scheme with a focus on the top two levels.

IV. DETAILED DESCRIPTION AND PERFORMANCE ANALYSIS

In this section, we concentrate in more detail on the scheme that proves Lemma 3.1 and Lemma 3.2. We first focus on Lemma 3.1 and then extend the proof to Lemma 3.2. As we have already seen in the previous section, we start by dividing the unit square into smaller squares of area $A_c = \frac{M}{n}$. Since the node density is n , there will be on average M nodes inside each of these small squares. The following lemma upper bounds the probability of having large deviations from the average. The proof, as well as those of following lemmas, is omitted due to space constraints.

Lemma 4.1: Let us partition a unit area network of size n into cells of area A_c . The number of nodes inside each cell is between $((1 - \delta)A_c n, (1 + \delta)A_c n)$ with probability larger than $1 - \frac{1}{A_c} e^{-\Lambda(\delta)A_c n}$ where $\Lambda(\delta) > 0$ for $\delta > 0$.

Applying Lemma 4.1 to the squares of area M/n , we see that all squares contain order M nodes with high probability.




Fig. 3. The time division in a hierarchical scheme as well as the salient features of the three phases are illustrated.

In the following discussion, we will need a stronger result, namely each of the 8 possible halves of a square should contain order $M/2$ nodes which again follows from the lemma together with union bound. This condition is sufficient for our below analysis on scaling laws to hold. However, in order to simplify the presentation, we assume that there are exactly $M/2$ nodes inside each half, thus exactly M nodes in each square. The clustering is used to realize a distributed MIMO system in three successive steps:

Phase 1: Setting Up Transmit Cooperation In this phase, source nodes distribute their data streams over their clusters and set up the stage for the long-range MIMO transmissions that we want to perform in the next phase. Clusters work in parallel according to the 9-TDMA scheme depicted in Figure 4, which divides the total time for this phase into 9 time-slots and assigns simultaneous operation to clusters that are sufficiently separated.

Let us focus on one specific source node s located in cluster S with destination node d in cluster D . Node s will divide a block of length LM bits of its data stream into M sub-blocks, each of length L bits, where L can be arbitrarily large but bounded. The destination of each sub-block in Phase 1 depends on the relative position of clusters S and D :

- (1) If S and D are either the same cluster or are not neighboring clusters: one sub-block is to be kept in s and the rest $M - 1$ sub-blocks are to be transmitted to the other $M - 1$ nodes located in S , one sub-block for each node.
- (2) If S and D are neighboring clusters: divide the cluster S into two halves, each of area $A_c/2$, one half located close to the border with D and the second half located farther to D . The M sub-blocks of source node s are to be distributed to the $M/2$ nodes located in the second half cluster (farther to D), each node gets two sub-blocks.

Since the above traffic is required for every source node in cluster S , we end up with a highly uniform traffic demand of delivering $M \times LM$ bits in total to their destinations. A key observation is that the problem can be separated into sub-problems, each similar to our original problem, but on a network size M and area A_c . More specifically, the traffic of transporting LM^2 bits can be handled by organizing M sessions and assigning M source-destination pairs for each session. (Note that due to the non-uniformity arising from point (2) above, one might be able to assign only $M/2$ source-destination pairs in a session and hence need to handle the traffic demand of transporting LM^2 bits by organizing up to $2M$ sessions in the extreme case instead of M .) The assigned source-destination pairs in each session can then communicate L bits. Since our channel model is scale invariant, the scheme in the hypothesis of Lemma 3.1 can be used to handle the traffic in each session, by simply scaling down the powers of the nodes by $(A_c)^{\alpha/2}$. Hence, the power used by each node will be bounded by $\frac{K_p(A_c)^{\alpha/2}}{M}$. The scheme is to be operated simultaneously inside all the clusters in the 9-TDMA scheme, so we need to ensure that the resultant inter-cluster interference satisfies the properties in Lemma 3.1.

Lemma 4.2: Consider clusters of size M and area A_c operating according to 9-TDMA scheme in Figure 4 in a network of size n . Let each node be constrained to an average power $\frac{K_p(A_c)^{\alpha/2}}{M}$. For $\alpha > 2$, the interference power received by a node from the simultaneously operating clusters is upperbounded by a constant K_{I_1} independent of n . For $\alpha = 2$, the interference power is bounded by $K_{I_2} \log n$ for K_{I_2} independent of n . Moreover, the interference signals received by different nodes in the cluster are uncorrelated.

Let us for now concentrate on the case $\alpha > 2$. By Lemma 4.2, the inter-cluster interference is bounded and is uncorrelated at different nodes. Thus, the strategy in the

hypothesis of Lemma 3.1 can achieve an aggregate rate $K_1 M^b$ in all the sessions for some $K_1 > 0$, with probability larger than $1 - e^{-M^{c_1}}$. Using the union bound, with probability larger than $1 - 2ne^{-M^{c_1}}$, the aggregate rate $K_1 M^b$ is achieved inside all sessions ($2M$ sessions in the extreme case) in all clusters in the network. With this aggregate rate, each session can be completed in at most $(L/K_1)M^{1-b}$ channel uses and $2M$ successive sessions are completed in $(2L/K_1)M^{2-b}$ channel uses. Using the 9-TDMA scheme, the phase is completed in less than $(18L/K_1)M^{2-b}$ channel uses all over the network with probability larger than $1 - 2ne^{-M^{c_1}}$.

Phase 2: MIMO Transmissions In this phase, we are performing the actual MIMO transmissions for all the source-destination pairs serially, i.e. one at a time. A MIMO transmission from source s to destination d involves the M (or $M/2$) nodes in the cluster S , where s is in (referred to as the source cluster for this MIMO transmission) to the M (or $M/2$) nodes of the cluster D , where d is in (referred to as the destination cluster of the MIMO transmission).

Let the distance between the mid-points of the two clusters be r_{SD} . If S and D are the same cluster, we skip the step for this source node s . Otherwise, we operate in two slightly different modes depending on the relative positions of S and D corresponding to the operations performed in the first phase: First consider the case where S and D are not neighboring clusters. In this case, the M nodes in cluster S independently encode the L bits-long sub-blocks they possess, originally belonging to node s , into C symbols by using a randomly generated Gaussian code \mathcal{C} that respects an average transmit power constraint $\frac{K_p(r_{SD})^\alpha}{M}$. The nodes then transmit their encoded sequences of length C symbols simultaneously to the M nodes in cluster D . The nodes in cluster D properly sample the signals they observe during the C transmissions and store these samples (that we will simply refer to as *observations* in the following text), without trying to decode the transmitted symbols. In the case where S and D are neighbors, the strategy is slightly modified so that the MIMO transmission is from the $M/2$ nodes in S , that possess the sub-blocks of s after Phase 1, to the $M/2$ nodes in D that are located in the farther half of the cluster to S . Each of these $M/2$ nodes in S possess two sub-blocks that come from s . They encode each sub-block into C symbols by again using a Gaussian code of power $\frac{K_p(r_{SD})^\alpha}{M}$. The nodes then transmit the $2C$ symbols to the $M/2$ nodes in D that in turn sample their received signals and store the observations. The observations accumulated at various nodes in D at the end of this step are to be conveyed to node d during the third phase.

After concluding the step for source node s of S , the phase continues by repeating the same step for source node $s + 1$ of S . The destination cluster for the new MIMO transmission will be, in general, a different cluster D' , which is the one that contains the ultimate destination node d' for the source node $s + 1$. The MIMO transmissions are repeated until the data originated from all source nodes in the network are transmitted to their respective destination clusters. Since the step for one source node takes either C or $2C$ channel uses, completing

the operation for all n source nodes in the network requires at most $2C \times n = 2Cn$ channel uses.

Phase 3: Cooperate to Decode In this phase, we aim to provide each destination node, the observations of the symbols that have been originally intended for it. With the MIMO transmissions in the second phase, these observations have been accumulated at the nodes of its cluster. As before, let us focus on a specific destination node d located in cluster D . Note that depending on whether the source node of d is located in a neighboring cluster or not, either each of the M nodes in D have C observations intended for d , or $M/2$ of the nodes have $2C$ observations each. Note that these observations are some real numbers that need to be quantized and encoded into bits before being transmitted. Let us assume that we are encoding each block of C observations into CQ bits, by using fixed Q bits per observation on the average. The situation is symmetric for all M destination nodes in D , since the cluster received M MIMO transmissions in the previous phase, one for each destination node. (The destination nodes that have source nodes in D are exception. Recall from Phase 1 and Phase 2 that in this case, each node in D possesses sub-blocks of the original data stream for the destination node, not MIMO observations. We will ignore this case by simply assuming $L \leq CQ$ in the below computation.) The arising traffic demand of transporting $M \times CQM$ bits in total is similar to Phase 1 and can be handled by using exactly the same scheme in less than $(2CQ/K_1)M^{2-b}$ channel uses. Recalling the discussion on the first phase, we conclude that the phase can be completed in less than $(18CQ/K_1)M^{2-b}$ channel uses all over the network with probability larger than $1 - 2ne^{-M^{c_1}}$.

Note that if it were possible to encode each observation into fixed Q bits without introducing any distortion, which is obviously not the case, the following lemma on MIMO capacity would suggest that with the Gaussian code \mathcal{C} used in Phase 2 satisfying $L/C \geq \kappa$ for some constant $\kappa > 0$, the transmitted bits could be recovered by an arbitrarily small probability of error from the observations gathered by the destination nodes at the end of Phase 3.

Lemma 4.3: The mutual information achieved by the $M \times M$ MIMO transmission between any two clusters grows at least linearly with M .

The following lemma states that there is actually a way to encode the observations using fixed number of bits per observation and at the same time, not to degrade the performance of the overall channel significantly, that is, to still get a linear capacity growth for the resulting *quantized MIMO* channel.

Lemma 4.4: There exists a strategy to encode the observations at a fixed rate Q bits per observation and get a linear growth of the mutual information for the resultant $M \times M$ quantized MIMO channel.

Although we do not discuss the proof of the above lemma, the following small lemma may provide motivation for the stated result. Lemma 4.5 points out a key observation on the

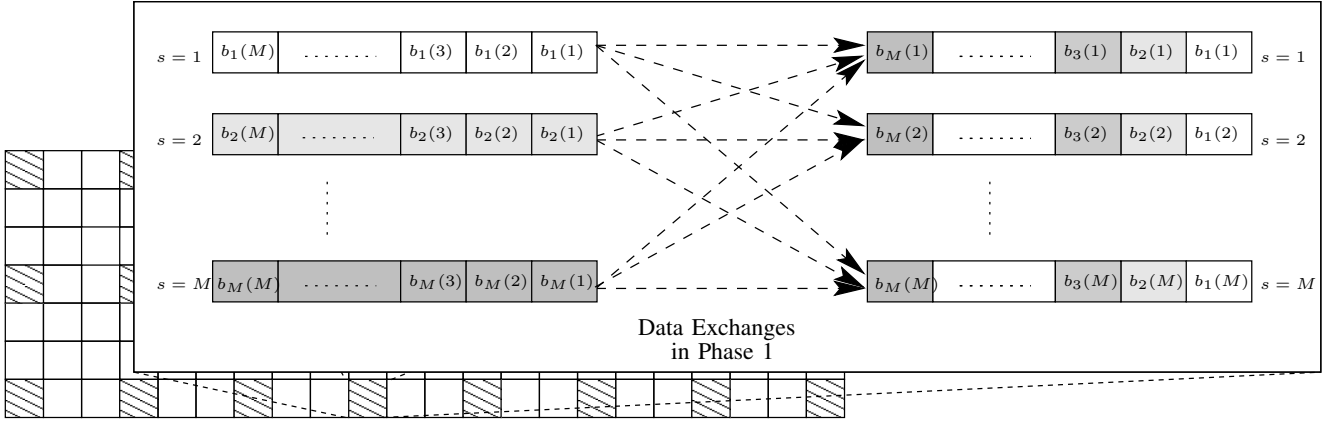


Fig. 4. Buffers of the nodes in a cluster are illustrated before and after the data exchanges in Phase 1. The data stream of the source nodes are distributed to the M nodes in the network as depicted. $b_s(j)$ denotes the j 'th sub-block of the source node s . Note the 9-TDMA scheme that is employed over the network in this phase.

way we choose our transmit powers in the MIMO phase. It is central to the proof of Lemma 4.4 and states that the observations have bounded power, that does not scale with M . This in turn suggests that one can use a fixed number of bits to encode them without degrading the scaling performance of the scheme.

Lemma 4.5: In Phase 2, the power received by each node in the destination cluster is bounded below and above by constants P_1 and P_2 respectively that are independent of M .

Putting it together, we have seen that the three phases described effectively realize virtual MIMO channels achieving spatial multiplexing gain M between the source and destination nodes in the network. Using these virtual MIMO channels, each source is able to transmit ML bits in

$$\begin{aligned} T_t &= T(\text{phase 1}) + T(\text{phase 2}) + T(\text{phase 3}) \\ &= \frac{18L}{K_1}M^{2-b} + 2Cn + \frac{18CQ}{K_1}M^{2-b} \end{aligned}$$

total channel uses where $L/C \geq \kappa$ for some $\kappa > 0$ independent of M (or n). This gives an aggregate throughput of

$$\begin{aligned} T(n) &= \frac{nML}{(18L/K_1)M^{2-b} + 2Cn + (18CQ/K_1)M^{2-b}} \\ &\geq K_2 n^{\frac{1}{2-b}} \end{aligned} \quad (5)$$

for some $K_2 > 0$ independent of n , by choosing $M = n^{\frac{1}{2-b}}$ with $0 \leq b < 1$, which is the optimal choice for the cluster size as a function of b . A failure arises if there are not order $M/2$ nodes in each half cluster or the scheme used in Phases 1 and 3 fails to achieve the promised throughput. Combining the result of Lemma 4.1 with the computed failure probabilities for Phases 1 and 3 yields

$$P_f \leq 4ne^{-M^{c_1}} + \frac{8n}{M}e^{-\Lambda(\delta)M/2} \leq e^{-n^{c_2}}$$

for some $c_2 > 0$.

Next, we show that the new scheme also satisfies the power constraint in (3): for Phases 1 and 3, we know that the scheme employed inside the clusters satisfies $P_j \leq K_p A_c^{\alpha/2}/M$. From this inequality, we deduce that

$$P_j \leq \frac{K_p}{n}. \quad (6)$$

Indeed, $A_c = M/n$, and for $\alpha \geq 2$ we have

$$P_j \leq \frac{K_p}{M} \left(\frac{M}{n}\right)^{\alpha/2} = \frac{K_p}{n} \left(\frac{M}{n}\right)^{\alpha/2-1} \leq \frac{K_p}{n}.$$

In Phase 2, each node is transmitting with power $\frac{K_p(r_{SD})^\alpha}{M}$ in at most fraction M/n of the total duration of the phase, while keeping silent during the rest of the time. This yields a per node average power $\frac{K_p(r_{SD})^\alpha}{n}$. Recall that r_{SD} is the distance between the mid-points of the source and destination clusters and $r_{SD} < 1$, which yields (6) also for the second phase.

In order to conclude the proof of Lemma 3.1, note that the new scheme achieves the same aggregate throughput scaling in the presence of external interference. In phases 1 and 3, this external interference with bounded power will simply add to the inter-cluster interference experienced by the nodes. For the MIMO phase, this will result in uncorrelated background-noise-plus-interference at the receiving nodes which is not necessarily Gaussian. However, it is well known that the achievable mutual information is lower bounded by assuming that the interference-plus-noise is i.i.d. Gaussian and Lemma 4.3 applies. This concludes the proof of Lemma 3.1. \square

Proof of Lemma 3.2: The scheme that proves Lemma 3.2 is completely similar to the one described above. Lemma 4.2 states that when $\alpha = 2$, the inter-cluster interference power experienced during Phases 1 and 3 is upperbounded by $K_{I_2} \log n = K_{I_2} \log M$. There is furthermore the external interference with power bounded by $K_I \log n$ that is adding to the inter-cluster interference. Under these conditions, the scheme in the hypothesis of Lemma 3.2 achieves an aggregate rate $K_1 \frac{M^b}{\log M}$ when used to handle the traffic in these phases.

For the second phase we have the following lemma which provides a lower bound on the spatial multiplexing gain of the quantized MIMO channel under the interference experienced.

Lemma 4.6: Let the MIMO signal received by the nodes in the destination cluster be corrupted by an interference of power $K_I \log M$, uncorrelated over different nodes and independent of the transmitted signals. There exists a strategy to encode these corrupted observations at a fixed rate Q bits per observation and get a $M/\log M$ growth of the mutual information for the resulting $M \times M$ quantized MIMO channel.

A capacity of $M/\log M$ for the resulting MIMO channel implies that there exists a code \mathcal{C} that encodes L bits-long sub-blocks into $C \log M$ symbols, where $L/C \geq \kappa'$ for a constant $\kappa' > 0$, so that the transmitted bits can be decoded at the destination nodes with arbitrarily small probability of error for L and C sufficiently large. Hence, starting again with a block of LM bits in each source node, the LM^2 bits in the first phase can be delivered in $(L/K_1)M^{2-b} \log M$ channel uses. In the second phase, the L bits-long sub-blocks now need to be encoded into $C \log M$ symbols, hence the transmission for each source-destination pair takes $C \log M$ channel uses, the whole phase taking $Cn \log M$ channel uses. Note that there are now $CM^2 \log M$ observations encoded into $CQM^2 \log M$ bits that need to be transported in the third phase. With the scheme of aggregate rate $K_1 \frac{M^b}{\log M}$, we need $(CQ/K_1)M^{2-b}(\log M)^2$ channel uses to complete the phase. Choosing $M = n^{\frac{1}{2-b}}$, gives an aggregate throughput of $K_2 n^{\frac{1}{2-b}} / (\log n)^2$ for the new scheme. This concludes the proof of Lemma 3.2. \square

V. CONCLUSIONS

In this paper, we have shown that the capacity of ad hoc wireless networks with n nodes in a fixed area actually scales *linearly* with n . This is a surprising result, as it suggests that interference is not a fundamental limitation in such a setting, provided that there is enough cooperation between nodes. It was known that the capacity of *mobile* ad hoc networks scale linearly with n [13], but the capacity-achieving scheme heavily exploits the mobility of the nodes and thus the delay in traffic delivery is of the order of the mobility time-scale. Here, we show that the same linear scaling can be achieved with *fixed* nodes as well, using more sophisticated physical layer techniques than just multihop. The key ideas behind this scheme are:

- using MIMO for long-range communication to achieve spatial multiplexing;
- local transmit and receive cooperation to maximize spatial reuse;
- setting up the intra-cluster cooperation such that it is yet another digital communication problem, but in a smaller network, thus enabling a hierarchical cooperation architecture.

Our result is based on only very weak assumptions about the channel. It is valid as long as long-range communication is possible and for any path loss exponent $\alpha \geq 2$. It holds regardless of whether there are multipaths, as long as nodal separation is much larger than the carrier wavelength, so that the phases of the channels are random. This is sufficient to enable MIMO. We have focused on the 2-D setting, where the nodes are on the plane, but our results generalize naturally to d -dimensional networks.

ACKNOWLEDGMENT

David Reed raised the question of whether a linear capacity scaling is possible, and that provides part of the impetus for this research. The authors would also like to thank to Emre Telatar, Shuchin Aeron and Venkatesh Saligrama for many helpful discussions. The work of Ayfer Özgür was supported by Swiss NSF grant Nr 200021-108089. Part of the work of Olivier Lévêque was performed when he was with the Electrical Engineering Department at Stanford University, supported by Swiss NSF grant Nr PA002-108976. The work of David Tse was supported by the U.S. National Science Foundation via an ITR grant: "The 3R's of Spectrum Management: Reuse, Reduce and Recycle".

REFERENCES

- [1] P. Gupta and P. R. Kumar, *The Capacity of Wireless Networks*, IEEE Transactions on Information Theory 42 (2), 2000, 388–404.
- [2] M. Franceschetti, O. Dousse, D. Tse and P. Thiran, *On the Throughput Capacity of Random Wireless Networks*, Preprint, 2006.
- [3] S. Aeron and V. Saligrama, *Wireless Ad hoc Networks: Strategies and Scaling Laws for the Fixed SNR Regime*, Preprint, 2006.
- [4] G. J. Foschini, *Layered Space-Time Architecture For Wireless Communication in a Fading Environment when Using Multi-Element Antennas*, AT&T Bell Labs Technical Journal 1 (2), October 1996, 41–59.
- [5] E. Telatar, *Capacity of Multi-Antenna Gaussian Channels*, European Transactions on Telecommunications 10 (6), November 1999, 585–596.
- [6] L.-L. Xie and P. R. Kumar, *A Network Information Theory for Wireless Communications: Scaling Laws and Optimal Operation*, IEEE Transactions on Information Theory 50 (5), May 2004, 748–767.
- [7] A. Jovicic, S. R. Kulkarni and P. Viswanath, *Upper Bounds to Transport Capacity of Wireless Networks*, IEEE Transactions on Information Theory 50 (11), November 2004, 2555–2565.
- [8] O. Lévêque and E. Telatar, *Information Theoretic Upper Bounds on the Capacity of Large Extended Ad Hoc Wireless Networks*, IEEE Transactions on Information Theory 51 (3), March 2005, 858–865.
- [9] L.-L. Xie and P. R. Kumar, *On the Path-Loss Attenuation Regime for Positive Cost and Linear Scaling of Transport Capacity in Wireless Networks*, IEEE Transactions on Information Theory 52 (6), June 2006, 2313–2328.
- [10] S. Ahmad, A. Jovicic and P. Viswanath, *Outer Bounds to the Capacity Region of Wireless Networks*, IEEE Transactions on Information Theory 52 (6), June 2006, 2770–2776.
- [11] A. Özgür, O. Lévêque and D. Tse, *Hierarchical Cooperation Achieves Optimal Capacity Scaling in Ad Hoc Networks*, submitted to the IEEE Transactions on Information Theory, 2006; on ArXiv: cs.IT/0611070.
- [12] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*, Cambridge University Press, 2005.
- [13] M. Grossglauser and D. Tse, *Mobility Increases the Capacity of Adhoc Wireless Networks*, IEEE/ACM Transactions on Networking 10 (4), August 2002, 477–486.