Fundamental Limits of Cooperation

Angel Lozano, Robert W. Heath Jr., *Fellow, IEEE,* and Jeffrey G. Andrews, *Fellow, IEEE*

Abstract—Cooperation is viewed as a key ingredient for inter-ference management in wireless networks. This paper shows that cooperation has fundamental limitations. First, it is established that in systems that rely on pilot-assisted channel estimation the spectral efficiency is upper-bounded by a quantity that does not depend on the transmit powers; in this framework, cooperation is possible only within clusters of limited size, which are subject to out-of-cluster interference whose power scales with that of the in-cluster signals. Second, an upper bound is also shown to exist if the cooperation extends to an entire (large) system operating as a single cluster; here, pilot-assisted transmission is necessarily transcended. Altogether, it is concluded that cooperation cannot in general change an interference-limited network to a noise-limited one. Consequently, existing literature that routinely as-sumes that the high-power spectral efficiency scales with the log-scale transmit power provides only a partial characterization. The complete characterization proposed in this paper subdivides the high-power regime into a *degrees-of-freedom regime*, where the scaling with the log-scale transmit power holds approximately, and a *saturation regime*, where the spectral efficiency hits a ceiling that is independent of the power. Using a cellular system as an example, it is demonstrated that the spectral efficiency saturates at power levels of operational relevance.

Index Terms—Wireless communications, interference, cooper-ative systems, wireless networks

I. INTRODUCTION

WIRELESS networks with many uncoordinated trans-mitters and receivers utilizing the same spectrum are interference-limited, meaning that an increase in the transmit powers does not improve the spectral efficiency once those powers are sufficiently high. In cellular systems, spectrum reuse leads to a large fraction of users having a low signal-to-interference-plus-noise ratio (SINR) unless bandwidth-wasting methods such as reuse patterns are implemented to relieve cell-edge users. Similar effects can be observed in other types of networks: in WiFi, inefficient contention-based medium access protocols protect receivers from interference by silencing nearby transmitters.

It has been persuasively argued in a by-now vast literature that this limitation is not fundamental, but rather an artifact

of each transmitter-receiver pair communicating autonomously rather than cooperatively (cf. [3]–[13] and references therein). If the various nodes could cooperate, the logic goes, the corresponding interference channel could be converted to a broadcast channel—for the downlink—or a multiple access channel—for the uplink—with all the transmitters (respec-tively receivers) jointly encoding (respectively decoding). In the cellular context, it would seem from this line of thinking that an arbitrary number of base stations (BSs) could cooperate to achieve enormous spectral efficiency gains over the lone-BS model, with the only limitation being the amount of coordination that can be afforded [14].

This is currently a problem of considerable theoretical and practical interest. Incomplete but significant cooperation between BSs has been attempted by industry and is still ongoing, broadly under the current moniker of "coordinated multipoint" (CoMP). The bit rate and latency of the backhaul links have restrained the benefits of such cooperation thus far, with disappointing improvements in average spectral efficiency typically not exceeding 30% [15]–[17]. In fact, one respected group has even observed a net loss from cooperative tech-niques when the various over-the-air overheads are accounted for [18]. Are these widespread observations a by-product of current technology limitations that could be overcome with better/more backhaul (e.g., over dedicated fiberoptic control channels), improved feedback and overhead techniques, and/or better joint encoding and decoding methods? Or is there a fundamental limitation lurking beneath the surface, one that is independent of the particular technology?

A. Background and Status Quo

There are currently about 3 million macrocellular BSs worldwide and that number is expected to reach 50 million by around 2015, once small cells (picocells and femtocells) are incorporated [19]. In a typical urban area, thousands to tens-of-thousands of BSs occupy the same spectrum. Because of the ubiquity and commercial importance of cellular networks, plus the clear scope for cooperation in them, and for the sake of concreteness, we focus the discussion around BS cooperation. However, the subsequent models and main results apply—at least qualitatively—to any generic network.

Evaluating the performance of cellular systems is an arduous task. Large-scale simulations are ultimately necessary to verify the performance of any specific technique, but they are hardly the best way to devise and probe ideas, build intuition, and glean insights. Most designs are therefore incubated in much simpler settings that represent a fragment of a network, and only eventually are they transplanted and trialed at the level of an entire system. A typical such controlled setting (at a given time epoch) is described by the following relationship.

Relationship 1 *The observation at receiver n is*

$$Y_n = \sum_{k=1}^{K} H_{nk}\sqrt{P}X_k + Z_n \qquad n = 1, \ldots, N \qquad (1)$$

where K and N are the number of transmitters and receivers, respectively, X_k is the signal generated by transmitter k, normalized such that P is its power, and H_{nk} is the channel from transmitter k to receiver n. Power differences at the K transmitters can be simply absorbed into the channel coefficients or the noise variances. The term Z_n is the noise at receiver n, typically white and Gaussian, with some normalized variance.

The signals and channel coefficients in Relationship 1 can be scalars, or else properly dimensioned vectors and matrices to accommodate multiple-input multiple-output (MIMO) techniques. For the sake of exposition, unless otherwise stated, we henceforth consider single-antenna transmitters and receivers. Thus, the elements of Relationship 1 are scalars.

Relationship 1 subsumes, in effect, most basic information-theoretic channel settings.

- Single-user channel if $K = N = 1$.
- Multiaccess channel (MAC) if $K > 1$ and $N = 1$.
- Broadcast channel (BC) if $K = 1$ and $N > 1$.
- Interference channel (IC) if $K = N$ with $K, N > 1$.

Combinations of these basic settings are also possible. Relationship 1 has led to characterizations of the single-user, MAC and BC capacity, as well as asymptotic notions such as the diversity-multiplexing tradeoff (DMT) [20] and practical techniques like multiuser MIMO.

Relationship 1 has been applied to study the IC both generically and in the context of cellular systems. Despite the fact that the capacity of the IC is yet to be determined, Relationship 1 has led to advances in its understanding.

- The definition of relevant quantities such as the number of degrees of freedom (DoF) [21].
- The development of centralized-decoding cooperative schemes such as Network MIMO [22].
- The genesis of distributed-decoding cooperative solutions such as interference alignment (IA) [23], Max-SINR [24], [25] and other forms of cooperative interference management [26].

Most developments on the IC take place in the high-power regime, which is where its nature comes to the fore. For any setting conforming to Relationship 1, the high-power behavior is as illustrated in Fig. 1. With the channel coefficients $\{H_{nk}\}$ known by the corresponding receivers and possibly also by the transmitters, the spectral efficiency grows, for $P \to \infty$, linearly with $\log P$ with a slope given by the number of DoF. Cooperative techniques such as Network MIMO or IA aim at maximizing the number of DoF.

The IC developments spawned by Relationship 1 have spurred an extensive amount of publications that invariably promise major improvements in spectral efficiency through cooperation. However, as noted at the outset, in subsequent system-level simulations these gains do not seem to materialize: DoF maximizations attained through IA translate only
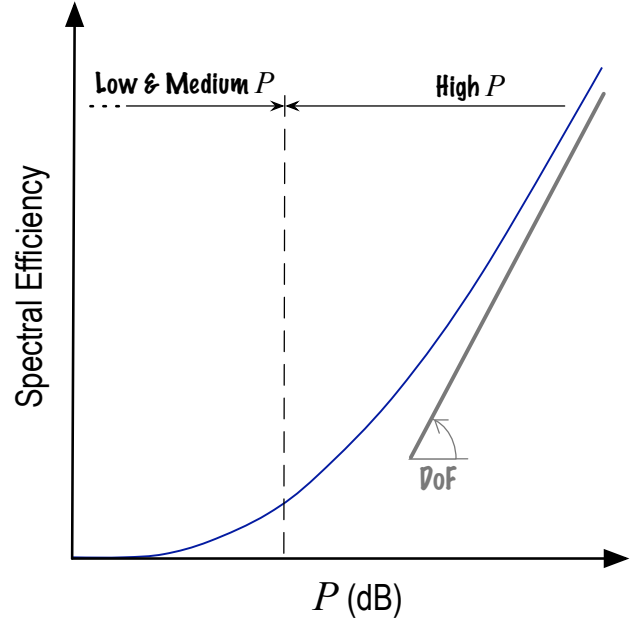

Fig. 1. Spectral efficiency v. P (in dB) according to Relationship 1.

to marginal gains or even outright losses in spectral efficiency [27]–[30]. In [31], for instance, it was shown that a 300% gain from using IA in an isolated 3-user setting shrinks to a mere 28% gain when the same exact IA scheme is applied to a 19-cell system. (The conditions are otherwise identical: transmitters and receivers have 4 antennas and the signal-to-noise ratio is 20 dB.) These serial many-fold discrepancies point to a disconnect, to some fundamental way in which a fragment of a cellular system is not properly modeled by Relationship 1.

B. Modeling a Cluster Within a System

An obvious problem with Relationship 1 is that there is a cutoff of K (potentially cooperating) transmitters, and all other interference is ignored. The K cooperating transmitters are typically geographical neighbors and are referred to as a *cluster*. It is habitually assumed that any interference from outside the cluster can be lumped into the $\{Z_n\}$ noise terms. It cannot, because each Z_n has a fixed variance that does not depend on P whereas the out-of-cluster interference power does scale with P. This scaling can be captured by modifying Relationship 1 as follows.

Relationship 2 *The observation at receiver n is*

$$Y_n = \sum_{k=1}^{K} H_{nk}\sqrt{P}X_k + \sum_{k=K+1}^{\tilde{K}} H_{nk}\sqrt{P}X_k + Z_n \qquad (2)$$

where $\tilde{K}$ and $\tilde{N}$ are the total numbers of transmitters and receivers in the system while K and N are the ones cooperating. Defining $Z'_n = \sum_{k=K+1}^{\tilde{K}} H_{nk}X_k$ as the out-of-cluster interference at receiver n,

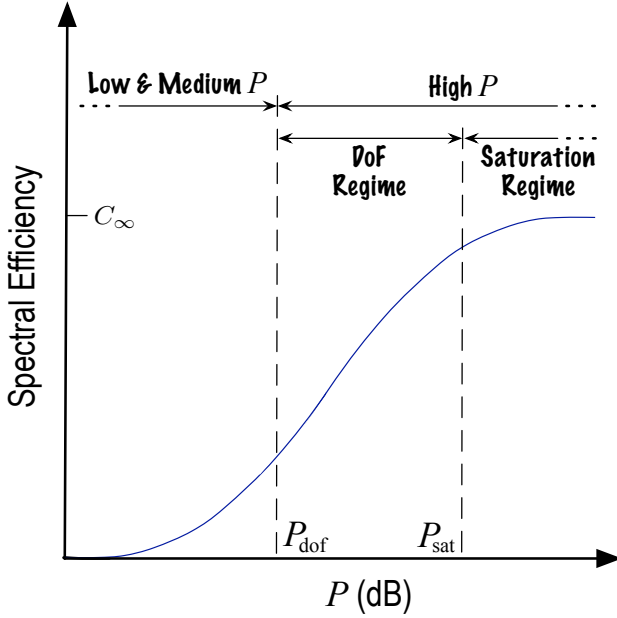$$Y_n = \sum_{k=1}^{K} H_{nk}\sqrt{P}X_k + \sqrt{P}Z'_n + Z_n. \qquad (3)$$

Fig. 2. Spectral efficiency v. P (in dB) according to Relationship 2.

As $P \to \infty$, the out-of-cluster interference does as well if the total network size is greater than K. This gives rise to a *saturation regime* as shown in Fig. 2, where further increasing the power P above some value P_{sat} does not noticeably improve the spectral efficiency because Z'_n dominates over the noise Z_n. Previous works in [32], [33] had acknowledged that the interference power scales with P, but the representation therein differs from Relationship 2 in at least one important way: in [32] and [33], the channel coefficients are known perfectly—at no cost—by all transmitters and receivers while, in Relationship 2, the channel gains outside the cluster of interest cannot be known (this is in fact what will come to define the boundaries of a cluster) and those within the cluster have to be learned, explicitly or not. Under the assumptions in [32] and [33], the spectral efficiency still scales indefinitely with $\log P$, only with a modified and rather intuitive DoF notion dubbed *Generalized DoF*. Under Relationship 2, in contrast, the aforementioned saturation is observed.

C. Summary of Contributions

The correct representation of a cluster within a large wireless system is given by Relationship 2 and not Relationship 1, and cooperation behaves very differently under the two representations. A possible objection to this claimed distinction could be that the high-power saturation occurs only because the cluster size K, i.e., the number of cooperating transmitters, is too small. If K were made large enough, wouldn't the saturation eventually disappear since interfering transmitters would be converted to cooperating ones? The answer this paper establishes is negative. Extending the cooperation scope to ever larger clusters can at best shift the saturation point, P_{sat}. This fundamental limitation stems from the selective nature of wireless channels, which inevitably vary in frequency and

in time (even if users are static, because of environmental motion), and from the dynamic nature of wireless networks, where users constantly switch from sleep-mode onto activity and vice versa.

The points made in the sequel, which are consistent with other recent works [34]–[36], descend from the model detailed in Section II and its subsequent analysis in Sections III and IV. The main observations from this analysis and the accompanying examples are as follows.

1) Because of out-of-custer interference, a cluster within a cellular system is described by Relationship 2 rather than Relationship 1 and the performance of any cooperative technique resembles Fig. 2 rather than Fig. 1.
2) For high P, two distinct sub-regimes can then be identified: a *DoF regime* where the noise dominates over the out-of-cluster interference and a *saturation regime* where the two become comparable and the spectral efficiency chokes. The notion of DoF is only meaningful within the DoF regime.
3) Two new quantities of interest emerge: P_{sat} and C_∞. The former is the power level where the transition to saturation is said to occur whereas the latter is the corresponding limiting spectral efficiency. These new quantities depend on the system topology, the propagation laws, and the degree of channel selectivity and user dynamics. In most cases, the transition to the saturation regime takes place well within the range of operational interest and thus studies conducted using Relationship 1 are at risk of being misleading in terms of system-level performance.
4) The saturation takes place regardless of whether the communication is assisted by pilots, or not. Therefore, the saturation does not occur simply because of explicit channel estimation or because of pilot overhead in the face of finite channel coherence. Nor could improved channel estimation avert the problem.

II. DETAILED SYSTEM MODEL

This section is devoted to refining Relationship 2 and to describing the models used to embody it in the remainder of the paper.

A. Large-Scale Modeling

Let G_{nk} be the average channel power gain (associated with distance decay, shadowing, building penetration losses, and antenna patterns) between transmitter k and receiver n and let us also define the normalized channel power gains

$$g_{nk} = \frac{G_{nk}}{\sum_{\kappa=1}^{K} G_{n\kappa}} \tag{4}$$

such that, for every n,

$$\sum_{k=1}^{K} g_{nk} = 1. \tag{5}$$

Thus, g_{nk} signifies the share of receiver n's signal power that corresponds to transmitter k and the set $\{g_{nk}\}$ completely and

compactly characterizes the degree of connectedness among in-cluster nodes. We term the set $\{g_{nk}\}$ the *geometry profile*.

We can further absorb the various normalizations for each receiver n into a signal-to-noise ratio (SNR_n) that scales with P and a signal-to-(out-of-cluster)-interference ratio (SIR_n) that does not scale with P. Specifically,

$$\mathsf{SNR}_n = \frac{\sum_{k=1}^{K} G_{nk} P}{N_0 B} \tag{6}$$

$$\mathsf{SIR}_n = \frac{\sum_{k=1}^{K} G_{nk}}{\sum_{k=K+1}^{\tilde{K}} G_{nk}} \tag{7}$$

where N_0 is the noise spectral density and B the bandwidth.

From the definitions in (5), (6) and (7), we can rewrite (3) as

$$y_n = \sqrt{\mathsf{SNR}_n} \sum_{k=1}^{K} \sqrt{g_{nk}} h_{nk} x_k + \sqrt{\frac{\mathsf{SNR}_n}{\mathsf{SIR}_n}} z_n' + z_n \tag{8}$$

where the noise terms $\{z_n\}$, the out-of-cluster interference terms $\{z_n'\}$, the signals $\{x_k\}$, and the fading channel coefficients $\{h_{nk}\}$ are all mutually independent random variables normalized to be unit-variance.

If both the noise and the out-of-cluster interference are Gaussian, then (8) becomes

$$y_n = \sqrt{\mathsf{SNR}_n} \sum_{k=1}^{K} \sqrt{g_{nk}} h_{nk} x_k + \sqrt{1 + \frac{\mathsf{SNR}_n}{\mathsf{SIR}_n}} z_n'' \tag{9}$$

where z_n'' is the aggregate noise-plus-interference, also Gaussian and unit variance, i.e., $z_n'' \sim \mathcal{N}_{\mathbb{C}}(0,1)$. Since the Gaussian distribution correctly models thermal noise, and the out-of-cluster interference is made up of a large number of independent terms and thus its distribution tends to be approximately Gaussian too, we focus on (9) with $z_n'' \sim \mathcal{N}_{\mathbb{C}}(0,1)$. Nevertheless, all of the points made henceforth apply—at least qualitatively—in the wider generality of (8), with z_n and z_n' having different distributions. Eq. (9) can be further rewritten into the following.

Relationship 3 *The observation at receiver n is given by*

$$y_n = \sqrt{\frac{\mathsf{SNR}_n \, \mathsf{SIR}_n}{\mathsf{SNR}_n + \mathsf{SIR}_n}} \sum_{k=1}^{K} \sqrt{g_{nk}} h_{nk} x_k + z_n'' \tag{10}$$

$$= \sqrt{\mathsf{SINR}_n} \sum_{k=1}^{K} \sqrt{g_{nk}} h_{nk} x_k + z_n'' \tag{11}$$

where SINR_n equals the harmonic mean of SNR_n and SIR_n, i.e.,

$$\frac{1}{\mathsf{SINR}_n} = \frac{1}{\mathsf{SNR}_n} + \frac{1}{\mathsf{SIR}_n}. \tag{12}$$

For small P, $\mathsf{SINR}_n \approx \mathsf{SNR}_n$ whereas, for large P, $\mathsf{SINR}_n \approx \mathsf{SIR}_n$. Although functionally equivalent to Relationship 2, the above formulation on the basis of SNR_n and SIR_n is more general in that it captures not only the scaling with P, but also with other parameters such as cell size or noise variance.

B. Small-Scale Modeling

Referring back to Relationship 3, the small-scale fading is modeled as Rayleigh distributed and thus the normalized fading coefficients satisfy $h_{nk} \sim \mathcal{N}_{\mathbb{C}}(0,1)$. (The unnormalized channel coefficients utilized in Relationship 2 were thus $H_{nk} = \sqrt{G_{nk}} h_{nk}$.) We consider frequency-selective fading with coherence bandwidth B_c, such that the value of each h_{nk} varies from subband to subband in an independent identically distributed (IID) fashion. In terms of time selectivity, both block- and continuous-fading are accommodated in our model.

- With block fading, the channels within each subband hold constant for some time and then change to a different value also in an IID fashion. If we denote by T_c the channel coherence in time, then, irrespective of how the signaling is arranged along the time and frequency dimensions, the number of symbols over which the channel remains coherent is roughly $L = B_c T_c$.
- With continuous fading, the channels within each subband are discrete-time stationary and ergodic random processes with a Doppler spectrum $S_h(\cdot)$ that is bandlimited, i.e., $S_h(\nu) = 0$ for $|\nu| > f_D$ with $f_D \leq 1/2$ some maximum Doppler frequency. Typically

$$f_D = \frac{v}{\lambda B_c} \tag{13}$$

where v is the velocity and λ is the carrier wavelength, yet our formulation holds regardless of how f_D is defined.

Block fading is a coarse but effective approximation to continuous fading, and remarkable equivalences between the two have been uncovered [37], [38]. In particular, both models have been shown to be equivalent in terms of channel estimation minimum mean-square error (MMSE) for the case of a rectangular Doppler spectrum, $S_h(\nu) = \frac{1}{2f_D}$ for $|\nu| \leq f_D$. With such a spectrum, the MMSE is equivalent to that of a block-fading channel having $L = \lceil \frac{1}{2f_D} \rceil$ (cf. Section III).

Example 1 *Application of (13) with the typical cellular values $\lambda = 0.15$ m and $B_c = 370$ KHz leads to $f_D = 2.5 \cdot 10^{-5}$ for $v = 5$ Km/h (pedestrian velocity) and $f_D = 5 \cdot 10^{-4}$ for $v = 100$ Km/h (vehicular velocity). This maps to $L = 2 \cdot 10^4$ and $L = 10^3$, respectively.*

C. Exemplary Cellular System

To close this section, let us introduce an exemplary system that shall be utilized throughout.

Example 2 *Consider a cellular system with tri-sector hexagonal cells of size R. Each sector's antenna has a uniform gain over the $120°$ span of the sector and a $Q|_{\mathrm{dB}}$-lower uniform gain outside that span. Orthogonal signaling resources (time slots and frequency subcarriers) are allocated to the users within each sector. On any given resource, thus, there is a single user per sector and hence $K = N$. Each user is centered in azimuth within its sector and at distance $2R/3$ from the BS. Depicted in Fig. 3 are an arbitrary reference cell and the two tiers around it, yet the system has infinitely many cells. The signals experience distance-dependent decay with an exponent γ and also Rayleigh fading.*
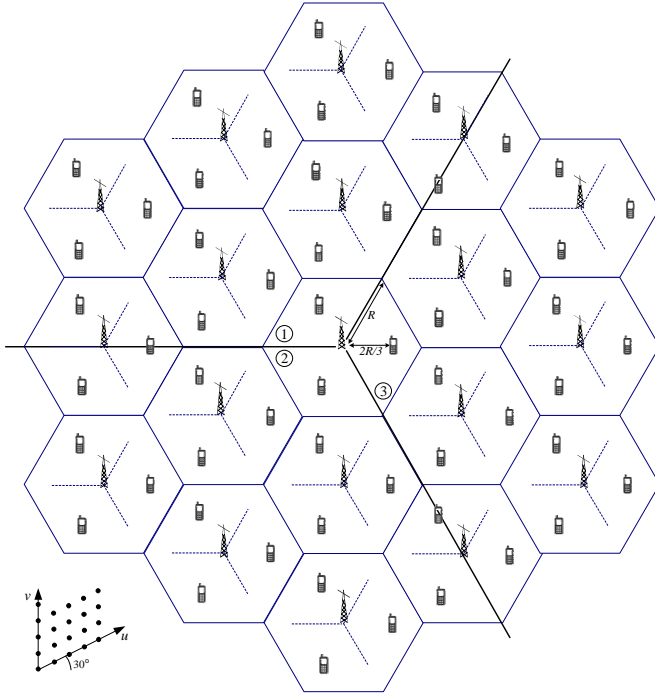
Fig. 3. Regular universe with tri-sector hexagonal cells of size R. Each user is centered in azimuth within its sector and at distance $2R/3$ from the BS.

Example 2 is representative of a cellular system while having the virtue of having regular user locations. Relationship 3 can then be embodied for any desired values of $K = N$, i.e., for clusters of arbitrary size, and the out-of-cluster interference can be easily summed for $\tilde{K}, \tilde{N} \rightarrow \infty$ thanks to the regular user locations.

III. CLUSTERED SYSTEM WITH PILOT-ASSISTED CHANNEL ESTIMATION

In this section, we establish that the spectral efficiency inevitably saturates for growing P whenever a system is arranged into cooperation clusters. Specifically, the focus is on clusters of dimension $K < L$, which renders it possible to estimate all the channel coefficients within each cluster on the basis of pilot symbol transmissions. The rationale for studying systems featuring explicit channel estimation followed by coherent detection of payload data [39] is that virtually every existing wireless system abides by this methodology. The procedure entails the transmission of pilot symbols regularly in time (and in frequency when multiples of B_c are spanned). We note the following.

- Pilot symbols are overhead.
- The periodicity of pilot transmission is determined by the channel coherence.

Only a finite number of pilot symbols can be transmitted within a given L-symbol coherence interval. At least one pilot symbol, and possibly many, must be devoted to enabling the estimation of each channel coefficient at the corresponding receiver within each coherence interval. If the channels were matrix-valued, as in MIMO, separate pilot symbols would be needed for every entry therein. The number of channel

coefficients that can be estimated is therefore limited by L, hence the restriction in this section that $K < L$. Having a limited size, each cluster is inevitably exposed to interference from all the transmitters beyond its borders.

Intuitively, excessively large clusters end up incurring a disproportionate pilot overhead and/or poor channel estimates, which, coupled with the diminishing returns in terms of useful signal recovered from faraway transmitters, ultimately curb the benefits of cooperation. To make this intuition specific, denote by α the share of symbols reserved for pilots; the rest, $(1-\alpha)$, is available for payload data. The pilot transmissions should be orthogonally multiplexed from each of the transmitters [40] and thus every receiver's estimation of each of the K channel coefficients corresponding to an in-cluster transmitter relies on a share α/K of the symbols. When only the kth transmitter is actively sending pilots, the SINR at receiver n is $g_{nk}\mathsf{SINR}_n$. With block fading, the estimation MMSE for h_{nk} is then [39]

$$\mathsf{MMSE}_{nk} = \frac{1}{1 + g_{nk}\mathsf{SINR}_n L\alpha/K} \tag{14}$$

whereas, with continuous fading [37], [41]

$$\mathsf{MMSE}_{nk} = 1 - \int_{-f_D}^{f_D} \frac{g_{nk}\mathsf{SINR}_n S_h^2(\nu)}{K/\alpha + g_{nk}\mathsf{SINR}_n S_h(\nu)} d\nu. \tag{15}$$

If the Doppler spectrum $S_h(\cdot)$ is rectangular, (15) becomes

$$\mathsf{MMSE}_{nk} = \frac{1}{1 + g_{nk}\mathsf{SINR}_n \frac{\alpha/K}{2f_D}} \tag{16}$$

which coincides with (14) for $f_D = \frac{1}{2L}$. The MMSE expression in (16) thus allows embracing both the block- and continuous fading models in a single framework [37]. Note that, to accommodate at least one pilot symbol per channel coefficient to be estimated, we need $\alpha \geq \alpha_{\min}$ where $\alpha_{\min} = K/L$ with block fading and $\alpha_{\min} = 2Kf_D$ with continuous fading.

Denoting the channel estimate by $\hat{h}_{nk}$ and the corresponding estimation error by $\tilde{h}_{nk}$, it follows that $h_{nk} = \hat{h}_{nk} + \tilde{h}_{nk}$ with $\hat{h}_{nk}$ and $\tilde{h}_{nk}$ uncorrelated and with $\mathbb{E}[|\tilde{h}_{nk}|^2] = \mathsf{MMSE}_{nk}$ and $\mathbb{E}[|\hat{h}_{nk}|^2] = 1 - \mathsf{MMSE}_{nk}$. Hence, from Relationship 3,

$$y_n = \sqrt{\mathsf{SINR}_n} \sum_{k=1}^{K} \sqrt{g_{nk}}\hat{h}_{nk}x_k + \sqrt{\mathsf{SINR}_n} \sum_{k=1}^{K} \sqrt{g_{nk}}\tilde{h}_{nk}x_k + z_n'' \tag{17}$$

The receivers customarily utilize the channel estimates as if they were correct, in which case the terms in the second summation in (17) play the role of additional Gaussian noise [42]. With that, the effective SINR at receiver n upon payload data detection is

$$\mathsf{SINR}_n^{\mathrm{eff}} = \frac{\mathsf{SINR}_n \sum_{k=1}^{K} g_{nk}(1 - \mathsf{MMSE}_{nk})}{1 + \mathsf{SINR}_n \sum_{k=1}^{K} g_{nk}\mathsf{MMSE}_{nk}} \tag{18}$$

and the average spectral efficiency (bits/s/Hz/user) that can be attained reliably is

$$C = (1-\alpha)\frac{1}{N}f\left(\mathsf{SINR}_1^{\mathrm{eff}}, \cdots, \mathsf{SINR}_N^{\mathrm{eff}}\right) \tag{19}$$

where the function $f(\cdot)$ depends on the type of cooperation among the $K = N$ users. The spectral efficiency is maximized by a proper choice of $\alpha_{\min} \leq \alpha < 1$. As K grows, $\alpha_{\min}$
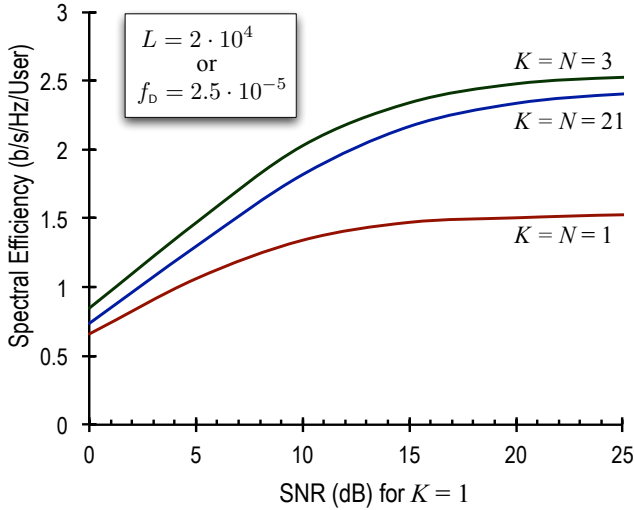
Fig. 4. Spectral efficiency as function of the single-user SNR, given by $PG_{kk}/(N_0 B)$, for varying cluster sizes with either block fading ($L = 2 \cdot 10^4$) or continuous fading ($f_D = 2.5 \cdot 10^{-5}$). Uplink with full cooperation.
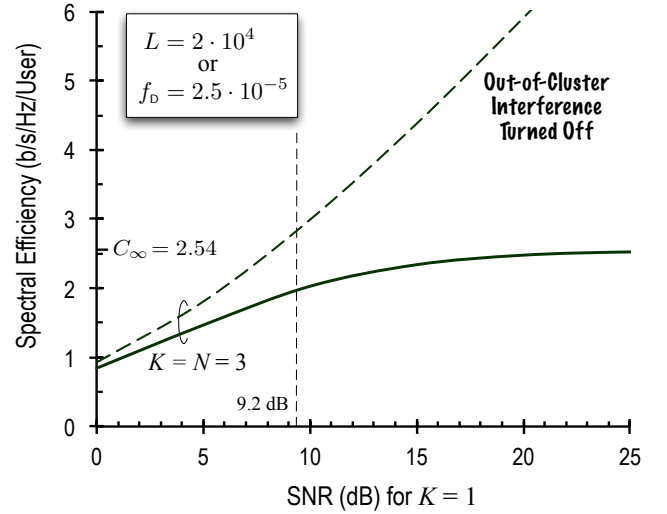

Fig. 5. Spectral efficiency as function of the single-user SNR, given by $PG_{kk}/(N_0 B)$, with either block fading ($L = 2 \cdot 10^4$) or continuous fading ($f_D = 2.5 \cdot 10^{-5}$). Uplink with full cooperation. In solid, with the out-of-cluster interference included. In dashed, with the out-of-cluster interference turned off.

grows with it squeezing α into higher values; additionally, α is pushed further beyond $\alpha_{\min}$ by the need to maintain effective SINRs with ever more distant transmitters and, as a result, the spectral efficiency is bound to peak at some cluster size and diminish thereafter.

Example 3 *Consider the uplink of the system in Example 2 with $K = N$ and with clusters of dimension $K < L$ where the in-cluster BSs fully cooperate via Network MIMO, i.e., they jointly decode the K signals received at the N sectors. The distance-decay exponent is $\gamma = 3.8$ whereas $Q|_{dB} = 20$ dB. Defining an $N \times K$ matrix $\mathbf{S}$ whose entries are independent and such that the (n, k)th entry is $\mathbf{S}_{nk} \sim \mathcal{N}_{\mathbb{C}}(0, \sigma_{nk}^2)$ where*

$$\sigma_{nk}^2 = \frac{g_{nk}\mathsf{SINR}_n(1 - \mathsf{MMSE}_{nk})}{1 + \mathsf{SINR}_n \sum_{\kappa=1}^K g_{n\kappa}\mathsf{MMSE}_{n\kappa}}, \quad (20)$$

the average spectral efficiency (bits/s/Hz/user) is

$$C = (1 - \alpha)\frac{1}{N}\mathbb{E}[\log_2 \det(\mathbf{I} + \mathbf{SS}^\dagger)]. \quad (21)$$

Note that, because of the regular user locations, $\mathsf{SNR}_n = \mathsf{SNR}$ for $n = 1, \ldots, N$. The fading is either block ($L = 2 \cdot 10^4$) or continuous with a rectangular Doppler spectrum ($f_D = 2.5 \cdot 10^{-5}$), which are equivalent in terms of channel estimation and correspond to a pedestrian velocity. Shown in Fig. 4 is the spectral efficiency for several cluster dimensions as a function of the single-user SNR, i.e., the SNR without cooperation given by $PG_{kk}/(N_0 B)$. The baseline curve for $K = N = 1$ corresponds precisely to single-user decoding, without cooperation. The case $K = N = 3$ corresponds to a cluster of 3 facing sectors as in Fig. 3. Finally, the case $K = N = 21$ corresponds to a cluster of 7 cells: one central cell plus the first tier around it. The pilot overhead α is equal for all users but otherwise optimized to maximize C, i.e., it is optimized for every cluster size and power level. For further details, see Appendix C.

The performance in Example 3 improves when the cluster size goes from 1 to 3, but then degrades slightly when the cluster size increases to 21 and continues to degrade for even larger sizes. To gauge the impact of out-of-cluster interference, we replot the spectral efficiency corresponding to $K = N = 3$ in Fig. 4 next to its counterpart in the same exact conditions except with all the transmitters outside the cluster of interest turned off. This comparison, presented in Fig. 5, evidences that without out-of-cluster interference we recover the traditional behavior (cf. Fig. 1). Out-of-cluster interference, however, drastically modifies that behavior (cf. Fig. 2). Note that modeling the out-of-cluster interference as additional fixed-variance noise cannot fix the representation in Relationship 1, as this would merely shift the spectral efficiency by a certain amount. Only the representation in Relationships 2–3 can properly reproduce the actual behavior, as illustrated next for Example 3.

Example 4 *For Example 3 with $K = N = 3$, the SIR that should be inserted into Relationship 3 equals*

$$\mathsf{SIR}_n = \frac{\sum_{k=1}^3 G_{nk}}{\sum_{k=4}^\infty G_{nk}} \quad (22)$$

$$= 9.2 \text{ dB} \qquad n = 1, 2, 3 \quad (23)$$

which, indeed, corresponds with the inflection point observed in Fig. 5. This SNR point, which we had defined earlier as P_{sat} in terms of transmit power, delineates the transition between the DoF regime and the saturation regime. In turn, the limiting spectral efficiency can be found to be $C_\infty = 2.54$ bits/s/Hz/user.

With different cluster arrangements, or simply with user locations different from the regular ones in Example 3, the values of $\mathsf{SIR}_1, \ldots, \mathsf{SIR}_N$ change and thus the optimum cluster size may vary and the transition between the DoF and the
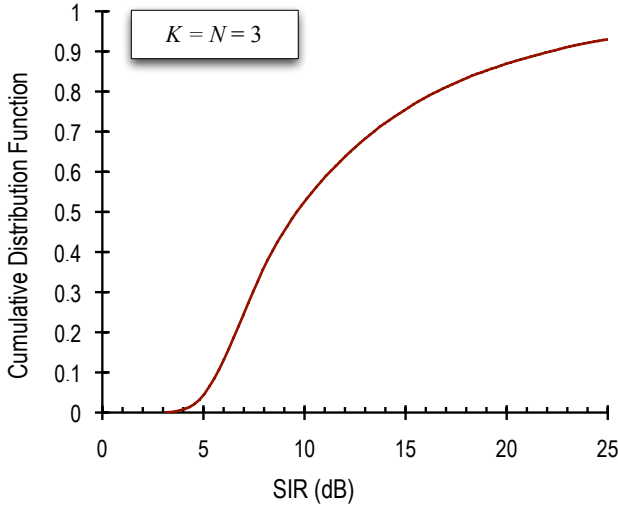
Fig. 6. Cumulative distribution of SIR over all in-cluster locations for $K = N = 3$. Uplink with $K = N = 3$ and full cooperation.

saturation regimes may take place at different points, but qualitatively speaking the behavior is unaltered.

Example 5 *Consider a variation of Example 3, with the out-of-cluster users still centered within their own sectors but with the location of each in-cluster user randomized to be uniform over its sector. Shown in Fig. 6 is the cumulative distribution of SIR at each of the three in-cluster receivers. In almost* 90% *of locations, the SIR is below* 20 *dB. The median SIR is at* 9.6 *dB, confirming that the value obtained in Example 4 is indeed quite representative.*

The precedings examples indicate that, in a majority of cases, the spectral efficiency with clusters implementing full cooperation saturates at SNR levels of operational interest. The final steps to fully generalize these figures would be to randomize the location of the out-of-cluster users and to incorporate shadow fading on all the links. Exactly adding up the out-of-cluster interference for an infinitely large network becomes challenging in that broad generality, but the computation is otherwise conceptually identical and the resulting SIR distribution is not expected to substantially depart from the one in Fig. 6.

Qualitatively similar observations to the ones made in this section for the uplink can be made for the downlink, which could be analyzed by extending the pilot-assisted scheme in [43] to the multicell realm.

IV. FULLY COOPERATIVE SYSTEM

The analysis in the previous section confirmed that, as a result of out-of-cluster interference, the spectral efficiency is upper-bounded by a quantity which does not depend on the transmit power. The analysis, however, relied on pilot-based channel estimation and thus had the restriction $K < L$ on the dimensionality of the clusters. Conceivably, if clusters could be overcome and the entire system could cooperate as one, then out-of-cluster interference would be eliminated altogether

thereby restoring—backhaul and hardware limitations aside—the unbounded growth of the spectral efficiency with the SNR. This section explores this possibility, admittedly quixotic but relevant from the vantage of a fundamental understanding. For $K > L$, there are not sufficient opportunities within each coherence interval to separately estimate all the channel coefficients and thus we resort to a pure information-theoretic characterization; this can be interpreted as pertaining to the realm of noncoherent communication, since no explicit channel estimation is presumed, but we hasten to emphasize that in actuality this characterization subsumes any feasible form of communication, noncoherent or otherwise.

In the absence of explicit channel-state information (CSI) at the receivers, the capacity-achieving signal distribution and the capacity itself are generally unknown. The spectral efficiency achievable with complex Gaussian signals is also unknown, save for certain single-user channels [44]. Given the difficulty in computing the exact spectral efficiency achievable by a fully cooperative system without CSI at the receivers, we focus on showing that its limiting value for $P \to \infty$ is bounded by above by a quantity that does not depend on P.

Consider the uplink again. With block fading, the transmit-receive relationship can be vectorized for the entire system and all the symbols in a channel coherence block as

$$\mathbf{Y} = \mathrm{diag}\left\{\sqrt{\mathsf{SNR}_1}, \ldots, \sqrt{\mathsf{SNR}_N}\right\}\mathbf{H}\mathbf{X} + \mathbf{Z} \qquad (24)$$

where $\mathbf{Y}$ and $\mathbf{Z}$ are $N \times L$, $\mathbf{H}$ is $N \times K$ and $\mathbf{X}$ is $K \times L$. The entries of $\mathbf{Z}$ are IID with $\mathbf{Z}_{n\ell} \sim \mathcal{N}_{\mathbb{C}}(0,1)$ whereas the entries of $\mathbf{H}$ are independent with $\mathbf{H}_{nk} \sim \mathcal{N}_{\mathbb{C}}(0, g_{nk})$; notice that the normalized channel power gains $\{g_{nk}\}$ are directly incorporated as the variances of the entries of $\mathbf{H}$. The entries of $\mathbf{X}$ are unit variance and, since the kth row of $\mathbf{X}$ contains the signal sequence transmitted by user k over the L symbols of a fading block, the rows of $\mathbf{X}$ are independent. The entire system is represented by (24) and there is no out-of-cluster interference because all the transmitters in the system participate in the cooperation.

Proposition 1 *Consider the uplink of a cellular system subject to block-fading with* $K > L$*. Define, for* $n = 1, \ldots, N$*, respective diagonal matrices* $\mathbf{G}_n = \mathrm{diag}\{g_{n1}, \ldots, g_{nK}\}$*. Each* $\mathbf{G}_n$ *contains along its diagonal the normalized power gains between the* K *transmitters and the* nth *receiver and thus* $\mathrm{Tr}\{\mathbf{G}_n\} = 1$*. If* $\mathbf{X}$ *is full rank w.p.1, the average spectral efficiency that a fully cooperative system can achieve reliably for* $P \to \infty$ *satisfies* $C_\infty \leq C_\infty^{\mathsf{UB}}$ *with*

$$C_\infty^{\mathsf{UB}} = -\frac{1}{K}\sum_{n=1}^{N}\frac{1}{L}\,\mathbb{E}\left[\log_2 \det\left(\mathbf{X}^\dagger \mathbf{G}_n \mathbf{X}\right)\right]. \qquad (25)$$

Proof: *See Appendix A.*

In light of (5), C_∞^{UB} remains bounded for $K, N \to \infty$ with a fixed ratio K/N. In large systems, therefore, sheer cooperation cannot restore the unbounded growth of the spectral efficiency with P.

The constraint on the rank of $\mathbf{X}$ is in principle mild, because wireless transmissions invariably amount to sequences of IID

symbols, but it does lead to a way out of the result that is discussed in detail in Section V.

If the signal sequence transmitted by each user is not only IID, but complex Gaussian, then C_∞^{UB} in (25) can be expressed using [45, Prop. 4] in a form that is closed but not particularly convenient to work with when K and L are large.

A more compact analytical handle on C_∞^{UB} can be obtained by resorting to large-dimensional results in random matrix theory, with the added advantage that the expressions then hold for IID signals regardless of the distribution from which the symbols are drawn, e.g., PSK or QAM in addition to complex Gaussian. The following result offers an approximation derived through random matrix theory.

Proposition 2 *Consider the uplink of a cellular system subject to block-fading with $K > L$. If every transmission consists of IID symbols drawn from a common distribution, then*

$$C_\infty^{\mathsf{UB}} \approx \frac{N}{K} \log_2 e + \frac{1}{K} \sum_{n=1}^{N} \left[\log_2 \frac{a_n}{L} - \frac{1}{L} \sum_{k=1}^{K} \log_2(1 + a_n g_{nk}) \right] \tag{26}$$

with each a_n, for $n = 1, \ldots, N$, the nonnegative solution to

$$\sum_{k=1}^{K} \frac{g_{nk}}{g_{nk} + 1/a_n} = L. \tag{27}$$

Proof: *See Appendix B.*

In certain special situations such as the one in Example 2, Propositions 1 and 2 simplify significantly.

Corollary 1 *If the system is isotropic in the sense that the set of channel power gains looks the same from the vantage of each receiver, i.e., the set $\{g_{nk}\}_{k=1}^{K}$ can be reordered for every n into a common set $\{g_k\}_{k=1}^{K}$, then we can define a unique matrix $\mathbf{G}$ containing such common set on its diagonal and rewrite (25) as*

$$C_\infty^{\mathsf{UB}} = -\frac{N}{K} \frac{1}{L} \mathbb{E} \left[\log_2 \det \left(\mathbf{X}^\dagger \mathbf{G} \mathbf{X} \right) \right]. \tag{28}$$

In turn, Proposition 2 then specializes to

$$C_\infty^{\mathsf{UB}} \approx \frac{N}{K} \left[\log_2 e + \log_2 \frac{a}{L} - \frac{1}{L} \sum_{k=1}^{K} \log_2(1 + a g_k) \right] \tag{29}$$

with the unique a being the nonnegative solution to

$$\sum_{k=1}^{K} \frac{g_k}{g_k + 1/a} = L. \tag{30}$$

In the context of continuous fading channels, the expressions in both Proposition 2 and Corollary 1 can be interpreted using the equivalence $L = \left\lceil \frac{1}{2f_D} \right\rceil$ put forth earlier in the paper.

Before applying the above corollary to the infinitely large system in Example 2, we test its accuracy on a reduced version thereof, with $K = N$ and L small enough that the exact value in (28) can be computed numerically.

Example 6 *Consider the uplink of a finite version of the system in Example 2 having 20×20 cells ($K = 1200$) and let $L = 100$. All the sectors in the system cooperate fully via Network MIMO. The distance-decay exponent is $\gamma = 3.8$ whereas $Q|_{\mathrm{dB}} = 20$ dB. Monte-Carlo evaluation of (28) gives $C_\infty^{\mathsf{UB}} = 5.183$ b/s/Hz/user whereas (29)–(30) give $C_\infty^{\mathsf{UB}} = 5.181$ b/s/Hz/user.*

Let us now apply Corollary 1 to the full system in Example 2 for values of L in the range of practical interest.

Example 7 *Reconsider Example 6 with $K = N \to \infty$. For $L = 2 \cdot 10^4$, $C_\infty^{\mathsf{UB}} = 11.86$ bits/s/Hz/user whereas, for $L = 10^3$, $C_\infty^{\mathsf{UB}} = 7.98$ bits/s/Hz/user. For further details on the computations, see Appendix C.*

The spectral efficiencies actually achievable for $P \to \infty$ may be substantially lower than the values in Example 7, not only because these correspond to (not necessarily tight) upper bounds but also because the values of L utilized for the computations, which are correct over the propagation distances encountered within a single cell or a small cluster of cells, are bound to shrink with propagation taking place among faraway cells. Precisely, the long propagation delays that arise when distant units cooperate is sure to lead to a longer delay spread and thus a smaller B_c, with the consequent reduction in L.

A question that can be posed at this point is the following: if one wants to assume perfect CSI at the receivers, which values for $\mathsf{SIR}_1, \ldots \mathsf{SIR}_N$ should be inserted into Relationship 3 to reproduce the results in Example 7? We can gauge this from the values of C_∞^{UB} obtained therein. Because of the isotropy in Example 7, every user operates at the same spectral efficiency and thus $\mathsf{SIR}_n = \mathsf{SIR}$ for $n = 1, \ldots, N$. With perfect CSI at the receivers, complex Gaussian signals are capacity-achieving and the uplink capacity of Relationship 3 with full cooperation approaches, for $P \to \infty$,

$$C_\infty^{\mathsf{CSI}} = \frac{1}{K} \mathbb{E} \left[\log_2 \det \left(\mathbf{I} + \mathsf{SIR}\, \mathbf{H} \mathbf{H}^\dagger \right) \right]. \tag{31}$$

Applying [46, Thm. 5] as a finite-dimensional approximation

$$C_\infty^{\mathsf{CSI}} \approx 2 \log_2 \left(\frac{1 + \sqrt{1 + 4\,\mathsf{SIR}}}{2} \right) - \frac{\log_2 e}{4\,\mathsf{SIR}} \left(\sqrt{1 + 4\,\mathsf{SIR}} - 1 \right)^2. \tag{32}$$

Solving for the SIR that equates (32) with the values for C_∞^{UB} in Example 7, we obtain

$$\begin{aligned} \mathsf{SIR} &= 39.96\,\mathrm{dB}, \quad L = 2 \cdot 10^4 \text{ (pedestrian)} \\ \mathsf{SIR} &= 28.02\,\mathrm{dB}, \quad L = 10^3 \text{ (vehicular)} \end{aligned} \tag{33}$$

which, deriving from (approximations to) spectral efficiency upper bounds, are themselves (approximate) upper bounds on the actual SIR that quantifies the amount of interference to which channel uncertainty is effectively equivalent. Although, with shadow fading and randomized user locations, these equivalent SIR values are likely to vary significantly, they are indicative: if an entire system were to fully cooperate, it would still face a fundamental performance ceiling corresponding to values of SIR within the range of interest in high-power analysis. This ceiling, furthermore, would depend exclusively on the coherence L, which relates to the degree of selectivity and dynamics, and on the geometry profile $\{g_{nk}\}$, which quantifies the degree of connectedness among users.

Besides computational convenience, an additional benefit of the approximations in Proposition 2 and Corollary 1 is that they cast light on how C_∞^{UB} depends on the geometry profile.

- If $\{g_{nk}\}_{k=1}^K$ is highly skewed for a given n, then most of the power received by BS n corresponds to a few nearby users. Intuition then says that, given their relative strength, the fading of these users' channels and the overlaying signals could be determined and most of the received power therein could be rendered useful. The approximate results confirm this intuition: with fewer than L nonnegligible terms, C_∞^{UB} can be arbitrarily large and thus a sustained increase of the spectral efficiency with P is feasible (cf. Example 8 below).

- Alternatively, if $\{g_{nk}\}_{k=1}^K$ for a given n contains a myriad minute terms, rather than a few strong ones, each of these terms is simply too weak relative to the aggregate rest. Intuitively, this should give rise to an ocean of terms that are fundamentally undecodable. Again, the approximate results confirm this intuition: if $g_{nk} = 1/K$ for $k = 1, \ldots, K$, then, as K grows without bound for fixed L, C_∞^{UB} is not only curbed but vanishingly small (cf. Example 9 below).

Example 8 *Consider a skewed geometry profile given by $g_{nk} = (1-\epsilon)/K'$ for $k = 1, \ldots, K'$ and $g_{nk} = \epsilon/(K-K')$ for $k = K'+1, \ldots, K$, with $K' < L$ and $\epsilon \ll 1$. This corresponds to having $K' < L$ nonnegligible interferers and $K - K'$ negligible ones with ϵ determining the degree of negligibility. From (30), $a = \mathcal{O}(1/\epsilon)$ and, expanding (29) with respect to ϵ,*

$$C_\infty^{\mathsf{UB}} \approx \frac{N}{K}\left(1 - \frac{K'}{L}\right)\log_2 \frac{1}{\epsilon} + \mathcal{O}(1) \tag{34}$$

which grows without bound as ϵ vanishes and the skewness becomes pronounced.

Example 9 *Consider an unskewed geometry profile given by $g_{nk} = 1/K$ for $k = 1, \ldots, K$. From (30),*

$$a = \frac{1}{1/L - 1/K} \tag{35}$$

which, plugged into (29), yields

$$C_\infty^{\mathsf{UB}} \approx \frac{N}{K}\left[\log_2 e - \log_2\left(1 - \frac{L}{K}\right)\right.$$
$$\left. - \frac{K}{L}\log_2\left(1 + \frac{1}{K/L - 1}\right)\right] \tag{36}$$

$$= \frac{N}{K}\frac{L}{2K}\log_2 e + \mathcal{O}\left(\frac{L}{K}\right)^2 \tag{37}$$

which is not only curbed for fixed L, but small for $K \gg L$.

In an actual system, there will be a few strong signals accompanied by progressively weaker ones. The result is a finite value for C_∞^{UB} that depends on the precise geometry profile. Note that, since the $\{g_{nk}\}$ are normalized, they are scale independent: cell size is therefore immaterial in terms of the geometry profile. Note also that the schedulers that determine which user(s) in each cell are allocated to a given signaling resource play a significant role in establishing the geometry profile. Subject to latency and quality-of-service constraints, the schedulers can shape the geometry profile. Dynamic definition of the cooperation clusters is also likely to be beneficial [47], [48]. Other aspects that may affect the geometry profile include power control, vertical antenna patterns and antenna downtilting [49].

For the downlink, the roles of transmission and reception are reversed relative to the uplink. The corresponding performance of a fully cooperative system can be upper-bounded by allowing all the receivers to cooperate, in which case the uplink derivations carry over and a performance ceiling is readily observed. Tighter upper bounds might be obtained by removing the premise of receiver cooperation while considering downlink transmission strategies other than IID signaling.

V. CONCLUDING DISCUSSION

As argued up to this point, Relationships 2 or 3 with the appropriate values for $\mathsf{SIR}_1, \ldots, \mathsf{SIR}_N$ are the correct representation of a large cellular system, or any fragment thereof. The traditional high-power regime for user n, characterized by $\mathsf{SNR}_n \gg 1$, splits into two sub-regimes (cf. Fig. 2).

1) The DoF regime, where $\mathsf{SNR}_n \ll \mathsf{SIR}_n$. In this regime, the out-of-cluster interference is negligible relative to the noise and the spectral efficiency grows approximately linearly with $\log P$ according to the number of DoF computed without out-of-cluster interference. The notion of DoF remains approximately valid.

2) The saturation regime, where SNR_n is comparable to or greater than SIR_n. In this regime, the spectral efficiency chokes as it approaches C_∞. The notion of DoF becomes meaningless; more precisely, the actual number of DoF is revealed to be zero.

The transition between these regimes takes place, for user n, at the value of SNR_n corresponding to P_{sat}. (In pilot-assisted communication, such SNR_n essentially equals SIR_n.) If this transitional value of SNR_n is not explicitly known but the corresponding C_∞ is known, the former can be ascertained by establishing where an interference-free high-power expansion of the spectral efficiency intersects with the latter. It is pointless to operate at values of SNR_n much above this point.

It has also been shown that, in a fully cooperative system, the transitional SNR_n and the corresponding C_∞ depend on the skewness of the geometry profile and it has been pointed out that a process of scheduling can modify such skewness. Could scheduling modify the skewness to the point of restoring the unbounded scaling of the spectral efficiency with P? The answer, implicitly given by Example 8, is affirmative. Specifically, the example shows how the spectral efficiency upper bound can be arbitrarily high if less than L transmitters have nonnegligible powers, and scheduling can certainly ensure this by silencing all but $K' < L$ transmitters throughout the entire system. This is consistent with the result derived in [50] for a MIMO setting similar to (24), only with $\mathbf{H}$ having IID entries, whereby the best asymptotic ($P \to \infty$) scaling is achieved if no more than $L/2$ transmitters are activated; the

ensuing scaling is $C = L/4 \log_2 P + \mathcal{O}(1)$. This result can be extended to our setting, where the entries of $\mathbf{H}$ can have radically distinct variances.[1] The cost of achieving those $L/4$ DoF is to allow only $L/2$ users to transmit concurrently and within the same bandwidth. The remaining $K - L/2$ users in the system, however, cannot be denied access and thus they must be moved to separate orthogonal channels (in either time or in frequency). It follows that the existing time and bandwidth resources would have to be divided into $K/(L/2)$ orthogonal channels, each to accommodate a subset of $L/2$ users. Therefore, the asymptotic ($P \to \infty$) scaling achieved through this process would be

$$C = \frac{L/2}{K} \frac{L}{4} \log_2 P + \mathcal{O}(1) \qquad (38)$$

$$= \frac{L^2}{8K} \log_2 P + \mathcal{O}(1) \qquad (39)$$

which is optimal for $P \to \infty$, but highly inefficient at finite power levels. In fact, one of the main evolutionary axes of wireless networks has been the move to ever more aggressive resource reuse patterns, nowadays approaching universal reuse whereby every unit of bandwidth and every time slot is always active at every sector. The channelization required to ensure (39) would go a full circle, countering what cooperation set out to do in the first place: avoid bandwidth-wasting reuse patterns.

Altogether then, the limits of cooperation are delineated by the geometry profile. Mechanisms such as scheduling can modify the geometry profile subject to a clear tradeoff: more skewness pushes the saturation regime to higher power levels, at the expense of less bandwidth reuse, while less skewness causes the saturation to occur at lower powers but increases the bandwidth utilization. The guideline that emerges is that, to establish the fundamental limits at a given operating point in terms of P, the geometry profile should be adjusted so as to push the saturation just beyond such P.

A. Benefits of Cooperation

We hasten to emphasize that the points made in this paper do not nullify the benefits of cooperation but, rather, they evidence that cooperation has fundamental limitations that cannot be overcome through faster backhaul, more sophisticated signal processing, or any other technological advance. Under Relationships 2-3, cooperation can still yield a markedly higher C_∞ than if all interference was simply ignored. Similarly, cooperation can provide an increased slope within the DoF regime, allowing C_∞ to be approached at lower power levels.[2] This is demonstrated in the following example.

Example 10 *Consider Relationship 3 with $K = N = 3$ and with perfect CSI at the receivers. Let every BS and user have two antennas. Further let* $\mathsf{SNR}_n = \mathsf{SNR}$, *for $n = 1, 2, 3$, which embodies a widely studied canonical setting. Shown in Fig.*

---

[1]The fact that the same number of DoF holds regardless of the variances of the channel matrix, as long as these are not identically zero, contextualizes the extent to which the number of DoF is informative [51], [52].

[2]Recall that, in this regime, the slope is well approximated by the number of DoF with the out-of-cluster interference neglected.
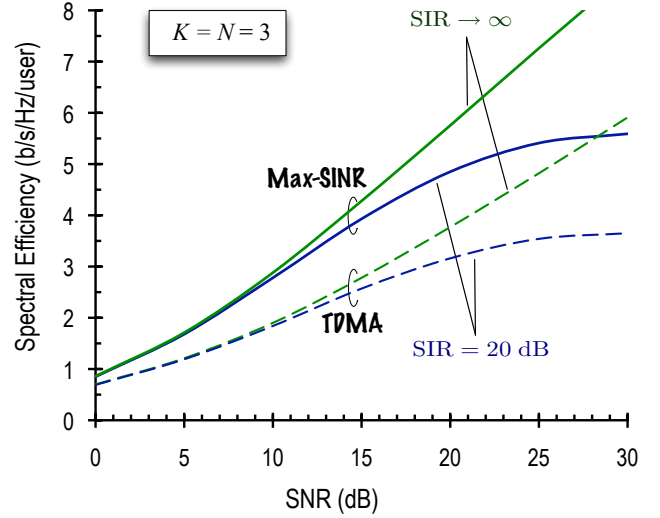

Fig. 7. Max-SINR v. TDMA utilizing Relationship 3 with $K = N = 3$ and with two antennas per transmitter and per receiver, for $\mathsf{SIR} \to \infty$ and $\mathsf{SIR} = 20$ dB.

7 are the spectral efficiencies achieved by distributed Max-SINR [24], [25] and by round-robin TDMA for both $\mathsf{SIR} \to \infty$ *(no out-of-cluster interference) and for* $\mathsf{SIR} = 20$ *dB. Note that, for* $\mathsf{SIR} \to \infty$, *Relationship 3 reverts to Relationship 1 and the theoretical number of DoF per user (1 with Max-SINR and $2/3$ with TDMA) are approached with growing* SNR. *For* $\mathsf{SIR} = 20$ *dB, the inflection that delineates the DoF and saturation regimes occurs around* $\mathsf{SNR} = 20$ *dB as expected.*

Apropos Example 10 with $\mathsf{SIR} = 20$ dB, we observe that the value of C_∞ is over 50% higher with Max-SINR relative to TDMA. Further, in the DoF regime—from about 5 to 20 dB—a substantial difference between Max-SINR and TDMA builds up. Altogether then, cooperation continues to provide a substantial advantage in this setting but the number of DoF is only a partial measure thereof.

B. Future Directions

A key insight of the present work is that interference in large wireless networks is made up of a multitude of terms, the majority of which are too minute to be tracked given that they are subject to independent fading. Beyond a few strong terms corresponding to nearby transmitters, the structure of the interference is too intricate to discern within the fading coherence yet the sum of all those minute terms is very substantial. Essentially then, the receivers are near-sighted. They can only focus on a few strong nearby transmitters, depending on the channel selectivity and system dynamics, and everything else in the distance looks fundamentally blurry. The result is a naturally clustered structure, with everything that ensues.

There is a certain sense of acceptance that near-sightedness is unavoidable, not only in the prevalence of clusters in the cooperation literature, but also in the more recent trend of studying partially connected networks [53]. The concept of partial connectedness can be interpreted as a quantization of the entries of the geometry profile to either zero or nonzero, a

simplification that facilitates the analysis while being faithful up to a certain SNR (beyond which the neglected terms would become relevant).

Undoubtedly, further work is needed to expand the observations in this paper, and we outline several directions of interest.

- The analysis focused on the uplink. The extension to the downlink is relatively straightforward in clustered systems with pilot-assisted transmission, but nontrivial for fully cooperative systems.
- Although the model did encompass MIMO transmissions, for the analysis only single-antenna transmitters and receivers were considered. It would be desirable to formalize the extension to MIMO, and especially to large-dimensional MIMO [54], [55]. Note that the upper bound in (25) grows sustainedly for $N \to \infty$ with fixed K, buttressing the potential of deploying base stations with massive numbers of antennas.
- Related to the above point, it would also be desirable to understand the tradeoffs between numbers of antennas at a single base station and cooperation of many base stations to gauge whether there is a fundamental preference for centralized or distributed antenna architectures [36].
- If a system is small in number of cells or if the clusters are highly isolated (e.g., through geographical separation, penetration losses, or millimeter-wave propagation losses), then P_{sat} can be large enough to render the saturation anecdotal. It would be useful to understand and classify systems based on the proximity of P_{sat} to operationally relevant values.
- The modeling in this paper also neglected propagation delay among distant units. Characterizing the impact of delay (as a function of bandwidth) on P_{sat} and C_∞ would likely further reduce both of these values, making them even more practically relevant.
- Determining how P_{sat} and C_∞ behave in heterogeneous networks with diverse cell sizes and user velocities would also be desirable, as these are the types of networks that are likely to predominate henceforth.

APPENDIX A
PROOF OF PROPOSITION 1

Using the chain rule, the mutual information between $\mathbf{X}$ and $\mathbf{Y}$ (in bits per coherence block) can be expressed as

$$I(\mathbf{X};\mathbf{Y}) = I(\mathbf{H}\mathbf{X};\mathbf{Y}) - I(\mathbf{Y};\mathbf{H}|\mathbf{X}). \tag{40}$$

Since the entries of $\mathbf{X}$ are zero-mean and those of $\mathbf{H}$ are independent, the entries of $\mathbf{H}\mathbf{X}$ are zero-mean. Furthermore, the variance of the (n,k)th entry of $\mathbf{H}\mathbf{X}$ is $\sum_{j=1}^{K} g_{nj}$ which, from (5), equals unity. It follows that the entries of $\mathbf{H}\mathbf{X}$ are zero-mean and unit-variance. The term $I(\mathbf{H}\mathbf{X};\mathbf{Y})$ is then upper-bounded by the value it would take if those entries were IID and $\mathcal{N}_{\mathbb{C}}(0,1)$. Thus,

$$I(\mathbf{X};\mathbf{Y}) \le \sum_{n=1}^{N} L \log_2(1 + \mathsf{SNR}_n) - I(\mathbf{Y};\mathbf{H}|\mathbf{X}). \tag{41}$$

Turning now our attention to $I(\mathbf{Y};\mathbf{H}|\mathbf{X})$, and denoting the differential entropy operator by $\mathfrak{h}(\cdot)$,

$$I(\mathbf{Y};\mathbf{H}|\mathbf{X}) = \mathfrak{h}(\mathbf{Y}|\mathbf{X}) - \mathfrak{h}(\mathbf{Y}|\mathbf{H},\mathbf{X}). \tag{42}$$

Conditioned on $\mathbf{X}$, the rows of $\mathbf{Y}$ are independent (both conditionally and unconditionally on $\mathbf{H}$) and thus both differential entropy terms in (42) can be computed row-wise and simply added. Hence,

$$I(\mathbf{Y};\mathbf{H}|\mathbf{X}) = \sum_{n=1}^{N} I(\mathbf{y}_n; \mathbf{h}_n|\mathbf{X}) \tag{43}$$

where $\mathbf{y}_n = \sqrt{\mathsf{SNR}_n}\, \mathbf{h}_n \mathbf{X} + \mathbf{z}_n$ with $\mathbf{h}_n$ and $\mathbf{z}_n$ the nth rows of $\mathbf{H}$ and $\mathbf{Z}$, respectively. Since $\mathbf{h}_n$ and $\mathbf{z}_n$ are complex Gaussian vectors,

$$I(\mathbf{Y};\mathbf{H}|\mathbf{X}) = \sum_{n=1}^{N} \mathbb{E}\left[\log_2 \det\left(\mathbf{I} + \mathsf{SNR}_n \mathbf{X}^\dagger \mathbf{G}_n \mathbf{X}\right)\right] \tag{44}$$

where, recall, $\mathbf{G}_n = \mathrm{diag}\{g_{n1}, \ldots, g_{nK}\}$.

Since $I(\mathbf{X};\mathbf{Y})$ increases monotonically with P, we concentrate on upper-bounding it for $P \to \infty$. Suppose first that we had $K < L$. Then, for large SNR_n, $n = 1, \ldots, N$, we would have

$$I(\mathbf{Y};\mathbf{H}|\mathbf{X}) = \sum_{n=1}^{N} \left(K \log_2 \mathsf{SNR}_n + \mathbb{E}\left[\log_2 \det\left(\mathbf{X}\mathbf{X}^\dagger \mathbf{G}_n\right)\right]\right) \\ + o(1) \tag{45}$$

and, from (41),

$$I(\mathbf{X};\mathbf{Y}) \le \sum_{n=1}^{N} \left((L - K) \log_2 \mathsf{SNR}_n \\ - \mathbb{E}\left[\log_2 \det\left(\mathbf{X}\mathbf{X}^\dagger \mathbf{G}_n\right)\right]\right) + o(1) \tag{46}$$

whose right-hand side grows unboundedly with P indicating that there is hope for $I(\mathbf{X};\mathbf{Y})$ to grow unboundedly with P.

Conversely, for $K \ge L$, and given the full-rank condition on $\mathbf{X}$,

$$I(\mathbf{Y};\mathbf{H}|\mathbf{X}) = \sum_{n=1}^{N} \left(L \log_2 \mathsf{SNR}_n + \mathbb{E}\left[\log_2 \det\left(\mathbf{X}^\dagger \mathbf{G}_n \mathbf{X}\right)\right]\right) \\ + o(1) \tag{47}$$

and thus (41) becomes, expanded with respect to P,

$$I(\mathbf{X};\mathbf{Y}) \le -\sum_{n=1}^{N} \mathbb{E}\left[\log_2 \det\left(\mathbf{X}^\dagger \mathbf{G}_n \mathbf{X}\right)\right] + o(1) \tag{48}$$

from which the upper bound in (25) follows.

APPENDIX B
PROOF OF PROPOSITION 2

If each user transmits an IID sequence with the symbols drawn from a common distribution, the entries of $\mathbf{X}$ are IID. It follows that, if the variances of the entries of $\mathbf{G}$ are uniformly bounded, the expression in Proposition 1 satisfies the conditions for the large-dimensional analysis in [46], [56].

Under these conditions, and couched in the notation of this paper, we have that [46, Section VI-A]

$$-\frac{1}{L}\mathbb{E}\left[\log_2 \det\left(\mathbf{X}^\dagger \mathbf{G}_n \mathbf{X}\right)\right] \to \log_2 e - \log_2\left(\sum_{k=1}^{K}\frac{g_{nk}}{1+a_{nk}}\right)$$
$$-\frac{1}{L}\sum_{k=1}^{K}\log_2(1+a_{nk}) \quad (49)$$

as $K, L \to \infty$ with $K \geq L$, and with a_{nk} the nonnegative solution to

$$a_{nk} = \frac{L\, g_{nk}}{\sum_{j=1}^{K}\frac{g_{nj}}{1+a_{nj}}}. \quad (50)$$

Defining

$$a_n = \frac{L}{\sum_{j=1}^{K}\frac{g_{nj}}{1+a_{nj}}} \quad (51)$$

we can write $a_{nk} = a_n g_{nk}$ and

$$-\frac{1}{L}\mathbb{E}\left[\log_2 \det\left(\mathbf{X}^\dagger \mathbf{G}_n \mathbf{X}\right)\right] \to \log_2 e + \log_2\left(\frac{a_n}{L}\right)$$
$$-\frac{1}{L}\sum_{k=1}^{K}\log_2(1+a_n g_{nk}) \quad (52)$$

with a_n the nonnegative solution to

$$\sum_{k=1}^{K}\frac{g_{nk}}{g_{nk}+1/a_n} = L. \quad (53)$$

It must be noted that, in a cellular system, the variances of the entries of $\mathbf{G}$ need not be uniformly bounded for $K, N \to \infty$. Hence, there are no guarantees of convergence to the right-hand side of (49) and the expression therein should in general be taken only as an approximation for finite K, N and L. All the tests conducted for finite but large values, e.g., Example 6, have indeed yielded very tight approximations to the exact expression in the left-hand-side of (49).

APPENDIX C
DETAILS OF EXAMPLES 3 AND 7

If the integers u and v denote the indices of a cell on the axes shown in the inset of Fig. 3, the cartesian coordinates of the corresponding BS are

$$x = \frac{3}{2}uR \quad (54)$$
$$y = \sqrt{3}\left(v+\frac{u}{2}\right)R. \quad (55)$$

Relative to its serving BS, the relative cartesian position of a user centered in azimuth in its sector and at distance $2R/3$ from the base is, for each of the sectors as labeled in Fig. 3,

$$\Delta x_1 = -\frac{R}{3} \qquad \Delta x_2 = -\frac{R}{3} \qquad \Delta x_3 = \frac{2R}{3} \quad (56)$$
$$\Delta y_1 = \frac{R}{\sqrt{3}} \qquad \Delta y_2 = -\frac{R}{\sqrt{3}} \qquad \Delta y_3 = 0. \quad (57)$$

Therefore, the distance between the BS at the origin ($u = v = 0$) and the users at the cell with indices u and v equals, for each of the sectors therein,

$$d_1(u,v) = R\sqrt{\left(\frac{3}{2}u - \frac{1}{3}\right)^2 + 3\left(v+\frac{u}{2}+\frac{1}{3}\right)^2} \quad (58)$$

$$d_2(u,v) = R\sqrt{\left(\frac{3}{2}u - \frac{1}{3}\right)^2 + 3\left(v+\frac{u}{2}-\frac{1}{3}\right)^2} \quad (59)$$

$$d_3(u,v) = R\sqrt{\left(\frac{3}{2}u + \frac{2}{3}\right)^2 + 3\left(v+\frac{u}{2}\right)^2}. \quad (60)$$

The corresponding normalized power gains are $g_1(u,v) = Dd_1^{-\gamma}$, $g_2(u,v) = Dd_2^{-\gamma}/Q$ and $g_3(u,v) = Dd_3^{-\gamma}/Q$, with D an appropriate constant and with Q the antenna front-to-back ratio. (Because of symmetries, applying the factor $1/Q$ to sectors 2 and 3 of each cell is equivalent to applying it to all three sectors in the slices of the system spanned by sectors 2 and 3 of the central reference cell.)

Then, for Example 3, we can compute (7) by having the appropriate terms in the numerator and denominator with $K \to \infty$.

For Example 7, D is the constant rendering

$$\sum_{u=-\infty}^{\infty}\sum_{v=-\infty}^{\infty}\left(g_1(u,v)+g_2(u,v)+g_3(u,v)\right) = 1 \quad (61)$$

which, for $\gamma = 3.8$, gives $D = 0.157R^{3.8}$. Then, a is the solution to

$$\sum_{u=-\infty}^{\infty}\sum_{v=-\infty}^{\infty}\left(\frac{g_1(u,v)}{g_1(u,v)+1/a} + \frac{g_2(u,v)}{g_2(u,v)+1/a}\right.$$
$$\left.+\frac{g_3(u,v)}{g_3(u,v)+1/a}\right) = L \quad (62)$$

while

$$C_\infty = \log_2\frac{ae}{L} - \frac{1}{L}\sum_{u=-\infty}^{\infty}\sum_{v=-\infty}^{\infty}\left[\log_2(1+a\,g_1(u,v))\right.$$
$$\left.+\log_2\left(1+ag_2(u,v)\right)+\log_2\left(1+ag_3(u,v)\right)\right]. \quad (63)$$

Numerical evaluation of (62) and (63) with $Q = 100$ yields the results in Example 7.

REFERENCES

[1] A. Lozano, J. G. Andrews, and R. W. Heath, "On the limitations of cooperation in wireless networks," in *Information Theory and Applications Workshop (ITA'12)*, 2012, pp. 123–130.
[2] A. Lozano, J. G. Andrews, and R. W. Heath, "Spectral efficiency limits in pilot-assisted cooperative communications," in *IEEE International Symposium on Information Theory (ISIT'12)*, 2012, pp. 1132–1136.
[3] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity. Part I. system description," *IEEE Transactions on Communications*, vol. 51, no. 11, pp. 1927 – 1938, nov 2003.
[4] N. Laneman, D. Tse, and G. Wornell, "Cooperative diversity in wireless networks: Efficient protocols and outage behavior," *IEEE Trans. on Inform. Theory*, vol. 50, no. 12, pp. 3062– 3080, Dec. 2004.

[5] M. Janani, A. Hedayat, T. E. Hunter, and A. Nosratinia, "Coded cooperation in wireless communications: space-time transmission and iterative decoding," *IEEE Transactions on Signal Processing*, vol. 52, no. 2, pp. 362 – 371, Feb. 2004.

[6] G. J. Foschini, K. Karakayali, and R. A. Valenzuela, "Enormous spectral efficiency of isolated multiple antenna links emerges in a coordinated cellular network," *IEE Proceedings - Communications*, vol. 153, no. 4, pp. 548–55, Aug. 2006.

[7] S. Venkatesan, A. Lozano, and R. Valenzuela, "Network MIMO: Overcoming intercell interference in indoor wireless systems," in *Asilomar Conference on Signals, Systems and Computers*, 2007, pp. 83–87.

[8] S. Venkatesan, H. Huang, A. Lozano, and R. Valenzuela, "A WiMAX-based implementation of network MIMO for indoor wireless systems," *EURASIP Journal on Advances in Signal Processing*, vol. 2009, pp. 9, Oct. 2009.

[9] H. Huang, M. Trivellato, A. Hottinen, M. Shafi, P. Smith, and R. Valenzuela, "Increasing downlink cellular throughput with limited network MIMO coordination," *IEEE Trans. on Wireless Communications*, vol. 8, no. 6, pp. 2983–2989, June 2009.

[10] S. W. Peters and R. W. Heath Jr., "User partitioning for less overhead in MIMO interference channels," *IEEE Trans. Wireless Communications (to appear), preprint at arXiv:1007.0512*, 2012.

[11] D. Gesbert, S. Hanly, H. Huang, S. Shamai, O. Simeone, and Wei Yu, "Multi-cell MIMO cooperative networks: A new look at interference," *IEEE Journal on Sel. Areas in Communications*, vol. 28, no. 9, pp. 1380–1408, Dec. 2010.

[12] P. Wang, H. Wang, L. Ping, and X. Lin, "On the capacity of MIMO cellular systems with base station cooperation," *IEEE Trans. on Wireless Communications*, vol. 10, no. 11, pp. 3720–3731, Nov. 2011.

[13] O. Simeone, N. Levy, A. Sanderovich, O. Somekh, B. M. Zaidel, H. V. Poor, and S. Shamai, "Information theoretic considerations for wireless cellular systems: The impact of cooperation," *Foundations and Trends in Communications and Information Theory*, vol. 7, 2012.

[14] O. Simeone, O. Somekh, H. V. Poor, and S. Shamai, "Local base station cooperation via finite-capacity links for the uplink of linear cellular networks," *IEEE Trans. Info. Theory*, vol. 55, no. 1, pp. 190–204, Jan. 2009.

[15] Coordinated multi-point operation for LTE physical layer aspects (Release 11), *3GPP TR 36.819*, 2011.

[16] A. Gorokhov, "Coordinated joint transmission in WWAN," *IEEE Communication Theory Workshop (CTW'10), available at "http://www.ieee-ctw.org/2010/mon/Gorokhov.pdf"*, May 2010.

[17] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H.-P. Mayer, L. Thiele, and V. Jungnickel, "Coordinated multipoint: Concepts, performance, and field trial results," *IEEE Communications Magazine*, vol. 49, no. 2, pp. 102–111, Feb. 2011.

[18] A. Barbieri, P. Gaal, S. Geirhofer, T. Ji, D. Malladi, Y. Wei, and F. Xue, "Coordinate downlink multi-point communications in heterogeneous 4G cellular networks," in *Information Theory and Applications Workshop (ITA)*, Feb. 2012.

[19] J. G. Andrews, H. Claussen, M. Dohler, S. Rangan, and M. C. Reed, "Femtocells: Past, present, and future," *IEEE Journal on Sel. Areas in Comm.*, vol. 30, no. 3, pp. 497–508, Apr. 2012.

[20] L. Zheng and D. Tse, "Diversity and multiplexing: A fundamental tradeoff in multiple antenna channels," *IEEE Trans. Inform. Theory*, vol. 49, pp. 1073–1096, May 2003.

[21] V. R. Cadambe and S. A. Jafar, "Interference alignment and the degrees of freedom for the K user interference channel," *IEEE Trans. Inform. Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.

[22] M. K. Karakayali, G. J. Foschini, and R. A. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Communications*, vol. 13, no. 4, pp. 56–61, Aug. 2006.

[23] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over MIMO X channels: Interference alignment, decomposition, and performance analysis," *IEEE Trans. Inform. Theory*, vol. 54, no. 8, pp. 3457–3470, Aug. 2008.

[24] V. R. Cadambe, K. Gomadam, and S. A. Jafar, "A distributed numerical approach to interference alignment and applications to wireless interference networks," *IEEE Trans. Inform. Theory*, vol. 57, no. 6, pp. 3309–3322, June 2011.

[25] S. W. Peters and R. W. Heath, Jr., "Algorithms for the MIMO interference channel," *submitted to IEEE Transactions on Signal Processing*, July 2009.

[26] Q. Shi, M. Razaviyayn, Z. Luo, and C. He, "An iteratively weighted MMSE approach to distributed sum-utility maximization for a MIMO interfering broadcast channel," *IEEE Trans. on Signal Processing*, vol. 59, no. 9, pp. 4331–4340, 2011.

[27] R. Tresch and M. Guillaud, "Clustered interference alignment in large cellular networks," in *Int'l Symp. on Personal, Indoor and Mobile Radio Communications (PIMRC'09)*, Sept. 2009, pp. 1024–1028.

[28] W. Mennerich and W. Zirwas, "Reporting effort for cooperative systems applying interference floor shaping," in *Int'l Symp. on Personal, Indoor and Mobile Radio Communications (PIMRC'11)*, Sept. 2011.

[29] R. Mungara, G. George, and A. Lozano, "System-level performance of distributed cooperation," *Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA*, Oct. 2012.

[30] D. Aziz, F. Boccardi, and A. Weber, "System-level performance study of interference alignment in cellular systems with base-station coordination," in *Int'l Symp. on Personal Indoor and Mobile Radio Communications (PIMRC'12)*, 2012, pp. 1155–1160.

[31] C. Suh, M. Ho, and D. N. C. Tse, "Downlink interference alignment," *IEEE Trans. on Communications*, vol. 59, no. 9, pp. 2616–2626, Sept. 2011.

[32] R. H. Etkin, D. N. C. Tse, and H. Wang, "Gaussian interference channel capacity to within one bit," *IEEE Trans. on Inform. Theory*, vol. 54, no. 12, pp. 5534–5562, 2008.

[33] S. A. Jafar and S. Vishwanath, "Generalized degrees of freedom of the symmetric Gaussian K user interference channel," *IEEE Trans. on Inform. Theory*, vol. 56, no. 7, pp. 3297–3303, 2010.

[34] G. Caire, S. A. Ramprashad, and H. C. Papadopoulos, "Rethinking network MIMO: cost of CSIT, performance analysis, and architecture comparisons," in *Information Theory and Applications Workshop (ITA)*, 2010, pp. 1–10.

[35] H. Huh, A. M. Tulino, and G. Caire, "Network MIMO with linear zero-forcing beamforming: Large system analysis, impact of channel estimation, and reduced-complexity scheduling," *IEEE Trans. on Inform. Theory*, vol. 58, no. 5, pp. 2911–2934, May 2012.

[36] H. Huh, G. Caire, H. C. Papadopoulos, and S. A. Ramprashad, "Achieving "massive MIMO" spectral efficiency with a not-so-large number of antennas," *IEEE Trans. on Wireless Communications*, vol. 11, no. 9, pp. 3226–3239, Sept. 2012.

[37] N. Jindal and A. Lozano, "A unified treatment of optimum pilot overhead in multipath fading channels," *IEEE Trans. on Communications*, vol. 58, no. 10, pp. 2939–2948, Oct. 2010.

[38] A. Lozano and D. Porrat, "Non-peaky signals in wideband fading channels: achievable bit rates and optimal bandwidth," *IEEE Trans. on Wireless Communications*, vol. 11, no. 1, pp. 246–257, Jan. 2012.

[39] B. Hassibi and B. M. Hochwald, "How much training is needed in multiple-antenna wireless links?," *IEEE Trans. Inform. Theory*, vol. 49, no. 4, pp. 951–963, Apr. 2003.

[40] T. L. Marzetta, "BLAST training: Estimating channel characteristics for high capacity space-time wireless," *Proc. of 37th Annual Allerton Conf. on Communication, Control and Computing, Monticello, IL*, vol. 37, pp. 958–966, Sept. 1999.

[41] A. Lozano, "Interplay of spectral efficiency, power and Doppler spectrum for reference-signal-assisted wireless communication," *IEEE Trans. Communications*, vol. 56, no. 12, Dec. 2008.

[42] A. Lapidoth and S. Shamai, "Fading channels: How perfect need 'perfect side information' be?," *IEEE Trans. Inform. Theory*, vol. 48, no. 5, pp. 1118–1134, May 2002.

[43] G. Caire, N. Jindal, M. Kobayashi, and N. Ravindran, "Multiuser MIMO achievable rates with downlink training and channel state feedback," *IEEE Trans. on Inform. Theory*, vol. 56, no. 6, pp. 2845–2866, 2010.

[44] F. Rusek, A. Lozano, and N. Jindal, "Mutual information of IID complex Gaussian signals on block Rayleigh-faded channels," *IEEE Trans. Inform. Theory*, vol. 58, no. 1, Jan. 2012.

[45] A. Lozano, A. M. Tulino, and S. Verdu, "High-SNR power offset in multiantenna communication," *IEEE Trans. Inform. Theory*, vol. 51, no. 12, pp. 4134–4151, Dec. 2005.

[46] A. M. Tulino, A. Lozano, and S. Verdú, "Impact of antenna correlation on the capacity of multiantenna channels," *IEEE Trans. Inform. Theory*, vol. 51, no. 7, pp. 2491–2509, Aug. 2005.

[47] A. Papadogiannis, D. Gesbert, and E. Hardouin, "A dynamic clustering approach in wireless networks with multi-cell cooperative processing," in *IEEE Int'l Conf. on Communications (ICC'08)*, 2008, pp. 4033–4037.

[48] I. D. Garcia, N. Kusashima, K. Sakaguchi, and K. Araki, "Dynamic cooperation set clustering on base station cooperation cellular networks," in *IEEE Int'l Symp. on Personal, Indoor and Mobile Radio Communications (PIMRC'10)*, 2010, pp. 2127–2132.

[49] L. Thiele, T. Wirth, K. Brner, M. Olbrich, V, Jungnickel, J. Rumold, and S. Fritze, "Modeling of 3D field patterns of downtilted antennas

and their impact on cellular systems," *International ITG Workshop on Smart Antennas (WSA 2009)*, Feb. 2009.

[50] L. Zheng and D. N. C. Tse, "Communication on the Grassman manifold: A geometric approach to the non-coherent multiple-antenna channel," *IEEE Trans. Inform. Theory*, vol. 48, no. 2, pp. 359–383, Feb. 2002.

[51] A. Lozano and N. Jindal, "Transmit diversity vs. spatial multiplexing in modern MIMO systems," *IEEE Trans. on Wireless Communications*, vol. 9, no. 1, pp. 186–197, Jan. 2010.

[52] A. Lozano and N. Jindal, "Are yesterday's information-theoretic fading models and performance metrics adequate for the analysis of today's wireless systems?," *IEEE Communications Magazine*, vol. 50, no. 11, pp. 210–217, 2012.

[53] M. Guillaud and D. Gesbert, "Interference alignment in partially connected interfering multiple-access and broadcast channels," in *IEEE Global Telecommunications Conference (GLOBECOM 2011)*, 2011, pp. 1–5.

[54] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. Wireless Communications*, vol. 9, no. 11, pp. 3590 –3600, Nov. 2010.

[55] A. Ashikhmin and T. L. Marzetta, "Pilot contamination precoding in multi-cell large scale antenna systems," in *Int. Symp. on Inform. Theory (ISIT'12)*, 2012, pp. 1137–1141.

[56] J. W. Silverstein and Z. D. Bai, "On the empirical distribution of eigenvalues of a class of large dimensional random matrices," *Journal Multivariate Analysis*, vol. 54, pp. 175–192, 1995.

Angel Lozano Angel Lozano is a Professor of Information and Communication Technologies at Universitat Pompeu Fabra (UPF) in Barcelona, Spain, where he teaches and conducts research on wireless communications. He is also UPF's Vice-Rector for Research. Prof. Lozano received the Master of Science and Ph.D. degrees in Electrical Engineering from Stanford University, USA, in 1994 and 1998, respectively. In 1999, he joined Bell Labs (Lucent Technologies, now Alcatel-Lucent) in Holmdel, USA, where he was a member of the Wireless Communications Research Department until 2008. Between 2005 and 2008 he was also an Adjunct Associate Professor of Electrical Engineering at Columbia University, NY, USA.

Prof. Lozano is an associate editor for the IEEE Transactions on Information Theory (since 2011), a former editor for the IEEE Transactions on Communications (1999-2009) and the Journal of Communications & Networks (2010-2012), has guest-edited various journal special issues, and is actively involved in committees and conference organization tasks for the IEEE. In particular, he is the Chair of the IEEE Communication Theory Technical Committee (2013-2014) and was elected to the Board of Governors of the IEEE Communications Society (2012-2014). He has further participated in standardization activities for 3GPP, 3GPP2, IEEE 802.20 and the IETF.

Prof. Lozano has published extensively, holds 15 patents, and has contributed to several books. His papers have received two awards: the best paper at the 2006 IEEE Intl Symposium on Spread Spectrum Techniques & Applications, and the Stephen O. Rice prize to the best paper published in the IEEE Transactions on Communications in 2008. He also received the Bell Labs Presidents Gold Award in 2002 and the ICREA Academia Award from the Catalan Government in 2011.

Robert W. Heath Jr. Robert W. Heath Jr. received the B.S. and M.S. degrees from the University of Virginia, Charlottesville, VA, in 1996 and 1997 respectively, and the Ph.D. from Stanford University, Stanford, CA, in 2002, all in electrical engineering. From 1998 to 2001, he was a Senior Member of the Technical Staff then a Senior Consultant at Iospan Wireless Inc, San Jose, CA where he worked on the design and implementation of the physical and link layers of the first commercial MIMO-OFDM communication system. Since January 2002, he has been with the Department of Electrical and Computer Engineering at The University of Texas at Austin where he is a Professor and Director of the Wireless Networking and Communications Group. He is also President and CEO of MIMO Wireless Inc. and Chief Innovation Officer at Kuma Signals LLC. His research interests include several aspects of wireless communication and signal processing: limited feedback techniques, multihop networking, multiuser and multicell MIMO, interference alignment, adaptive video transmission, manifold signal processing, and millimeter wave communication techniques.

Dr. Heath has been an Editor for the IEEE Transactions on Communication, an Associate Editor for the IEEE Transactions on Vehicular Technology, lead guest editor for an IEEE Journal on Selected Areas in Communications special issue on limited feedback communication, and lead guest editor for an IEEE Journal on Selected Topics in Signal Processing special issue on Heterogenous Networks. He currently serves on the steering committee for the IEEE Transactions on Wireless Communications. He was a member of the Signal Processing for Communications Technical Committee in the IEEE Signal Processing Society. Currently he is the Chair of the IEEE COMSOC Communications Technical Theory Committee. He was a technical co-chair for the 2007 Fall Vehicular Technology Conference, general chair of the 2008 Communication Theory Workshop, general co-chair, technical co-chair and co-organizer of the 2009 IEEE Signal Processing for Wireless Communications Workshop, local co-organizer for the 2009 IEEE CAMSAP Conference, technical co-chair for the 2010 IEEE International Symposium on Information Theory, the technical chair for the 2011 Asilomar Conference on Signals, Systems, and Computers, general chair for the 2013 Asilomar Conference on Signals, Systems, and Computers, general co-chair for the 2013 IEEE GlobalSIP conference, and is technical co-chair for the 2014 IEEE GLOBECOM conference.

Dr. Heath was a co-author of best student paper awards at IEEE VTC 2006 Spring, WPMC 2006, IEEE GLOBECOM 2006, IEEE VTC 2007 Spring, and IEEE RWS 2009, as well as co-recipient of the Grand Prize in the 2008 WinTech WinCool Demo Contest. He was co-recipient of the 2010 EURASIP Journal on Wireless Communications and Networking best paper award and the 2012 IEEE Signal Processing Society Magazine Best Paper Award. He was a 2003 Frontiers in Education New Faculty Fellow. He is the recipient of the David and Doris Lybarger Endowed Faculty Fellowship in Engineering, is a licensed Amateur Radio Operator, and is a registered Professional Engineer in Texas.

Jeffrey G. Andrews Jeffrey Andrews (S'98, M'02, SM'06, F'13) received the B.S. in Engineering with High Distinction from Harvey Mudd College in 1995, and the M.S. and Ph.D. in Electrical Engineering from Stanford University in 1999 and 2002, respectively. He is a Professor in the Department of Electrical and Computer Engineering at the University of Texas at Austin, where he was the Director of the Wireless Networking and Communications Group (WNCG) from 2008-12. He developed Code Division Multiple Access systems at Qualcomm from 1995-97, and has consulted for entities including Verizon, the WiMAX Forum, Intel, Microsoft, Apple, Samsung, Clearwire, Sprint, and NASA. He is a member of the Technical Advisory Boards of Accelera and Fastback Networks.

Dr. Andrews is co-author of two books, Fundamentals of WiMAX (Prentice-Hall, 2007) and Fundamentals of LTE (Prentice-Hall, 2010), and holds the Earl and Margaret Brasfield Endowed Fellowship in Engineering at UT Austin, where he received the ECE departments first annual High Gain award for excellence in research. He is a Distinguished Lecturer for the IEEE Vehicular Technology Society, served as an associate editor for the IEEE Transactions on Wireless Communications from 2004-08, was the Chair of the 2010 IEEE Communication Theory Workshop, and is the Technical Program co-Chair of ICC 2012 (Comm. Theory Symposium) and Globecom 2014. He is an elected member of the Board of Governors of the IEEE Information Theory Society and an IEEE Fellow.

Dr. Andrews received the National Science Foundation CAREER award in 2007 and has been co-author of five best paper award recipients, two at Globecom (2006 and 2009), Asilomar (2008), the 2010 IEEE Communications Society Best Tutorial Paper Award, and the 2011 Communications Society Heinrich Hertz Prize. His research interests are in communication theory, information theory, and stochastic geometry applied to wireless cellular and ad hoc networks.