How Typical is "Typical"?
Characterizing Deviations Using the Meta Distribution

Martin Haenggi

University of Notre Dame, IN, USA

SpaSWiN 2017 Keynote

May 19, 2017

Available at http://www.nd.edu/~mhaenggi/talks/spaswin17.pdf

Growth of articles on IEEE Xplore over 1.5 decades

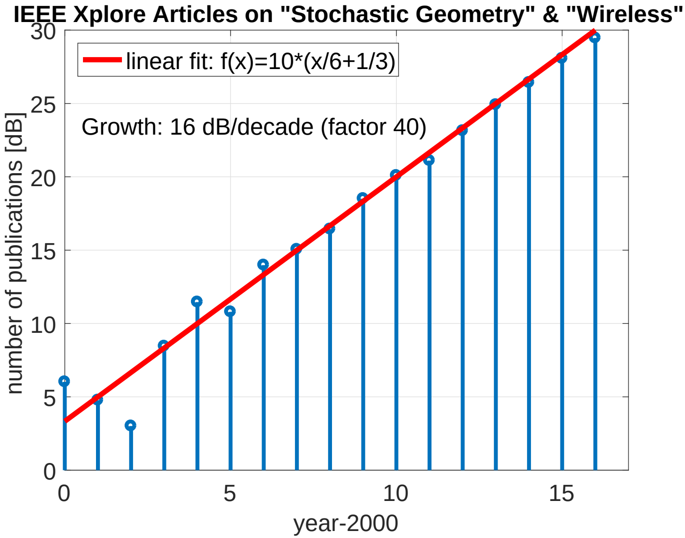

Menu

Overview

- What is "typical"?
- From spatial averages to the meta distribution
- Poisson bipolar networks
- Poisson cellular networks
  - Downlink
  - Uplink
  - D2D
  - What is coverage?
- Spatial outage capacity
- Ergodic spectral efficiency
- Conclusions

The typical person

The typical French person (in numbers)

- Lives 82 years.
- Makes USD 29,759 per year (disposable income).
- Lives in a household whose wealth is USD 53,851.
- Lives in 1.8 rooms.

The typical American person (in numbers)

- Lives 79 years.
- Makes USD 41,071 per year (disposable income).
- Lives in a household whose wealth is USD 163,268.
- Lives in 2.4 rooms.

The (stereo)typical French person

The (stereo)typical American tourist

The globally typical person

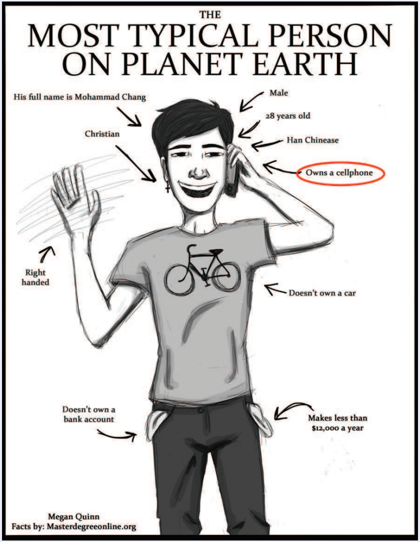

The typical user!

Inequality

Inequality in France

Social inequality (household income and wealth): 4.67.

1% income: \$242,000.
0.1% income: \$720,000.
0.01% income: \$2,252,000.

Top 20% earns 5× as much as the bottom 20%.

Inequality in the USA

Social inequality (household income and wealth): 8.19.

1% income: \$465,000.
0.1% income: \$1,695,000.
0.01% income: \$9,141,000.

Top 20% earns 8× as much as the bottom 20%.

Typicality and inequality

- In many situations, merely considering the typical entity reveals only limited information.
- The satisfaction of people in a country depends as much (or more) on the inequality than the absolute level of income, wealth, number of rooms, etc.
- Similarly, in a cellular network, whether a user is happy or not with 1 Mb/s strongly depends on whether other users get 100 kb/s or 10 Mb/s.
- Industry often focuses on the performance of the "5% user", which is the performance that the top 95% of the users experience. Increasing the performance of the typical user may decrease the performance of the 5% user.

Stochastic geometry: From spatial averages to the meta distribution

Spatial/ensemble average

Let Φ be a point process and $f : \mathcal{N} \to \mathbb{R}^+$ a performance function.
Ensemble average:

$$\bar{f} = \mathbb{E}^o f(\Phi)$$

- User at o is the typical user. In an ergodic setting, its performance is the performance averaged over all users (spatial average).
- But in a realization of Φ, no user is typical. All users have (much) better/worse performance than $\bar{f}$.
- To quantify this, we need to calculate other properties of the random variable $f(\Phi)$ than just its mean.

Spatial distribution

Refined analysis:

$$\bar{F}(x) = \mathbb{E}^o \mathbf{1}(f(\Phi) > x) \equiv \mathbb{P}^o(f(\Phi) > x)$$

- In an ergodic setting, this yields the fraction of users that achieve performance at least x (spatial distribution). This is more informative.
- This is still an average (but so is any distribution of any random variable), and it is again evaluated at the typical user.
- Of course, $\bar{f} = \int \bar{F}(x)\mathrm{d}x$.
- *Key example in wireless networks:* The SIR distribution.

SIR distribution

Let $f_\theta(\Phi) = \mathbf{1}(\mathrm{SIR}(\Phi) > \theta)$. The SIR distribution at the typical user is

$$\mathbb{P}^o(\mathrm{SIR}(\Phi) > \theta) \equiv \mathbb{E}^o f_\theta(\Phi),$$

which is a family of spatial averages since the performance function has an extra parameter θ. It yields the fraction of users that achieve an SIR of θ. It is also interpreted as the success probability of the typical user.

Does this give us complete information on the network performance, such as the performance of the 5% user?

Interpretation of the SIR distribution

Let us assume that $\mathbb{P}^o(\mathrm{SIR}(\Phi) > 1/10) = 0.95$ and consider a realization of Φ representing the node locations for a certain period of time.

- 95% of the users achieve an SIR of -10 dB, at any given time.
- However, the set of users that achieve this SIR changes in each coherence interval. Hence each user is likely to belong to the bottom 5% and to the top 95% many times in a short period. (This is why the SIR distribution is not the coverage—details to follow.)
- Moreover, this does not mean that an individual user achieves -10 dB SIR 95% of the time.
- In fact, nothing can be said about the SIR at an individual user.

  It could be that for each group of 100 users, 5 never achieve -10 dB (100% outage for 5 users, 0% outage for the rest).
  Or it could be that the 5 who do not achieve -10 dB are picked uniformly at random every 10ms (5% outage for all users.)

How to get more information?

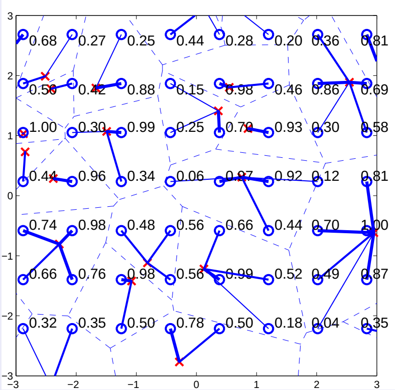

- To capture user performance, we need to adopt a longer-term viewpoint. This way, we can talk consistently about the 5% user.
- This means we need to average over the fading (and channel access).
- So lets assign to each user a personal SIR distribution (success probability):

$$\mathbb{P}(\mathrm{SIR}_{\mathrm{u}}(\Phi) > \theta \mid \Phi)$$

SIR distribution: $\mathbb{E}^o(\mathbb{P}(\mathrm{SIR}(\Phi) > \theta \mid \Phi))$

The meta distribution of the SIR

SIR distribution: $\mathbb{E}^o(\mathbb{P}(\mathsf{SIR}(\Phi) > \theta \mid \Phi))$

So the SIR distribution is just the spatial average of the conditional success probability random variable

$$P_{\mathrm{s}}(\theta) \triangleq \mathbb{P}(\mathsf{SIR}_o(\Phi) > \theta \mid \Phi_o).$$

But, as before, instead of considering only the average, let's consider the distribution!

The meta distribution of the SIR is the ccdf ⟨Haenggi, 2016⟩

$$\bar{F}(\theta, x) = \bar{F}_{P_{\mathrm{s}}(\theta)}(x) \triangleq \mathbb{P}^o(P_{\mathrm{s}}(\theta) > x), \quad \theta \in \mathbb{R}^+,\ x \in [0,1].$$

$\bar{F}(\theta, x)$ is the fraction of users that achieve an SIR of θ with probability at least x, in each realization of Φ. Those users do not change over time.

Meta distribution of the SIR

Using the previous notation, we have $f_\theta(\Phi) = \mathbb{P}(\mathsf{SIR}_o > \theta \mid \Phi)$ and

$$\bar{F}(\theta, x) = \mathbb{E}^o \mathbf{1}(f_\theta(\Phi) > x) = \mathbb{E}^o \mathbf{1}\Big(\mathbb{E}\mathbf{1}(\mathsf{SIR}_o(\Phi) > \theta \mid \Phi_o) > x\Big).$$

It is the distribution of the conditional SIR distribution, hence the term "meta". The standard SIR distribution (mean success probability) is

$$p_s(\theta) = \mathbb{P}^o(\mathsf{SIR} > \theta) = \int_0^1 \bar{F}(\theta, x)\mathrm{d}x.$$

Performance of the 5% user:
Rate (spectral efficiency) is determined by θ, e.g., through $\log(1+\theta)$.
The 5% user achieves the rate-reliability trade-off pairs (θ, x) given by $\bar{F}(\theta, x) = 0.95$.

Example contour plot of $\bar{F}(\theta, x)$: Trading off rate and reliability

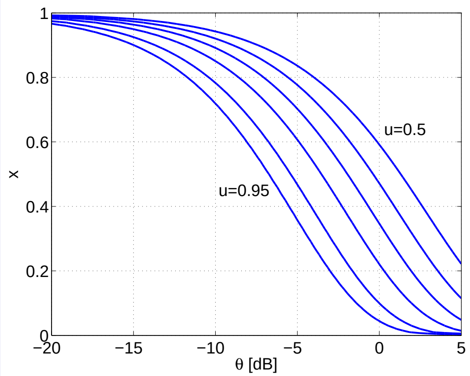

Contours $\bar{F}(\theta, x) = u$ for $u \in \{0.5, 0.6, 0.7, 0.8, 0.9, 0.95\}$.

The bottom curve $u = 0.95$ gives the performance of the 5% user.

For example, this user achieves an SIR of -10 dB with reliability 0.72 or an SIR of -4.3 dB with reliability 0.3.

Remarks on the meta distribution

- The standard SIR distribution (mean success probability) is an expectation over the point process and the fading, treating both sources of uncertainty the same.
- The meta distribution separates fading (time averaging) and location (spatial averaging). This makes sense since the coherence time of the large-scale path loss is much longer than that of the small-scale fading.
- For stationary and ergodic Φ, the ccdf of P_{s} can be alternatively written as the limit

$$\bar{F}_{P_{\mathrm{s}}(\theta)}(x) = \lim_{r\to\infty} \frac{1}{\lambda p \pi r^2} \sum_{\substack{y\in\Phi \\ \|y\|<r}} \mathbf{1}(\mathbb{P}(\mathrm{SIR}_{\tilde{y}} > \theta \mid \Phi) > x),$$

where $\tilde{y}$ is the receiver of transmitter y.

How to determine the meta distribution

A direct calculation does not seem feasible, but we may be able to calculate the moments

$$M_b \triangleq \mathbb{E}^o(P_\mathrm{s}(\theta)^b), \quad b \in \mathbb{C}.$$

- Even just M_2 is valuable, since the variance is a first important step towards characterizing the discrepancies between the users, i.e.,

$$\mathbb{E}^o\left((P_\mathrm{s}(\theta) - p_\mathsf{s}(\theta))^2\right) = M_2 - M_1^2,$$

  and we can use standard bounding techniques.
- If we know M_{jt}, $j \triangleq \sqrt{-1}$, $t \in \mathbb{R}^+$, we can use the Gil-Pelaez theorem to determine the entire distribution exactly!

$$\bar{F}(\theta, x) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \frac{\Im(e^{-jt\log x} M_{jt})}{t}\mathrm{d}t.$$

Prior work

- M_1 has a very rich history, see, *e.g.*, ⟨Zorzi and Pupolin, 1995; Baccelli et al., 2006; Andrews et al., 2011; Mukherjee, 2012; Nigam et al., 2014; DiRenzo, 2015; Deng et al., 2015; Madhusudhanan et al., 2014⟩.
- M_{-1} is the mean local delay ⟨Baccelli and Blaszczyszyn, 2010; Haenggi, 2013⟩. It is the mean number of transmission attempts needed until success.

  Conditioned on Φ, the transmission success events are independent and occur with probabality $P_{\rm s}(\theta)$. Hence the conditional local delay is geometric and $\mathbb{E}(D) = \mathbb{E}(P_{\rm s}(\theta)^{-1})$.
- For Poisson bipolar networks (without MAC), M_b, $b \geq 0$, is derived in ⟨Ganti and Andrews, 2010⟩.
- M_k, $k \in \mathbb{N}$, is the joint success probability of succeeding k times in a row, which is calculated for Poisson bipolar and cellular networks in ⟨Haenggi and Smarandache, 2013⟩ and ⟨Zhang and Haenggi, 2014a⟩.

Poisson bipolar networks

Conditional success probabilities

For a realization of a Poisson bipolar network, attach to each link the probability

$$P_{\mathrm{s}}^{(x)}(\theta) \triangleq \mathbb{P}(\mathrm{SIR}_x > \theta \mid \Phi, \mathrm{tx}),$$

taken over fading and ALOHA.

$P_{\mathrm{s}}^{(x)}(\theta)$ are random variables that capture the individual link performance.

Alternative interpretation of $P_{\mathrm{s}}(\theta)$ (thanks to Steven Weber): If node x has full knowledge of Φ, $P_{\mathrm{s}}^{(x)}(\theta)$ is its estimated link success probability.

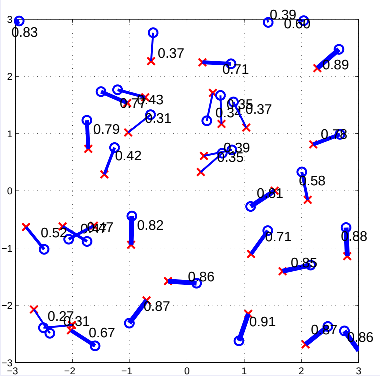

The histogram of all $P_{\mathrm{s}}^{(x)}$ gives very fine-grained information about the network performance.

Conditional success probability histograms

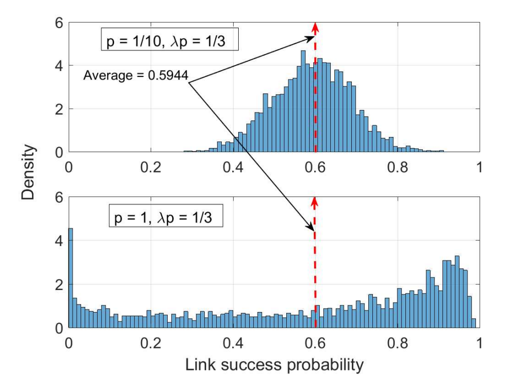

The mean (standard success probability) M_1 is the same for all pairs (λ, p) with the same λp. For Rayleigh fading, we have the well-known result $\mathbb{P}(\mathrm{SIR} > \theta) = \exp\left(-\frac{\lambda p \pi r^2 \theta^\delta}{\operatorname{sinc}\delta}\right)$, where $\delta = 2/\alpha$.

But the disparity between the links depends strongly on both λ and p.

Example (Meta distribution for Poisson bipolar network with ALOHA)

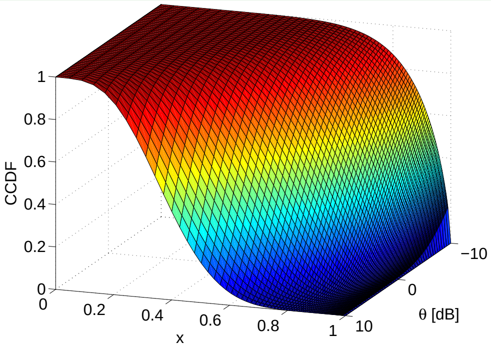

$$\lambda = 1,\ p = 1/4,\ \alpha = 4, \text{ and } r = 1/2$$

For each realization of Φ, the meta distribution yields the fraction of links that achieve a success probability at least x.

Moments for Poisson bipolar networks with ALOHA

The moments $M_b \triangleq \mathbb{E}(P_{\mathrm{s}}(\theta)^b)$, $b \in \mathbb{C}$, are given by ⟨Haenggi, 2016⟩

$$M_b = \exp\left(-\lambda\pi r^2\theta^\delta\Gamma(1-\delta)\Gamma(1+\delta)D_b(p,\delta)\right), \quad b \in \mathbb{C},$$

where

$$D_b(p,\delta) = pb\,{}_2F_1(1-b, 1-\delta; 2; p).$$

- $M_1 = \int_0^1 \bar{F}(\theta, x)\mathrm{d}x$ is the standard success probability.
- The variance var $P_{\mathrm{s}}(\theta) = M_1^2(M_1^{p(\delta-1)} - 1)$ quantifies the link disparity and yields the concentration result

$$\lim_{\substack{p\to 0\\ \lambda p=\tau}} P_{\mathrm{s}}(\theta) = M_1.$$

- Using M_{jt}, $t \in \mathbb{R}$, Gil-Pelaez inversion gives an integral expression of the exact meta distribution.

Classical bounds

For $x \in [0,1]$, the ccdf $\bar{F}_{P_\mathrm{s}}$ is bounded as

$$1 - \frac{\mathbb{E}^o((1-P_\mathrm{s}(\theta))^b)}{(1-x)^b} < \bar{F}_{P_\mathrm{s}}(x) \le \frac{M_b}{x^b}, \quad b > 0.$$

Illustrations for $\theta = 1$, $r = 1/2$ and $\lambda p = 1/4 \Rightarrow p_\mathrm{s} = M_1 = 0.735$:

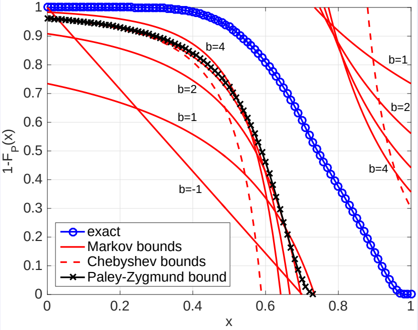

$\lambda = 1,\ p = 1/4,\ \mathrm{var}(P_\mathrm{s}) = 0.0212$

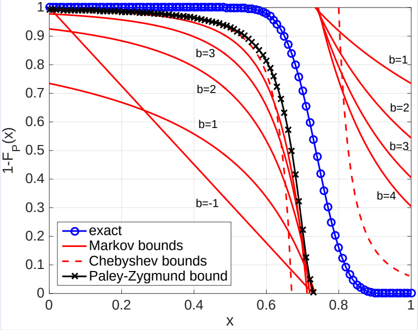

$\lambda = 5,\ p = 1/20,\ \mathrm{var}(P_\mathrm{s}) = 0.00418$

Best bounds using M_1 through M_4

Illustrations for $\alpha = 4$, $\theta = 1$, $r = 1/2$, and $p = 1/2$:

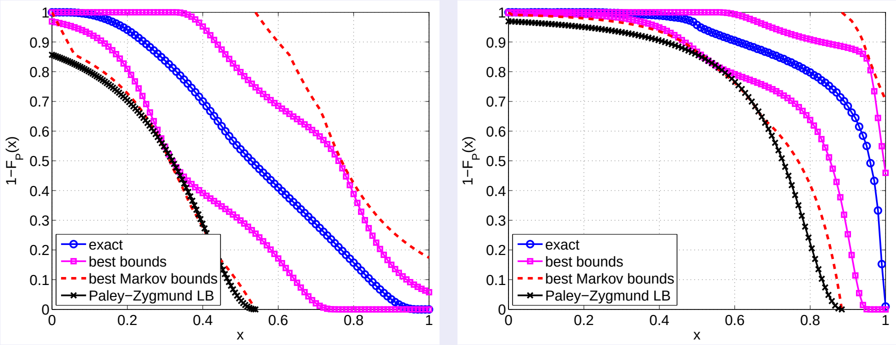

$\lambda = 1 \Rightarrow p_s = 0.54$, $\mathrm{var}(P_s) = 0.049$ $\qquad$ $\lambda = 1/5 \Rightarrow p_s = 0.88$, $\mathrm{var}(P_s) = 0.024$

Approximation with beta distribution

Since P_s is supported on $[0, 1]$, it is natural to approximate it as a beta random variable.

Beta approximation of the meta distribution

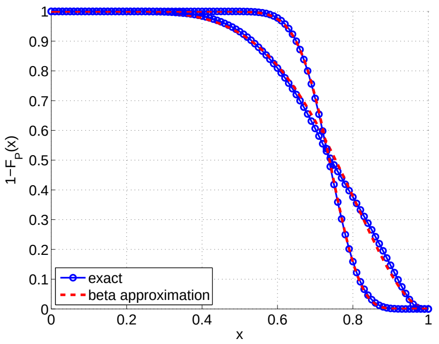

Exact ccdf and beta approximation for $\theta = 1$, $r = 1/2$, $\alpha = 4$, and $\lambda p = 1/4$.

Two cases: (1) $\lambda = 1$, $p = 1/4 \quad \rightarrow \text{var}\, P_{\rm s} = 0.02$.
(2) $\lambda = 5$, $p = 1/20 \rightarrow \text{var}\, P_{\rm s} = 0.004$.

For both cases, $M_1 = 0.735$. The standard analysis does not distinguish between the two networks.

Cellular networks

Downlink Poisson cellular networks

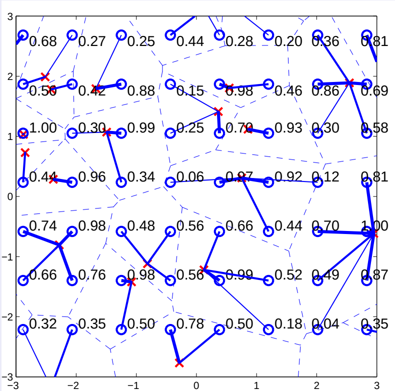

$\alpha = 4,\ \theta = 1$

Base stations (BSs) form a homogeneous Poisson point process (PPP) Φ of density λ.

A user connects to nearest BS, while all others interfere.

The received power at user u is $S_u = h_u \|x_u - u\|^{-\alpha}$, where $x_u = \text{argmin}\{x \in \Phi \colon \|x - u\|\}$ and h_u is exponential.

For each user u, calculate $\quad P_{\mathrm{s}}^{(u)} = \mathbb{P}(\mathrm{SIR}_u > \theta \mid \Phi) = \mathbb{E}_h \mathbf{1}(\mathrm{SIR}_u > \theta)$.

Basic result for downlink

For the typical user, we have ⟨Andrews et al., 2011⟩

$$p_{\mathsf{s}}(\theta) \triangleq \mathbb{P}(\mathsf{SIR} > \theta) = \bar{F}_{\mathsf{SIR}}(\theta) = \frac{1}{{}_2F_1(1, -\delta; 1-\delta; -\theta)}, \quad \delta \triangleq 2/\alpha.$$

$$\text{For } \delta = 1/2 \ (\alpha = 4): \quad p_{\mathsf{s}}(\theta) = \frac{1}{1 + \sqrt{\theta} \arctan \sqrt{\theta}}$$

Remarkably, the same result holds for a multi-tier Poisson model (HIP model), where each tier can have a different density and transmit power ⟨Nigam et al., 2014; Madhusudhanan et al., 2016⟩.

Focusing on the user at o, we are interested in the meta distribution

$$\bar{F}(\theta, x) = \bar{F}_{P_{\mathrm{s}}(\theta)}(x) \triangleq \mathbb{P}(P_{\mathrm{s}}(\theta) > x), \quad \theta \in \mathbb{R}^+, \ x \in [0,1],$$

and the moments $M_b \triangleq \mathbb{E}(P_{\mathrm{s}}(\theta)^b)$. Again $M_1 = p_{\mathsf{s}}$.

Example (PPP, Rayleigh fading, $\alpha = 4$)

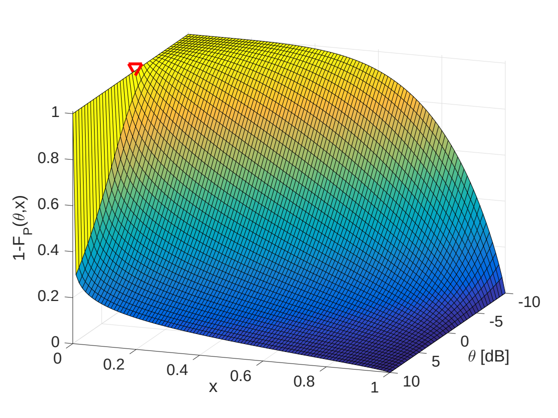

$$\bar{F}(\theta, x) = \mathbb{P}(P_{\mathrm{s}}(\theta) > x)$$

$=$ Fraction of users who achieve an SIR of θ with probability at least x.

Theorem (Moments of P_{s} for Rayleigh fading ⟨Haenggi, 2016⟩)

For Poisson cellular networks with nearest-BS association and Rayleigh fading,

$$M_b = \frac{1}{{}_2F_1(b, -\delta; 1-\delta; -\theta)}, \quad b \in \mathbb{C}.$$

Remark

Alternatively,

$$M_b = \frac{(1+\theta)^b}{{}_2F_1(b, 1; 1-\delta; \theta/(1+\theta))}.$$

This way, we can write the hypergeometric function as a series and obtain

$$M_b = \left[(1-z)^b \sum_{n=0}^{\infty} \frac{b(b+1)\cdot\ldots\cdot(b+n-1)}{(1-\delta)\cdot\ldots\cdot(n-\delta)} z^n\right]^{-1},$$

where $z = \theta/(1+\theta) < 1$.

M_1, M_2, and variance

$$M_1 = \left[(1-z)\sum_{n=0}^{\infty} \frac{n!}{(1-\delta)\cdot\ldots\cdot(n-\delta)} z^n\right]^{-1}$$

$$M_2 = \left[(1-z)^2\sum_{n=0}^{\infty} \frac{(n+1)!}{(1-\delta)\cdot\ldots\cdot(n-\delta)} z^n\right]^{-1}$$

Letting $g_n(\delta)$ the polynomial of order n with roots $[n]$ and $g_n(0)=1$, i.e.,

$$g_n(\delta) \triangleq \frac{(1-\delta)\cdot\ldots\cdot(n-\delta)}{n!},$$

we have

$$\operatorname{var}(P_{\mathrm{s}}) = \frac{1}{(1-z)^2}\left[\frac{1}{\sum \frac{n+1}{g_n(\delta)} z^n} - \frac{1}{\left(\sum \frac{1}{g_n(\delta)} z^n\right)^2}\right].$$

This can be used to obtain rational approximations and bounds.

Proof (Moments M_b)

Let $x_0 = \operatorname{argmin}\{x \in \Phi\colon \|x\|\}$ be the serving BS. Given the BS process Φ, the success probability is

$$P_{\mathrm{s}}(\theta) = \mathbb{P}\Big(h > \|x_0\|^\alpha \theta \sum_{x \in \Phi \setminus \{x_0\}} h_x \|x\|^{-\alpha} \;\Big|\; \Phi\Big)$$
$$= \prod_{x \in \Phi \setminus \{x_0\}} \frac{1}{1 + \theta(\|x_0\|/\|x\|)^\alpha}.$$

The b-th moment follows as

$$M_b = \mathbb{E} \prod_{x \in \Phi \setminus \{x_0\}} \frac{1}{(1 + \theta(\|x_0\|/\|x\|)^\alpha)^b}.$$

To evaluate this, we use the relative distance process (RDP), defined as

$$\mathcal{R} \triangleq \{x \in \Phi \setminus \{x_0\}\colon \|x_0\|/\|x\|\} \subset [0,1].$$

Proof (cont'd)

The pgfl of the RDP for a PPP is ⟨Ganti and Haenggi, 2016a⟩

$$G_{\mathcal{R}}[v] \triangleq \mathbb{E} \prod_{x \in \mathcal{R}} v(x) = \frac{1}{1 + 2\int_0^1 (1 - v(x)) x^{-3} \mathrm{d}x},$$

hence we obtain

$$M_b = \frac{1}{1 + 2\int\limits_0^1 \left(1 - \frac{1}{(1+\theta r^\alpha)^b}\right) r^{-3} \mathrm{d}r} = \frac{1}{{}_2F_1(b, -\delta; 1-\delta; -\theta)}, \quad b \in \mathbb{C}.$$

Remark

Using the RDP provides us with a more direct way of calculating quantities such as $p_{\mathrm{s}}(\theta)$ and $P_{\mathrm{s}}(\theta)$. It combines the two steps of first conditioning on the distance to the nearest BS and then taking an expectation w.r.t. that distance into one.

Meta distribution and bounds

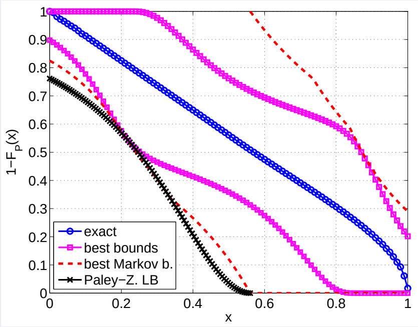

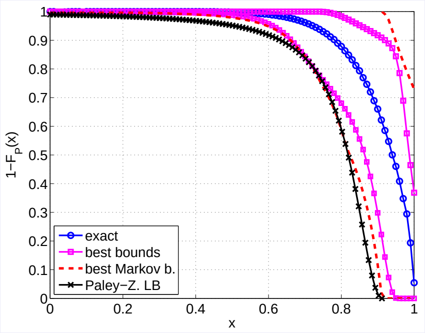

$\alpha = 4,\ \theta = 1$
$\to p_s = 0.56,\ \mathrm{var}(P_s) = 0.098$

$\alpha = 4,\ \theta = 1/10$
$\to p_s = 0.91,\ \mathrm{var}(P_s) = 0.0086$

As before, "best bounds" here means the best possible bounds that can be obtained given the first four moments.

Approximation with beta distribution

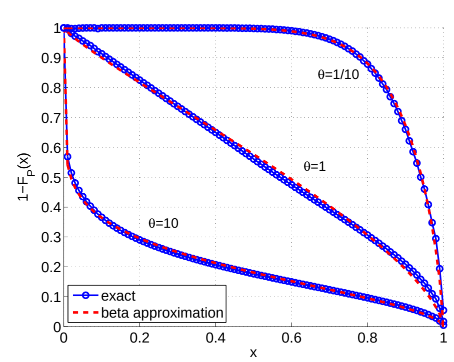

Exact ccdf and beta approximation for $\theta = 1/10, 1, 10$ for $\alpha = 4$.

The beta distribution tightly approximates the meta distribution.

Coverage

What is "coverage"?

- $\mathbb{P}(\mathrm{SIR} > \theta)$ gives, in each realization and each time slot, the fraction of users who happen to succeed. Some because of good fading from the BS, some because of bad fading from an interfering BS, some because they are close to the BS.
  In the next time slot, some previously successful users won't succeed, and vice versa.
- This is not a robust metric for coverage. Declaring a user "covered" or not on a 10 ms time scale is impractical. We would have to redraw coverage maps 100 times/s, at a spatial scale of cm.
- We need a metric that does not depend on the instantaneous channel realization, but still takes into account the fading statistics.

What is *coverage*—solution

- For each user u, calculate

$$P_{\mathrm{s}}^{(u)} = \mathbb{P}(\mathsf{SIR}_u > \theta \mid \Phi) = \mathbb{E}_h \mathbf{1}(\mathsf{SIR}_u > \theta).$$

  This averages over the fading (and random access).
- Then declare those user covered for whom $P_{\mathrm{s}}^{(u)} > x$, where $x \in [0, 1]$ is a reliability constraint.
  This gives a robust coverage map and reflects true user satisfaction.
- It also achieves a time scale separation between the time scales of fading and changes in the network geometry.

Coverage means to consistently achieve a certain SIR.

The next talk has nice illustrations of coverage and per-user SINR ccdfs.

Uplink in cellular networks

Uplink with power control ⟨Wang et al., 2017⟩

Often, the benefits of a transmission technique are not reflected in the mean success probability. Uplink power control is an important example.

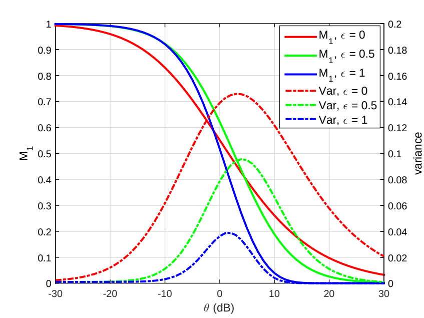

Power control: For link distance R, user transmits at power $R^{\alpha\epsilon}$.

$\epsilon \in [0, 1]$: no power control to full inversion of large-scale path loss.

For a target SIR of around 0 dB, $p_s(1) \approx$ 50–60%, irrespective of ϵ. So what ϵ is best?

The variance $M_2 - M_1^2$ shows a gain of at least a factor of 3 for $\epsilon = 1$. Hence power control reduces the inequality in the user experiences.

Meta distribution with D2D underlay ⟨Salehi et al., 2017⟩

- Network modeled as superposition of a Poisson cellular network and an independent Poisson bipolar network (D2D users).
- Base stations transmit with probability p_{BS} and D2D users with probability p_{D2D}.
- For both types of users, the moments M_b can be calculated exactly.

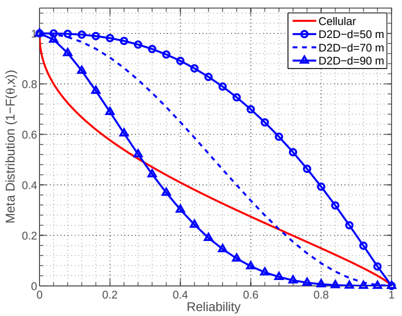

$\lambda_{\mathrm{BS}} = 2\ \mathrm{km}^{-2}$, $\lambda_{\mathrm{D2D}} = 50\ \mathrm{km}^{-2}$. $\theta = 1$, $P_{\mathrm{BS}}/P_{\mathrm{D2D}} = 100$, $p_{\mathrm{BS}} = 0.7$, $p_{\mathrm{D2D}} = 0.3$, $\alpha = 4$.

Using the meta distribution, we can calculate the density of D2D links that can be accommodated such that both types of users maintain a target reliability. Again the beta approximation is very accurate.

Other work

Other "per-user" and "meta" work

- SIR and throughput improvement for cellular networks using rateless codes ⟨Rajanna and Haenggi, 2017⟩
- Predicting transmission success in Poisson bipolar networks ⟨Weber, 2017⟩
- Millimeter wave networks ⟨Deng and Haenggi, 2017⟩
- Downlink cellular networks with base station cooperation
- Bipolar networks with interference cancellation
- Spatial outage capacity in bipolar networks ⟨Kalamkar and Haenggi, 2017a⟩
- Ergodic spectral efficiency in cellular networks ⟨George et al., 2017⟩

Spatial outage capacity

Joint work with S. Kalamkar

Density of links given an outage constraint

Fundamental question: What is the maximum density of concurrent transmissions given an outage constraint?

For a stationary and ergodic point process Φ of potential transmitters,

$$\lambda_\varepsilon \triangleq \lim_{r\to\infty} \frac{1}{\pi r^2} \sum_{\substack{y\in\Phi \\ \|y\|<r}} \mathbf{1}(\mathbb{P}(\mathrm{SIR}_{\tilde{y}} > \theta \mid \Phi) > 1-\varepsilon)$$

is the density of transmissions satisfying an outage constraint ε.

The outage constraint results in a static dependent thinning of Φ to a point process of density λ_ε.

Goal: Maximize λ_ε over λ and p to obtain the spatial outage capacity.

Key observation

λ_ε can be expressed using the meta distribution as

$$\lambda_\varepsilon = \lambda p \bar{F}(\theta, 1-\varepsilon),$$

where λ is the density of Φ and p is the fraction of concurrently active transmitters.

Definition (Spatial outage capacity (SOC) ⟨Kalamkar and Haenggi, 2017b⟩)

For a stationary and ergodic point process model and parameters $\theta > 0$ and $\varepsilon \in (0,1)$, the spatial outage capacity is defined as

$$S(\theta, \varepsilon) \triangleq \sup_{\lambda>0, p\in(0,1]} \lambda_\varepsilon = \sup_{\lambda>0, p\in(0,1]} \lambda p \bar{F}(\theta, 1-\varepsilon).$$

λ is the density of the point process, p is the fraction of links that are concurrently active, and $\bar{F}$ is the SIR meta distribution.

Comparison with transmission capacity ⟨Weber et al., 2010⟩

Let $P_s(\theta)$ be the conditional success probability given the point process.

$$\text{SOC:} \quad S(\theta,\varepsilon) \triangleq \sup_{\lambda>0, p\in(0,1]} \left\{\lambda p\, \mathbb{P}\big(P_s(\lambda,p,\theta) > 1-\varepsilon\big)\right\}$$

$$\text{TC:} \quad c(\theta,\varepsilon) \triangleq (1-\varepsilon)\sup\{\lambda p > 0 \colon \mathbb{E}P_s(\lambda,p,\theta) > 1-\varepsilon\}$$

The mean $p_s(\lambda p,\theta) \triangleq \mathbb{E}P_s(\lambda,p,\theta)$ only depends on the product λp and is monotonic, hence the TC can be written as $c(\theta,\varepsilon) \triangleq (1-\varepsilon)p_s^{-1}(1-\varepsilon)$.

The TC yields the maximum density of links such that the typical link satisfies an outage constraint. The supremum is taken only over one parameter, namely λp.

In the SOC, the outage constraint is applied at each individual link. It yields the maximum density of links that satisfy an outage constraint. This means that λ and p need to be considered separately.

Example (SOC vs. TC)

For the Poisson bipolar network with ALOHA, we set $r = 1$, $\alpha = 4$ and consider $\varepsilon = \theta = 1/10$.

- Transmission capacity:
  $c(1/10, 1/10) = 0.061$, achieved at $\lambda p = 0.0675$. By design, $p_s = 0.9$. But at $p = 1$, only 82% of the transmissions satisfy the 10% outage. Hence the spatial density of links that achieve 10% outage is only 0.055.
- Spatial outage capacity:
  $S(1/10, 1/10) = 0.092$, achieved at $\lambda = 0.23$ and $p = 1$, resulting in $p_s = 0.7$.

Hence the maximum spatial density of links given the 10% outage constraint is more than 50% larger than the TC.

3D plot for Poisson bipolar networks with ALOHA

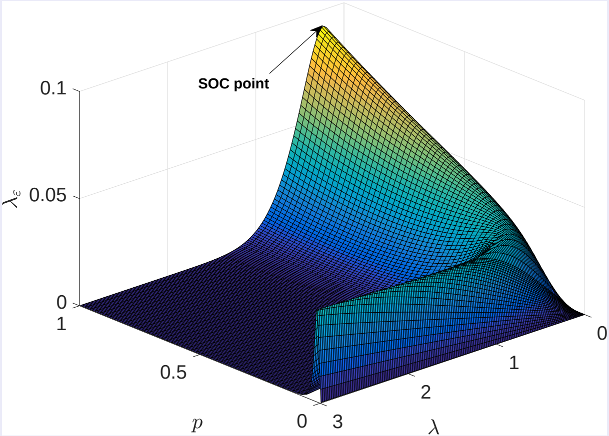

3D plot for $\varepsilon = 1/10$, $\theta = 1/10$, $r = 1$, $\alpha = 4$.

Poisson bipolar networks in the high-reliability regime

Using de Bruijn's Tauberian theorem, it can be shown that ⟨Kalamkar and Haenggi, 2017b⟩

$$\lambda_\varepsilon \sim \lambda p \exp\left(-\left(\frac{\theta p}{\varepsilon}\right)^\kappa \frac{\left(\lambda\delta\pi r^2\Gamma(1-\delta)\right)^{\kappa/\delta}}{\kappa}\right), \quad \varepsilon \to 0,$$

where $\kappa = \delta/(1-\delta) = 2/(\alpha-2)$.

- Interestingly, only the ratio of ε to θ matters.
- λ and p have different exponents, hence not only their product matters.

Outage-constrained density for Poisson bipolar network

For non-asymptotic values of ε, λ_ε can be approximated using the beta distribution.

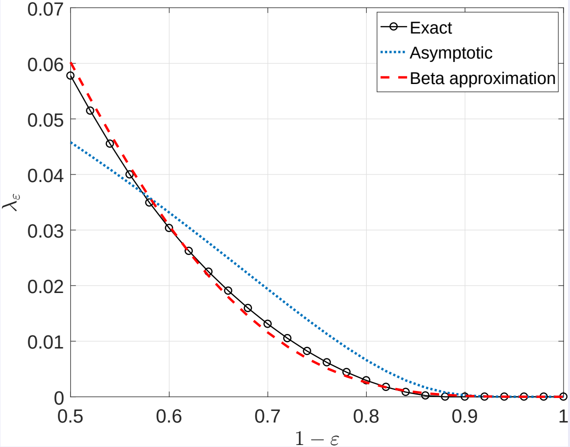

λ_ε for $\varepsilon = 1 - x$, $\theta = 1$, $r = 1$, $\alpha = 4$, $\lambda = 1/2$, $p = 1/3$.

SOC in the high-reliability regime

It follows from the high-reliability result for λ_ε that

$$S(\theta,\varepsilon) \sim \left(\frac{\varepsilon}{\delta\theta}\right)^{\delta} \frac{e^{-(1-\delta)}}{\pi r^2 \Gamma(1-\delta)}, \quad \varepsilon \to 0.$$

- The SOC is achieved at $p = 1$. (This holds also for Rayleigh distributed link distances.)
- The ratio ε/θ shows an interesting rate-reliability trade-off: At low rates, $\log(1+\theta) \sim \theta$, so a $10\times$ higher reliability can be achieved by lowering the rate by a factor 10.
- Alternative form:

$$S\pi r^2 \sim \left(\frac{\varepsilon}{\theta}\right)^{\delta} f(\delta)$$

- For $r = 1$ and $\alpha = 4$, $S \sim 0.154\sqrt{\varepsilon/\theta}$, and $M_{1,\mathrm{opt}} \sim 1 - 1.2533\sqrt{\varepsilon}$.

Ergodic spectral efficiency

Joint work with G. Georgie, R. Mungara, and A. Lozano

Motivation

- The outage-based framework of the meta distribution is useful for short messages and low-latency situations.
- For longer messages (codewords) transmitted over larger bandwidths or many antennas or using hybrid ARQ, an ergodic point of view is warranted.
- As before, we aim at a clean time-scale separation. Ergodicity applies to the time scale of small-scale fading, with the network geometry fixed. Then stochastic geometry is applied to capture different network configurations.
- This approach lends itself to MIMO extensions and sectorization.

Ergodic spectral efficiency ⟨George et al., 2017⟩

Let
$$\rho \triangleq \frac{\|x_0\|^{-\alpha}}{\sum_{x\in\Phi\setminus\{x_0\}}\|x\|^{-\alpha}}$$

be the SIR (of the user at the origin) without fading. It captures the network geometry. Next, let

$$C(\rho) \triangleq \mathbb{E}_h(\log(1+h\rho))$$

be the ergodic spectral efficiency given the point process. For Rayleigh fading, $C(\rho) = e^{1/\rho}E_1(1/\rho)$, where E_1 is the exponential integral.

Q: Why not include the fading of the interferers' channels?

A: Because the user does not know them.

Ignoring the fading of the interferers yields a tight lower bound, while including it in the expression would yield a looser upper bound.

Distribution of the ergodic spectral efficiency

- The conditional ergodic spectral efficiency $C(\rho)$ is a random variable. In a realization, each user u has her/his personal ρ_{u}.
- Naturally, we are interested in the ccdf $F_{C(\rho)}(\gamma) = \mathbb{P}(C(\rho) \leq \gamma)$. To evaluate it, we need the distribution of ρ.
  - For $\theta \geq 1$, $F_\rho(\theta) = 1 - \mathrm{sinc}(\delta)\theta^{-\delta}$ ⟨Zhang and Haenggi, 2014b; Madhusudhanan et al., 2014⟩.
  - For $\theta < 1$ an exact integral expression can be given that gets increasingly cumbersome for $\theta \to 0$ ⟨Blaszczyszyn and Keeler, 2015⟩.
  - As $\theta \to 0$, $\log F_\rho(\theta) = s^*/\theta + o(1)$, where $s^* < 0$ is given by $s^{*\delta}\bar{\Gamma}(-\delta, s^*) = 0$ ⟨Ganti and Haenggi, 2016b⟩.
- Lastly, we use

$$F_{C(\rho)}(\gamma) = F_\rho(C^{-1}(\gamma)).$$

SISO with Rayleigh fading

Using an invertible approximation of $C(\rho) = e^{1/\rho} E_1(1/\rho)$:

$$F_C(\gamma) \approx F_\rho\left(\frac{e^{\frac{\gamma}{1.4}}-1}{0.82}\right)$$

$$= \begin{cases} e^{\frac{0.82\, s^\star}{\exp(\gamma/1.4)-1}} & 0 \le \gamma < 1.4\log\left(1+\frac{0.82\, s^\star}{\log A_\delta}\right) \\ A_\delta & 1.4\log\left(1+\frac{0.82\, s^\star}{\log A_\delta}\right) \le \gamma < 0.48 \\ 1-\frac{\operatorname{sinc}\delta}{\left(\frac{e^{\gamma/1.4}-1}{0.82}\right)^\delta} + B_\delta\left(\frac{e^{\gamma/1.4}-1}{1.82-e^{\gamma/1.4}}\right) & 0.48 \le \gamma < 0.84 \\ 1-\frac{\operatorname{sinc}\delta}{\left(\frac{e^{\gamma/1.4}-1}{0.82}\right)^\delta} & \gamma \ge 0.84, \end{cases}$$

where

$$A_\delta = 1 - 2^\delta \operatorname{sinc}\delta + B_\delta(1)$$

$$B_\delta(x) = \frac{\delta\operatorname{sinc}^2(\delta)\,\Gamma^2(\delta+1)\,{}_2F_1(1,\delta+1;2\,\delta+2;-1/x)}{x^{1+2\,\delta}\,\Gamma(2\,\delta+2)}$$

Distribution of the ergodic spectral efficiency

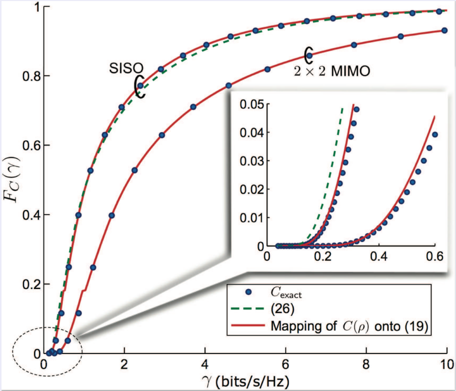

With SISO, essentially no user gets less than 0.18 bps/Hz. With 2x2 MIMO, no user gets less than 0.3 bps/Hz.

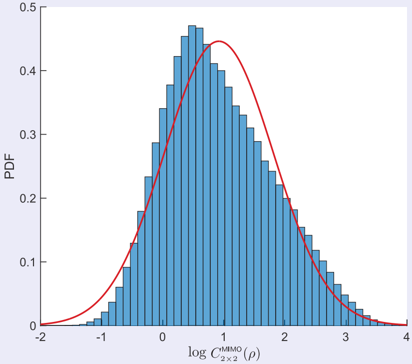

Interesting observation: Spectral efficiencies are essentially lognormal.

Scaling behavior at high user reliabilities

What is the spectral efficiency achieved by a fraction $1-\xi$ of the users, for $\xi \ll 1$? Setting $\xi = \mathbb{P}(C(\rho) \leq \gamma) = F_C(\gamma) \approx e^{1.15s^*/\gamma}$, we obtain

$$\gamma \approx \frac{1.15s^*}{\log \xi}, \quad \xi \ll 1.$$

- For $\alpha = 4$ and $\xi < 0.15$, this simplifies to $\gamma \approx -1/\log \xi$.
- For $\xi = 1/100$, for example, we obtain $\gamma \approx 0.22$ bps/Hz, while the exact value is 0.24 bps/Hz.
- For comparison, if we used $\bar{F}_{\mathrm{SIR}}(\theta) = 0.99$, we would get $\theta = -20$ dB and $\gamma = 0.014$ bps/Hz.

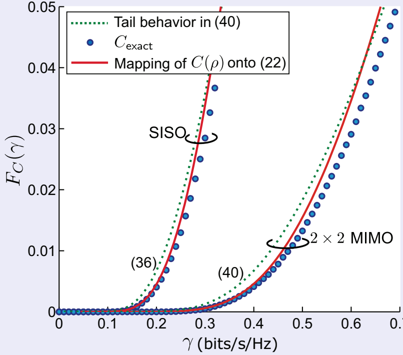

Conclusions

- Spatial distributions of the form $\mathbb{P}(\mathbb{E}_h f(\Phi) > t)$ achieve a clean separation of temporal and spatial randomness.
  - $f(\Phi) = \mathbf{1}(\mathrm{SIR} > \theta)$ yields the meta distribution of the SIR (outage-based performance).
  - $f(\Phi) = \log(1 + h\rho)$, where ρ is the SIR without fading, yields the distribution of the ergodic spectral efficiency.

  This yields the area/user/link fraction that achieves performance t and thus the performance of the 5% user. Classical averages for the typical user/link are obtained by integration over t.
- The meta distribution can be bounded and calculated using the moments. A beta approximation yields simple yet accurate results.
- The ergodic spectral efficiency distribution can be well approximated in closed-form, also for MIMO, using results on the SIR distribution without fading. It is close to lognormal.

Slides available at: `www.nd.edu/~mhaenggi/talks/spaswin17.pdf`

References I

M. Haenggi. The Meta Distribution of the SIR in Poisson Bipolar and Cellular Networks. *IEEE Transactions on Wireless Communications*, 15(4):2577–2589, April 2016.

M. Zorzi and S. Pupolin. Optimum Transmission Ranges in Multihop Packet Radio Networks in the Presence of Fading. *IEEE Transactions on Communications*, 43(7):2201–2205, July 1995.

F. Baccelli, B. Blaszczyszyn, and P. Mühlethaler. An ALOHA Protocol for Multihop Mobile Wireless Networks. *IEEE Transactions on Information Theory*, 52(2):421–436, February 2006.

J. G. Andrews, F. Baccelli, and R. K. Ganti. A Tractable Approach to Coverage and Rate in Cellular Networks. *IEEE Transactions on Communications*, 59(11):3122–3134, November 2011.

S. Mukherjee. Distribution of Downlink SINR in Heterogeneous Cellular Networks. *IEEE Journal on Selected Areas in Communications*, 30(3):575–585, April 2012.

G. Nigam, P. Minero, and M. Haenggi. Coordinated Multipoint Joint Transmission in Heterogeneous Networks. *IEEE Transactions on Communications*, 62(11):4134–4146, November 2014.

M. DiRenzo. Stochastic Geometry Modeling and Analysis of Multi-Tier Millimeter Wave Cellular Networks. *IEEE Transactions on Wireless Communications*, 14(9):5038–5057, September 2015.

N. Deng, W. Zhou, and M. Haenggi. The Ginibre Point Process as a Model for Wireless Networks with Repulsion. *IEEE Transactions on Wireless Communications*, 14(1):107–121, January 2015.

References II

P. Madhusudhanan, J. G. Restrepo, Y. Liu, T. X. Brown, and K. R. Baker. Analysis of Downlink Connectivity Models in a Heterogeneous Cellular Network via Stochastic Geometry. *IEEE Transactions on Wireless Communications*, 13(12):6684–6696, December 2014.

F. Baccelli and B. Blaszczyszyn. A New Phase Transition for Local Delays in MANETs. In *IEEE INFOCOM'10*, San Diego, CA, March 2010.

M. Haenggi. The Local Delay in Poisson Networks. *IEEE Transactions on Information Theory*, 59(3):1788–1802, March 2013.

R. K. Ganti and J. G. Andrews. Correlation of Link Outages in Low-Mobility Spatial Wireless Networks. In *44th Asilomar Conference on Signals, Systems, and Computers (Asilomar'10)*, Pacific Grove, CA, November 2010.

M. Haenggi and R. Smarandache. Diversity Polynomials for the Analysis of Temporal Correlations in Wireless Networks. *IEEE Transactions on Wireless Communications*, 12(11): 5940–5951, November 2013.

X. Zhang and M. Haenggi. A Stochastic Geometry Analysis of Inter-cell Interference Coordination and Intra-cell Diversity. *IEEE Transactions on Wireless Communications*, 13 (12):6655–6669, December 2014a.

P. Madhusudhanan, J. G. Restrepo, Y. Liu, and T. X. Brown. Analysis of Downlink Connectivity Models in a Heterogeneous Cellular Network via Stochastic Geometry. *IEEE Transactions on Wireless Communications*, 15(6):3895–3907, June 2016.

References III

R. K. Ganti and M. Haenggi. Asymptotics and Approximation of the SIR Distribution in General Cellular Networks. *IEEE Transactions on Wireless Communications*, 15(3):2130–2143, March 2016a.

Y. Wang, M. Haenggi, and Z. Tan. The Meta Distribution of the SIR for Cellular Networks with Power Control. ArXiv, `http://arxiv.org/abs/1702.01864v1`, February 2017.

M. Salehi, A. Mohammadi, and M. Haenggi. Analysis of D2D Underlaid Cellular Networks: SIR Meta Distribution and Mean Local Delay. *IEEE Transactions on Communications*, 2017. Accepted. Available at `http://www.nd.edu/~mhaenggi/pubs/tcom17b.pdf`.

A. Rajanna and M. Haenggi. Enhanced Cellular Coverage and Throughput using Rateless Codes. *IEEE Transactions on Communications*, 65(5):1899–1912, May 2017.

S. Weber. The value of observations in predicting transmission success in wireless networks under slotted Aloha. In *15th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'17)*, Paris, France, May 2017.

N. Deng and M. Haenggi. A Fine-Grained Analysis of Millimeter-Wave Device-to-Device Networks. *IEEE Transactions on Communications*, 2017. Submitted.

S. S. Kalamkar and M. Haenggi. The Spatial Outage Capacity of Wireless Networks. *IEEE Transactions on Wireless Communications*, 2017a. Submitted.

G. George, R. K. Mungara, A. Lozano, and M. Haenggi. Ergodic Spectral Efficiency in MIMO Cellular Networks. *IEEE Transactions on Wireless Communications*, 16(5):2835–2849, May 2017.

References IV

S. S. Kalamkar and M. Haenggi. Spatial Outage Capacity of Poisson Bipolar Networks. In *IEEE International Conference on Communications (ICC'17)*, Paris, France, May 2017b. Available at http://www.nd.edu/~mhaenggi/pubs/icc17.pdf.

S. Weber, J. G. Andrews, and N. Jindal. An Overview of the Transmission Capacity of Wireless Networks. *IEEE Transactions on Communications*, 58(12):3593–3604, December 2010.

X. Zhang and M. Haenggi. The Performance of Successive Interference Cancellation in Random Wireless Networks. *IEEE Transactions on Information Theory*, 60(10):6368–6388, October 2014b.

B. Blaszczyszyn and H. P. Keeler. Studying the SINR process of the typical user in Poisson networks by using its factorial moment measures. *IEEE Transactions on Information Theory*, 61(12):6774–6794, December 2015.

R. K. Ganti and M. Haenggi. SIR Asymptotics in Poisson Cellular Networks without Fading and with Partial Fading. In *IEEE International Conference on Communications (ICC'16)*, Kuala Lumpur, Malaysia, May 2016b.