Modeling and Analysis of Advanced Architectures for 5G Networks

Martin Haenggi

University of Notre Dame, IN, USA

ET5G Keynote

Dec. 4, 2016

Available at http://www.nd.edu/~mhaenggi/talks/globecom16.pdf

Cellular system evolution

nG, $n \leq 4$

5G+: The ultra-era

- ultra-fast
- ultra-dense
- ultra-low latency
- ultra-high reliability
- ultra-heterogeneous
- ultra-MIMO (well, massive)
- ultra-high frequency
- ultra-flexible
- ultra-software-defined
- ultra-virtualized

$\Rightarrow$ ultra-happy users (?)

Modeling networks using point processes

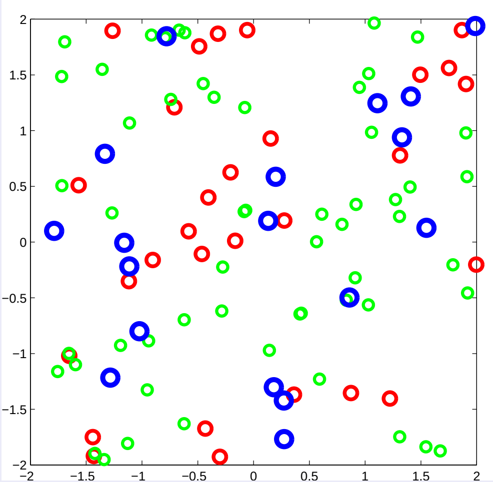

Lots of points—base stations at different tiers, users, RRHs, etc.

Modeling without point processes

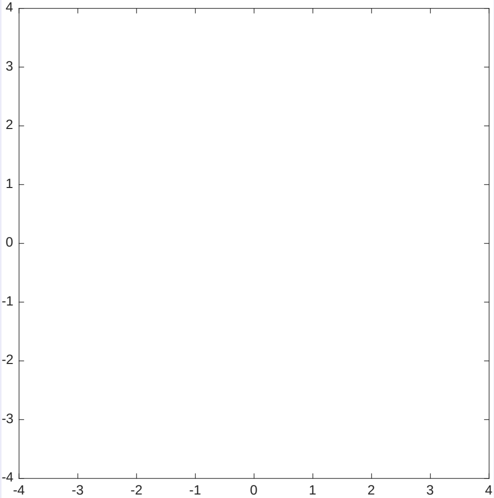

Modeling without stochastic geometry is point-less.

Menu

Overview

- What if networks are not HIP (homogeneous independent Poisson)?
- Revisiting coverage: the meta distribution of the SIR
  - Downlink
  - Uplink
  - D2D
- Diverse messages and ergodic spectral efficiency
- Conclusions

Ultra-Menu

Overview in ultra-speak

- Analyzing ultra-HetNets
- Ultra-reliability at ultra-low latency
- Ultra-dense networks
- Ultra-fast and ultra-reliable transmission of long messages
- Conclusions

HIP—and beyond

The HIP (homogeneous independent Poisson) model

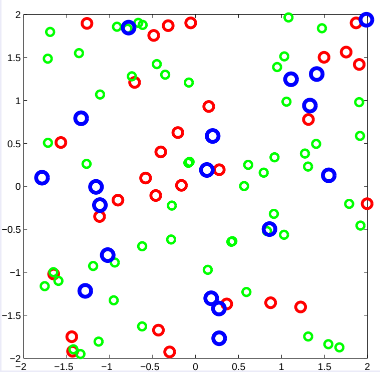

Base stations in each tier form independent Poisson point processes (PPPs) of densities λ_i and transmit powers P_i.

Here $\lambda_i = 1, 2, 3$ for the blue, red, and green tiers.

This model is doubly independent and thus highly tractable.

A user connects to the BS that is strongest on average (not including fading), while all others interfere. We are interested in the SIR.

Basic result for HIP downlink

Assumptions:

- Homogeneous path loss law $\ell(r) = r^{-\alpha}$
- Rayleigh fading

Remarkably, the SIR distribution at the typical user[a] is independent of the number of tiers, their densities, and their power levels [ZH14, NMH14]:

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

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

What can be said about non-HIP models?

[a] If users form an independent stationary point process, this equals the average over all users assuming they are served.

Beyond HIP

SIR gain [GH15, Hae14]

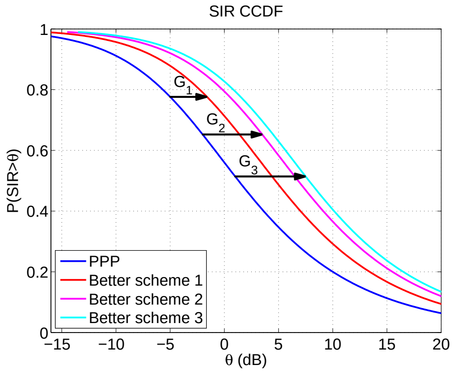

When comparing different models or transmission schemes, the resulting SIR ccdfs appear to be horizontally shifted versions of each other.

The nearly constant gap is the SIR gain.

If the SIR ccdfs were indeed just shifted:

$$p_{\mathrm{s,PPP}}(\theta) \triangleq \mathbb{P}(\mathrm{SIR}_{\mathrm{PPP}} > \theta) \quad \Rightarrow \quad p_{\mathrm{s}}(\theta) = p_{\mathrm{s,PPP}}(\theta/G).$$

ASAPPP: "Approximate SIR Analysis based on the PPP". Or simply "as a PPP".

Deployment gain [GH15]

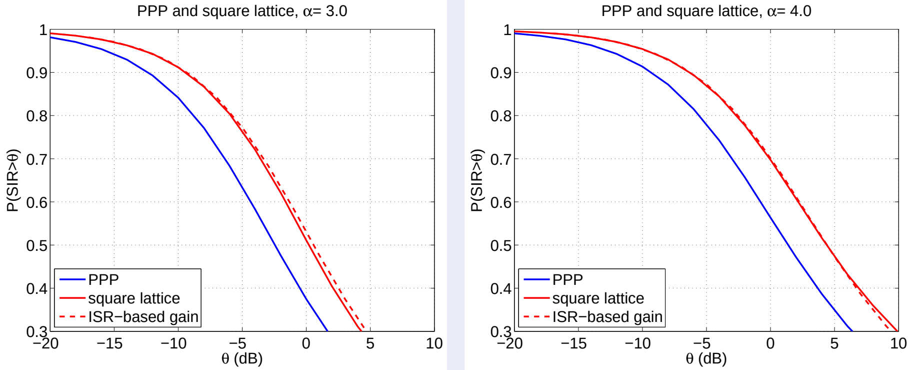

- For the square lattice, the gap (deployment gain) is quite exactly 3 dB—irrespective of α! For $\alpha = 4$, $p_s^{\text{sq}} = (1 + \sqrt{\theta/2}\arctan\sqrt{\theta/2})^{-1}$.
- For the triangular lattice, it is 3.4 dB. This is the maximum achievable.

Definition (Interference-to-(average)-signal ratio $\overline{\text{ISR}}$)

$$\overline{\text{ISR}} \triangleq \frac{I}{\mathbb{E}_h(S)} = \frac{I}{\ell(x_0)}$$

Relevance [Hae14]

- The SIR ccdf is given by $p_s(\theta) = \mathbb{E}\bar{F}_h(\theta\,\overline{\text{ISR}})$.
- For Rayleigh fading, $p_s(\theta) = \mathcal{L}_{\overline{\text{ISR}}}(\theta)$ and, letting $\text{MISR} = \mathbb{E}(\overline{\text{ISR}})$,

$$p_s(\theta) \sim \theta\,\text{MISR}, \quad \theta \to 0.$$

- Hence the asymptotic gain G_0 at $\theta \to 0$ is $G_0 = \text{MISR}_{\text{PPP}}/\text{MISR}_2$, where

$$\text{MISR}_{\text{PPP}} = \frac{2}{\alpha - 2} = \frac{\delta}{1-\delta}.$$

- The gain G_0 is surprisingly accurate over the entire ccdf, and it barely depends on α and and the fading distribution [GH16].

The bandgap of SIR distributions ($\alpha = 4$)

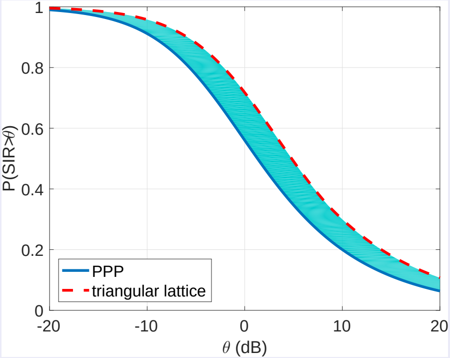

All (repulsive) deployments have SIRs that fall into this relatively thin band—even if there are many tiers.

Higher gains can only be achieved using interference-mitigating and/or signal-boosting schemes.

Extensions to HetNets

General K-tier HetNet model

- Let Φ_k, $k \in [K]$, be a family of independent simple stationary point processes modeling the locations of the base stations of tier k.
- Let λ_k and P_k denote the density and transmit power of tier k.
- Let G_k denote the SIR gain (relative to the PPP) of the single-tier network consisting only of tier k.
- A user is served by the BS providing the highest power (on average).

Example (3-tier HetNet)

Let $K = 3$, and take tier 1 to be a square lattice with $\lambda = 1/10$ and $P = 100$, tier 2 to be a Ginibre point process (GPP) with $\lambda = 1/2$ and $P = 10$, and tier 3 to be a PPP with $\lambda = 1$ and $P = 1$.

Result [WDZH16]

Define the effective gain as

$$G_{\text{eff}} \triangleq 1 + \sum_{k \in [k]} w_k^2 (G_k - 1),$$

where

$$w_k = \frac{\lambda_k P_k^\delta}{\sum_{i \in [K]} \lambda_i P_i^\delta}, \quad \delta = 2/\alpha.$$

In the HIP model, w_k is the probability that the typical user is served by tier k.

Then the SIR distribution for the HetNet is (ultra-)tightly approximated by

$$\bar{F}_{\text{SIR}}(\theta) \approx \bar{F}_{\text{SIR}}^{\text{PPP}}(\theta / G_{\text{eff}}).$$

Discussion

Let $\tilde{G}_k = G_k - 1$ and $\tilde{G}_{\text{eff}} = G_{\text{eff}} - 1$. Then

$$\tilde{G}_{\text{eff}} = \sum_{k \in [K]} w_k^2 \tilde{G}_k.$$

- Since the superposition of many independent stationary point processes results (under mild conditions) in a PPP, we have $\tilde{G}_{\text{eff}} \to 0$ as $K \to \infty$, no matter what the tiers are.
- For example, if all tiers are square lattices with the same intensity and transmit power, $\tilde{G}_k = 1$ and $w_k = 1/K$. It follows that

$$\tilde{G}_{\text{eff}} = \sum_{k \in [K]} \frac{1}{K^2} = \frac{1}{K}.$$

- Generally, for K identical tiers, $\tilde{G}_{\text{eff}} = \tilde{G}/K$.

Example: Numerical results for 3-tier HetNets

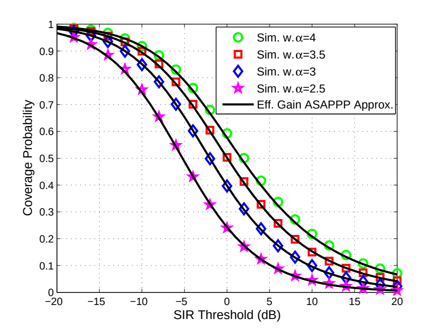

1-GPP / 0.5-GPP / PPP

$\lambda_1 = 10^{-5},\ P_1 = 1,$

$\lambda_2 = 2\lambda_1,\ P_2 = 1/5,$

$\lambda_3 = 5\lambda_1,\ P_3 = 1/25$

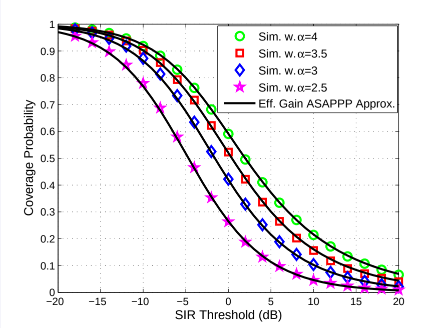

square lattice / 1-GPP / PPP

$\lambda_1 = 10^{-5},\ P_1 = 1,$

$\lambda_2 = 2\lambda_1,\ P_2 = 1/5,$

$\lambda_3 = 5\lambda_1,\ P_3 = 1/25$

Basketball

Who is a good free throw shooter?

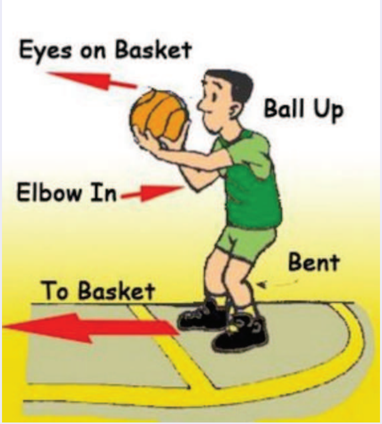

To find out which of 100 players are good free throw shooters, let them shoot one free throw. The good ones are those who succeed—are they?

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"?—cont'd

- So let the BS transmit 100 packets to the user. Declare the user satisfied—covered—if (s)he successfully receives 80 of them (in the first attempt). And ultra-happy if it is 99/100.
- Analogy: Let basketball players shoot 100 free throws. Declare those who make 80/100 good free throw shooters.

What is "coverage"—solution

- For each user u, calculate

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

  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.

Distribution of the conditional SIR distribution P_s

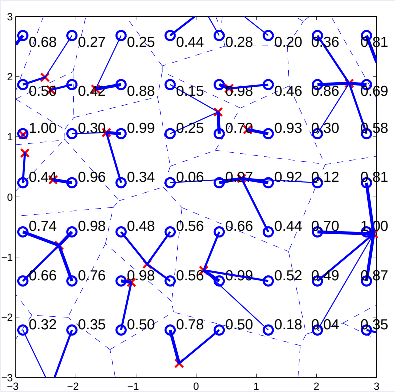

Per-user success probabilities $P_s^{(u)}$ for a PPP with $\theta = 1$, $\alpha = 4$.

The mean of all these random variables $P_s^{(u)}$ is the (standard) success probability $p_s(\theta) = \mathbb{P}(\mathrm{SIR} > \theta)$. Here we have $p_s(1) = 0.56$.

But we need to know more: What fraction of users achieve $P_s^{(u)} > x$, i.e., what is $\mathbb{P}(P_s^{(u)} > x)$?

So instead of just considering the mean of P_s, we need its ccdf. This is the meta distribution.

Definition

Definition (Meta distribution of the SIR [Hae16])

Let $\Phi \subset \mathbb{R}^2$ be a stationary and ergodic point process of base stations and let $\mathsf{SIR} = \mathsf{SIR}_o$ be the signal-to-interference ratio at the origin. The conditional success probability is the random variable

$$P_{\rm s}(\theta) \triangleq \mathbb{P}(\mathsf{SIR} > \theta \mid \Phi),$$

and its distribution is the meta distribution

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

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

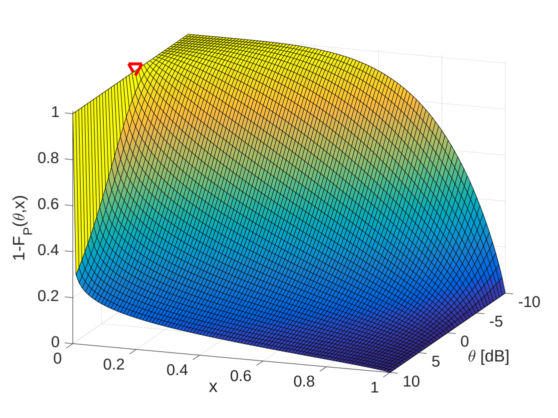

$$\bar{F}(\theta, x) = \mathbb{P}(P_s(\theta) > x)$$

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

Moments

The moments

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

reveal interesting properties of the meta distribution.

- $M_1 = \int_0^1 \bar{F}(\theta, x)\mathrm{d}x$ is the standard success probability.
- $\operatorname{var} P_{\mathrm{s}}(\theta) = M_2 - M_1^2$ gives basic information about the disparity of the user experiences.
- M_{-1} is the mean number of transmissions until success (aka local delay).

And they can be calculated in closed-form—at least in some cases.

Using M_{jt}, Gil-Pelaez inversion yields an integral expression of the exact meta distribution.

Theorem (Moments of P_{s} for Rayleigh fading (Hae16))

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

- For $b \in \mathbb{N}$, M_b is the joint success probability of b transmissions [ZH14].

Approximation with beta distribution

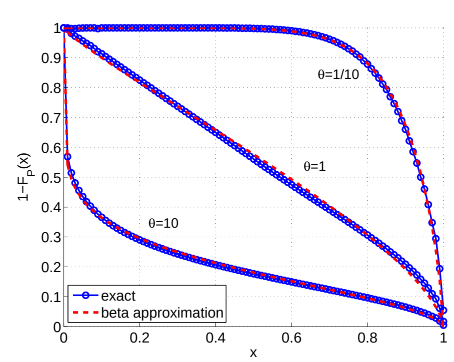

Exact ccdf and beta approximation for $\theta = 1/10, 1, 10$ for $\alpha = 4$.

The beta distribution tightly approximates the meta distribution.

Uplink with power control [WH16]

Often, the benefits of a transmission technique are not reflected in the mean success probability. Example: Uplink power control.

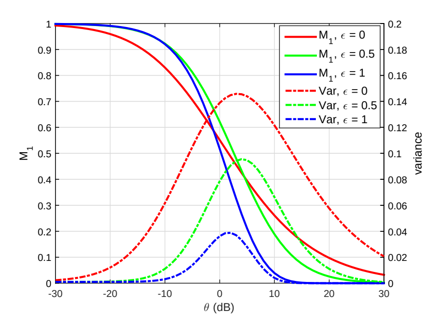

Power control: For link distance R, user transmits at power $R^{\alpha\epsilon}$.

$\epsilon \in [0, 1]$: 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 answer is given by the variance $M_2 - M_1^2$! It shows a gain of at least a factor of 3 for $\epsilon = 1$.
Hence power control leads to a concentration in the user experiences.

Meta distribution with D2D underlay [SMH16]

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

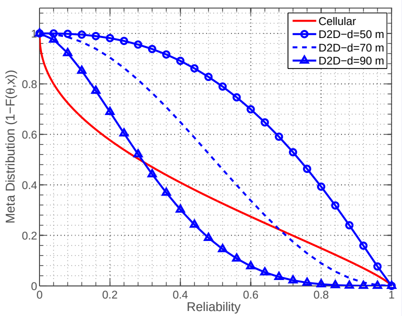

$\lambda_{\rm BS} = 2\ {\rm km}^{-2}$, $\lambda_{\rm D2D} = 50\ {\rm km}^{-2}$. $\theta = 1$, $P_{\rm BS}/P_{\rm D2D} = 100$, $p_{\rm BS} = 0.7$, $p_{\rm D2D} = 0.3$, $\alpha = 4$.

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.

Spatial outage capacity [KH17]

A fundamental question—in view of ultra-dense networks: What is the maximum density of concurrent links that achieve a reliability x?

Definition (Spatial outage capacity):

$$S(\theta, x) = \sup_{\lambda, p} \lambda p \bar{F}(\theta, x), \quad \lambda > 0, p \in (0, 1]$$

p is the fraction of links that are concurrently active.

Example: Poisson bipolar network in the (ultra-)high reliability regime

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

is achieved at $p = 1$. 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.

Ergodic spectral efficiency

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.

Long (diverse) messages: Ergodic spectral efficiency [GMLH16]

Let
$$\rho \triangleq \frac{\ell(x_0)}{\sum_{x\in\Phi\setminus\{x_0\}} \ell(x)}$$

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.

Ergodic spectral efficiency distribution

Using stochastic geometry, we can evaluate the expectation w.r.t. point process and find (approximately) the distribution of C.

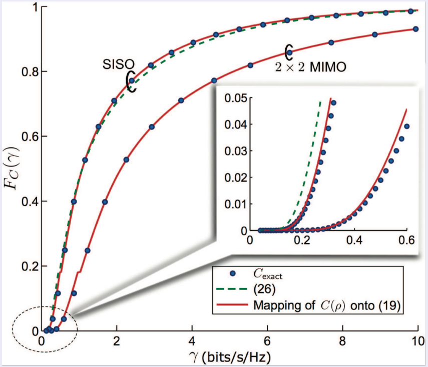

With SISO, essentially no user gets less than 0.18 bps/Hz. With 2x2 MIMO, no user gets less than 0.3 bps/Hz.

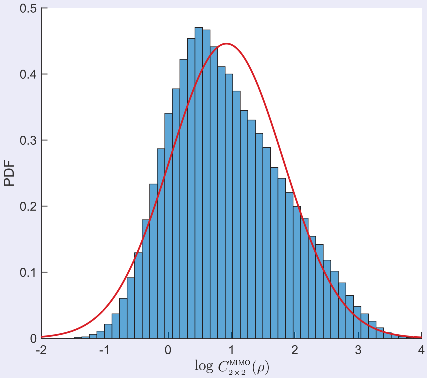

Observation: Spectral efficiencies are essentially lognormal.

Conclusions

- We have discussed several ultras: Density, reliability, latency, heterogenity.
- With ASAPPP, models do not have to be HIP—they can be quite geeky. The basic idea is to shift the SIR ccdf of the PPP horizontally, by a gain factor that is obtained asymptotically.
- The meta distribution captures the per-user experience and has a natural interpretation as a metric for coverage. It captures the disparity of user experiences and yields the performance of user percentiles, say the "5% user".
  It also provides a mathematical foundation for questions of network densification under strict reliability constraints.
- Ergodic spectral efficiency is the long-packet (or diverse packet) counterpart. Its distribution can be well approximated in closed-form, also for MIMO.

Slides available at: `www.nd.edu/~mhaenggi/talks/globecom16.pdf`

References I

A. Guo and M. Haenggi.
Asymptotic Deployment Gain: A Simple Approach to Characterize the SINR Distribution in General Cellular Networks.
IEEE Transactions on Communications, 63(3):962–976, March 2015.

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 2016.

G. George, R. K. Mungara, A. Lozano, and M. Haenggi.
Ergodic Spectral Efficiency in MIMO Cellular Networks.
IEEE Transactions on Wireless Communications, 2016.
Submitted. Available at https://arxiv.org/pdf/1607.04352v2.

M. Haenggi.
The Mean Interference-to-Signal Ratio and its Key Role in Cellular and Amorphous Networks.
IEEE Wireless Communications Letters, 3(6):597–600, December 2014.

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.

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

References II

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. Salehi, A. Mohammadi, and M. Haenggi.
Analysis of D2D Underlaid Cellular Networks: SIR Meta Distribution and Mean Local Delay.
IEEE Transactions on Communications, 2016.
Submitted.

H. Wei, N. Deng, W. Zhou, and M. Haenggi.
Approximate SIR Analysis in General Heterogeneous Cellular Networks.
IEEE Transactions on Communications, 64(3):1259–1273, March 2016.

Y. Wang and M. Haenggi.
The Meta Distribution of the SIR for Cellular Networks with Power Control.
In preparation., 2016.

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 2014.