Analysis and Design of Wireless Networks

A Stochastic Geometry Approach

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and a Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop analysis of Poisson networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Analysis and Design of Wireless Networks

Lecture 1: Introduction and a Key Result

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and a Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop analysis of Poisson networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 1 Overview

1 Introduction

2 Modeling

3 Interference and Outage in Poisson Networks

4 Summary

Section Outline

1 Introduction
- Approaches to wireless network analysis
- The network geometry
- Spatial reuse
- How to manage spatial reuse?
- Abstraction
- Key result

2 Modeling

3 Interference and Outage in Poisson Networks

4 Summary

Analysis of Wireless Networks

Comparison of analytical approaches

scaling laws

very limited design insight

analysis of networks with fixed geometry

concrete results but no generality

stochastic geometry analysis of random networks

generality by spatial averaging *and* design insight

The Critical Role of the Network Geometry

Wireless transmissions are separated in space, time, or frequency

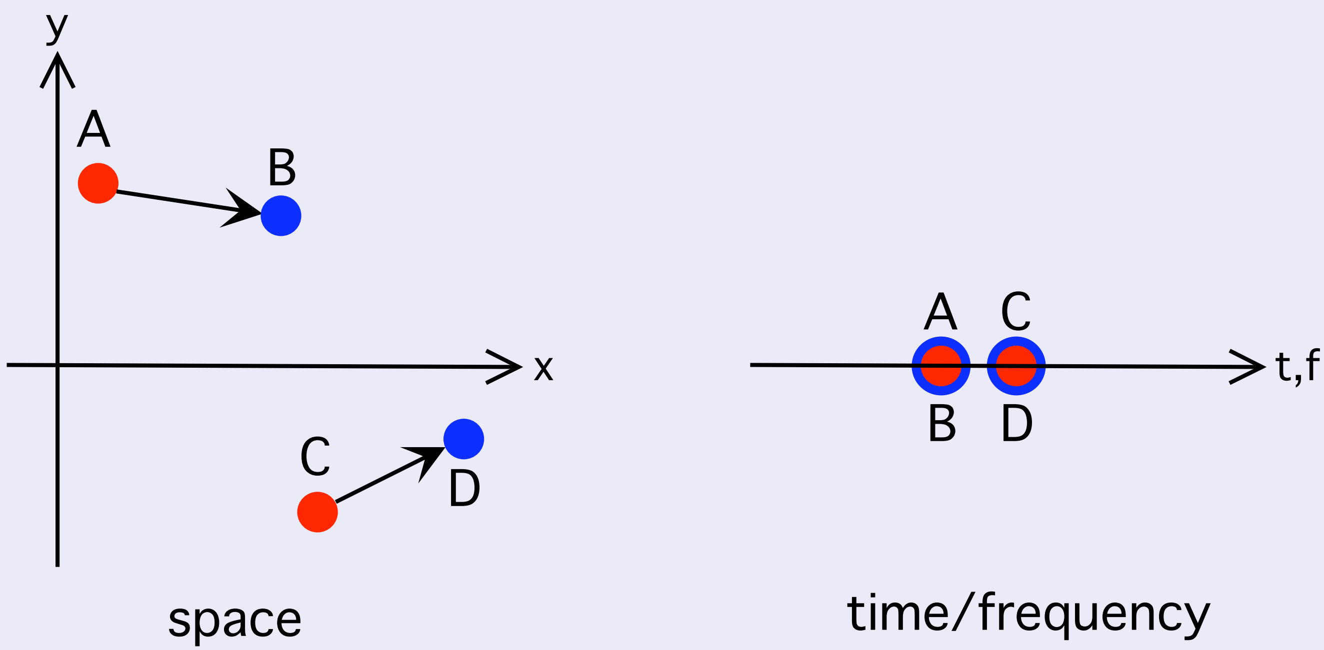

- Separation in time and frequency not sufficient for wireless networks.
- Need for *spatial reuse*. But separation in space is much more challenging.

Why is spatial reuse hard?

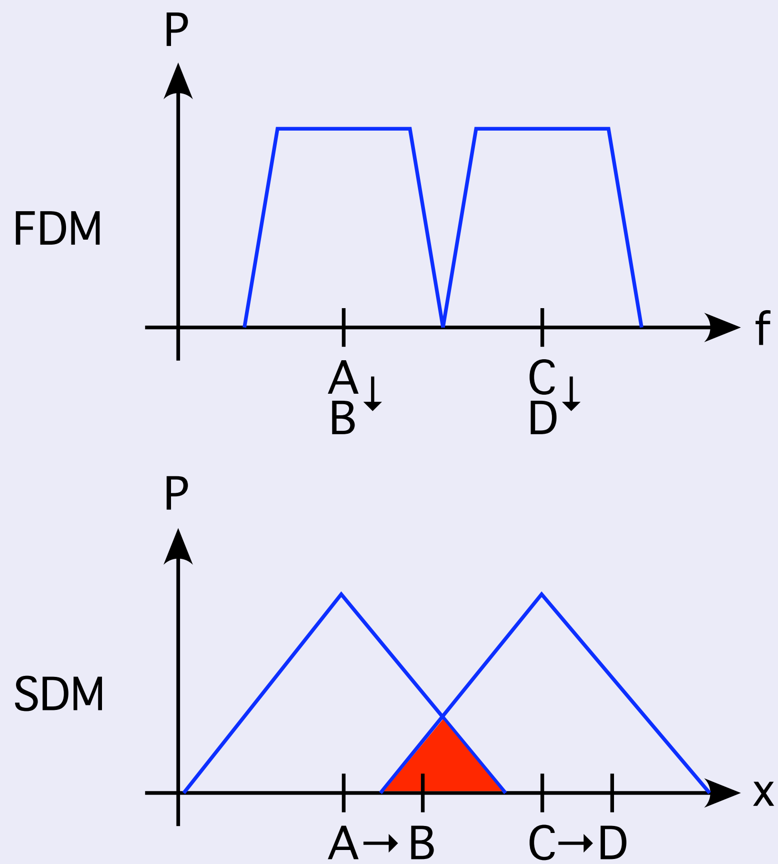

>100dB/decade
Tx, Rx colocated
Larger P ⇒ higher R

20-40dB/decade (dist.)
Tx, Rx separated
SIR independent of P

- There is interference between concurrent transmissions.
- Transmitter and receiver have a different picture of the situation.

The cellular solution

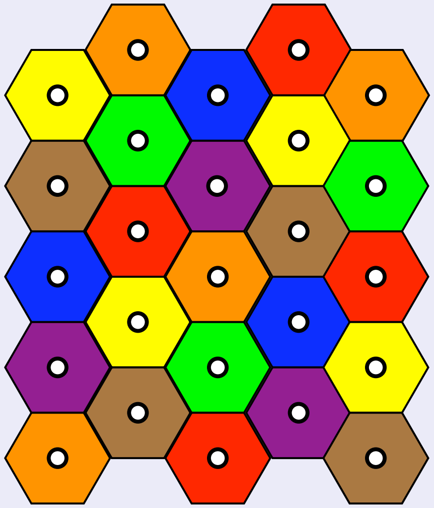

Cellular system with frequency reuse factor 1/7

A *sensible* solution: CSMA

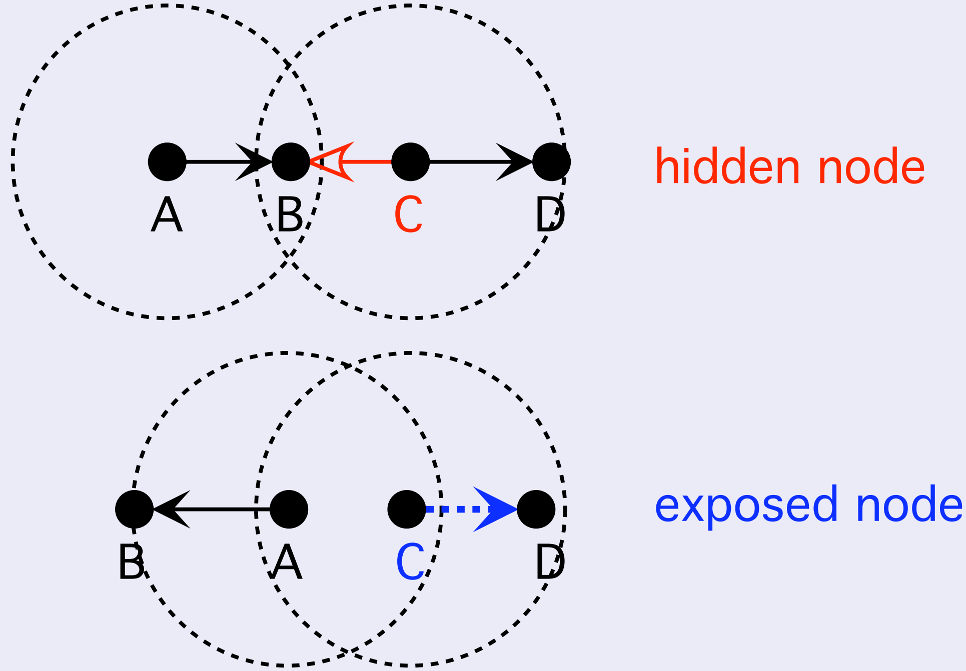

The simplest solution: ALOHA

Let nodes transmit independently with probability p.

Performance analysis and protocol design

We need to analyze first the interference and the outage probabilities and then determine optimum protocols (MAC, routing).

Abstraction

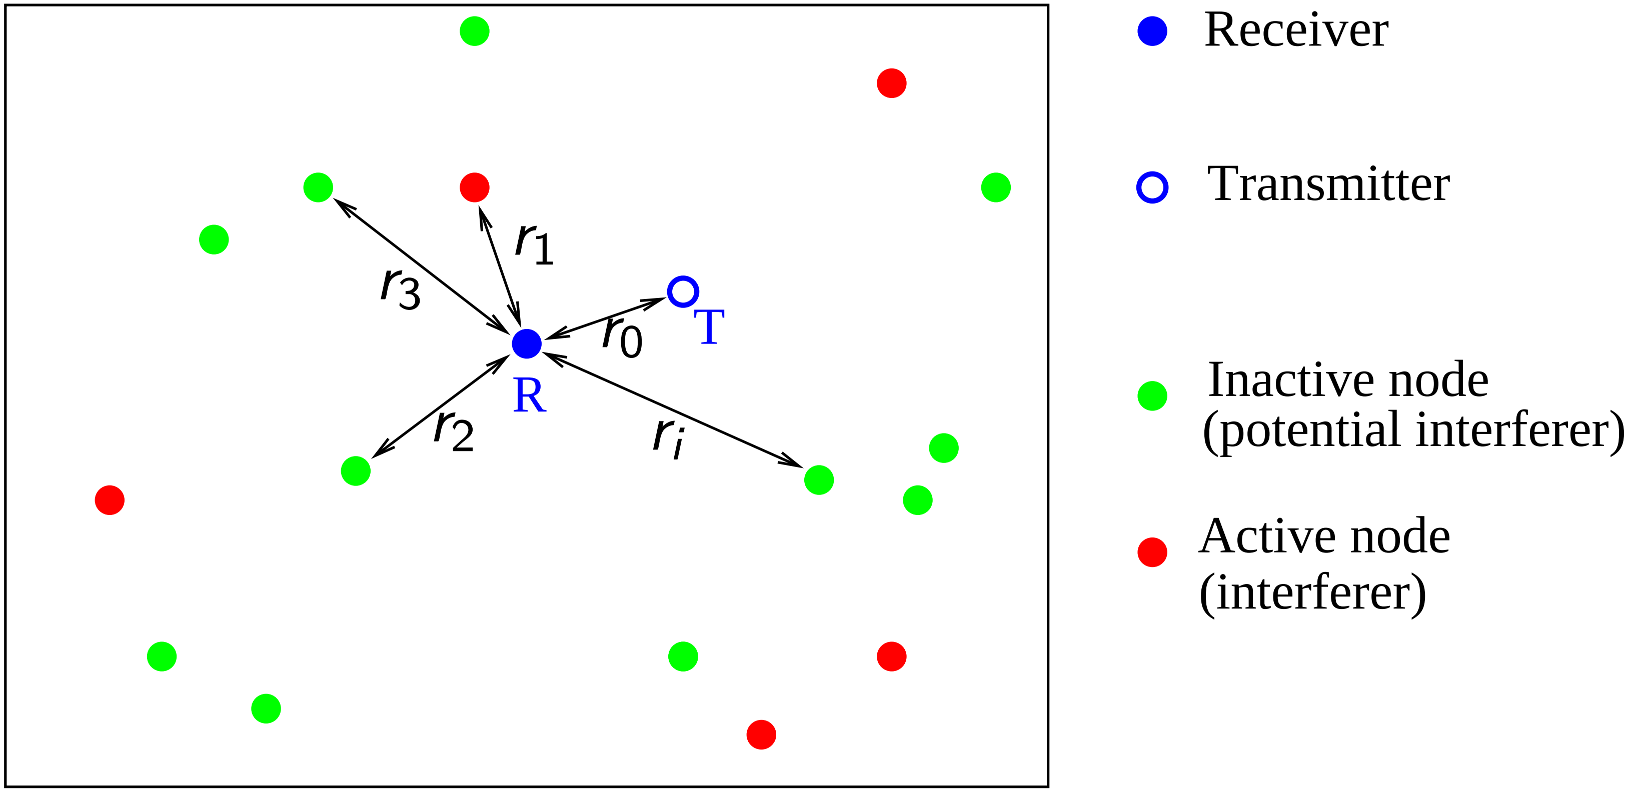

Questions of interest

- What is the interference at R? How likely is the transmission T$\longrightarrow$ R to succeed?
- How to regulate channel access and perform routing?

Key Result

Setup

- Infinitely many nodes are distributed randomly on the plane, with density λ, and a fraction p transmits, all at power P.
- Channels are Rayleigh fading.
- Path loss is $g(r) = r^{-\alpha}$.
- Noise power is W.
- The transmission is successful if the signal-to-interference-plus-noise ratio exceeds a threshold θ.

Result

Probability of successful transmission over distance R [BBM06]:

$$p_s = \mathbb{P}(\mathrm{SINR} > \theta) = \exp\left(-p\lambda\theta^{2/\alpha} C(\alpha) R^2 - \theta W R^\alpha / P\right) ,$$

where $C(\alpha) = \pi\Gamma(1 - 2/\alpha)\Gamma(1 + 2/\alpha) = 2\pi^2/(\alpha \sin(2\pi/\alpha))$.

Section Outline

1 Introduction

2 Modeling
- Propagation and physical layer model
- Uncertainty cube
- Network model

3 Interference and Outage in Poisson Networks

4 Summary

Propagation and Physical Layer

Path loss and fading

If a node transmits at power P over a distance r, the received power is

$$S = Phg(r)\,,$$

where:

- $g(r)$ is the large-scale (or mean) path loss law, assumed monotonically decreasing. Typically $g(r) = r^{-\alpha}$, where α is the path loss exponent.
- h is the power fading coefficient. We always have $\mathbb{E}h = 1$.
  We usually assume a *block fading model*, where h changes from one transmission to the next.
  Often we consider Rayleigh fading, where h is exponential:

$$F_h(x) = 1 - \exp(-x)\,, \quad x \geqslant 0.$$

  The amplitude $\sqrt{h}$ is Rayleigh distributed.

SINR

With thermal noise of variance W, the signal-to-noise ratio (SNR) is $S/W = Phg(r)/W$.
The interference I is the cumulative power from all undesired transmitters.

$$I = \sum_{i \in \mathcal{I}} P_i h_i g(r_i) .$$

This leads to the signal-to-interference-plus-noise ratio (SINR)

$$\mathrm{SINR} = \frac{Phg(r)}{W + I} .$$

Model for transmission success

$$p_s \triangleq \mathbb{P}(\mathrm{SINR} > \theta) .$$

The (bandwidth-normalized) rate of transmission is smaller than (but close to) $\log_2(1+\theta)$ (bits/s/Hz).

Example (Rayleigh block fading with power path loss law)

With k interferers at known distances r_i and path loss law $r^{-\alpha}$:

$$p_s(r) = \mathbb{P}(S > \theta(W+I)) = \underbrace{\exp\left(-\frac{\theta W}{P} r^{\alpha}\right)}_{p_s^N} \cdot \underbrace{\prod_{i=1}^{k} \frac{1}{1+\theta \frac{P_i}{P}\left(\frac{r}{r_i}\right)^{\alpha}}}_{p_s^I}$$

Proof

Let $S = Phr^{-\alpha}$ be the received power, $\bar{S} = Pr^{-\alpha}$, and $I = \sum_{i=1}^{k} P_i h_i r_i^{-\alpha}$.

$$\begin{aligned} P_s = \mathbb{P}[S > \theta(W+I)] &= \mathbb{E}_I\left\{\exp\left(-\frac{\theta(I+W)}{\bar{S}}\right)\right\} \\ &= \exp\left(-\frac{\theta W}{Pr^{-\alpha}}\right) \cdot \mathbb{E}_I\left\{\exp\left(-\frac{\theta I}{\bar{S}}\right)\right\} \end{aligned}$$

These are Laplace transforms! $p_s = \mathcal{L}_W(\theta r^{\alpha}/P) \cdot \mathcal{L}_I(\theta/\bar{S})$.

Remarks

- In a wireless network, there is a lot more uncertainty than fading: k, r_i, perhaps P_i. There is a need to model uncertainty in the locations of the nodes.
- Let I_1 denote the interference at the receiver. We have

$$\mathrm{SINR}_1 = \frac{Phg(r)}{W + I_1}\,.$$

  Now assume all nodes scale their power by a factor a. Then $I_a = aI_1$, and

$$\mathrm{SINR}_a = \frac{aPhg(r)}{W + I_a} = \frac{Phg(r)}{W/a + I_1}$$

  So, increasing the power improves the SINR, since the noise power W is reduced by a.
- The noise term $\exp(-\theta W r^\alpha / P)$ is less interesting, so we often focus on the SIR only.

The Uncertainty Cube

Three dimensions of uncertainty

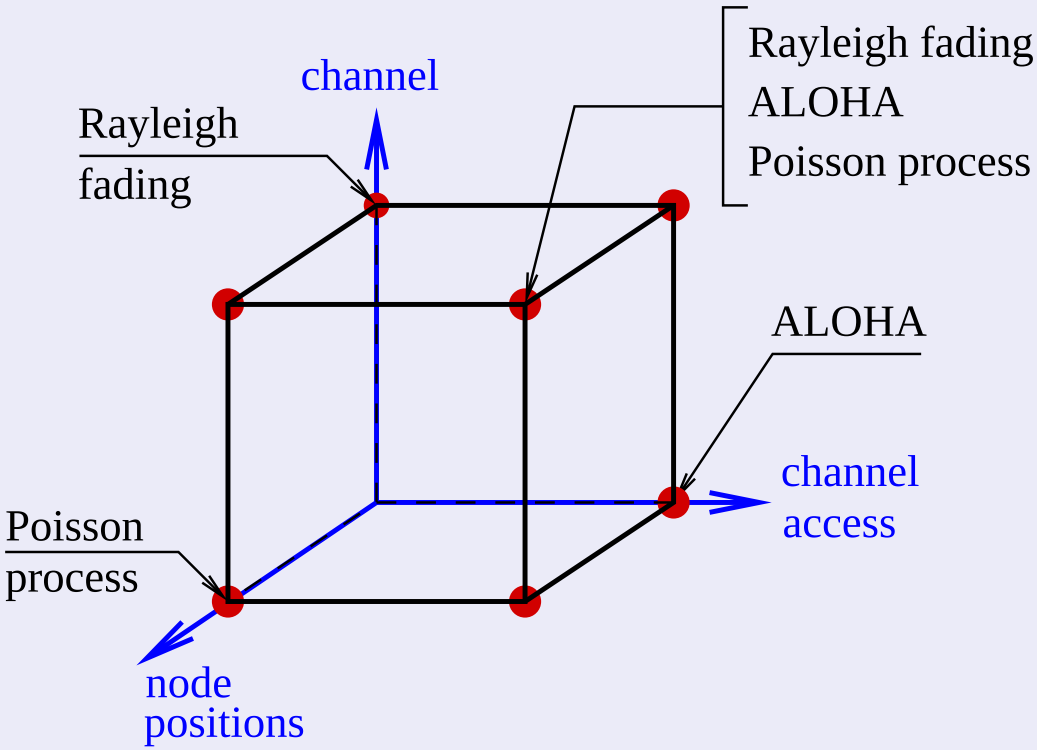

The interferer geometry is determined by the point process (node distribution) and the MAC scheme.

- The goal is to characterize the *average* network, using suitable averaging over the uncertainty.
- We will focus on some of the corner points in Lectures 1 and 2 and talk about the interior in Lecture 6.

Node Locations

The Poisson point process

The (homogeneous) Poisson point process (PPP) is the reference model for the distribution of nodes in a wireless network.

Example (PPP of intensity λ)

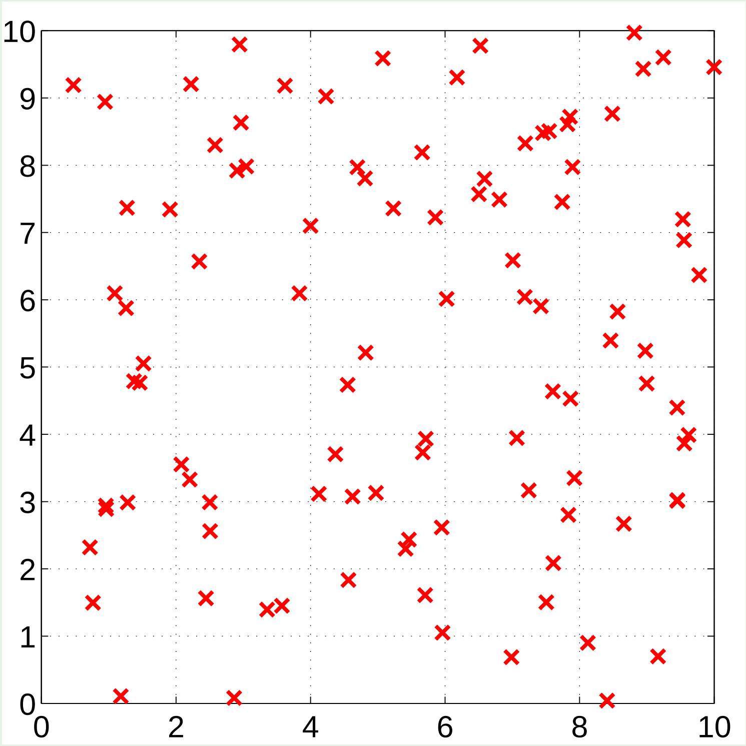

Take a Poisson process $\Phi = \{x_1, x_2, \ldots\}$ of constant intensity λ in a square or disk of area A.

Often, $A \to \infty$ to avoid boundary issues (or use toroidal boundary conditions).

Properties of the PPP

Let $A_i \subset A$ be non-overlapping areas. Then

1. $$\mathbb{P}[\Phi(A_i) = n] = \exp(-\lambda|A_i|)\frac{(\lambda|A_i|)^n}{n!}$$
   where $\Phi(\cdot)$ is the counting measure (number of points), $|\cdot|$ the Lebesque measure (area), and λ is the density of the (homogeneous) PPP.
2. $\Phi(A_1)$, $\Phi(A_2)$, ... are independent.

One of these properties is actually a consequence of the other.

Rényi (1967)

The Poisson distribution implies independence. This is not the case if Property 1 only holds for convex A_i.

Advantages of the PPP model

- Often viewed as worst case (maximum entropy).
- Analytical tractability. Due to the independence property, conditioning on having a point somewhere does not affect the rest of the network. As a consequence, there are many nice results (distances, interference, outage, percolation, connectivity, coverage, ...).

Disadvantages of the PPP model

- Independence property $\Longleftrightarrow$ zero interaction between nodes' positions. This is not a good model for many networks.
- Often, the set of transmitters or active nodes is to be modeled. Except for pure ALOHA, the transmitters do not form a PPP even if the underlying process containing all nodes is Poisson.

Simple result for the PPP: Internode distances [Hae05]

The pdf of the distance to the n-th nearest neighbor is

$$p_{R_n}(r) = r^{2n-1}(\lambda\pi)^n \frac{2}{\Gamma(n)} \exp(-\lambda\pi r^2)\,, \quad r \geqslant 0\,.$$

For $n = 1$ (nearest neighbor), this is a Rayleigh distribution. The mean distance is $1/(2\sqrt{\lambda})$. It does not matter whether we measure from an arbitrary point of the plane or from a point of the process.

Proof

The probability that the n-th nearest neighbor of a point is further away than r is the probability that there are less than n points in the area $\lambda\pi r^2$.

$$\mathbb{P}[R_n > r] = \exp(-\lambda\pi r^2) \sum_{k=0}^{n-1} \frac{(\lambda\pi r^2)^k}{k!}$$

This is the ccdf, so $p_{R_n}(r) = -\mathrm{d}\mathbb{P}[R_n > r]/\mathrm{d}r$

Channel Access

ALOHA

In ALOHA, each node makes the decision to transmit independently and randomly: In each time slot, each node decides to transmit with probability p and to stay quiet (listen) with probability $1-p$. This is (slotted) ALOHA.

CSMA

With carrier sensing, there is a minimum separation between concurrent transmitters. This usually also imposes a minimum spacing between *receivers* and interferers.

Deterministic scheduling

A better performance can usually be achieved if nodes' transmissions are scheduled deterministically. This is referred to as time division multiple access (TDMA). But TDMA requires a centralized scheduler.

Section Outline

1 Introduction

2 Modeling

3 Interference and Outage in Poisson Networks
- Setup
- Mapping of a PPP
- Mean interference
- Laplace transform
- Outage for Rayleigh fading
- Effect of individual interferers
- Effect of path loss law

4 Summary

Setup

PPP plus a desired transmitter

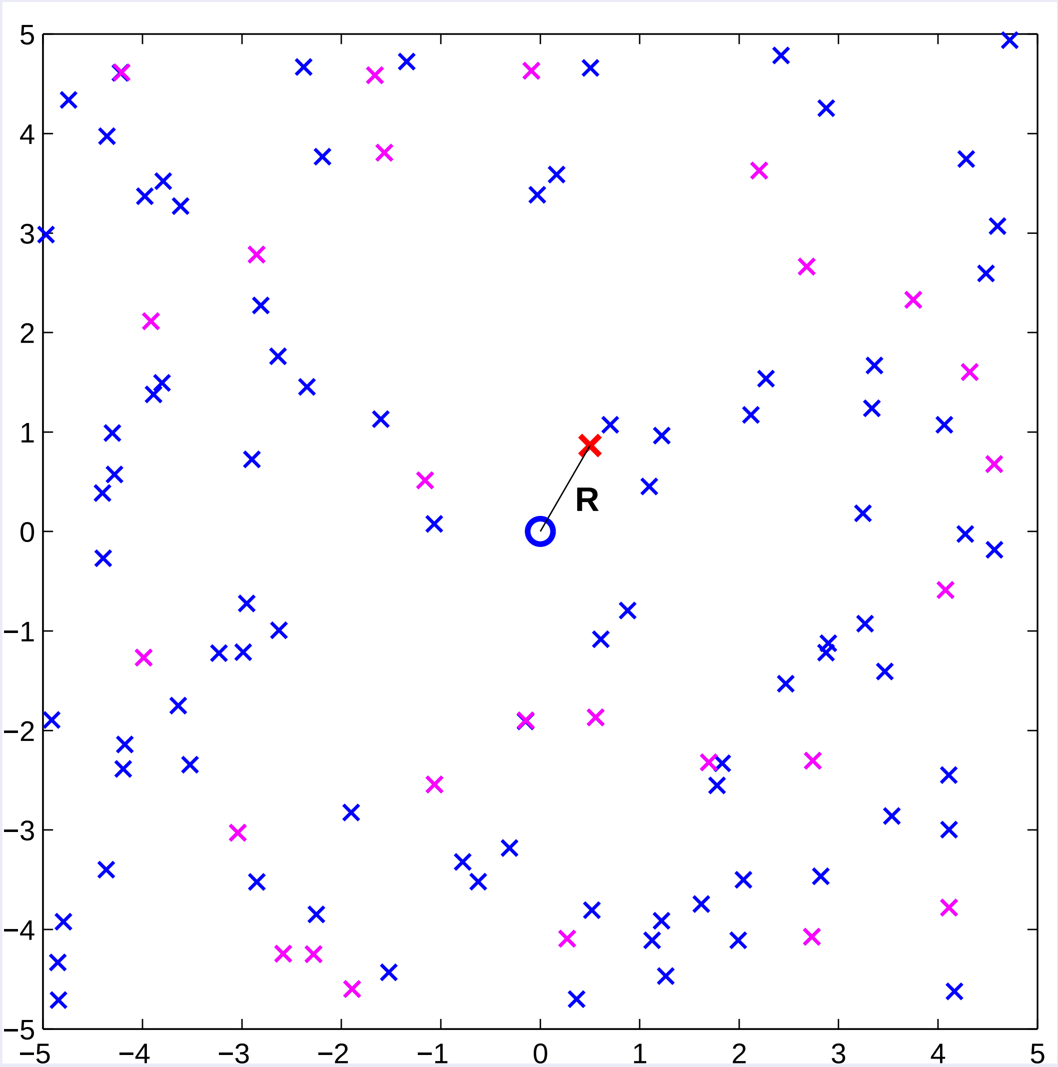

- PPP {x,x} of intensity λ.
- ALOHA with probability p. Active nodes (interferers) x form a PPP of intensity λp.
- o: Receiver under consideration, assumed at origin.
- x: Desired transmitter (not part of the point process), at distance R.

Questions: interference at o? Success prob. $\mathbb{P}(\mathrm{SINR} > \theta)$?

Intensity measure

For a point process $\Phi = \{x_1, x_2, \ldots\}$, the number of nodes in a subset $A \subset \mathbb{R}^2$ is

$$\Phi(A) \triangleq |\Phi \cap A| = \sum_{x \in \Phi} \mathbf{1}(x \in A)\,.$$

$\Phi(A)\colon \mathbb{R}^2 \to \mathbb{N}_0$ is a counting measure for the number of points in A.
The intensity or mean measure is the expected number of nodes $\mathbb{E}\Phi(A)$:

$$\Lambda(A) \triangleq \mathbb{E}\Phi(A)\,.$$

For a stationary PPP of intensity λ: $\Lambda(A) = \lambda|A|$.

Non-homogeneous PPP

A PPP whose intensity depends on the location is a non-homogeneous or non-stationary PPP. In this case,

$$\Lambda(A) = \int_A \lambda(x)\mathrm{d}x\,.$$

The number of points in A is Poisson with mean $\Lambda(A)$.

Mapping

Let $f : \mathbb{R}^2 \to \mathbb{R}$ be a mapping function. Then $\Phi^* = \{f(x_1), f(x_2), \ldots\}$ is a one-dimensional point process. By the mapping theorem [Kin93], Φ^* is a Poisson process with $\Lambda^*(B) = \Lambda(f^{-1}(B))$.

The one-dimensional PPP of distances

Let Φ be stationary with intensity λ and $f(x) = \|x\|$. For $B = [0, r]$, $f^{-1}(B) = b(o, r)$, the ball of radius r at the origin. We obtain

$$\Lambda^*(B) = \Lambda(b(o, r)) = \lambda \pi r^2$$

and

$$\lambda^*(r) = 2\lambda\pi r\,, \quad r \geq 0\,.$$

So the distances of the points of a PPP that is homogeneous on the plane form a non-homogenous PPP on $\mathbb{R}^+$ with linearly increasing density.
The *squared distances* $\{\|x_1\|^2, \|x_2\|^2, \ldots\}$ form again a homogeneous PPP, with intensity $\lambda^* = \lambda\pi$.

Analogous: If x is uniformly randomly distributed on the disk $b(o, R)$, the radius $\|x\|$ has probability density $f_{\|x\|}(r) = 2r/R^2$, $0 \leq r \leq R$.

Example (Mapping from $x \in \mathbb{R}^2$ to $\|x\| \in \mathbb{R}$)

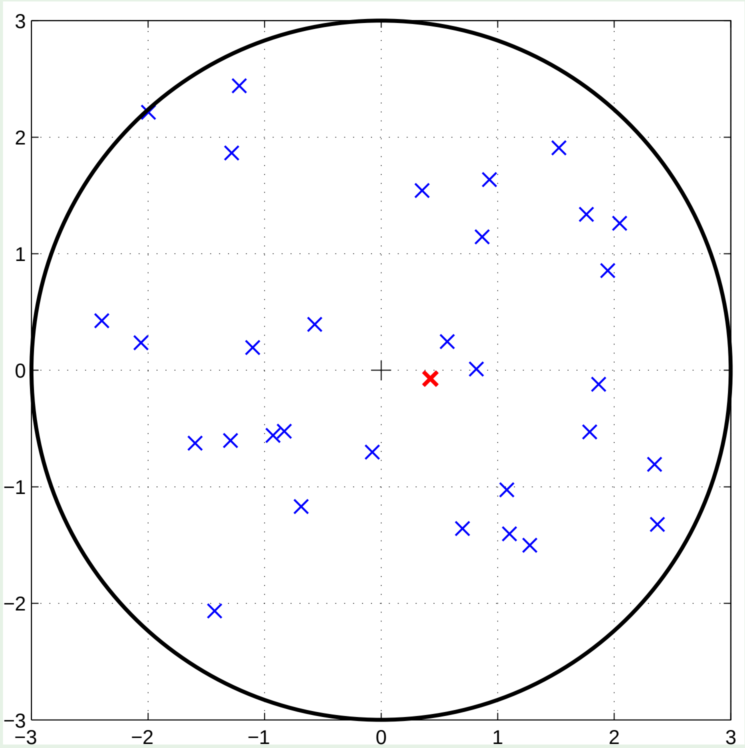

PPP on $b(o,3)$ with $\lambda = 1$. $\Lambda(b(o,3)) = \pi 3^2 \approx 28.2$.

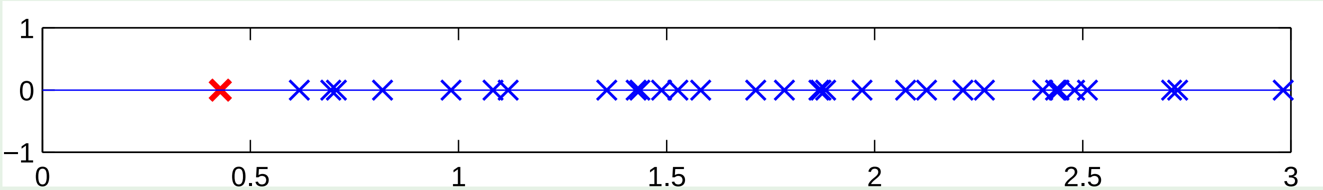

PPP on $[0,3]$ with $\lambda^*(r) = 2\pi r$. $\Lambda^*([0,3]) = \Lambda(b(o,r))$.

Interference

Assume the transmitting nodes form a stationary PPP Φ of intensity λ in $\mathbb{R}^2$. All nodes transmit at unit power, and the path loss $g(r) = r^{-\alpha}$.
Interference at origin:

$$I \triangleq \sum_{x \in \Phi} h_x \|x\|^{-\alpha} .$$

Due to the stationarity of the PPP, the *distribution* of I is the same everywhere.
Equivalently,

$$I = \sum_{r \in \Phi^*} h_r r^{-\alpha} ,$$

where $\Phi^* = \{\|x_1\|, \|x_2\|, \dots\}$ is the PPP of the distances.

Mean interference

Since $\mathbb{E}h = 1$,

$$\mathbb{E}(I) = \mathbb{E}\sum_{r\in\Phi^*} h_r r^{-\alpha} = \mathbb{E}\sum_{r\in\Phi^*} r^{-\alpha}\,.$$

Campbell's theorem

Let f be non-negative function. Then

$$\mathbb{E}\sum_{x\in\Phi} f(x) = \int_{\mathbb{R}^d} f(x)\Lambda(\mathrm{d}x)\,.$$

In our case,

$$\mathbb{E}\sum_{r\in\Phi^*} r^{-\alpha} = \int_{\mathbb{R}^+} r^{-\alpha} 2\pi\lambda r \mathrm{d}r = \frac{2\pi\lambda}{2-\alpha} r^{2-\alpha}\Big|_0^\infty\,, \quad \alpha \neq 2\,.$$

This diverges for all α! So $\mathbb{E}(I) = \infty$.

Mean interference

From previous slide: $\mathbb{E}(I)=\frac{2\pi\lambda}{2-\alpha}r^{2-\alpha}\Big|_0^\infty$.

If $\alpha<2$, the upper integration bound is the culprit. There is too much interference from all the far nodes.
If $\alpha>2$, the lower integration bound is the culprit. The nodes near the origin make $\mathbb{E}(I)$ diverge, since $r^{-\alpha}$ grows too quickly as $r\downarrow 0$ if $\alpha>2$.

A bounded path loss model would solve the problem for $\alpha>2$. More on that later.
Similarly, if it can be ensured that no node is close to the origin, $\mathbb{E}(I)$ remains finite for $\alpha>2$. Replace the lower integration bound by $\rho>0$ to obtain

$$\mathbb{E}(I)=\frac{2\pi\lambda}{\alpha-2}\rho^{2-\alpha}\,.$$

This can be used to model CSMA.

Laplace transform

We would like to find the Laplace transform of I to learn more about the interference.

$$\mathcal{L}_I(s) = \mathbb{E}(e^{-sI}) = \mathbb{E}_{\Phi^*,h}\left(e^{-s\sum_{r\in\Phi^*} h_r r^{-\alpha}}\right)$$
$$= \mathbb{E}_{\Phi^*}\Big(\prod_{r\in\Phi^*} \underbrace{\mathbb{E}_h(e^{-sh_r r^{-\alpha}})}_{v(r)}\Big)\,.$$

We need to calculate the expectation of a product over the point process. In the Poisson case, this is easy, thanks to the probability generating functional: For functions $v\colon \mathbb{R}^d \mapsto [0,1]$,

$$\mathbb{E}\prod_{x\in\Phi^*} v(x) = G[v]\,.$$

Probability generating functional (pgfl)

The pgfl of a PPP Φ with intensity measure Λ is

$$G[v] = \mathbb{E} \prod_{x \in \Phi} v(x) = \exp\left(- \int_{\mathbb{R}^2} [1 - v(x)] \Lambda(\mathrm{d}x) \right) .$$

Laplace transform

Applied to the interference:

$$\mathcal{L}_I(s) = G[v] = \exp\left(- \int_{\mathbb{R}^+} [1 - \mathbb{E}_h(e^{-sh_r r^{-\alpha}})] \Lambda^*(\mathrm{d}r) \right) ,$$

where

$$\Lambda^*(\mathrm{d}r) = \lambda^*(r)\mathrm{d}r = 2\pi\lambda r \mathrm{d}r .$$

Laplace transform

From previous slide: $\mathcal{L}_I(s) = \exp\left(-\pi\lambda \int_{\mathbb{R}^+} 2[1 - \mathbb{E}_h(e^{-sh_r r^{-\alpha}})] r \mathrm{d}r\right)$

Swapping expectation and integral and conditioning on h, the integral is:

$$
\begin{aligned}
2\int_0^\infty \left[1-\exp(-shr^{-\alpha})\right] r \mathrm{d}r &\overset{(a)}{=} \int_0^\infty \left[1-\exp(-sh/y)\right] \delta y^{\delta-1} \mathrm{d}y \\
&\overset{(b)}{=} \int_0^\infty \left[1-\exp(-shx)\right] \delta x^{-\delta-1} \mathrm{d}x \\
&\overset{(c)}{=} \int_0^\infty x^{-\delta} sh \exp(-shx) \mathrm{d}x \\
&= (sh)^\delta \Gamma(1-\delta)\,, \quad 0<\delta<1.
\end{aligned}
$$

(a): $y \leftarrow r^{1/\alpha}$, $\delta = 2/\alpha$. (b): $x \leftarrow y^{-1}$. (c): Integration by parts.

Laplace transform

Taking the expectation over h, we have

$$\mathcal{L}_I(s) = \exp\left(-\lambda\pi\mathbb{E}(h^\delta)\Gamma(1-\delta)s^\delta \right), \quad 0 < \delta < 1\,.$$

- The interference has a stable distribution with characteristic exponent δ and dispersion $\lambda\pi\mathbb{E}(h^\delta)\Gamma(1-\delta)$.
- If $\delta \uparrow 1$ (or $\alpha \downarrow 2$), we have $\mathcal{L}_I(s) \downarrow 0$, for all $s > 0$, so $I \uparrow \infty$ almost surely (a.s.). So we need $\alpha > 2$ for finite interference.
- I does not have any finite moments.
- The interference (power) is very far from Gaussian. The *amplitude* may be, though.
- For ALOHA with probability p, replace λ by λp.

Interference distribution

Closed-form expressions for the distribution only exist for $\alpha = 4$. In this case, without fading,

$$\mathcal{L}_I(s) = \exp\left(-\lambda\sqrt{s}\frac{\pi^2}{2}\right).$$

The corresponding distribution is the Lévy distribution

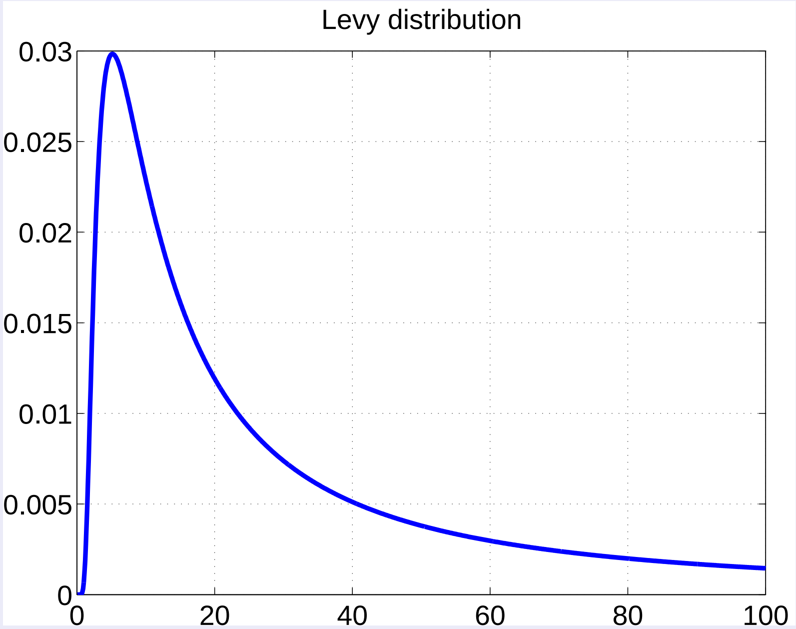

$$f_I(y) = \frac{\pi}{2}\lambda y^{-3/2} e^{-\pi^3\lambda^2/4y}, \qquad F_I(y) = 1 - \operatorname{erf}\left(\frac{\pi^{3/2}\lambda}{2\sqrt{y}}\right)$$

for $y \geqslant 0$.

It has a heavy tail, as expected from the fact that it does not have a mean.

Remarks on interference distribution

- This interference is not Gaussian, even for $\lambda \to \infty$.
- The tail of a Lévy-stable density with $\delta < 2$ decays like a power function.

$$f_I(y) \sim \lambda\pi\delta\,\mathbb{E}(h^\delta)y^{-(1+\delta)}\,,\quad y \to \infty$$

  The long tail is due to the singularity of the path loss law.
  Equivalently, $\mathbb{P}[I > x] = \Theta(x^{-\delta})$ as $x \to \infty$.
- Let $\tilde{I}$ denote the interference without the closest interferer. The distribution of $\tilde{I}$ has a different tail of smaller order:

$$\mathbb{P}[\tilde{I} > x] = o(x^{-\delta}) \quad \text{as } x \to \infty\,.$$

- The type of fading is irrelevant, only $\mathbb{E}(h^\delta)$ matters.

Laplace transform for Rayleigh fading

Since $\mathbb{E}(h^\delta) = \Gamma(1+\delta)$:

$$\mathcal{L}_i(s) = \exp\left(-\lambda\pi\Gamma(1+\delta)\Gamma(1-\delta)s^\delta\right) = \exp\left(-\lambda\pi s^\delta \frac{\pi\delta}{\sin(\pi\delta)}\right) .$$

Outage for Rayleigh fading and ALOHA

For a transmission over distance R, the received signal power is $S = hR^{-\alpha}$. The success probability is

$$p_s = \mathbb{P}(hR^{-\alpha} > I\theta) = \mathbb{E}(e^{-\theta R^\alpha I}) = \exp\left(-\lambda\pi\Gamma(1+\delta)\Gamma(1-\delta)\theta^\delta R^2\right)$$

with $\delta = 2/\alpha$. The Laplace transform gives us the success probability in Rayleigh fading! This is our key result!
So while we do not have a distribution of I in general, we have a distribution of the SIR: $p_s = \mathbb{P}(\text{SIR} > \theta)$ is its ccdf.

Interference from the nearest node

$$\mathbb{P}(I_1 \leqslant x) = \mathbb{P}(R^{-\alpha} \leqslant x) = \mathbb{P}(R > x^{-1/\alpha}) = \exp(-\lambda\pi x^{-\delta}) \,,$$

where $\delta = 2/\alpha$. We get

$$\mathbb{E}I_1 = \pi^{1/\delta}\Gamma(1 - 1/\delta) \,, \quad \delta > 1 \,.$$

If $\delta < 1$ then $\mathbb{E}(I_1)$ does not exist. Generally, $\mathbb{E}(I_1^p)$ exists for $p < \delta$ since

$$\mathbb{P}(I_1 > x) \sim \lambda\pi x^{-\delta} \qquad x \to \infty \,.$$

Interference from n-nearest node

The ccdf of the distance to the n-th nearest neighbor R_n is

$$\mathbb{P}(R_n > r) = \frac{\Gamma_{\mathrm{ic}}(n, \lambda\pi r^d)}{\Gamma(n)} \,.$$

So for $n = 2$,

$$\mathbb{P}(I_2 < x) = \exp(-\lambda\pi x^{-\delta})(1 + \lambda\pi x^{-\delta})$$

and

$$\mathbb{P}(I_2 > x) \sim \frac{1}{2}(\lambda\pi)^2 x^{-2\delta} \,.$$

So we need $2\delta > 1$ for $\mathbb{E}I_2$ to exist.

Interference from n-nearest node (cont.)

For general n:

$$\mathbb{P}(I_n < x) = \exp(-\lambda\pi x^{-\delta}) \sum_{i=0}^{n-1} \frac{(\lambda\pi x^{-\delta})^i}{i!}$$

For the tail probability we need to sum from n to ∞, so the dominant term will be the one for $i = n$ as $x \to \infty$. Therefore

$$\mathbb{P}(I_n > x) \sim \frac{1}{n!}(\lambda\pi)^n x^{-n\delta} .$$

This means that $\mathbb{E}(I_n^p)$ exists for $p < n\delta$. For example, we would need to cancel $k > \alpha$ interferers to have a finite second moment.
On the other hand, canceling all nodes within distance $\epsilon > 0$ or using a bounded path loss law would ensure finite moments.

Effect of Path Loss Law

Interference with bounded path loss

Consider a path loss law $g(r) = \min\{1, r^{-\alpha}\}$, and let the diameter of the network be $D > 1$. So we consider a PPP of intensity λ on $b(o, D)$.
In this case, from Campbell's theorem,

$$\begin{aligned}\mathbb{E}(I_D) &= \int_{\mathbb{R}^+} g(r) 2\lambda\pi r \mathrm{d}r = \int_0^1 2\lambda\pi r \mathrm{d}r + \int_1^D r^{-\alpha} 2\lambda\pi r \mathrm{d}r \\ &= \lambda\left(\pi + \frac{2\pi}{\alpha - 2}(1 - D^{2-\alpha})\right).\end{aligned}$$

The Laplace transform can also be calculated; it involves incomplete gamma functions. Moments can be obtained by

$$\mathbb{E}(I^m) = (-1)^m \frac{\mathrm{d}^m}{\mathrm{d}s^m} \log(\mathcal{L}_I(s))\Big|_{s=0}.$$

Section Outline

1 Introduction

2 Modeling

3 Interference and Outage in Poisson Networks

4 Summary
- Interference and outage
- Campbell's theorem and the pgfl

Lecture 1 Summary

Interference and outage

- The boundedness of the path loss model has a drastic impact on the distribution of the interference.
- On the other hand, the success probability is not affected significantly by the path loss model. Since for large interference, there is an outage anyway, the heavy tail does not matter much.
- In fact, the SIR has a very benign distribution. For Rayleigh fading, it is just a Weibull distribution:

$$\mathbb{P}(\mathrm{SIR} \leqslant x) = 1 - \exp(-cx^{\delta}) .$$

The mean is

$$\mathbb{E}(\mathrm{SIR}) = \frac{1}{R^{\alpha}} \frac{\Gamma(1 + 1/\delta)}{(\lambda p C(\alpha))^{1/\delta}} , \qquad C(\alpha) = \pi \Gamma(1 + \delta) \Gamma(1 - \delta) .$$

Outage

Since

$$p_s = \exp(-\lambda\pi R^2\theta^\delta C(\alpha))\,,$$

the success probability equals the void probability that there is no node in $b(o, R')$ with

$$R' = R\theta^{\delta/2}\sqrt{C(\alpha)}\,.$$

Generalization to d dimensions

Most results generalize to d dimensions in a very straightforward manner. There are only two changes needed in all results:

- Replace $\pi\lambda$ by $c_d\lambda$, where c_d is the volume of the d-dimensional unit ball: $c_d = |b(o, 1)| = \pi^{d/2}/\Gamma(1 + d/2)$.
- Replace $\delta = 2/\alpha$ by $\delta = d/\alpha$.

Campbell's theorem

Let f be non-negative function. Then

$$\mathbb{E}\sum_{x\in\Phi} f(x) = \int_{\mathbb{R}^d} f(x)\Lambda(\mathrm{d}x)\,.$$

Example (Number of points in A)

Take a stationary PPP and let $f(x) = \mathbf{1}(x \in A)$. Then we know that

$$\mathbb{E}\sum_{x\in\Phi} f(x) = \mathbb{E}\Phi(A)\,.$$

Campbell's theorem gives

$$\mathbb{E}\sum_{x\in\Phi} f(x) = \lambda\int_A \mathrm{d}x = \lambda|A|\,,$$

as expected.

Probability generating functional (pgfl)

The pgfl of a PPP Φ with intensity measure Λ is

$$G[v] = \mathbb{E} \prod_{x \in \Phi} v(x) = \exp\left(- \int_{\mathbb{R}^2} [1 - v(x)] \Lambda(\mathrm{d}x) \right) .$$

Example (Void probability of A)

Let $v(x) = 1 - \mathbf{1}(x \in A)$. Then $G[v] = 1$ only if all points in Φ are outside of A, *i.e.*, $G[v]$ is the void probability. We have

$$\mathbb{P}(A \text{ empty}) = \mathbb{E} \prod_{x \in \Phi} v(x) = G[v] = \exp(-\lambda \int_A \mathrm{d}x) = \exp(-\lambda |A|) ,$$

as expected from the void probability in the Poisson process.

References I

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.

Martin Haenggi.
On Distances in Uniformly Random Networks.
IEEE Trans. on Information Theory, 51(10):3584–3586, October 2005.
Available at `http://www.nd.edu/~mhaenggi/pubs/tit05.pdf`.

J. F. C. Kingman.
Poisson Processes.
Oxford Science Publications, 1993.

Analysis and Design of Wireless Networks

Lecture 2: Throughput Analysis and Design Aspects

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop analysis of Poisson networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 2 Overview

1. Throughput
2. Other Applications
3. A Geometric Interpretation of Fading
4. Summary

Section Outline

1. Throughput
   - Probabilistic throughput
   - Shannon throughput
   - Spatial Shannon throughput
   - Transmission capacity
   - Comparison with regular networks
2. Other Applications
3. A Geometric Interpretation of Fading
4. Summary

Throughput

Unconditional success probability

Previously we assumed that the desired transmitter transmits and the receiver listens. The *unconditioned success probability* is

$$\text{ALOHA:} \qquad p_T \triangleq p(1-p)\,p_s(p)\,.$$

Optimum ALOHA transmit probability

Let $p_s(p) = e^{-p\gamma}$, where γ is the *spatial contention* [Hae09]. For the PPP with Rayleigh fading, *e.g.*,

$$\gamma = \lambda\pi R^2\theta^\delta \frac{\pi\delta}{\sin(\pi\delta)}\,, \quad \delta = 2/\alpha.$$

We find

$$p_{\rm opt}(\gamma) = \frac{1}{\gamma} - \underbrace{\frac{1}{2}\left(\sqrt{1+\frac{4}{\gamma^2}}-1\right)}_{\text{half-duplex penalty}}\,.$$

Remark on optimum transmit probability p

Consider the success probability as a function of γ with the optimum transmit probability:

$$p_s^{\mathrm{opt}}(\gamma) \triangleq p_s(p_{\mathrm{opt}}(\gamma))$$

The resulting $p_s^{\mathrm{opt}}(\gamma)$ decays quickly with γ. For example, $p_s^{\mathrm{opt}}(\gamma) < 1/2$ for

$$\gamma > \frac{\log 2(2 - \log 2)}{1 - \log 2} \approx 2.9 \,.$$

Although throughput-optimum, such low success probabilities are not acceptable for many applications, and they waste energy. We will get back to this later.

Definition (Probabilistic throughput)

If the transmission rate is $\log(1+\theta)$ (nats/s/Hz), there is an outage if SIR $< \theta$. So for fixed-rate transmission, it is natural to set the rate to $\log(1+\theta)$. The probabilistic throughput is

$$T(\theta) \triangleq p_T(\theta)\log(1+\theta)$$

Optimum SIR threshold for full-duplex operation [Hae09]

The probabilistic throughput $T^f = p\exp(-p\gamma)\log(1+\theta)$ is maximized at the rate (spectral efficiency)

$$R_T^{\mathrm{opt}} = \log(1+\theta^{\mathrm{opt}}) = \mathcal{W}\left(-\frac{1}{\delta}e^{-1/\delta}\right) + \frac{1}{\delta} \quad \text{(nats/s/Hz)},$$

where $\mathcal{W}$ is the principal branch of the Lambert W function, *i.e.*, $\mathcal{W}(x)e^{\mathcal{W}(x)} = x$ with $\mathcal{W}(x) \geq -1$.
Tight bound: $R_T^{\mathrm{opt}}(\alpha) \lessapprox \alpha - 2$.

Optimum threshold and throughput for PPP ALOHA network with Rayleigh fading

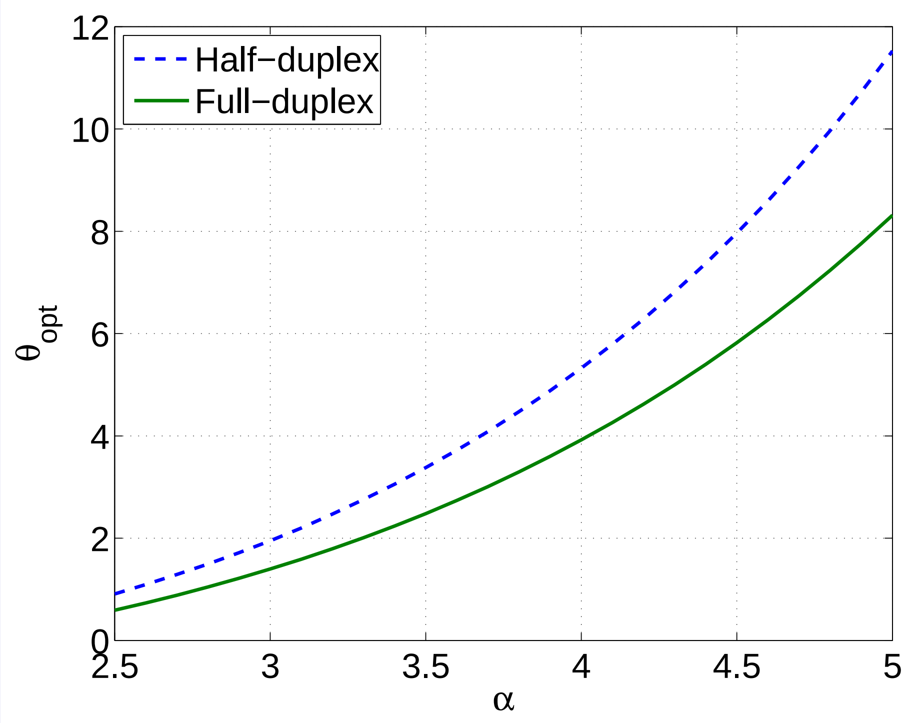

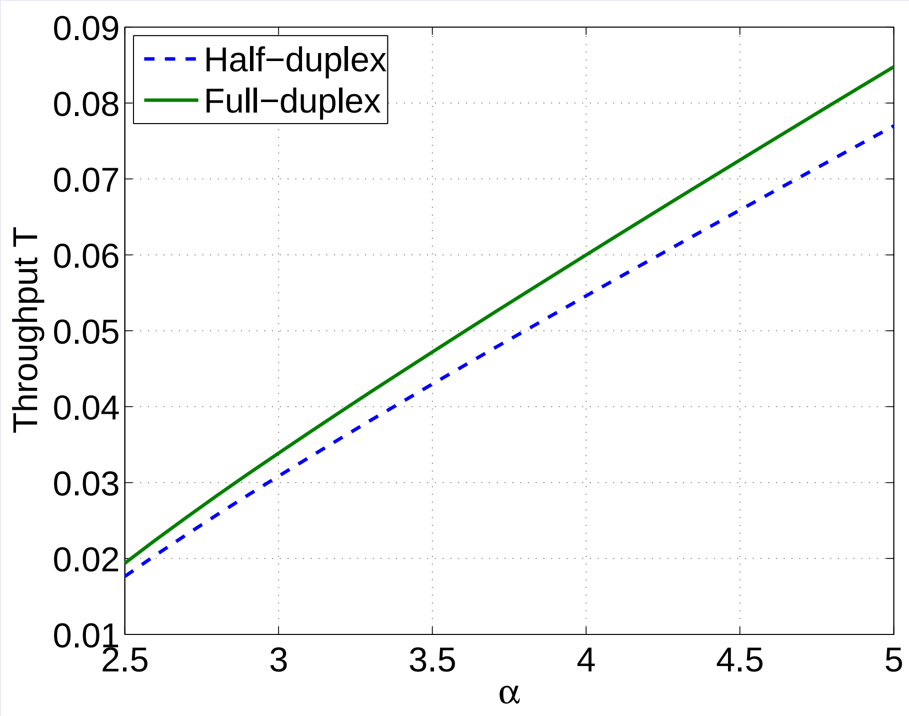

- Transmissions at relatively low rate are optimum.
- The throughput increases linearly with α.
- There is a fixed small penalty factor for half-duplex operation.

Shannon Throughput

Shannon throughput

If the transmitter has full knowledge of S and I, it can adjust its rate of transmission accordingly.

Alternatively, if there is enough time diversity in S and I, *e.g.*, through fast FH, the transmitter can signal at $\mathbb{E}\log(1+\mathrm{SIR})$.

Either way, the resulting throughput is the Shannon throughput $\mathbb{E}\log(1+\mathrm{SIR})$.

Definition (Shannon throughput C)

$$C \triangleq \mathbb{E}\log(1+\mathrm{SIR}) = \int_0^\infty -\log(1+\theta)\mathrm{d}p_s(\theta)\,,$$

where $p_s(\theta)$ is the ccdf of the SIR.

Alternate expression for Shannon throughput

Let $p_s(\theta)$ be the success probability as a function of the SINR threshold. The Shannon throughput can also be calculated as follows:

$$\mathbb{E}\log(1+\mathsf{SIR}) = \int_0^\infty \mathbb{P}(\log(1+\mathsf{SIR}) > \theta)\mathrm{d}\theta$$
$$= \int_0^\infty \mathbb{P}(\mathsf{SIR} > e^\theta - 1)\mathrm{d}\theta$$
$$= \int_0^\infty p_s(e^\theta - 1)\mathrm{d}\theta\,.$$

This is sometimes easier to evaluate or bound.

Example (PPP ALOHA network with Rayleigh fading)

For $\alpha = 4$ and $R = 1$:

$$C = 2\Re\{q\}\cos(p\lambda\pi^2/2) - 2\Im\{q\}\sin(p\lambda\pi^2/2)\,, \qquad q \triangleq \mathrm{Ei}(1, jp\lambda\pi^2/2)\,,$$

where $\mathrm{Ei}(1,z) = \int_1^\infty \exp(-xz)x^{-1}\mathrm{d}x$ is the exponential integral.
For general α:

$$C > \int_1^\infty -\log(\theta)\mathrm{d}p_s(\theta) = \frac{\alpha}{2}\,\mathrm{Ei}(1, p\lambda C(\alpha))\,.$$

Spatial Shannon throughput

Since $C(p) \to \infty$ as $p \to 0$, the Shannon throughput itself does not give insight on how to choose p.
Instead, as before, use the *spatial Shannon throughput* or *throughput density*:

$$p(1-p)C(p)$$

Spatial Shannon throughput for PPP ALOHA network

The *SIR*-based capacity diverges as $p \to 0$, so a better metric is the spatial capacity $p(1-p)C$ (half-duplex) and pC (full-duplex).

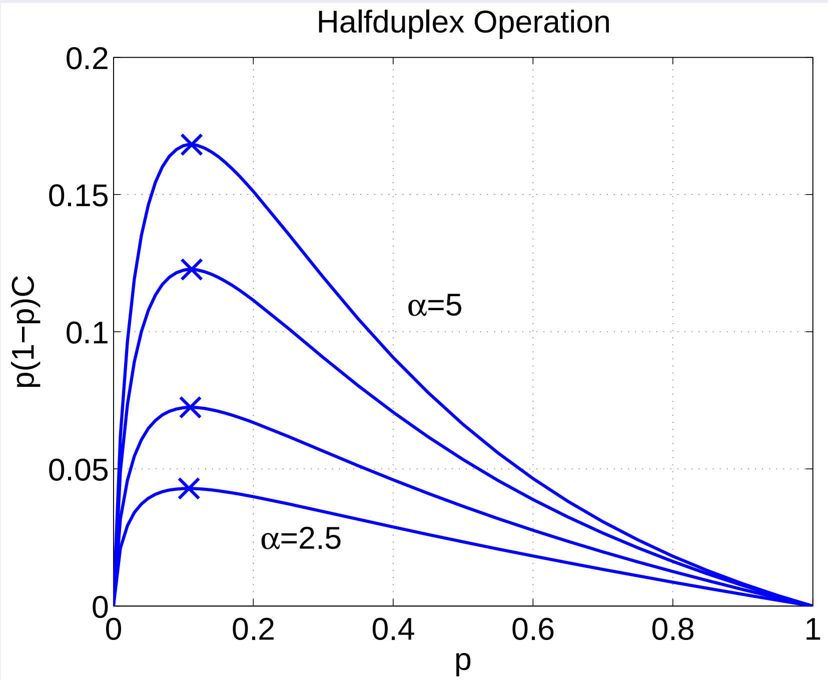

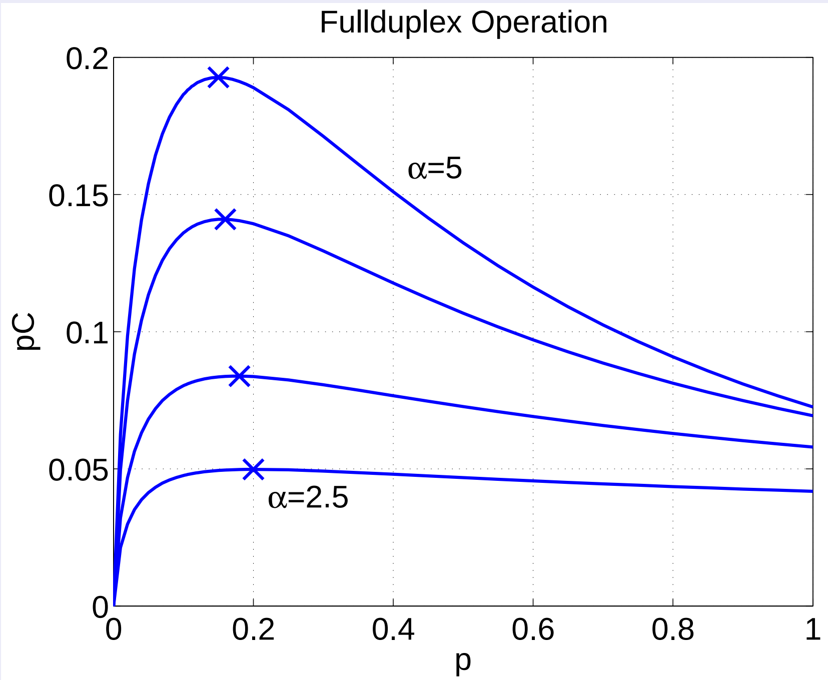

- The optimum p is independent of α. For half-duplex operation, $p_{\text{opt}} \approx 1/9$.
- The maxima are about 2.5 times higher than for the probabilistic throughput. So rate adaptation can result in a significant gain.

Transmission Capacity

Definition (Transmission capacity)

The maximum density of successful transmissions subject to an outage constraint ϵ, multiplied by $1-\epsilon$ [WYAdV05]:

$$\mathrm{TC}(\epsilon) = (1-\epsilon)\sup\{\lambda \colon p_s(\lambda) > 1-\epsilon\}\,.$$

Or, if p_s is viewed as just a function of λ and the inverse is p_s^{-1},

$$\mathrm{TC}(\epsilon) = (1-\epsilon)p_s^{-1}(1-\epsilon)\,.$$

It is available in closed-form whenever p_s is known and invertible.

Advantage over throughput as a metric

The transmission capacity metric imposes a maximum outage probability and is thus very useful in applications that cannot tolerate high packet loss rates.

Example (Transmission capacity for PPP with Rayleigh fading)

In this case, the success probability is invertible, and we have

$$\mathrm{TC}(\epsilon) = (1-\epsilon)\frac{-\log(1-\epsilon)}{R^2\theta^{2/\alpha}C(\alpha)} = \frac{\epsilon}{R^2\theta^{2/\alpha}C(\alpha)} + \Theta(\epsilon^2)\,, \quad \epsilon \to 0\,.$$

This is *linear* in ϵ and inversely proportional to the area of the disk $b(o, R)$.

Transmission capacity in more general settings [WAJ10]

- Approximations for general fading are possible.
- If there is no fading and $\alpha = 4$, the TC follows from the Lévy distribution.
- Considering only the nearest interferer yields an upper bound on the TC. The bound is tight when α is not too close to 2.
- MIMO: With N_t transmit and N_r receive antennas used for diversity,

$$\Omega(\max\{N_t, N_r\}^\delta) = \mathrm{TC}(\epsilon) = O((N_t N_r)^\delta)\,, \quad N_t, N_r \to \infty\,.$$

- SIC with $N_t = 1$ and $N_r > 1$: Canceling the strongest $N_r - 1$ interferers yields $\mathrm{TC}(\epsilon) = \Theta(N_r^{1-\delta})$. This is much better for larger α — but requires knowledge of the interferers' channels.

Line Networks with TDMA

Line network with spatial reuse parameter *m*

Take a regular line network with $\{r_i\} = |\mathbb{Z}|$ and have every m-th node transmit concurrently.

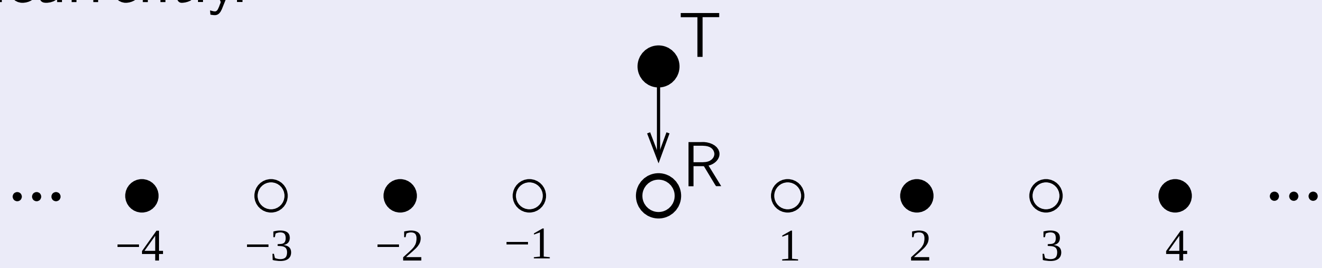

TDMA line network with $m = 2$.

Success probability with Rayleigh fading

$$\alpha = 2: \quad p_s = \left(\frac{y}{\sinh y}\right)^2, \quad \textit{where } y \triangleq \frac{\pi\sqrt{\theta}}{m} \qquad \text{[MM95]}$$

$$\alpha = 4: \quad p_s = \left(\frac{2y^2}{\cosh^2 y - \cos^2 y}\right)^2, \quad \textit{where } y \triangleq \frac{\pi\theta^{1/4}}{\sqrt{2}m}.$$

Probabilistic throughput

The probabilistic throughput is the unconditioned success probability.
ALOHA:

$$\text{full-duplex:} \quad p_T^f \triangleq p\, p_s(p)\,; \qquad \text{half-duplex:} \quad p_T^h \triangleq p(1-p)\, p_s(p)\,.$$

TDMA:

$$p_T \triangleq p_s(m)/m$$

Result for TDMA

From bounds on $p_s(m)$ we can obtain a very good estimate on the optimum m:

$$\hat{m}_{\mathrm{opt}} = \left\lceil \left(\zeta(\alpha)\theta(2\alpha - 1/2)\right)^{1/\alpha} \right\rfloor$$

Theorem (Capacity of TDMA line networks [Hae09])

For $\alpha = 2$,

$$2\log\left(\frac{2m}{\pi}\right) < C < \log\left(1 + \frac{7\zeta(3)}{\pi^2}m^2\right)$$

and

$$\mathbb{E}\sqrt{\mathrm{SIR}} = \frac{\pi}{4}m\,; \qquad \mathbb{E}\mathrm{SIR} = \frac{7\zeta(3)}{\pi^2}m^2\,.$$

For general $\alpha > 1$,

$$C > e^{\zeta(\alpha)/m^\alpha}\,\mathrm{Ei}(1, \zeta(\alpha)/m^\alpha)\,; \quad \mathbb{E}\mathrm{SIR} > \frac{m^\alpha}{\zeta(\alpha)}\,.$$

Comparison with 1-dim. PPP ALOHA network

$$\mathbb{E}\mathrm{SIR}_{\mathrm{ALOHA}} = \frac{\Gamma(1+\alpha)}{(C_1(\alpha)p)^\alpha}\,; \qquad \alpha = 2: \ \frac{\mathbb{E}\mathrm{SIR}_{\mathrm{TDMA}}}{\mathbb{E}\mathrm{SIR}_{\mathrm{ALOHA}}} \approx 4.2$$

TDMA line network capacity

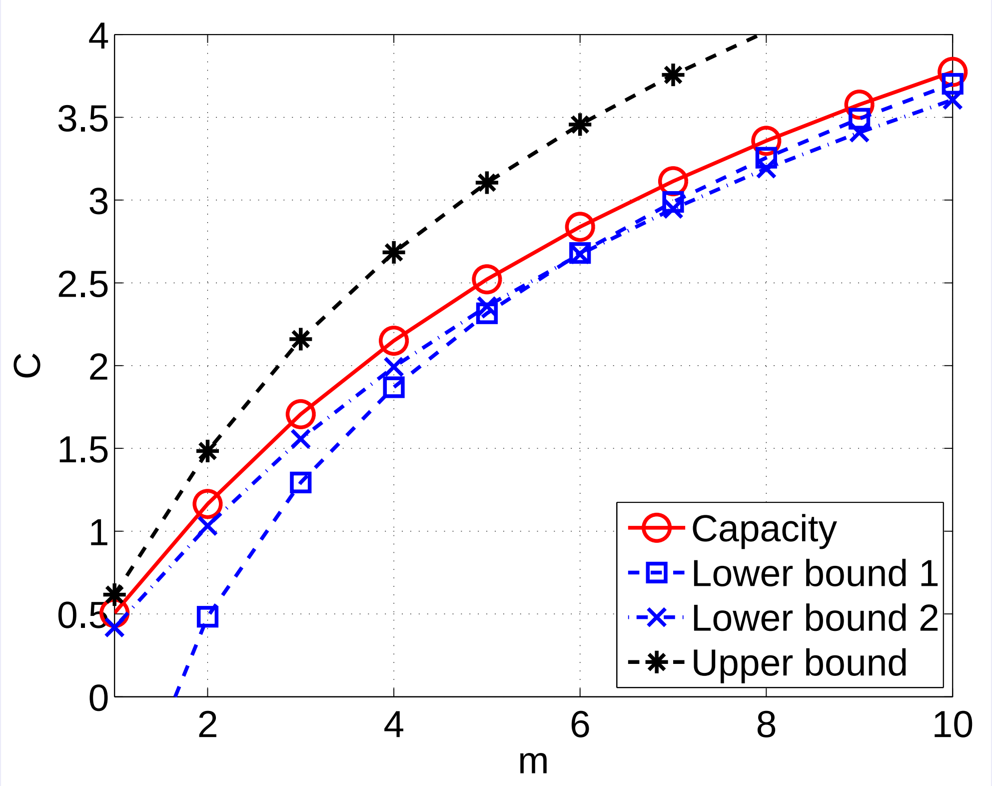

Link capacity of TDMA line network for $\alpha = 2$.

The maximum spatial capacity $C/m \approx 0.6$ is achieved at $m = 2$. This is about 75% higher than the (optimized) PPP line network with ALOHA. For $\alpha = 4$, $m = 3$ achieves a slightly higher capacity.

Section Outline

1 Throughput

2 Other Applications
- CSMA
- Spread-spectrum techniques
- Opportunistic ALOHA
- Power control
- Bandwidth partitioning

3 A Geometric Interpretation of Fading

4 Summary

Other Applications

CSMA emulation using guard zone

With a path loss law $g(r) = r^{-\alpha}\mathbf{1}(r > \rho)$, the success probability with CSMA can be calculated as a function of the guard zone radius ρ.

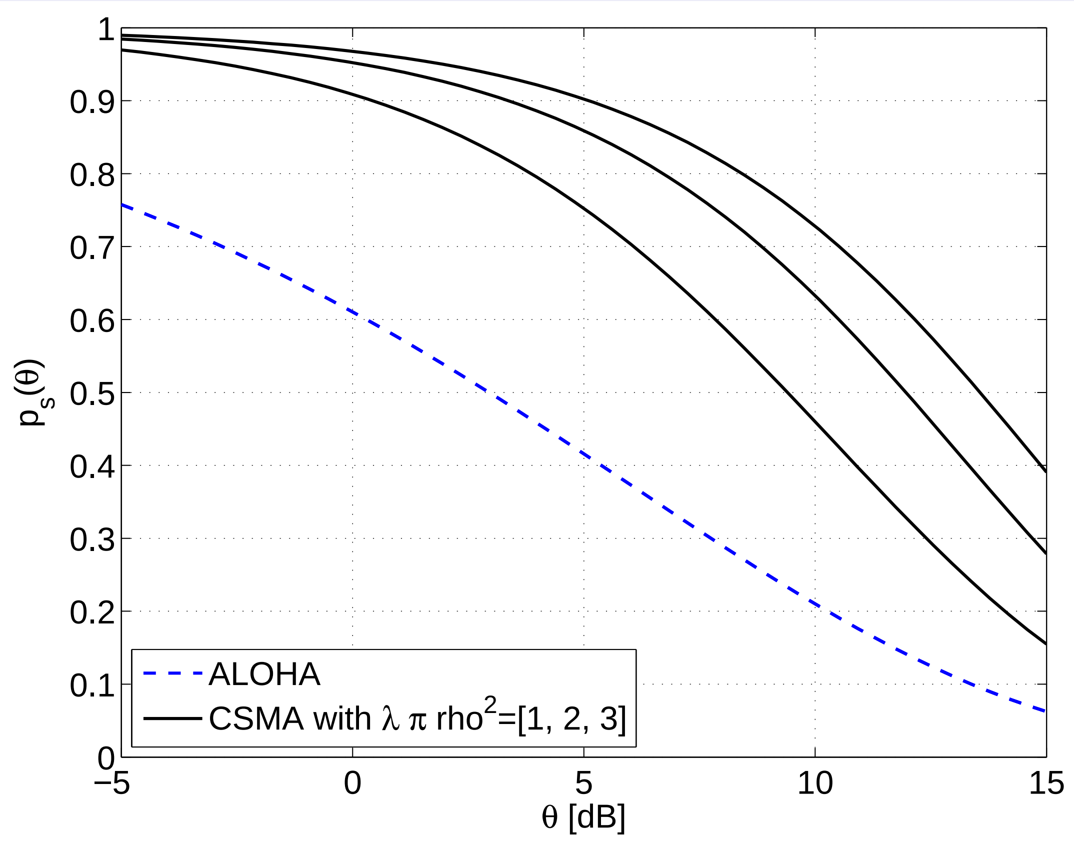

PPP with $\lambda = 1/10$ and $\alpha = 4$.

The exclusion radii are chosen such that 1, 2, and 3 interferers are muted on average.

Direct-sequence vs. frequency hopping spread spectrum [AWH07]

Consider spreading signals by a factor M, either using DS-SS or FH-SS. With DS-SS, the density of interferers stays the same, but the *interference is reduced by a factor M*.

$$p_s^{\mathrm{DS}}(\theta, M) = \mathbb{E}(e^{-\theta I/M}) = p_s(\theta/M)\,.$$

$$\text{DH-SS:} \quad \log p_s(\theta) \propto \theta^\delta \implies \frac{\log p_s(\theta/M)}{\log p_s(\theta)} = M^{-\delta}\,.$$

With FH-SS, the *density of interferers is reduced by a factor M*:

$$\text{FH-SS:} \quad \log p_s(\theta) \propto \lambda \implies \frac{\log p_s^{\mathrm{FH}}(\theta)}{\log p_s(\theta)} = M^{-1}\,.$$

Since $\delta < 1$, the benefit of FH-SS is larger; the difference is more drastic for small δ, *i.e.*, for large α.

Opportunistic ALOHA

Exploiting CSI

If the transmitter knows the channel, it can transmit whenever the channel state is good instead of transmitting "blindly" with probability p as in regular ALOHA.
For a threshold ν, let each node x transmit when $h_x > \nu$. For all other nodes, this is the same as ALOHA with $p = \mathbb{P}(h > \nu)$. We have

$$p_s(\nu) = \mathbb{P}(S > I\theta \mid h > \nu)\,.$$

For Rayleigh fading, we needed $\mathbb{E}(h^{-\delta}) = \Gamma(1-\delta)$. With conditioning,

$$\mathbb{E}(h^{-\delta} \mid h > \nu) = \Gamma(1-\delta, \nu)\,,$$

where $\Gamma(\cdot,\cdot)$ is the upper incomplete gamma function

$$\Gamma(a,z) = \int_z^\infty t^{a-1} \exp(-t) \mathrm{d}t\,.$$

Opportunistic ALOHA in Rayleigh fading

Upper bound [WAJ06]: $p_s < \exp\left(-p\lambda\pi\Gamma(1+\delta)\Gamma(1-\delta,\nu)\theta^\delta R^2\right)$,

where $p = \mathbb{P}(h > \nu) = \exp(-\nu)$.

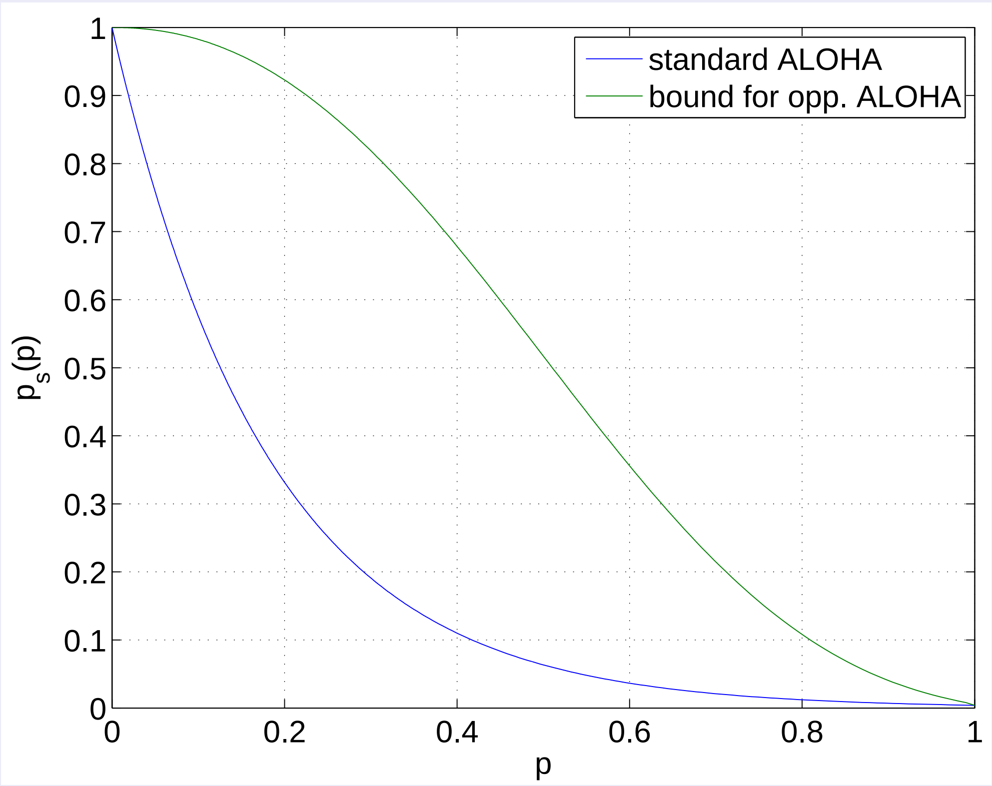

Comparison of $p_s(\theta)$ for $\lambda = 1/2$, $R = 1$, $\theta = 5$, and $\alpha = 4$ $(\delta = 1/2)$.

Power Control

Channel inversion without fading

Assume each transmitter talks to its nearest neighbor at distance R. Since R is Rayleigh with mean $1/(2\sqrt{\lambda})$, the transmitter power is Weibull distributed:

$$\mathbb{P}(P \leq x) = 1 - \exp(-\lambda\pi x^{\delta})$$

Power control at the transmitters acts like fading. Since $\mathbb{E}(P^{\delta}) = 1/(\pi\lambda)$,

$$\mathcal{L}_I(s) = \exp(-p\Gamma(1-\delta)s^{\delta})\,.$$

This does not depend on the density of the network.
Since the "fading" is not Rayleigh, the success probability cannot be derived from the Laplace transform.

Channel inversion with fading

Let each node have its destination at distance 1. If the fading is fully compensated, the received signal power $S \equiv 1$. The interference is

$$I = \sum_{x \in \Phi} \frac{h_{xo}}{h_{xz}} \|x\|^{-\alpha} ,$$

where h_{xo} is the fading from node x to the receiver at the origin, and h_{xz} is the fading from node x to its own receiver. The success probability

$p_s = \mathbb{P}(S > I\theta) = \mathbb{P}(I < \theta^{-1})$ cannot be calculated in closed-form.

If all fading is Rayleigh, an immediate problem is that $\mathbb{E}(h^{-1}) = \infty$, *i.e.*, finite power is not sufficient for full channel inversion.

Channel inversion with fading

For Rayleigh fading, $H = h_{xo}/h_{xz}$ is distributed as

$$F_H(x) = \mathbb{P}(H \le x) = \frac{x}{x+1} .$$

The relevant metric for the success probability is the δ-th moment H^δ. It turns out that full channel inversion *decreases* the success probability.

However, fractional channel inversion helps. Let the transmit power at each transmitter be $P_x = h_x^{-s}$ for $0 \le s \le 1$. $s = 0$ means no power control, while $s = 1$ means full channel inversion.
It is shown in [JWA08b] that $s = 1/2$ is optimal. This also solves the problem of infinite mean transmit power, since in this case $\mathbb{E}(P) = \sqrt{\pi}$.

This value of s minimizes $\mathbb{E}(X^{-s})\mathbb{E}(X^{s-1})$ for all non-negative RVs.

Bandwidth Partitioning

Optimum number of subbands [JWA08a]

Given a total bandwidth B, what number of subbands N should be chosen to maximize the number of concurrent links in the network?
Given a rate of transmission R_T, the corresponding SIR threshold is:

$$R_T = \frac{B}{N}\log(1+\theta(N)) \quad \Longrightarrow \quad \theta(N) = \exp(NR_T/B) - 1\,.$$

Let $b = NR_T/B$ be the spectral efficiency. Using the transmission capacity framework, we can find the b_{opt} that maximizes the total density of concurrent transmissions given an outage constraint:

$$\lambda(b,\epsilon) \propto \frac{b}{(e^b-1)^\delta} \quad \Longrightarrow \quad b_{\mathrm{opt}} = \mathcal{W}\left(-\frac{1}{\delta}e^{-1/\delta}\right) + \frac{1}{\delta} \quad (\mathrm{nats/s/Hz})\,,$$

which is exactly the spectral efficiency that maximized the probabilistic throughput on slide 7! Hence $N_{\mathrm{opt}} \approx (\alpha-2)B/R_T$.

Section Outline

1. Throughput
2. Other Applications
3. A Geometric Interpretation of Fading
   - Introduction
   - Local Connectivity
   - Broadcast Transport Capacity
4. Summary

Path Loss Processes with Fading [Hae08]

Basic idea

Have "distances" indicate path loss (with fading).

Definition (Path loss process with fading (PLPF))

$\{y_i\}$: PPP of intensity 1 in $\mathbb{R}^d$, ordered according to $\|y_i - o\|$.
Define a one-dimensional PPP $\{r_i \triangleq \|y_i - o\|\}$ on $\mathbb{R}^+$.
Let α be the path loss exponent of the network and let

$$\Phi = \{x_i \triangleq r_i^\alpha\}$$

be the path loss process (before fading) (PLP).
Let $\{f, f_1, f_2, \ldots\}$ be iid with distribution F and $\mathbb{E}f = 1$, and let

$$\Xi = \{\xi_i \triangleq x_i / f_i\}$$

be the path loss process with fading (PLPF).

Path Loss Process with Fading

Illustration: From PPP to PLPF

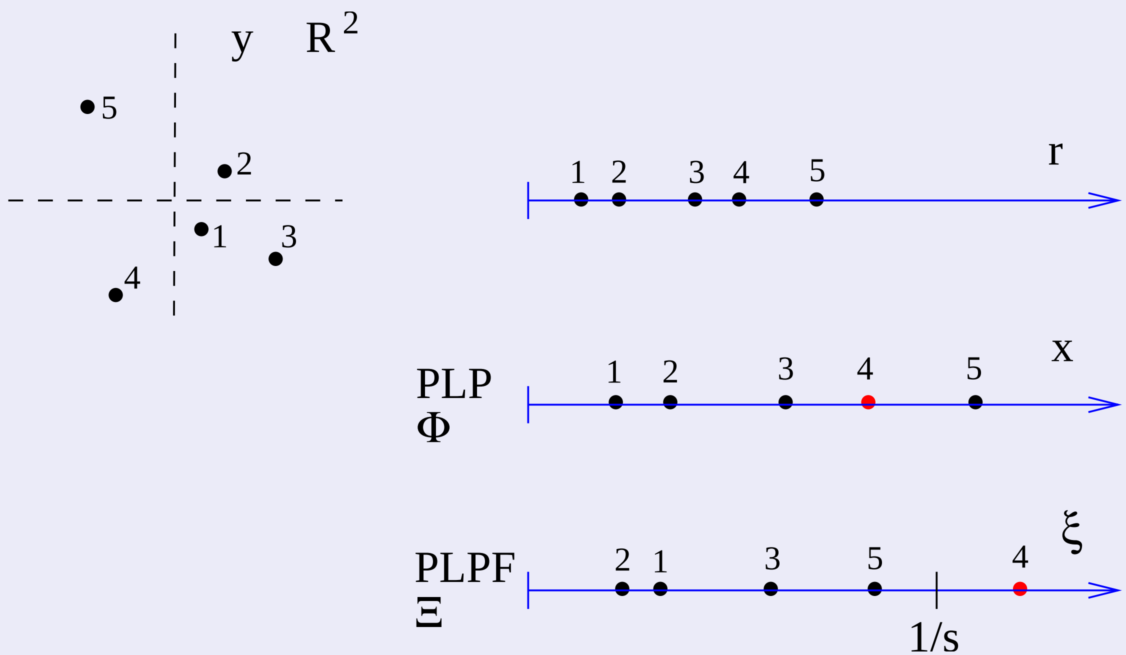

Notation

$\{y_i\}$: PPP with intensity 1

$\{r_i \triangleq \|y_i - o\|\}$

$\Phi = \{x_i \triangleq r_i^{\alpha}\}$

$\Xi = \{\xi_i \triangleq x_i/f_i\}$

Connected nodes

Node i is connected (to o) if $\xi_i < 1/s$.

Processes of connected nodes are $\hat{\Phi} = \{x_i : \xi_i < 1/s\}$ and $\hat{\Xi} = \Xi \cap [0, 1/s)$.

In the example PLPF, Node 4 is disconnected.

Basic Properties of the PLP, PLPF

- Φ, Ξ, and $\hat{\Xi}$ are Poisson (generally inhomogeneous).
- $\mathbb{E}\Phi([0,x)) = c_d x^\delta$, $\lambda(x) = c_d \delta x^{\delta-1}$, where $\delta = d/\alpha$.
  For $\delta = 1$, Φ is uniform (on $\mathbb{R}^+$). Density increasing with x if $d > \alpha$.
- For $\delta = 1$ and Rayleigh fading,

  $$F_{\xi_i}(x) = \frac{(c_d x)^i}{(c_d x + 1)^i}\,.$$

  Note that $\mathbb{E}\xi_i$ does not exist.
- For $\delta = 1$ and arbitrary fading (with unit mean),

  $$\Xi(B) \stackrel{d}{=} \Phi(B) \qquad \forall\, B \subset \mathbb{R}\,,$$

  i.e., fading is distribution-preserving.

Note that ξ_i are *neither ordered nor independent*.

Local Connectivity

Impact of fading

The process of connected nodes is

$$\hat{\Xi} = \{\xi_i : \xi_i < 1/s\} = \Xi \cap [0, 1/s).$$

The effect of fading is location-dependent (but node-independent) thinning:

$$\hat{\lambda}(x) = \lambda(x)(1 - F(sx))$$

since
$\mathbb{P}[x/f < 1/s] = \mathbb{P}[f > sx] = 1 - F(sx)$.

Without fading, $F(x) = u(x-1)$, and we obtain the "disk model".

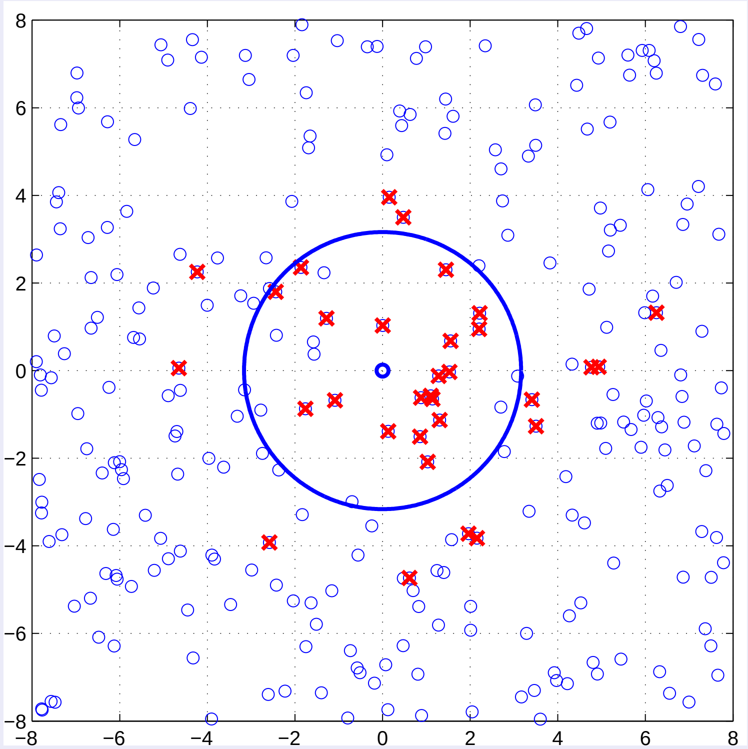

$s = 0.1$. The disk has radius $1/\sqrt{s}$.
Red nodes are connected under Rayleigh fading.

Fading as a stochastic mapping

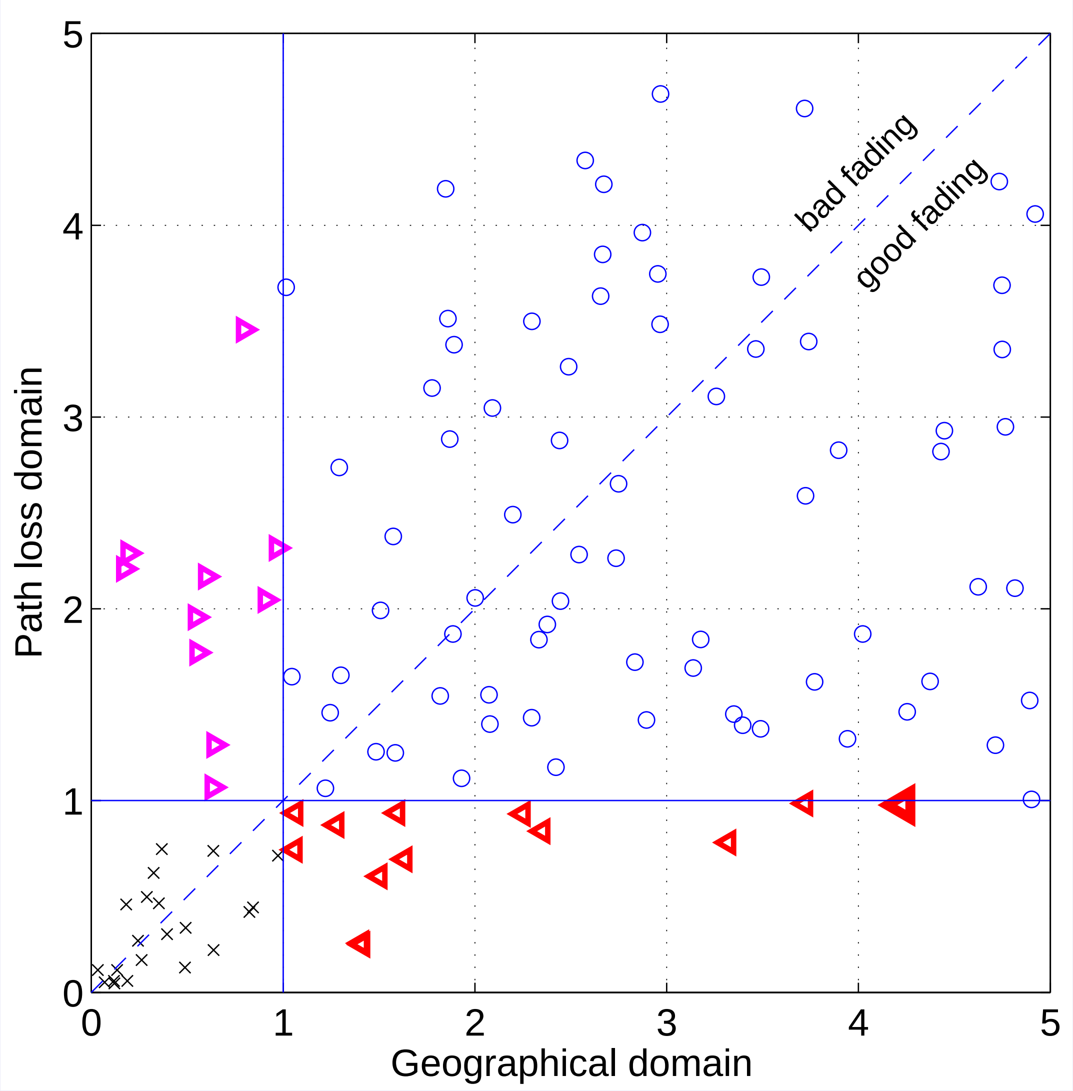

Connectivity for $s = 1$.

Red nodes benefit from fading, purple ones suffer from it.

A More General Fading Model

Definition (Nakagami-m fading)

$$p_f(x) = \frac{m^m}{\Gamma(m)} x^{m-1} \exp(-mx), \qquad m \geqslant 0$$

Remarks

- $\mathbb{E}f = 1$, and $\operatorname{var} f = 1/m$. $\mathbb{P}[f < x] = 1 - \Gamma(m, mx)/\Gamma(m)$.
- For $m = 1$, this is exponential (or Rayleigh in the amplitude). The sum of m iid Rayleigh-fading signals is a Nakagami distributed signal.
- Describes the amplitude of received signal after m-branch maximum ratio diversity combining (MIMO).
- For $m \to \infty$, $p_f(x) \to \delta(x-1)$ and thus $f = 1$ (no fading), so a non-negligible LOS component in the received signal can be modeled, and the case of no fading is a special case.

Connectivity under Nakagami fading

The number $\hat{N} = \hat{\Phi}(\mathbb{R}^+)$ of connected nodes is Poisson with mean

$$\mathbb{E}\hat{N}_m = \frac{c_d}{(ms)^\delta}\frac{\Gamma(\delta+m)}{\Gamma(m)}$$

The *connectivity fading gain* is

$$\frac{\mathbb{E}\hat{N}_m}{\mathbb{E}\hat{N}_\infty} = \frac{1}{m^\delta}\frac{\Gamma(\delta+m)}{\Gamma(m)} = \mathbb{E}(f^\delta)\,.$$

The *connectivity fading gain* equals the δ-th moment of the fading distribution.

Proof: Calculate $\int_0^\infty \hat{\lambda}(x)\mathrm{d}x$.

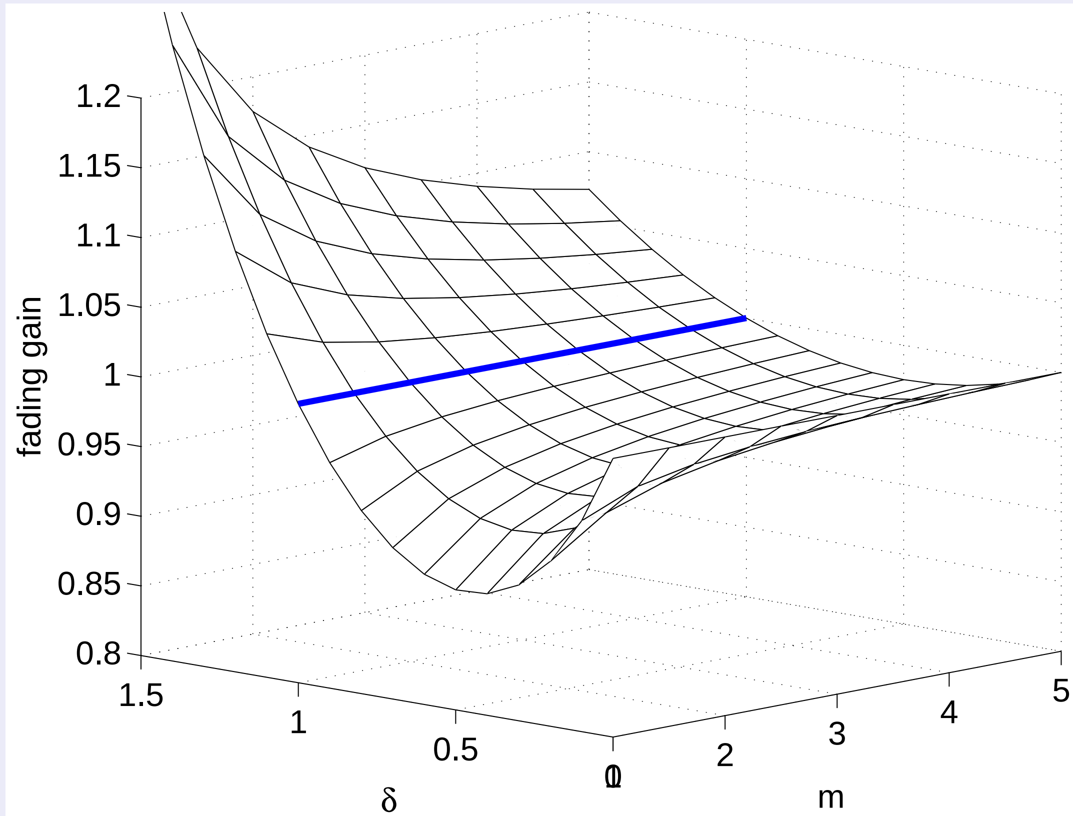

$\delta = d/\alpha$.

So only if $d > \alpha$, fading helps.

This is not common.

Broadcast Transport Capacity

Definition (Broadcast transport sum-distance)

This is the sum of the (geographical) distances to all the connected nodes:

$$D \triangleq \mathbb{E}\Big(\sum_{x\in\hat{\Phi}} x^{1/\alpha}\Big)$$

For Nakagami-m fading:

$$D_m = c_d \frac{\delta}{\Delta} \frac{1}{(ms)^{\Delta}} \frac{\Gamma(m+\Delta)}{\Gamma(m)},$$

where $\delta \triangleq d/\alpha$, $\Delta \triangleq (d+1)/\alpha$.

Proof.

Using Campbell's theorem:

$$D_m = \int_0^\infty x^{1/\alpha} \hat{\lambda}_m(x) \mathrm{d}x$$

□

Broadcast fading gain

The (broadcast) *fading gain* D_m/D_∞ is

$$\frac{D_m}{D_\infty} = \frac{1}{m^\Delta} \frac{\Gamma(m+\Delta)}{\Gamma(m)} = \mathbb{E}(f^\Delta) \,.$$

$\delta \triangleq d/\alpha$, $\Delta \triangleq (d+1)/\alpha$. The gain is the Δ-th moment of f.

Definition (Broadcast transport capacity)

Include the (maximum) rate of transmission $R = \log_2(1+s)$ and define

$$C \triangleq \max_{R>0}\{R \cdot D(R)\} = \max_{s>0}\{\log_2(1+s)D(s)\}.$$

as the broadcast transport capacity.

For Nakagami-m fading:

- For $\Delta \in (0,1]$, the broadcast transport capacity is achieved for

$$R_{\text{opt}} = \frac{\mathcal{W}\left(-\frac{e^{-1/\Delta}}{\Delta}\right) + \Delta^{-1}}{\log 2}, \quad \Delta \in (0,1],$$

  where $\mathcal{W}$ is the (principal branch) of the Lambert W function.
  The corresponding C_m is smallest for $\Delta = 1/(2\log 2) \approx 0.72$.
  For $d = 2$, this means $\alpha_{\min} = 4.16$.

Broadcast transport capacity for Nakagami-m fading, cont.

- For $\Delta > 1$ ($\alpha < 3$ for $d = 2$), the broadcast transport capacity increases without bounds as $R \to 0$ (or $s \to 0$), independent of the transmit power.

Intuition

$D(s) \propto s^{-\Delta}$, and for small s, $R \approx s$. So $C_m \approx s^{1-\Delta}$. Reaching more nodes more than offsets the reduced rate.
Impact on broadcast protocols? Is this the right metric?

With (fancy) superposition coding:

If all nodes could decode at their SNR, we would have (without fading)

$$\tilde{C} = \mathbb{E}\left[\sum_{x \in \Phi} x^{1/\alpha} \log_2(1 + x^{-1})\right] .$$

In this case, the singularity of the path loss model would become significant.

Section Outline

1. Throughput
2. Other Applications
3. A Geometric Interpretation of Fading
4. Summary

Lecture 2 Summary

Throughput analysis and design of Poisson networks

- In Lecture 1, we have derived a success probability result. For MAC design, a form of throughput is needed.
- Regular networks provide a larger throughput. This *regularity gain* is partially achieved by CSMA.
- The effects of power control and spread-spectrum can be analyzed. With power control, care is needed since the effect of all other receivers needs to be factored in. Opportunistic schemes are great—if CSI is available.
- Fading can be interpreted in a geometric fashion. The resulting point process represents path loss including fading. The Poisson property is preserved by this mapping.

References I

J. G. Andrews, S. Weber, and M. Haenggi.
Ad Hoc Networks: To Spread or not to Spread?
IEEE Communications Magazine, 45(12):84–91, December 2007.
Available at http://www.nd.edu/~mhaenggi/pubs/commag07.pdf.

M. Haenggi.
A Geometric Interpretation of Fading in Wireless Networks: Theory and Applications.
IEEE Trans. on Information Theory, 54(12):5500–5510, December 2008.
Available at http://www.nd.edu/~mhaenggi/pubs/tit08.pdf.

M. Haenggi.
Outage, Local Throughput, and Capacity of Random Wireless Networks.
IEEE Transactions on Wireless Communications, 8(8):4350–4359, August 2009.
Available at http://www.nd.edu/~mhaenggi/pubs/twc09.pdf.

N Jindal, S Weber, and J.G Andrews.
Bandwidth Partitioning in Decentralized Wireless Networks.
IEEE Transactions on Wireless Communications, 7(12):5408–5419, December 2008.

N Jindal, S Weber, and J.G Andrews.
Fractional power control for decentralized wireless networks.
Wireless Communications, IEEE Transactions on, 7(12):5482–5492, December 2008.

R. Mathar and J. Mattfeldt.
On the distribution of cumulated interference power in Rayleigh fading channels.
Wireless Networks, 1(1):31–36, February 1995.

References II

S. Weber, J. G. Andrews, and N. Jindal.
Throughput and Transmission Capacity of Ad Hoc Networks with Channel State Information.
In *2006 Allerton Conference*, Monticello, IL, September 2006.

S. Weber, J. G. Andrews, and N. Jindal.
An Overview of the Transmission Capacity of Wireless Networks.
IEEE Transactions on Communications, 2010.
Submitted. Available at `http://arxiv.org/abs/0809.0016`.

Steven Weber, Xiangying Yang, Jeffrey G. Andrews, and Gustavo de Veciana.
Transmission Capacity of Wireless Ad Hoc Networks with Outage Constraints.
IEEE Transactions on Information Theory, 51(12):4091–4102, December 2005.

Analysis and Design of Wireless Networks

Lecture 3: Random Graphs and Percolation Theory

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop analysis of Poisson networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 3 Overview

1. Introduction
2. Bond Percolation on Lattice
3. Percolation on the Disk Graph
4. Secrecy Graphs
5. The SINR Graph
6. Summary

Section Outline

1 Introduction
- Motivation
- Gilbert's disk graph

2 Bond Percolation on Lattice

3 Percolation on the Disk Graph

4 Secrecy Graphs

5 The SINR Graph

6 Summary

Introduction

Random graphs and geometric graphs

- Graph models have a long tradition in networking. In wired networks, there is a direct mapping from wires between nodes to edges in a graph.
- In mathematics, random graphs, where the existence of edges is subject to randomness, are well-studied objects [Bol01]. The basic model is: $G = (V, E)$ where $V = [n] = \{1, \ldots, n\}$ and $E \subset [n]^2$ with $M = |E|$ edges. If $N = \binom{n}{2}$, then there are $\binom{N}{M}$ such graphs, and each one occurs with probability $\binom{N}{M}^{-1}$.
  In sociology, such models are useful when interactions are possible between all individuals.
- If there is a notion of distance between vertices (points, nodes) and vertices that are further away are less likely to interact, the graph turns into a random geometric graph.

Random graphs and geometric graphs

- Although there are similarities between abstract random graphs and geometric random graphs, the latter are more challenging in many cases since there is a triangular relationship: If x is connected to y, and y to z, then it is quite likely that z is also connected to x.
- In the wireless setting, random geometric graphs are often good models, since distances are critical and edges may be subject to uncertainty.
- In a random geometric graph, each node has a location, usually in $\mathbb{R}^d$, and the probability of an edge $x \to y$ depends on $\|x - y\|$.
- The most basic model is Gilbert's disk graph.

Gilbert's Disk Graph [Gil61]

Definition (Gilbert's disk graph)

Take a stationary PPP of intensity λ as the vertices of a random geometric graph and connect two vertices by an edge if they are within distance r of each other. The resulting graph $G_{\lambda,r}$ is called a disk graph.

Example: $\lambda = 1$, $r = 1$

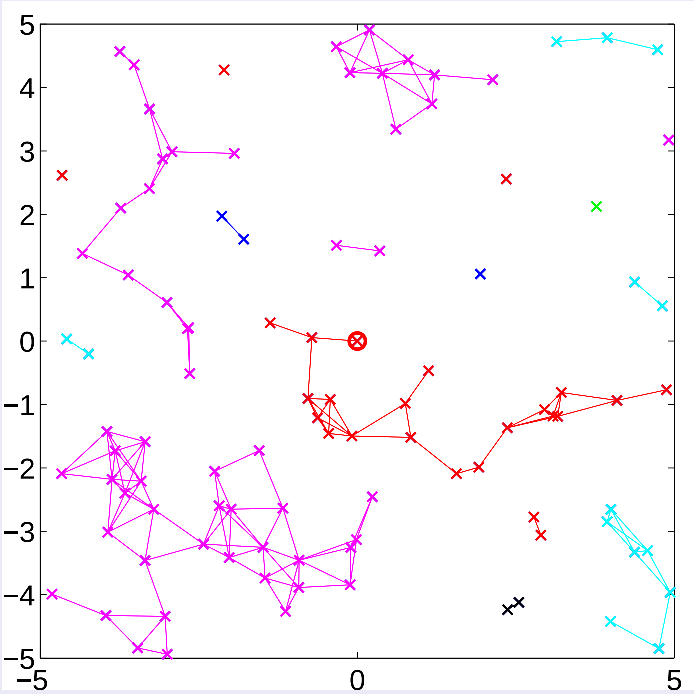

Interpretation

In the absence of interference, the condition SNR $> \theta$ defines a maximum *communication radius r*

$$r = \left(\frac{P}{\theta W}\right)^{1/\alpha}.$$

Connectivity

Let Φ be a PPP of intensity 1 on $[0, \sqrt{n}]^2$ so that $\mathbb{E}N = n$.

What communication radius r_c guarantees that $G_r(n)$ is connected whp as $n \to \infty$?
Formally, we want

$$\lim_{n\to\infty} \mathbb{P}[G_{r_c}(n) \text{ connected}] = 1\,.$$

Examples: $n = 100$

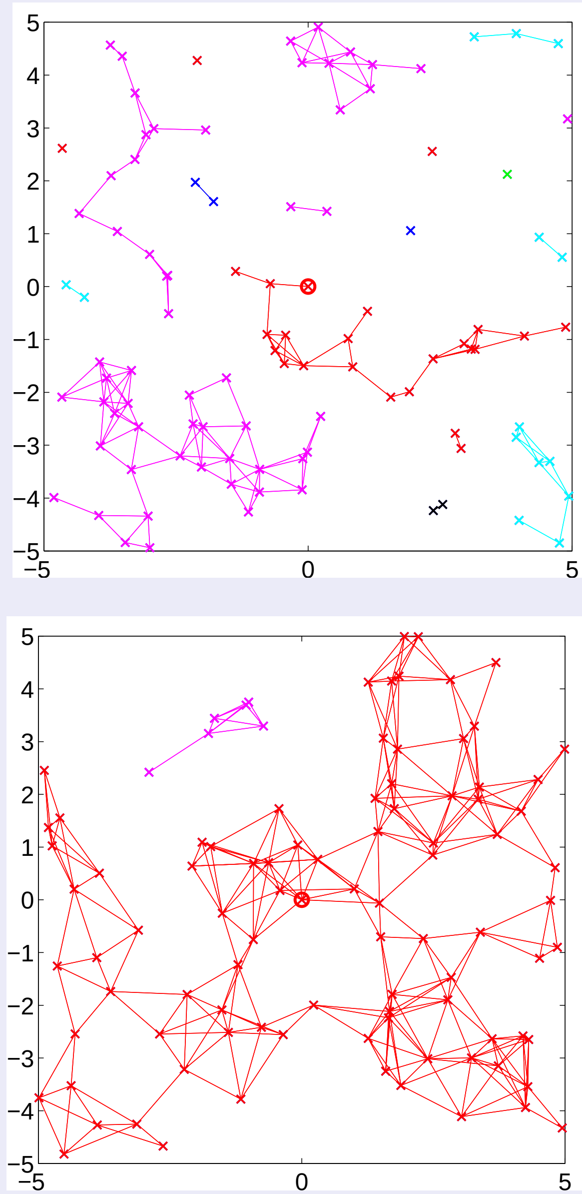

$r = 1$

$r = 3/2$

Minimum transmission radius for connectivity

A necessary condition for connectivity is that no node is isolated.
The expected number of isolated nodes

$$\mathbb{E}N_{\text{isol}} = n\,\mathbb{P}(\text{typical node is isolated}) = n\exp(-\pi r^2)$$

needs to go to 0 as n grows. So we need $\pi r^2 \gtrapprox \log n$ for connectivity.

[Pen97] showed that indeed the isolated nodes determine the connectivity such that

$$\pi r^2 = \log n + \omega(1)\,,$$

where $\omega(1)$ is any function f for which $f(n) \to \infty$ (arbitrarily slowly) as $n \to \infty$.

Setting $\pi r^2 = \log n + c$ would result in a $N_{\text{isol}} \sim \text{Po}(\exp(-c))$.

So r needs to grow with $\sqrt{\log n}$ to keep the network connected!

But a constant power is enough to keep an infinite number of nodes connected. Percolation theory gives bounds on this critical threshold.

Section Outline

1 Introduction

2 Bond Percolation on Lattice
- Model
- Critical probability

3 Percolation on the Disk Graph

4 Secrecy Graphs

5 The SINR Graph

6 Summary

Bond Percolation on Lattice

Lattice with open and closed bonds

Take the set of vertices to be the points in $\mathbb{Z}^2$, and put edges among all nearest neighbors. Make edges *open* (passable) with probability p or *closed* (blocked) with probability $1-p$.

Example (Bond percolation)

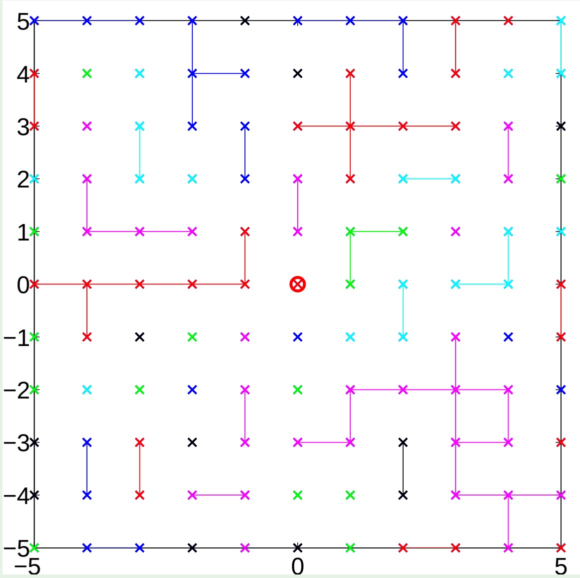

$p = 0.3$

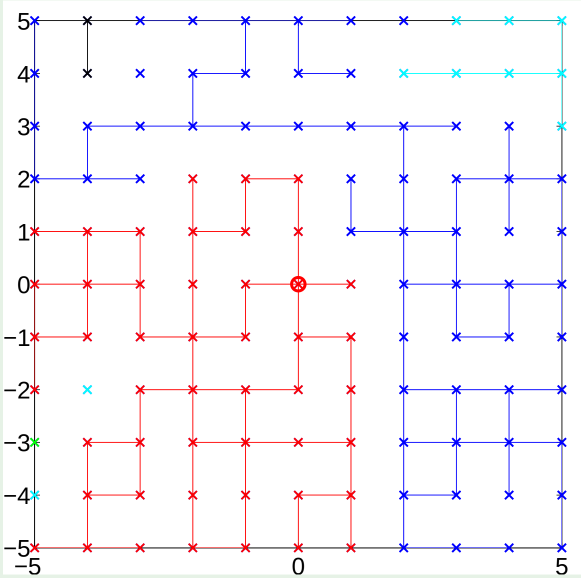

$p = 0.6$

Critical probability

Let $u \leftrightarrow v$ stand for the existence of an (open) path between $u, v \in \mathbb{Z}^2$.
The *open cluster* $C(v)$ is the set of all vertices that are connected to v by an open path:

$$C(v) = \{u \in \mathbb{Z}^2 \colon u \leftrightarrow v\}$$

The central quantity is the percolation probability

$$\psi(p) = \mathbb{P}(o \leftrightarrow \infty) = \mathbb{P}(|C(o)| = \infty) .$$

The lattice model exhibits a phase transition, *i.e.*, there exists a critical value p_c such that $\psi = 0$ for $p < p_c$ and $\psi > 0$ for $p > p_c$.

The critical probability is defined as

$$p_c \triangleq \sup\{p \colon \psi(p) = 0\} .$$

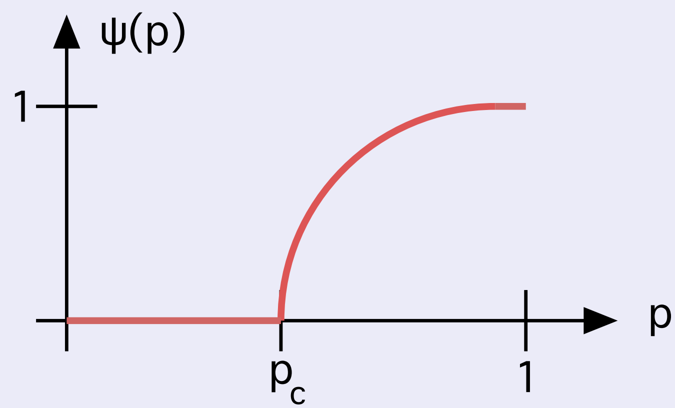

By Kolmogorov's 0-1 law, there exists an infinite component w.p. 1 as soon as $p > p_c$.

Two larger examples

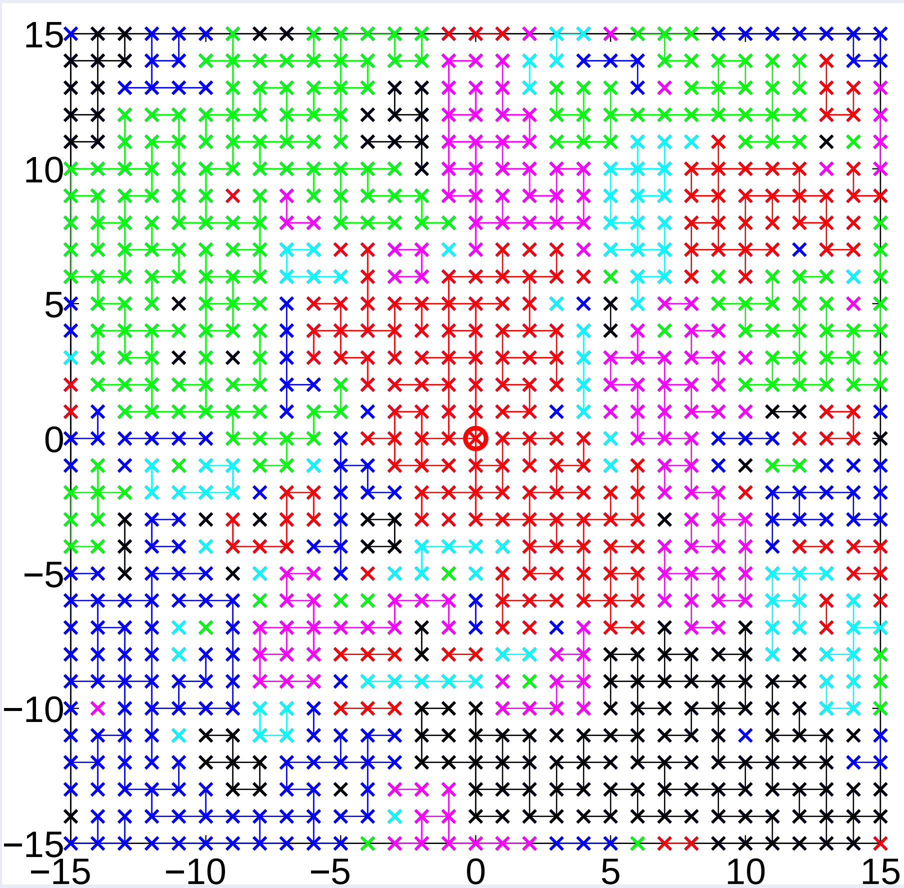

$p = 0.45$

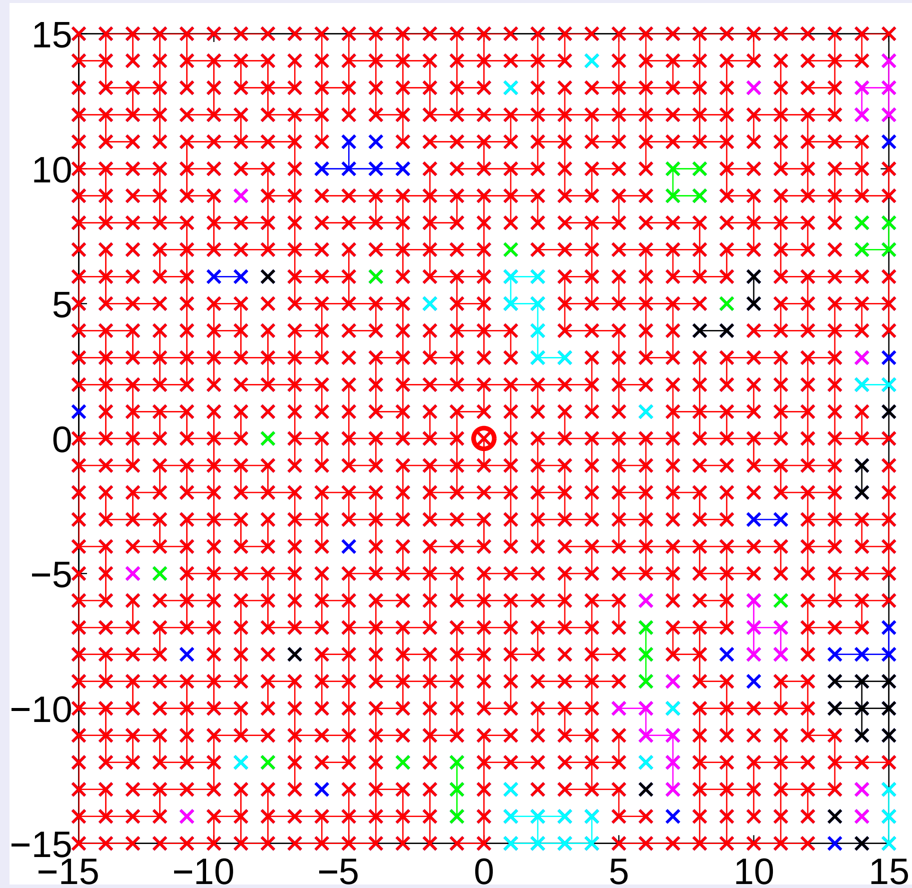

$p = 0.55$

This indicates that the critical probability is near 1/2.

Lower bounding the critical probability

If there is an infinite cluster, then for any n, there exists a (self-avoiding) path of length n:

$$\psi(p) \leq \mathbb{P}(\exists \text{ a path of length } n \text{ starting at } o) \quad \forall n \in \mathbb{N}\,.$$

If all edges were open, the number of $\kappa(n)$ of paths of length n is smaller than $4 \cdot 3^{n-1}$.

Each path exists with probability p^n, so by the union bound

$$\mathbb{P}(\exists \text{ a path of length } n \text{ starting at } o) \leq 4 \cdot 3^{n-1} p^n\,.$$

If $p < 1/3$, this goes to 0 as $n \to \infty$. So $p_c \geq 1/3$.

Upper bounding the critical probability I

- Take the dual lattice whose vertices are at $(\mathbb{Z}+1/2)^2$. Place an edge if it does not intersect an open edge in the original lattice.
- If a component is finite in the original lattice, it must be surrounded by a circuit in the dual lattice.
- If we can show that for some p, there is a positive probability of having no such circuit, we obtain an upper bound on p_c.

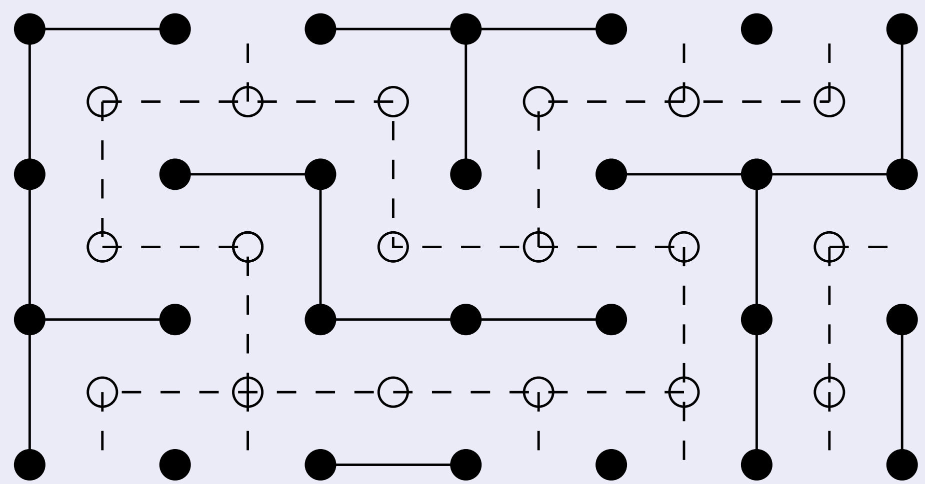

Bond percolation model and dual lattice.

Upper bounding the critical probability II

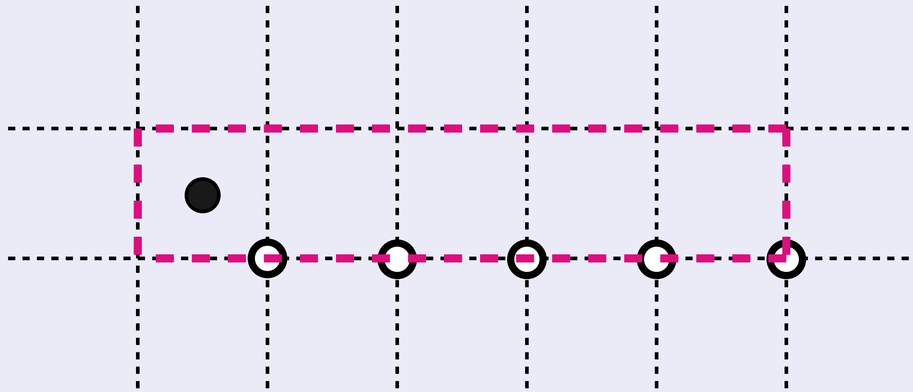

A circuit in the dual lattice of length $2n = 12$ around the origin has to go through one of the $n - 1 = 5$ dual vertices indicated.

The number $\sigma(n)$ of possible circuits of length $2n$ that surround the origin is bounded as

$$\sigma(n) \leqslant (n-1) \cdot 3^{2(n-1)},$$

since for the first and last edge, the direction is given. So:

$$\mathbb{P}(\text{closed circuit}) \leqslant \sum_{n=2}^{\infty} (1-p)^{2n} \sigma(n) = \frac{9(1-p)^4}{[1-9(1-p)^2]^2}.$$

When $p > 1 - 1/(2\sqrt{3}) \approx 0.71$, this is less than one, which means that there is a positive probability that the origin belongs to an infinite cluster. So $p_c \leqslant 1 - 1/(2\sqrt{3})$.

Critical probability

Harry Kesten showed in 1980 that the critical probability for bond percolation on the square lattice is $p_c = 1/2$ — 25 years after this was conjectured.

Section Outline

1 Introduction

2 Bond Percolation on Lattice

3 Percolation on the Disk Graph
- Definition
- Galton-Watson branching processes
- Bounding the critical radius

4 Secrecy Graphs

5 The SINR Graph

6 Summary

Percolation on the Disk Graph

The critical radius

Take a PPP of intensity λ and add a node at the origin o. Let

$$\psi(r) = \mathbb{P}(o \leftrightarrow \infty) = \mathbb{P}(|C(o)| = \infty)\,.$$

and define $$r_c \triangleq \sup\{r \colon \psi(r) = 0\}\,.$$

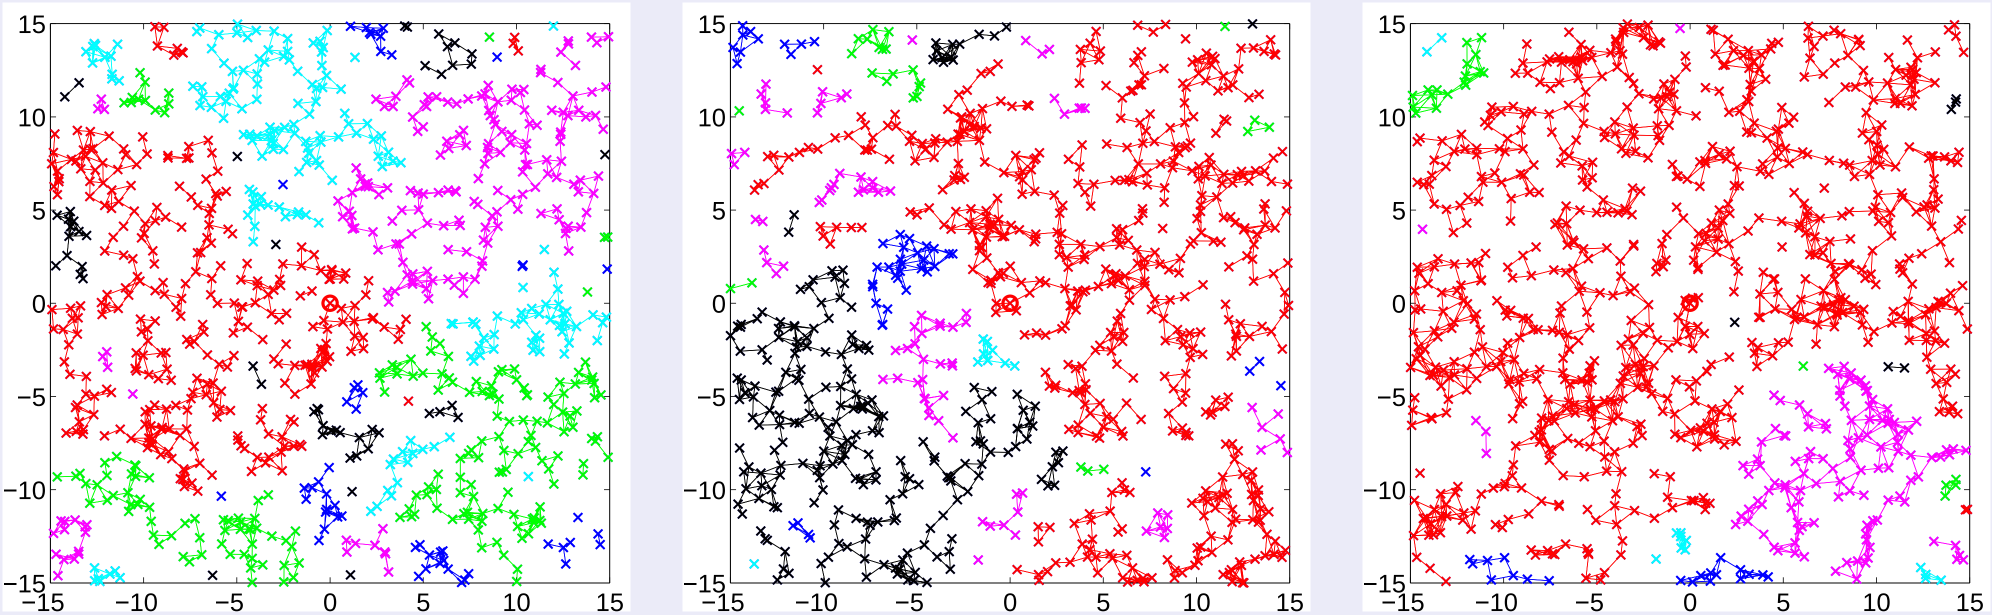

$r = 1.1$ $\quad$ $r = 1.2$ $\quad$ $r = 1.3$

This indicates that the critical radius is near 1.2 when $\lambda = 1$.

Galton-Watson branching processes

Let $Z_0 = 1$ and recursively define the stochastic process

$$Z_{n+1} = \sum_{i=1}^{Z_n} X_{n+1,i}\,, \qquad \text{where } X_{n,i} \text{ are iid for all } n, i \in \mathbb{N}.$$

Z_n can be viewed as the number of members in the n-th generation. Each member of this generation gives birth to a random number of children, $X_{n,i}$, which are the members of the $(n+1)$-th generation.

The process $\{Z_n\}$ is a Galton-Watson branching process. If $X \in \mathbb{N}_0$, the process can be represented by a random tree.

An important result is:
The probability of eventual extinction is 1 if $\mathbb{E}(X) \leq 1$ (unless $\mathbb{P}(X = 1) = 1$), whereas for $\mathbb{E}(X) > 1$, the process may live forever.

Lower bounding the critical radius

Populate the set $C(o)$ of nodes connected to the origin step by step. Start with $C = \{o\}$. Then at each step add all nodes that share an edge with an element of C.

Each node has on average $\lambda\pi r^2$ edges, so the process can be compared with a Galton-Watson branching process. If $\lambda\pi r^2 < 1$, the (independent) branching process dies out w.p. 1, so our process does too.

So $r_c > 1/\sqrt{\lambda\pi}$. For $\lambda = 1$, $r_c > 0.5642$.

Upper bounding the critical radius

For the upper bound, we use the result for bond percolation.

Divide the plane into squares of size $c = r/(2\sqrt{2})$.

Each square corresponds to a potential edge in the bond percolation model.

The bond is open if there is at least one point of the PPP in the square, which happens with probability $p = 1 - \exp(-\lambda c^2)$.

Upper bounding the critical radius

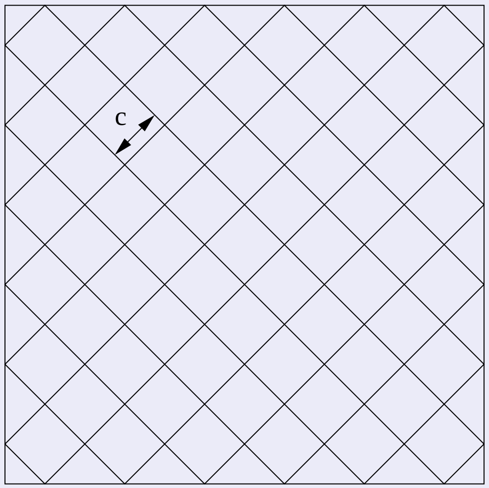

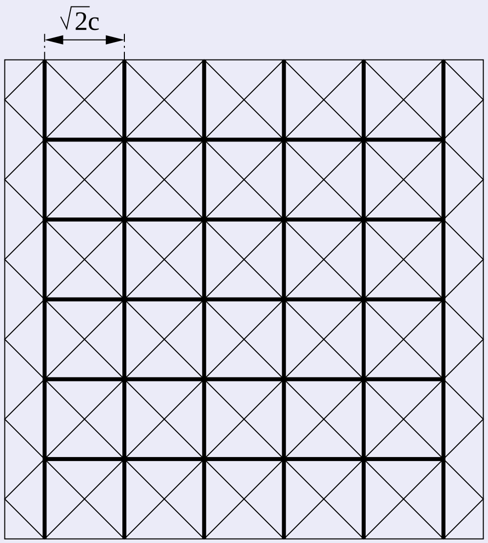

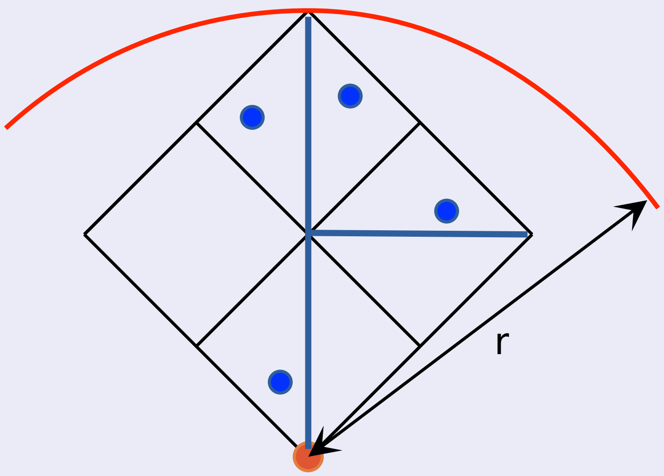

If $p = 1 - \exp(-\lambda c^2) > 1/2$, or $\lambda > \log 2/c^2$, the bond model percolates. If two edges are adjacent in the bond model, then two points of the PPP are located in squares that touch in at least a corner. Since the distance between them is at most $2\sqrt{2}c = r$, the two points are connected in $G_{\lambda,r}$. So we have $\lambda r_c^2 < 8 \log 2$, or

$$r_c < \sqrt{8 \log 2/\lambda}; \qquad \text{for } \lambda = 1: \quad r_c < 2.355\,.$$

The critical radius

We have shown (for $\lambda = 1$):

$$0.564 < r_c < 2.355\,.$$

The best known analytical bounds are $0.833 < r_c < 1.83$, so we're not too far off.
In [BBW05], the bounds

$$1.1979 < r_c < 1.1988$$

were established with 99.99% confidence (using MC integration of a complicated integral). This corresponds to a mean number of neighbors of $\lambda\pi r_c^2 \approx 4.51$ per node.

So, a positive fraction of nodes can be connected at constant power level.

At this level of connectivity, the fraction of isolated nodes is

$$\exp(-4.51) \approx 0.011\,.$$

Section Outline

1 Introduction

2 Bond Percolation on Lattice

3 Percolation on the Disk Graph

4 Secrecy Graphs
- Definition
- Power- and secrecy-limited regimes
- Percolation in the Poisson model

5 The SINR Graph

6 Summary

Secrecy Graphs [Hae08]

The Poisson-Poisson secrecy graph

Let Φ be a PPP of users or "good guys" of intensity 1, and let Ψ be a PPP of eavesdroppers of intensity λ. The secrecy graph $\vec{G}_{\lambda,r} = (\Phi, \vec{E})$ includes all directed edges for which $\overrightarrow{xy}$ if $\|x - y\| < r$ and y is closer to x than any eavesdropper.

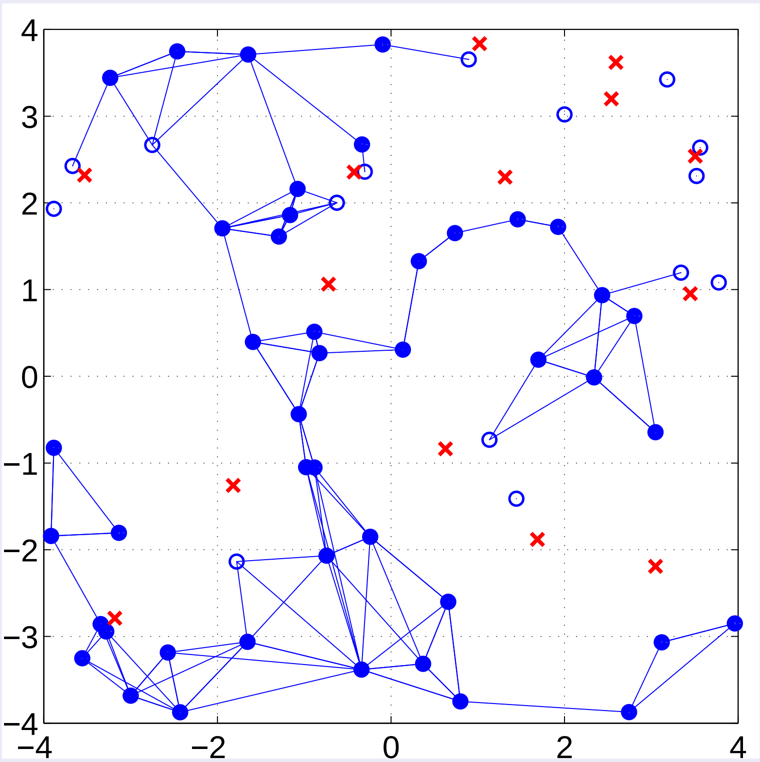

This graph contains only edges along which secure communication is possible.

good guys. **o** are receivers only.
eavesdroppers. $\lambda = 0.3$.
$r = \infty$ (no power constraint, only secrecy contraints)

Power- and secrecy-limited regimes

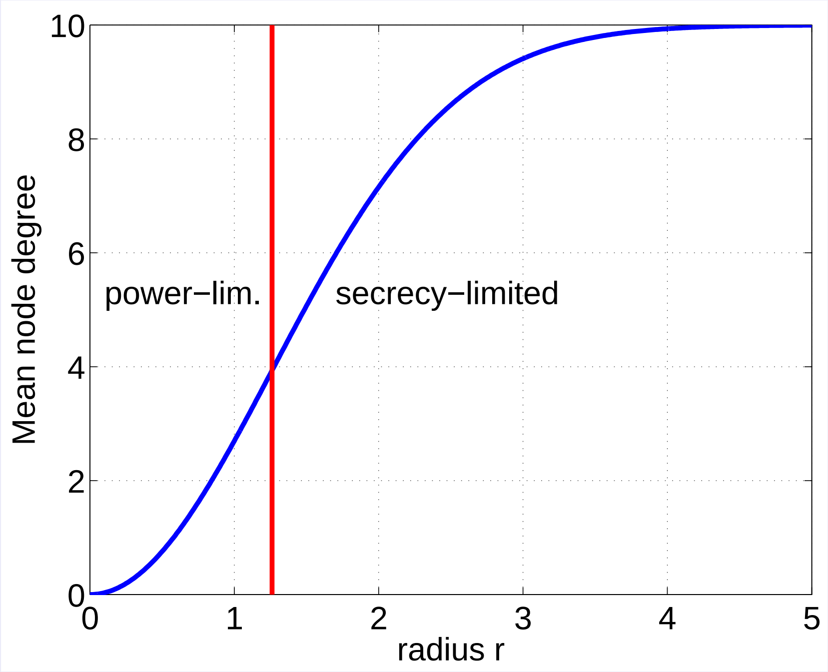

$$\mathbb{E}N(r) = \frac{1}{\lambda}(1 - \exp(-\lambda\pi r^2))\,.$$

Mean out-degree of $\vec{G}_{1/10,r}$.

The inflection point in $\mathbb{E}N(r)$ marks the boundary between the power-limited and the secrecy-limited regime.

In the power-limited regime, the degree distribution is close to Poisson.

At the inflection point, $r = r_T = (2\pi\lambda)^{-1/2} = \sqrt{5/\pi} \approx 1.26$.

Percolation in the Poisson model

Critical radius and density

- With $\psi(\lambda, r)$ being the probability that the component containing the origin (or any arbitrary fixed node) is infinite, the percolation threshold radius for $\hat{G}_r$ is [BBW05]

$$r_{\mathrm{G}} \triangleq \sup\{r : \psi(0, r) = 0\} \approx 1.19\,,$$

- For radii larger than r_{G}, we define

$$\lambda_c(r) \triangleq \inf\{\lambda : \psi(\lambda, r) = 0\}\,, \quad r > r_{\mathrm{G}}.$$

This is the smallest density of eavesdroppers that ensures that the network is partitioned into many small components.

Oriented Percolation

Fact (Out-percolation of $\vec{G}_{\lambda,r}$)

$\lambda_c(r)$ is monotonically increasing for $r > r_{\mathrm{G}}$, and we have

$$0 < \lim_{r\to\infty} \lambda_c(r) < \infty\,.$$

In other words, there exists a λ_∞ such that for $\lambda > \lambda_\infty$, $\vec{G}_{\lambda,r}$ does not out-percolate for any r.

This follows from the fact that for fixed r the mean degree $\mathbb{E}N(\lambda)$ is continuously decreasing to 0. For intensities smaller than λ_∞, we define

$$r_c(\lambda) \triangleq \sup\{r \,:\, \psi(\lambda,r) = 0\}\,, \quad \lambda \leqslant \lambda_\infty\,.$$

Fact (Percolation radius)

The percolation radius $r_c(\lambda)$ is monotonically increasing with λ and has a vertical asymptote at λ_∞.

- This follows from the monotonic decrease of the mean degree in λ. For $\mathbb{E}N^{\mathrm{out}} < 1$, for sure there is no percolation by the Galton-Watson result.
- Without a power constraint, the node out-degrees are geometric with mean $1/\lambda$:
  Start at a good guy with a ball of radius 0. Let the ball grow until it hits a node in $\Phi \cup \Psi$. The probability that it is a good guy (in Φ) is $1/(1+\lambda)$. If it is, let the ball grow further, until it hits the next point. Again the probability that this is a good guy is $1/(1+\lambda)$. So:

$$\mathbb{P}(\mathrm{out\,deg} = n) = (1-p)p^n\,, \quad p = 1/(1+\lambda)$$

- So $\lambda_\infty < 1$. This is not a tight bound, as expected.

Numerical investigation

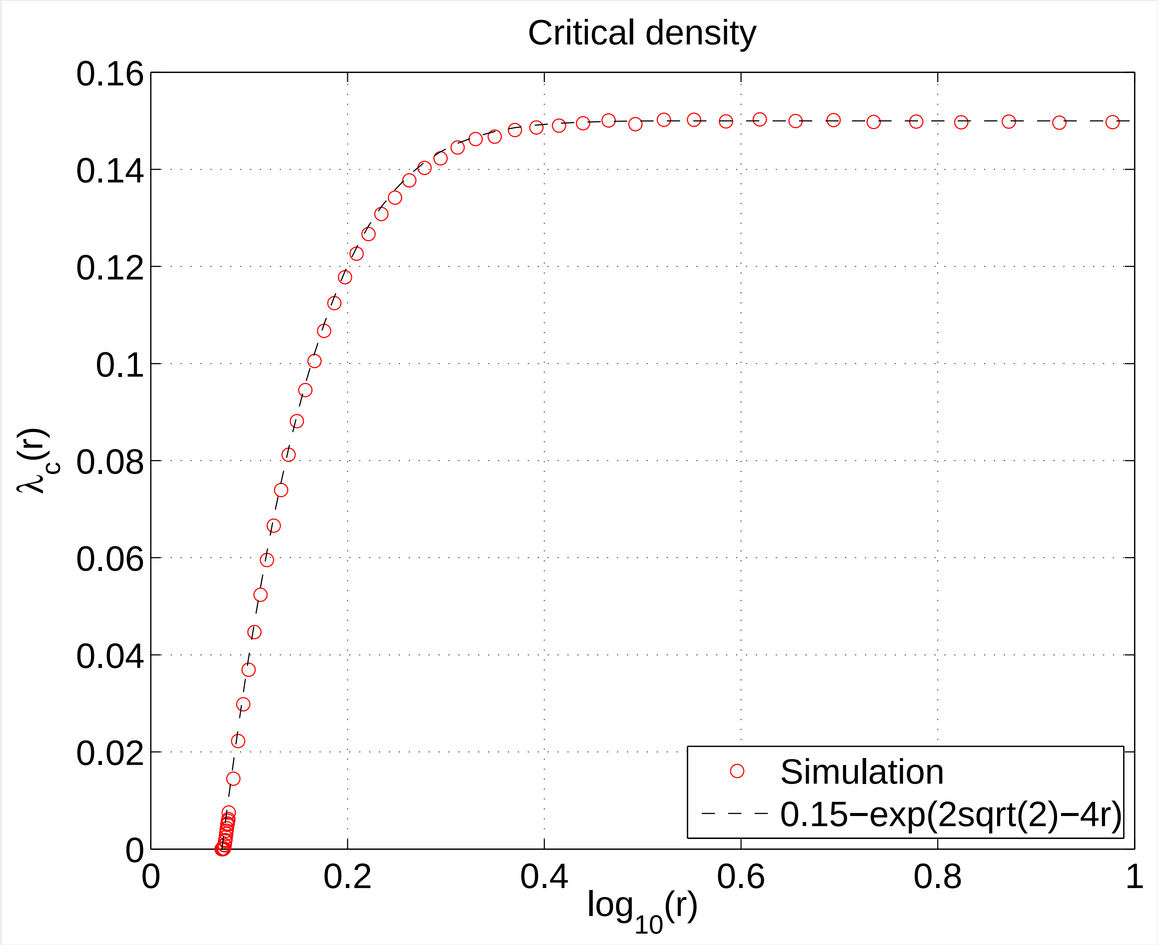

$$\lambda_c(r) \approx \lambda_\infty - \exp(a - br)\,, \quad r > r_{\mathrm{G}}$$

where: $\lambda_\infty \approx 0.1499 \quad a = 2\sqrt{2}\,, \quad b = 4$

Two larger examples ($r = \infty$)

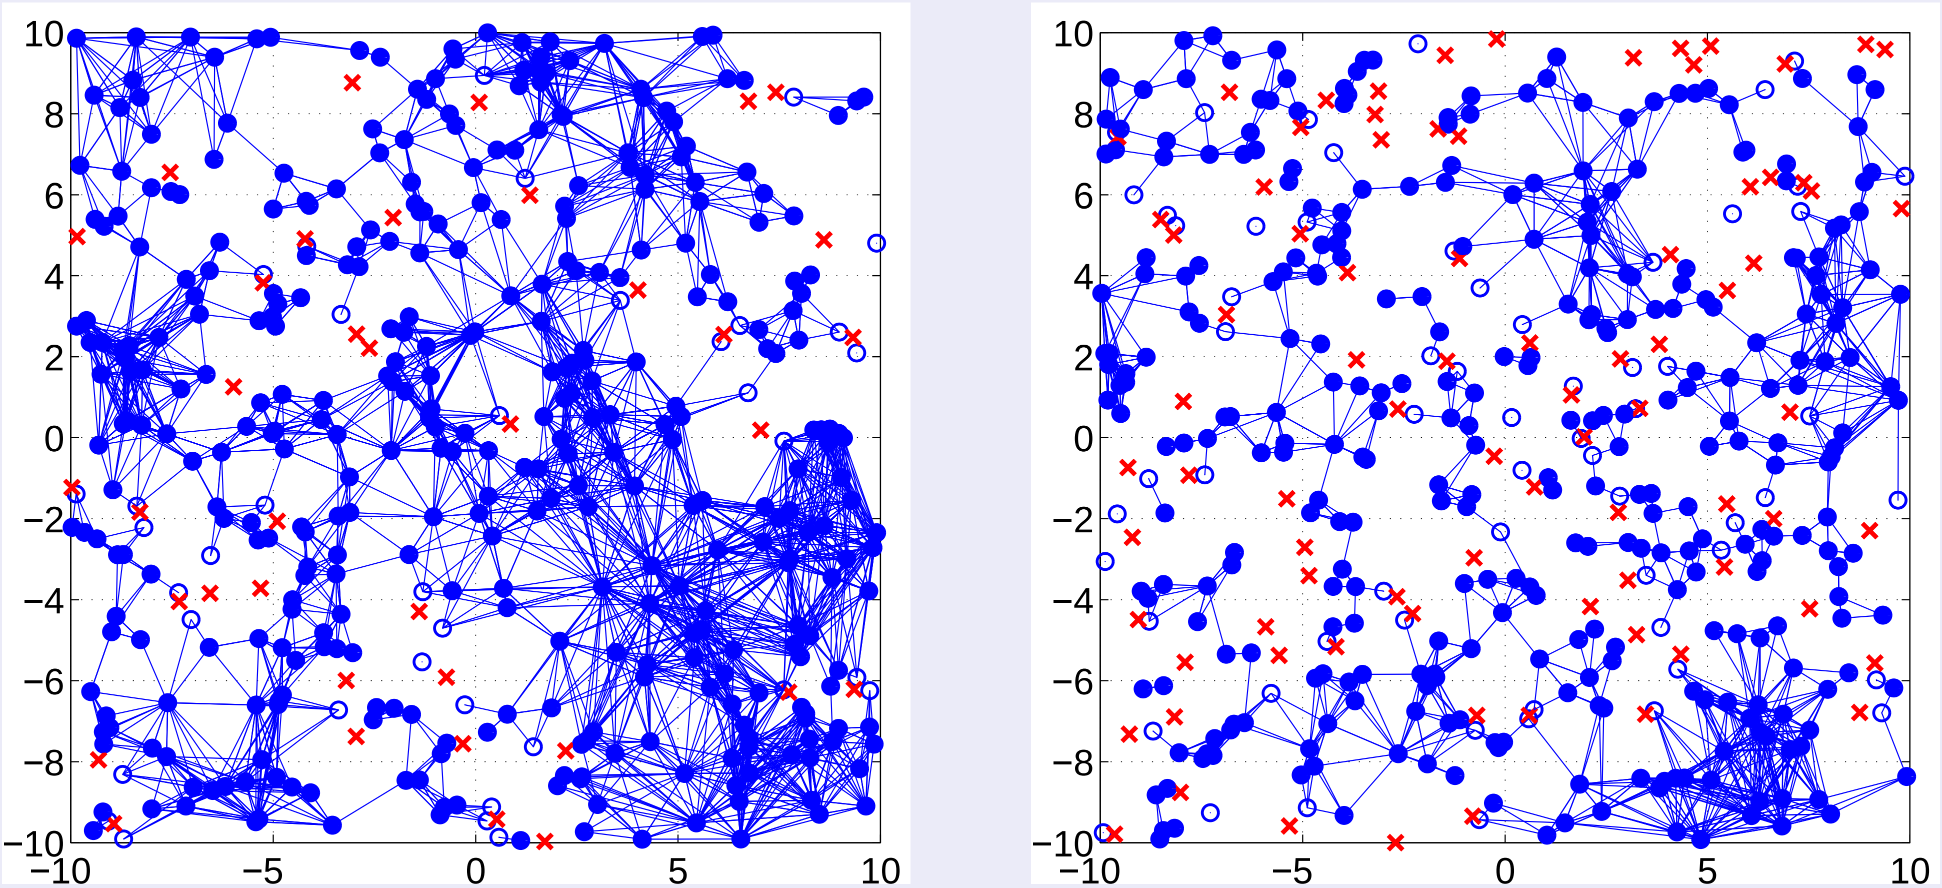

$\lambda = 0.1$ (percolates) $\lambda = 0.2$ (does not percolate)

Section Outline

1. Introduction
2. Bond Percolation on Lattice
3. Percolation on the Disk Graph
4. Secrecy Graphs
5. The SINR Graph
6. Summary

SINR-based Graph [DFM+06]

Definition

A link $x_1 \to x_2$ exists if

$$\frac{g(\|x_1 - x_2\|)}{W + \beta \sum_{i>2} g(\|x_i - x_2\|)} \geqslant \theta \,.$$

β is the interference reduction factor (processing gain). Two nodes are (bidirectionally) connected if there is a link $1 \to 2$ and a link $2 \to 1$.

Theorem [DFM+06]

There is a λ_c such that for all $\lambda > \lambda_c$ there exists a $\beta_c(\lambda)$ such that for $\beta < \beta_c(\lambda)$, there is a non-zero probability that a node belongs to an infinite component.

Outline of proof

The main idea is to couple the model to a bond percolation model on the lattice. In contrast to the disk graph, the edges here are dependent over a long range. Nonetheless is can be shown that by choosing β appropriately small, the squares in the lattice model can be crossed with a probability large enough so that bond percolation occurs.

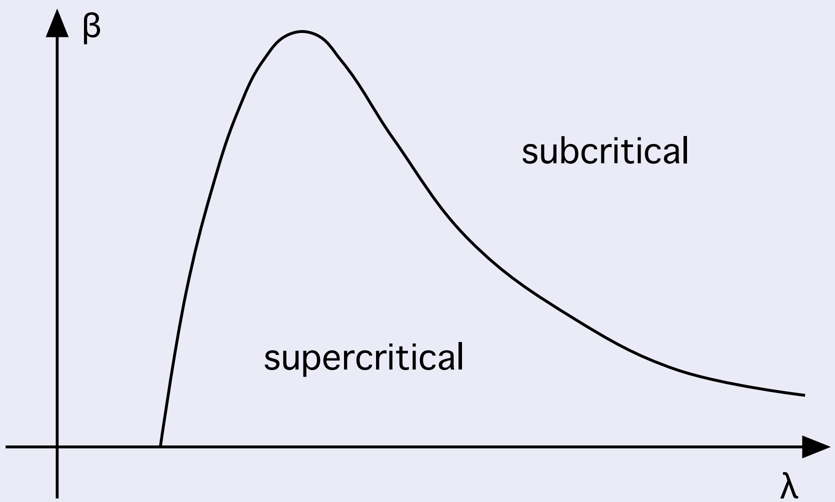

Qualitative behavior of the percolation domain for SINR model.

Comments

- First model that includes interference.
- Apparently $\beta \ll 1$ is needed.
- Open problem: Uniqueness of infinite component.
- Everybody transmits. Half-duplex constraint is violated.
- This is still a static graph (no fading).
- Interference reduction by β is non-trivial. It comes at the expense of bandwidth or time.

Section Outline

1 Introduction

2 Bond Percolation on Lattice

3 Percolation on the Disk Graph

4 Secrecy Graphs

5 The SINR Graph

6 Summary

Lecture 3 Summary

Random geometric graphs and percolation

- Full connectivity in Poisson networks requires increasing transmit power with growing network size. The main obstacle to connectivity are isolated nodes.
- Percolation can be achieved at finite power. This means that there exists an infinite component of connected nodes somewhere in the network. In the Poisson case, this component is unique.
- Critical radii or densities for percolation are unknown in most cases, but good bounds can usually be given.
- Percolation models have been extended to include interference [DFM+06] and, more recently, to secrecy.

References I

P. Balister, B. Bollobás, and M. Walters.
Continuum percolation with steps in the square or the disc.
Random Structures and Algorithms, 26(4):392–403, July 2005.

Béla Bollobás.
Random Graphs.
Cambridge University Press, 2001.
Second Edition. ISBN 0 521 80920 7.

O. Dousse, M. Franceschetti, N. Macris, R. Meester, and P. Thiran.
Percolation in the Signal-to-Interference Ratio Graph.
Journal of Applied Probability, 43(2):552–562, June 2006.

E.N. Gilbert.
Random plane networks.
Journal of the Society for Industrial Applied Mathematics, 9:533–543, 1961.

Martin Haenggi.
The Secrecy Graph and Some of its Properties.
In *2008 IEEE International Symposium on Information Theory (ISIT'08)*, Toronto, Canada, July 2008.
Available at http://www.nd.edu/~mhaenggi/pubs/isit08.pdf.

M. Penrose.
The Longest Edge of the Random Minimal Spanning Tree.
Annals of Applied Probability, 7:340–361, 1997.

Analysis and Design of Wireless Networks

Lecture 4: Connectivity and Coverage

Martin Haenggi

WIRELESS INSTITUTE
UNIVERSITY OF NOTRE DAME

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and a Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop Analysis of Poisson Networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 4 Overview

1. Connectivity
2. Coverage
3. Sentry Selection
4. Secrecy Coverage
5. Summary

Section Outline

1. Connectivity
   - Gilbert's disk graph
   - Single connectivity
   - k-connectivity
   - Nearest-neighbor graph
2. Coverage
3. Sentry Selection
4. Secrecy Coverage
5. Summary

Definition (Gilbert's disk graph [Gil61])

Take a stationary PPP of intensity λ as the vertices of a random geometric graph and connect two vertices by an edge if they are within distance r of each other. The resulting graph $G_{\lambda,r}$ is called a disk graph.

Example: $\lambda = 1$, $r = 1$

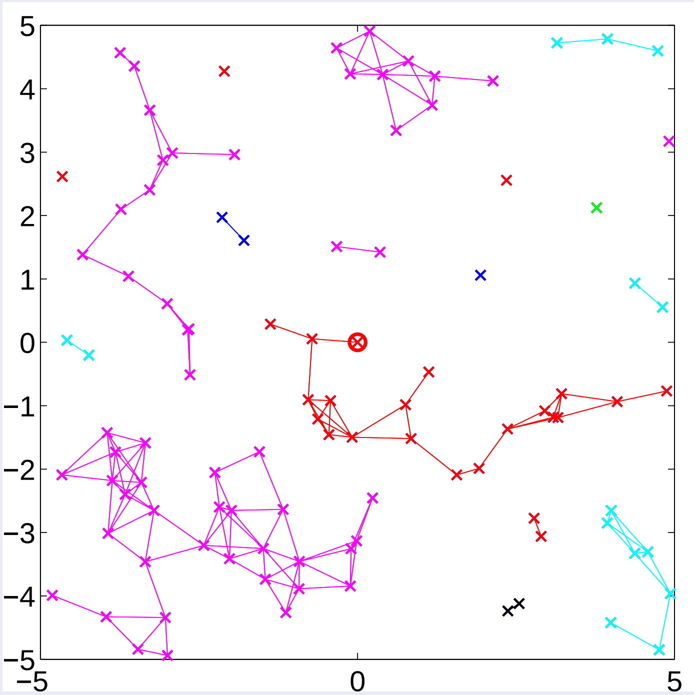

Interpretation

In the absence of interference and fading, the condition SNR $> \theta$ defines a maximum *communication radius*

$$r = \left(\frac{P}{\theta W}\right)^{1/\alpha} .$$

Connectivity

Let Φ be a PPP of intensity 1 on $[0, \sqrt{n}]^2$ so that $\mathbb{E}N = n$.

What communication radius r_c guarantees that $G_r(n)$ is connected whp as $n \to \infty$?

Formally, we want

$$\lim_{n\to\infty} \mathbb{P}[G_{r_c}(n) \text{ connected}] = 1 \,.$$

r_c will be a function of n.

Examples: $n = 100$

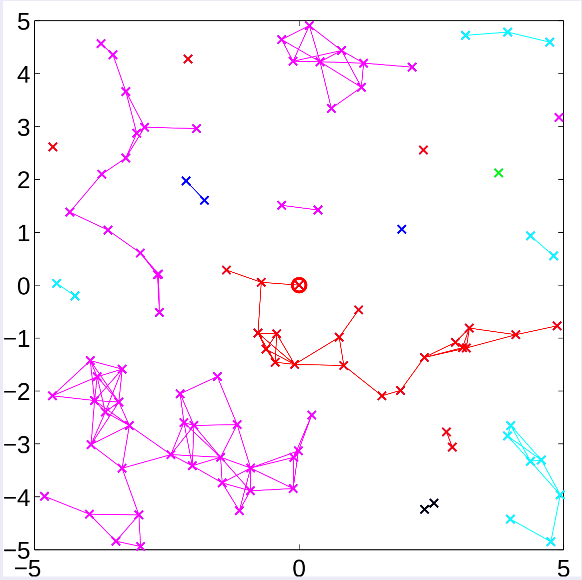

$r = 1$

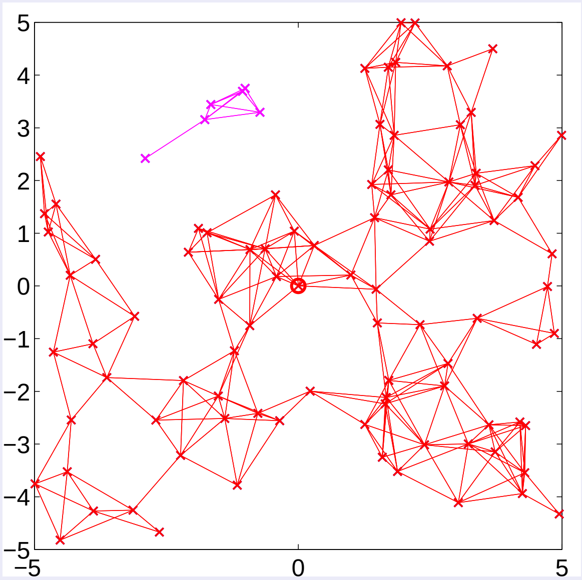

$r = 3/2$

Minimum transmission radius for connectivity

A necessary condition for connectivity is that no node is isolated.
The expected number of isolated nodes

$$\mathbb{E}N_{\text{isol}} = n\,\mathbb{P}(\text{typical node is isolated}) = n\exp(-\pi r^2)$$

needs to go to 0 as n grows. So we need $\pi r^2 \gtrsim \log n$ for connectivity.

[Pen97] showed that indeed the isolated nodes determine the connectivity such that

$$\pi r^2 = \log n + c(n)\,,$$

for $c(n) \to \infty$, *i.e.*, $c(n) = \omega(1)$, (arbitrarily slowly) is enough.

More precisely, setting $\pi r^2 = \log n + c$ results in a $N_{\text{isol}} \sim \text{Poi}(\exp(-c))$ and thus

$$\mathbb{P}(G_{r(n)}(n)\text{ connected}) \to \exp(-e^{-c})\,, \quad n \to \infty\,.$$

Different regimes of $G_{r(n)}(n)$

If $r(n) = $ const. (constant mean degree): This is the thermodynamic limit.
If $r(n)^2 \to 0$, the graph is sparse.
If $r(n)^2 \to \infty$, the graph is dense.
A special case of dense graphs are the ones in the connectivity regime, where

$$r(n) = \frac{c}{\pi}\sqrt{\log n}\,.$$

In this regime, the probability of a node being isolated is

$$P_{\mathrm{isol}} = \exp(-\pi r(n)^2) = n^{-c}\,,$$

and the number of isolated points is $nP_{\mathrm{isol}} = n^{1-c}$.
So if $c > 1$, there will be no isolated nodes as $n \to \infty$, whereas for $c < 1$, the number of isolated nodes goes to ∞.
If $r(n)^2/\log n \to \infty$, $G_{r(n)}(n)$ is in the superconnectivity regime.
If $r(n)^2/\log n \to 0$, $G_{r(n)}(n)$ is in the subconnectivity regime.

Bounds on isolation probability

What if $r(n)^2/\log n \to 1$?
Let

$$\pi r^2 = \log n + c(n)\,.$$

Then the mean number of isolated nodes is $e^{-c(n)}$, and the graph is disconnected with positive probability if $\limsup_n c(n) < \infty$. More precisely, the probability of being disconnected is bounded as

$$e^{-c}(1 - e^{-c}) \leq \lim_{n\to\infty} \mathbb{P}(\text{disconnected}) \leq 4e^{-c}\,,$$

where $c = \limsup_{n\to\infty} c(n)$ [GK98].

The proof of the upper bound is based on a result from continuum percolation, which says that the typical point lies either in an infinite component or is isolated a.s. as $n \to \infty$.

k-Connectivity

Condition for k-connectivity

k-connectivity means that any $k-1$ nodes can be removed and the graph is still connected. It implies the existence of k node-disjoint paths in the network (by Menger's theorem).

$$\pi r^2 = \log n + (k-1)\log\log n - \log(\underbrace{(k-1)!}_{\Gamma(k)}) + c(n)$$

results in k-connectivity a.a.s. if $c(n) \to \infty$ as $n \to \infty$, *i.e.*, $c(n) = \omega(1)$.

Penrose has shown that the graph is k-connected as soon as the smallest node degree is k [Pen99]. Equivalently, as soon as the number of nodes with degree $k-1$ goes to zero as $n \to \infty$.

Let's calculate the mean number of nodes with degree $k-1$.

k-connectivity: minimum node degree

Let N_{k-1} be the number of nodes with degree $k-1$. For $\pi r^2 = \log n + (k-1)\log\log n - \log(\Gamma(k))$,

$$\begin{aligned}\mathbb{E}(N_{k-1}) &= n\exp(-\pi r^2)\frac{(\pi r^2)^{k-1}}{\Gamma(k)}\\ &= n \cdot \frac{(\log n)^{1-k}\Gamma(k)}{n} \cdot \frac{[\log n + (k-1)\log\log n - \log\Gamma(k)]^{k-1}}{\Gamma(k)}\\ &= 1 + \Theta\left(\frac{\log\log n}{\log n}\right) = 1 + o(1)\,.\end{aligned}$$

If we make πr^2 a little larger, by $c(n) = \omega(1)$, the number of nodes with degree $k-1$ (or smaller) goes to zero.

So k-connectivity is achieved at small additional cost compared with simple connectivity. For $\pi r^2 = (1+\epsilon)\log n$, k-connectivity is achieved asymptotically for any k, for all $\epsilon > 0$.

Nearest-Neighbor Graph

k-nearest neighbor connectivity

Again let Φ be a PPP of intensity 1 on $[0, \sqrt{n}]^2$ so that $\mathbb{E}N = n$.
Assume a (bidirectional) edge exists from each node to its k nearest neighbors and denote the resulting graph by G_k.
Note that $\mathbb{E}(\deg G_k) > k$ since most nodes will have a degree larger than k.
What is the minimum k that guarantees connectivity whp?

Theorem [BBSW05]

$$0.3043 \log n < k < 0.5139 \log n$$

Remark

Easy to establish that $k = \Theta(\log n)$. Let $k = \lfloor c \log n \rfloor$. It was proven later that there is a sharp transition at some constant c_{crit}, such that for $c < c_{\text{crit}}$, the graph is disconnected a.a.s., and for $c > c_{\text{crit}}$, the graph is connected a.a.s.

Proof idea for lower bound

Careful analysis of situations that prevent a cluster of nodes from being connected. A likely scenario for disconnectedness is:

- The smallest disk contains a component of size at least $k + 1$.
- The annulus R_1 is empty.
- None of the nearest k neighbors of a node in R_3 is in the smallest disk.
  A node at x must have k neighbors inside A_3.

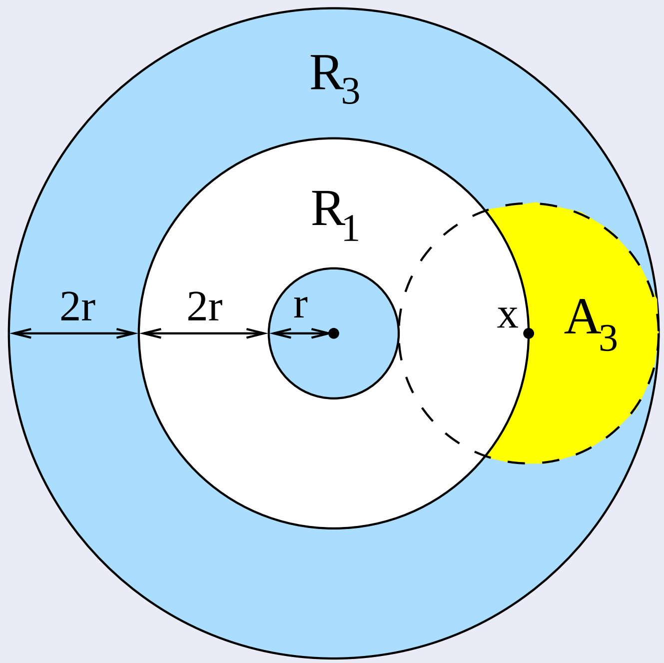

If $k = \log n/8$ and $\pi r^2 = k + 1$, this scenario is bound to occur as n grows. This establishes a lower bound of $\log n/8 < k$.

Section Outline

1 Connectivity

2 Coverage
- Single coverage
- k-coverage
- Cooperative coverage

3 Sentry Selection

4 Secrecy Coverage

5 Summary

Coverage

Coverage process

Let $\Phi = \{x_1, x_2, \ldots\} \subset \mathbb{R}^d$ be a point process, and $\{S_1, S_2, \ldots\}$ a collection of non-empty, possibly random sets. Then

$$\mathcal{C} = \{x_i + S_i : i = 1, 2, \ldots\}$$

is a coverage process.

The union of all sets in the coverage process is a germ-grain model, where x_i are the germs and S_i the grains.

Boolean models

If Φ is a stationary PPP and the S_i's are iid (and independent of Φ), then $\mathcal{C}$ is a Boolean model.

For coverage in sensor networks, often the S_i are just disks of radius r, *i.e.*, $S_i = S = \{x \in \mathbb{R}^2 \mid \|x\| < r\}$.

Concept of vacancy

The vacancy is the area of the part of the region of interest $\mathcal{R}$ that is not covered, *i.e.*,

$$V = \int_{\mathcal{R}} \chi(x)\mathrm{d}x\,,$$

where
$$\chi(y) = \mathbf{1}\{y \notin S_i + x_i\,,\ \forall i\} = \prod_i \mathbf{1}\{y \notin S_i + x_i\}\,.$$

Condition for single coverage

Consider the basic Boolean model with a PPP of intensity 1 and fixed disks of radius r. The probability that the origin is not covered is

$$\mathbb{E}\chi(o) = \exp(-\pi r^2)\,.$$

This holds for any point in $\mathbb{R}^2$, so the expected vacancy of a square of area n is

$$\mathbb{E}V(n) = n\exp(-\pi r^2)\,.$$

Condition for single coverage

For coverage, we need $\mathbb{E}V(n) \to 0$ as $n \to \infty$. This is guaranteed by

$$\pi r^2 = \log n + \omega(1)\,.$$

This is the same condition as for connectivity!

However, $\mathbb{E}V(n) \to 0$ does not guarantee (complete) coverage, for which we require $\mathbb{P}(V(n) = 0) \to 1$. (Convergence in mean does not imply a.s. convergence.)
To find the condition where $\mathbb{P}(V(n) = 0) \to 1$, we need an observation by Gilbert [Gil65]. Assuming disks are open sets, a necessary and sufficient condition for coverage is:

The area is covered if all intersections of disk boundaries are covered, and all intersections between disk boundaries and the border of the square of area n.

Let's look at the more general case of k-coverage.

k-Coverage

Gilbert's condition for k-coverage

Each intersection of boundaries must be covered k times.

Example ($r = 1/3$, $k = 2$)

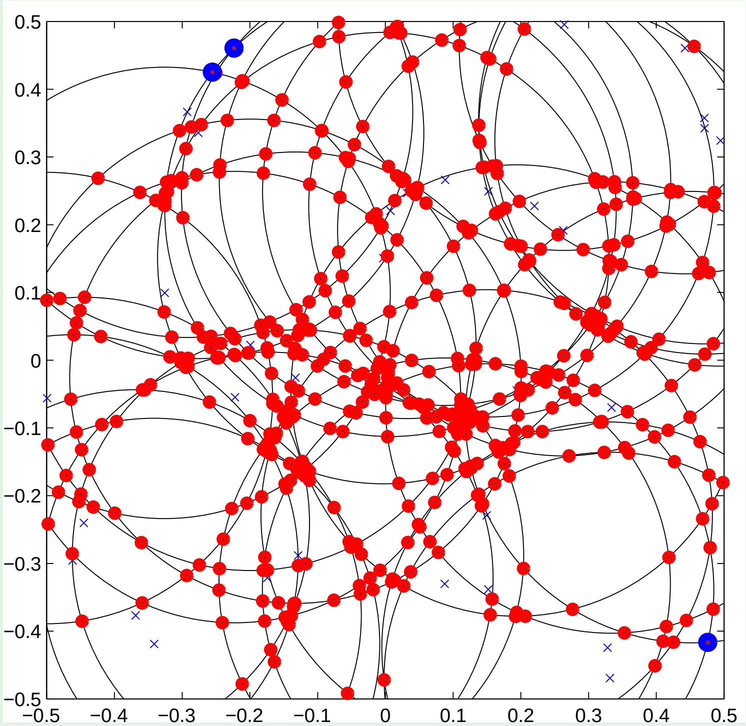

Here, disks of radius 1/3 were placed uniformly at random until 2-coverage was achieved.
33 disks were needed, giving rise to 545 intersections, most of them covered many times, on average 8.6 times.
The 3 blue intersections are covered exactly twice.

Condition for k-coverage

- The density of intersections is $4\pi r^2$, since a disk boundary ∂D intersects each disk boundary within distance $2r$.
- The expected number of intersections in the square of area n is $4\pi r^2 n = (4 + o(1))n \log n$.
- In a disk graph of radius r containing the Poisson points and the intersections, each intersection point needs to have degree at least k. The number of intersections with degree $k-1$ (or smaller) needs to go to zero, *i.e.*,

$$\mathbb{E}(I_{k-1}) = 4n \log n \exp(-\pi r^2) \frac{(\pi r^2)^{k-1}}{\Gamma(k)}$$

  should go to zero as $n \to \infty$.
- Similar to the connectivity problem, with an additional factor $\log n$. So we need to add a term $\log \log n$ in the expression for πr^2.

Condition for k-coverage

$$\pi r^2 = \log n + k \log\log n + \omega(1)$$

is the necessary and sufficient condition for $\mathbb{P}(V(n) \to 0) \to 1$, *i.e.*, k-coverage [Hal85].

More precisely, for $\pi r^2 = \log n + k \log\log n + x$,

$$\mathbb{P}(V(n) = 0) \to \exp\left(-\frac{e^{-x}}{\Gamma(k)}\right), \quad n \to \infty\,.$$

As in the connectivity problem, k-coverage requires only a slightly larger radius than single coverage.

Example ($n = 20^2$)

- $\mathbb{E}(V) \to 0$ at $\pi r^2 = \log n + c(n)$ or $r \approx 1.38$.
- $\mathbb{P}(V = 0) \to 1$ at $\pi r^2 = \log n + \log \log n + c(n)$ or $r \approx 1.57$.

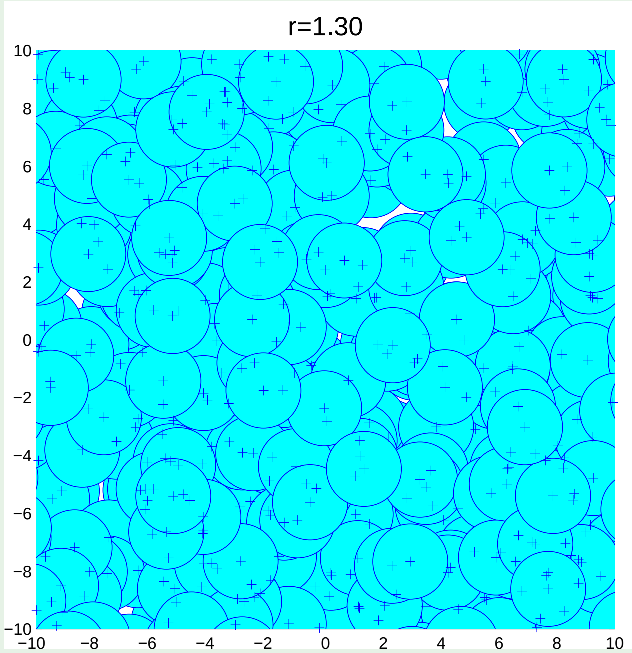

$r = 1.30$

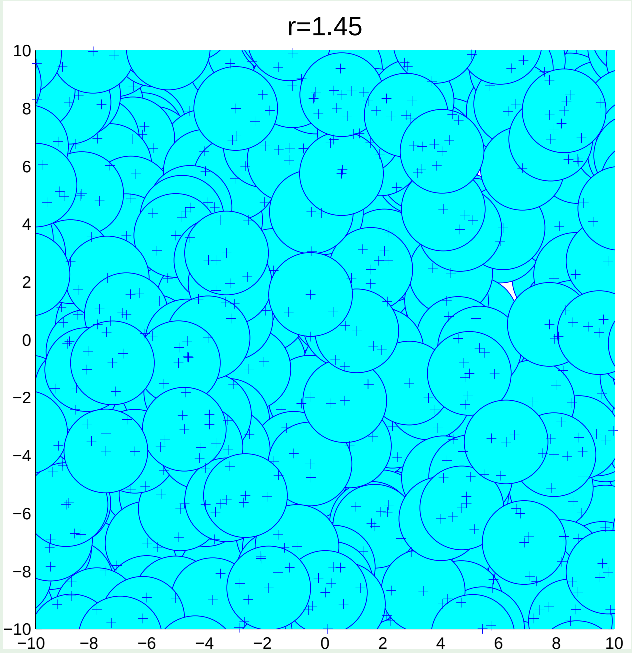

$r = 1.45$

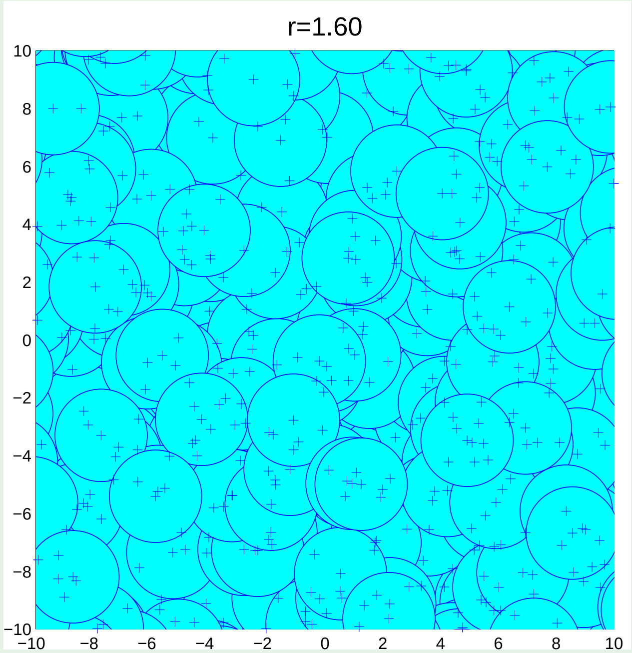

$r = 1.60$

Remark on single coverage

Let $c(n) = \Theta(\log\log\log n)$. In the intermediate regime

$$\log n + c(n) < \pi r^2 < \log n + \log\log n + c(n)\,,$$

we have $\mathbb{E}(V) \to 0$ but $\mathbb{P}(V = 0) \to 0$.

Since $V = V\mathbf{1}\{V > 0\}$, we have from Cauchy-Schwarz

$$\mathbb{P}(V > 0) \geq \frac{(\mathbb{E}V)^2}{\mathbb{E}(V^2)} \quad \Longrightarrow \quad \mathbb{P}(V = 0) \leq \frac{\operatorname{var} V}{\mathbb{E}(V^2)}\,.$$

Vacancy for general Boolean model

Hall showed that the vacancy result generalizes to [Hal88]

$$\mathbb{E}(V) = n \exp\big(-\lambda\mathbb{E}(\|S\|)\big)$$

for a PPP of intensity λ and iid S_i.

Cooperative coverage

The cooperative coverage problem

- In sensor networks, nodes may cooperate to jointly achieve coverage. A node further away from a phenomenon will contribute less to the joint coverage than a nearby node.
- If ξ occurs at location y and is to be measured, the contribution from a node x at distance $r = \|x - y\|$ is $\xi r^{-\alpha}$.
- Assume sensor nodes form a PPP of intensity λ, and that the phenomenon occurs at the origin o. The total cooperative measurement is

$$G(\lambda) = \sum_{x \in \Phi} \xi \|x\|^{-\alpha}$$

- We would like to find the minimum λ such that $\mathbb{P}(G(\lambda) > \beta) > 1 - \epsilon$.

Duality

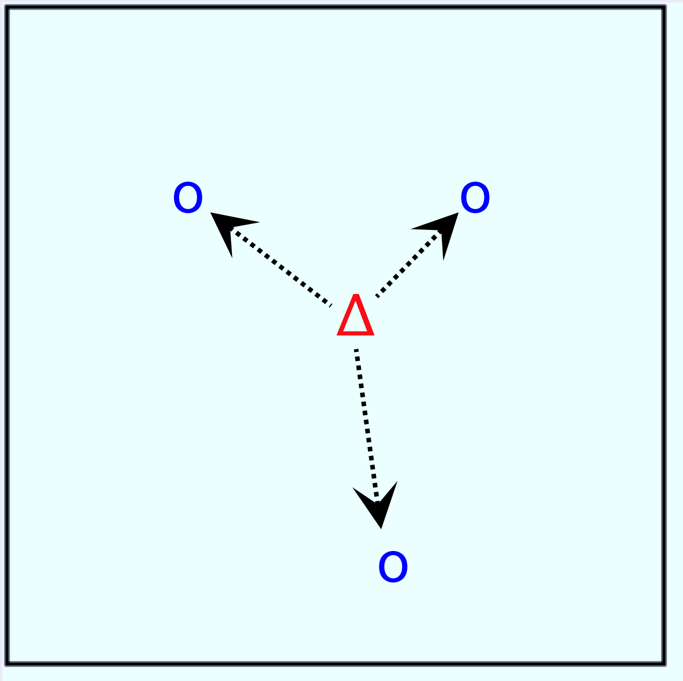

Cooperative coverage

Sensor nodes o combine their signals.

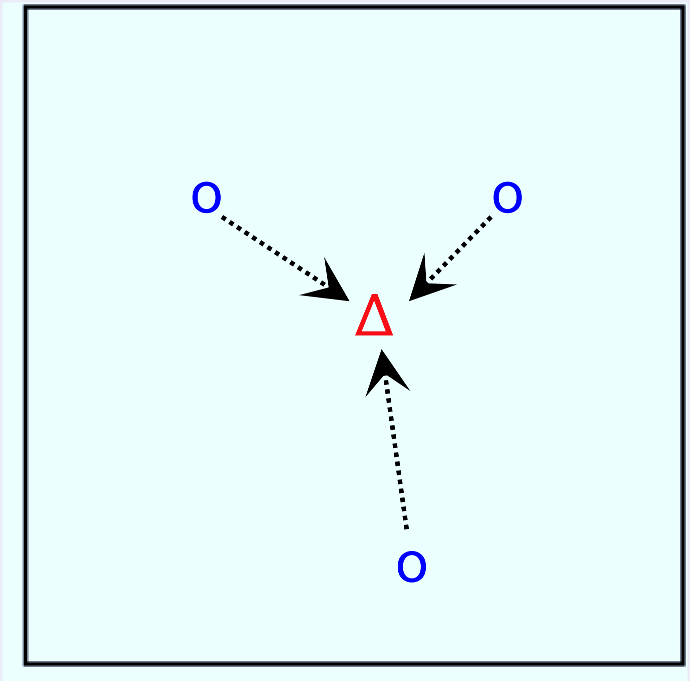

Interference

At the phenomon location Δ, measure the total interference.

If the signal decay law is the same, then the cumulative measurement G at the sensor nodes equals the interference power measured at the location of the phenomenon if all nodes transmitted.

The cooperative coverage problem

- The problem is exactly dual to the interference problem!
- The cooperative coverage question is the same as asking:
  What is the minimum λ such that the interference at o exceeds β with probability at least $1-\epsilon$ if all nodes transmit at power ξ?
- We can apply all we know about interference in Poisson networks [VHC06].

Section Outline

1 Connectivity

2 Coverage

3 Sentry Selection
- Problem formulation
- k-cover vs. k single covers
- Main results
- Proof sketch
- Algorithmic and distributed aspects

4 Secrecy Coverage

5 Summary

What is Sentry Selection?

Standing guard means that a few soldiers (or animals) monitor the surroundings so that the rest can sleep (or eat). Those guards are thus acting as sentries so that the rest can save energy and recover.

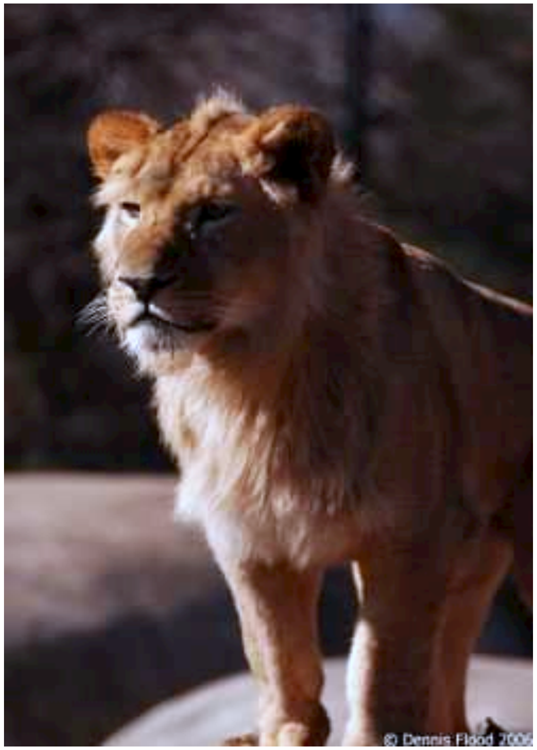

Sentry selection is:

$\Longrightarrow$

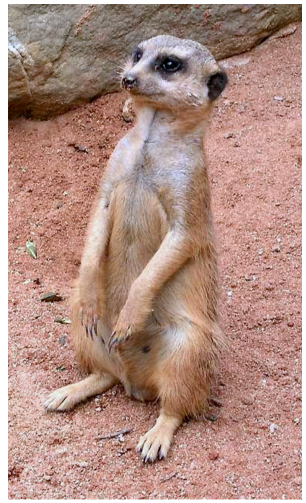

Sentry selection in sensor networks

Putting nodes to sleep is the only efficient way to save energy.
In surveillance applications, how do we choose the subset of nodes acting as sentries (standing guard) in a given period?
- The sentries should be able to monitor the entire area.
- After a certain period, a disjoint set of sensors shall assume sentry duty.

Problem Formulation

Setting

Consider a Poisson point process $\Phi = \{x_i\} \subset \mathbb{R}^2$ of unit intensity, and let S_n be the square of area n centered at the origin. So $\mathbb{E}|\Phi \cap S_n| = n$.
Each point x_i is the center of an open disk of radius r, and we focus on the set of disks $\mathcal{C}_r(n)$ that intersects S_n.
Note: This avoids boundary problems.

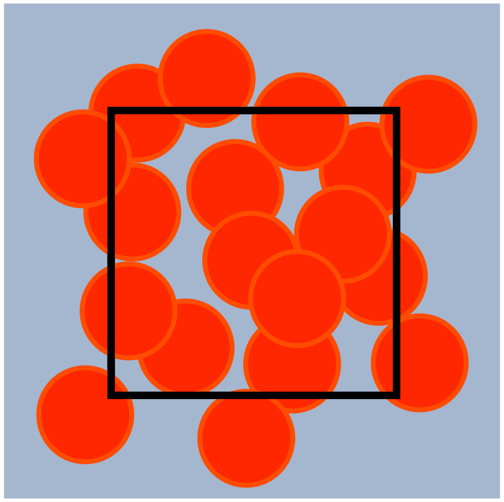

Question

What is the smallest r such that whp we may partition $\mathcal{C}_r(n)$ into k classes such that each point $y \in S_n$ is contained in a disk of each class:

$$r(n,k) = \inf_r \left\{ \lim_{n\to\infty} \mathbb{P}[\mathcal{C}_r(n) \text{ is } k\text{-partitionable}] = 1 \right\}.$$

Example ($r = 1/5$, $k = 2$)

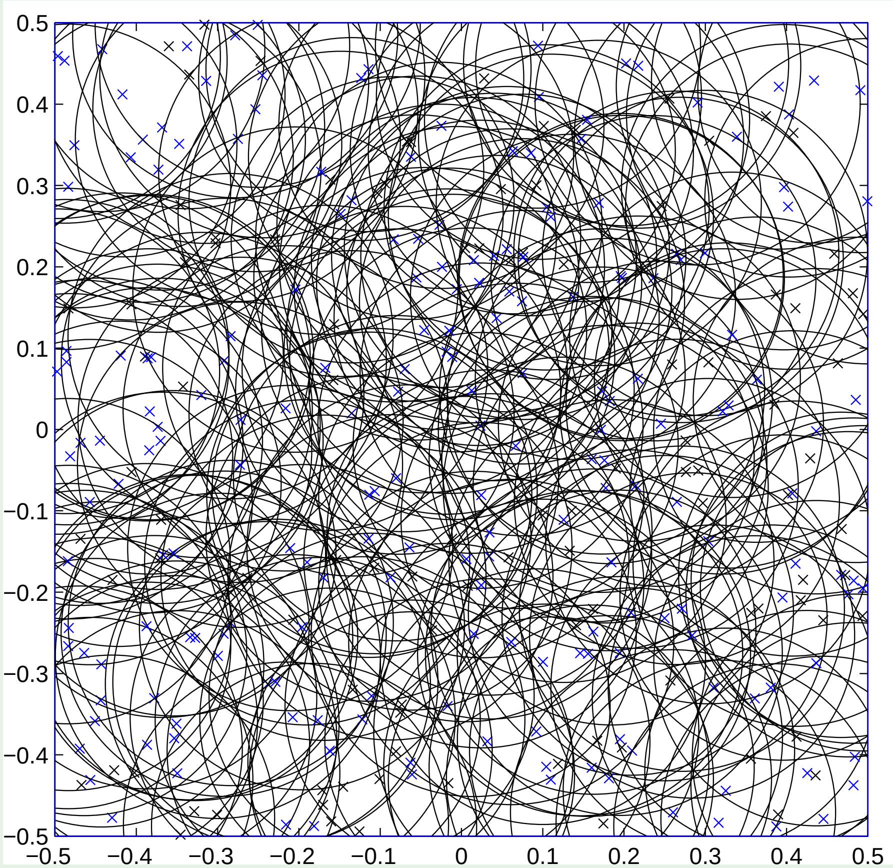

100 disks on the unit square.

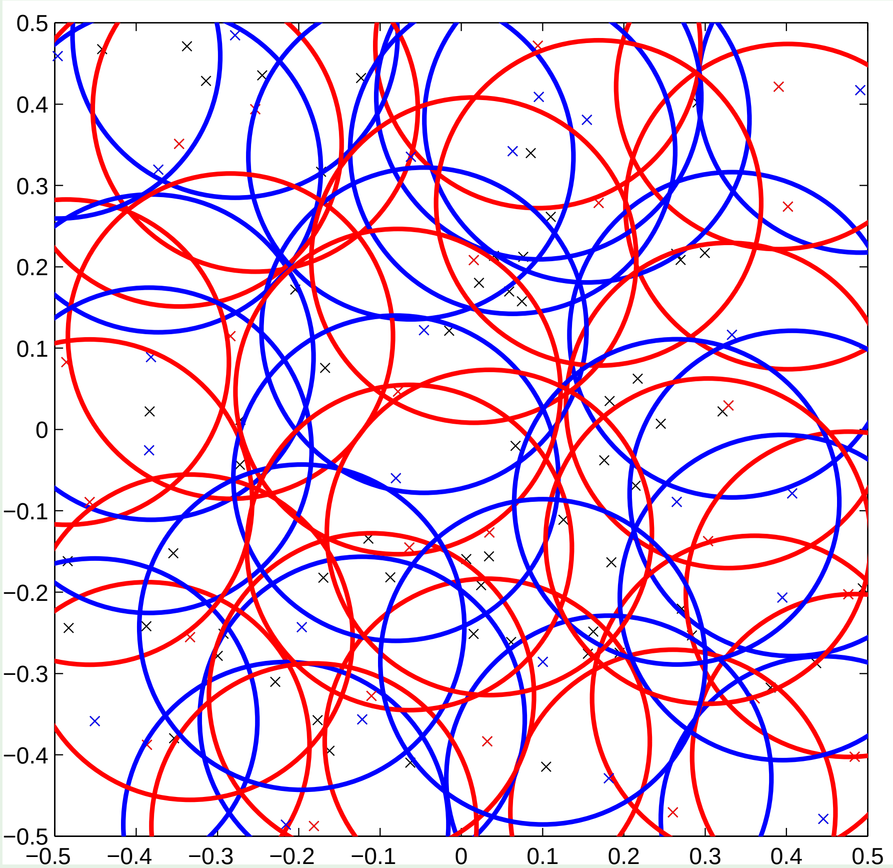

Both red and blue disks cover the square. Only 46 disks are needed.

k-Cover vs. k Single Covers

Definition (k-cover)

Consider a set of disks $\mathcal{C}_r(n) = \{C_1, \ldots, C_n\}$ of radius r, with $C_i \subset \mathbb{R}^2$.
$\mathcal{C}_r(n)$ forms a k-cover of $\mathcal{S} \subset \mathbb{R}^2$ iff

$$\min_{x \in \mathcal{S}} \left\{ \sum_{i=1}^{n} \mathbf{1}_{C_i}(x) \right\} = k\,.$$

Necessity vs. sufficiency

Clearly, a k-cover is necessary. But in general, it is not sufficient.
For instance, let $\mathcal{S}$ be the set of all subsets of $A = \{1, 2, \ldots, n\}$ of size k.
The n sets $S_i = \{B \in \mathcal{S} \,:\, i \in B\}$, $i \in A$, form a k-cover of $\mathcal{S}$ which cannot even be partitioned into two single covers if $2k \geqslant n$.
So a solution to our problem must make use of its geometric setting.

Example (A k-cover does not imply k single covers)

Let $n = 4$, $k = 2$.
The set of subsets of $\{1, 2, 3, 4\}$ of size $k = 2$ has $\binom{4}{2} = 6$ elements:

$$\mathcal{S} = \Big\{\{1,2\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\}\Big\}$$

Then let

$$S_1 = \Big\{\{1,2\},\{1,3\},\{1,4\}\Big\}$$
$$S_2 = \Big\{\{1,2\},\{2,3\},\{2,4\}\Big\}$$
$$S_3 = \Big\{\{1,3\},\{2,3\},\{3,4\}\Big\}$$
$$S_4 = \Big\{\{1,4\},\{2,4\},\{3,4\}\Big\}$$

Each element of $\mathcal{S}$ is an element in exactly $k = 2$ sets S_i, but there is no way to partition S_i into two single covers.

Main Results

Theorem (k-cover vs. k single covers [BBSW10])

Let E_r^k be the event that the disks of radius r form a k-cover, and F_r^k the event that they are partitionable into k single covers. Then

$$\mathbb{P}(E_r^k \setminus F_r^k) \leq \frac{c_k}{\log n} \,.$$

Interpretation

This implies that asymptotically, we have k single covers as soon as we have a k-cover. The critical radius is the same. This does not mean that all k-covers are partitionable, but it means that such configurations appear with probability tending to zero.

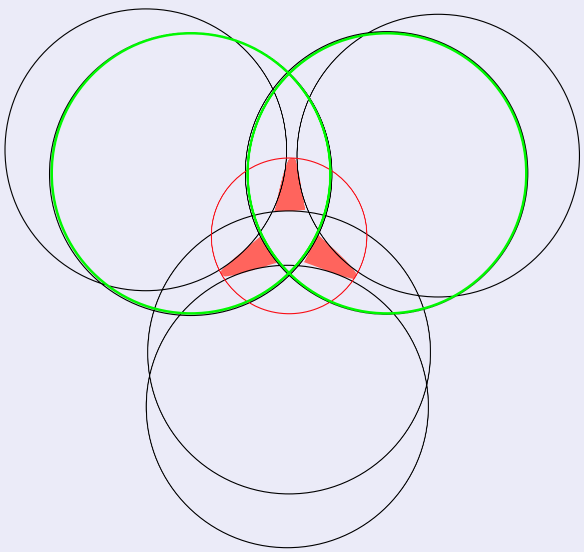

A non-partitionable 2-cover.

Theorem (Sufficient condition)

Let $k \in \mathbb{N}$ and r be given by

$$\pi r^2 = \log n + (2k+1)(\log\log n)^2 + \omega(1)\,,$$

where $\omega(n) \to \infty$ as $n \to \infty$. Then whp $\mathcal{C}_r(n)$ can be partitioned into k single covers of S_n.

Remarks

- $\log n$ is needed "anyway" (for simple coverage or connectivity).
- This is sufficient but not necessary.

Proof Sketch

1. Gilbert's observation

Let the disks be open disks. Then $\mathcal{C}_r(n)$ is (at least) a k-cover if every intersection of disk boundaries is covered (at least) k times.
Expected number of intersections within S_n: $4\pi r^2 n = (4 + o(1))n \log n$.

Example ($r = 1/3$, $k = 2$)

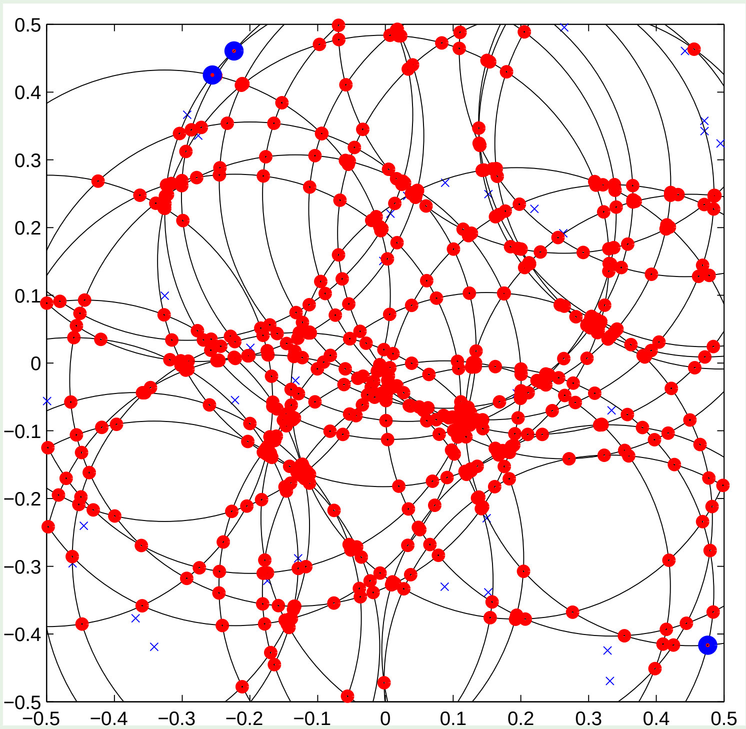

33 disks, 545 intersections, most of them covered many times, on average 8.6 times. The 3 blue intersections are covered exactly twice.

2. s-cover

With r given in the theorem, $\pi r^2 = \log n + (2k+1)(\log\log n)^2 + \omega(1)$, $\mathcal{C}_r(n)$ forms an s-cover of S_n whp, where $s = \lfloor (2k+1)\log\log n \rfloor$.

Idea: Show that $\mathbb{P}(\text{Po}(\pi r^2) \leqslant s-1) = o(1/n\log n)$ for this choice of s.

3. Bounding the number of nearby intersections

There are at most $677(\log n)^2$ intersections within distance $2r$ of each intersection. Any intersection within distance $2r$ of a fixed intersection x is defined by two points within distance $3r$. It can be shown that:

$$\mathbb{P}(\text{Po}(9\pi r^2) \geqslant 26\log n) = o(1/n\log n)\,.$$

So there are at most $26\log n + 2$ points within $3r$, giving rise to no more than $677(\log n)^2$ intersections.

4. Use the Probabilistic Method

Pick a fixed instance of the random cover. Assume it forms an s-cover.
Randomly color the disks with one of k colors.
If there is a positive probability that each intersection is covered by a disk of each color, we're done.
For a given intersections x, let A_x be the *bad* event that x is not covered by a disk of each color.
If two intersections y and z are at least $2r$ apart, the two sets of disks covering y and z are disjoint. Thus A_y is independent of the σ-algebra generated by the events $\{A_z : \|y - z\| > 2r\}$.
So A_y is independent of all but at most $d = 677(\log n)^2$ events. Since y is covered by at least s disks,

$$\mathbb{P}(A_y) \leqslant k\left(1 - \frac{1}{k}\right)^s \qquad \text{(union bound)}\,,$$

where $s = \lfloor (2k+1)\log\log n \rfloor$.

5. Apply the Lovász Local Lemma

Putting the pieces together, we can show

$$\mathbb{P}(A_y) < \frac{1}{ed}\,,$$

where $d = 677(\log n)^2$.
With the Lovász Local Lemma, this implies that the probability of the intersection of all *positive* events $\bar{A}_j$ is positive.
Note that $[(d-1)/d]^{d-1} > 1/e$ for $d \geqslant 2$.

The fact that *random coloring* achieves the goal with positive probability proves the existence of a coloring scheme.

The (symmetric) Lovász Local Lemma

Let $A_1, \ldots, A_m$ be events whose dependence graph has maximal degree $d > 2$. If

$$\mathbb{P}(A_i) \leqslant \frac{(d-1)^{d-1}}{d^d}$$

then

$$\mathbb{P}\left(\bigcap_{j=1}^{m} \bar{A}_j\right) > 0\,.$$

Existence of Efficient Distributed Algorithms

Due to its probabilistic nature, the proof is not constructive.
A standard hypergraph coloring algorithm (intersections are hyperedges) does a good job. But not in a distributed fashion.

Theorem (Tradeoff between coverage radius and local decisions)

Fix $k \in \mathbb{N}$. Let $A \triangleq \log n + k(\log n)^{2/3}$ and consider a Poisson point set of density one in the $\sqrt{n} \times \sqrt{n}$ box. Place a disk of area A centered at each point. Then whp the point (disk) set can be split into k covers. Moreover, each point can decide which cover it is in by looking only at the other points within $(\log n)^{1/6}$ of it.

Since the coverage radius is about $\sqrt{\log n}$, each point only needs to know about its very near neighbors. However, for nodes in a dense neighborhood, it should need much less, maybe a node only needs to know (at most) its k nearest neighbors.

Result

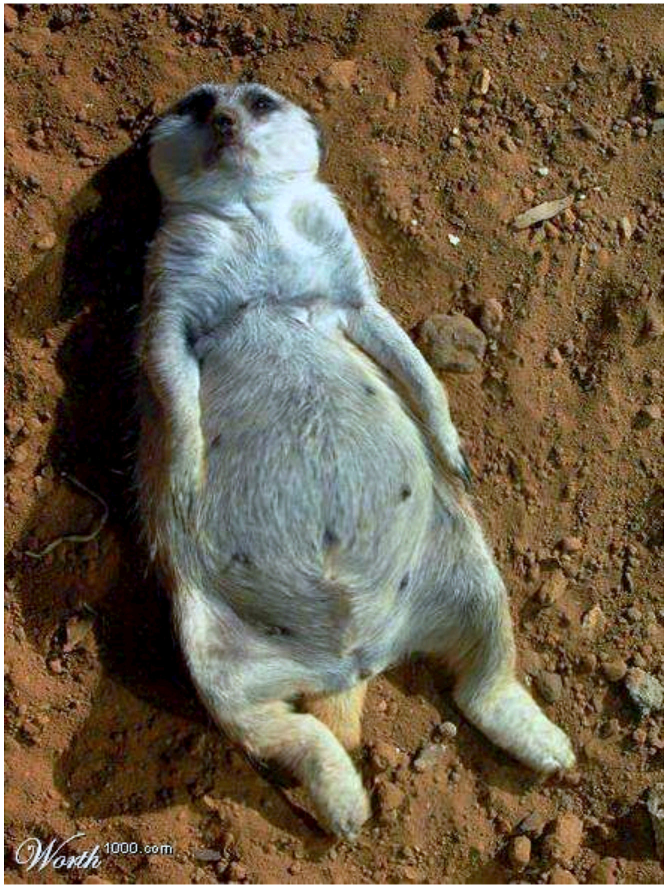

Section Outline

1. Connectivity
2. Coverage
3. Sentry Selection
4. Secrecy Coverage
   - Setup
   - In one dimension
   - In two dimensions
5. Summary

Setup

- Take a PPP of intensity 1 of base stations, and another, independent PPP of intensity λ of eavesdroppers.
- Each base station covers a disk-shaped area whose radius is determined by the nearest eavesdropper.

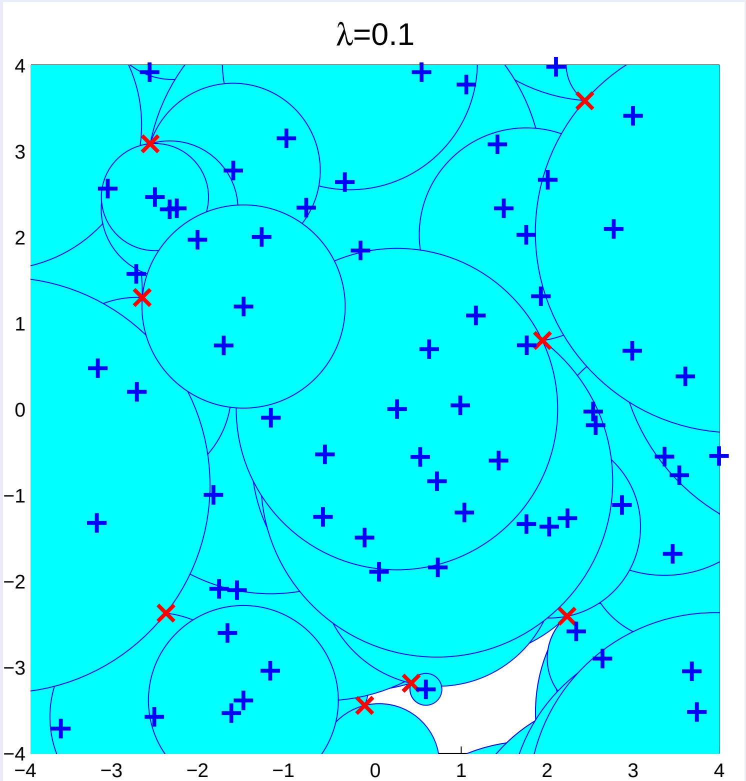

Focusing on the square of area n, what $\lambda(n)$ guarantees coverage?

The One-Dimensional Case

Covered volume fraction

For $\lambda > 0$,

$$C^1(\lambda) = \mathbb{P}(0 \text{ is covered})$$

is the *covered volume fraction*.

It is not difficult to show that

$$C^1(\lambda) = \frac{1+4\lambda}{(1+2\lambda)^2}\,.$$

Vacancy

So, if $V(n)$ is the total length of the vacant parts in an interval of length n,

$$\mathbb{E}V(n) = n(1 - C^1(\lambda)) = \frac{4\lambda^2 n}{1+4\lambda+4\lambda^2}\,.$$

For $\mathbb{E}V(n) \to 0$ as $n \to \infty$, we need $\lambda^2 n \to 0$ or $\lambda = o(1/\sqrt{n})$.

Secrecy coverage result in one dimension [SH10]

If there are two consecutive red points, the interval in between is not covered. If X is the number of consecutive red points, we have

$$\mathbb{E}X \sim \lambda^2 n$$

and $\mathbb{P}(V(n)=0) \leq \mathbb{P}(X=0)$. Bounding $\mathbb{P}(X=0)$ and applying Janson's inequality yields that for $\lambda^2 n \to c$, $\mathbb{P}(V(n)=0) \leq P_n \to e^{-c}$.
So for $\lambda^2 n \to \infty$, $\mathbb{P}(V(n)=0) \to 0$.

A more detailed analysis reveals that

$$\mathbb{P}(V(n)=0) \sim e^{-4\lambda^2 n},$$

so that for $\lambda^2 n \to 0$, $\mathbb{P}(V(n)=0) \to 1$.

So in this case, $\mathbb{E}(V) \to 0$ and $\mathbb{P}(V=0) \to 1$ happen simultaneously.

In Two Dimensions

Asymptotic results [SH10]

The two-dimensional case requires significantly more work.

The currently best bounds are:

- For $\lambda^3 n(\log n)^3 \to 0$, $\mathbb{P}(V=0) \to 1$.
- For $\lambda^3 n \to \infty$, $\mathbb{P}(V=0) \to 0$.

There is hope that the gap between the two $\lambda(n)$ can be narrowed. At any rate, $\lambda \approx n^{-1/3}$ is the right order.

This model can be viewed as a germ-grain model with correlated germ size. The correlation increases the required mean disk size significantly compared to the independent case, where $\lambda = o(1/\log n)$ is sufficient.

Section Outline

1 Connectivity

2 Coverage

3 Sentry Selection

4 Secrecy Coverage

5 Summary

Lecture 4 Summary

Connectivity and coverage

- Full connectivity in Poisson networks requires increasing transmit power with growing network size. The main obstacle to k-connectivity are nodes with degree $k - 1$.
- The coverage problem is quite different in nature. It only requires a local analysis. However, it turns out that in the Poisson case, the critical radii are closely related.
- Sentry selection and secrecy coverage are examples of ongoing research problems. They can be formulated cleanly as pure mathematical problems, yet are practically motivated.

References I

P. Balister, B. Bollobás, A. Sarkar, and M. Walters.
Connectivity of Random k-nearest-neighbor Graphs.
Adv. in Appl. Probab., 37:1–24, 2005.

P. Balister, B. Bollobás, A. Sarkar, and M. Walters.
Sentry Selection in Wireless Networks.
Advances in Applied Probability, 42(1):1–25, 2010.

E.N. Gilbert.
Random plane networks.
Journal of the Society for Industrial Applied Mathematics, 9:533–543, 1961.

E. N. Gilbert.
The probability of covering a sphere with n circular caps.
Biometrika, 56:323–330, 1965.

P. Gupta and P. R. Kumar.
Stochastic Analysis, Control, Optimization and Applications, chapter "Critical Power for Asymptotic Connectivity in Wireless Networks", pages 547–566.
Birkhauser, Boston, 1998.
ISBN 0-8176-4078-9.

P. Hall.
On the coverage of k-dimensional space by k-dimensional spheres.
The Annals of Probability, 13(3):991–1002, August 1985.

Peter Hall.
Introduction to the Theory of Coverage Processes.
Wiley Series in Probability and Mathematical Statistics, 1988.

References II

M. Penrose.
The Longest Edge of the Random Minimal Spanning Tree.
Annals of Applied Probability, 7:340–361, 1997.

M. Penrose.
On k-Connectivity for a Geometric Random Graph.
Random Structures and Algorithms, 15(2):145–164, 1999.

A. Sarkar and M. Haenggi.
Secrecy Coverage.
Internet Mathematics, 2010.
Submitted. Available at http://www.nd.edu/~mhaenggi/pubs/im10.pdf.

Jagadish Venkataraman, Martin Haenggi, and Oliver Collins.
Shot Noise Models for the Dual Problems of Cooperative Coverage and Outage in Random Networks.
In *44th Annual Allerton Conference on Communication, Control, and Computing (Allerton'06)*, Monticello, IL, September 2006.
Available at http://www.nd.edu/~mhaenggi/pubs/allerton06.pdf.

Analysis and Design of Wireless Networks

Lecture 5: Multi-hop Analysis of Poisson Networks

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop Analysis of Poisson Networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 5 Overview

1. Routing in Poisson Networks
2. Correlation in Poisson Networks
3. Local Delay in Poisson Networks
4. Information Propagation in Poisson Networks
5. Summary

Section Outline

1. Routing in Poisson Networks
   - Network model
   - Random access transport capacity
   - Delay analysis with queueing
2. Correlation in Poisson Networks
3. Local Delay in Poisson Networks
4. Information Propagation in Poisson Networks
5. Summary

Multihop Network Model

PPP network model of M-hop routes

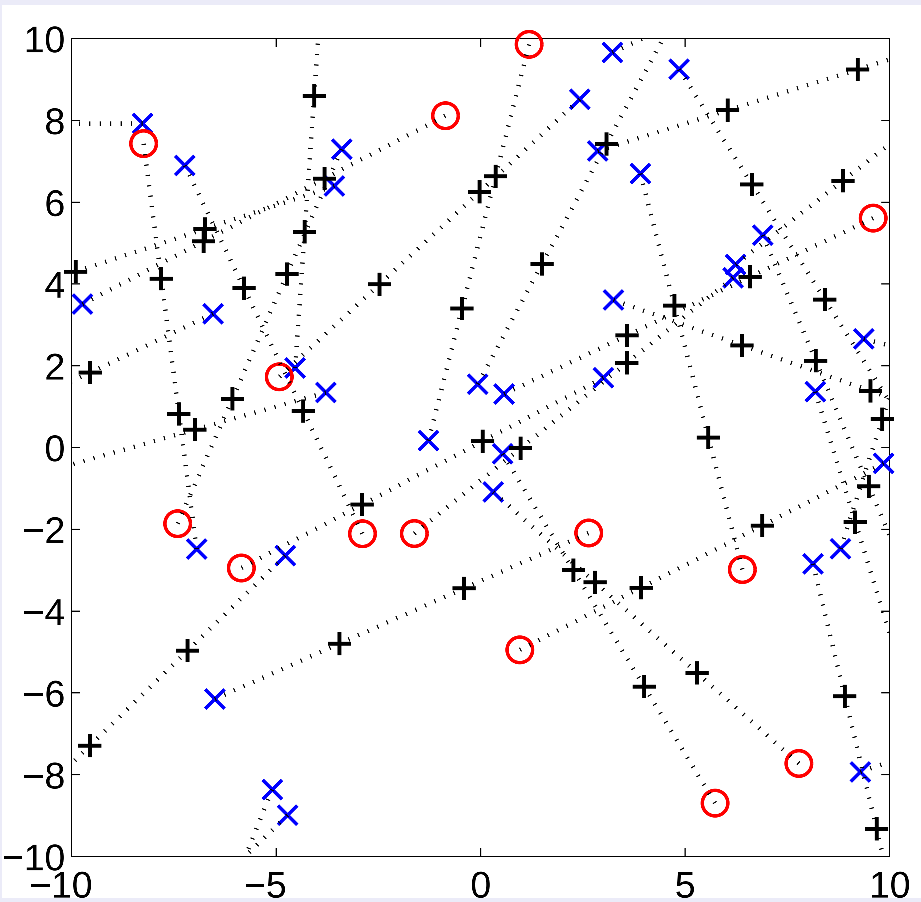

- x: PPP of intensity λ of sources.
- o: Destinations, at fixed distance R in random direction.
- +: $M-1=2$ relays per route placed equidistantly on the SD line.

Channel access: Only one node per route transmits. This means the set of transmitters forms a PPP at all times.

Random access transport capacity [AWKH10]

Retransmission scheme and delay bound

- At each hop, the packet is transmitted until received successfully.
- The number of transmissions at hop m is geometric and denoted by $T_m(M)$. The total number of transmissions is

$$T(M) = \sum_{m=1}^{M} T_m(M) \,.$$

- The maximum number of transmissions is restricted to A. Hence if $T(M) \le A$, the packet is successfully delivered. If $T(M) > A$, there is an outage, and the *actual* number of transmissions is $\min\{T(M), A\}$. Thus the effective rate per route is

$$\log(1+\theta)\mathbb{P}(T(M) \le A)/\min\{T(M), A\}.$$

- M can be chosen from $\{1, 2, \dots, A\} = [A]$. $M = 1$ means A single-hop transmissions, while $M = A$ means no retransmissions are possible.

Definition (Random access transport capacity)

$$C(A) = \max_{M\in[A]} \lambda R\, \mathbb{P}(T(M) \le A) \frac{\log(1+\theta)}{\mathbb{E}\min\{T(M), A\}}$$
$$= \lambda R \log(1+\theta) \max_{M\in[A]} \frac{\mathbb{P}(T(M) \le A)}{\mathbb{E}\min\{T(M), A\}}$$

Upper bound

Since

$$\frac{\mathbb{P}(T(M) \le A)}{\mathbb{E}\min\{T(M), A\}} \le \frac{1}{\mathbb{E}\,T(M)}$$

and $\mathbb{E}\,T(M) = M/p_s(M)$,

$$C(A) < \lambda R \log(1+\theta) \max_{M\in[A]} \frac{p_s(M)}{M}\,.$$

Since $p_s(M) = \exp(-c(R/M)^2)$ there is an optimum M.

Optimum number of hops

$$M_{\text{opt}} = \arg\max_{M \in [A]} \frac{p_s(M)}{M}$$

For $\alpha = 4$ (and without noise),

$$M_{\text{opt}} = \min\{A, R\sqrt{\lambda}\theta^{1/4}\pi\}\,.$$

The resulting upper bound on the random access transport capacity is

$$C < \frac{\sqrt{\lambda}\log(1+\theta)}{\pi\sqrt{e}\theta^{1/4}}\,.$$

Noise can be included in the analysis [AWKH10].

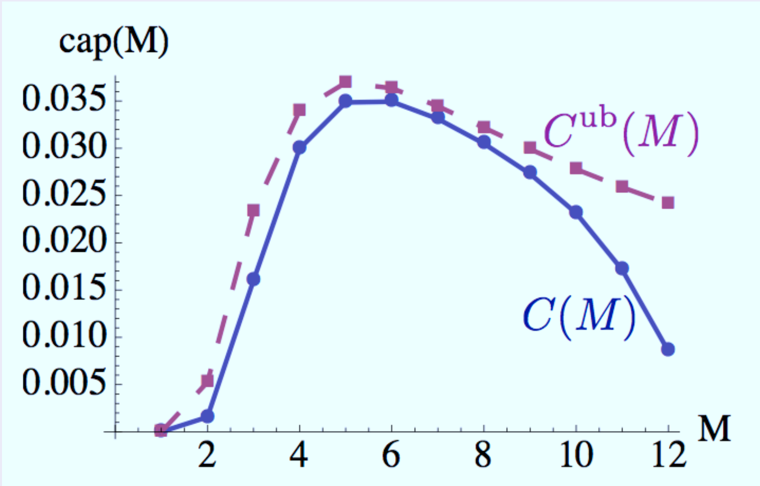

Comparison of actual $C(M)$ and upper bound.

Delay Analysis with Queueing [SH10]

The TDMA-ALOHA channel access scheme

The previous model did not include queueing delays, since there is no queueing.
If nodes are not permitted to re-transmit in each slot until successful, queueing delays become important.

Use the same network model, but relax the assumption of equal hop length and change the channel access scheme to TDMA-ALOHA: In each route, a token is passed from node to node, and the node with the token is allowed to transmit with probability p.

But this node only transmits when it has a packet. So even if scheduled to transmit, the node may not contribute to the interference.
So we do not make the "heavy-traffic assumption", where all nodes always have packets. Only the source is backlogged.

Success probabilities

Let $p_s(n)$ be the success probability at the n-th hop. Packet arrivals are geometric with parameter $pp_s(1)$ (source traffic intensity), provided that $p_s(1) < p_s(n)$, $\forall n > 1$. Then the queue of the $(n-1)$-th relay is not empty with probability $p_s(1)/p_s(n)$.

It follows that the point process of interferers is a PPP with intensity

$$\lambda_I = \lambda p_I = \frac{\lambda p}{M} \sum_{n=1}^{M} \frac{p_s(1)}{p_s(n)},$$

where p_I is the probability that a node is allowed to transmit *and has a packet*.

So λ_I depends on $p_s(n)$, $n \in [M]$. But

$$p_s(n) = \mathbb{P}(\mathrm{SIR}_n > \theta) = \exp(-\lambda_I c r_n^2),$$

which introduces an intricate inter-dependence between λ_I and $p_s(n)$!

End-to-end delay

We obtain

$$\lambda_I = \frac{\lambda p}{M} \sum_{n=1}^{M} \exp(-\lambda_I c(r_1^2 - r_n^2) \,,$$

which is a fixed point equation for λ_I given r_n, $n \in [M]$.
The end-to-end delay is

$$D = H_1 + \sum_{n=2}^{M} H_n + Q_n = \frac{M}{pp_s(1)} + M \sum_{n=2}^{M} \frac{1 - pp_s(n)}{pp_s(n) - pp_s(1)} \,,$$

where H_n is the service time and Q_n is the waiting time at node n.

To minimize the delay, necessarily $r_2 = r_3 = \ldots = r_M = (R - r_1)/(M-1)$.
So the problem is to find M_{opt} and $r_{1,\mathrm{opt}}$, such that $D(M, r_1)$ is minimized.
Direct analytical optimization is not possible, but numerical evaluation is straightforward.

Numerical results ($\alpha = 4$, $\lambda = 10^{-4}$, $p = 0.05$)

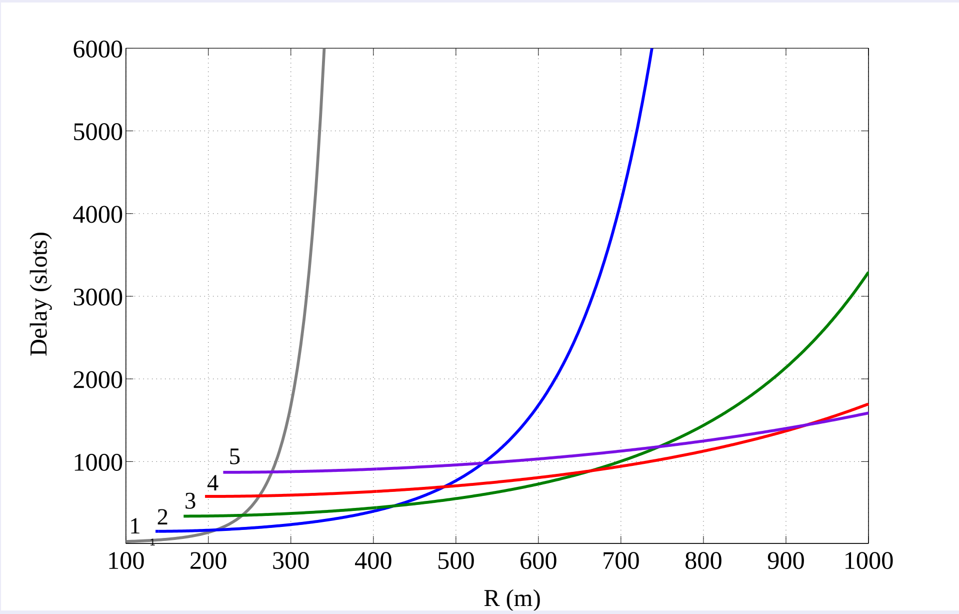

End-to-end delay

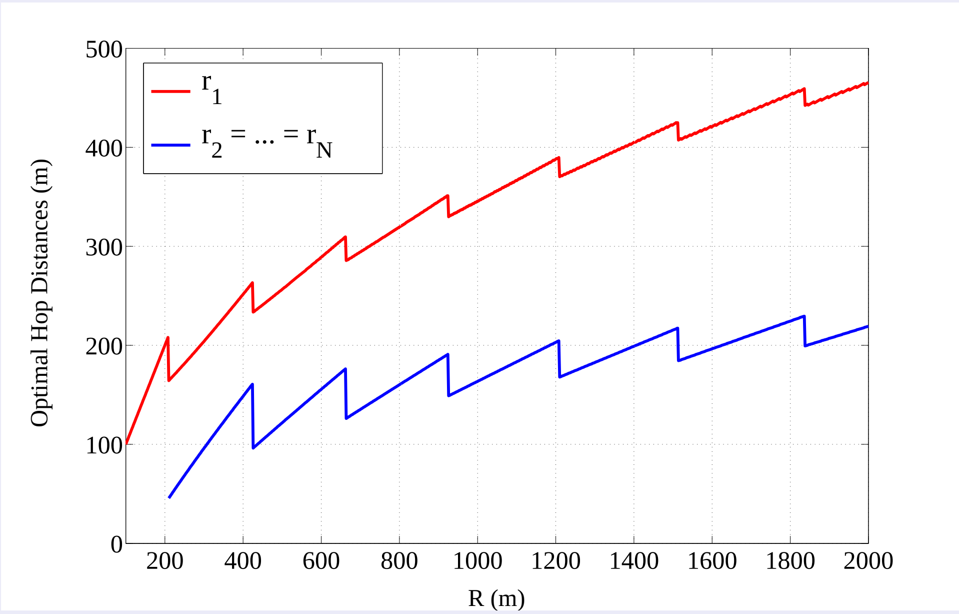

Hop distances

Each jump in the right plot corresponds to a crossover point in the left curve. These are the points when $M+1$ hops are better than M hops.

These results provide optimum relay locations and thus give guidelines for routing. In practice, relays will have to be chosen from a point process, so there is some deviation between the optimum and actual relay location.

A critical assumption

In both models discussed, a critical assumption was made:

The transmission success events in subsequent time slots and across hops are independent.

Strictly speaking, this is never completely true.
If there is significant mobility, this assumption is more likely to hold. But in static networks, there is correlation between outage events.

Such correlations are the topic of the next section.

Section Outline

1. Routing in Poisson Networks
2. Correlation in Poisson Networks
   - Introduction
   - Spatiotemporal correlation
   - Outage correlation
3. Local Delay in Poisson Networks
4. Information Propagation in Poisson Networks
5. Summary

Correlation in Poisson Networks

Intuition (PPP with ALOHA probability p)

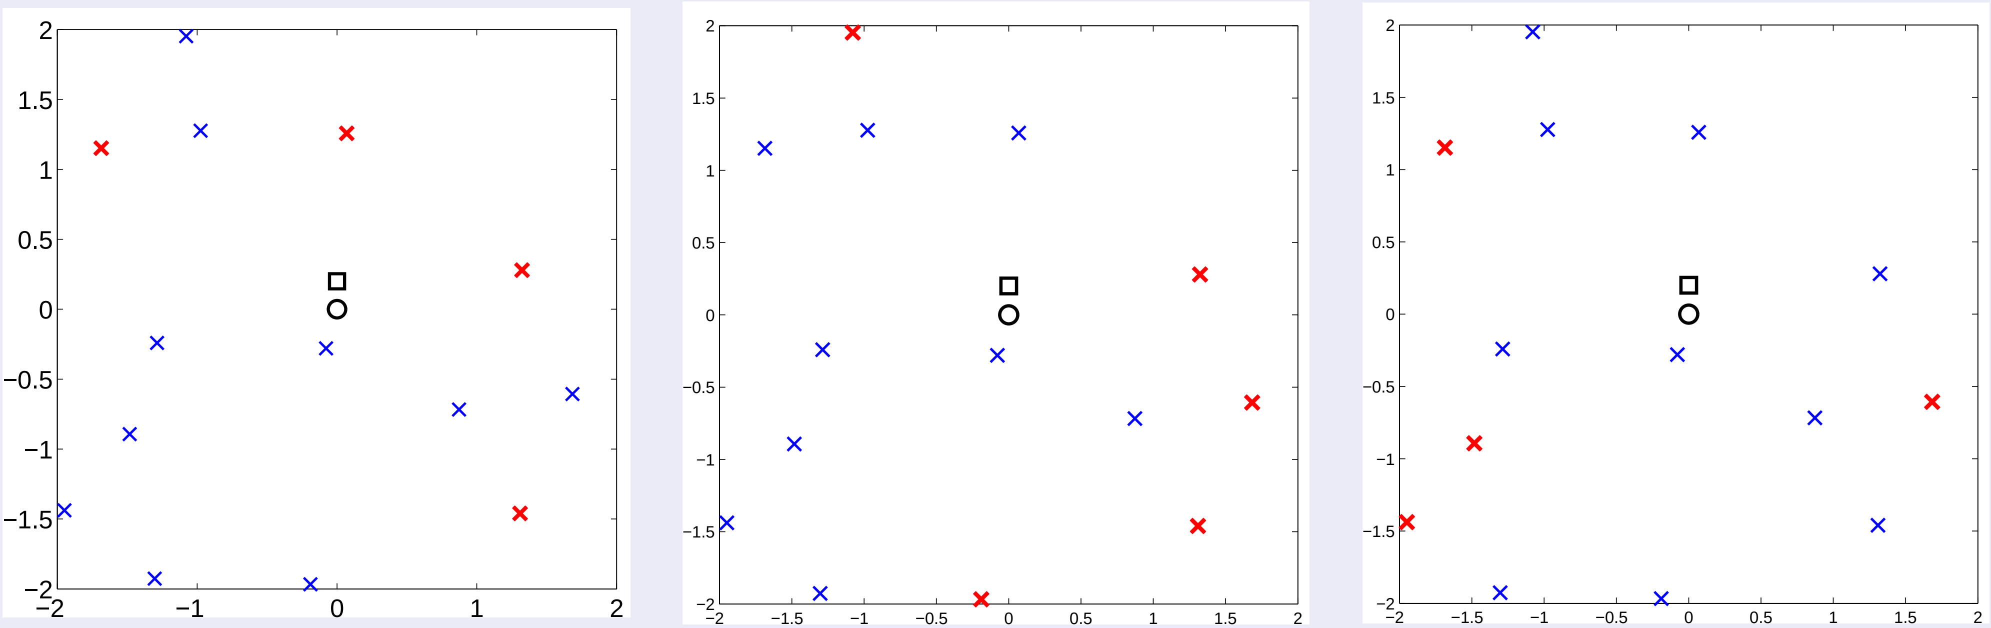

Since the PPP is static (common randomness), there is temporal correlation of the interference at o in different time slots.

There is also spatial correlation between the interference measured at nearby points ○ and □.

Interference correlation: Setup

- A PPP $\Phi \subset \mathbb{R}^2$ with ALOHA with transmit prob. p and iid fading.
- Let $I_k(u)$ be the interference measured at u in time slot k.

The distribution of $I_k(u)$ is the same for all $k \in \mathbb{Z}$ and $u \in \mathbb{R}^d$, but the common randomness Φ introduces dependence. The random MAC and fading help de-correlate the interference, but not fully.

For example: Let's assume $p = 1$ and no fading. Then $I_k(u)$ and $I_\ell(u)$ would be perfectly correlated, for all $k, \ell \in \mathbb{Z}$.

Definition (The spatio-temporal correlation coefficient)

For path loss laws $g(x) \colon \mathbb{R}^2 \to \mathbb{R}^+$ for which the interference has a finite second moment and $k \neq \ell$,

$$\zeta(u, v) \triangleq \frac{\mathbb{E}[I_k(u) I_\ell(v)] - \mathbb{E}[I_k(u)]^2}{\mathbb{E}[I_k(u)^2] - \mathbb{E}[I_k(u)]^2} .$$

Calculation of the moments

For all $k \in \mathbb{Z}$ and $u \in \mathbb{R}^2$, $I_k(u) \stackrel{d}{=} I_0(o)$.
The first moment, $\mathbb{E}I_k(o)$, follows directly from Campbell's theorem:

$$\mathbb{E}I_k(o) = p\lambda \int_{\mathbb{R}^2} g(x)\mathrm{d}x \,.$$

The second moment is

$$\begin{aligned}\mathbb{E}(I_k(o)^2) &= \mathbb{E}\left[\left(\sum_{x \in \Phi_k} h_{xo} g(x)\right)^2\right] \\ &= p\mathbb{E}(h^2)\lambda \int_{\mathbb{R}^2} g^2(x)\mathrm{d}x + p^2\mathbb{E}(h^2)\lambda^2 \int_{\mathbb{R}^2}\int_{\mathbb{R}^2} g(x)g(y)\mathrm{d}x\mathrm{d}y \,,\end{aligned}$$

which follows from the *second-order product density* of the PPP (see Lecture 6).

Spatio-temporal correlation [GH09]

Spatio-temporal correlation coefficient of $I_k(u)$ and $I_\ell(v)$, $k \neq \ell$:

$$\zeta(u,v) = \frac{p \int_{\mathbb{R}^2} g(x) g(x - \|u - v\|) \mathrm{d}x}{\mathbb{E}(h^2) \int_{\mathbb{R}^2} g^2(x) \mathrm{d}x} \,.$$

Temporal correlation

Setting $u = v$ yields the temporal correlation coefficient. For Nakagami-m fading, it is simply

$$\zeta_t = p \frac{m}{m+1} \,.$$

- The correlation is proportional to the transmit probability p.
- Fading helps decorrelate the interference. In Rayleigh fading, the correlation coefficient is $p/2$.
- Different MAC schemes and channels with memory exhibit stronger correlation, so this is a lower bound.

Impact of interference correlation on outage

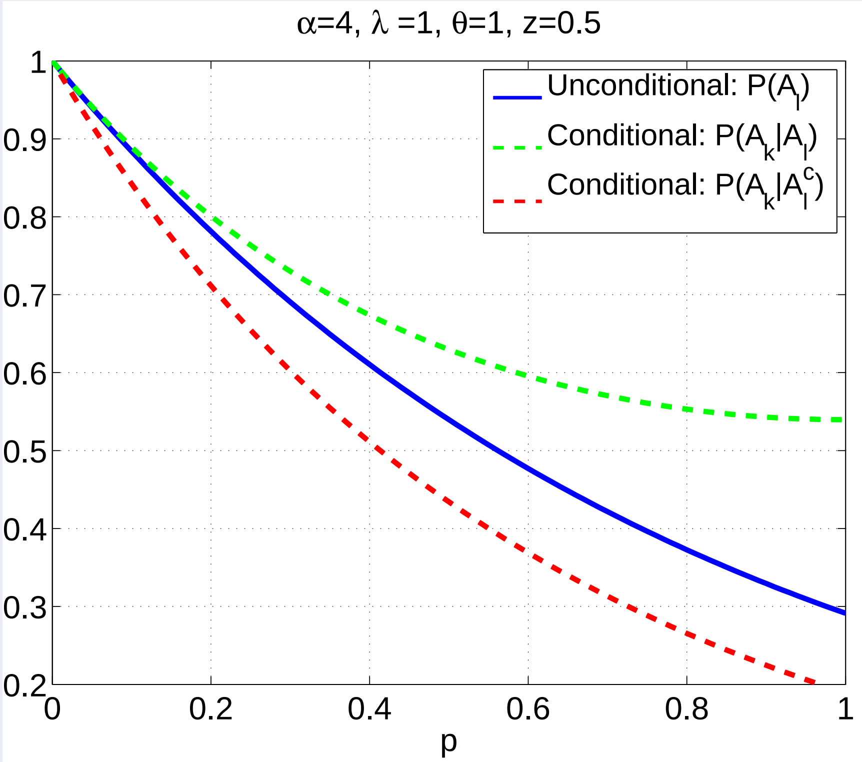

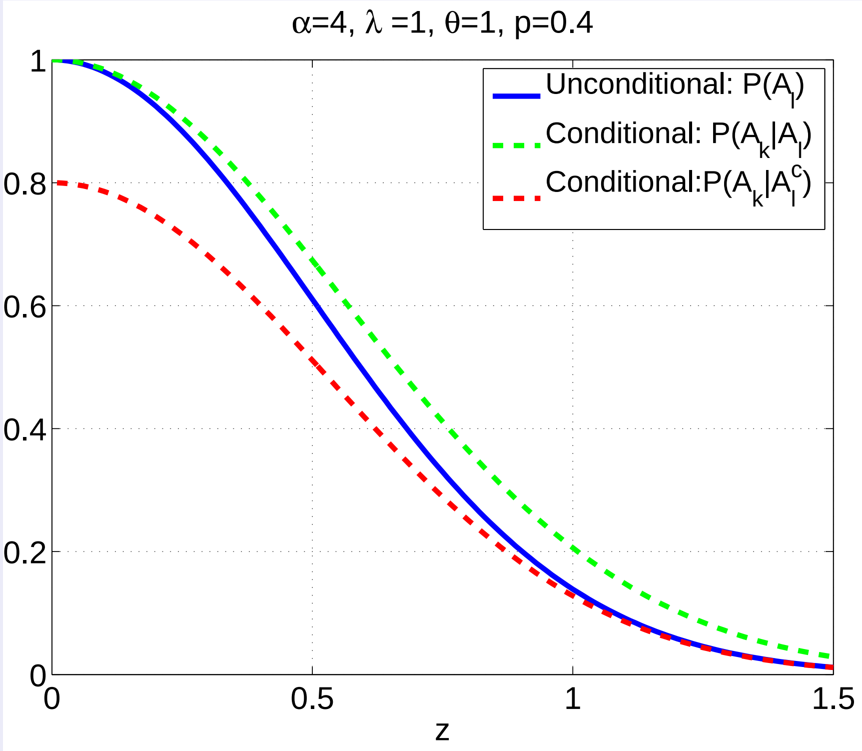

$g(x) = \|x\|^{-4}$, $\theta = 1$. Follows from the joint success probability (pgfl)

$$\mathbb{P}(A_u, A_v) = \exp\left(-\lambda \int_{\mathbb{R}^2} 1 - \left(\frac{p}{1+\theta g(x)/g(z)} + 1 - p\right)^2 \mathrm{d}x\right).$$

This has an impact on retransmission schemes and delays.

Section Outline

1 Routing in Poisson Networks

2 Correlation in Poisson Networks

3 Local Delay in Poisson Networks
- Setup and problem formulation
- High-mobility networks
- Static networks
- Summary of results
- Remarks on the noisy case

4 Information Propagation in Poisson Networks

5 Summary

Local Delay

Basic question

How long does it take for a node in a (Rayleigh fading) Poisson network with ALOHA to successfully communicate with its nearest neighbor?

Scenarios

High-mobility network:
For each attempt, a new realization of the PPP is drawn.

Static network:
The realization stays the same over time.

Nearest-neighbor transmission:
Transmit to nearest node or to nearest receiver. (Or receive from nearest node or nearest transmitter.)

Nearest-neighbor transmission (NNT) in a static network

3 time slots:

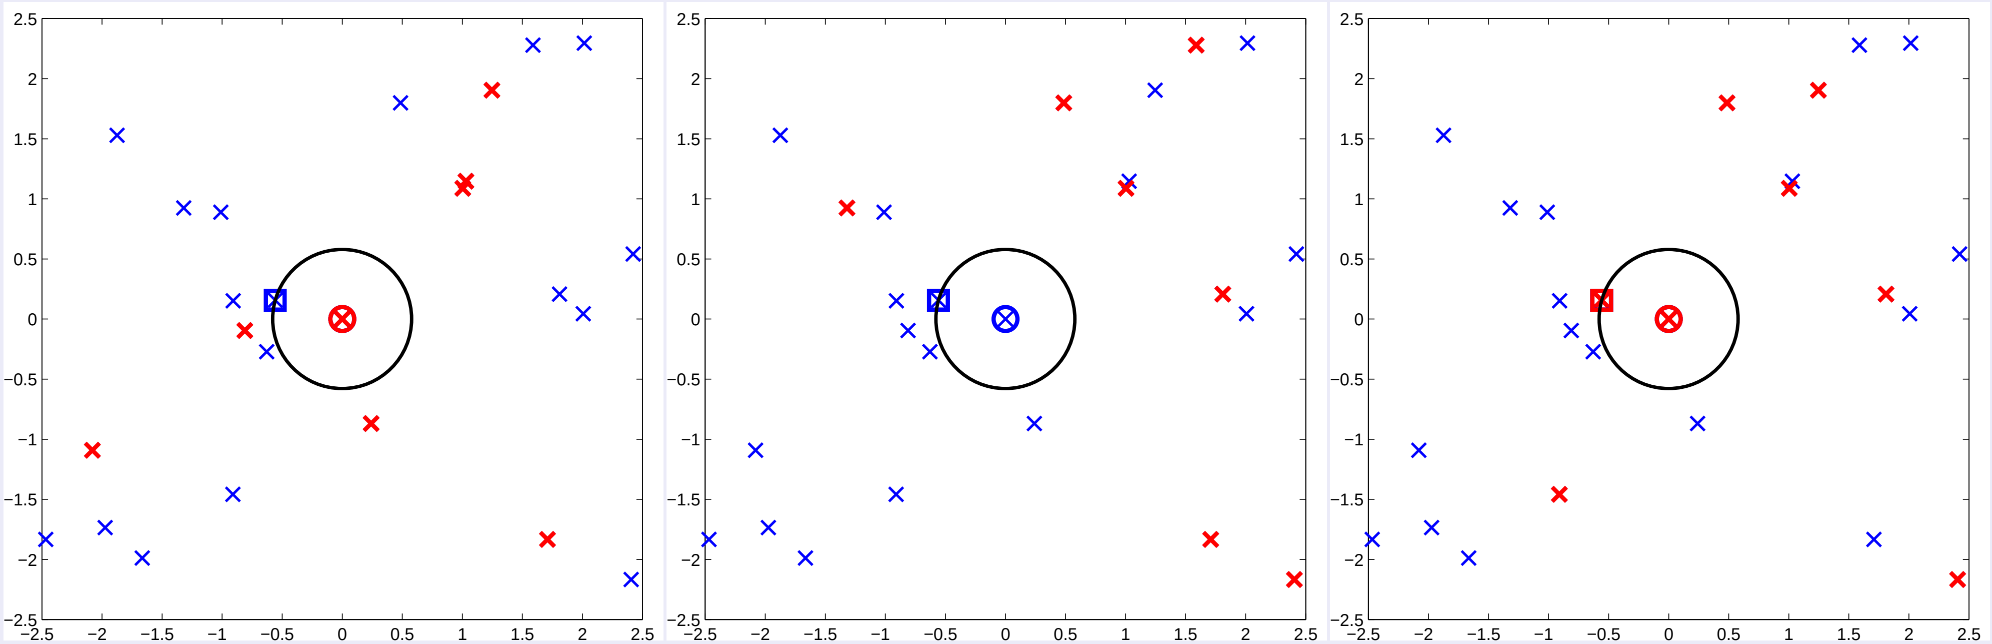

x Transmitters. x Receivers. **o** Source node under consideration.
☐ Destination node under consideration.

The black disk is necessarily free of interferers! This means we need to calculate the conditional interference.

Nearest-neighbor transmission (NNT)

Outage conditioned on nearest-neighbor distance

- Nearest-neighbor distance:
  $f_R(r) = 2\lambda\pi r \exp(-\lambda\pi r^2)$
- Having the nearest neighbor at distance R implies that there is no interferer in the ball $B_o(R)$ centered at o with radius R. So the nearest node sees the conditional interference, conditioned on the disk $B_o(R)$ being empty.
- By stationarity of Φ, the situation is statistically the same if the transmitter is located at $(R, 0)$ and its nearest neighbor at o.

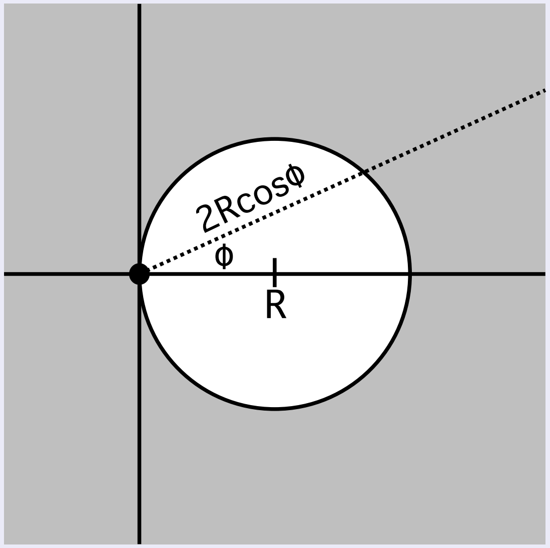

Nearest-receiver transmission (NRT) in a static network

3 time slots:

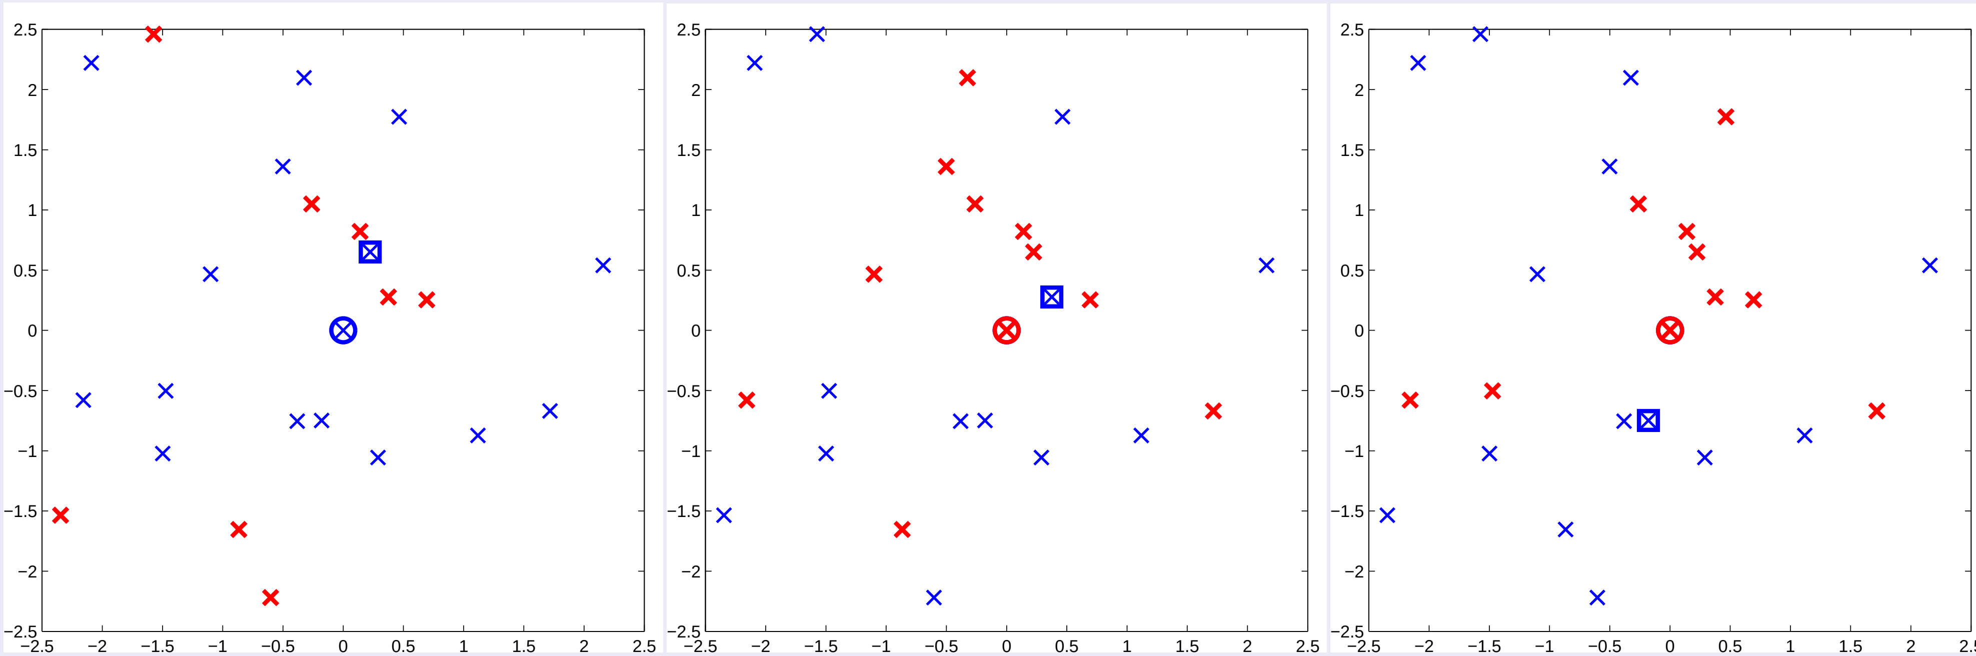

x Transmitters. **x** Receivers. **o** Source node under consideration.
□ Destination node under consideration.

High-Mobility Networks

A lemma

Let $\mathcal{H} \subset \mathbb{R}^2$ and

$$I = \sum_{x \in \Phi} t_x h_x \|x\|^{-\alpha} ,$$

where $t \in \{0, 1\}$ are the Bernoulli transmit marks.
The conditional Laplace transform of I given that $\mathcal{H}$ does not contain any nodes of Φ is $\mathcal{L}_I(s \mid \mathcal{H})$=

$$\mathcal{L}_I(s \mid \mathcal{H} \cap \Phi = \emptyset) = \exp\left(-\lambda p \int_{\mathbb{R}^2 \setminus \mathcal{H}} \frac{s}{s + \|x\|^\alpha} \mathrm{d}x \right) .$$

This follows from the probability generating functional for (non-stationary) PPPs.

Nearest-receiver Transmission in High-Mobility Networks

Outage conditioned on nearest-neighbor distance

If $\mathcal{H} = \emptyset$, the success probability of a transmission $p_s(R)$ between two nodes at distance R is

$$\begin{aligned} p_s(R) &= \mathbb{P}(S > \theta I) = \mathbb{P}(h > \theta R^\alpha I) \\ &= pq\mathbb{E}(e^{-\theta R^\alpha I}) = pq\mathcal{L}_I(\theta R^\alpha) \\ &= pq \exp(-\gamma p \lambda R^2)\,, \end{aligned}$$

where

$$\gamma \triangleq \theta^{2/\alpha} C(\alpha) \quad \text{and} \quad C(\alpha) \triangleq 2\pi^2/(\alpha \sin(2\pi/\alpha))\,.$$

NRT in High-Mobility Networks

Outage and local delay

R is the distance from the origin to the nearest *receiver* in Φ, which is Rayleigh distributed with mean $1/(2\sqrt{q\lambda})$. Hence

$$\begin{aligned} p_s &= p \int \exp(-\gamma p r^2) 2q\lambda\pi r \exp(-q\lambda\pi r^2) \mathrm{d}r \\ &= \frac{p\pi}{\pi + \gamma p q^{-1}} \end{aligned}$$

and

$$D^{\mathrm{NRT}} = \frac{1}{\mathbb{P}(\mathcal{C})} = \frac{1}{p} + \frac{\gamma}{\pi q} \,.$$

The optimum transmit probability is

$$p_{\mathrm{opt}}(\gamma) = \frac{\pi - \sqrt{\pi\gamma}}{\pi - \gamma} \,.$$

NNT in High-Mobility Networks

Outage

Here we can apply the lemma with $\mathcal{H} = B_{(R,0)}(R)$. The left half plane is not affected by the empty disk, so we can write

$$\mathcal{L}_I(s \mid \mathcal{H}) = \exp(-\lambda p C(\alpha) s^{2/\alpha}/2) \exp(-\lambda p A(R,s)),$$

where

$$A(R,s) = \int_{-\pi/2}^{\pi/2} \underbrace{\int_{2R\cos\phi}^{\infty} \frac{rs}{r^\alpha + s} \mathrm{d}r}_{A'(R,s,\phi)} \,\mathrm{d}\phi$$

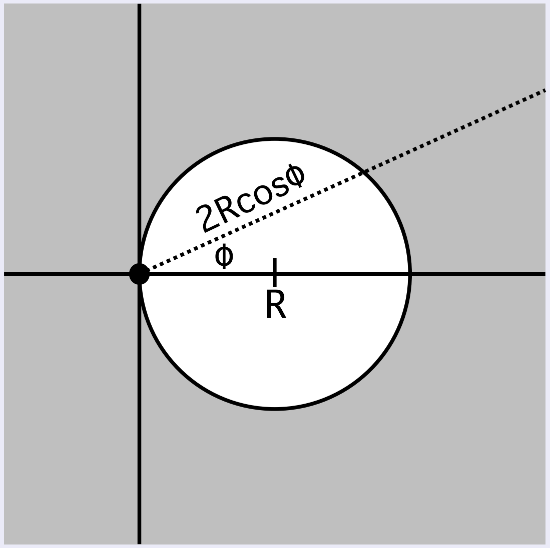

is the integral over the right half plane with the hole, expressed in polar coordinates.

$A(R,s)$ can be well approximated. Then decondition with respect to R.

NNT in High-Mobility Networks

Success probability for $\alpha = 4$

$$\mathbb{P}(\mathcal{C}) \begin{cases} \lessapprox \dfrac{2pq}{\pi p\sqrt{\theta}+2q}\,, & \theta \geqslant 16\,, \\ \approx \dfrac{8pq}{2\pi p\sqrt{\theta}+p\theta^{3/4}(\pi-1)+8}\,, & \theta \leqslant 16\,. \end{cases}$$

Local delay for $\alpha = 4$

$$D^{\text{NNT}} \begin{cases} \gtrapprox \dfrac{\pi\sqrt{\theta}}{2q}+\dfrac{1}{p}\,, & \theta \geqslant 16 \\ \approx \dfrac{1}{pq}+\dfrac{\pi\sqrt{\theta}}{4q}+\dfrac{(\pi-1)\theta^{3/4}}{8q}\,, & \theta \leqslant 16 \end{cases}$$

In the high-rate case ($\theta > 16$), this is the same as for NRT!

Static Networks

Key idea

Transmission success events are conditionally independent given Φ.

Conditioned on Φ, the number of transmissions until success is again geometric with parameter

$$p_s(R \mid \Phi) = \mathcal{L}_I(\theta R^\alpha \mid \Phi) = \mathbb{E}(\exp(-\theta R^\alpha I \mid \Phi))\,.$$

So:

$$D(R) = \mathbb{E}_\Phi\left(\frac{1}{\mathcal{L}_I(\theta R^\alpha \mid \Phi)}\right)\,.$$

The local delay is again obtained by de-conditioning on R: $D = \mathbb{E}_R(D(R))$.

Need to calculate the conditional Laplace transform.

Static Networks

Lemma

Let I denote the interference as defined before, $\mathcal{H} \subset \mathbb{R}^2$, and let

$$\mathcal{L}_I(s \mid \Phi, \mathcal{H}) = \mathbb{E}(\exp(-sI \mid \Phi, \Phi \cap \mathcal{H} = \emptyset))$$

be the conditional Laplace transform given Φ and given that there is no transmitter in $\mathcal{H}$. Then

$$\mathbb{E}\left(\frac{1}{\mathcal{L}_I(s \mid \Phi, \mathcal{H})}\right) = \exp\left(\lambda \int_{\mathbb{R}^2 \setminus \mathcal{H}} \frac{ps}{sq + \|x\|^\alpha} \mathrm{d}x\right) ,$$

which for $\mathcal{H} = \emptyset$ evaluates to

$$= \exp\left(\frac{p\lambda C(\alpha) s^{2/\alpha}}{q^{1-2/\alpha}}\right) .$$

Nearest-receiver transmission

Local delay

$$D^{\mathrm{NRT}} = \frac{1}{p}\frac{\pi}{\pi - \gamma p q^{2/\alpha - 2}}$$

if $pq^{2/\alpha-2} < \pi/\gamma$. Hence there is a phase transition in the local delay. If p is too large, the local delay is infinite.

Bounds

So we have $D_u \geqslant D \geqslant D_l$ for

$$D_u = \frac{1}{p}\frac{\pi}{\pi - \gamma p q^{-2}}\,, \quad \gamma p < q^2\pi$$

$$D_l = \frac{1}{p}\frac{\pi}{\pi - \gamma p q^{-1}}\,, \quad \gamma p < q\pi$$

The optimum transmit probability can be bounded using these bounds.

Summary of Results

Mean delay for nearest-receiver transmission [Hae10c]

High-mobility networks: $$D = \frac{1}{p} + \frac{\gamma}{\pi(1-p)}$$

Static networks: $$D = \frac{1}{p}\frac{\pi}{\pi - \gamma p(1-p)^{2/\alpha - 2}}$$

!

$\gamma = \theta^{2/\alpha} 2\pi^2 / (\alpha \sin(2\pi/\alpha))$ is the spatial contention parameter.

"High mobility" means that a new realization of the PPP is drawn in each time slot.

Comparison ($\alpha = 4$)

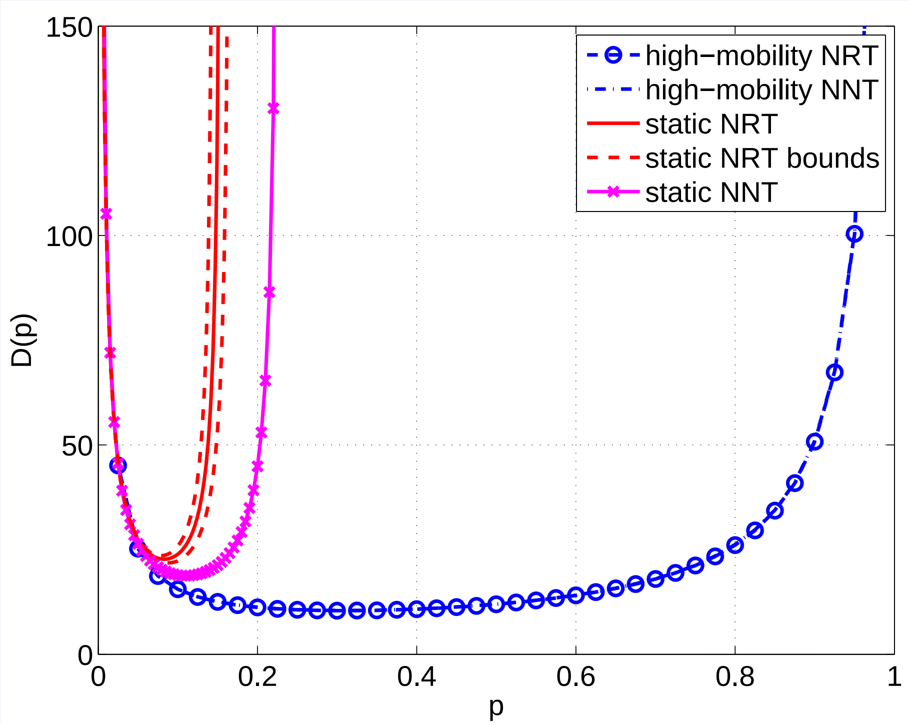

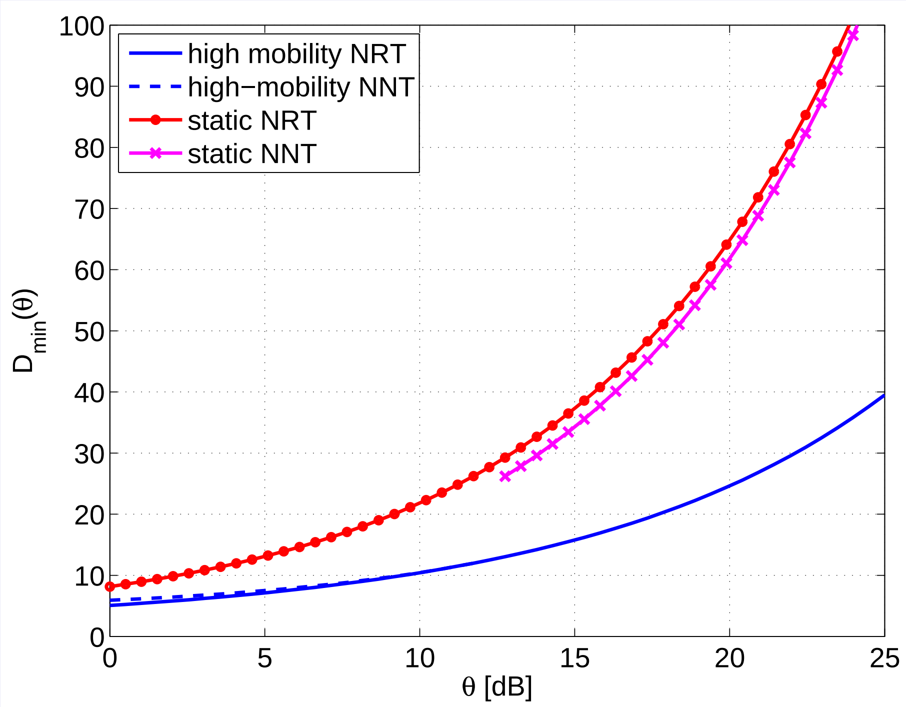

NRT: Nearest-receiver transmission. NNT: Nearest-neighbor transmission.

- In the high-mobility case, the delay is insensitive to p.
- Static networks suffer from a significantly increased delay (due to correlation or lack of diversity). The min. delay is 4 times larger asymptotically.

Local Delay with Noise [Hae10b]

Networks without interference, just noise

This case is non-trivial also. Without interference, the network is a collection of independent links. Over a link of distance R, the received power is

$$P_r = PhR^{-\alpha}\,,$$

where P is the transmit power and h is the (power) fading coefficient. Given R, we have

$$p_{s|R} = \mathbb{P}(h > \theta R^\alpha / P) = 1 - F_h(\theta R^\alpha / P)\,.$$

Fading is assumed iid over time, thus the mean delay of successful transmission, conditioned on R, is $p_{s|R}^{-1}$.
So we have:

iid R: $D = 1/\mathbb{E}_R(p_{s|R})\,;$ static R: $D = \mathbb{E}_R(1/p_{s|R})$ (ensemble avg.) .

The Rayleigh/Rayleigh case

If R is Rayleigh distributed (as a nearest-neighbor link in a Poisson networks), then for constant transmit power P,

$$D = 2\pi\lambda \int_0^\infty \exp(\theta r^\alpha / P) r \exp(-\lambda\pi r^2) \mathrm{d}r$$

diverges to infinity as soon as $\alpha > 2$. Hence power control is necessary for $\alpha > 2$ for finite delay.

Let

$$P \triangleq aR^{\alpha-2+b}\,, \quad \text{for } a, b \geq 0\,.$$

and consider $D(a, b)$. We find

$$D(a, 0) = \frac{\lambda\pi}{\lambda\pi - \theta/a}\,, \quad \theta < a\lambda\pi\,; \qquad D(a, 2) = \exp\left(\frac{\theta}{a}\right)\,.$$

Induced fading

Assume there is no channel fading, but the sender uses randomized power control.
Is it possible to achieve finite local delay for all $\theta > 0$ even with $b = 0$?

Answer: Yes—using a polynomial-tail distribution for power control.

We use the Pareto distribution with complementary distribution

$$\mathbb{P}(H > x) = \left(\frac{k-1}{kx}\right)^k, \quad k > 1,\ x \geq 1 - 1/k\,,$$

parametrized with a single parameter k such that $\mathbb{E}(H) = 1$ for all $k > 1$.

Pareto random power control

The transmit power is chosen to be $P = HR^{\alpha-2+b}$, with H temporally independently Pareto.
It follows that

$$p_s(R) = \begin{cases} \left(\frac{k-1}{k\theta R^{2-b}/a}\right)^k & \text{for } R^{2-b} > \frac{a(k-1)}{\theta k} \\ 1 & \text{otherwise} \end{cases}$$

The local delay with Rayleigh R is minimized for $k = 2$ (heaviest tail).
In this case,

$$D(a, 0) = 1 + (4\xi + 8\xi^2)\exp(-1/(2\xi)),$$

where $\xi \triangleq \theta/(\lambda\pi a)$.

This is finite for all choices of θ and a, and $D(a, 0) = \Theta(\theta^2)$, $\theta \to \infty$!

Section Outline

1 Routing in Poisson Networks

2 Correlation in Poisson Networks

3 Local Delay in Poisson Networks

4 Information Propagation in Poisson Networks
- The SIR multigraph
- Propagation delay
- The delay graph

5 Summary

Information Propagation in Poisson Networks

The SIR multigraph

Let Φ be a PPP on $\mathbb{R}^2$, partitioned at each time $k \in \mathbb{N}$ into a transmitter process $\Phi_t(k)$ and a receiver process $\Phi_r(k)$ by a ALOHA.
Let $\mathbf{1}_k(x \to y) = 1$ if $x \in \Phi_t(k)$ and $y \in \Phi_r(k)$ and the following conditions hold:

- Interference: The disk $b(y, \beta\|x - y\|)$, $\beta > 0$ is free from other transmitters.
- Noise: $\|x - y\| < \eta$.

Otherwise $\mathbf{1}_k(x \to y) = 0$.

These two conditions approximate the condition SINR $> \theta$. They are known as the *protocol model* for communication [GK00].

Definition (The SINR multigraph)

The connectivity at time k is captured by the weighted and directed random geometric graph $G(k) = (\Phi, \vec{E}_k)$ with

$$\vec{E}_k = \{(x, y) \colon \mathbf{1}_k(x \to y) = 1\}\,.$$

A weight k is attached to all these directed edges.
The SINR multigraph $\mathcal{G}$ is the edge-union of these snapshot graphs:

$$\mathcal{G}(0, n) = \left(\Phi, \bigcup_{k=0}^{n} \vec{E}_k\right)$$

$\mathcal{G}(0, n)$ captures all information about the network from time 0 to time n.

Definition (Causal path)

A causal path is a directed path $\{x_0, \vec{e}_0, x_1, \vec{e}_1, \ldots, \vec{e}_{q-1}, x_q\}$ with *strictly increasing edge weights*.

Example (SIR multigraph)

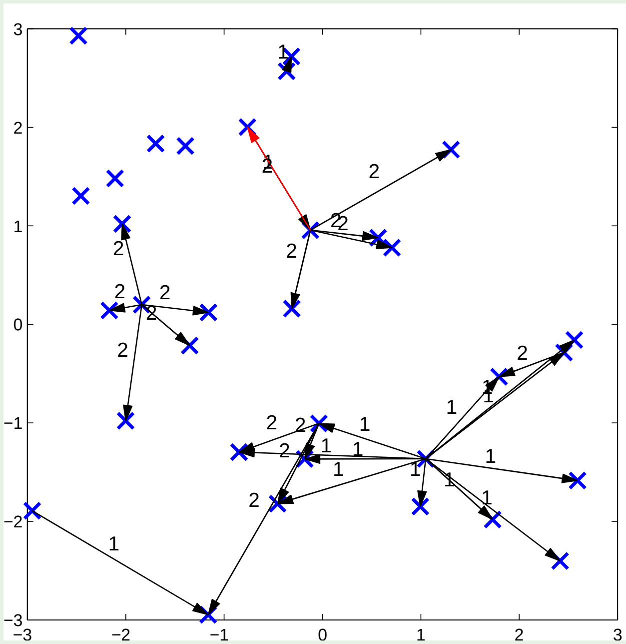

Time slots $1, 2$

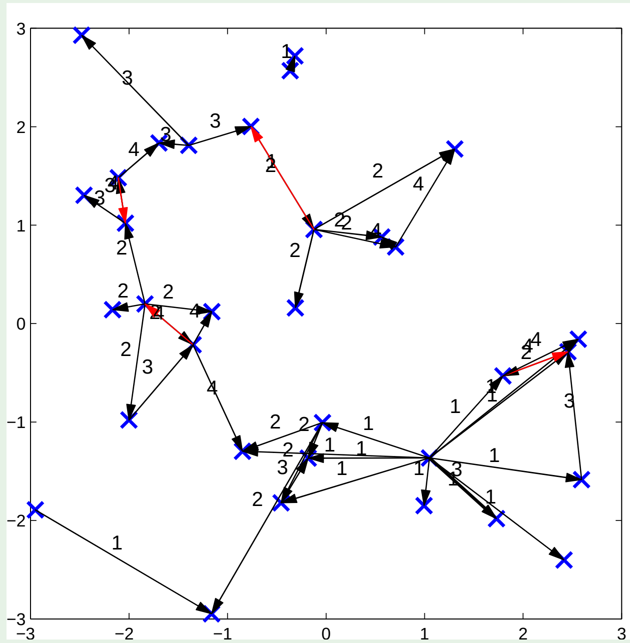

Time slots $1-4$

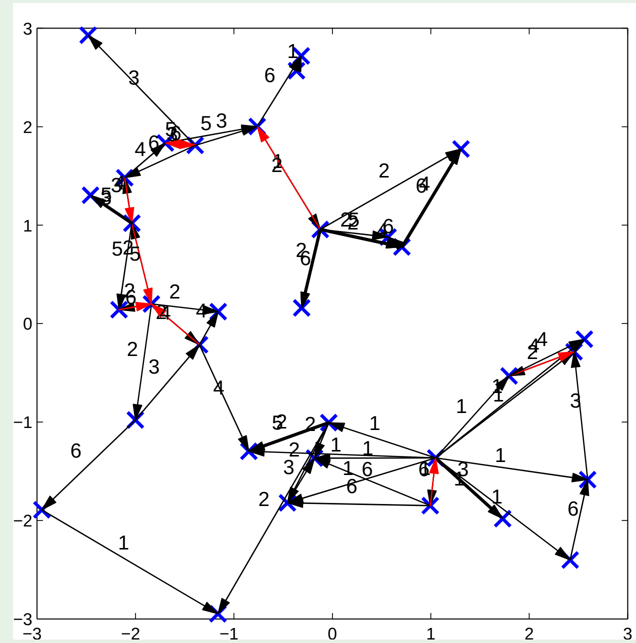

Time slots $1-6$

Example (SIR multigraph after 10 and 20 time slots)

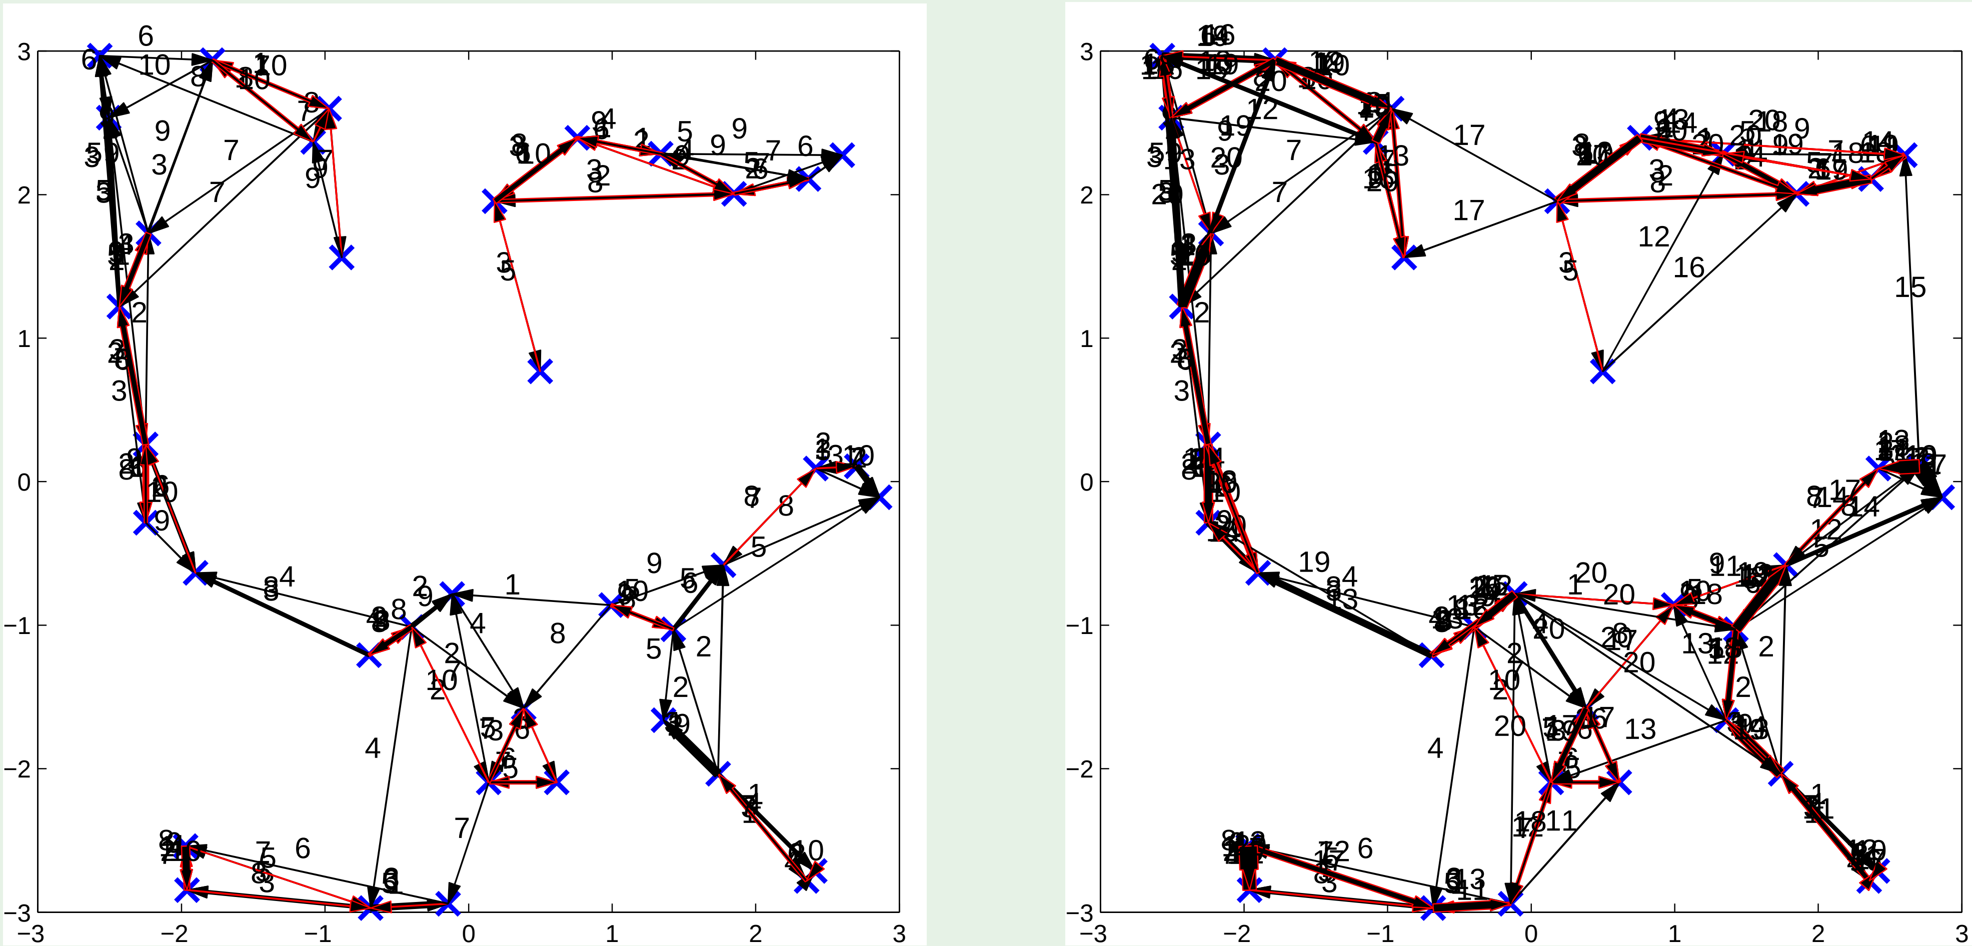

Edge thickness increases with multiplicity, and red edges indicate bidirectionality.

Information propagates along causal paths in this graph. Causal here means edge weights are strictly increasing.

Propagation Delay [GH10]

Single-hop delay

Add a node at the origin, and let

$$T_O = \min\{k\colon \sum_{x\in\Phi} \mathbf{1}_k(o\to x) > 0\}$$

be the number of time slots to connect to any node.

If $\eta < \infty$, $\mathbb{E}(T_O) = \infty$, since the origin is too far from any node with probability $\exp(-\lambda\pi\eta^2)$.
So focus on the interference-limited regime where $\eta = \infty$. In this case, $\mathbb{E}(T_O)$ is finite and can be lower bounded.

Path formation time

The path formation time from x to y is

$$T(x,y) = \min\{k\colon \mathcal{G}(0,k) \text{ has a path from } x \text{ to } y\}\,.$$

Similarly, define

$$T_n(x,y) = \min_{k>n}\{k-n\colon \mathcal{G}(n,k) \text{ has a path from } x \text{ to } y\}\,.$$

Since T is *sub-additive*, for $0 < p < 1$, *i.e.*,
$T(o,y) \leq T(o,x) + T_{T(o,x)}(x,y)$, the propagation time constant

$$\mu = \lim_{x\to\infty} \frac{\mathbb{E}T(o,x)}{x}$$

is finite. It is infinite with noise, but if the disk graph $G_{\lambda,\eta}$ percolates and we only consider nodes in the infinite component, again μ is finite.

Numerical results

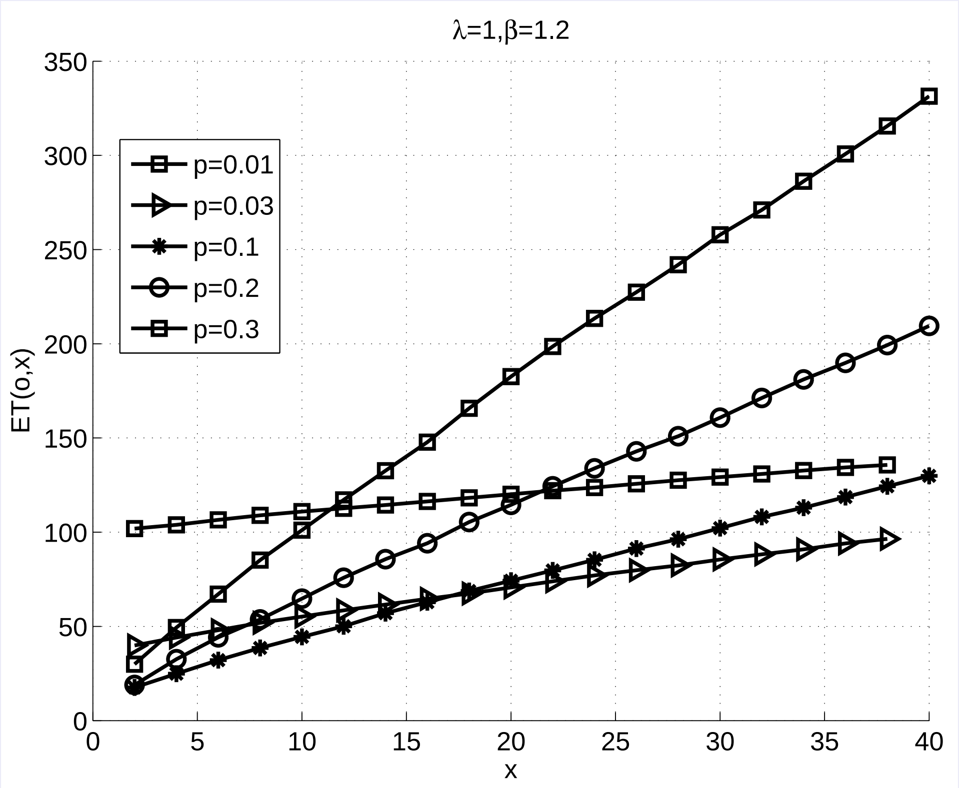

Mean delay from o to x

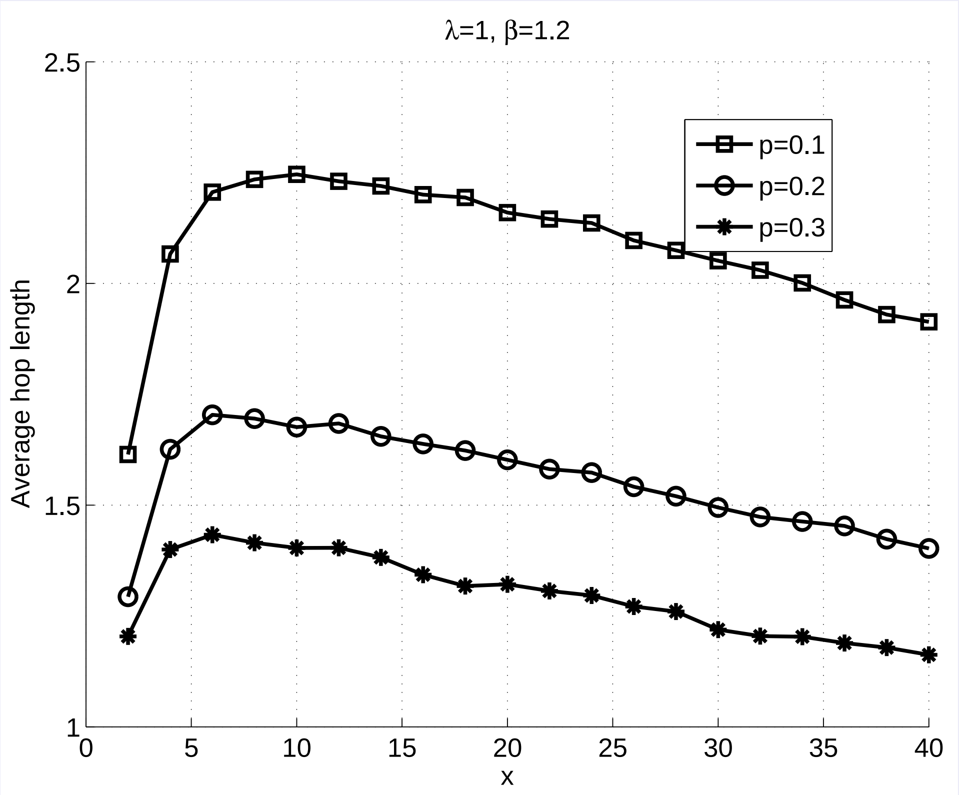

Mean hop lengths of fast routes

Depending on the distance, a different p is optimum. The mean hop length of fast routes increases with $\|x\|$.

This framework allows *reverse engineering* to find good routing protocols.

The Delay Graph [Hae10a]

Definition (Single-hop delay)

Given a point process Φ and a MAC scheme that, at any given moment k, partitions the network into transmitters and receivers, the single-hop delay $D\colon \Phi^2 \to \mathbb{R}$ is defined as

$$D(x,y) \triangleq \mathbb{E}\left[\min_{k\in\mathbb{N}} \mathbf{1}_k(x \to y)\right]$$

where the expectation is taken with respect to the MAC scheme and the fading.

Definition (Delay graph $\mathcal{G}_\tau$)

The delay graph is the random geometric digraph $\mathcal{G}_\tau = (\Phi, \vec{E}_\tau)$, where $(x,y) \in \vec{E}_\tau$ if $D(x,y) \le \tau$.

Properties of the delay graph

- A source and a destination node are connected by a directed edge if the source can be expected to reach the destination (in a single hop) in at most τ time slots.
- The delay graph is related to the SIR multigraph as follows: In the delay graph, the edge $\overrightarrow{xy}$ is present if the expected smallest edge weight in the SIR graph is at most τ.
- ALOHA: For $\tau < 4$, all nodes are isolated, while for $\tau \to \infty$, the graph is fully connected. So the connectivity exhibits a *phase transition* with respect to τ, in the sense that there exists a finite critical value τ_c, such that G_τ a.s. has an infinite out-component for $\tau > \tau_c$ (*i.e.*, there is a node from which an infinite number of nodes can be reached).
- The delay graph, averaged over Φ, can used to determine the delay-minimizing hop length.

Delay-optimum number of hops for ALOHA

Let

$$h(n) \triangleq \frac{n(n+1)\sqrt{\log(1+1/n)}}{\sqrt{b(2n+1)}}, \quad n > 0,$$

where

$$b = \frac{p\lambda\pi\Gamma(1+2/\alpha)\Gamma(1-2/\alpha)\theta^{2/\alpha}}{q^{1-2/\alpha}}.$$

This function yields the distance $R_n = h(n)$ for which $n\tilde{D}(R_n/n) = (n+1)\tilde{D}(R_n/(n+1))$. So at distance R_n, the delay for n hops is the same as the delay for $n+1$ hops, hence for smaller distances, n hops is better than $n+1$. It follows that

$$n \text{ hops is optimum} \iff h(n-1) < R \leq h(n).$$

As $n \to \infty$, $h(n) \sim n/\sqrt{2b}$, so to cover a distance R, the optimum number of hops $n_{\mathrm{opt}} \approx \lceil R\sqrt{2b} \rceil$.

Example (Delay graphs for constant $p = 0.25$)

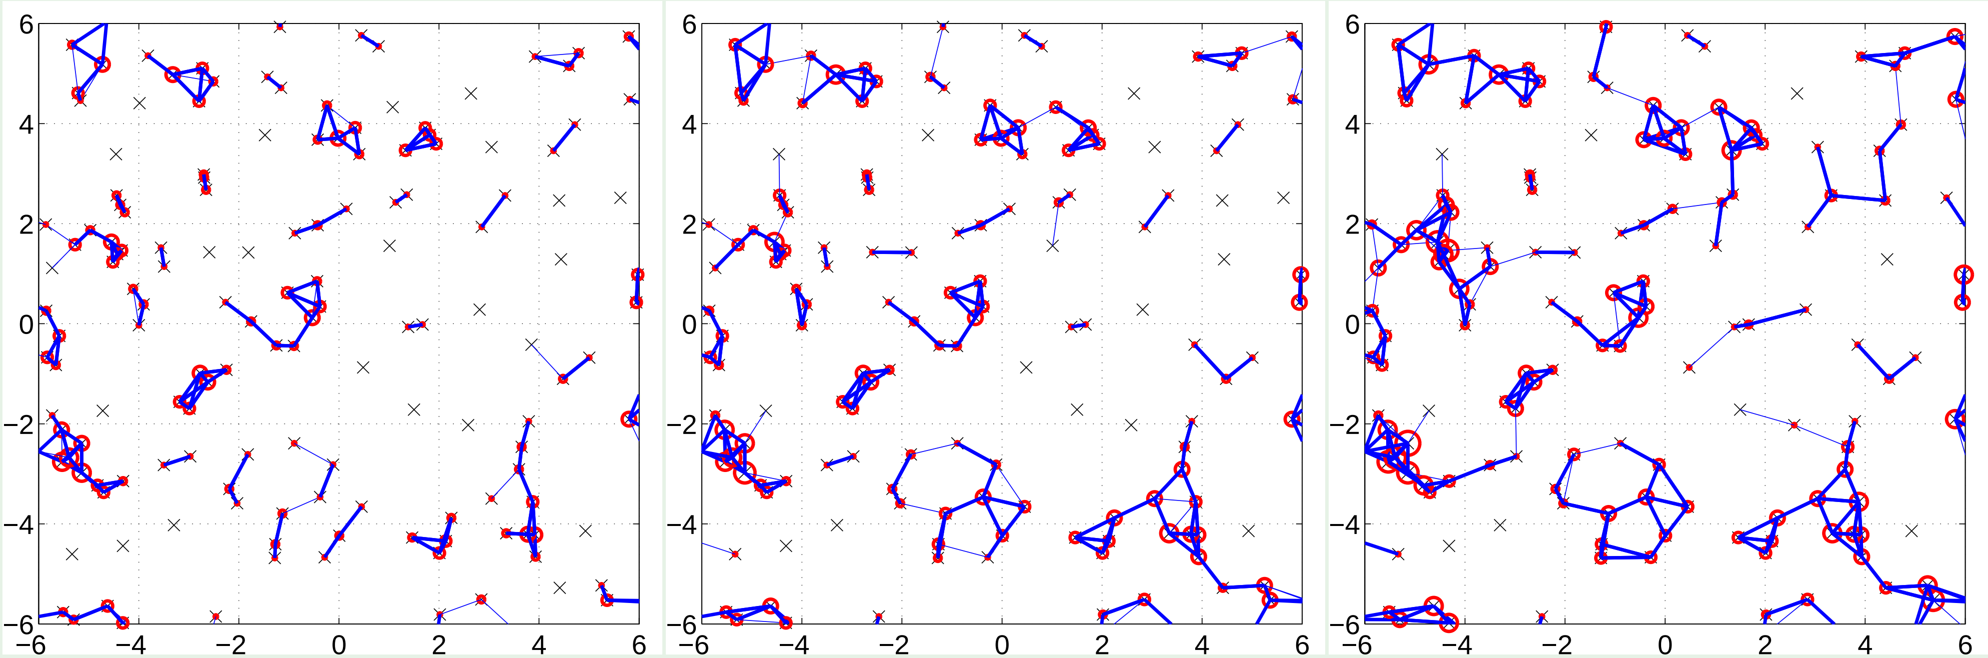

Delay graphs for $\tau = 50, 100, 200$ for $\lambda = 1$, transmit probability $p = 0.25$, path loss exponent $\alpha = 4$, and SIR threshold $\theta = 10$. Bidirectional edges are bold.

Mean out-degrees: 1.81, 2.28, and 2.89.

Example (Delay graphs for constant $p = 0.05$)

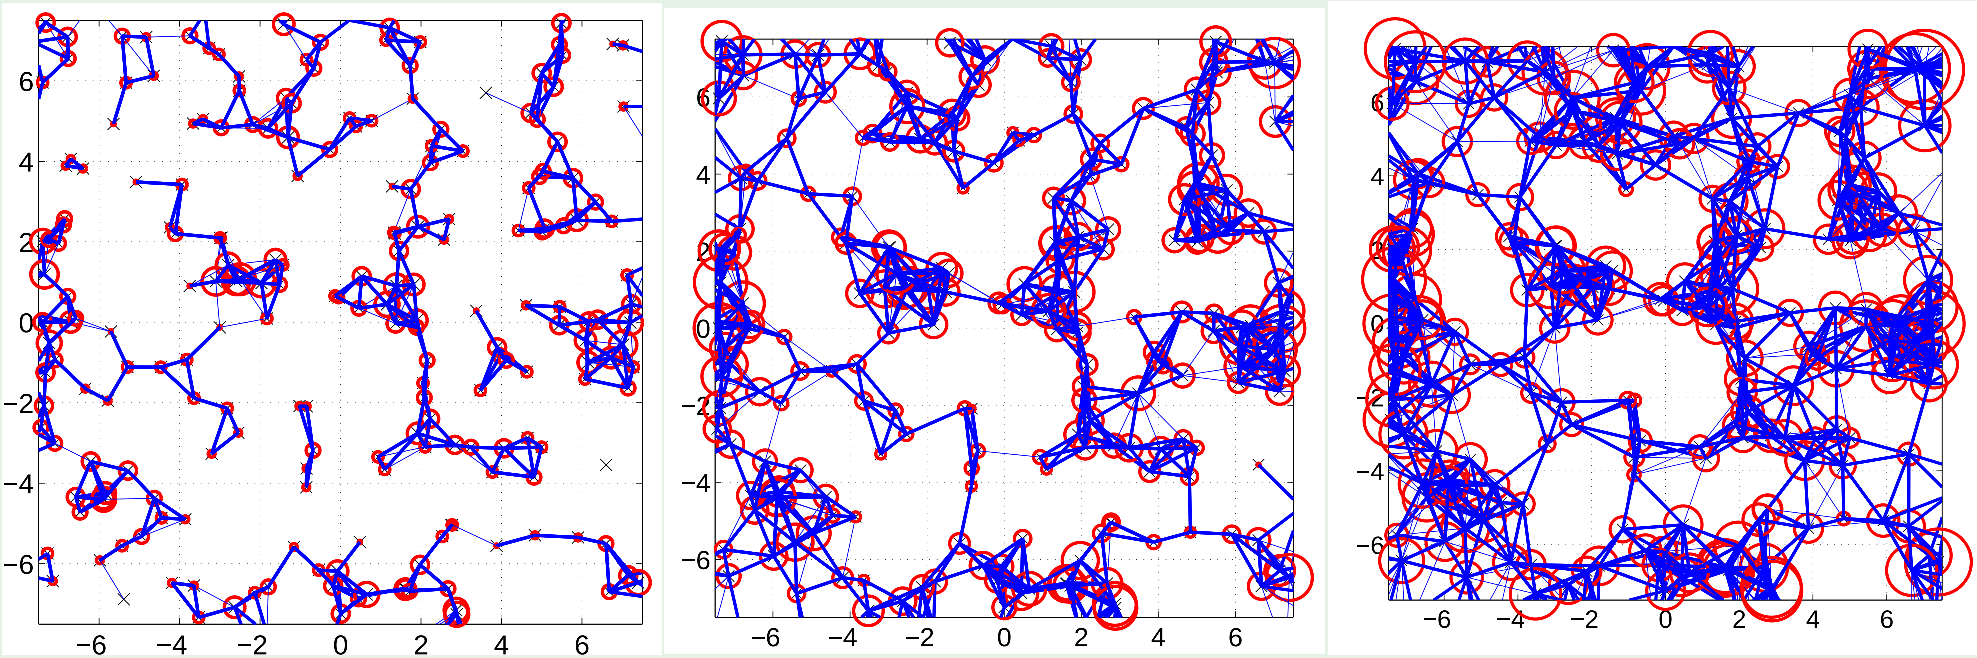

Delay graphs for $\tau = 50, 100, 200$ for $\lambda = 1$, transmit probability $p = 0.05$, path loss exponent $\alpha = 4$, and SIR threshold $\theta = 10$. The radius of the circles at each node is proportional to the node degree.
Mean out-degrees: 3.85, 7.45, and 12.5.

Example (Variable p)

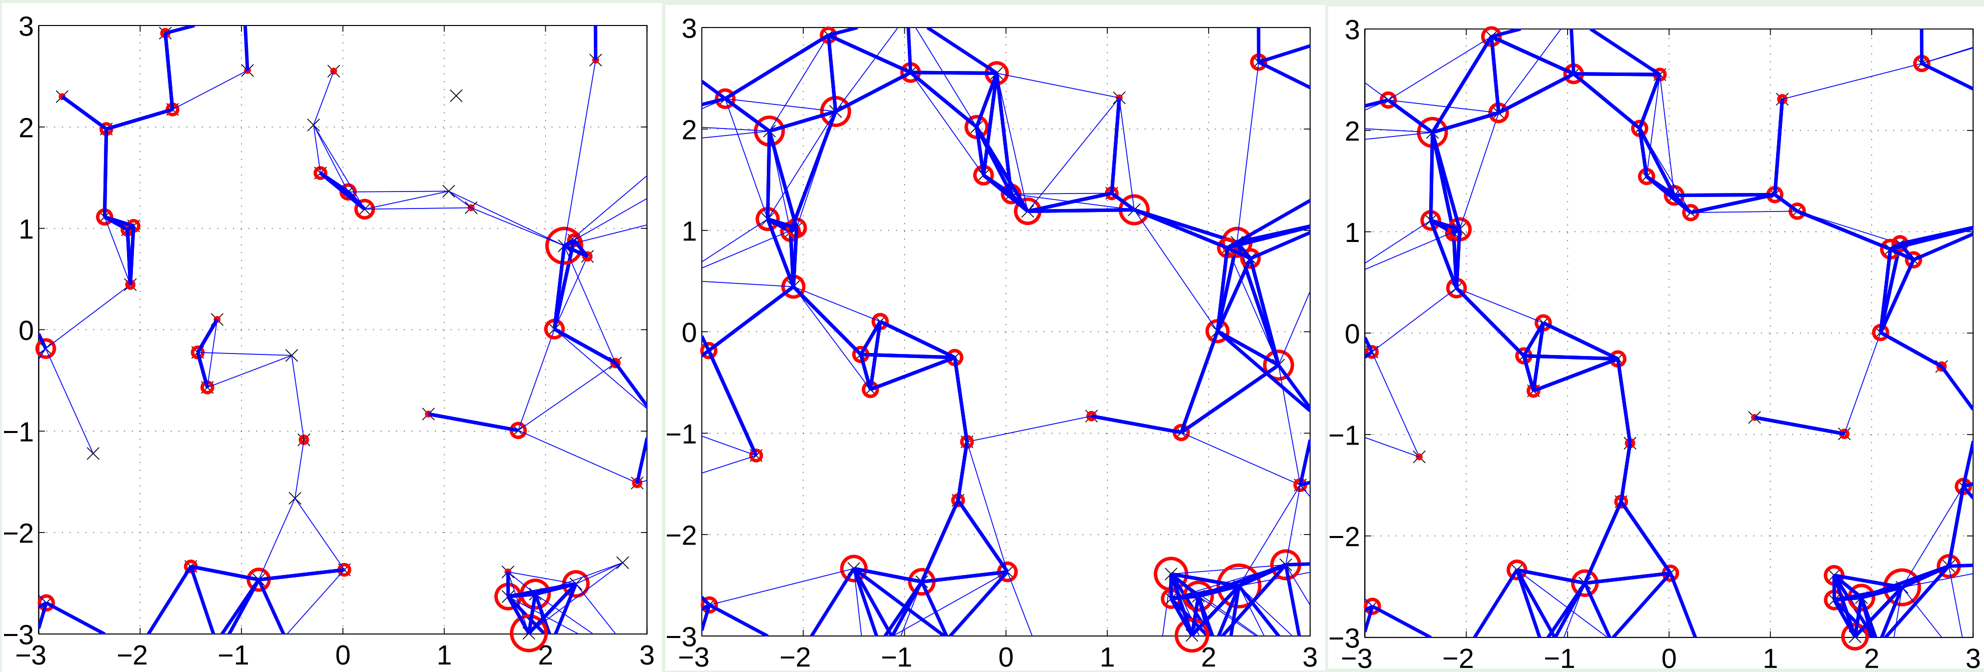

$p = (1/2)p_{\mathrm{opt}}$ $\qquad$ $p = p_{\mathrm{opt}}$ $\qquad$ $p = 2p_{\mathrm{opt}}$

Delay graphs for $\tau = 50$ for $\lambda = 1$, path loss exponent $\alpha = 4$, SIR threshold $\theta = 10$, with $p_{\mathrm{opt}} = e/\tau \approx 0.0544$

Section Outline

1 Routing in Poisson Networks

2 Correlation in Poisson Networks

3 Local Delay in Poisson Networks

4 Information Propagation in Poisson Networks

5 Summary

Lecture 5 Summary

Multi-hop analysis

- Multihop extensions and end-to-end analyses are possible in the Poisson case. They are based on the single-hop success probabilities.
- A common assumption is that transmission success events are spatially and temporally independent.
- A calculation of the correlation coefficient shows that this assumption is reasonable for small transmit probabilities and with fading. For larger p, there is enough dependence in static networks that the local delay becomes infinite.
- Another common assumption is that nodes are always backlogged. A more careful analysis considers queues and the fact that nodes do not transmit if they do not have packets.

Multi-hop analysis

- The SINR multigraph captures the dynamic connectivity of the network. Using methods from first-passage percolation, the propagation speed of a prioritized packet can be bounded. A simpler version is the delay graph, which still captures the effects of interference.

References I

J. G. Andrews, S. Weber, M. Kountouris, and M. Haenggi.
Random Access Transport Capacity.
IEEE Transactions on Wireless Communications, 9(6):2101–2111, June 2010.
Available at http://www.nd.edu/~mhaenggi/pubs/twc10.pdf.

R. K. Ganti and M. Haenggi.
Spatial and Temporal Correlation of the Interference in ALOHA Ad Hoc Networks.
IEEE Communications Letters, 13(9):631–633, September 2009.
Available at http://www.nd.edu/~mhaenggi/pubs/commletter09.pdf.

R. K. Ganti and M. Haenggi.
Dynamic Connectivity in ALOHA Ad Hoc Networks.
IEEE Transactions on Information Theory, 2010.
Submitted. Available at http://www.nd.edu/~mhaenggi/pubs/tit10a.pdf.

Piyush Gupta and P. R. Kumar.
The Capacity of Wireless Networks.
IEEE Transactions on Information Theory, 46(2):388–404, March 2000.

M. Haenggi.
Delay-Based Connectivity of Wireless Networks.
In *International Symposium on Mathematical Theory of Networks and Systems (MTNS'10)*, Budapest, Hungary, August 2010.
Available at http://www.nd.edu/~mhaenggi/pubs/mtns10.pdf.

M. Haenggi.
Local Delay in Poisson Networks with and without Interference.
In *2010 Allerton Conference on Communication, Control and Computing*, Monticello, IL, September 2010.
Available at http://www.nd.edu/~mhaenggi/pubs/allerton10a.pdf.

References II

M. Haenggi.
Local Delay in Static and Highly Mobile Poisson Networks with ALOHA.
In *2010 IEEE International Conference on Communications (ICC'10)*, Cape Town, South Africa, May 2010.
Available at http://www.nd.edu/~mhaenggi/pubs/icc10a.pdf .

K. Stamatiou and M. Haenggi.
Delay-Minimizing Routing for Random Wireless Networks.
IEEE/ACM Transactions on Networking, 2010.
In preparation.

Analysis and Design of Wireless Networks

Lecture 6: Analysis of General Networks

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop Analysis of Poisson Networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 6 Overview

1. General Point Processes
2. Palm Theory
3. Analysis of Poisson Cluster Processes
4. Outage Probability in General Networks
5. Summary

Section Outline

1. General Point Processes
   - Attraction and repulsion
   - Examples
   - Motion-invariant point processes
   - Formal definition
2. Palm Theory
3. Analysis of Poisson Cluster Processes
4. Outage Probability in General Networks
5. Summary

General Point Processes

Simon Denis Poisson, 1781-1840.

Is the analytical treatment of wireless networks restricted to his model?

Motivation

The transmitter process is only a PPP if the process of all nodes is PPP and ALOHA is used as the MAC protocol. In all other cases, the Poisson model is at best an approximation.

Two typical cases:

- Nodes form a PPP, but a CSMA-type MAC is used. This leads to repulsion among the transmitters, *i.e.*, a more regular process.
- Nodes form a cluster process. This usually leads to attraction between the transmitters as well, *i.e.*, a more clustered process.

From regular to clustered processes

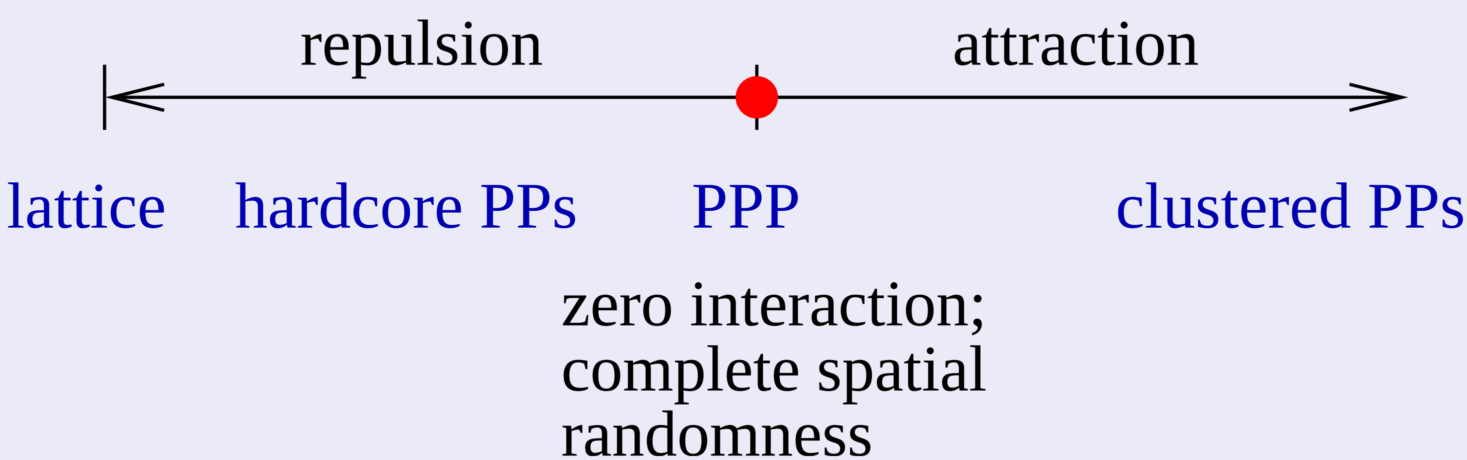

Examples

Example (Matern Hard-Core Process of Type I)

Take a homogeneous PPP of intensity λ_p and remove all points that are within distance r of each other. The resulting process has intensity $\lambda = \lambda_p \exp(-\lambda_p \pi r^2)$.

Remarks:

- Imposes a minimum distance r.
- Process is stationary.
- Obtained by dependent thinning.
- Repulsion or inhibition.
- Possible model for CSMA.

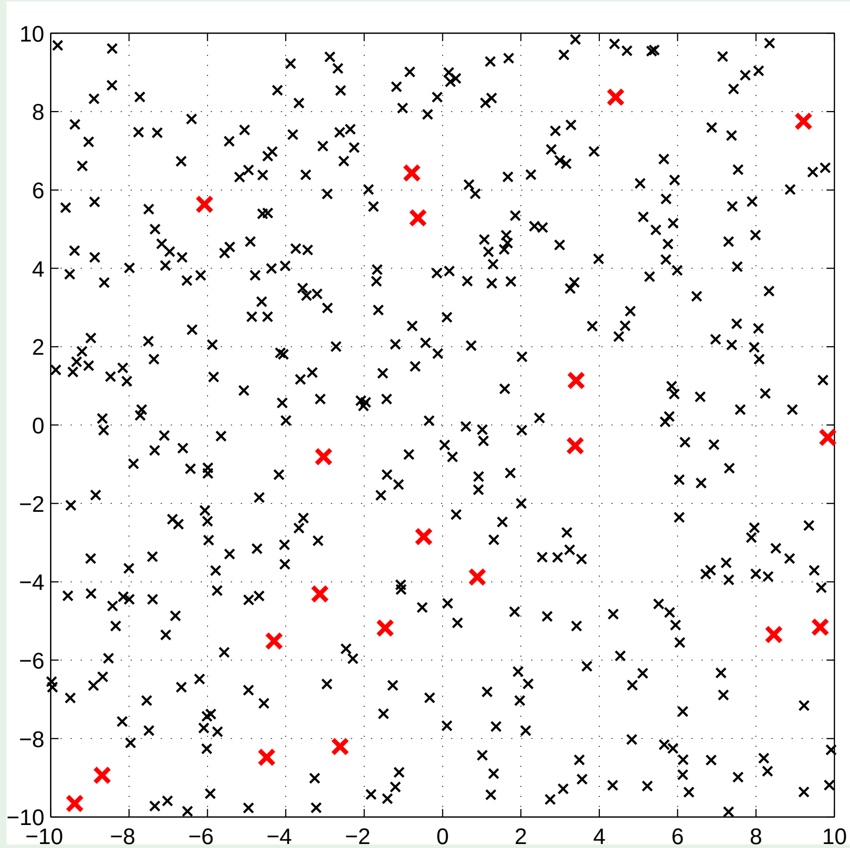

$\lambda_p = 1$, $r = 1$. Red points are retained.

Example (Neyman-Scott Cluster Processes)

Poisson cluster processes resulting from homogeneous independent clustering applied to a PPP.

- Parent points $\Phi_p = \{x_1, x_2, \ldots\}$ form a PPP of intensity λ_p.
- Clusters N^x are of the form $N^{x_i} = N_i + x_i$ for each $x_i \in \Phi_p$.
- The N_i are a family of iid finite point sets with distribution $F(x)$ independent of the parent process. The complete process is given by

$$\Phi = \bigcup_{x \in \Phi_p} N^x \,.$$

- The intensity of the cluster process is $\lambda = \lambda_p \bar{c}$, where $\bar{c}$ is the average number of points per cluster.
- If the number of points per cluster is Poisson, the process is called a *Poisson cluster process*.

Example (Special Neyman-Scott processes: Matern and Thomas cluster processes)

Matern cluster process (parameters λ_p, $\bar{c}$, and a):
Daughter points are iid *uniformly* distributed in a ball of radius a around the parent.
Thomas cluster process (parameters λ_p, $\bar{c}$, and σ):
Daughter points are iid *symmetrically normally* distributed with variance σ^2 around the parent, *i.e.*, each child cluster forms an inhomogeneous PPP with intensity

$$\lambda(x) = \frac{\bar{c}}{2\pi\sigma^2} \exp(-\|x\|^2/2\sigma^2)\,, \quad \text{(2-dim.)}$$

so that the mean number of children per parent is $\bar{c}$.
Useful to model tactical networks (soldiers, troops, platoons, ...), human cocktail parties, networks with closely cooperating nodes (virtual MIMO).

Comparison of Thomas cluster process and PPP on $[-5,5]^2$:

Thomas process

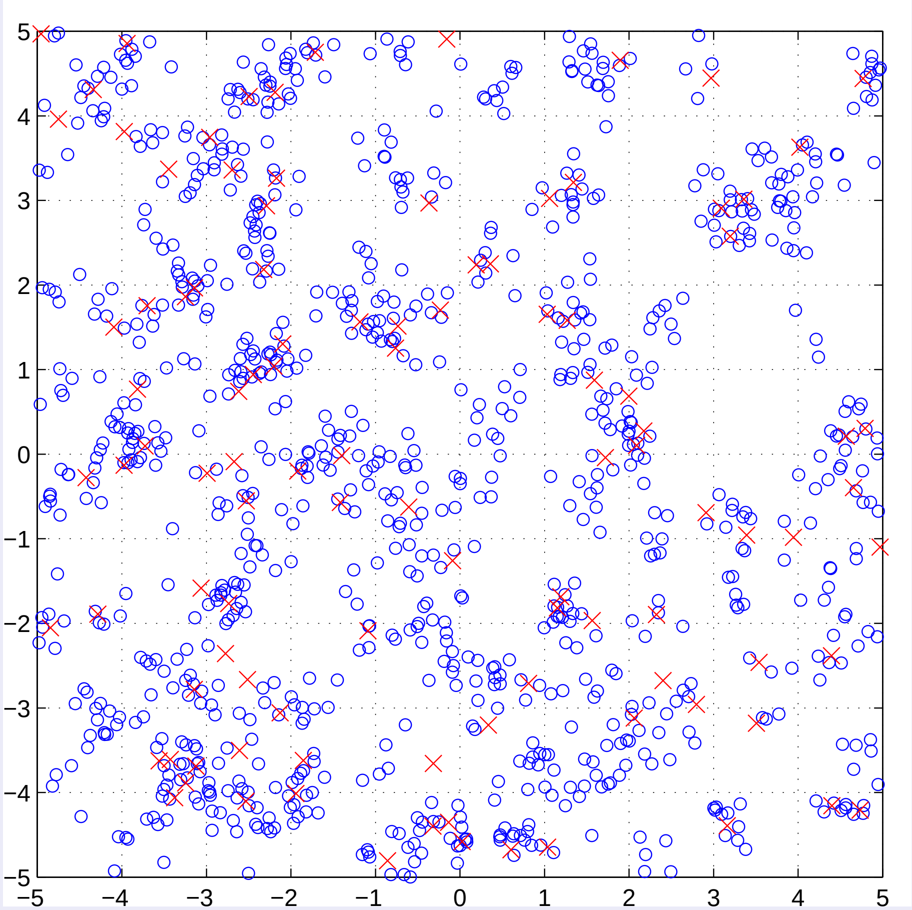

$\lambda_p = 1, \bar{c} = 5$, and $\sigma = 0.2$

PPP

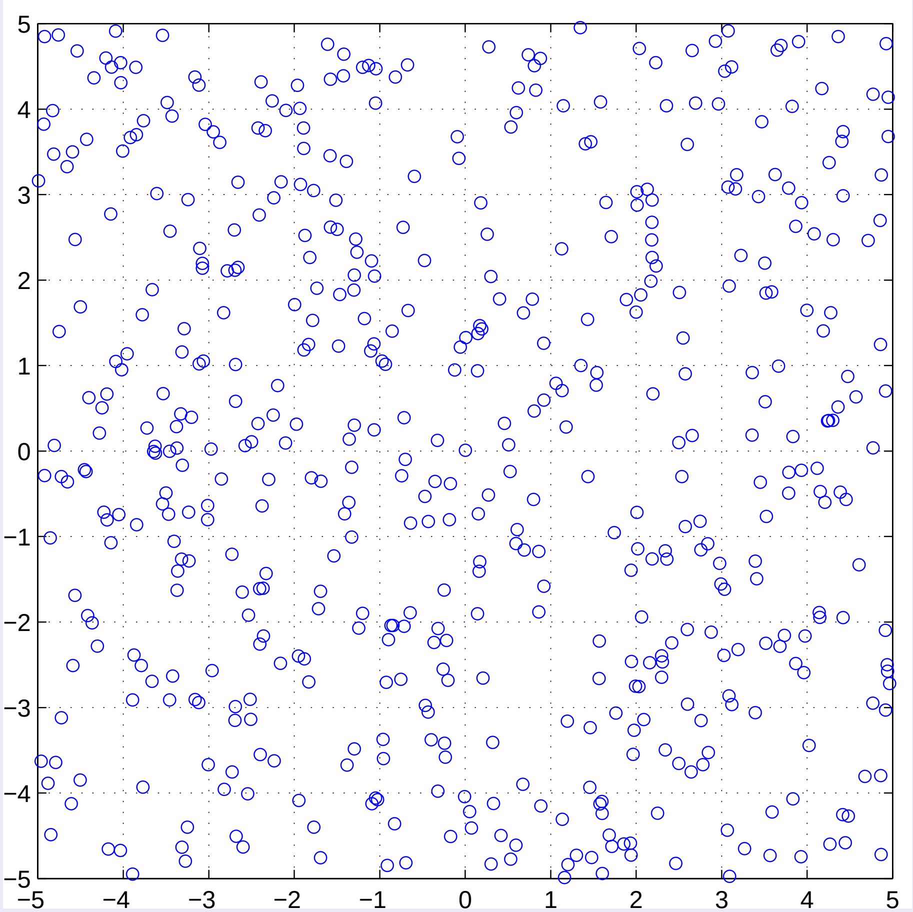

$\lambda = 5.$

Comparison of Thomas cluster process and PPP on $[-5,5]^2$:

Almost a Thomas process

$\lambda, \bar{c}, \sigma = ?$

PPP

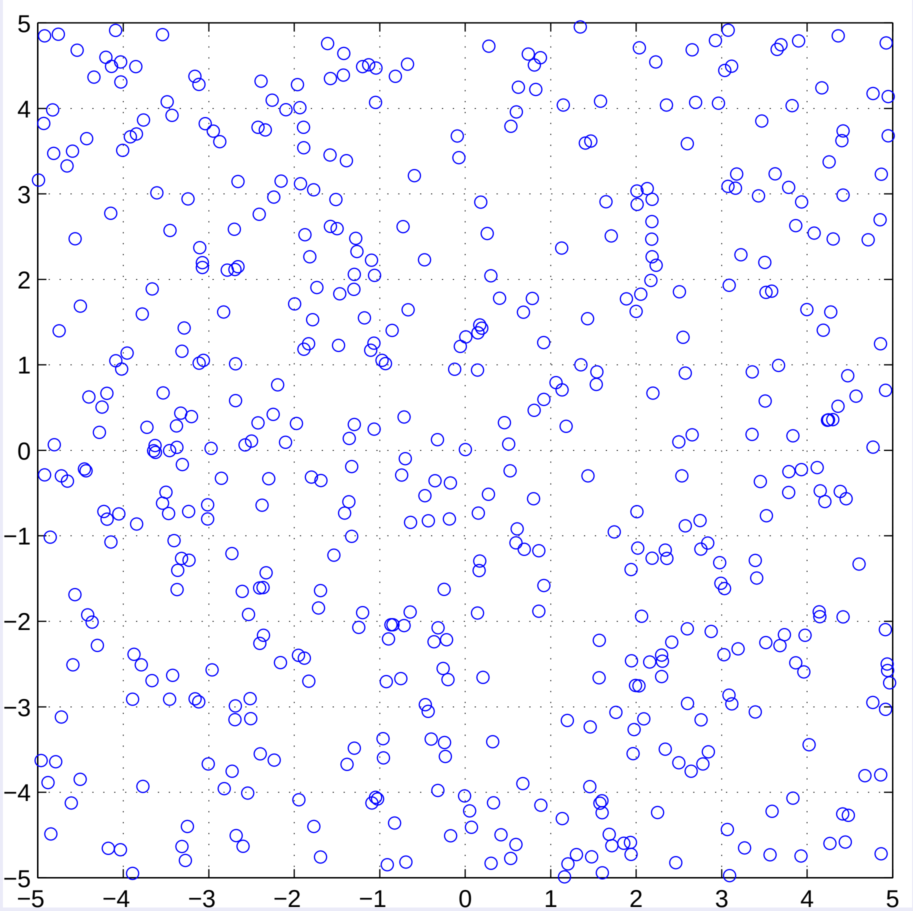

$\lambda = 5.$

Example (Poisson hole process)

Take a homogeneous PPP Ψ of intensity λ_p and a second PPP Φ of intensity 1. Remove all points in Φ that are within distance r of any point in Ψ. The resulting Poisson hole process has intensity $\lambda = \exp(-\lambda_p \pi r^2)$.

- Process is stationary.
- Obtained by dependent thinning.
- Possible model for cognitive networks. Ψ are the primary and Φ the secondary users. The secondary users who are allowed to transmit form the hole process.

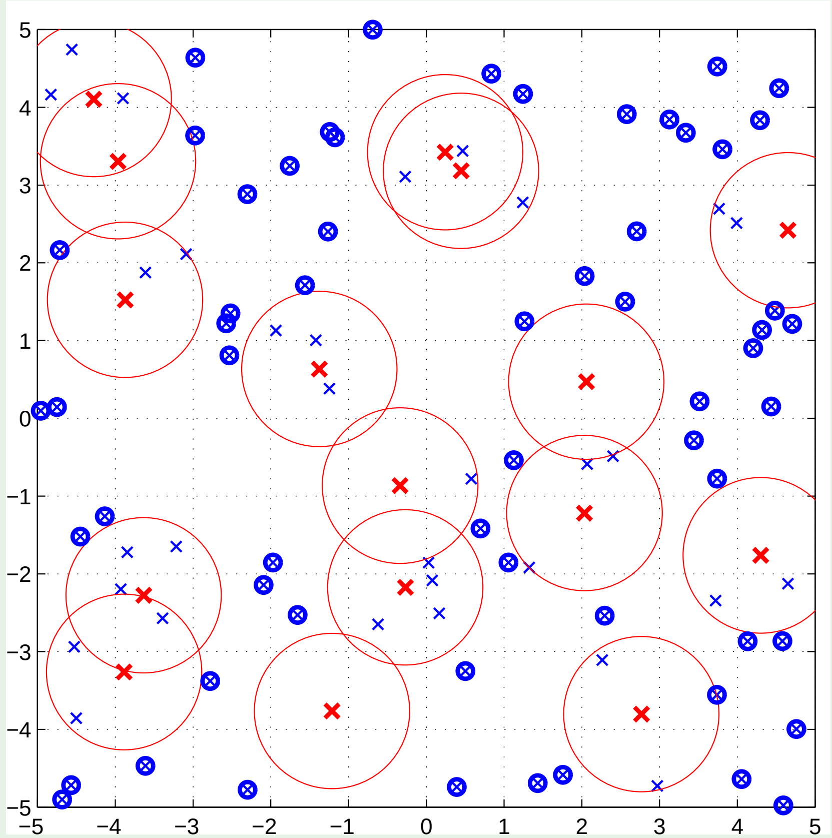

$\lambda_p = 0.2$, $r = 1$. Blue points **o** form the hole process.

Motion-invariant Point Processes

Some definitions

- Stationarity: $\{x_i\}$ and $\{x_i + x\}$ have the same distribution $\forall x \in \mathbb{R}^d$.
- Isotropy: The same holds for all rotations about the origin.
- Motion-invariance: Stationarity plus isotropy. The stationary (homogeneous) PPP is motion-invariant.

All point processes we consider are motion-invariant.

Palm theory

The analysis of all non-Poisson point processes requires proper conditioning on the process having a node somewhere, typically at the origin. Such events have probability 0, so some care is required. This is the topic of Palm theory.

Formal Definition

Definition (Point process)

N. A point process Φ on $\mathbb{R}^d$ is a random element taking values in a measurable space $(N, \mathcal{N})$, where N is the family of all sequences ϕ of points of $\mathbb{R}^d$ satisfying two regularity conditions:

1. The sequence ϕ is *locally finite*.
2. The sequence is *simple*, *i.e.*, $x_i \neq x_j$ if $i \neq j \; \forall i, j \in \phi$.

σ-algebra $\mathcal{N}$: Smallest σ-algebra on N such that $\phi(A)$ measurable for all bounded Borel A.

So $\Phi \colon (\Omega, \mathcal{A}, \mathbb{P}) \mapsto (N, \mathcal{N})$, and the distribution of Φ is

$$P(Y) = \mathbb{P} \circ \Phi^{-1}(Y) = \mathbb{P}(\Phi \in Y), \quad Y \in N.$$

Intuition

Intuitively, a point process Φ is a random choice of one of the ϕ in N.

Two points of view

The space of outcomes N can be interpreted as

- The family of simple and countably finite point sets.
- The family of random *counting measures* counting the number of points in $B \subset \mathbb{R}^d$.

Accordingly, we can write

- $\Phi = \{x_1, x_2, \ldots\} = \{x_i\}$.
- $\Phi(B) = |\{\Phi \cap B\}|$.

Intensity measure

$$\Lambda(A) = \mathbb{E}(\Phi(A)) = \int_N \phi(A) P(\mathrm{d}\phi) \qquad \text{for Borel } A\,,$$

where P is the *distribution* of the point process Φ.
If Φ is stationary, $\Lambda(A) = \lambda|A|$, *i.e.*, λ is the ratio of the intensity measure to the Lebesgue measure.

Comparison with numerical random variables

	Numerical random variable	Point process
Probability space	$(\Omega, \mathcal{A}, \mathbb{P})$	$(\Omega, \mathcal{A}, \mathbb{P})$
Measurable space	$(\mathbb{R}, \mathcal{B})$	$(N, \mathcal{N})$
Random element	$X \in \mathbb{R}$	$\Phi \in N$
Events	$B \in \mathcal{B}$	$Y \in \mathcal{N}$
Distribution	$P(B) = \mathbb{P} \circ X^{-1}(B)$	$P(Y) = \mathbb{P} \circ \Phi^{-1}(Y)$
Measurability	$X^{-1}(B) = \{\omega \in \Omega \colon X(\omega) \in B\} \in \mathcal{A}$	$\Phi^{-1}(Y) = \{\omega \in \Omega \colon N^{\omega} \in Y\} \in \mathcal{A}$
Measure space	$(\mathbb{R}, \mathcal{B}, P)$	$(N, \mathcal{N}, P)$
Distr. function	$F(x) = P((-\infty, x])$	

Section Outline

1. General Point Processes

2. Palm Theory
   - Motivation
   - Palm distribution
   - Second-order statistics

3. Analysis of Poisson Cluster Processes

4. Outage Probability in General Networks

5. Summary

Palm Theory

Motivation

Take a motion-invariant point process Φ_t *of transmitters*. Assume we are interested in interference. Where do we measure? We could take an arbitrary point $z \in \mathbb{R}^2$:

$$I(z) = \sum_{x \in \Phi_t} h_x g(\|x - z\|) \,.$$

Since Φ_t is stationary, the distribution of $I(z)$ does not depend on z. And $\mathbb{E}(I(z))$ is the same for all point processes with the same intensity (Campbell's theorem).

However, we are not interested in an arbitrary location, but either

- in a point of a process $\Phi \supset \Phi_t$, or
- a point near a transmitter, which would be the desired transmitter and thus does not contribute to the interference.

Motivation: Matern hard core process

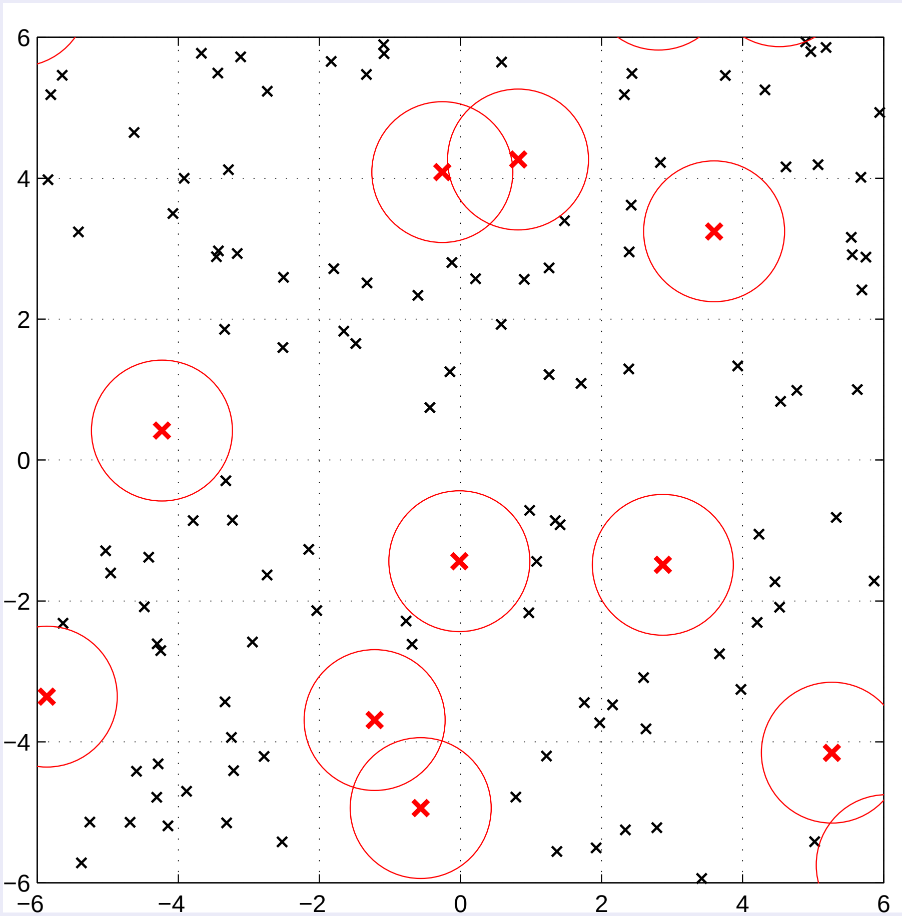

The mean interference measured at a point x is smaller than at an arbitrary point since there are no nodes nearby.

But measuring at $x \in \Phi$ means conditioning that Φ has a point at x.

Since Φ is stationary, we can always condition on $o \in \Phi$.

Motivation: Cluster process

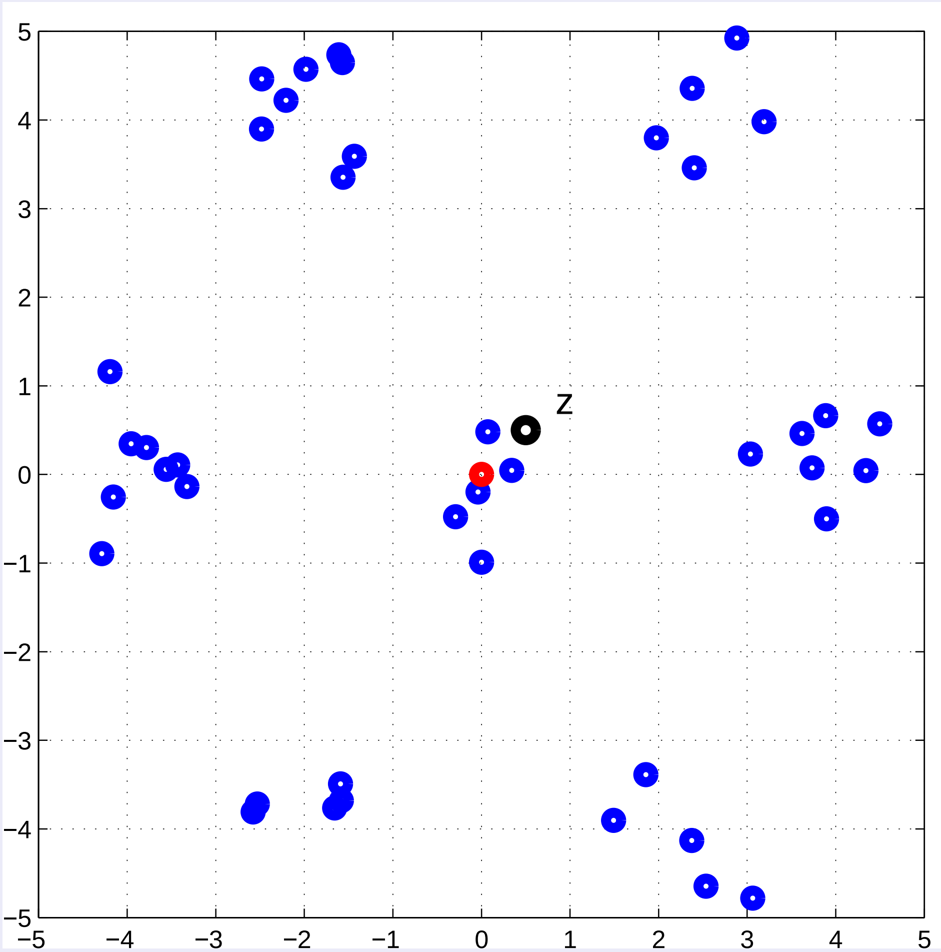

Conditioned on $o \in \Phi$, this means that there is a cluster near the origin.

Measuring $I(z)$ for small $\|z\|$ means that the interference is likely to be high, as there are more transmitters close.

The node at the origin could be the desired transmitter, so it does not contribute to the interference.

Conditioning on a node at o but disregarding its impact yields the reduced Palm distribution.

Palm Distribution

Definition

Let $Y \in \mathcal{N}$ be a point processes event, *i.e.*, a certain PP property, such as having no point in $b(o, r)$. $\Phi \in Y$ means that Φ has this property.

The Palm distribution $\mathbb{P}_o$ is defined as

$$\mathbb{P}_o(\Phi \in Y) \triangleq \mathbb{P}(\Phi \in Y \parallel o)\,,$$

and the reduced Palm distribution $\mathbb{P}^!_o$ is

$$\mathbb{P}^!_o(\Phi \in Y) \triangleq \mathbb{P}(\Phi \setminus \{o\} \in Y \parallel o)\,.$$

The corresponding expectations are $\mathbb{E}_o$ and $\mathbb{E}^!_o$.

Slivnyak's theorem

For a PPP: $\quad \mathbb{P}^!_o \equiv \mathbb{P} \quad \Rightarrow \quad \mathbb{E}^!_o \equiv \mathbb{E}$

Second-order Statistics

Reduced second moment measure

The first-order statistic of a stationary point process is its intensity λ. The second moment measure plays a role similar to the variance.

The reduced second moment measure $\lambda\mathcal{K}_2(B)$ is the mean number of points in $B \setminus \{o\}$ given that $o \in \Phi$: $\lambda\mathcal{K}_2(B) = \mathbb{E}^!_o\Phi(B)$.

There is a corresponding density, the second-order product density $\varrho^{(2)}$:

$$\lambda\mathcal{K}_2(B) = \frac{1}{\lambda}\int_B \varrho^{(2)}(x)\mathrm{d}x$$

$\varrho^{(2)}$ measures the probability that there are two points separated by x; it is the density pertaining to the second-order factorial moment measure:

$$\alpha^{(2)}(A \times B) = \mathbb{E}\left(\sum_{x,y\in\Phi}^{\neq} \mathbf{1}_A(x)\mathbf{1}_B(y)\right) = \int_A\int_B \varrho^{(2)}(x-y)\mathrm{d}y\mathrm{d}x$$

Second-order factorial moment measure

- The name *factorial moment measure* comes from the fact that

$$\alpha^{(2)}(A \times A) = \mathbb{E}(\Phi(A)^2) - \mathbb{E}(\Phi(A)) = \mathbb{E}(\Phi(A)(\Phi(A) - 1))\,.$$

- If Φ is motion-invariant, then $\varrho^{(2)}(x)$ depends only on $\|x\|$.
- For the uniform PPP, $\varrho^{(2)}(x) \equiv \lambda^2$, and

$$\alpha^{(2)}(A \times B) = \lambda^2 |A|\,|B|\,.$$

Section Outline

1 General Point Processes

2 Palm Theory

3 Analysis of Poisson Cluster Processes

- Interference in Poisson cluster processes
- Outage in Poisson cluster processes

4 Outage Probability in General Networks

5 Summary

Interference in Poisson Cluster Processes

Main idea.

Use Ripley's K-function to estimate the interference.

Definition (Reduced second moment function ("Ripley's K-function"))

$$\lambda K(r) \triangleq \int_N \phi(b(o,r)) P^!_o(\mathrm{d}\phi)$$

$$K(r) \triangleq \lambda^{-1}\mathbb{E}[\text{number of extra points within distance } r \text{ of a randomly chosen point}]$$

For PPP, $K(r) = \pi r^2$. For small r:

- For regular processes, $K(r)$ is smaller than πr^2.
- For clustered processes, $K(r)$ is larger than πr^2.

Asymptotically $K(r)$ approaches πr^2 for all motion-invariant point processes.

The L-Function

$$L(r) = \sqrt{\frac{K(r)}{\pi}}$$

To quantify regularity (deviation from CSR), the derivative $L'(r)$ can be used.
$L'(r) \equiv 1$ for the stationary PPP.
For small r:
- $L'(r) < 1 \Rightarrow$ regular process.
- $L'(r) > 1 \Rightarrow$ clustered process.

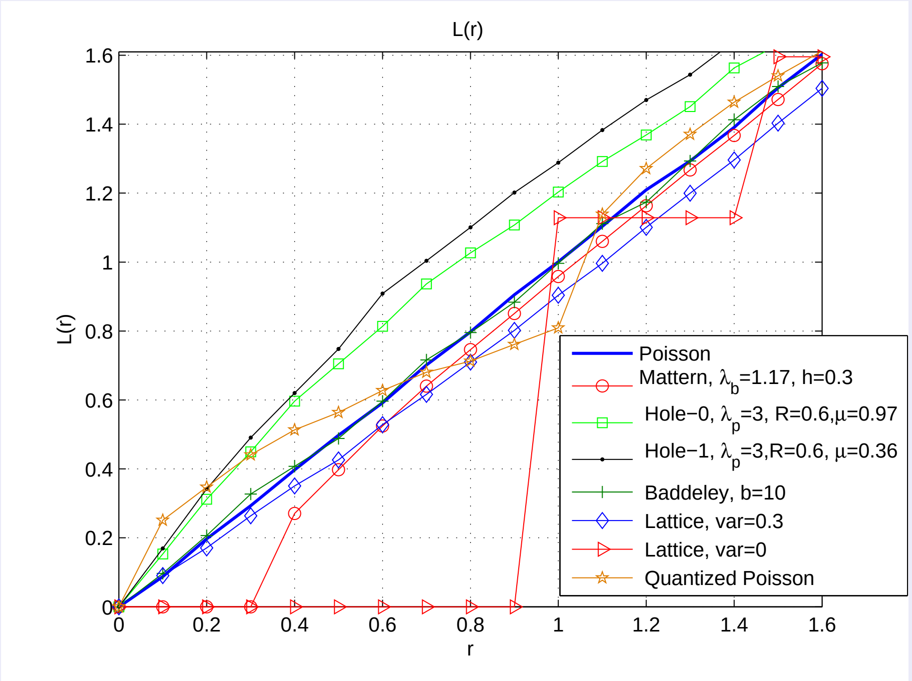

Can be used to estimate interference. We expect that the interference decreases with increasing regularity.

The K function and interference

Let Φ be a motion-invariant point process on $\mathbb{R}^2$. Conditioned on $o \in \Phi$ (desired transmitter), the interference at point z is given by

$$I_o^!(z) = \sum_{x \in \Phi \| o \setminus \{o\}} h_x g(x - z) \,.$$

The distribution of $I_o^!$ will depend on $\|z\|$ only, since Palm distributions of motion-invariant PPs are not stationary but isotropic.
Let $\mathcal{K}_n(\mathbf{B})$ denote the reduced n-th factorial moment measure of Φ.
$\mathbf{B} = \mathbf{B}_1 \times \ldots \times \mathbf{B}_{n-1}$, $\mathbf{B}_i \in \mathbb{R}^2$.

$$\mathcal{K}_n(\mathbf{B}) = \int_N \sum_{x_1,\ldots,x_{n-1} \in \phi}^{\neq} \mathbf{1}_{\mathbf{B}}(x_1, \ldots, x_{n-1}) P_o^!(\mathrm{d}\phi) \,.$$

We can determine $\mathbb{E}_o^! I(z)$ and $\mathbb{E}_o^! I^2(z)$ using $\mathcal{K}_2$ and $\mathcal{K}_3$. Let us focus on the mean.

Mean interference

We have

$$\mathbb{E}^!_o I(z) = \mathbb{E}^!_o \Big[\sum_{x \in \Phi} h_x g(x - z) \Big]$$

$$= \mathbb{E}[h] \lambda \int_{\mathbb{R}^2} g(x - z) \mathcal{K}_2(\mathrm{d}x)$$

$$= \mathbb{E}[h] \lambda \int_{\mathbb{R}^2} g(x - z) L'(\|x\|) \mathrm{d}x \,,$$

where $L'(r) = \mathrm{d}L(r)/\mathrm{d}r$.
Since the process is stationary, $\mathcal{K}_2(B)$ can be expressed as

$$\mathcal{K}_2(B) = \frac{1}{\lambda^2} \int_B \varrho^{(2)}(x) \mathrm{d}x \,,$$

where $\varrho(x)$ is the second order product density.
We expect $\mathbb{E}^!_o I(z)$ to increase with decreasing regularity for small $\|z\|$.

Mean interference for Thomas cluster process [GH09]

It is known that [SKM95]: $$\frac{\varrho^{(2)}(x)}{\lambda^2} = 1 + \frac{1}{4\pi\lambda_p\sigma^2}\exp\left(\frac{-\|x\|^2}{4\sigma^2}\right)$$

where $\lambda = \lambda_p\bar{c}$. We obtain

$$\begin{aligned}\mathbb{E}_o^! I(z) &= \mathbb{E}(I_{\mathrm{PPP}}) + \frac{\bar{c}}{4\pi\sigma^2}\int_{\mathbb{R}^2} g(x-z)\exp\left(\frac{-\|x\|^2}{4\sigma^2}\right)\mathrm{d}x \\ &= \mathbb{E}(I_{\mathrm{PPP}}) + \bar{c}\,\mathbb{E}g(X-z)\,,\end{aligned}$$

where X is a 2D Gaussian with variance $\sqrt{2}\sigma$ in both directions—cf. Slide 9.
Note: $g(x)$ must be bounded as $\|x\| \to 0$ for this to be finite.

As expected, the interference is larger in the clustered case. The difference gets smaller as $\|z\| \to \infty$.

Outage in Poisson Cluster Processes

Sketch of derivation

- Assume the receiver is not part of the PP, but the desired transmitter is.
- Assuming Rayleigh fading, the success probability is given by the conditional Laplace transform of the interference.
- Express the conditional LT using the conditional generating functional.
- Derive the conditional generating functional using the refined Campbell theorem and Slivnyak's theorem.

Outage for clustered processes [GH09]

$$p_s = \exp\left\{ -\lambda_p \int_{\mathbb{R}^2} \Big[1 - \exp(-\bar{c}\beta(R,y))\Big] \mathrm{d}y \right\} \times \int_{\mathbb{R}^2} \exp(-\bar{c}\beta(R,y)) f(y) \mathrm{d}y$$

where

$$\beta(R,y) = \int_{\mathbb{R}^2} \frac{g(x-y-R)}{\frac{g(R)}{\theta} + g(x-y-R)} f(x) \mathrm{d}x \,.$$

where $f(x) = \lambda_N(x)/\bar{c}$ is the pdf of the points in each cluster.

To emphasize isotropy, we (ab)use $R = (R, 0) \in \mathbb{R}^2$.

The first term does not depend on the cluster distribution, while the second does not depend on the overall intensity.

Comparison PPP vs. Matern cluster process ($a = 0.6$)

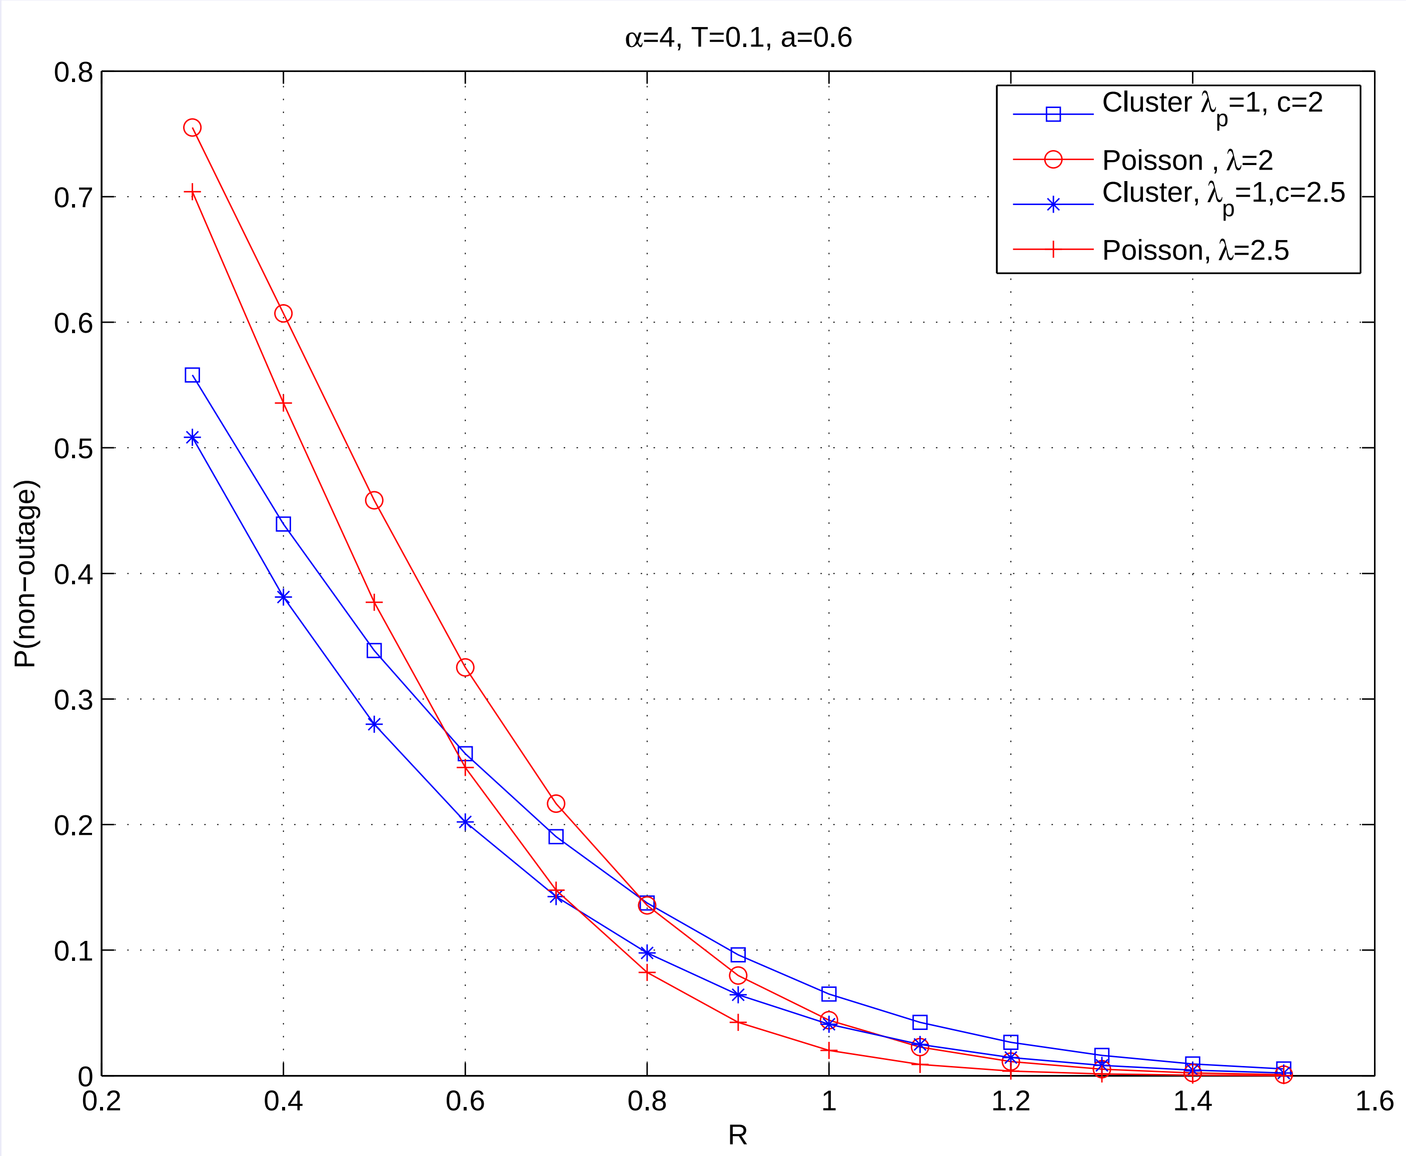

There is an R^* where the curves cross!

Section Outline

1. General Point Processes
2. Palm Theory
3. Analysis of Poisson Cluster Processes
4. Outage Probability in General Networks
   - Problem formulation
   - Outage in the high-SIR regime
   - Extension to general fading statistics
5. Summary

Outage Probability in General Networks

Setup (1)

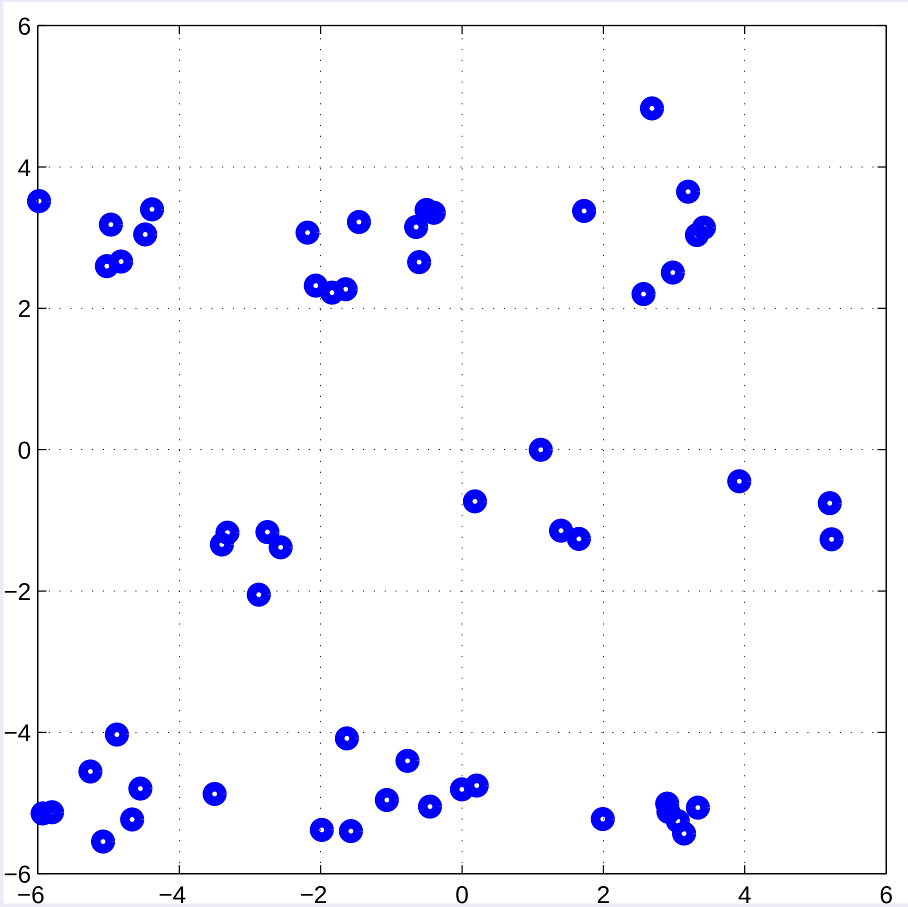

Start with a motion-invariant point process of density λ.

Setup (2)

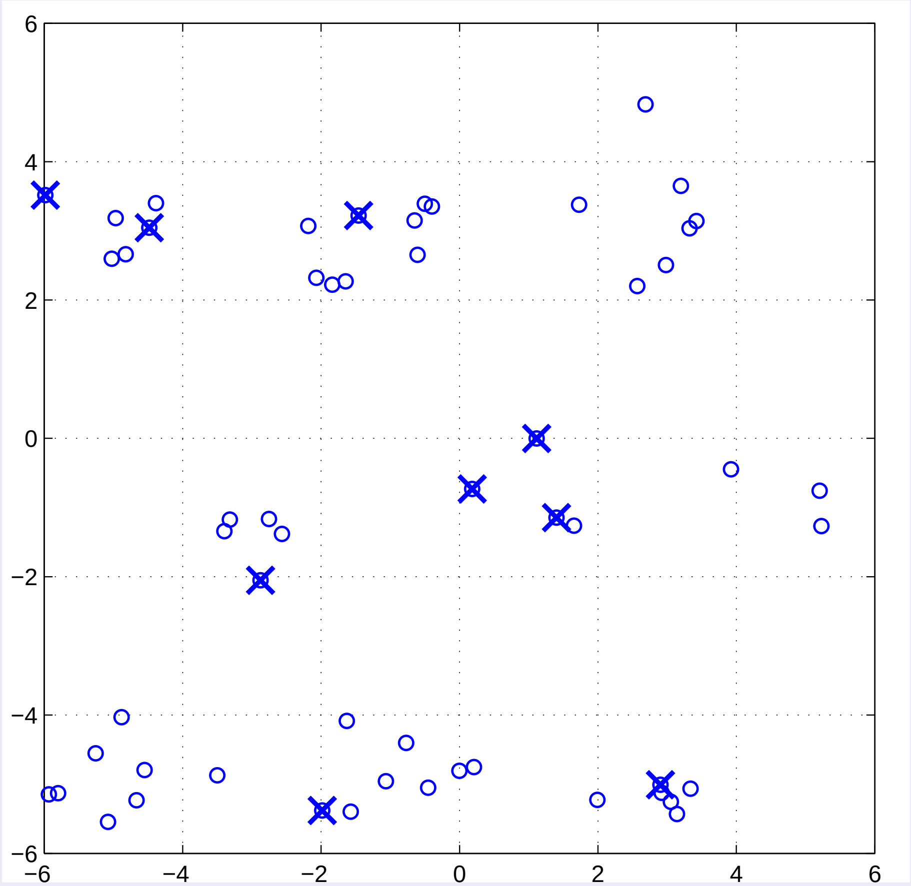

The MAC scheme selects a subset of nodes as transmitters Φ_η, for $0 \leq \eta \leq 1$ s.t. the density of the transmitter point process is $\lambda_t = \eta\lambda$.

Setup (3)

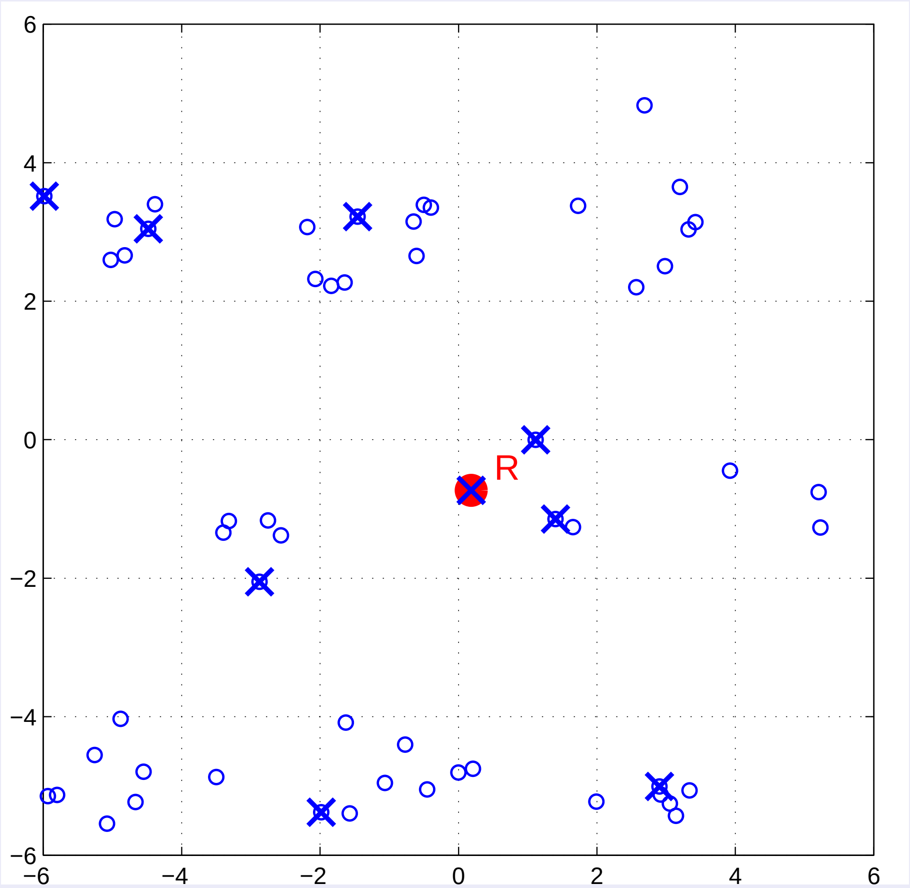

Let one transmitter be the receiver under consideration.

Setup (4)

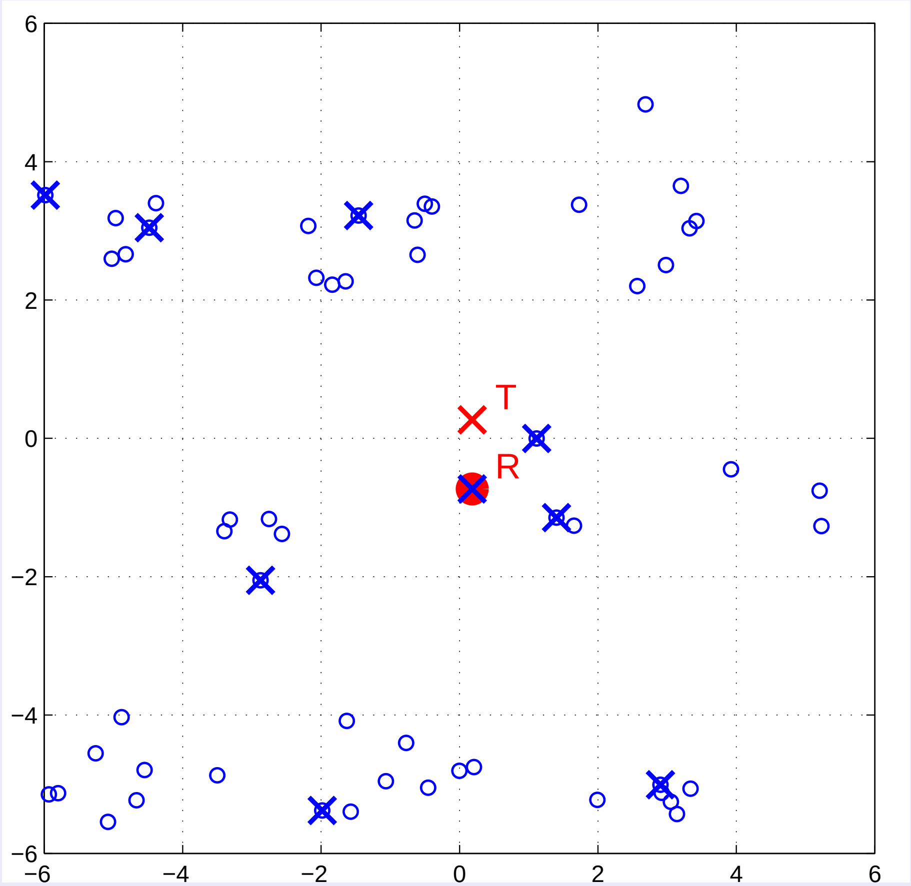

Add a virtual transmitter at unit distance, with unit transmit power.
$\Longrightarrow$ What is the outage probability from T to R as $\eta \to 0$?

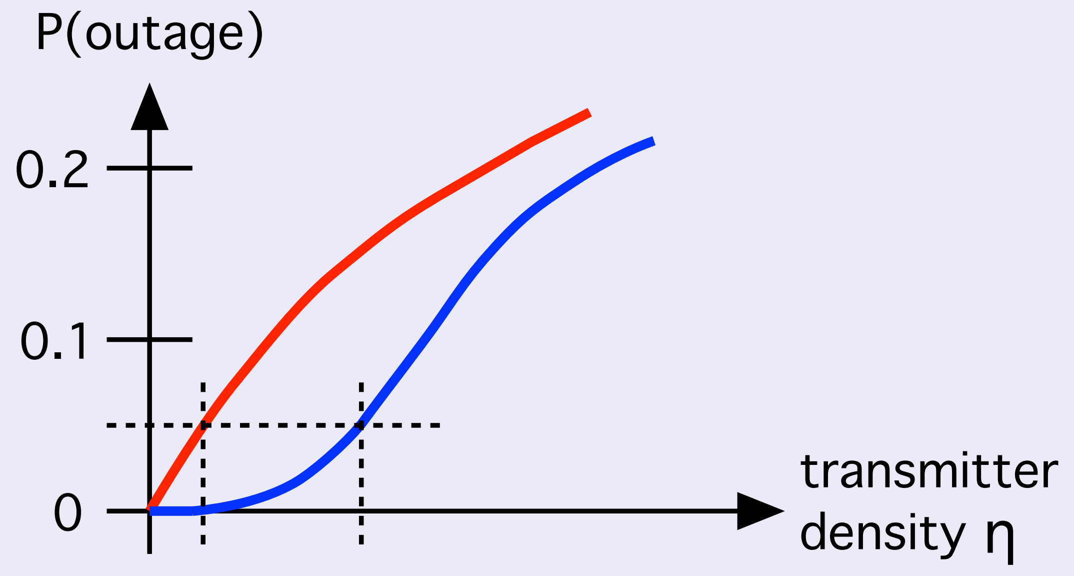

Questions

Is the outage probability near $\eta = 0$ convex or concave?
Can the network accommodate some spatial reuse without affecting the outage probability?

Problem formulation

- Take a general motion-invariant PP of intensity λ and a MAC scheme that can tune the intensity of transmitters λ_t from 0 to λ.
- Let $\eta \triangleq \lambda_t/\lambda$. What is $p_s(\eta) = \mathbb{P}(\mathrm{SIR} > \theta)$ for Rayleigh fading as $\eta \to 0$ (high-SIR asymptotics)?

Outage in the High-SIR Regime

Result [GGH10]

For all *reasonable* MAC schemes, $\exists$ unique parameters $\gamma > 0$ and $1 \leqslant \kappa \leqslant \alpha/2$ s.t.

$$p_s(\eta) \sim 1 - \gamma\eta^\kappa\,, \qquad \eta \to 0\,.$$

Moreover, $p_s(\eta) \geqslant 1 - \gamma\eta^\kappa$.
A MAC scheme is reasonable iff $\lim_{\eta\to 0} p_s(\eta) = 1$.

$\gamma(\alpha, \theta)$ is the spatial contention parameter that captures the spatial reuse capability of a network. The smaller the better.
$\kappa(\alpha)$ is the interference scaling parameter and measures the *coordination level* of the MAC. The larger the better.

Result (from previous slide)

$$p_s(\eta) \sim 1 - \gamma\eta^{\kappa} \qquad (\eta \to 0)$$

Discussion

- For all networks that use ALOHA, $\kappa = 1$. We know the result for the PPP:
$$p_s = \exp(-\eta\gamma) \qquad \Longrightarrow \kappa = 1\,.$$
- For lattices with TDMA, $\kappa = \alpha/2$.
- CSMA with sensing range $\Theta(\eta^{-1/2})$ also achieves $\kappa = \alpha/2$ (hard-core process).
- *Conjecture:* For Rayleigh fading, for all $0 \leq \eta \leq 1$,
$$1 - \gamma\eta^{\kappa} \leq p_s(\eta) \leq \frac{1}{1 + \gamma\eta^{\kappa}}\,.$$

Reasonable and *unreasonable* ALOHA for clustered point process

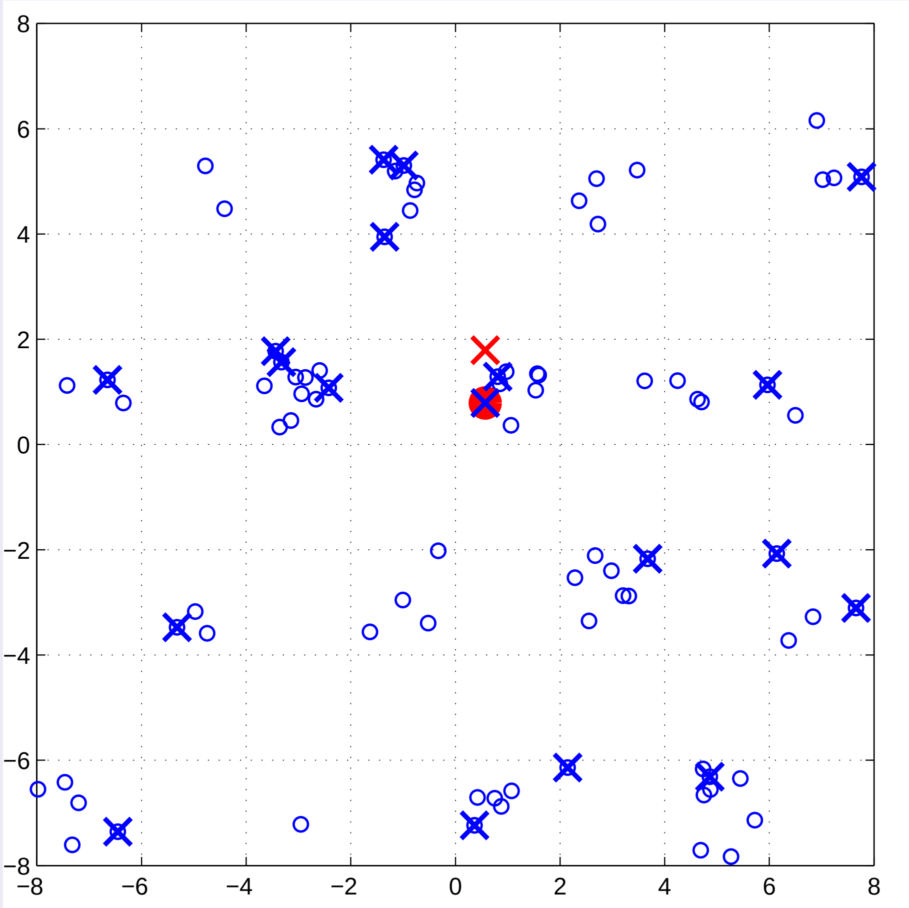

Reasonable:
20% of nodes transmit.

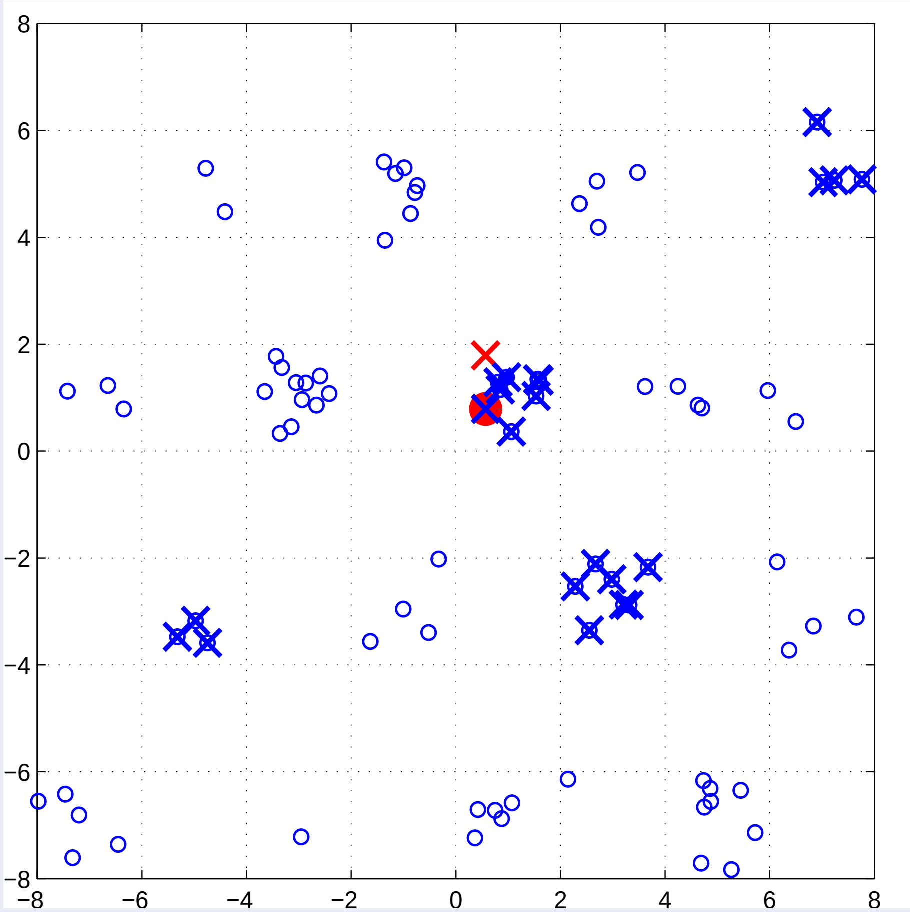

Unreasonable:
20% of *clusters* transmit.

CSMA ($\alpha = 4$)

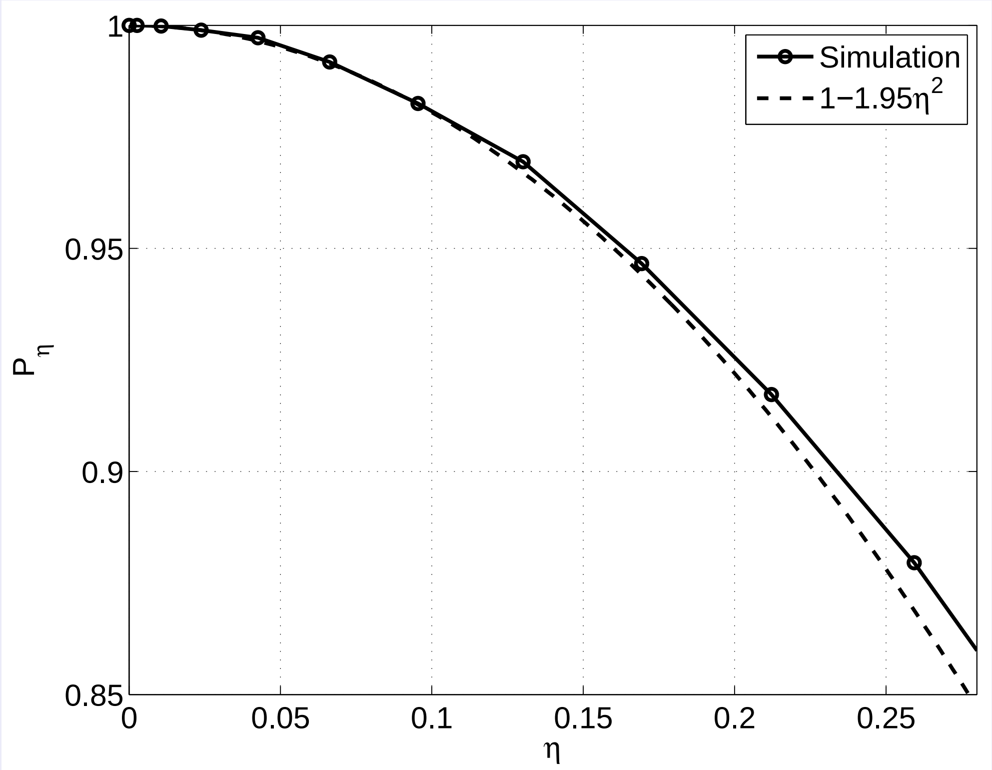

Starting with a PPP of intensity $\lambda = 0.3$, the hard-core distance is adjusted to get $\lambda_t = 0.3\eta$. $\gamma \approx 1.95$ can be analytically determined, and $\kappa = \alpha/2 = 2$.

The asymptotic expression provides a good bound for $\eta \in [0, 0.3]$.

Reasonable TDMA on square lattice

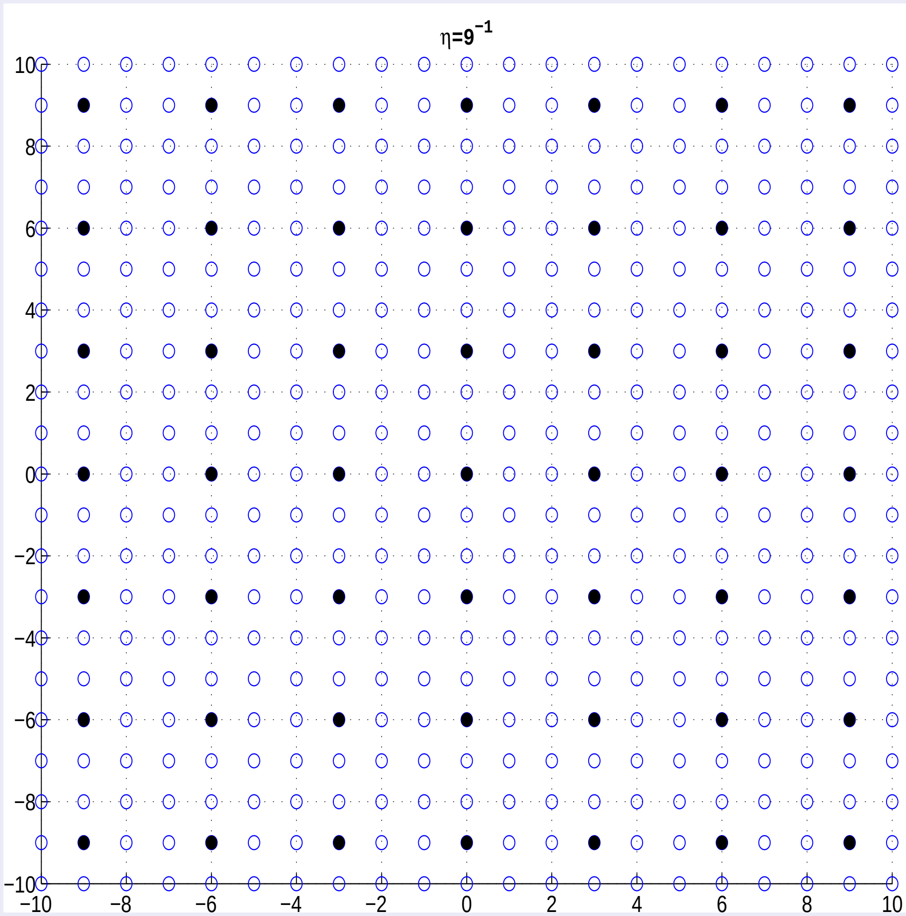

$\eta = 1/9$

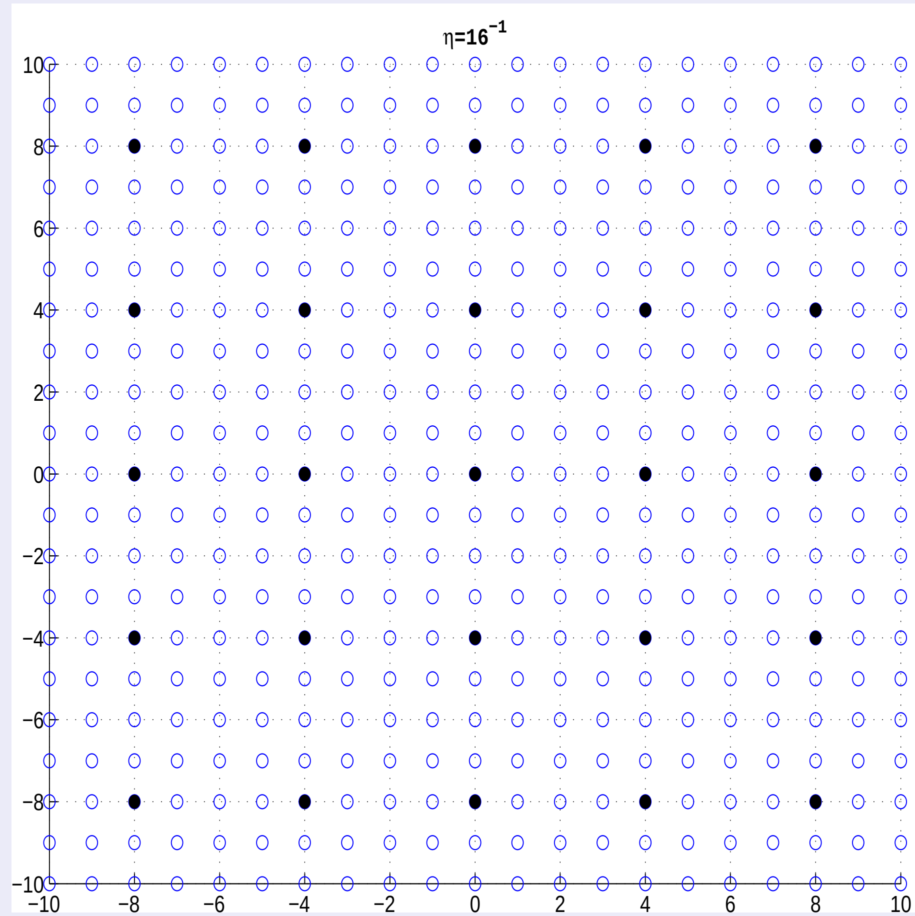

$\eta = 1/16.$

Minimum distance increases with $\eta^{-1/2}$.

Unreasonable TDMA on square lattice

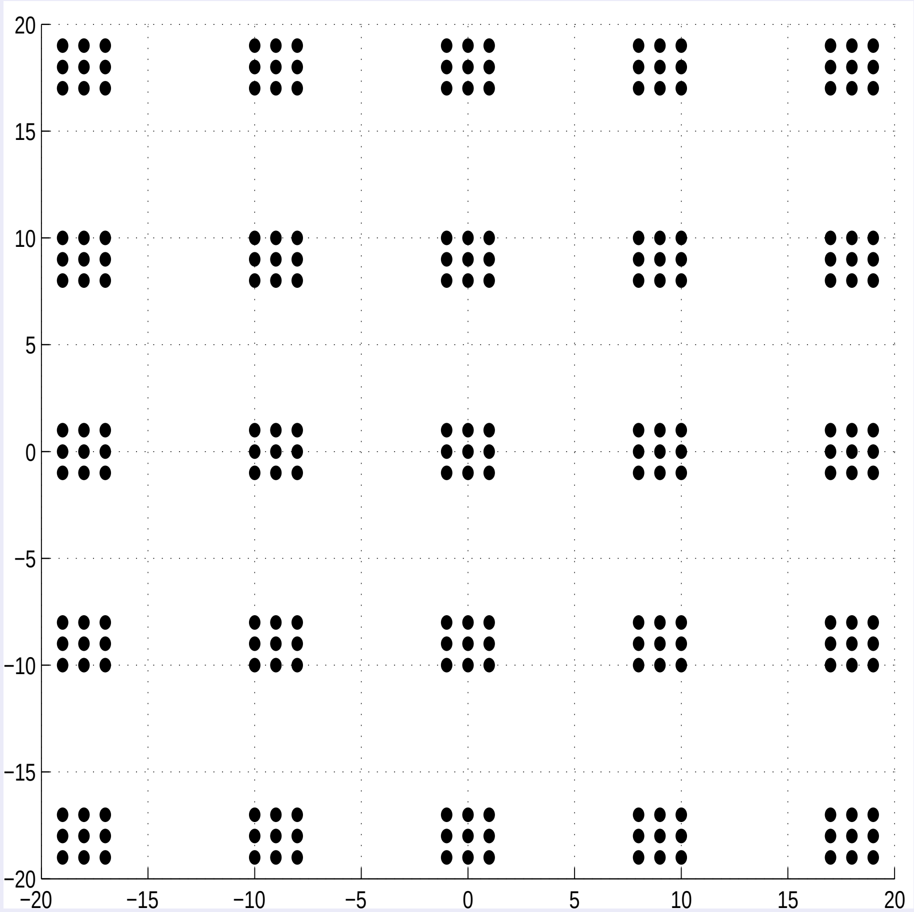

$\eta = 1/9$

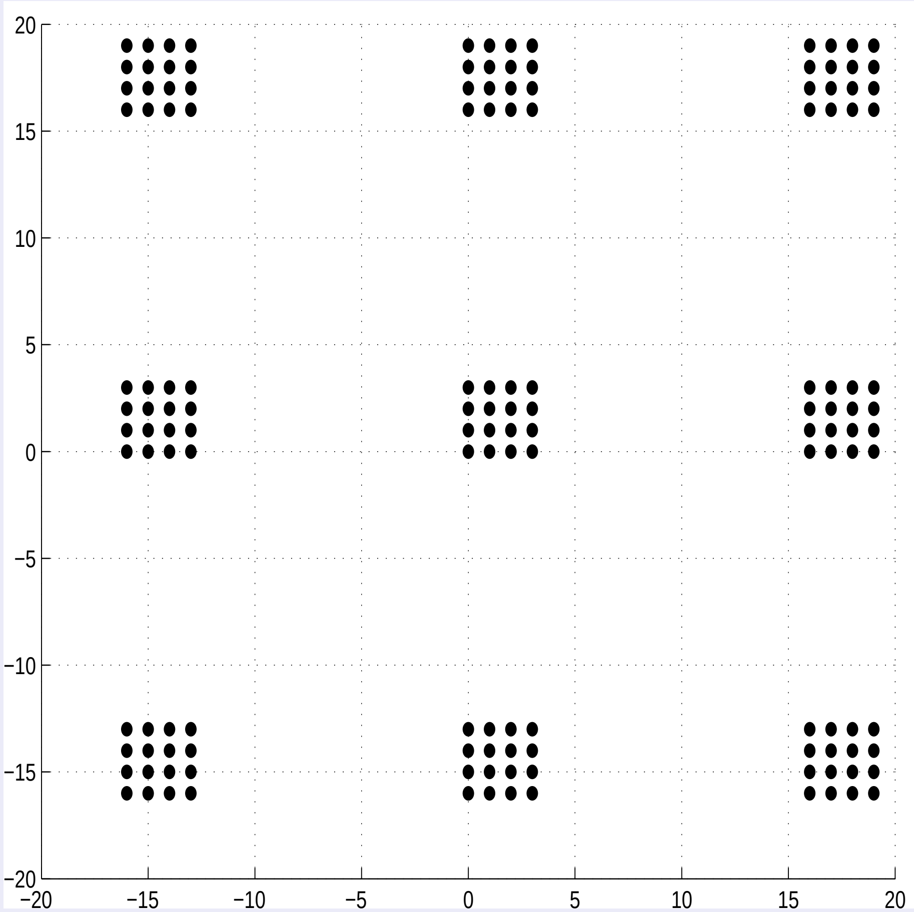

$\eta = 1/16.$

Minimum distance does not increase with decreasing η. $\lim_{\eta\to 0} p_s(\eta) < 1$.
Actually, $p_s(\eta)$ *decreases* with decreasing η.

Extension to General Fading Statistics

Asymptotic success probability with general fading [GAH10]

Let $\bar{F}(x)$ be the complementary cdf of the fading random variables. We assume that $\bar{F}$ has a Taylor series expansion given by

$$\bar{F}(x) = 1 - c_o x^{\nu} + \sum_{k=1}^{\infty} \frac{c_k}{k!} x^{k+\nu} , \quad \nu \in \mathbb{N}, \; c_0 > 0 .$$

We still have

$$p_s(\eta) \sim 1 - \gamma \eta^{\kappa} , \quad \eta \to 0 ,$$

but now

$$1 \leq \kappa \leq \alpha\nu/2 .$$

- For Rayleigh fading, $\bar{F}(x) = 1 - x + \Theta(x^2)$, so $\nu = 1$ as expected.
- For Nakagami-m fading, $\nu = m$.
- For fading distributions with smaller likelihood of very small values, ν is larger.

Section Outline

1. General Point Processes
2. Palm Theory
3. Analysis of Poisson Cluster Processes
4. Outage Probability in General Networks
5. Summary

Lecture 6 Summary

Analysis of general point processes

- The analysis of non-Poisson point processes is difficult due to the dependence among the node locations.
- The analysis requires the use of Palm theory and higher-order statistics such as the reduced second moments measures and second-order product densities.
- By considering the right asymptotic regimes, sharp statements are still possible. In particular, the success probability is in great generality

$$p_s(\eta) \sim 1 - \gamma\eta^{\kappa}\,, \quad \eta \to 0\,,$$

for a spatial contention parameter γ and an interference scaling exponent κ.

References I

R. K. Ganti, J. G. Andrews, and M. Haenggi.
High-SIR Transmission Capacity of Wireless Networks with General Fading and Node Distribution.
IEEE Transactions on Information Theory, 2010.
Submitted. Available at http://www.nd.edu/~mhaenggi/pubs/tit10b.pdf.

R. Giacomelli, R. K. Ganti, and M. Haenggi.
Outage Probability of General Ad Hoc Networks in the High-Reliability Regime.
IEEE/ACM Transactions on Networking, 2010.
Submitted. Available at http://www.nd.edu/~mhaenggi/pubs/ton10.pdf.

R. K. Ganti and M. Haenggi.
Interference and Outage in Clustered Wireless Ad Hoc Networks.
IEEE Trans. on Information Theory, 55(9):4067–4086, September 2009.
Available at http://www.nd.edu/~mhaenggi/pubs/tit09.pdf.

Dietrich Stoyan, Wilfrid S. Kendall, and Joseph Mecke.
Stochastic Geometry and its Applications.
John Wiley & Sons, 1995.
2nd Ed.

Analysis and Design of Wireless Networks

Lecture 7: Emerging Architectures

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop Analysis of Poisson Networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 7 Overview

1. Cognitive Networks
2. Cellular Networks
3. Femtocells
4. Summary

Section Outline

1. Cognitive Networks
   - Regulations
   - Unlicensed access
   - Interference
   - Summary
   - Interference in cognitive networks
   - Application to TV white space
   - Cognitive peer-to-peer networks
   - Outlook
   - Conclusions
   - References

2. Cellular Networks

3. Femtocells

Cognitive Networking

Ingredients

- A wireless network operated by an *incumbent user*
- A secondary or cognitive user who wishes to operate a network in the same frequency band
- Software-defined radios
- Maxwell's equations
- Government regulations and spectrum policies

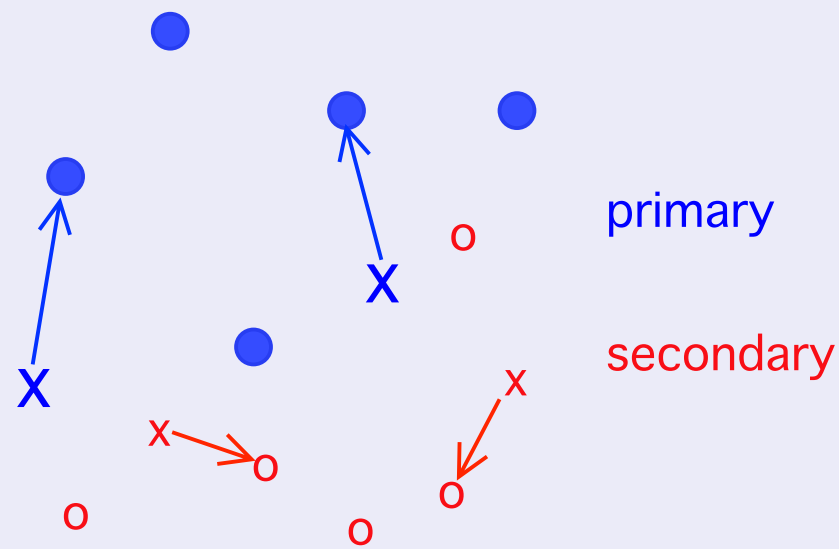

Regulations

US Government Agencies

- **NTIA**: National Telecommunications and Information Administration (www.ntia.doc.gov). Part of US Dept. of Commerce. Manages federal use of spectrum.
  OSM: Office of Spectrum Management (www.ntia.doc.gov/osmhome/Osmhome.html).
- **FCC**: Federal Communications Commission (www.fcc.gov). Manages all other uses of spectrum.
  Wireless Telecommunications Bureau (wireless.fcc.gov).
  Spectrum Policy Task Force (http://www.fcc.gov/sptf/).

Unlicensed Access

2008 FCC Report and Order and Memorandum (FCC 08-260)

Permits "unlicensed operation in the TV broadcast bands" and promises "additional spectrum for unlicensed devices below 900 MHz and in the 3 GHz band". (Nov. 4, 2008).

Accessing a database of all fixed devices

All devices, except personal/portable devices operating in client mode, must include a geolocation capability and provisions to access over the Internet a database of protected radio services and the locations and channels that may be used by the unlicensed devices at each location.

Sensing

Alternatively, unlicensed users may sense the presence of primary users and transmit if they do not detect any primary transmission they could interfere with.

Spectrum Sensing (FCC 08-260)

We will permit applications for certification of devices that do not include the geolocation and database access capabilities, and instead rely on spectrum sensing to avoid causing harmful interference, subject to a much more rigorous set of tests by our Laboratory in a process that will be open to the public. These tests will include both laboratory and field tests to fully ensure that such devices meet a "Proof of Performance" standard that they will not cause harmful interference.
Devices (operating in either mode) will be required to sense TV signals, wireless microphone signals, and signals of other services that operate in the TV bands, including those that operate on intermittent basis, at levels as low as -114 dBm.

Sensing difficulty

Detecting digital TV signals is easy due to their embedded pilot tones.
Detecting wireless microphones, however, is difficult.

Wireless microphone usage

"Going digital would destroy the soul of the music!"

Sensing wireless microphones (FCC 08-260)

Wireless microphones will be protected in a variety of ways. The locations where wireless microphones are used, such as entertainment venues and for sporting events, can be registered in the database and will be protected as for other services. In addition, channels from 2—20 will be restricted to fixed devices, and we anticipate that many of these channels will remain available for wireless microphones that operate on an itinerant basis. In addition, in 13 major markets where certain channels between 14 and 20 are used for land mobile operations, we will leave 2 channels between 21 and 51 free of new unlicensed devices and therefore available for wireless microphones. Finally, as noted above, we have required that devices also include the ability to listen to the airwaves to sense wireless microphones as an additional measure of protection for these devices.

Quote (graduate student trying to sense a wireless microphone signal)

"Detecting a wireless microphone is like finding a needle in a haystack. Its signal is very narrow, and it can be anywhere in the spectrum."

TV White Space DSA

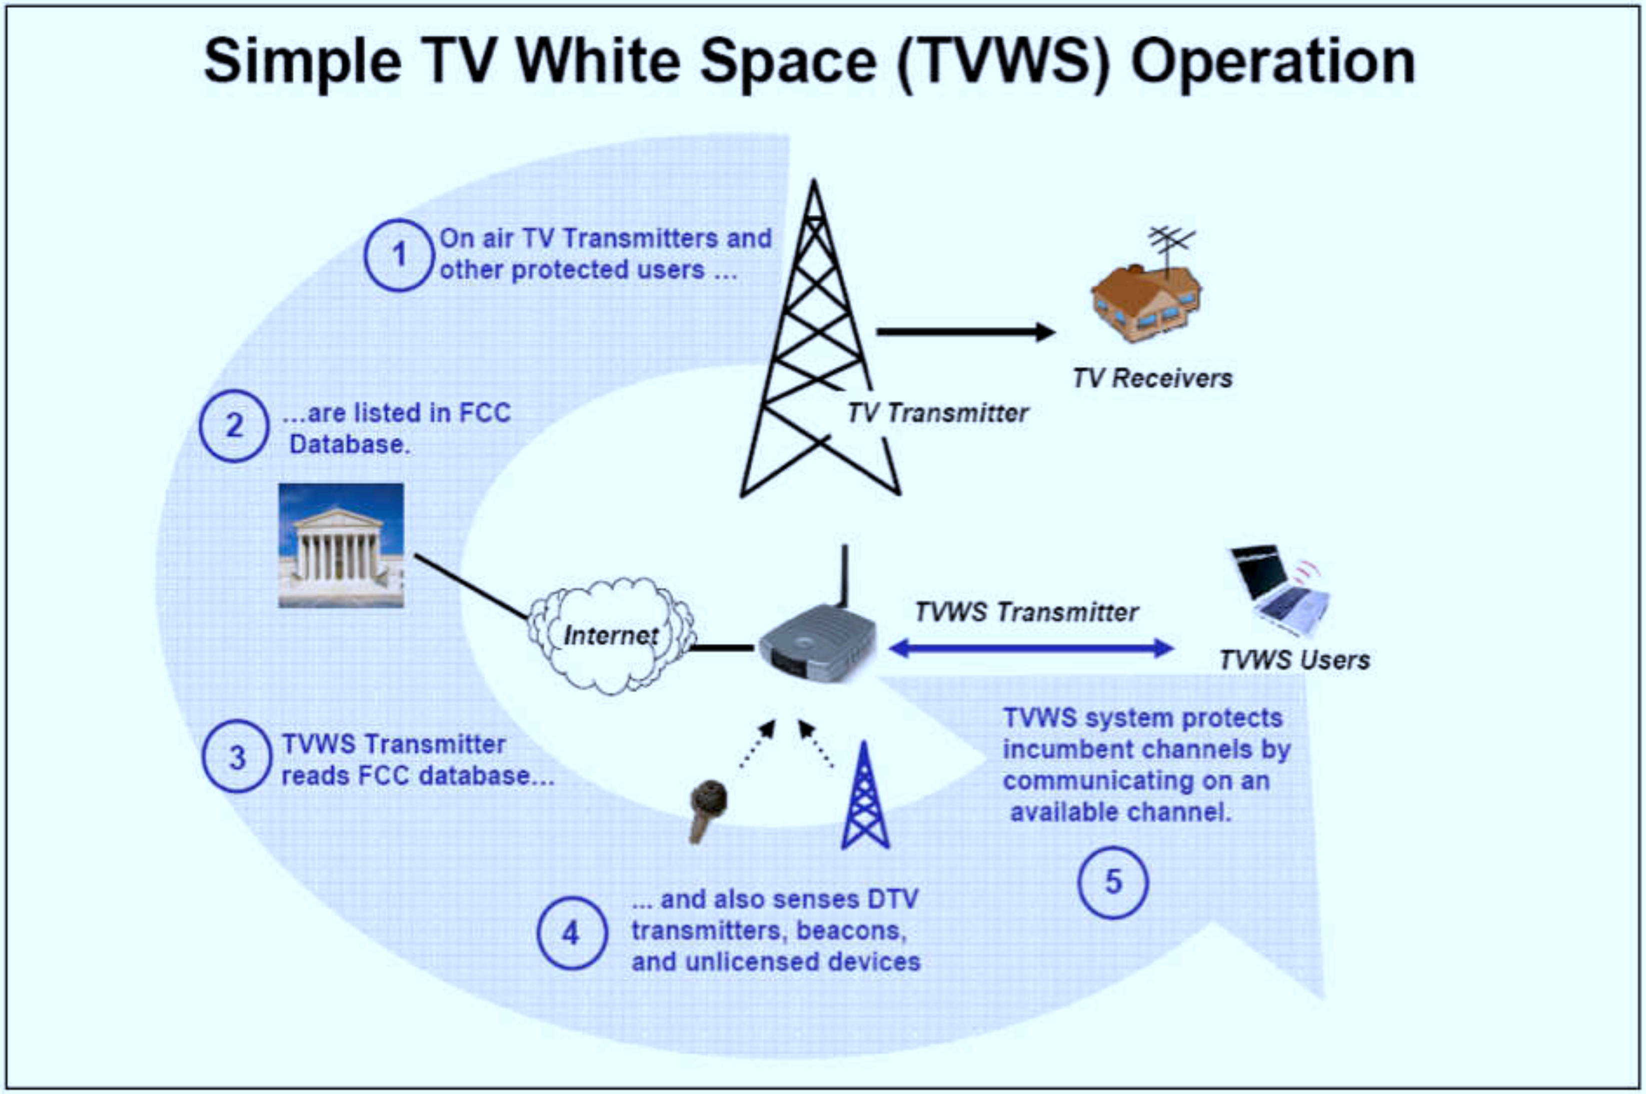

(From "Considerations for Successful Cognitive Radio Systems in US TV White Space", D. Borth et al., Motorola Inc, DySPAN 2008.)

The database catch 22

Short distance secondary link:

- The database can only be accessed over a wired connection
- If both secondary Tx and Rx need to access the database, they may also communicate over the wired link
- If only one does (can), how does it tell its partner node what frequency to use?

Long-distance secondary link:

- Tx and Rx may have different pictures of the primary user activity. How do they negotiate?
- If the Rx is in a rural area, it may not have database access, at least not very dynamically.

In both cases, CUs may not be aware of other CUs. The cumulative interference is not known.

What is Interference?

Definition (Interference)

The effect of unwanted energy due to one or a combination of emissions, radiations, or inductions upon reception in an RF communications system, manifested by any performance degradation, misinterpretation, or loss of information which could be extracted in the absence of such unwanted energy.

Permissible vs. harmful interference

Permissible interference: Defined as any interference allowed by the FCC. On the other hand, harmful interference is prohibited.

Harmful interference

Topic of heated discussion.
Google July 26, 2010: 263,000 hits for "harmful interference" (in USA).
Google July 30, 2010: 285,000 hits

Two cases with a clear definition:

- UWB: Maximum emission is limited (-48.5dBm/MHz). More than that is harmful.
- Direct Broadcast Satellite: An increase in unavailability of up to 10% is tolerable (from 0.02% to 0.022%).

But in general?

Definition (HI – http://www.its.bldrdoc.gov/fs-1037/dir-017/_2541.htm)

Any emission, radiation, or induction interference that endangers the functioning or seriously degrades, obstructs, or repeatedly interrupts a communications system, such as a radio navigation service, telecommunications service, radio communications service, search and rescue service, or weather service, operating in accordance with approved standards, regulations, and procedures.

Note: To be considered harmful interference, the interference must cause serious detrimental effects, such as circuit outages and message losses, as opposed to interference that is merely a nuisance or annoyance that can be overcome by appropriate measures.

HI—European Union (Nov. 29, 2007)

Harmful Interference means interference which degrades or interrupts radiocommunication to an extent beyond that which would reasonably be expected when operating in accordance with the applicable EU or national regulations.

EU Spectrum Management

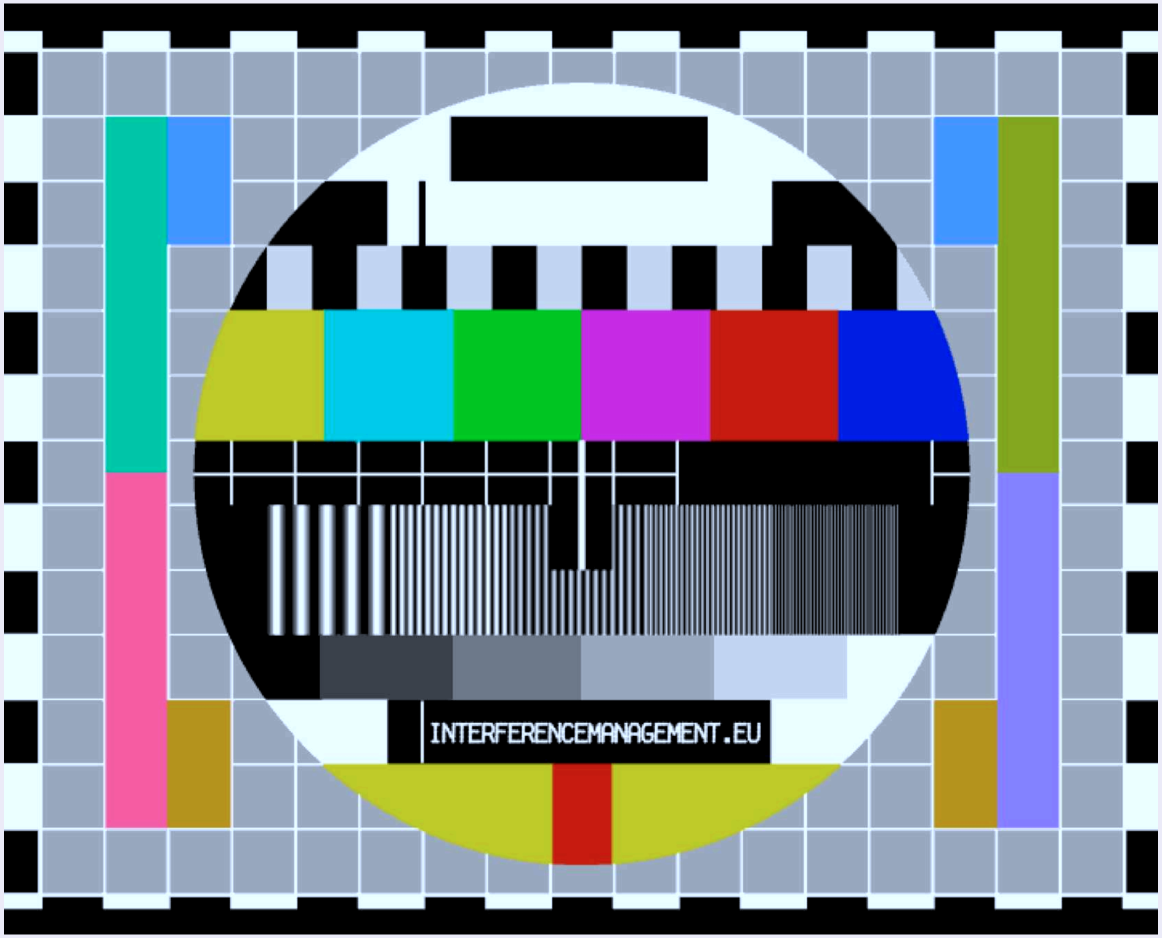

Check spectrumtalk.blogspot.com/2007/10/european-commission-workshop-on.html.

UK: Ofcom at www.ofcom.org.uk/.

Summary

Use of White Space

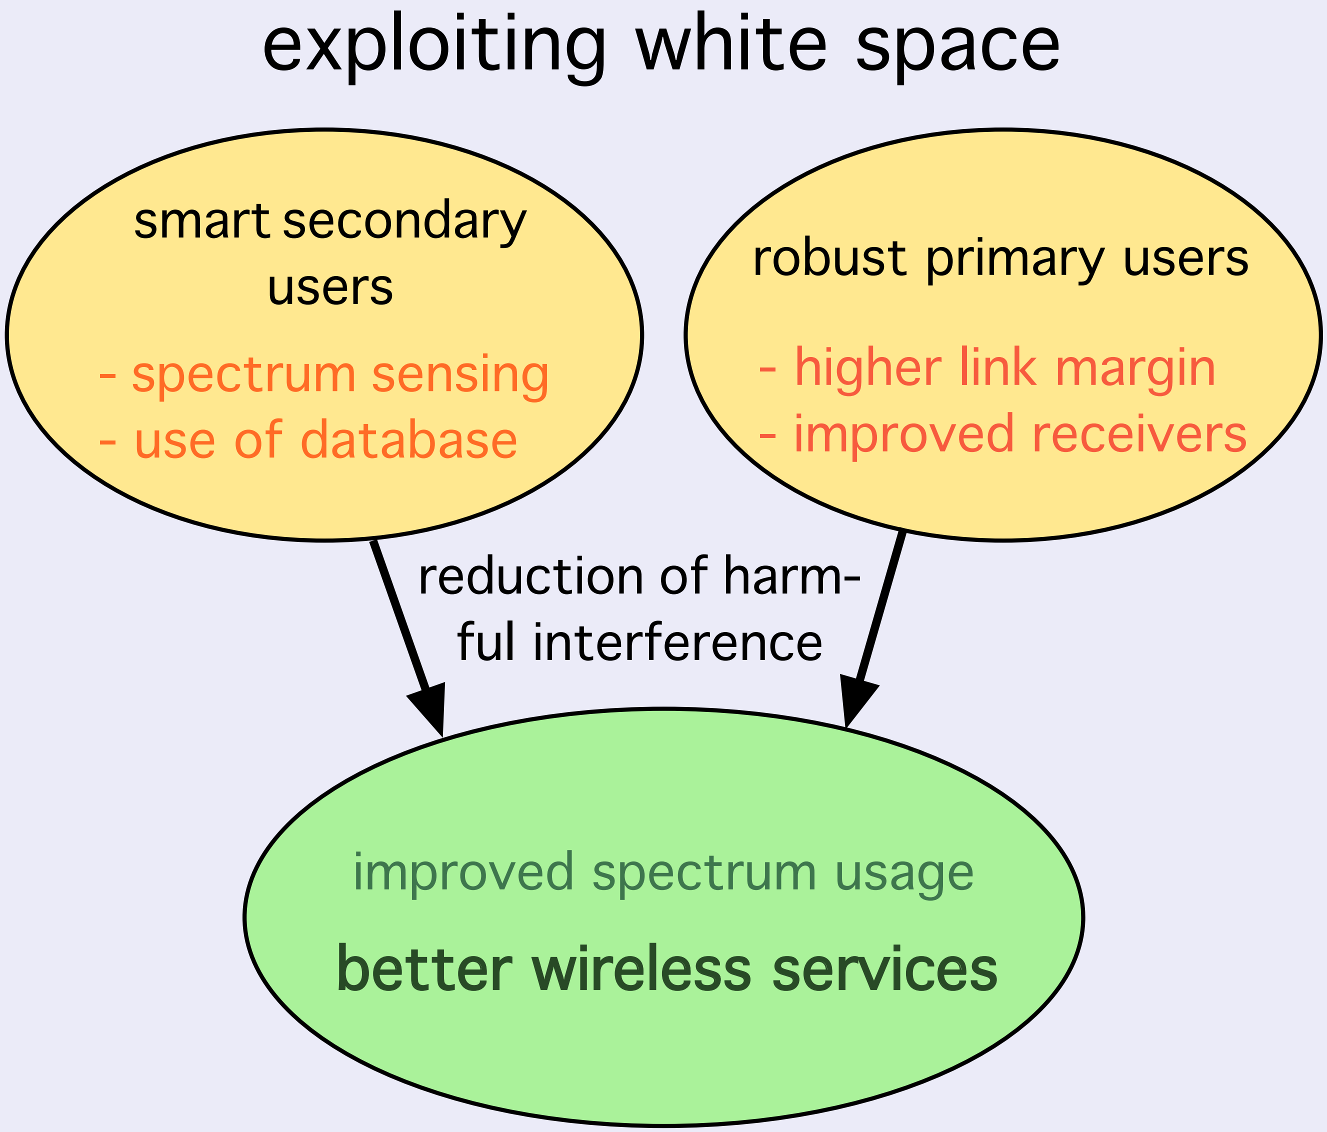

Summary

Use of White Space

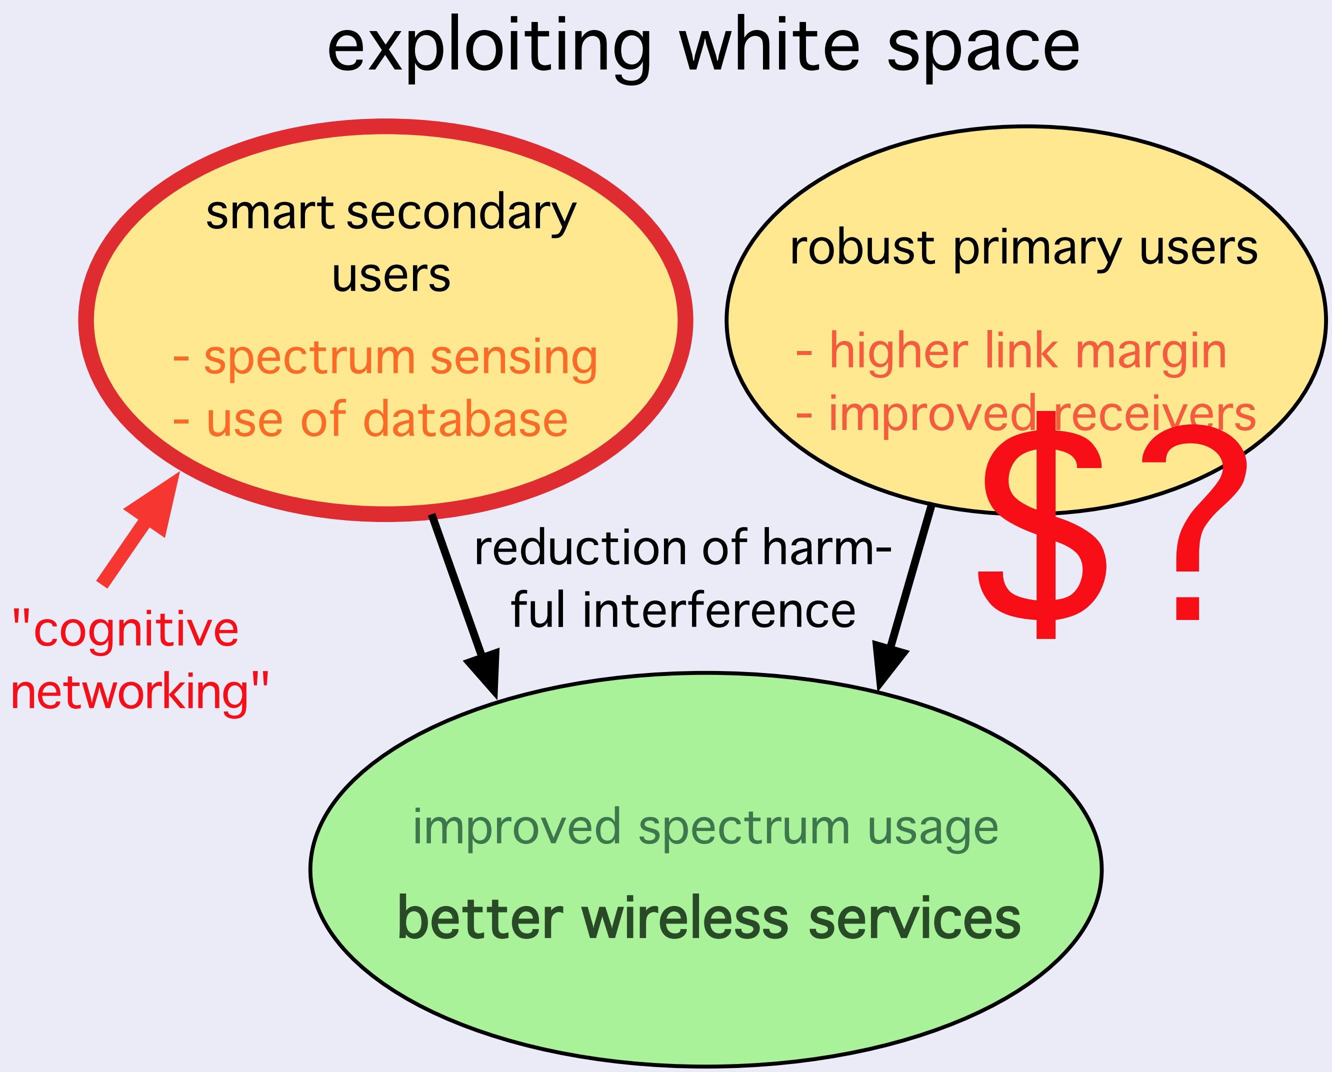

Types of interference

In a cognitive network, there are four types of interference.
Example with two primary and secondary links each:

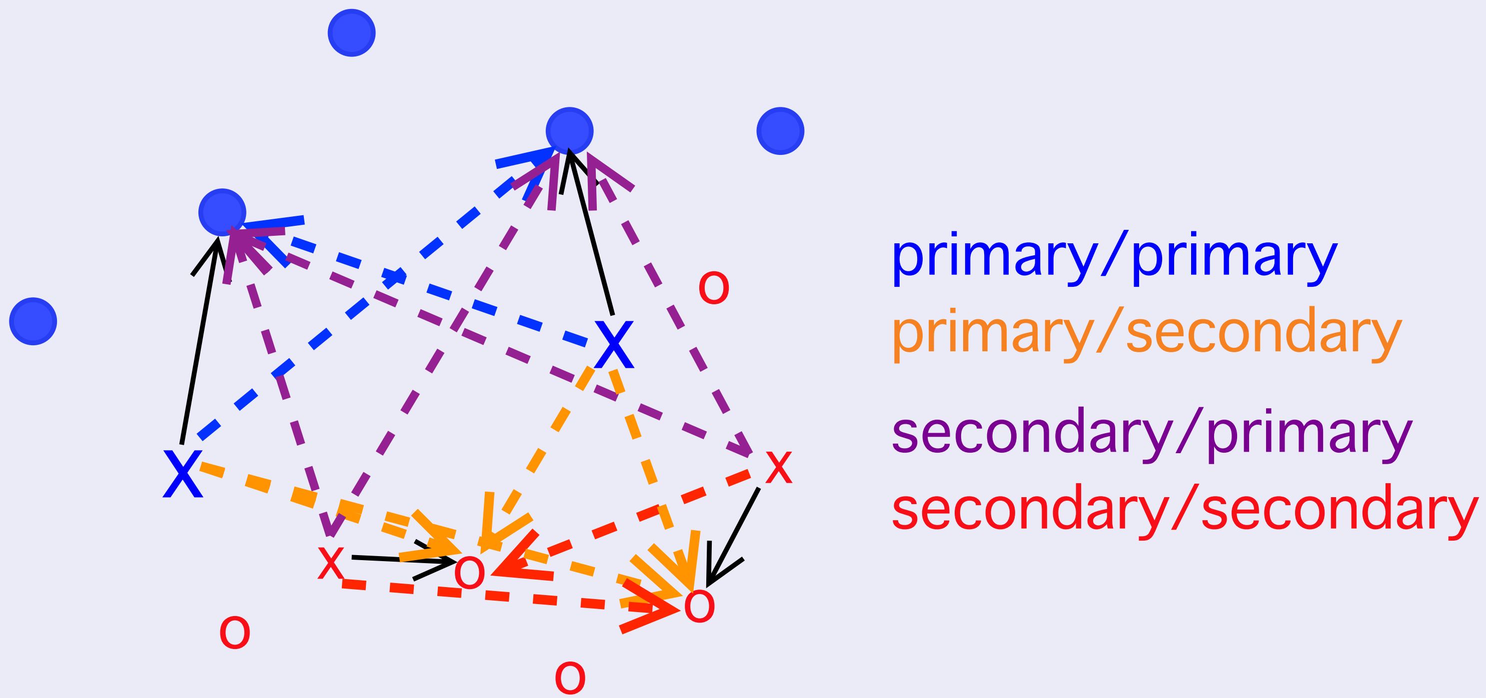

We denote the four types as I_{pp}, I_{ps}, I_{sp}, I_{ss}. The potentially harmful one I_{sp}.
Need to characterize these interferences, in the presence of unknown node locations and fading. Stochastic geometry is the right tool!

Application to TV White Space

Setup

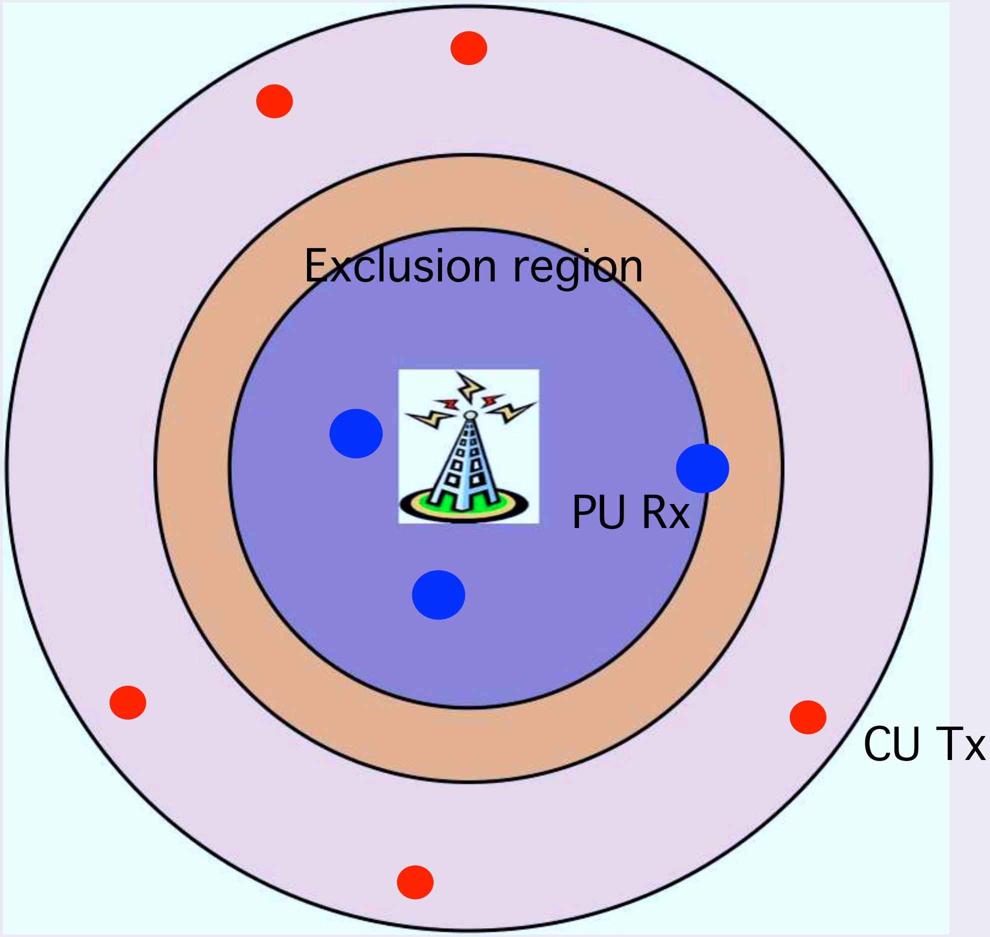

Assume CUs are uniformly randomly distributed in the red annulus with density λ (PPP).

Analysis

Goal: Satisfy the worst-case PU's interference constraint.

Distance between PU and CU at position (r, ϕ):

$$d^2(r,\phi) = r^2 + R^2 - 2Rr\cos\phi$$

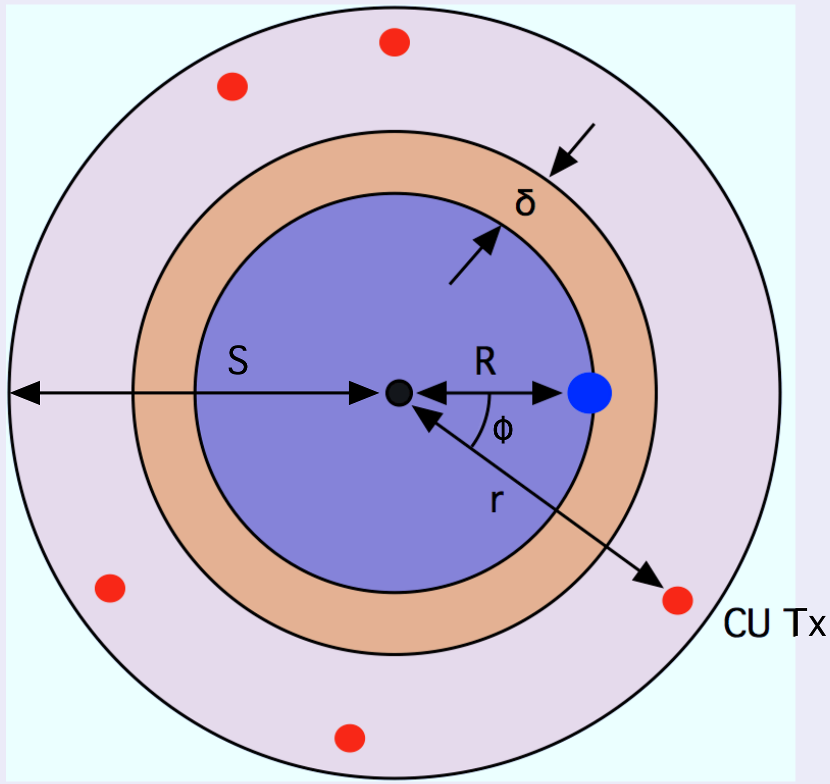

The CUs are distributed with radial pdf

$$f(x) = \frac{2x}{S^2 - (R+\delta)^2}, \quad R+\delta \le x \le S,$$

and the mean number of CUs is

$$n = \lambda\pi(S^2 - (R+\delta)^2).$$

Analysis

The mean interference is thus, by Campbell's theorem,

$$\mathbb{E}(I) = \lambda P \int_{R+\delta}^{S} \int_{0}^{2\pi} \frac{r\mathrm{d}r\mathrm{d}\phi}{(r^2 + R^2 - 2Rr\cos\phi)^{\alpha/2}},$$

which, for $\alpha = 4$, is

$$\mathbb{E}(I) = P\lambda\pi \left[\frac{(R+\delta)^2}{\delta^2(2R+\delta)^2} - \frac{S^2}{(S^2 - R^2)^2} \right].$$

The success probability is

$$p_s = \mathbb{P}(P_{\mathrm{TV}} R^{-\alpha}/I \geq \theta)$$

Using Markov's inequality, we obtain

$$p_s \geq 1 - \frac{\mathbb{E}(I)\theta R^{\alpha}}{P_{\mathrm{TV}}}$$

Example

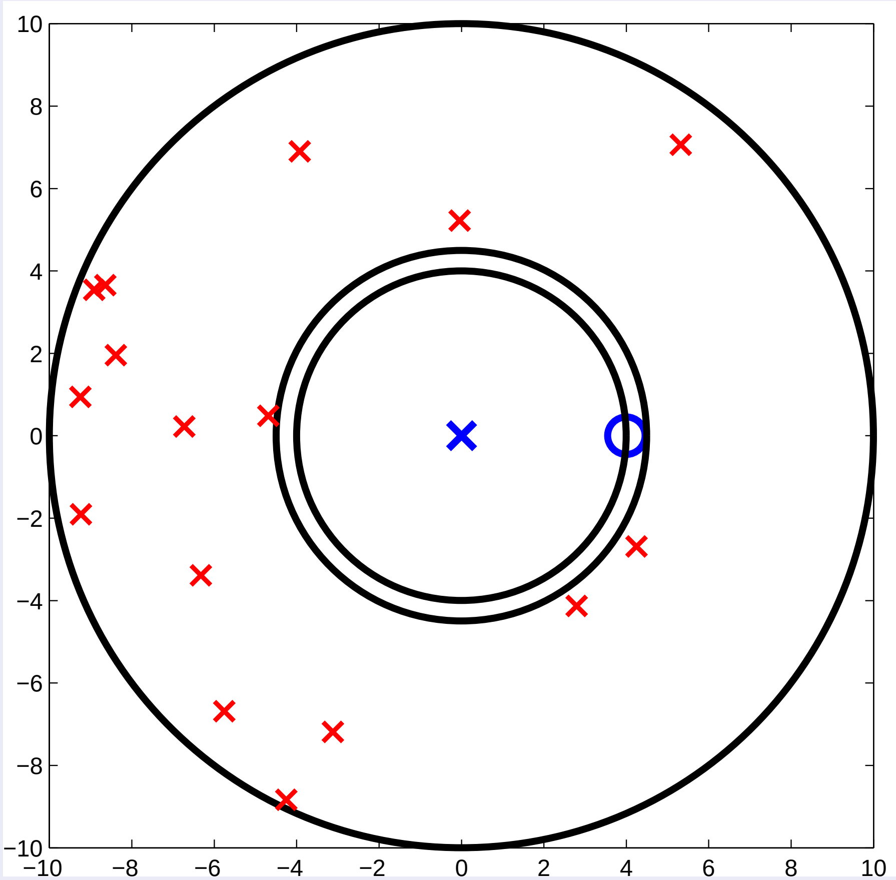

$P_{\text{TV}} = 100,\ P = 0.1,\ \lambda = 0.05,$
$R = 4,\ S = 10,\ \alpha = 4,\ \theta = 4;$
$n \approx 13.$

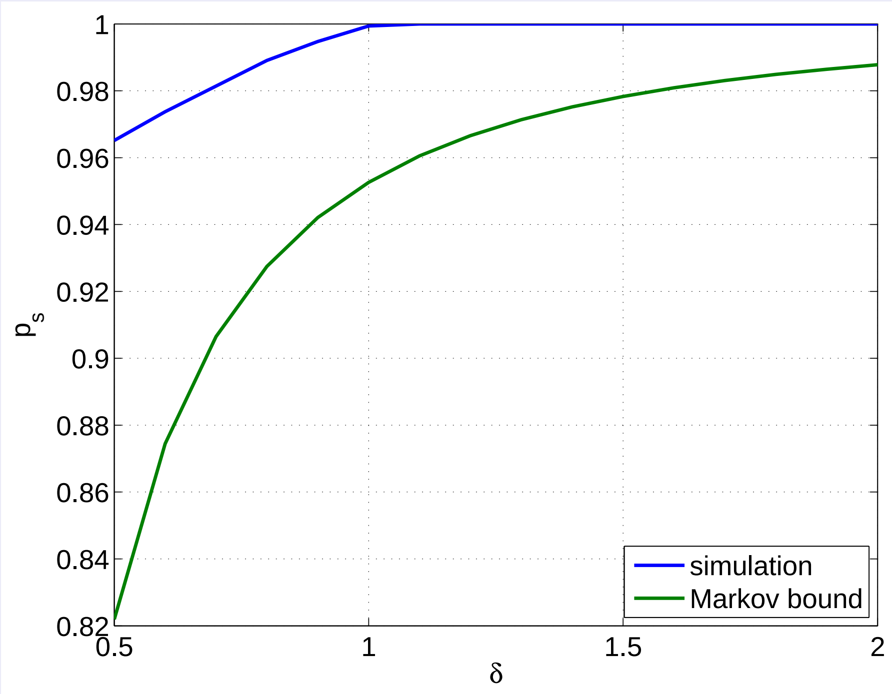

Simulation result and Markov bound as a function of the guard zone width δ.

So far so good...

The white space box

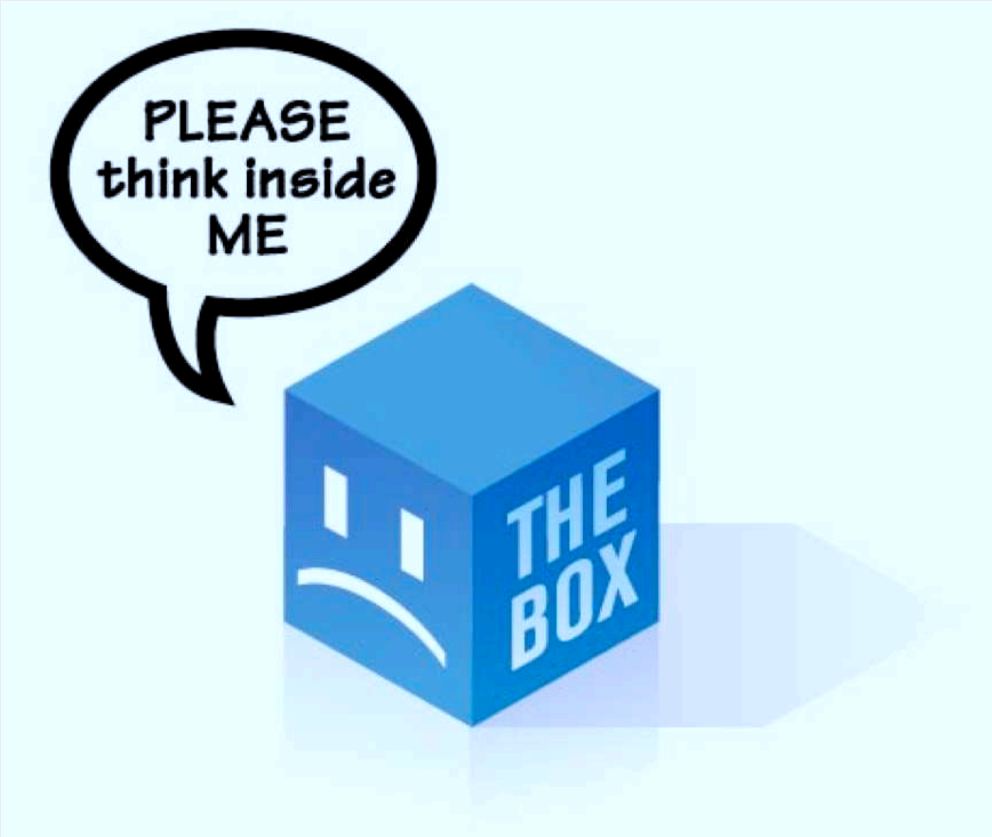

How about...

...thinking outside the white space box?

Is the wireless world just black and white?

Is there white space inside the blue space?

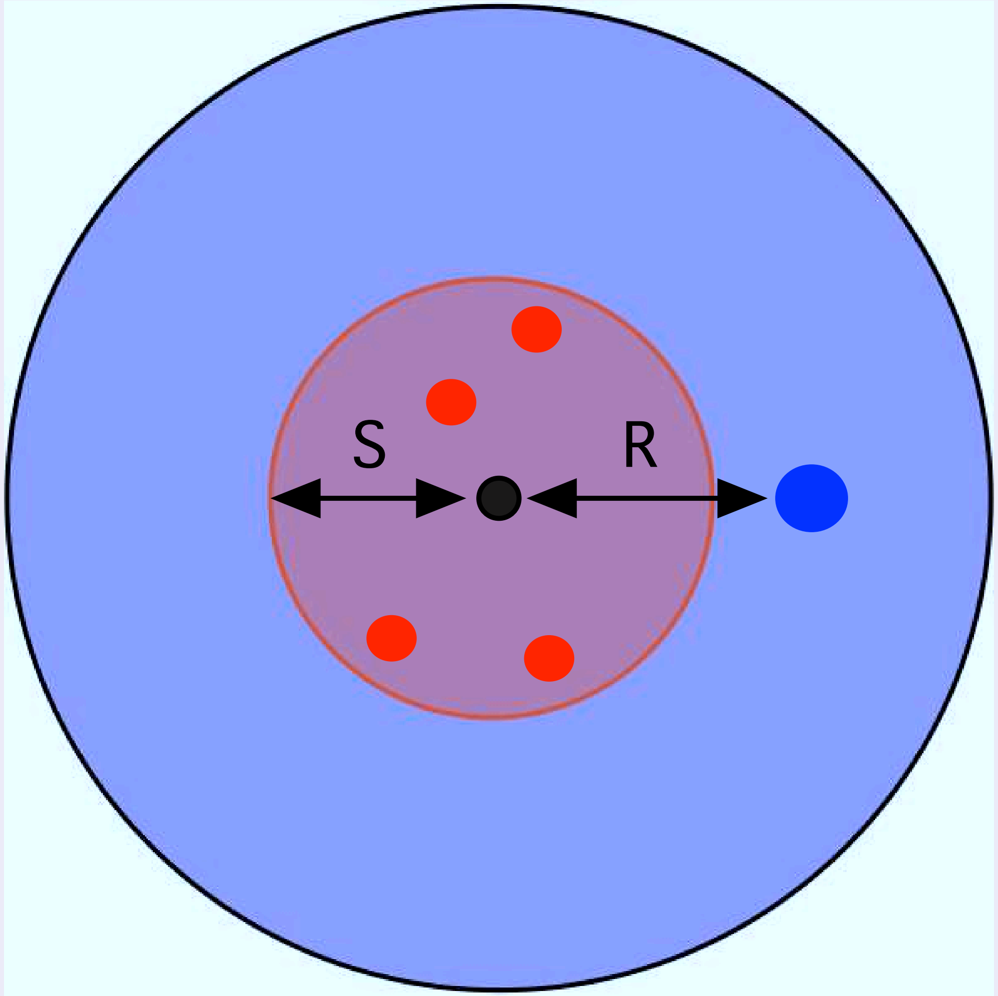

Thinking inside the blue disk...

...but why would we want to put CUs right at the TV station's epicenter??

Why does it work?

Check the SIR condition!

- Inside the disk of radius S, the PU's received signal is strong.
- Outside the disk of radius S, the interference from the CUs is weak.

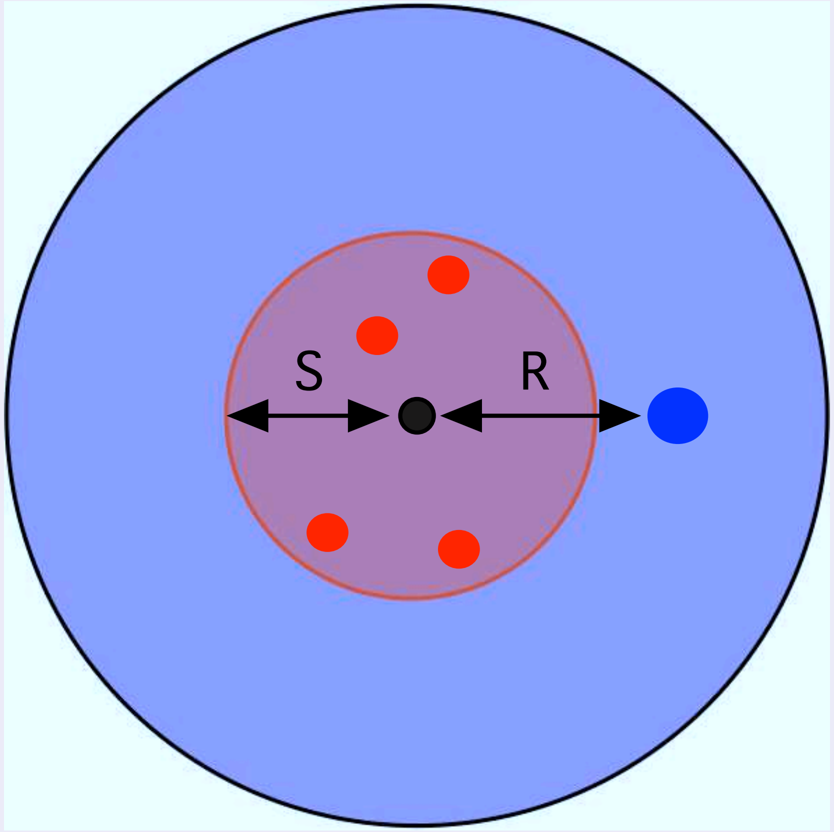

$\Longrightarrow$ Either way, the SIR condition at the PU Rx is met!

Example

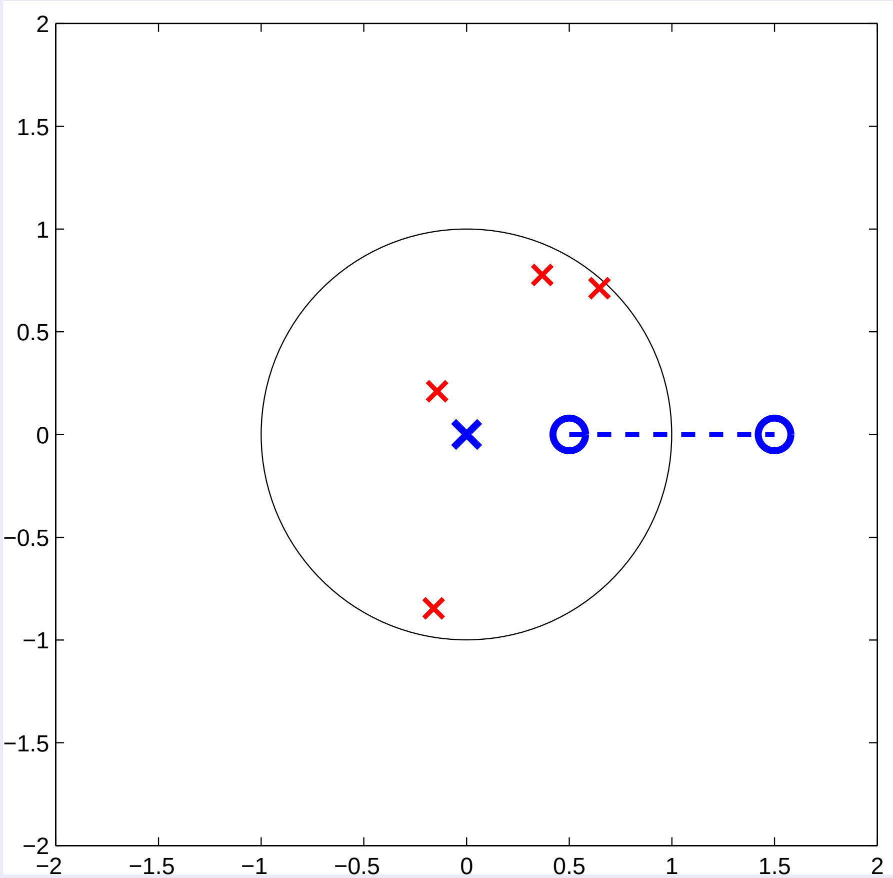

$P_{\mathrm{TV}} = 100,\ P = 0.1,\ \lambda = 1,$
$R = [1/2, 3/2],\ S = 1,\ \alpha = 4,$
$\theta = 4;\ n \approx 3.$

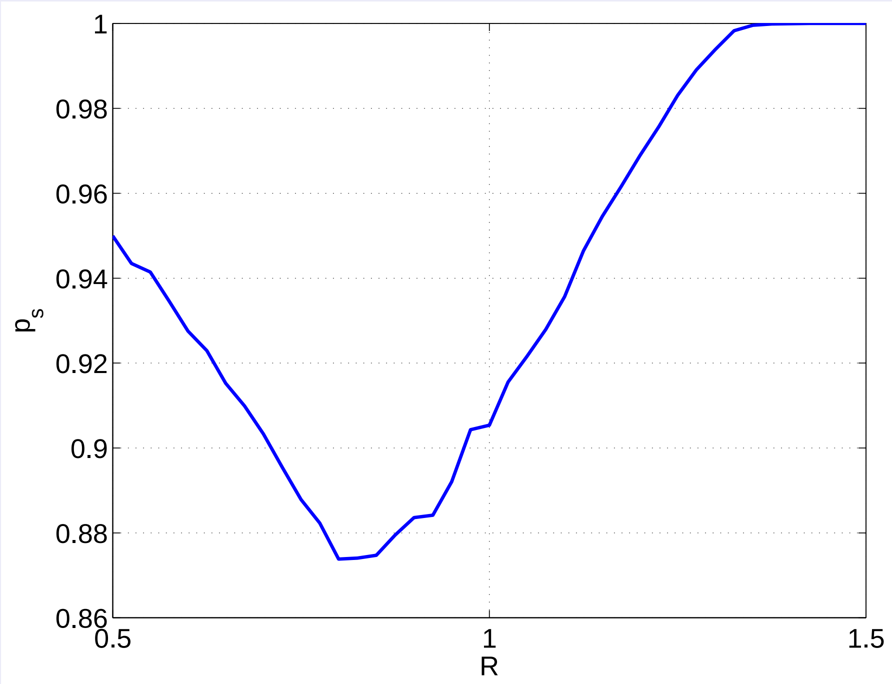

p_s a function of the PU link distance R.

How about the secondary receiver?

How is it ensured that the SIR at the secondary receiver is large enough?

- Use small link distances
- Much better: Use interference canceling techniques! The TV signal is strong and has a well-defined structure, so it can be subtracted at the secondary receiver, so that there is vanishing interference.

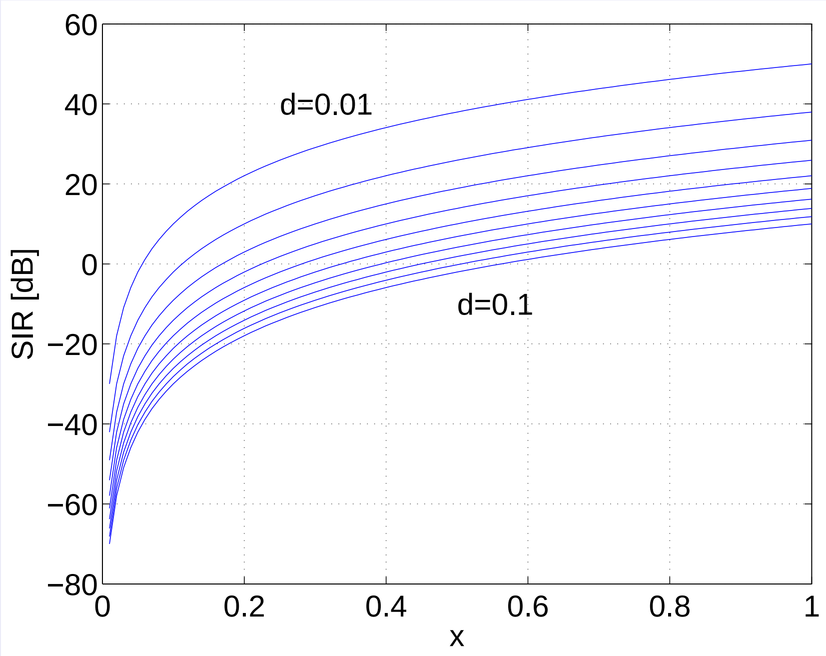

SIR at CU without IC

Interference cancellation is only possible if the interfering signal is stronger. So it is preferable to place CUs near the strong TV transmitter!

Cognitive Peer-to-Peer Networks [LH10]

Bipolar model: Setup

- PU transmitters form a PPP of intensity λ_p.
- CU *potential* transmitters form a PPP of intensity λ_s.
- PU receivers are at distance r_p.
- CU receivers are at distance r_s.
- CUs cannot be active if within distance D of a primary receiver.

The active CUs form a Poisson hole process.

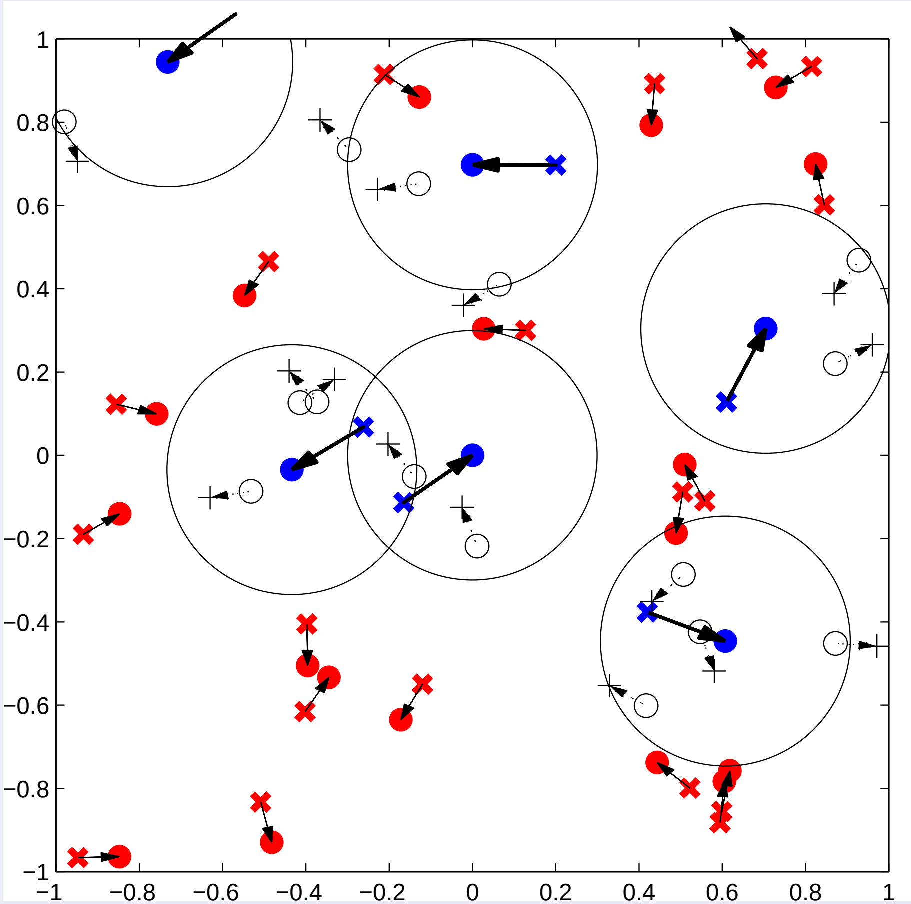

Poisson hole process

- The Poisson hole process with fixed guard zone models a cognitive bipolar peer-to-peer network.
- It is a stationary and isotropic point process.
- Interference compared to the Poisson/Poisson case without guard zone:

  - $I_{\rm pp}$ is unchanged.
  - $I_{\rm ps}$ is smaller, since there is a minimum distance $D - r_p - r_c$ between a primary Tx and a secondary Rx.
  - $I_{\rm sp}$ is (much) smaller, due to the guard zone D.
  - $I_{\rm ss}$ changes only due to the smaller intensity of secondary transmitters. $\lambda'_s = \lambda_s \exp(-\lambda_p \pi D^2)$.

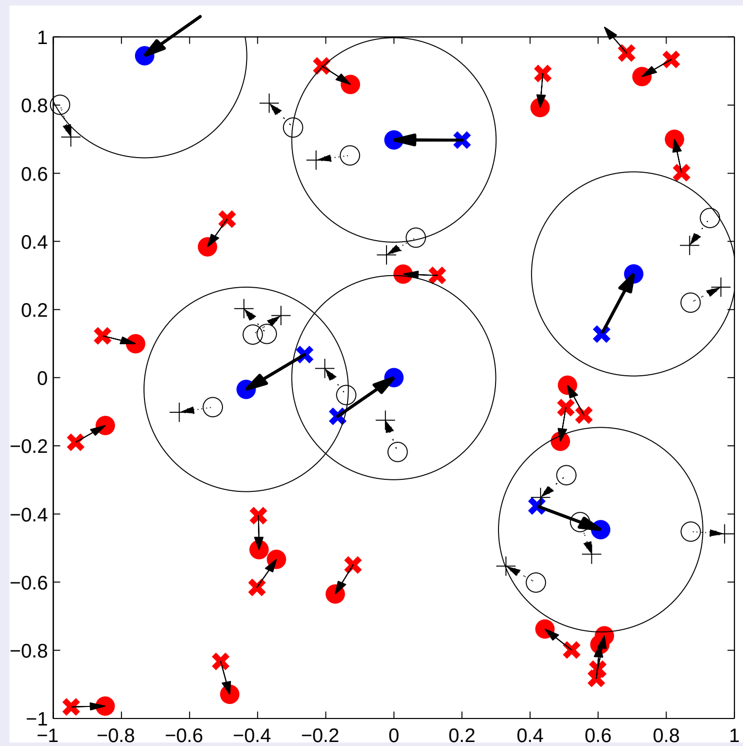

Interference and outage

The total interference at the typical PU Rx is $I = I_{\rm pp} + I_{\rm sp}$. Let $\delta \triangleq 2/\alpha$.

$$I_{\rm pp} \triangleq \sum_{x\in\Phi_p} P h_x \|x\|^{-\alpha}$$

$$\mathcal{L}_{I_{\rm pp}}(s) = \mathbb{E}\exp(-sI) = \exp\left(-\lambda_p \frac{\pi^2\delta}{\sin(\pi\delta)} P^\delta s^\delta\right).$$

Success probability within PUs:

$$\mathbb{P}(S/I_{\rm pp} > \theta) = \mathcal{L}_{I_{\rm pp}}(\theta r_p^\alpha/P) = \exp\left(-\lambda_p r_p^2 \frac{\pi^2\delta}{\sin(\pi\delta)}\theta^\delta\right)$$

Total success probability: Since $I_{\rm pp}$ and $I_{\rm sp}$ are negatively correlated:

$$\mathbb{P}(\mathrm{SIR} > \theta) \leq \mathcal{L}_{I_{\rm pp}}(\theta r_p^\alpha/P)\cdot\mathcal{L}_{I_{\rm sp}}(\theta r_p^\alpha/P) \qquad \text{(by FKG)}.$$

But we don't know $I_{\rm sp}$.

Interference and outage

The critical interference term is $I_{\rm sp}$. The point process of transmitting CUs is the Poisson hole process. There are three possibilities to approximate of bound $I_{\rm sp}$ and the outage probability:

1. Approximate the Poisson hole process with a Poisson cluster process by matching first- and second-order statistics. Use known results for Poisson cluster processes to proceed.
2. Upper bound the interference by only excluding the CUs outside the reference receiver.
3. Approximate the interference by a PPP of secondary transmitters of intensity $\lambda_s \exp(-\lambda_p \pi D^2)$ outside the guard zone.

We focus on Methods 2 and 3. In both cases, the approximate interference $\hat{I}_{\rm sp}$ is independent of $I_{\rm pp}$, *i.e.*, we're restoring independence.

Interference and outage

Let $\hat{I}_{\rm sp}$ be the interference at the typical PU Rx stemming from a PPP of intensity λ_s outside the guard zone.

$$\mathcal{L}_{\hat{I}_{\rm sp}}(s) = \exp\left\{-\lambda_s\pi\Big(s^\delta\mathbb{E}_h(h^\delta\gamma(1-\delta, sh\rho^{-\alpha})) - D^2\mathbb{E}_h\big(1-\exp(-shD^{-\alpha})\big)\Big)\right\}.$$

We know that

$$\hat{I}_{\rm sp} \prec I_{\rm sp}$$

and thus

$$\mathbb{P}(\mathrm{SIR} > \theta) > \mathcal{L}_{I_{\rm pp}}(\theta r_p^\alpha)\cdot\mathcal{L}_{\hat{I}_{\rm sp}}(\theta r_p^\alpha)$$

(assuming $P=1$). Thus the additional outage caused by the presence of the CUs is at most $1-\mathcal{L}_{\hat{I}_{\rm sp}}(\theta r_p^\alpha)$.

Results

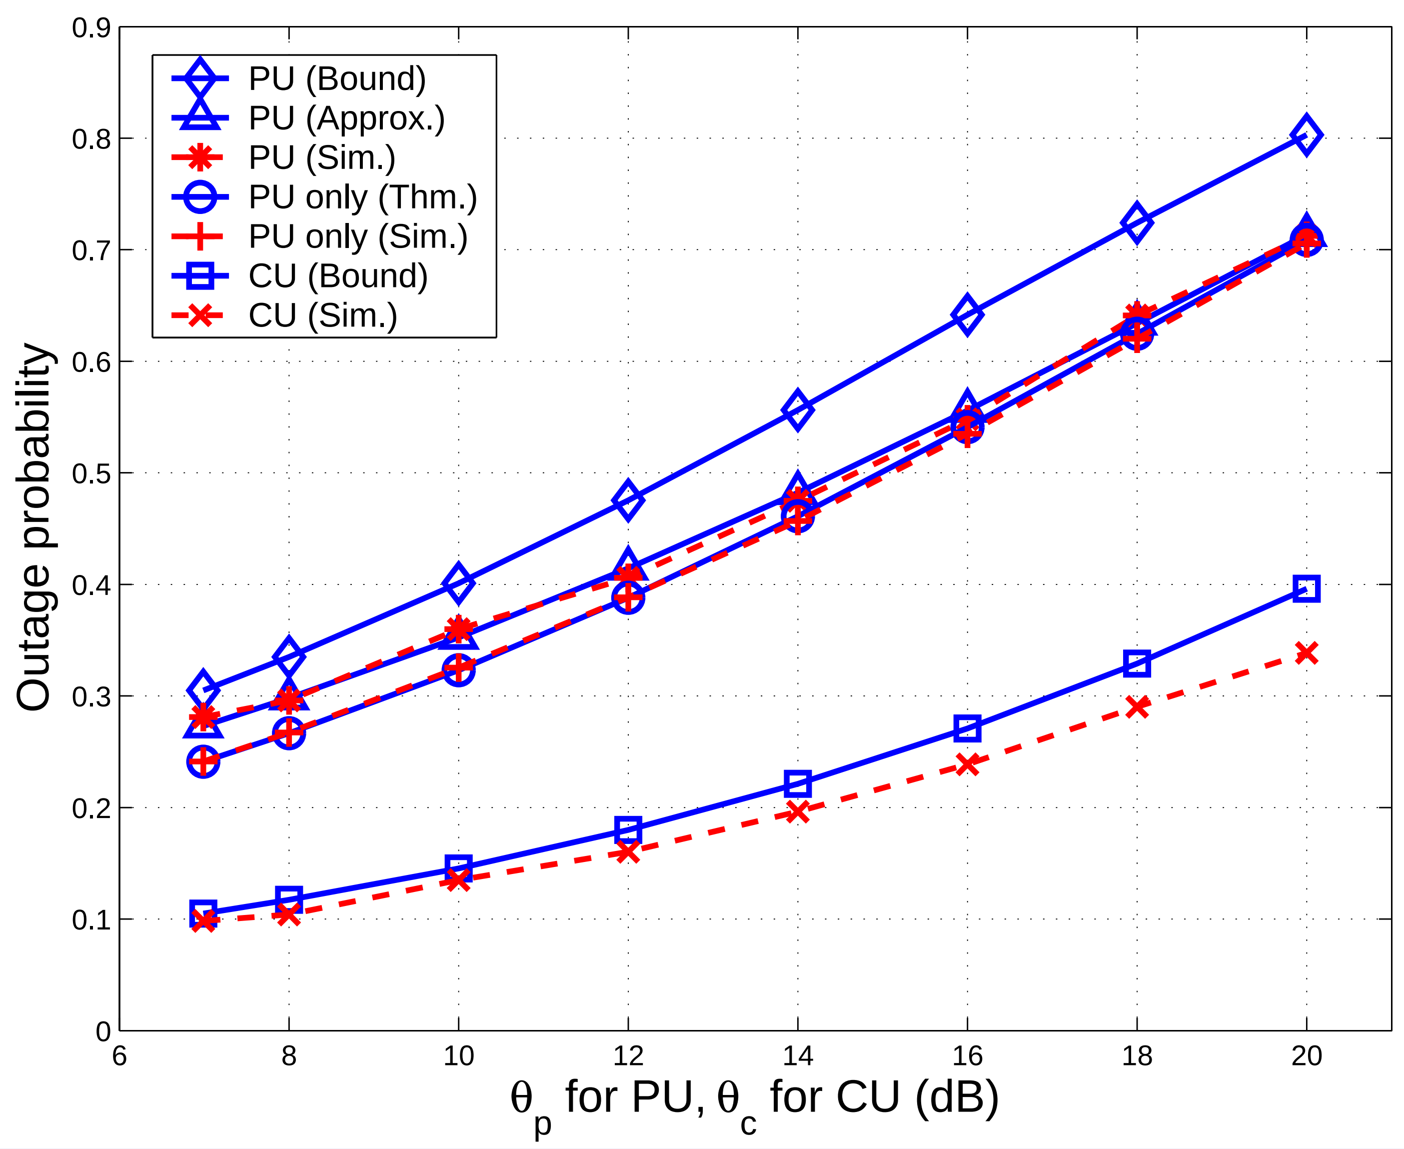

PU (Bound)
PU (Approx.)
PU (Sim.)
PU only (Thm.)
PU only (Sim.)
CU (Bound)
CU (Sim.)
Outage probability
θp for PU, θc for CU (dB)
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
6
8
10
12
14
16
18
20

Nearest-neighbor model: Setup

- PUs form a PPP of intensity λ_p.
- CUs form a PPP of intensity λ_s.
- PUs apply ALOHA with prob. p_p. Tx finds nearest node as its receiver.
- CUs cannot be active if within distance D_i of a primary receiver.
- Other CUs use ALOHA with prob. p_c and transmit to nearest neighbor.

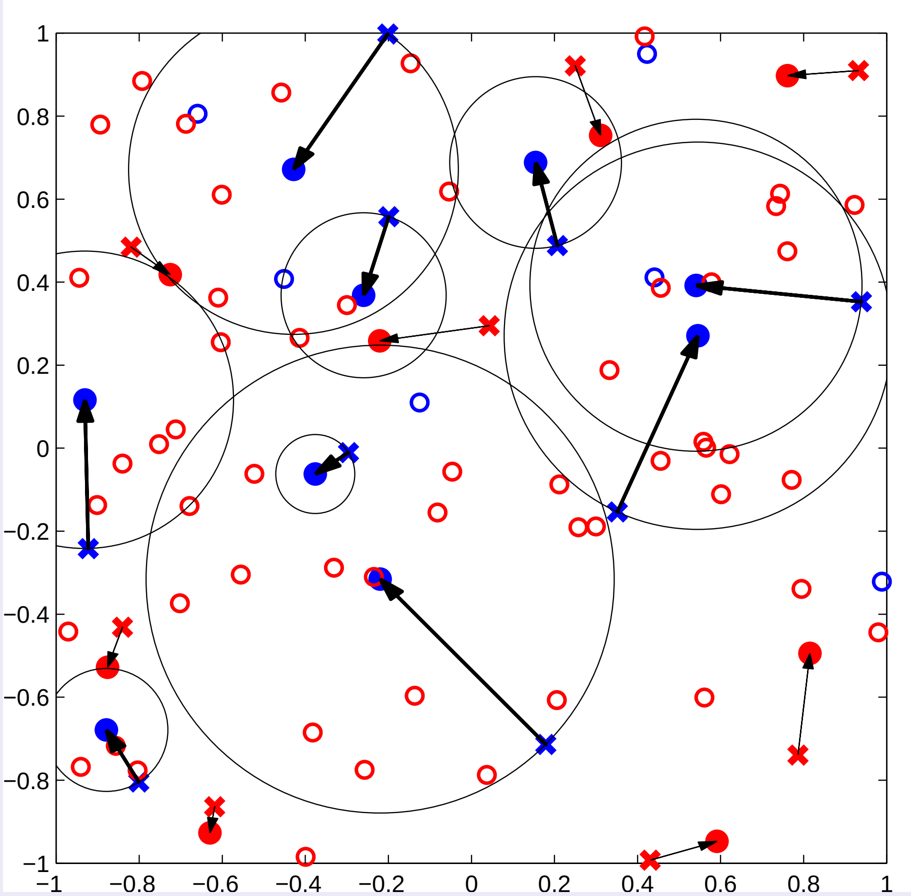

The guard zone D_i is a random variable with known distribution.

Interference and outage

From the probability generating functional for PPPs it follows that:
The intensity of secondary transmitters is $\exp(-p_p)$.
This is independent of λ_p, since a larger λ_p implies smaller guard zones. In fact, $\mathbb{E}(D^2) = \lambda_p^{-1}$.

Similar approximations as in the bipolar case lead to good bounds.

Exclusion regions around transmitters

- Exclusion regions around receivers can make sense if their locations are known (database).
- With a sensing-based approach, only transmitters can be detected.
- With guard zones around the primary transmitters, the primary receivers suffer from increased interference I_{sp}, as the effective guard zone radius reduces to $D - r_p$. I_{pp} and I_{pp} and I_{ss} remain the same, and I_{ps} decreases.
- If a receiver acknowledges packet reception, its presence can also be detected. A CU can match transmitter-receiver pairs and transmit concurrently with a PU transmitter if the PU receiver is on the other side.

The mutual nearest-neighbor model

- In the previous nearest-neighbor model, the receiver may not be able to acknowledge, since there may be another node nearby.
- To prevent ACK collision, the mutual-nearest-neighbor transmission protocol may be applied. Here, nodes form nearest-neighbor pairs if they are mutual nearest neighbors. The fraction of nodes thus paired is 62%.
- The resulting point process of transmitters thus has maximum density 31%, and it is more regular than a PPP.

Outlook

Ongoing and future work

- Software-defined radio
- (Collaborative) detection and learning
- Standardization (IEEE 802.22)
- Economic aspects (spectrum leasing, pricing) and game theory
- Legal aspects: how to detect and punish cheaters? The "hit and run" radio problem.
- Database issues
- Ruling on TV white space
- Network protocols, in particular for CUs (including Tx-Rx coordination)

Concluding remarks

- Cognitive radio enables the transition from "spectrostatics" to "spectrodynamics".
- Space is the critical resource; the network geometry greatly affects the interference and thus the performance of cognitive networks.
- Need to consider all potential CUs, not just one.
- Stochastic geometry permits the analysis of interference and outages in many scenarios where nodes are randomly distributed.
- The problem of white spaces is not a black and white problem. Wireless transmissions offer many gray areas, especially if advanced receiver technologies are available.
- "FCC rules are like Maxwell's equations"

Cognitive networks pose multi-faceted challenges: Technical, economic, legal, and policy issues.

Cognitive Radio Policy and Regulations

- U.S. National Broadband Plan (www.broadband.gov)
- Ofcom Statement on Cognitive Devices (stakeholders.ofcom.org.uk/binaries/consultations/cognitive/statement/statement.pdf)
- IEEE 802.22 WG on Enabling Rural Broadband Wireless Access Using Cognitive Radio Technology (www.ieee802.org/22/)
- Proceedings of the Dynamic Spectrum Access (DySPAN) conferences

Section Outline

1 Cognitive Networks

2 Cellular Networks
- Coverage and outage
- Relaying

3 Femtocells

4 Summary

Cellular Networks

Coverage

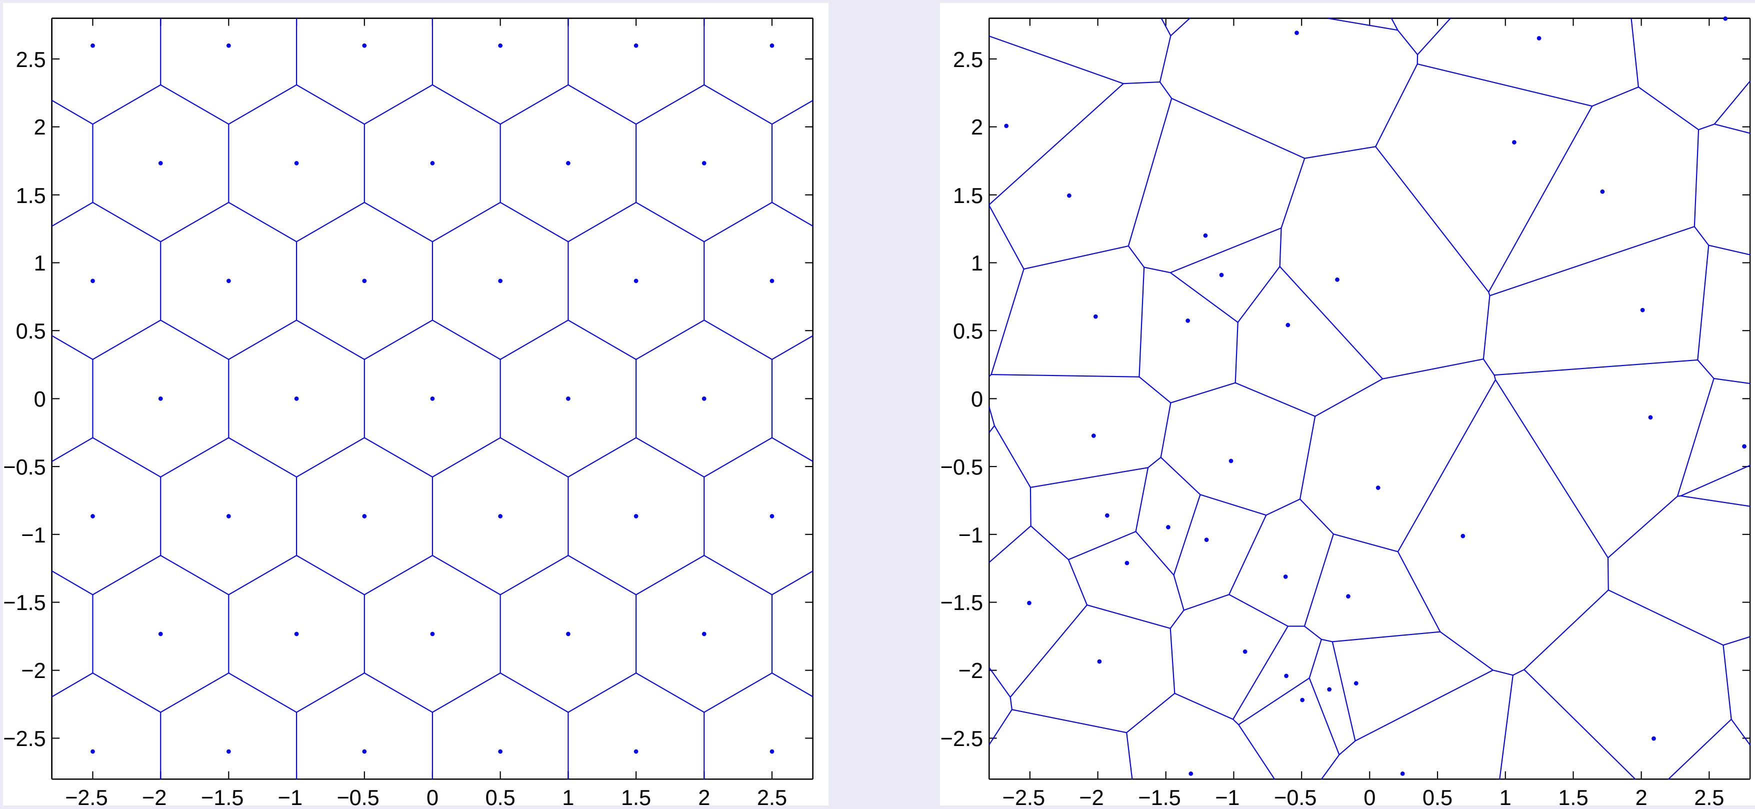

Voronoi diagrams for base stations arranged as triangular lattice (hexagonal cells) and as a Poisson process (Poisson Voronoi cells).

Definition (SINR cell)

For a point process $\Phi \subset \mathbb{R}^2$, the SINR cell of point $x \in \Phi$ is

$$C(x, \Phi) = \{y \in \mathbb{R}^2 \colon Pg(\|x - y\|) \geq \theta(I_\Phi(y) + W)\} ,$$

where g is the path loss law and

$$I_\Phi(y) = \sum_{z \in \Phi \setminus \{x\}} Pg(\|y - z\|)$$

is the interference.

Interpretation

The SINR cell of node x is comprised of all points, where the condition SINR $\geq \theta$ is met when x is the desired transmitter.

Fading can be included. The coherence length (or spatial correlation properties) of the fading process needs to be specified.

Coverage by SINR cells

Each SINR cell is a random closed set. For the standard path loss model and a PPP Φ, the volume of each cell is [BB09, Sec. 5.3]

$$|C(o,\Phi)| = \frac{1}{\lambda\theta^{2/\alpha}} \cdot \underbrace{\frac{\alpha}{2\Gamma(2/\alpha)\Gamma(1-2/\alpha)}}_{f(\alpha)}$$

The Poisson Voronoi cell has volume $1/\lambda$. If $\theta < 1$ (spread-spectrum), cells can overlap and be larger than $1/\lambda$. Consider two scenarios:

1. What would happen if there was no interference?
2. Let $\theta = 1$. What happens if $W = 0$ and $\alpha \to \infty$?

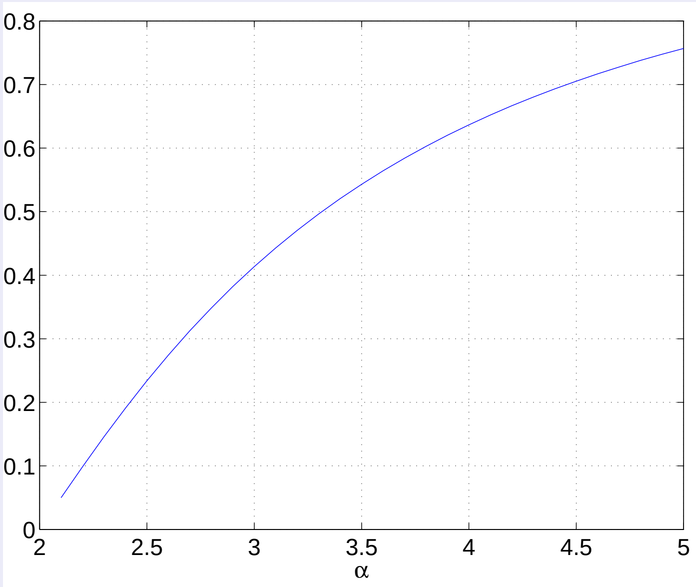

Function $f(\alpha)$

Coverage by SINR cells

1. In the interference-free scenario, the SNR cells are the grains of a Boolean model with radius $r = (P/(\theta W))^{1/\alpha}$.
2. In the noise-free case with $\alpha \to \infty$ (and $\theta = 1$), the SIR cells coincide with the Voronoi cells. Indeed, $\lim_{\alpha\to\infty} f(\alpha) = 1$.

In both cases, convergence can be rigorously proven.

The union of SINR cells forms a coverage process with dependent grains. Interesting questions are what fraction of the plane is covered, and how often each point in the plane is covered. Integral expressions can be derived for these problems [BB09, Sec. 7.5].

Why consider Poisson distributed base stations?

- Base station distributions may be irregular (different terrain, micro-BSs, femtocells)
- Analysis is simplified in some cases. Results are pessimistic compared to a regular arrangement.

Relaying

A two-hop downlink scheme

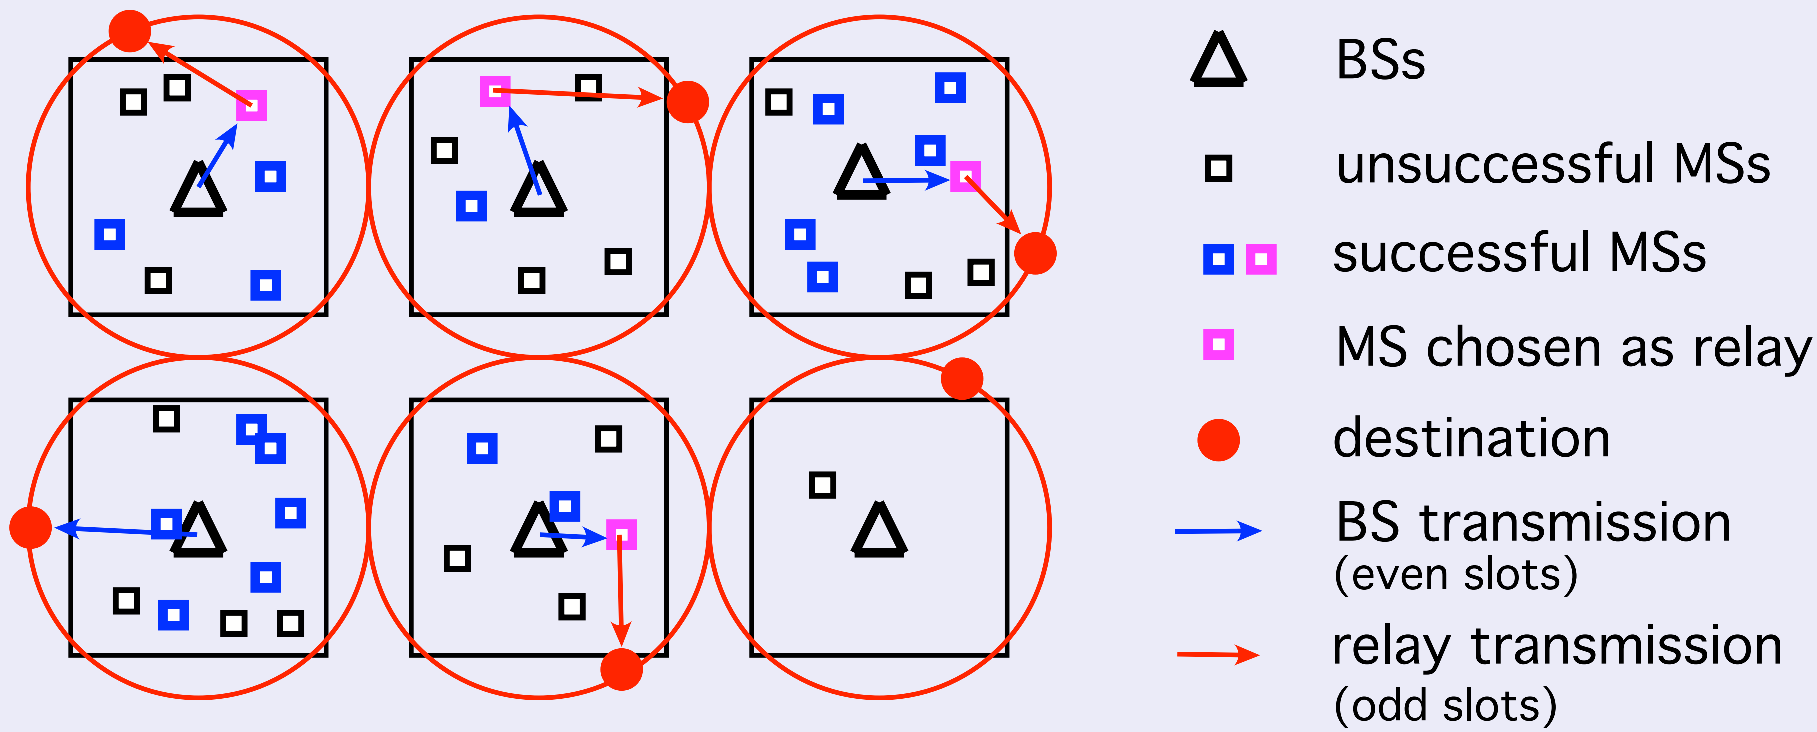

Relaying

A two-hop downlink scheme

Let

$$\Phi_b = \left\{ x \in \mathbb{Z}^2 \colon \frac{x}{\sqrt{\lambda_b}} \right\}$$

be the base station point process. The density is λ_b. The mobile users assisting the BS $x \in \Phi_b$ as relays form a PPP Φ_x of intensity $\lambda_x(y) = \eta(y - x)$ with mean

$$N = \int_{\mathbb{R}^2} \eta(x) \mathrm{d}x < \infty \,.$$

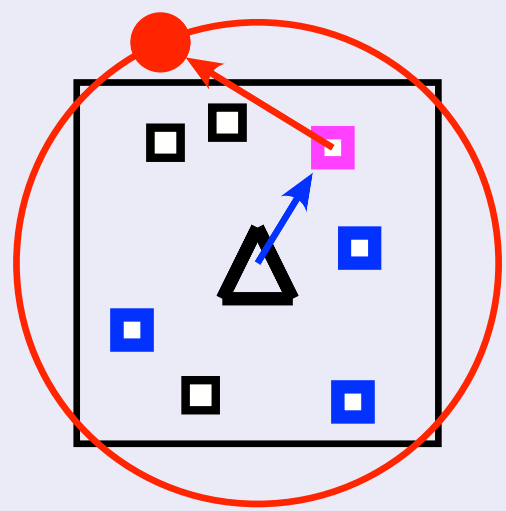

The probability that a cell is not empty (of mobile users) is $\mu = 1 - \exp(-N)$.

For example, we may choose $\eta(y) = \mathbf{1}_y([-1/2, 1/2]^2)$ and $\lambda_b = 1$, which leads to a square coverage area per BS.

A two-hop downlink scheme

Further assumptions:

- Independent fading between all pairs of transmitters and receivers and the standard SINR condition.
- BSs transmit in even, all at equal power, MSs transmit in odd slots, synchronized across all cells.
- The *destination* in cell x is an additional node $r(x)$ at distance $R = \|x - r(x)\|$.

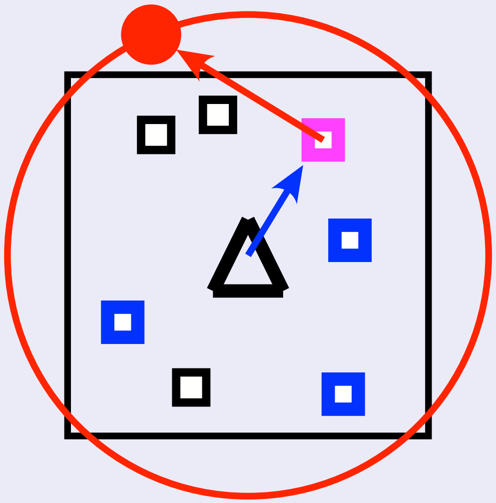

Operation:

- In the even slot, some relays will be able to decode the packet from their BS.
- A relay selection scheme selects a subset of these successful relays as transmitters in the next slot. We consider the method where the relay with the best channel to the destination transmits.

Analysis I

Let

$$\mathbf{1}(x \to y \mid \Phi)$$

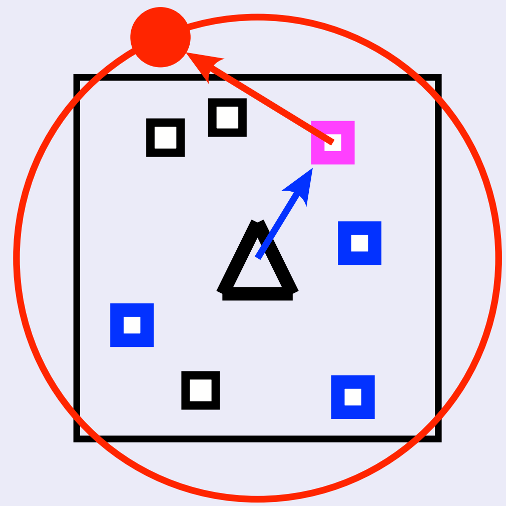

be the indicator that y can receive from x when the interferers are the nodes in Φ.

Consider the cell at o and let P_d be the probability that the BS can connect to the destination in a single hop:

$$P_d = \mathbb{E}\mathbf{1}(o \to r(o) \mid \Phi_b \setminus \{o\}).$$

Let $u_x \in \Phi_x$ be the relay selected in cell x and P_r the probability that the selected relay can connect to the destination:

$$P_r = \mathbb{E}\mathbf{1}(u_o \to r(o) \mid \Psi_r \setminus \{x\}),$$

where Ψ_r is the set of all selected relays for the second hop.

Analysis II

The success probability for the two-hop scheme is

$$P_s = 1 - (1 - P_d)(1 - P_r)\,.$$

We compare with a scheme where the base station transmits twice (time diversity). The gain is the ratio of the outages

$$G(\mathrm{SNR}, \lambda_b) = \frac{(1 - P_d)^2}{(1 - P_d)(1 - P_r)} = \frac{1 - P_d}{1 - P_r}\,.$$

The SNR is defined as $\mathrm{SNR} = PR^{-\alpha}/W$, where P is the total transmit power for both slots.

The diversity gain is

$$D(\lambda_b) = -\lim_{\mathrm{SNR}\to\infty} \frac{\log(1 - P_s)}{\log(\mathrm{SNR})}\,.$$

Analysis III

For the asymptotic evaluation, we tie the BS density λ_b and the SNR together by

$$\lambda_b = \mathrm{SNR}^{-\beta}, \quad \beta \geq 0\,.$$

For $\beta = 0$, the density is independent of β.
For $\beta < 2/\alpha$, the system is interference-limited.
For $\beta > 2/\alpha$, the system is noise-limited.
For the direct transmission,

$$P_d \sim \begin{cases} 1 - \theta\mathrm{SNR}^{-1} & \alpha\beta > 2 \text{ (noise)} \\ 1 - \theta(1 + cR^\alpha)\mathrm{SNR}^{-1} & \alpha\beta = 2 \text{ (noise/interference)} \\ 1 - \theta cR^\alpha \mathrm{SNR}^{-\alpha\beta/2} & 0 < \alpha\beta < 2 \text{ (interference)}\,. \end{cases}$$

Analysis IV

For the analysis of the two-hop performance, we assume the cell at o is not empty.

- The number of relays decoding the message from BS is $k = |\hat{\Phi}_o|$. k is Poisson.
- There is an outage if the best relay's (from the k) channel does not satisfy the SINR condition.
- Conditioned on k, the k relays are iid in the cell (but not uniform).

The resulting integrals need to be evaluated numerically.

The diversity order is

$$D(\mathrm{SNR}^{-\beta}) = \min\left\{1, \frac{\alpha\beta}{2}\right\} .$$

So there is no diversity gain in the two-hop system. This is due to the Poisson number of nodes. No loss as soon as $\alpha\beta > 2$ (noise-limited regime). But the outage gain $G = (1 - P_d)/(1 - P_r)$ can be 10 or larger.

Result

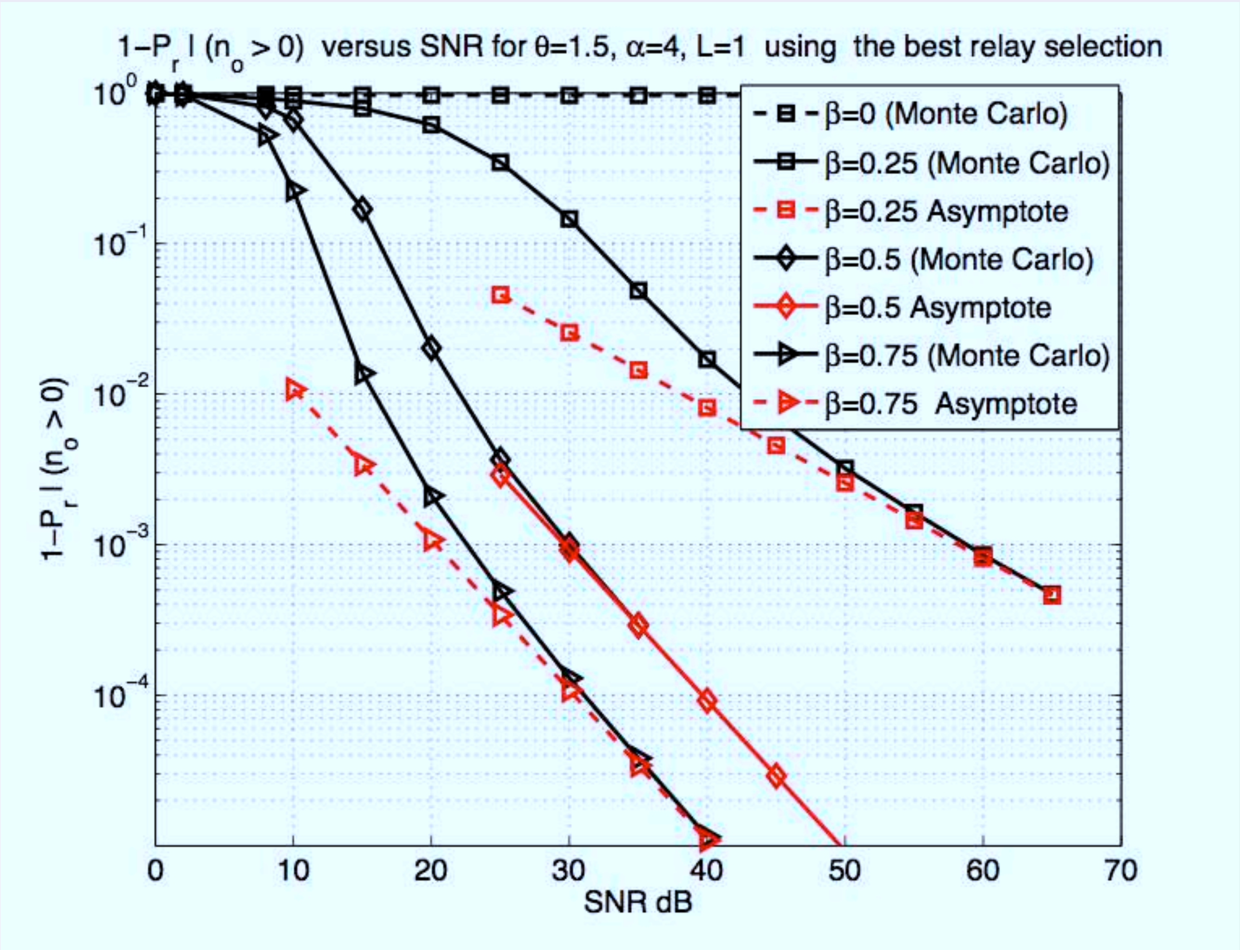

Outage probability vs. SNR for $\lambda_b = \text{SNR}^{-\beta}$ for $\alpha = 4$, $\eta(y) = 5\mathbf{1}_y([-0.5, 0.5]^2)$ and destination at $z = (0.5, 0.5)$.

Section Outline

1 Cognitive Networks

2 Cellular Networks

3 Femtocells

4 Summary

Femtocells

Architecture [CAG08]

- The capacity of a cellular system is increased by decreasing link distances. However, adding base stations is expensive.
- A cheaper alternative are femtocells, aka "home base stations", where mini-BSs are installed by home users for improved coverage.
- The device communicates with the cellular network by DSL or cable.

Architecture

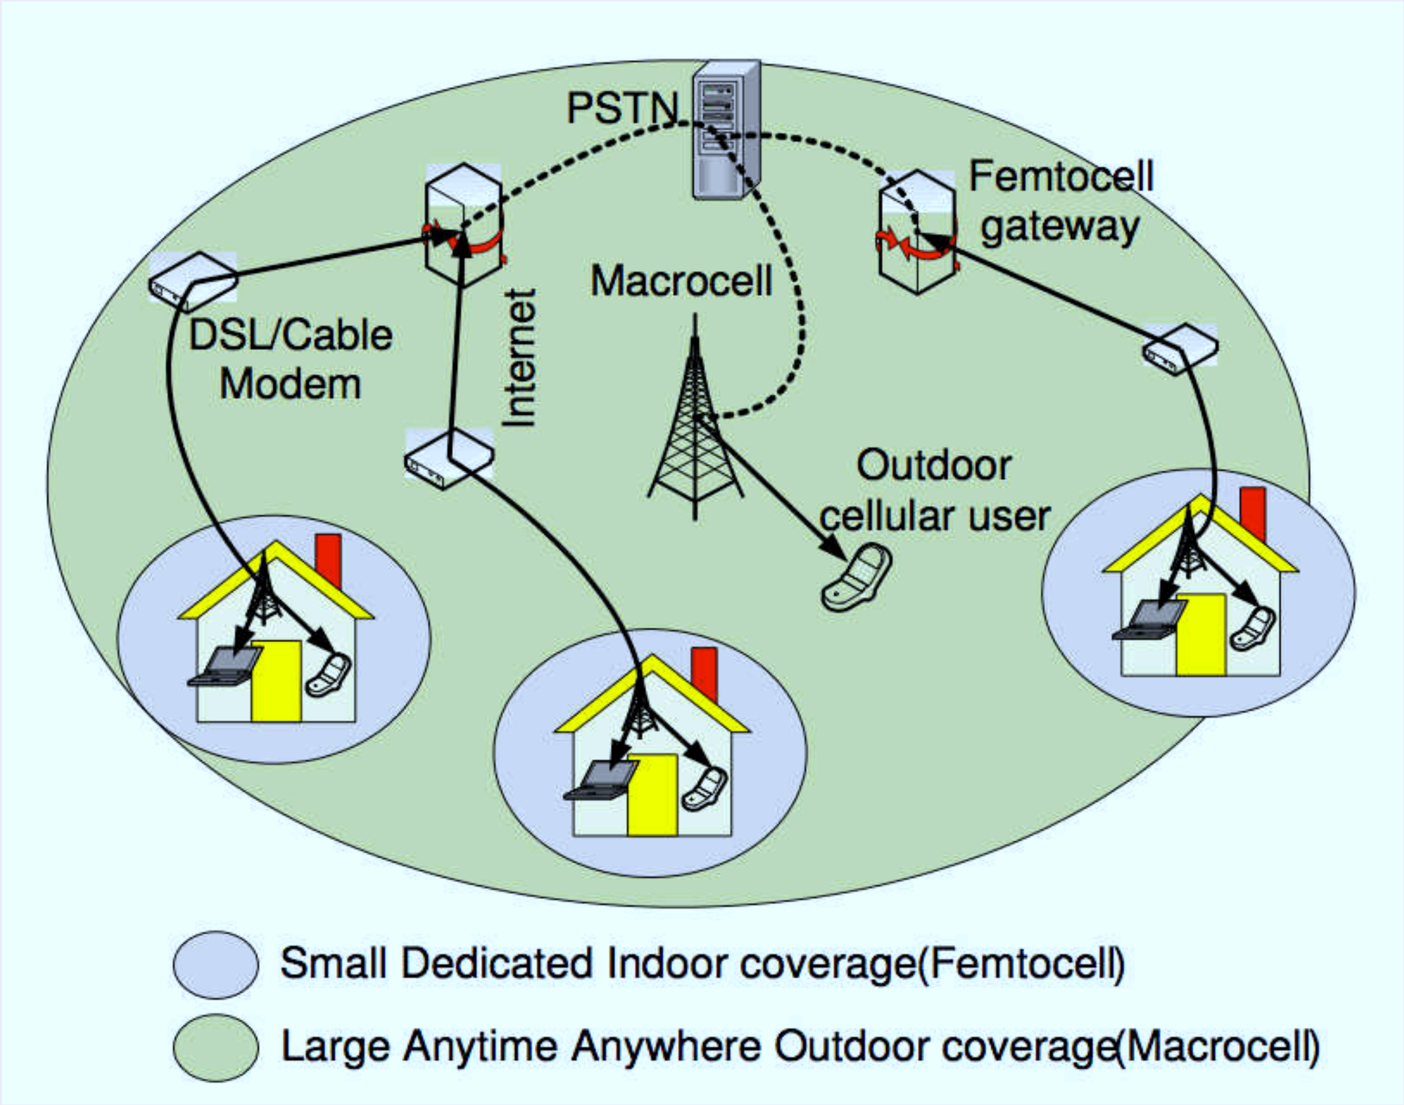

Femtocell architecture. Figure taken from PhD thesis "Coexistence in Femtocell-aided Cellular Architectures" by V. Chandrasekhar (UT Austin).

Challenges

- Interference (near-far effect):
  - Macrocell to femtocell. A user near a femtocell talking to the BS transmits at high power.
  - Femtocell to femtocell. Different femtocells may be located nearby.
  - Femtocell to macrocell. The femto-BS may cause excessive interference to a nearby user who is connected to a BS.
- Synchronization. Over the IP backhaul, the necessary $1\mu s$ timing accuracy is hard to achieve.
- QoS. How to guarantee low latency in the order of 15ms?

Analytical approach

Due to the irregular deployment of the femtocells, the locations of the femto-BSs needs to be modeled as random—as a point process. Different interference management and power control schemes can be compared in terms of their coverage and achievable rates.

Section Outline

1 Cognitive Networks

2 Cellular Networks

3 Femtocells

4 Summary

Lecture 7 Summary

Emerging Architectures

- All emerging wireless architectures add randomness and uncertainty to the system. Their analysis greatly benefits from the tools provided by stochastic geometry.
- Cognitive networks and femtocell-aided cellular networks are heterogeneous networks with different types of self- and cross-interference that need to be characterized.
- Relaying schemes give raise to non-Poisson point processes that require Palm theory and higher-order moments.

References I

F. Baccelli and B. Blaszczyszyn.
Stochastic Geometry and Wireless Networks.
Foundations and Trends in Networking. NOW, 2009.

V. Chandrasekhar, J. G. Andrews, and A. Gatherer.
Femtocell Networks: A Survey.
IEEE Communications Magazine, 46(9):59–67, September 2008.

C.-H. Lee and M. Haenggi.
Interference and Outage in Doubly Poisson Cognitive Networks.
In *2010 International Conference on Computer Communication Networks (ICCCN'10)*, Zurich, Switzerland, August 2010.
Available at http://www.nd.edu/~mhaenggi/pubs/icccn10b.pdf.

Analysis and Design of Wireless Networks

Lecture 8: Modeling and Managing Uncertainty

Martin Haenggi

INFORTE Educational Program
Oulu, Finland
Sep. 23-24, 2010

Overview

Contents of the Short Course

- Lecture 1: Introduction and a Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop Analysis of Poisson Networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Lecture 8 Overview

1. Noise vs. Interference
2. Mobility
3. Overall Summary and Conclusions

Section Outline

1. Noise vs. Interference
   - Introduction
   - Success probabilities
   - Balanced design
   - Rate control
   - Transport density
2. Mobility
3. Overall Summary and Conclusions

Noise vs. Interference

What is "interference-limited"?

Often, we read "In this paper, we assume that the network is interference-limited". What does this really mean?

- Is it that there is no noise, *i.e.*, $W \equiv 0$?
- Is it that $I(x) \gg W$ for all node locations x in the network? Since $I(x)$ varies tremendously over x and *over time*, we would have to say $\mathbb{E}I(x) \gg W$, $\forall x \in \Phi$. But where is the signal here?
- Is it that the link success probabilities are dominated by interference in the sense that $p_s^N \approx 1$, or $1 - p_s^N \ll 1 - p_s^I$? How does this depend on the link distance?
- Is it $\mathbb{E}(\mathrm{SIR}) \ll \mathrm{SNR}$?

Distances are random and vary greatly.

Power control

If there is no noise, only interference, we can reduce all power levels to arbitrarily small powers—without affecting the performance at all. But in doing so, the noise has to become relevant at some point.

We will get back to this later.

Success probability with noise and interference

For a PPP with Rayleigh fading and ALOHA with transmit probability p:

$$p_s(r) = \underbrace{\exp(-C_N r^\alpha)}_{p_s^N(r)} \cdot \underbrace{\exp(-C_I r^2)}_{p_s^I(r)},$$

where $C_N = \theta W/P$ and $C_I = \gamma p = \theta^{2/\alpha} C(\alpha) p$.

The noise and interference part are affected differently by the link distance!

Definition (Critical distance)

The critical distance ρ is given by $p_s^N(\rho) = p_s^I(\rho)$:

$$\rho \triangleq \left(\frac{C_I}{C_N}\right)^{\frac{1}{\alpha-2}} = \theta^{-1/\alpha} \left(\frac{C(\alpha)pP}{W}\right)^{\frac{1}{\alpha-2}}, \quad \alpha > 2\,.$$

Success probability at $r = \rho$

$p_s(\rho)$ is independent of θ. For $r < \rho$, we have $p_s^I < p_s^N$.

The success probabilities are given by

$$p_s^N(\rho) = p_s^I(\rho) = \exp\left(-(C(\alpha)p)^{\frac{\alpha}{\alpha-2}} \cdot (P/W)^{\frac{2}{\alpha-2}}\right) .$$

$$p_s(\rho) = p_s^N(\rho)^2 = p_s^I(\rho)^2$$

Also, we have

$$\rho^\alpha \propto \theta^{-1} .$$

With increasing θ (rate of transmission), the critical distance decreases.

Definition (Noise- and interference-limited links)

A link is noise-limited if $r \geqslant \rho$ and interference-limited if $r \leqslant \rho$.

Remarks

- This definition does not say anything about $p_s(r)$ or $p_s(\rho)$.
- Long links tend to be noise-limited.

Example ($\alpha = 4$)

$$\frac{C_I}{C_N} = \frac{\pi^2 P p}{2W\sqrt{\theta}}\,; \qquad \rho = \sqrt{\frac{\pi^2}{2}\frac{P}{W}\frac{p}{\sqrt{\theta}}}$$

Some numbers: For $P/W = 20\text{dB}$, $\theta = 10$, $p = 2\sqrt{\theta}/(10\pi^2) \approx 1/16$, we get $C_N = 1/10$, $C_I = 1$, and $\rho = \sqrt{10} \approx 3.16$.

Example (Short distances; $\alpha = 4$)

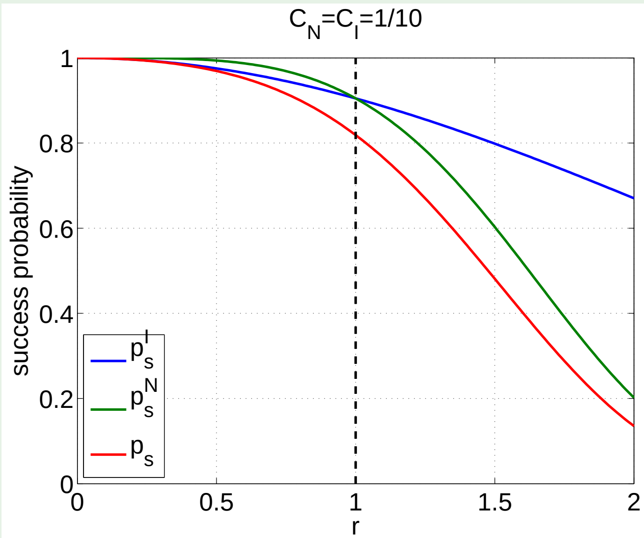

$\rho = 1$

$p_s(\rho) = 0.82$

Ok.

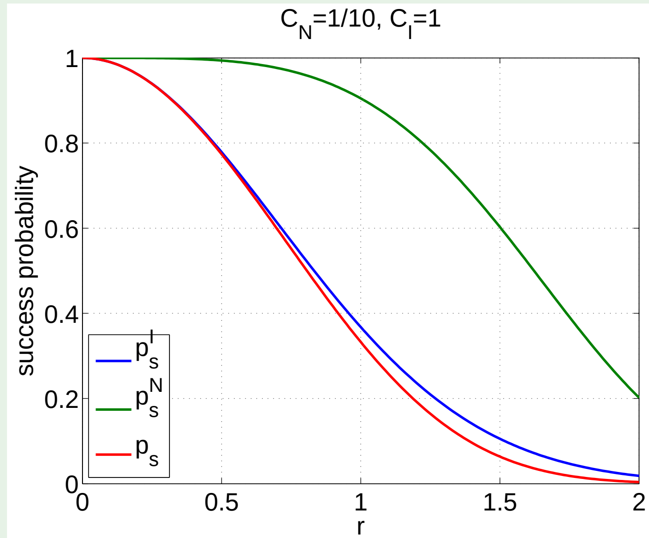

$\rho = \sqrt{10}$.

$p_s(\rho)$ abysmal.

Waste of power.

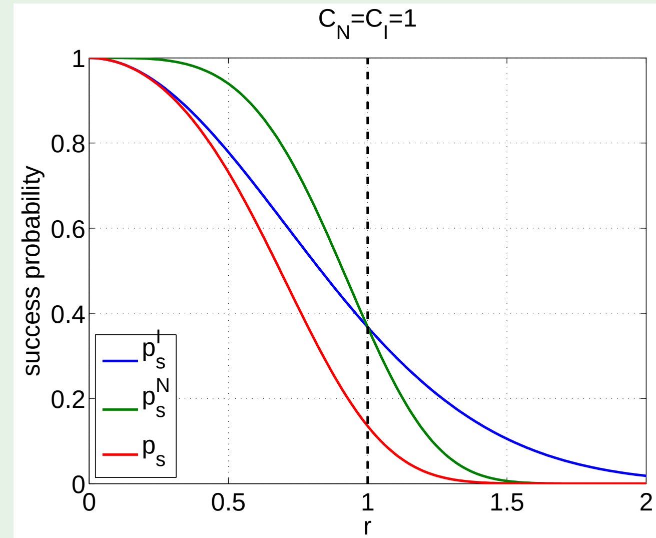

$\rho = 1$. 10× less power.

Better but $p_s(\rho)$ still low.

Discussion

- If there are no noise-limited links in an ad hoc network, power is wasted.
  In other words: We need the link distance to reach the critical distance ρ (while keeping $p_s(\rho)$ acceptably high).
- Long-range communication is noise-limited. In a connected large network (with limited power), there will always be noise-limited links.
- Noise- or interference-limitedness is not a network property, but a link property!

How to find the right balance between noise and interference?

Per-node power control changes the picture. However, channel inversion increases the interference (and may not be feasible even). Also, constant power is preferred from a PA standpoint.

Balanced design

Given:

- Mean link distance $\bar{R}$
- Target success probability $p_s(\bar{R}) = (1-\epsilon)^2$
- W, θ.

From $C_I \bar{R}^2 = -\log(1-\epsilon)$ and $C_N \bar{R}^\alpha = -\log(1-\epsilon)$ we have

$$C_N = \frac{-\log(1-\epsilon)}{\bar{R}^\alpha} \approx \frac{\epsilon}{\bar{R}^\alpha}; \quad C_I = \frac{-\log(1-\epsilon)}{\bar{R}^2} \approx \frac{\epsilon}{\bar{R}^2}$$

and

$$P = \frac{\bar{R}^\alpha \theta W}{\epsilon}; \qquad p = \frac{\epsilon}{\bar{R}^2 \theta^{2/\alpha} C(\alpha)}.$$

For this P, p, the design is balanced in the sense that some links will be interference- and some will be noise-limited since $C_I/C_N = \bar{R}^{\alpha-2}$ and $\rho = \bar{R}$.

Example ($\alpha = 4$)

Let $\bar{R} = 3$, $W = 1$, $\theta = 10$, and $(1-\epsilon)^2 = 0.9$.
Then $\epsilon = 1 - \sqrt{0.9}$ and

$$C_I = \frac{\epsilon}{\bar{R}^2} \approx 0.0059$$
$$C_N = \frac{\epsilon}{\bar{R}^4} \approx 0.00065$$
$$P \approx 15800$$
$$p \approx 0.00037\,.$$

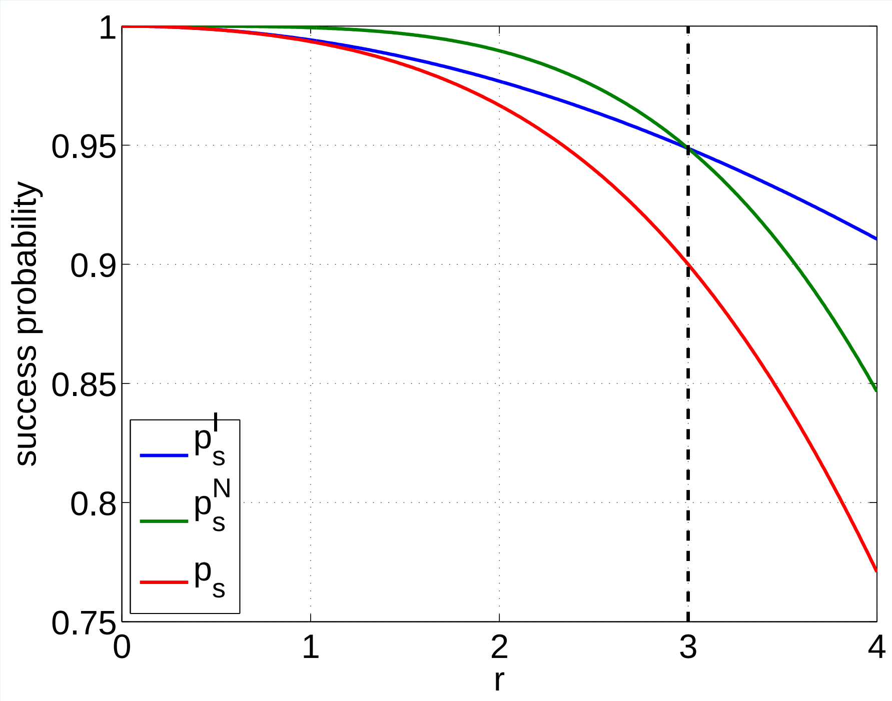

The mean SNR (at dist. $\bar{R}$) is about 23dB.
The mean SIR is 39dB. Mean dist. between transmitters is 26.
$p_s^I = p_s^N = 0.95$.

Dependence on θ

For all $\alpha > 2$, $p_s(\rho)$ is independent of θ, and $\rho \propto \theta^{-1/\alpha}$.

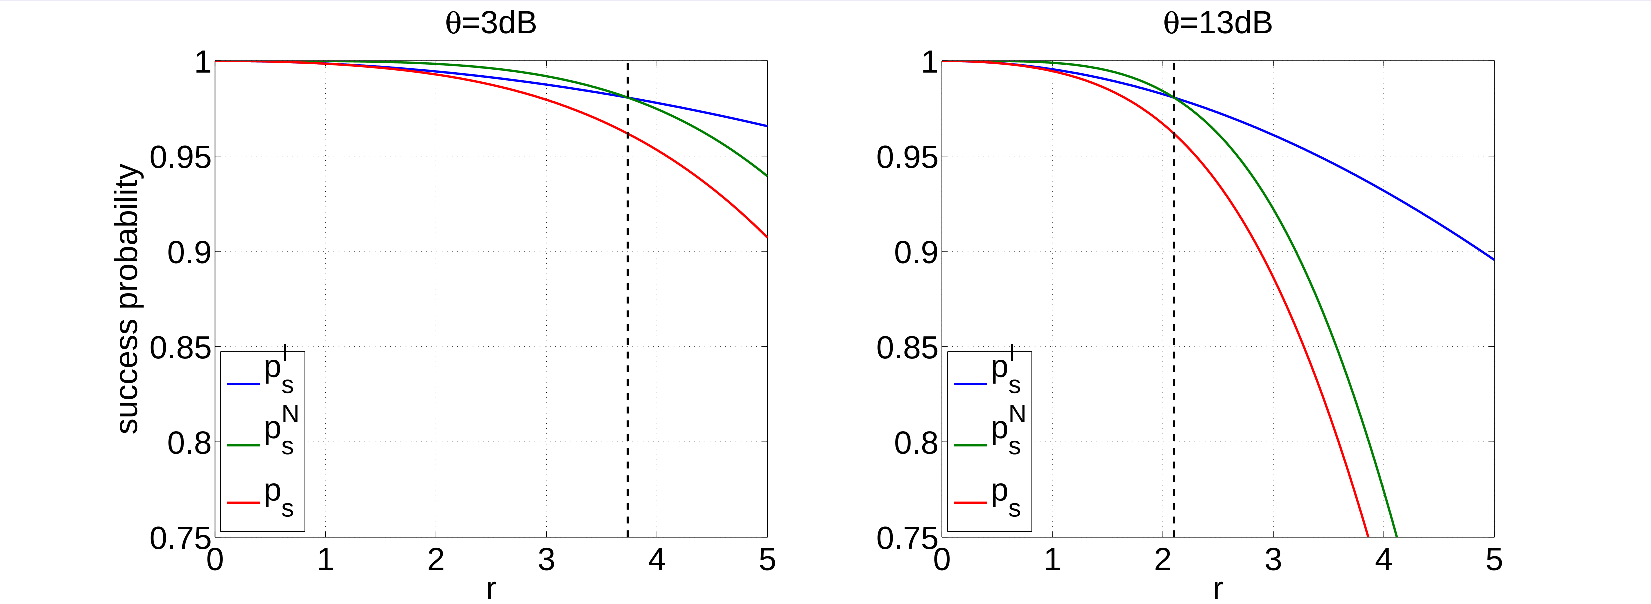

$$\alpha = 4,\ d = 2.\ P/W = 43\text{dB},\ p = 1/5000.$$

Maximization of "transport capacity": Maximize $\theta^{-1/\alpha}\log(1+\theta)$:

$$\theta_{\text{opt}}(\alpha) = \frac{-\alpha}{\mathcal{W}(-\alpha e^{-\alpha})},$$

where $\mathcal{W}$ is the Lambert W function. This is rate control.

Optimum rate for transport capacity

For $\alpha \in (2,5]$, the optimum rate of transmission

$$R_{T\mathrm{opt}}(\alpha) \approx \frac{5}{2} + (\alpha - 2)\log 5\,.$$

This is essentially linear (affine) in α.

Is this the best overall?

What is the most meaningful overall metric, and how to we maximize it? We need to include the density of transmission also.

When doing so, how interference-limited is the network?

Definition (Transport density)

Assuming the transmitter density is λ and links have distance R, the transport density is

$$T \triangleq p_s(R, \theta, P, \lambda) \log(1+\theta) R\lambda \,.$$

Optimizing the transport density

For the PPP with Rayleigh fading, $W \equiv 1$ and $\alpha = 4$,

$$\begin{aligned} T(R, \theta, P, \lambda) &= p_s^N(R, \theta, P) p_s^I(R, \theta, \lambda) \log(1+\theta)\lambda \\ &= \exp(-\theta R^4/P) \exp(-c\sqrt{\theta} R^2 \lambda) \log(1+\theta)\lambda \,, \quad c = \pi^2/2 \,. \end{aligned}$$

We could try to maximize this quantity by optimizing in the four-dimensional space (R, θ, P, λ). For $\alpha = 4$, we find the relationship

$$\lambda_{\mathrm{opt}} = \frac{2}{\pi^2 \sqrt{\theta} R^2} \,.$$

Optimizing the transport density (first attempt)

Now let $R \to 0$ while increasing $\lambda_{\text{opt}} \propto R^{-2}$. We have

$$p_s^N \to 1, \quad p_s^I \to e^{-1}, \quad \lambda \to \infty\,,$$

so the transport density grows without bounds at fixed θ, P (for the singular path loss law).

Does that mean the transport density is a useless metric?
Not quite—if we consider that we cannot expect to let $\lambda \to \infty$. The density of nodes is typically given and cannot be made arbitrarily large. Also, the path loss law breaks down at very small distances.

So: We fix $\lambda = 1$ and play with the ALOHA parameter p. This means that $R = 1$ is a lower bound on the distance, since the nearest-neighbor distance is $1/(2\sqrt{\lambda}) = 1/2$.

Optimizing the transport density (second attempt)

With $R = 1$, $\lambda = 1$ and ALOHA,

$$p_{\mathrm{opt}} = \frac{2}{\pi^2\sqrt{\theta}}\,,$$

and

$$T = \frac{2}{\pi^2 e\sqrt{\theta}}\exp(-\theta/P)\log(1+\theta)\,.$$

Letting $P \to \infty$,

$$T_\infty = \frac{2}{\pi^2 e}\frac{\log(1+\theta)}{\sqrt{\theta}}\,,$$

and

$$\theta_{\mathrm{opt}} = \exp(2+\mathcal{W}(-2/e^2)) - 1 \approx 3.912 \approx 5.93\text{ dB}\,.$$

With infinite power and θ_{opt},

$$T \to 0.06000\text{ nats/(s Hz m}^2)\,.$$

Results

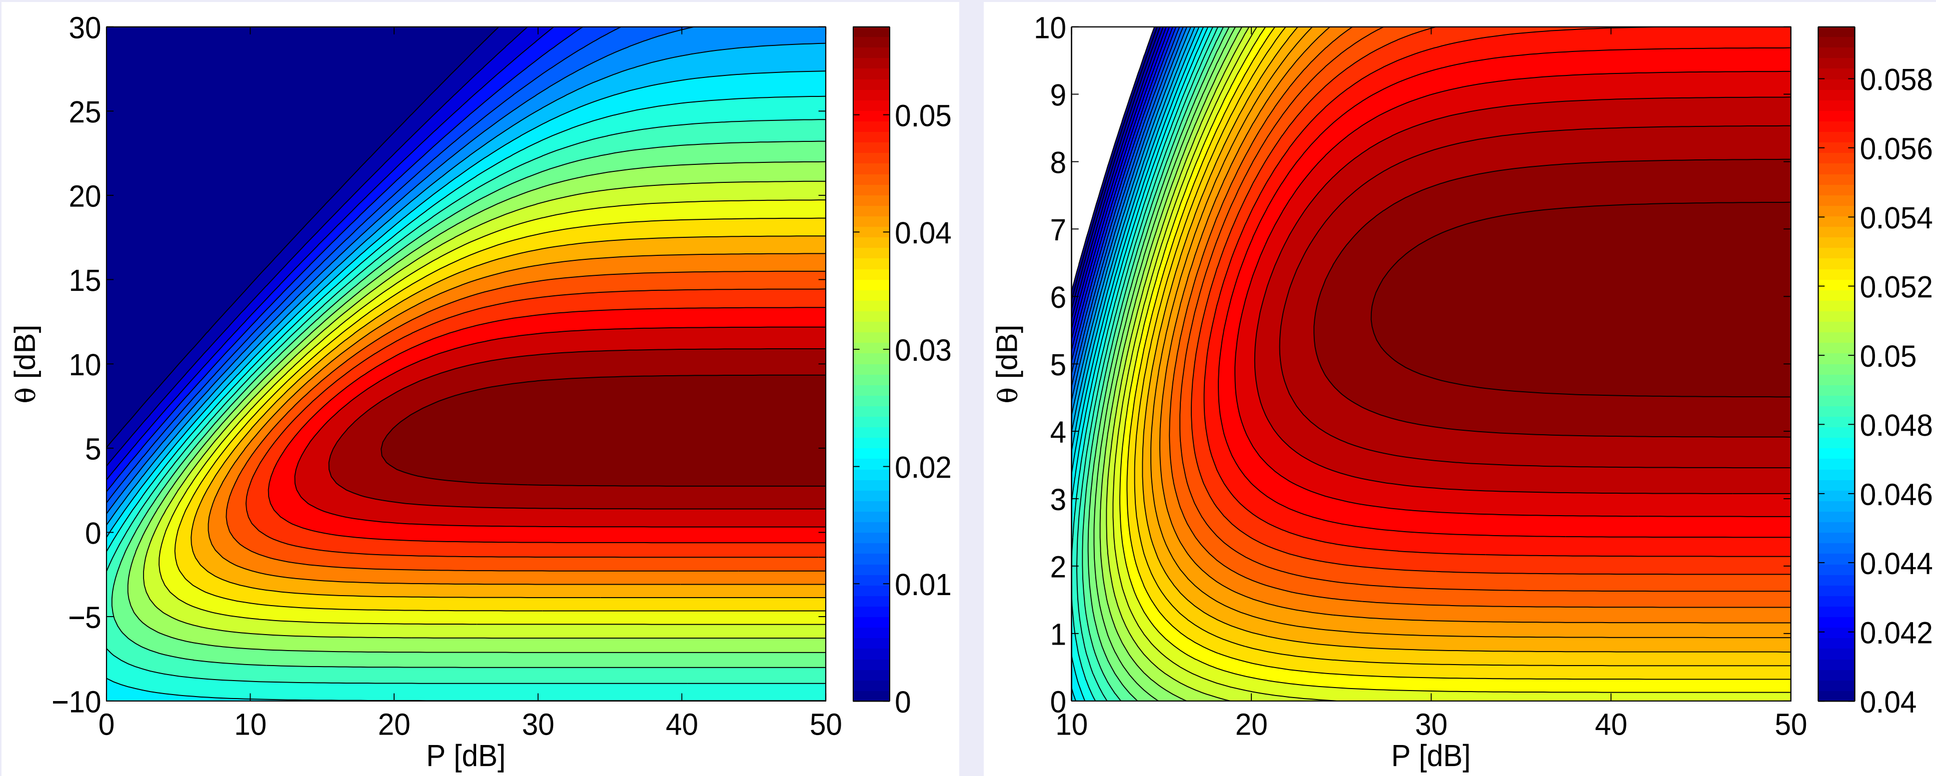

Left plot: The blue region in the top left corner is the power-limited regime. The bottom region is the rate-limited regime, and the top right part is "overrated". At $P = 10\theta$, 90% of the transport density is achieved ($\exp(-1/10) \approx 0.9$).

Right plot: Zoomed in on the range of optimum θ. $\theta = 5.93$ dB is asymptotically optimum. $\theta = 6$ dB and $P = 26$dB are sensible choices that achieve 0.0594, which is 99% of the asymptotic optimum.

Optimizing the transport density vs. balancing

This design procedure is balanced if $\theta = P$. Then $p_s = e^{-2} = 0.135$, which is abysmal.
For $P = 100\theta$, $p_s^N = 0.99$, and $p_s \approx p_s^I = e^{-1} = 0.368$. This is *interference-limited*.

As with ALOHA, maximizing a throughput metric may lead to poor reliability and, consequently, delay performance. Better reliability is achieved with less aggressive spatial reuse; balanced design permits a good trade-off between power and density of transmissions.

Our previous (balanced) example adapted to $R = 1$ yields $T = 0.007$, about 12% of the optimum. There is a high cost incurred by high reliability.

Transport density with balanced design for $\theta = 2, 10$

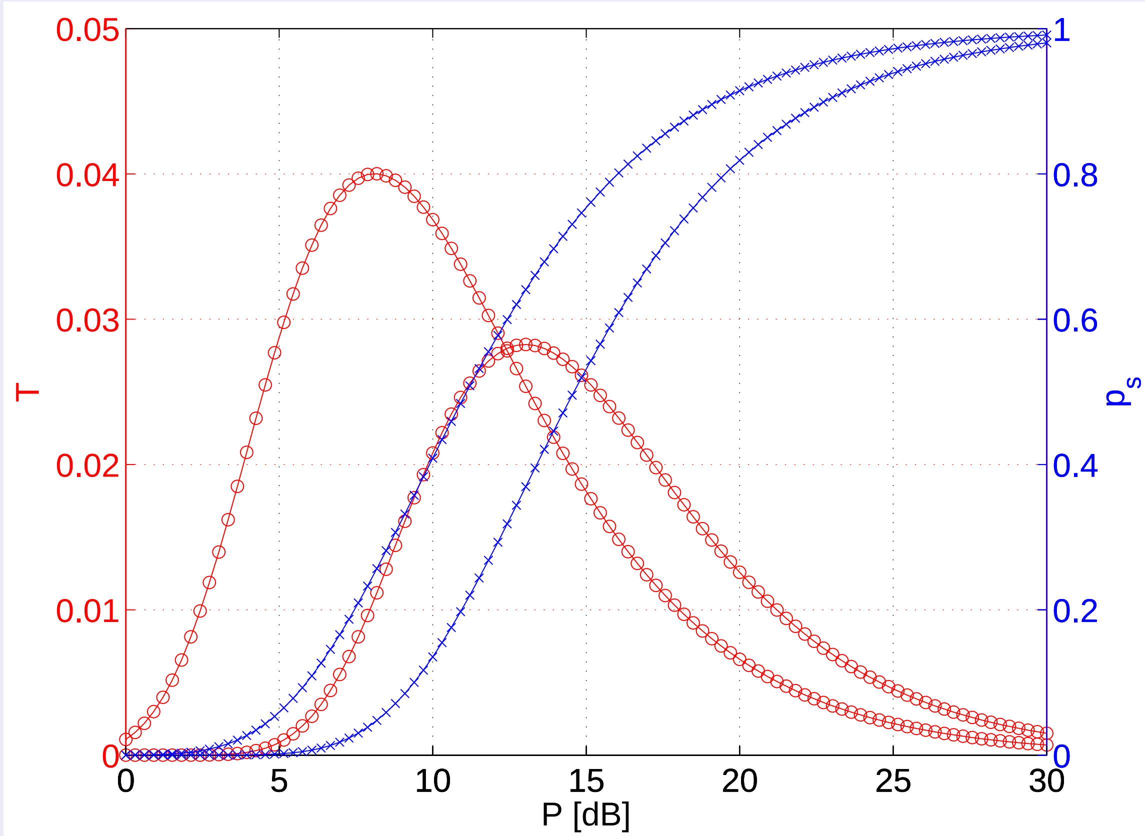

Parameters: $W = 1$, $R = 1$, $\alpha = 4$.

$\theta = 2$ achieves higher (near-optimum) performance. Again, maximizing T results in $p_s \approx e^{-1}$.

A definition of interference-limitedness

Let a be a power scaling factor, by which all transmit power levels are scaled. Let $M(a)$ be the (scalar) metric of interest, say throughput density or delay. Then the network is ϵ-interference-limited if

$$\lim_{a\to\infty}\left|\frac{M(1)-M(a)}{M(a)}\right| \leqslant \epsilon\,.$$

I.e., if the performance is already within a fraction ϵ of the performance achievable with infinite power.

From $M(1)$ to $M(a)$, only the transmit powers change, nothing else. So this definition is not applicable together with balancing, since with balancing, the metric may *decrease* with increasing power.

Section Outline

1 Noise vs. Interference

2 Mobility

3 Overall Summary and Conclusions

Mobility

Point processes and mobility

- Mobility adds a temporal dimension to the node locations.
- We have considered only two cases so far:
  - "infinite mobility": A new realization of the point process is drawn at each time slot.
  - No mobility: There is only one realization, which stays the same for all time.
- Realistic mobility models fall in between.
- If nodes move independently, the statistics of the uniform Poisson point process do not change, but mobility introduces temporal correlations.

Mobility and fading

If nodes have a home location and make excursions within a certain region around the location, a suitable model is

$$x_i(t) = x_i^H + w_i(t)\,, \quad \forall x_i \in \Phi\,,$$

where $w(t)$ is iid across nodes and time on some region $B \subset \mathbb{R}^2$.

The channel variations due to such a mobility model can also be interpreted as fading!
Hence many of the tools we have to deal with fading channels are applicable. Conversely, we may view fading as a geometric phenomenon [Hae08].

The reason why this is possible is that the mobility model is stationarity. For each i, $\mathbb{E}(x_i(t))$ is not a function of time. In many cases, $\mathbb{E}(x_i(t)) \equiv x_i^H$.

Other mobility models where $x_i(t) = x_i(t-1) + w_i(t)$ (Brownian motion, random walk) are more difficult to analyze.

Mobility and interference

The analogy between mobility and fading may be extended to include interference.
In the infinite mobility case, consider the interference I from the nearest node only, say at distance R. The SIR $\gamma = 1/I = R^\alpha$. Averaged over R, we have for the pdf of the SIR

$$f_\gamma(x) = \delta\lambda\pi x^{\delta-1}\exp(-\lambda\pi x^\delta)\,, \quad \delta = 2/\alpha\,,$$

which is a Weibull distribution. For $\delta = 1$, γ is exponentially distributed—quite exactly like the SNR in a Rayleigh fading environment!

This is further explored in [GH10].

It is also important to analyze the correlations in the interference structure. If nodes move according to some mobility model, how are the interference $I_x(t)$ and $I_x(t+1)$ correlated? Correlated interference makes outages correlated, which affects ARQ schemes and routing.

Mobility and connectivity

Consider a random walk, where nodes choose a random direction every m time slots and take steps of size s in that direction in each slot.
Putting n such nodes on the unit torus $[0,1]^2$ and connecting each pair of nodes within distance r, we obtain a dynamic random geometric graph.
Some of its properties were studied in [DMPG09]. In particular, they derived the expected lengths of the periods of connectivity and disconnectivity.

Let $\mu = n\exp(-\pi r^2 n)$ denote mean number of isolated nodes. If $\mu = \Theta(1)$, for each time, $\mathbb{P}(\text{conn.}) \sim \exp(-\mu)$ as $n \to \infty$.

If $r = \Theta(\sqrt{\log n/n})$ and $s = \Theta(1/\sqrt{n\log n})$, then

$$T_C \sim (1-\exp(-\mu(1-e^{-4srn/\pi}))^{-1}\,, \qquad T_D \sim (\exp(\mu)-1)\,T_C\,.$$

Section Outline

1 Noise vs. Interference

2 Mobility

3 Overall Summary and Conclusions
- Summary and concluding remarks
- Outlook

Looking back

Contents of the Short Course

- Lecture 1: Introduction and a Key Result
- Lecture 2: Throughput Analysis and Design Aspects
- Lecture 3: Random Graphs and Percolation Theory
- Lecture 4: Connectivity and Coverage
- Lecture 5: Multi-hop Analysis of Poisson Networks
- Lecture 6: Analysis of General Networks
- Lecture 7: Emerging Architectures
- Lecture 8: Modeling and Managing Uncertainty

Comparison of analytical approaches

scaling laws

very limited design insight

analysis of networks with fixed geometry

concrete results but no generality

stochastic geometry analysis of random networks

generality by spatial averaging *and* design insight

Stochastic geometry

- Stochastic geometry is a tool to **model and manage** uncertainty.
- There is tremendous uncertainty in modern wireless systems. The one due to locations is critical.
- Fading can be incorporated into the geometry.
- The MAC thins the point process of all nodes to a point process of transmitters. The result is a geometry that includes all sources of uncertainty.
- Dependences and correlations need to be modeled. They depend strongly on the relative time scales.
- Stochastic geometry and random graph theory result in random *geometric* graph theory.

The uncertainty cube

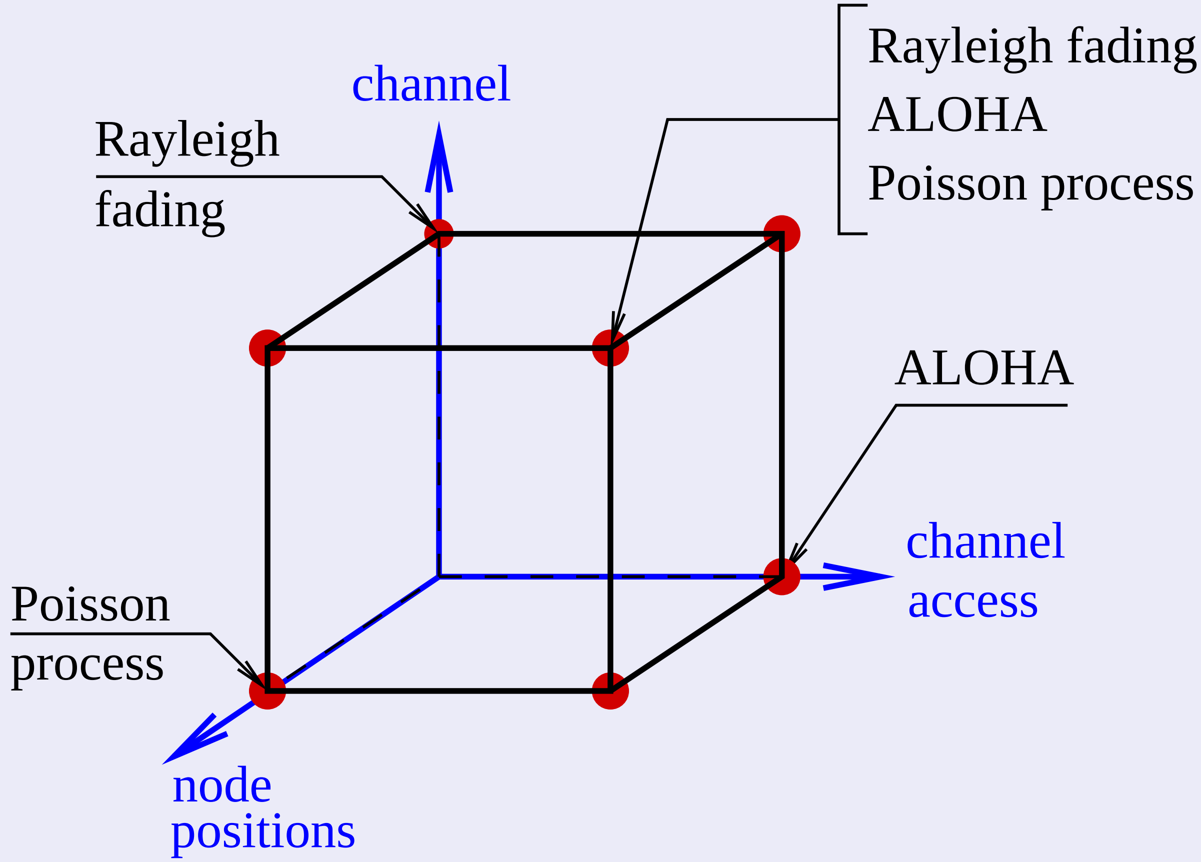

We may project all three axes onto one to obtain a geometry of path loss with fading.

Concluding remarks

- Space is the critical resource; the network geometry greatly affects the performance of wireless networks.
- The uniform PPP is great to work with, but it is time to consider other, often more realistic node distributions.
- Stochastic geometry, in particular Palm theory, offers the tools to analyze more general networks.
- Interference correlation affects efficiency of ARQ and routing but is a greatly under-investigated topic.
- The theory is applicable to wireless networks with infrastructure, such as multihop extensions of cellular systems.

Roadmap

The road to the analysis of general networks

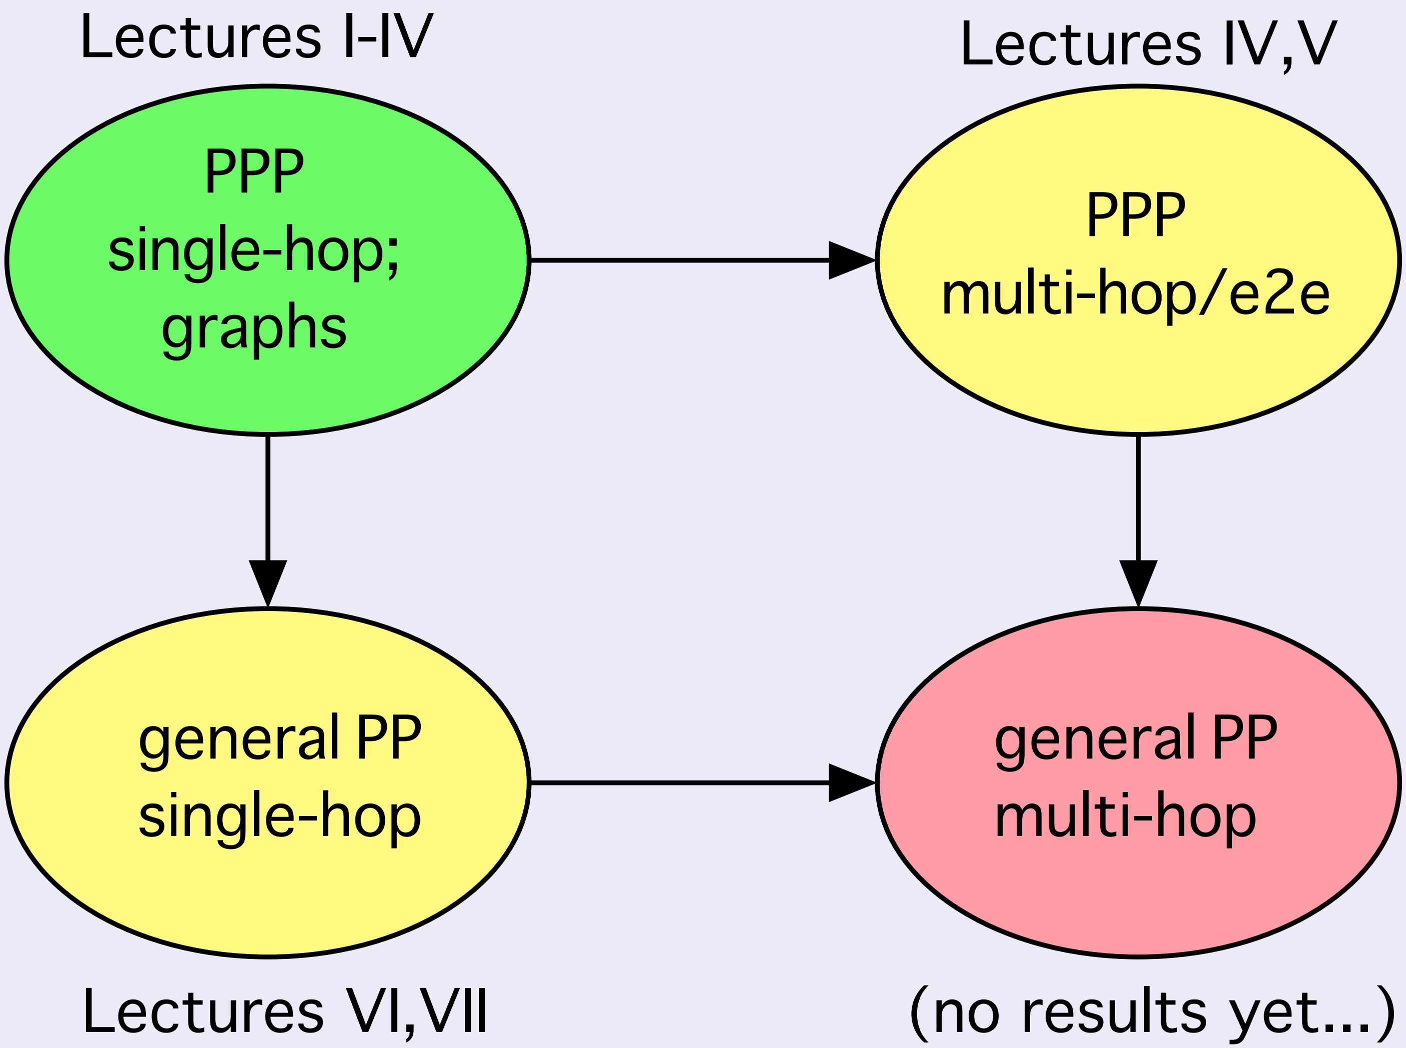

From green to red: Increasing **dependence** in time and space.

Outlook

Ongoing work

- Analysis of end-to-end throughput and delay (including queueing delays)
- From analysis to synthesis: Routing, MAC, power control?
- MIMO networks: How to use antennas?
- Cooperative communications in large networks: Interference cancellation, information-theoretic relaying, broadcasting and multiple-access
- Cognitive radio networks
- Femtocells
- Inclusion of secrecy constraints (secrecy graph, secrecy coverage)
- Mobility: Temporal coherence of the point process

References I

J. Diaz, D. Mitsche, and X. Perez-Gimenez.
Large Connectivity of Dynamic Random Geometric Graphs.
IEEE Transactions on Mobile Computing, 8(6):821–835, June 2009.

Z. Gong and M. Haenggi.
Mobility and Fading: Two Sides of the Same Coin.
In *IEEE Global Communications Conference (GLOBECOM'10)*, Miami, FL, December 2010.
Available at http://www.nd.edu/~mhaenggi/pubs/globecom10c.pdf .

M. Haenggi.
A Geometric Interpretation of Fading in Wireless Networks: Theory and Applications.
IEEE Trans. on Information Theory, 54(12):5500–5510, December 2008.
Available at http://www.nd.edu/~mhaenggi/pubs/tit08.pdf .

Books and Tutorial and Survey Articles I

Stochastic geometry

- D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and its Applications*. John Wiley & Sons, 1995. 2nd Ed.
- D. J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, Vol. I and II, Springer, 2003. 2nd Ed.
- J. F. C. Kingman, *Poisson Processes*, Oxford University Press, 1993.

Random Graphs and Percolation

- M. Penrose, *Random Geometric Graphs*, Oxford University Press, 2003.
- R. Meester and R. Roy, *Continuum Percolation*, Cambridge University Press, 1996.
- G. Grimmett, *Percolation*, Springer, 1999. 2nd Ed.
- B. Bollobás and O. Riordan, *Percolation*, Cambridge University Press, 2006.

Books and Tutorial and Survey Articles II

Applications to wireless networks

- F. Baccelli and B. Blaszczyszyn, *Stochastic Geometry and Wireless Networks*, Foundations and Trends in Networking, NOW Publishers, 2009.
- M. Haenggi and R. K. Ganti, *Interference in Large Wireless Networks*, Foundations and Trends in Networking, NOW Publishers, 2008. Available at http://www.nd.edu/~mhaenggi/pubs/now.pdf.
- M. Haenggi, J. G. Andrews, F. Baccelli, O. Dousse, and M. Franceschetti, *Stochastic Geometry and Random Graphs for the Analysis and Design of Wireless Networks*, IEEE Journal on Selected Areas in Communications, vol. 27, pp. 1029-1046, Sept. 2009, http://www.nd.edu/~mhaenggi/pubs/jsac09.pdf.
- S. Weber, J. G. Andrews, and N. Jindal, *An Overview of the Transmission Capacity of Wireless Networks*, http://arxiv.org/abs/0809.0016.

See also http://users.ece.utexas.edu/~jandrews/stochgeom/.

List of Symbols

$[n]$	The set $\{1, 2, \ldots, n\}$
$b(x, r)$	ball of radius r centered at x
Φ	Point process (and counting measure)
λ	Intensity of (stationary) point process
d	number of network dimensions
o	Origin in $\mathbb{R}^2$ or $\mathbb{R}^d$
$\mathcal{L}_X(s)$	Laplace transform $\mathbb{E}(\exp(-sX))$ of random variable X
α	Path loss exponent
δ	$\triangleq d/\alpha$
$g(r)$ or $g(x)$	Path loss law; typically $g(r) = r^{-\alpha}$
P	Transmit power
S	Received power
I	Interference power
W	Noise power
h	Fading coefficient
R	Transmission distance
p	Transmit probability in slotted ALOHA
θ	SINR threshold for successful reception
p_s	Success probability of a transmission
R_T	Transmission rate (spectral efficiency)