Stochastic geometry and wireless ad-hoc networks

from the coverage probability

to the asymptotic end-to-end delay on long routes

B. Błaszczyszyn (INRIA/ENS Paris, France)

based on joint works with F. Baccelli, O. Mirsadeghi and P. Mühlethaler

Spatial Network Models for Wireless Communications

Isaac Newton Institute, Cambridge, 6-9 April 2010

Ad-hoc Network

Network made of nodes “arbitrarily” repartitioned in some region, exchanging packets either transmitting or receiving them on a common frequency, use intermediary retransmissions by nodes lying on the path between the packet source node and its destination nodes.

Ad-hoc $\equiv$ Poisson

Nodes "arbitrarily" repartitioned $\equiv$ given network nodes are modeled as an instance of a Poisson point process (p.p.).

Ad-hoc = Poisson

Nodes "arbitrarily" repartitioned $\equiv$ given network nodes are modeled as an instance of a Poisson point process (p.p.).

Recall: Φ is a (homogeneous) Poisson p.p. of intensity λ (points per unit of surface) if:

- number of points of Φ in any set A, $\Phi(A)$, is Poisson random variable with mean λ times the surface of A.
- numbers of points of Φ in disjoint sets are independent random variables.

Medium Access Control (MAC)

The Medium Access Control (MAC) layer is a part of the data communication protocol organizing simultaneous packet transmissions in the network.

Aloha MAC = Independent Thinning

In our talk we will consider the, perhaps most simple, algorithm used in the MAC layer, called Aloha:

at each time slot (we will consider only slotted; i.e., discrete, time case), each potential transmitter independently tosses a coin with some bias p; it accesses the medium (transmits) if the outcome is heads and it delays its transmission otherwise.

Aloha MAC $=$ Independent Thinning

In our talk we will consider the, perhaps most simple, algorithm used in the MAC layer, called Aloha:

at each time slot (we will consider only slotted; i.e., discrete, time case), each potential transmitter independently tosses a coin with some bias p; it accesses the medium (transmits) if the outcome is heads and it delays its transmission otherwise.

Thus, (slotted) Aloha $\equiv$ (independent) thinning of the pattern of nodes willing to emit.

Aloha MAC $\equiv$ Independent Thinning

In our talk we will consider the, perhaps most simple, algorithm used in the MAC layer, called Aloha:

at each time slot (we will consider only slotted; i.e., discrete, time case), each potential transmitter independently tosses a coin with some bias p; it accesses the medium (transmits) if the outcome is heads and it delays its transmission otherwise.

Thus, (slotted) Aloha $\equiv$ (independent) thinning of the pattern of nodes willing to emit.

Thinning is a nice operation on a p.p.

In particular, thinning of Poisson p.p. of intensity λ leads to Poisson p.p. of intensity $p\lambda$.

Tuning Aloha Parameter p

In Aloha algorithm it is important to tune the value of the Medium Access Probability (MAP) p, so as to realize a compromise between two contradicting types of wishes:

- a "social one" to have as many concurrent transmissions as possible in the network and
- an "individual one" to have high chances that authorized transmissions be successful and/or efficient.

Tuning Aloha Parameter p

In Aloha algorithm it is important to tune the value of the Medium Access Probability (MAP) p, so as to realize a compromise between two contradicting types of wishes:

- a "social one" to have as many concurrent transmissions as possible in the network and
- an "individual one" to have high chances that authorized transmissions be successful and/or efficient.

The contradiction between these two wishes stems from the fact that the very nature of the "medium" in which the transmissions take place (Ethernet cable or electromagnetic field in the case of wireless communications) imposes some constraints on the maximal number and configuration of successful concurrent transmissions.

Signal to Interference Ratio (SIR)

A given transmission is successful if the power of the received signal is sufficiently large with respect to the interference and possibly some extra noise, where

interference is the sum of the powers of signals received from all other concurrent transmissions.

Signal to Interference Ratio (SIR)

A given transmission is successful if the power of the received signal is sufficiently large with respect to the interference and possibly some extra noise, where

interference is the sum of the powers of signals received from all other concurrent transmissions.

Interference created at $\boldsymbol{y}$ by transmissions of $\boldsymbol{\Phi} \equiv$ Shot-Noise (SN) $\boldsymbol{I(y) = \sum_{X \in \Phi} 1/l(|X - y|)}$, where $\boldsymbol{l(r)}$ is the power attenuation (path-loss) function on the distance $\boldsymbol{r}$.

Poisson Shot-Noise

Fact: If Φ is homogeneous Poisson p.p. than the Laplace transform (LT) $\mathcal{L}_I$ of the SN $I(y)$ is

$$\mathcal{L}_I(s) = \exp\Big[-2\lambda\pi \int_0^\infty r(1 - e^{1/l(r)})\, dr\Big].$$

Can be extended to joint LT of vectors $(I(y_1), \ldots, I(y_2))$.

Our Setting

In the remaining part of this talk we will show some simple (?) models and results regarding ad-hoc networks assuming

- Poisson repartition of nodes on the plane,
- Shot-Noise interference,
- Aloha MAC.

Related Works

The are now quite many works on various wireless communications problems using the stochastic geometry setting I have just mentioned (say Poisson p.p. network $+$ Shot-Noise interference).

Related Works

The are now quite many works on various wireless communications problems using the stochastic geometry setting I have just mentioned (say Poisson p.p. network $+$ Shot-Noise interference).

Among them, most related to what I will be talking about are by: J. Andrews, O. Dousse, M. Franceschetti, M. Haenggi, Ph. Jacquet, M. Kountouris, P. Thiran, E. Yeh, and many others ...

Related Works

The are now quite many works on various wireless communications problems using the stochastic geometry setting I have just mentioned (say Poisson p.p. network $+$ Shot-Noise interference).

Among them, most related to what I will be talking about are by: J. Andrews, O. Dousse, M. Franceschetti, M. Haenggi, Ph. Jacquet, M. Kountouris, P. Thiran, E. Yeh, and many others ...

In a broader sense, many outstanding theoreticians of stochastic geometry, random graphs, percolation theory were and are also interested in communication technology problems ...

I will not be able to pay tribute to the work they have done ...

Outline

- COVERAGE PROBABILITIES
- LOCAL DELAYS
- END-TO-END DELAYS ON LONG ROUTS

Outline

- **COVERAGE PROBABILITIES**
- **LOCAL DELAYS** ← A phase transition
- **END-TO-END DELAYS ON LONG ROUTS**

Outline

- **COVERAGE PROBABILITIES**
- **LOCAL DELAYS** $\leftarrow$ A phase transition
- **END-TO-END DELAYS ON LONG ROUTS** $\leftarrow$ A first passage percolation problem

COVERAGE
(or SUCCESSFUL TRANSMISSION) PROBABILITY IN A SPATIAL ALOHA MODEL

BASIC BIPOLAR AD-HOC NETWORK MODEL WITH ALOHA

Model

Independently marked Poisson point process (p.p.) $\widetilde{\Phi} = \{(X_i, e_i, y_i, \mathbf{F}_i)\}$, where

1. $\Phi = \{X_i\}$ denotes the locations of the nodes (the potential transmitters); Φ is always assumed Poisson with positive and finite intensity λ;

Model

Independently marked Poisson point process (p.p.) $\widetilde{\Phi} = \{(X_i, e_i, y_i, \mathbf{F}_i)\}$, where

1. $\Phi = \{X_i\}$ denotes the locations of the nodes (the potential transmitters); Φ is always assumed Poisson with positive and finite intensity λ;
2. $\{e_i\}$ is the MAC indicator of node i; ($e_i = 1$ if node i is allowed to transmit and 0 otherwise).
   Aloha principle:The random variables e_i are i.i.d. and independent of everything else, with $\mathsf{P}(e_i = 1) = p$ (p is the MAP).

Model

Independently marked Poisson point process (p.p.) $\widetilde{\Phi} = \{(X_i, e_i, y_i, \mathbf{F}_i)\}$, where

1. $\Phi = \{X_i\}$ denotes the locations of the nodes (the potential transmitters); Φ is always assumed Poisson with positive and finite intensity λ;
2. $\{e_i\}$ is the MAC indicator of node i; ($e_i = 1$ if node i is allowed to transmit and 0 otherwise).
   Aloha principle:The random variables e_i are i.i.d. and independent of everything else, with $\mathsf{P}(e_i = 1) = p$ (p is the MAP).
   Consequence of Aloha: the set of nodes that transmit $\Phi^1 = \{X_i : e_i = 1\}$ is a Poisson p.p. with intensity $\lambda_1 = \lambda p$ (as an independent thinning of Φ).

Receivers in Bipolar Model

3. $\{y_i\}$ denotes the location of the receiver for node X_i (we assume here that no two transmitters have the same receiver). We assume that $\{X_i - y_i\}$ are i.i.d random vectors with $|X_i - y_i| = r$; i.e. each receiver is at distance r from its transmitter.

Receivers in Bipolar Model

3. $\{y_i\}$ denotes the location of the receiver for node X_i (we assume here that no two transmitters have the same receiver). We assume that $\{X_i - y_i\}$ are i.i.d random vectors with $|X_i - y_i| = r$; i.e. each receiver is at distance r from its transmitter.

This is an (acceptable at this stage) simplification.
Later in this talk we will show extensions.

Wireless Channel Conditions — Fading

4. $\{\mathbf{F}_i = (F_i^j : j)\}$ where F_i^j denotes the virtual power emitted by node i (provided $e_i = 1$) towards receiver y_j; by this we understand the product of the (effective) power of transmitter i and of the random fading from this node to receiver y_j.

Wireless Channel Conditions — Fading

4. $\{\mathbf{F}_i = (F_i^j : j)\}$ where F_i^j denotes the virtual power emitted by node i (provided $e_i = 1$) towards receiver y_j; by this we understand the product of the (effective) power of transmitter i and of the random fading from this node to receiver y_j.

The random vectors $\{\mathbf{F}_i\}$ are assumed to be i.i.d. and the components (F_i^j, j) are assumed to be i.i.d. as a generic r.v. denoted by F with mean $1/\mu$ assumed finite.

Wireless Channel Conditions — Fading

4. $\{\mathbf{F}_i = (F_i^j : j)\}$ where F_i^j denotes the virtual power emitted by node i (provided $e_i = 1$) towards receiver y_j; by this we understand the product of the (effective) power of transmitter i and of the random fading from this node to receiver y_j.

The random vectors $\{\mathbf{F}_i\}$ are assumed to be i.i.d. and the components (F_i^j, j) are assumed to be i.i.d. as a generic r.v. denoted by F with mean $1/\mu$ assumed finite.
A spacial important case consists in assuming constant emitted power and Rayleigh fading which implies exponential F.

Omnidirectional Path-loss

Select some omnidirectional path-loss (OPL) model $l(\cdot)$. The receiver of node i receives the transmitter located at node j with a power equal to $F_i^j / l(|X_j - y_i|)$, where $|\cdot|$ denotes the Euclidean distance on the plane.

Omnidirectional Path-loss

Select some omnidirectional path-loss (OPL) model $l(\cdot)$. The receiver of node i receives the transmitter located at node j with a power equal to $F_i^j/l(|X_j - y_i|)$, where $|\cdot|$ denotes the Euclidean distance on the plane.

An important special case consists in taking

(1) $$l(u) = (Au)^\beta \quad \text{for } A > 0 \text{ and } \beta > 2,$$

which we call in what follows OPL 3.

Omnidirectional Path-loss

Select some omnidirectional path-loss (OPL) model $l(\cdot)$. The receiver of node i receives the transmitter located at node j with a power equal to $F_i^j / l(|X_j - y_i|)$, where $|\cdot|$ denotes the Euclidean distance on the plane.

An important special case consists in taking

(1) $$l(u) = (Au)^{\beta} \quad \text{for } A > 0 \text{ and } \beta > 2,$$

which we call in what follows OPL 3.

Note that $1/l(u)$ has a pole at $u = 0$, and thus in particular *is not* correct for small distances. Despite it, the OPL 3 path-loss model (1), we will use it as our default model, because it is precise enough for large enough values of u, it simplifies analysis and reveals important scaling laws.

Coverage (Successful Transmission)

We will say that transmitter $\{X_i\}$ covers its receiver y_i in the reference time slot if

(2) $$\mathrm{SINR}_i = \frac{F_i^i / l(|X_i - y_i|)}{W + I_i^1} \geq T \,,$$

where

- $I_i^1 = \sum_{X_j \in \widetilde{\Phi}^1,\, j \neq i} F_j^i / l(|X_j - y_i|)$ is the SN of $\widetilde{\Phi}^1$ and models the interference,
- $W > 0$ is the external (thermal) noise — a r. v. independent of everything else.
- and where T is some SINR threshold.

We say equivalently that x_i is successfully received by y_i.

Coverage Indicator as a New Mark

Denote by δ_i the indicator that transmitter X_i covers its receiver y_i; i.e., that the SINR condition (2) holds. We will consider δ_i as a new mark of X_i.

Coverage Indicator as a New Mark

Denote by δ_i the indicator that transmitter X_i covers its receiver y_i; i.e., that the SINR condition (2) holds. We will consider δ_i as a new mark of X_i.

The marked point process $\widetilde{\Phi}$ enriched by δ_i is stationary; i.e., its distribution is invariant with respect to any transition.

Coverage Indicator as a New Mark

Denote by δ_i the indicator that transmitter X_i covers its receiver y_i; i.e., that the SINR condition (2) holds. We will consider δ_i as a new mark of X_i.

The marked point process $\widetilde{\Phi}$ enriched by δ_i is stationary; i.e., its distribution is invariant with respect to any transition. However, in contrast to the original marks $e_i, y_i, \mathbf{F}_i$, given the points of Φ, the random variables $\{\delta_i\}$ are neither independent nor identically distributed given Φ. Indeed, the points of Φ lying in dense clusters have a smaller probability of coverage than more isolated points due to interference; in addition, the shot noise variables I_i^1 make that δ_i's dependent.

Coverage Indicator as a New Mark

Denote by δ_i the indicator that transmitter X_i covers its receiver y_i; i.e., that the SINR condition (2) holds. We will consider δ_i as a new mark of X_i.

The marked point process $\widetilde{\Phi}$ enriched by δ_i is stationary; i.e., its distribution is invariant with respect to any transition. However, in contrast to the original marks $e_i, y_i, \mathbf{F}_i$, given the points of Φ, the random variables $\{\delta_i\}$ are neither independent nor identically distributed given Φ.
Indeed, the points of Φ lying in dense clusters have a smaller probability of coverage than more isolated points due to interference; in addition, the shot noise variables I_i^1 make that δ_i's dependent.

Do we have some typical node?

Coverage Probability for the Typical Node

By probability of coverage of the typical node, given it is a transmitter, we understand

$$\mathsf{P}^{0}\{\,\delta_0 = 1 \,|\, e_0 = 1\} = \mathsf{E}^{0}[\delta_0 \,|\, e_0 = 1],$$

where P^{0} is the Palm probability associated to the (marked) stationary point process $\widetilde{\Phi}$ and where δ_0 is the mark of the point $X_0 = 0$ a.s. located at the origin 0 under P^{0}.

Palm Probability — a Reminder

This Palm probability P^0 is derived from the original (stationary) probability P by the following relation

$$\mathsf{P}^0\{\,\delta_0 = 1 \,|\, e_0 = 1\} = \frac{1}{\lambda_1 |B|} \mathsf{E}\Big[\sum_{X_i \in \Phi^1} \delta_i \mathbf{1}(X_i \in B)\Big]\,;$$

B is an arbitrary subset of the plane and $|B|$ is its surface.

Palm Probability — a Reminder

This Palm probability P^0 is derived from the original (stationary) probability P by the following relation

$$\mathsf{P}^{0}\{\,\delta_0 = 1 \,|\, e_0 = 1\} = \frac{1}{\lambda_1|B|}\mathsf{E}\Big[\sum_{X_i\in\Phi^1}\delta_i 1(X_i \in B)\Big]\,;$$

B is an arbitrary subset of the plane and $|B|$ is its surface.

Knowing that $\lambda_1|B|$ is the expected number of transmitters in B, the typical node coverage probability is the mean number of transmitters which cover their receivers in any given window B in which we observe our network. Note that this mean is based on a double averaging: a mathematical expectation – over all possible realizations of the network and, for each realization, a spatial averaging – over all nodes in B.

Palm Probability — Ergodic Interpretation

If the underlying point process is ergodic (as it is the cased for our i.m. Poisson p.p. $\tilde{\mathbf{\Phi}}$) the typical node coverage probability can also be interpreted as a spatial average of the number of transmitters which cover their receiver in almost every given realization of the network and large $\boldsymbol{B}$ (tending to the whole plane).

Palm Probability — Poisson p.p. Case

For a stationary i.m. Poisson p.p. the probability P^0 can easily be constructed due to Slivnyak's theorem:
under P^0, the nodes of our Poisson network and their marks follow the distribution

$$\widetilde{\Phi} \cup \{(X_0 = 0, e_0, y_0, \mathbf{F}_0)\},$$

where $\widetilde{\Phi}$ is the original stationary i.m. Poisson p.p. (i.e. that seen under the original probability P) and $(e_0, y_0, \mathbf{F}_0)$ is a new copy of the mark independent of everything else and distributed like all other i.i.d. marks $(e_i, y_i, \mathbf{F}_i)$ of $\widetilde{\Phi}$ under P.

Palm Probability — Poisson p.p. Case

For a stationary i.m. Poisson p.p. the probability $\mathsf{P}^{\mathbf{0}}$ can easily be constructed due to Slivnyak's theorem:
under $\mathsf{P}^{\mathbf{0}}$, the nodes of our Poisson network and their marks follow the distribution

$$\widetilde{\boldsymbol{\Phi}} \cup \{(\boldsymbol{X_0 = 0, e_0, y_0, \mathrm{F}_0})\} \, ,$$

where $\widetilde{\boldsymbol{\Phi}}$ is the original stationary i.m. Poisson p.p. (i.e. that seen under the original probability P) and $(\boldsymbol{e_0, y_0, \mathrm{F}_0})$ is a new copy of the mark independent of everything else and distributed like all other i.i.d. marks $(\boldsymbol{e_i, y_i, \mathrm{F}_i})$ of $\widetilde{\boldsymbol{\Phi}}$ under P.

Under $\mathsf{P}^{\mathbf{0}}$, the node at the origin $\boldsymbol{X_0 = 0}$ is called the the typical node. Note that the typical node, is not necessarily a transmitter; $\boldsymbol{e_0}$ is equal to 1 or 0 with probability $\boldsymbol{p}$ and $\boldsymbol{1 - p}$ respectively.

Back to the Coverage Probability

Denote by $p_c(r, \lambda_1, T) = \mathsf{E}^0[\delta_0 \,|\, e_0 = 1]$ the probability of coverage of the typical node given it is a transmitter.
It follows from the above construction (Slivnyak's theorem) that this probability only depends on the density of effective transmitters $\lambda_1 = \lambda p$, on the distance r and on the SINR threshold T; it can be expressed using three independent generic random variables F, I^1, W by the following formula:

$$
\begin{aligned}
p_c(r, \lambda_1, T) &= \mathsf{P}^0\{\, F_0^0 > l(r)T(W + I_0^1) \mid e_0 = 1 \,\} \\
&= \mathsf{P}\{\, F \geq Tl(r)(I^1 + W) \,\}\,. \qquad (3)
\end{aligned}
$$

Back to the Coverage Probability

Denote by $p_c(r, \lambda_1, T) = \mathsf{E}^0[\delta_0 \,|\, e_0 = 1]$ the probability of coverage of the typical node given it is a transmitter.
It follows from the above construction (Slivnyak's theorem) that this probability only depends on the density of effective transmitters $\lambda_1 = \lambda p$, on the distance r and on the SINR threshold T; it can be expressed using three independent generic random variables F, I^1, W by the following formula:

$$\begin{aligned} p_c(r, \lambda_1, T) &= \mathsf{P}^0\{\, F_0^0 > l(r)T(W + I_0^1) \mid e_0 = 1 \,\} \\ &= \mathsf{P}\{\, F \geq Tl(r)(I^1 + W) \,\}\,. \end{aligned} \tag{3}$$

First goal: evaluate $p_c(r, \lambda_1, T)$.

Coverage Probability with Rayleigh Fading

Proposition 1 *In Poisson bipolar network model with exponential* F

$$p_c(r,\lambda_1,T) = \exp\left\{-\mu W T l(r) - 2\pi\lambda_1 \int_0^\infty \frac{u}{1 + l(u)/(Tl(r))}\, du\right\}. \tag{4}$$

In particular if $W \equiv 0$ *and that the path-loss model (1) is used then*

$$p_c(r,\lambda_1,T) = \exp(-\lambda_1 r^2 T^{2/\beta} K(\beta))\,, \tag{5}$$

where $K(\beta)$, *called spatial contention parameter is equal*

$$K(\beta) = \frac{2\pi\Gamma(2/\beta)\Gamma(1-2/\beta)}{\beta} = \frac{2\pi^2}{\beta \sin(2\pi/\beta)}\,. \tag{6}$$

Proof of Proposition 1

From (3) with exponential F (of parameter μ) by independence we obtain

$$
\begin{aligned}
p_c(r,\lambda_1,T) &= \mathsf{E}\Big[\exp[-\mu T l(r)(I^1+W)]\Big] \\
&= e^{-\mu W T l(r)}\mathsf{E}[e^{-\mu T l(r) I^1}] .
\end{aligned}
$$

The second factor in the above expression is just the Laplace transform of the Poisson Shot-noise $\mathcal{L}_{I^1}(s)$ evaluated at $s=\mu T l(r)$. It admits the following closed form expression

$$
\mathcal{L}_{I^1}(s) = \mathsf{E}[e^{-I^1 s}] = \exp\Big\{-\lambda_1 2\pi \int_0^\infty t\Big(1-\mathcal{L}_F(s/l(t))\Big)\,\mathrm{d}t\Big\} , \tag{7}
$$

where $\mathcal{L}_F$ is the Laplace transform of F (here exponential).

Tuning Aloha to Guarantee Coverage Prob.

Example 1 Assume one wants to operate a network with Aloha MAC where each transmitter-receiver distance is r and a successful transmission is guaranteed with a probability at least $1-\varepsilon$, where ε is a predefined QoS. Then, the MAP p parameter of Aloha should be such that $p_c(r, \lambda p, T) = 1-\varepsilon$. In particular, assuming the path-loss setting (1), one should take

$$p = \min\left(1, \frac{-\ln(1-\varepsilon)}{\lambda r^2 T^{2/\beta} K(\beta)}\right) \approx \min\left(1, \frac{\varepsilon}{\lambda r^2 T^{2/\beta} K(\beta)}\right).$$

For example, for $T = 10$dB [a] and OPL 3 model with $\beta = 4$, $r = 1$, one should take $p \approx \min(1, 0.064\,\varepsilon/\lambda)$.

---

[a] A positive real number x is $10\log_{10}(x)$ dB.

Coverage Probability with General Fading

The results of Proposition 1 can be extended to a general case of F using Plancherel-Parseval theorem.

Coverage Probability with General Fading

The results of Proposition 1 can be extended to a general case of $\boldsymbol{F}$ using Plancherel-Parseval theorem.

Proposition 2 *Consider the Poisson bipolar network model with fading variables* $\boldsymbol{F}$ *such that*

- $\boldsymbol{F}$ *has a finite first moment and admits a square integrable density;*
- *Either* $\boldsymbol{I}^1$ *or* $\boldsymbol{W}$ *admit a density which is square integrable.*

Then the probability of a successful transmission is equal to

$$p_c(r, \lambda_1, T) = \int_{s=-\infty}^{\infty} \mathcal{L}_{I^1}(2i\pi l(r)Ts)\, \mathcal{L}_W(2i\pi l(r)Ts)\, \frac{\mathcal{L}_F(-2i\pi s) - 1}{2i\pi s} ds\,.$$

Impact of Fading on Coverage Probability

Proposition 2 allows to compare analytically the impact of fading on coverage probability. A general observation is of this sort:

Stronger fading is beneficial in for larger transmission distances and detrimental for smaller ones.

We skip the details.

Mean Packet Progress

In view of multi-hop routing one might be interested in finding the transmission distance r which maximizes the mean packet progress

(8) $$prog(r, \lambda_1, T) = r\mathsf{E}^0[\delta_0] = rp_c(r, \lambda_1, T)$$

given all other parameters (including λ, p) fixed in our simple model.

Mean Packet Progress

In view of multi-hop routing one might be interested in finding the transmission distance $\boldsymbol{r}$ which maximizes the mean packet progress

$$prog(r, \lambda_1, T) = r\mathsf{E}^0[\delta_0] = r p_c(r, \lambda_1, T) \tag{8}$$

given all other parameters (including $\boldsymbol{\lambda}, \boldsymbol{p}$) fixed in our simple model.

Obviously small $\boldsymbol{r}$ makes the transmissions more sure but involves more relaying nodes to communicate on some given (large) distance. On the other hand large $\boldsymbol{r}$ reduces the number of hops but might increase the number faults and retransmissions on a given hop.

Optimizing the Mean Packet Progress

Simple analysis (that we skip here) shows that in the case of the power-law path loss function (OPL 3) the optimal transmission distance r for the mean packet progress is of the following order

$$r_{\text{max}}(\lambda p) = \frac{\text{const}}{T^{1/\beta}\sqrt{\lambda p}}$$

and in the case of Rayleigh fading

$$\text{const} = \frac{1}{2K(\beta)},$$

where $K(\beta)$ is the spatial contention parameter.

Optimizing the Mean Packet Progress

Simple analysis (that we skip here) shows that in the case of the power-law path loss function (OPL 3) the optimal transmission distance r for the mean packet progress is of the following order

$$r_{\max}(\lambda p) = \frac{\text{const}}{T^{1/\beta}\sqrt{\lambda p}}$$

and in the case of Rayleigh fading

$$\text{const} = \frac{1}{2K(\beta)},$$

where $K(\beta)$ is the spatial contention parameter.

Further optimization, in p, degenerates: $p_{\max} = 0$.

Density of Successful Transmissions

In contrast to $\boldsymbol{prog}$ this is a network (social) performance metric defined as

$$d_{suc}(r, \lambda_1.T) = \frac{1}{|B|} \mathsf{E}\Big[\sum_i e_i \delta_i \mathbf{1}(X_i \in B)\Big].$$

Density of Successful Transmissions

In contrast to $prog$ this is a network (social) performance metric defined as

$$d_{suc}(r, \lambda_1.T) = \frac{1}{|B|}\mathsf{E}\Big[\sum_i e_i \delta_i \mathbf{1}(X_i \in B)\Big].$$

By stationarity of the model, does not depend on the particular choice of set B and by Campbell's formula it can be expressed in therms of coverage probability

(9) $$d_{suc}(r, \lambda_1.T) = \lambda_1 p_c(r, \lambda_1, T) = \lambda p\, p_c(r, \lambda p, T).$$

Optimizing Density d_{suc} in p

Density of successful transmissions d_{suc} can be explicitly optimized in p.

Optimizing Density d_{suc} in p

Density of successful transmissions d_{suc} can be explicitly optimized in p.

Simple analysis (that we skip here) shows that in the case of the power-law path loss function (OPL 3) and Rayleigh fading the optimal MAP p for the mean packet progress is

$$p_{\max}(\lambda, r) = \min(1, \lambda_{\text{max}}/\lambda)$$

where

$$\lambda_{\text{max}}(\lambda, r) = \frac{1}{K(\beta) r^2 T^{2/\beta}}$$

and $K(\beta)$ is the spatial contention parameter.

Optimizing Density d_{suc} in p

Density of successful transmissions d_{suc} can be explicitly optimized in p.
Simple analysis (that we skip here) shows that in the case of the power-law path loss function (OPL 3) and Rayleigh fading the optimal MAP p for the mean packet progress is

$$p_{\max}(\lambda, r) = \min(1, \lambda_{\text{max}}/\lambda)$$

where

$$\lambda_{\text{max}}(\lambda, r) = \frac{1}{K(\beta) r^2 T^{2/\beta}}$$

and $K(\beta)$ is the spatial contention parameter.

> Note: for small density of nodes ($\lambda < \lambda_{\max}$) no MAC in needed ($p_{\max} = 1$)!

Optimizing Density d_{suc} in p

Density of successful transmissions d_{suc} can be explicitly optimized in p.
Simple analysis (that we skip here) shows that in the case of the power-law path loss function (OPL 3) and Rayleigh fading the optimal MAP p for the mean packet progress is

$$p_{\max}(\lambda, r) = \min(1, \lambda_{\text{max}}/\lambda)$$

where

$$\lambda_{\text{max}}(\lambda, r) = \frac{1}{K(\beta) r^2 T^{2/\beta}}$$

and $K(\beta)$ is the spatial contention parameter.

Note: for small density of nodes ($\lambda < \lambda_{\max}$) no MAC in needed ($p_{\max} = 1$)!

Further optimization, in r, again degenerates: $r_{\max} = 0$.

Other Spatial/Social Performance Metrics

The following characteristics can also be expressed in terms of the coverage probability $p_c(r, \lambda_1, T)$.

- spatial density of progress, d_{prog}, the mean number of meters progressed by all transmissions taking place per unit surface unit;
- spatial density of Shannon throughput, d_{throu}, the mean throughput per unit surface unit;
- spatial density of transport, d_{trans}, the mean number of bit-meters transported per second and per unit of surface.

We skip the details.

Bipolar Model — Conclusions

Simple yet not simplistic model. Allows for

- closed form expression for the successful transmission probability.

Bipolar Model — Conclusions

Simple yet not simplistic model. Allows for

- closed form expression for the successful transmission probability.
- pertinent optimization of many network performance metrics in Aloha parameter p and transmission distance r.

Bipolar Model — Conclusions

Simple yet not simplistic model. Allows for

- closed form expression for the successful transmission probability.
- pertinent optimization of many network performance metrics in Aloha parameter p and transmission distance r.

A better receiver model is needed to avoid degenerate joint optimization in r and p. We will propose such models in what follows.

Bipolar Model — Conclusions

Simple yet not simplistic model. Allows for

- closed form expression for the successful transmission probability.
- pertinent optimization of many network performance metrics in Aloha parameter p and transmission distance r.

A better receiver model is needed to avoid degenerate joint optimization in r and p. We will propose such models in what follows.

Before changing the receiver model, let us briefly visit some two other extensions of the Bipolar model.

EXTENSION 1: OPPORTUNISTIC ALOHA

The Idea

In the basic Aloha scheme, each node tosses a coin to access the medium independently of the channel conditions.

It is clear that something more clever can be done by combining the random selection of transmitters with the occurrence of good channel conditions.

The Idea

In the basic Aloha scheme, each node tosses a coin to access the medium independently of the channel conditions. It is clear that something more clever can be done by combining the random selection of transmitters with the occurrence of good channel conditions.

The general idea of Opportunistic Aloha is to select the nodes with the channel fading larger than a certain threshold as transmitters in the reference time slot.

The Idea

In the basic Aloha scheme, each node tosses a coin to access the medium independently of the channel conditions. It is clear that something more clever can be done by combining the random selection of transmitters with the occurrence of good channel conditions.

The general idea of Opportunistic Aloha is to select the nodes with the channel fading larger than a certain threshold as transmitters in the reference time slot.

This is a kind of ad-hoc MAC version of the HDR (HSDPA) protocol implemented in cellular networks.

Model Modification

Opportunistic Aloha can be described by $\widetilde{\Phi} = \{(X_i, \theta_i, y_i, \mathrm{F}_i)\}$, where $\{(X_i, y_i, \mathrm{F}_i)\}$ is as in the basic Poisson Bipolar Model (1)–(4), with item (2) replaced by:

(2') Opportunistic Aloha principle: The MAC indicator e_i of node i ($e_i = 1$ if node i is allowed to transmit and 0 otherwise) is the following function of the channel condition to its receiver F_i^i: $e_i = 1(F_i^i > \theta_i)$, where $\{\theta_i\}$ are new random i.i.d. marks, with a generic mark denoted by θ.

Model Modification

Opportunistic Aloha can be described by $\widetilde{\boldsymbol{\Phi}} = \{(\boldsymbol{X}_i, \boldsymbol{\theta}_i, \boldsymbol{y}_i, \mathbf{F}_i)\}$, where $\{(\boldsymbol{X}_i, \boldsymbol{y}_i, \mathbf{F}_i)\}$ is as in the basic Poisson Bipolar Model (1)–(4), with item (2) replaced by:

(2') Opportunistic Aloha principle: The MAC indicator $\boldsymbol{e_i}$ of node $\boldsymbol{i}$ ($\boldsymbol{e_i = 1}$ if node $\boldsymbol{i}$ is allowed to transmit and 0 otherwise) is the following function of the channel condition to its receiver $\boldsymbol{F_i^i}$: $\boldsymbol{e_i = 1(F_i^i > \theta_i)}$, where $\boldsymbol{\{\theta_i\}}$ are new random i.i.d. marks, with a generic mark denoted by $\boldsymbol{\theta}$.

Special cases of interest are that where $\boldsymbol{\theta}$ is constant, and that where $\boldsymbol{\theta}$ is exponential with parameter $\boldsymbol{\nu}$. (allows for close-form expression for the coverage probability).

Model Modification

Opportunistic Aloha can be described by $\widetilde{\mathbf{\Phi}} = \{(X_i, \theta_i, y_i, \mathbf{F}_i)\}$, where $\{(X_i, y_i, \mathbf{F}_i)\}$ is as in the basic Poisson Bipolar Model (1)–(4), with item (2) replaced by:

(2') Opportunistic Aloha principle: The MAC indicator e_i of node i ($e_i = 1$ if node i is allowed to transmit and 0 otherwise) is the following function of the channel condition to its receiver F_i^i: $e_i = 1(F_i^i > \theta_i)$, where $\{\theta_i\}$ are new random i.i.d. marks, with a generic mark denoted by θ.

Special cases of interest are that where θ is constant, and that where θ is exponential with parameter ν. (allows for close-form expression for the coverage probability).
As in Aloha $\{e_i\}$ are again i.i.d. marks of the point process $\widetilde{\mathbf{\Phi}}$, which now depend on $\{\theta_i, F_i^i\}$.

Coverage Probability in Opp'c Aloha

Proposition 3 *Assume Rayleigh fading (exponential* F *with parameter* μ*), exponential distribution of the threshold* θ *with parameter* ν*, and (for simplicity)* $W \equiv 0$ *and the OPL 3 model (1). Then*

$$\begin{aligned} &\widehat{p}_c(r, \lambda_1, \nu) \\ &\quad = \frac{\mu+\nu}{\nu}\exp\{-\lambda_1 T^{2/\beta} r^2 K(\beta)\} \\ &\qquad -\frac{\mu}{\nu}\exp\Big\{-\lambda_1\left(\frac{(\mu+\nu)T}{\mu}\right)^{2/\beta} r^2 K(\beta)\Big\}, \end{aligned}$$

with $\lambda_1 = \lambda\nu/(\mu+\nu)$.

Opportunistic vs Plain Aloha

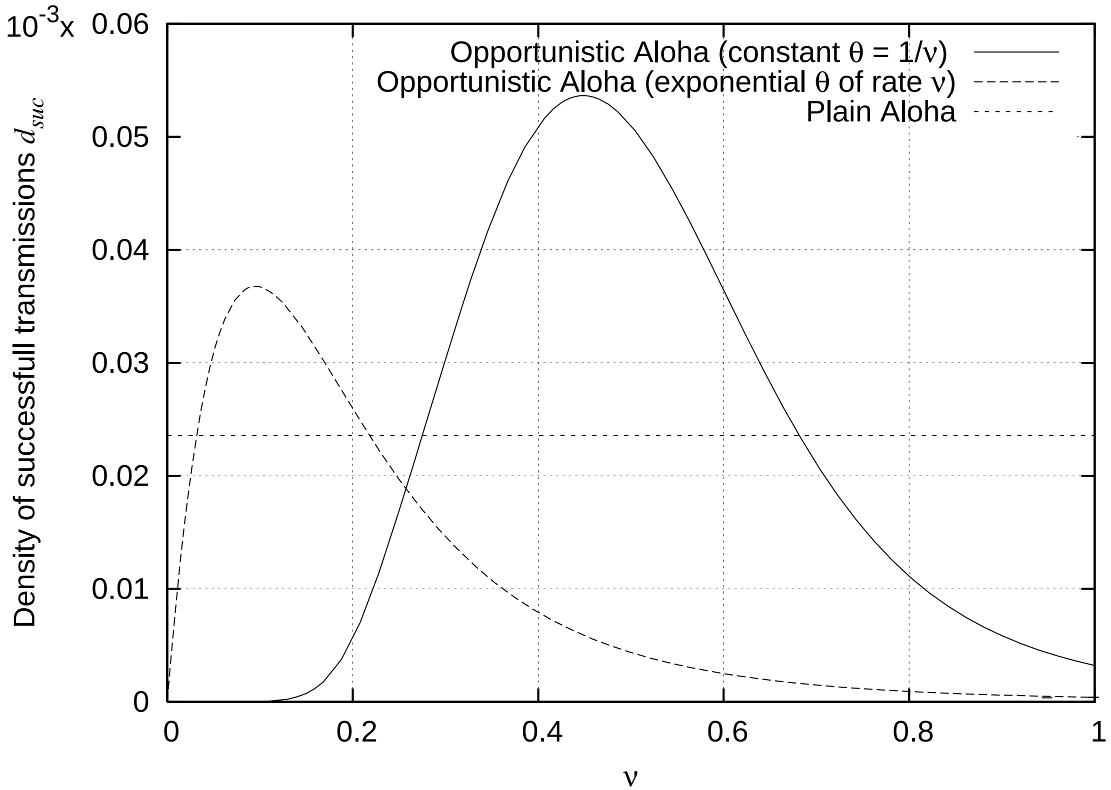

The density of successful transmissions d_{suc} of Opportunistic Aloha for various choices of θ. The propagation model is (1). We assume Rayleigh fading with mean 1 and $W = 0$, $\lambda = 0.001$, $T = 10$dB, $r = \sqrt{1/\lambda}$ and $\beta = 4$. For comparison the constant value $\lambda_{\text{max}} p_c(r, \lambda_{\text{max}})$ of plain Aloha is plotted.

EXTENSION 2:
NON-SLOTTED ALOHA

Asynchronous Transmissions

All nodes transmit a packet of length B and back-off for some random time before the next transmission asynchronously (no common notion of time slots).

Asynchronous Transmissions

All nodes transmit a packet of length B and back-off for some random time before the next transmission asynchronously (no common notion of time slots).

We have proposed two models of this protocol, all require space-time modeling. Here we present the simpler one.

Asynchronous Transmissions

All nodes transmit a packet of length B and back-off for some random time before the next transmission asynchronously (no common notion of time slots).

We have proposed two models of this protocol, all require space-time modeling. Here we present the simpler one.

The objective is to revisit the “classical” result saying that the slotted Aloha outperforms the non-slotted one by the factor of 2 with respect to the fraction of successful transmissions, when both are optimally tuned. This classical result being obtained in for a geometry-less collision model which assumes that simultaneous transmissions are never successful.

Poisson Rain Model for Non-slotted Aloha

- $\Psi = \{(X_i, T_i)\}$ time-space Poisson point process with density λ_s transmission initiations per km^2 and par unit time. (Indexing by i is arbitrary and in particular does *not* mean successive emissions over time).
- $e_i(t) = 1(T_i \le t < T_i + B)$ on-off process of the MAC state of the node X_i at (real) time t.

Poisson Rain Model for Non-slotted Aloha

- $\Psi = \{(X_i, T_i)\}$ time-space Poisson point process with density λ_s transmission initiations per km^2 and par unit time. (Indexing by i is arbitrary and in particular does *not* mean successive emissions over time).
- $e_i(t) = 1(T_i \leq t < T_i + B)$ on-off process of the MAC state of the node X_i at (real) time t.

We may think of nodes “born” at time T_i at location X_i, transmitting a packet during time B and “disappearing” immediately after. This can be naturally motivated by mobility of nodes.

All other assumptions are the same as for the slotted Aloha (including the fixed distance r to the receiver) except that in the SINR capture condition (2) the interference (that is not constant) is averaged out over the packet reception time B

$$I_i^{mean} = 1/B \int_{T_i}^{T_i+B} I(t).$$

All other assumptions are the same as for the slotted Aloha (including the fixed distance r to the receiver) except that in the SINR capture condition (2) the interference (that is not constant) is averaged out over the packet reception time B

$$I_i^{mean} = 1/B \int_{T_i}^{T_i+B} I(t).$$

This is a reasonable assumption if some coding with repetition and interleaving on the whole packet duration is used.

Coverage Prob. in the non-Slotted Model

Proposition 4 *Assume Rayleigh fading and SINR condition with averaged interference. The coverage probability is*

$$p_{rain}^{mean} = \mathcal{L}_W(\mu T l(r)) \times \exp\left\{ -4\pi\lambda_s B \int_0^\infty u\left(1 - \frac{l(u)}{l(r)} \log\left(1 + \frac{l(r)T}{l(u)}\right)\right) du\right\}.$$

In particular for $W \equiv 0$ *and power-law path-loss*

$$p_{rain}^{mean} = \exp(-\lambda_s B r^2 T^{2/\beta} K'(\beta)),$$

with spatial contention parameter

$$K'(\beta) = \frac{4\pi}{\beta} \int_0^\infty u^{2/\beta-1}(1 - u\log(1 + u^{-1}))\, du.$$

Slotted vs Non-slotted Aloha

In the simplest case (power law path loss function, Rayleigh fading) the expressions for the coverage probability in slotted and non-slotted Aloha differ only by the spatial contention parameters $\boldsymbol{K(\beta)}$, $\boldsymbol{K'(\beta)}$ with slotted Aloha having smaller spatial contention parameter

Slotted vs Non-slotted Aloha

In the simplest case (power law path loss function, Rayleigh fading) the expressions for the coverage probability in slotted and non-slotted Aloha differ only by the spatial contention parameters $\boldsymbol{K(\beta)}$, $\boldsymbol{K'(\beta)}$ with slotted Aloha having smaller spatial contention parameter

Proposition 5 *The ratio of the spatial contention parameters*

$$\boldsymbol{1 > \frac{K(\beta)}{K'(\beta)} > 0.5 \quad \textit{for } 2 < \beta < \infty}$$

is equals the ratio of the density of successful transmissions optimized respectively in both models.

$K(\beta)/K'(\beta)$

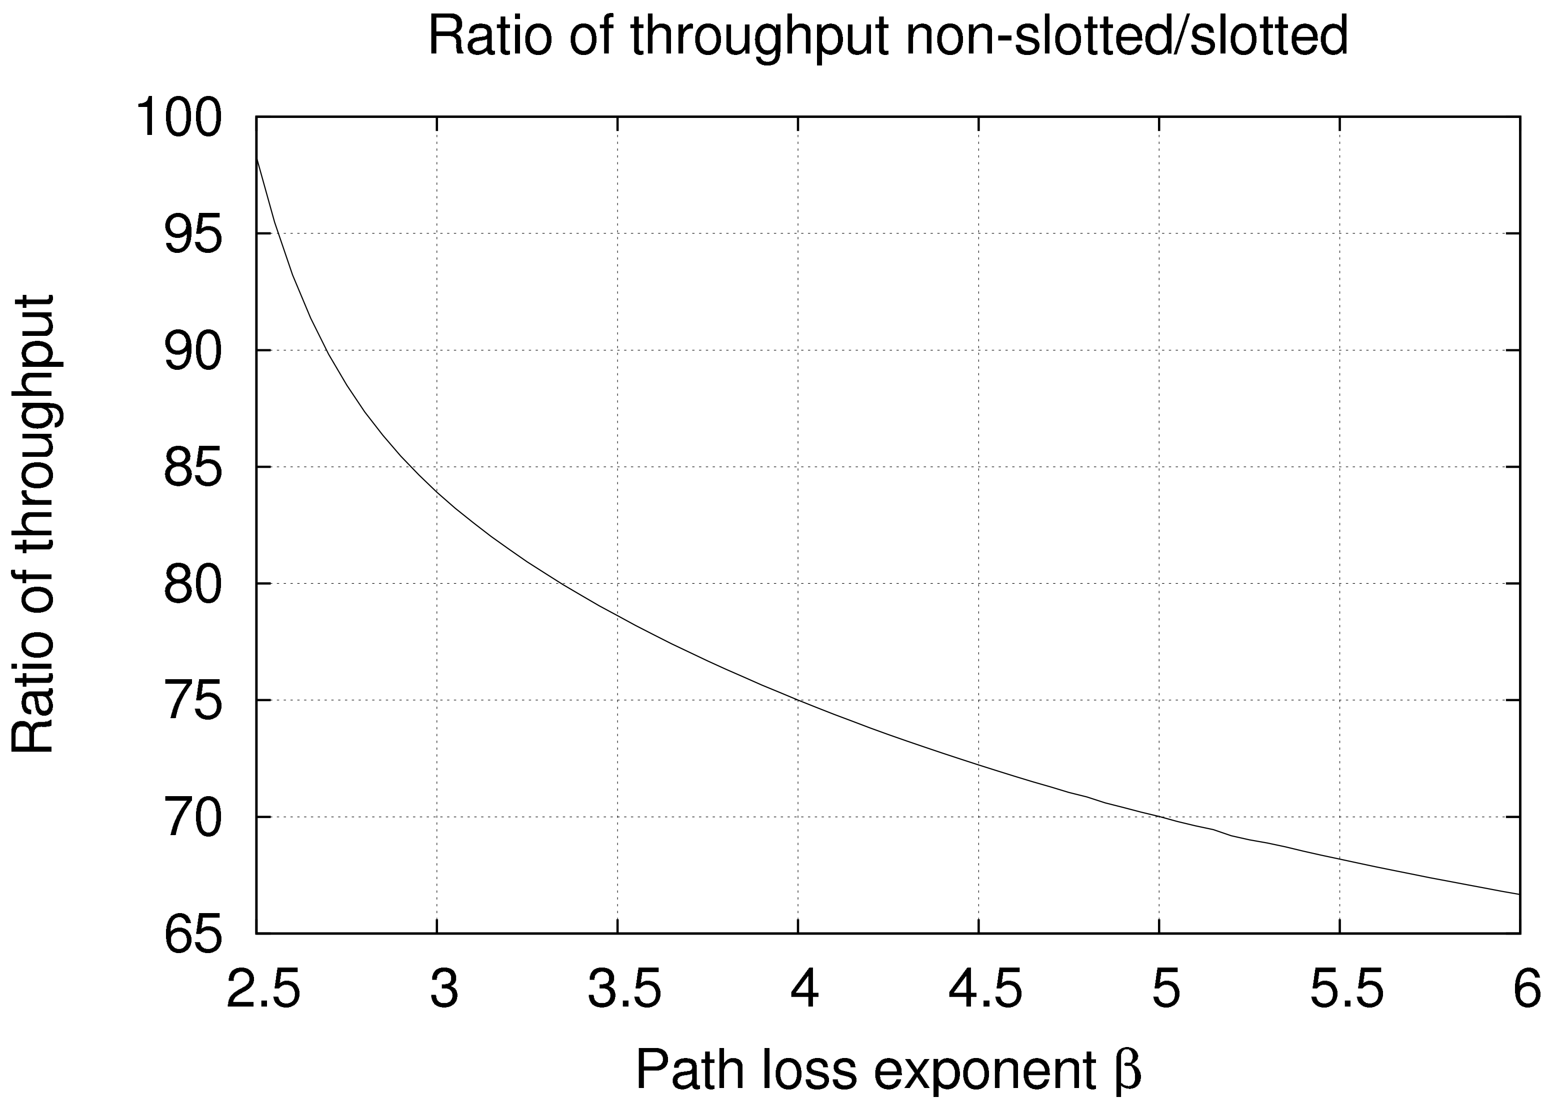
Ratio of throughput non-slotted/slotted
100
95
90
85
80
75
70
65
Ratio of throughput
2.5
3
3.5
4
4.5
5
5.5
6
Path loss exponent β

The Classical Comparison Result Revisited

- For small values of path-loss exponent β (close to 2) the performances of optimized slotted and non-slotted Aloha are similar.
- For large β (approaching ∞) the good-put ratio goes to 0.5 — the value predicted by the widely used simplified model with the simplified collision model.
- However, e.g. for $\beta = 4$ this ratio is still 75% and even for $\beta = 6$ the ratio still remains significantly larger than 50%.

The Classical Comparison Result Revisited

- For small values of path-loss exponent β (close to 2) the performances of optimized slotted and non-slotted Aloha are similar.
- For large β (approaching ∞) the good-put ratio goes to 0.5 — the value predicted by the widely used simplified model with the simplified collision model.
- However, e.g. for $\beta = 4$ this ratio is still 75% and even for $\beta = 6$ the ratio still remains significantly larger than 50%.

When trying to explain the above asymptotic value of 50%, one may argue that in the presence of a very strong path-loss only very local interactions exist and the model becomes "geometryless".

A NEW PHASE TRANSITION FOR LOCAL DELAYS IN AD-HOC NETWORKS

RESUME

A Spatial (In-)Stability

Communication systems typically have bounded stability regions: throughput is non-null only if the offered traffic is small enough.

A Spatial (In-)Stability

Communication systems typically have bounded stability regions: throughput is non-null only if the offered traffic is small enough.

This story is on a $\boxed{\text{spatial}}$ stability of wireless networks —

- a new notion of stability,
- intrinsically related to spatial reuse of wireless spectrum,
- observed here in mobile ad-hoc networks (MANETs).

Setting

Network: Emitters and their (next hop) receivers randomly located on the plane

MAC: Aloha

Successful transmission:
SINR larger than some threshold

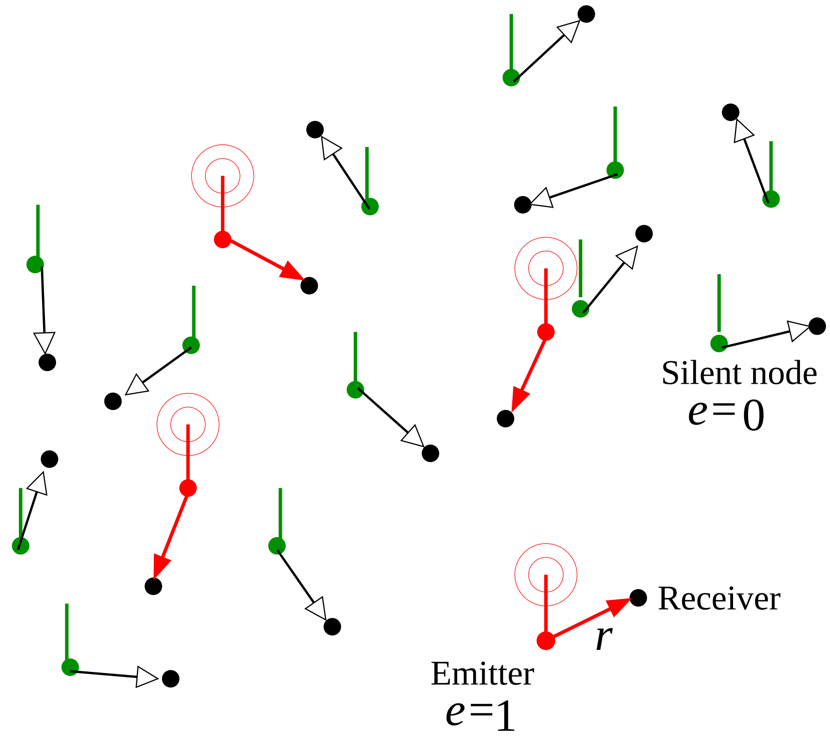

Setting

Network: Emitters and their (next hop) receivers randomly located on the plane

MAC: Aloha

Successful transmission:
SINR larger than some threshold

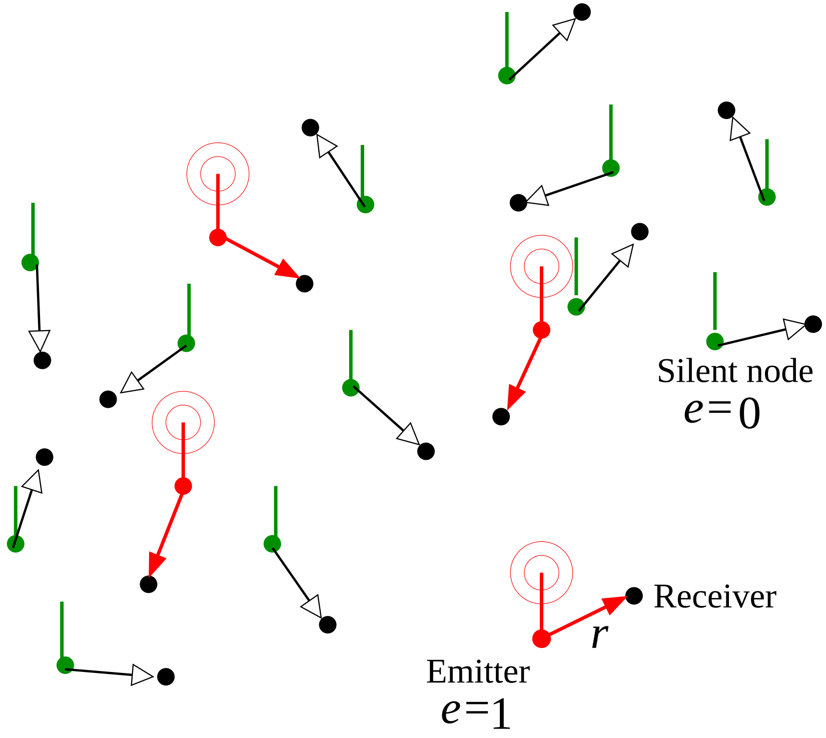

We analyze the local delays: number of times slots required for nodes to transmit a packet to their receivers.

Results

- Locally the network works well: for a given realization of nodes, of each node has a positive probability of successful transmission (with respect to channel and MAC variability), finite mean local delay and thus a positive next hop throughput.

Results

- Locally the network works well: for a given realization of nodes, of each node has a positive probability of successful transmission (with respect to channel and MAC variability), finite mean local delay and thus a positive next hop throughput.
- Still macroscopically (spatially) the network might not be stable: large node-population averaging of the finite individual mean delays (in several practical cases) gives infinite values in several practical cases.

Results

- Locally the network works well: for a given realization of nodes, of each node has a positive probability of successful transmission (with respect to channel and MAC variability), finite mean local delay and thus a positive next hop throughput.

- Still macroscopically (spatially) the network might not be stable: large node-population averaging of the finite individual mean delays (in several practical cases) gives infinite values in several practical cases.

- Sometimes network exhibits interesting/dangerous phase transition: a slight change of certain model parameters (receiver distance, thermal noise power, medium access probability) may take the network from spatial stability to instability.

What actually means the spatial instability?

What actually means the spatial instability?

The usual (temporal) instability: even if each individual packet is eventually transmitted, the average (over a large number of packets) time required to transmit one packets is getting infinite.

What actually means the spatial instability?

The usual (temporal) instability: even if each individual packet is eventually transmitted, the average (over a large number of packets) time required to transmit one packets is getting infinite.

Spatial network instability: even if each individual node in the network has some finite mean transmission delay, the average (over a large number of nodes) mean transmission delay per node is getting infinite.

What actually means the spatial instability?

The usual (temporal) instability: even if each individual packet is eventually transmitted, the average (over a large number of packets) time required to transmit one packets is getting infinite.

Spatial network instability: even if each individual node in the network has some finite mean transmission delay, the average (over a large number of nodes) mean transmission delay per node is getting infinite.

Spatial instability: the MAC protocol performance does not scale with the network size.

MODEL DESCRIPTION
&
FIRST RESULTS

Previous Model, Time Dimension Added

- Static Poisson MANET of density λ nodes/km^2.

Previous Model, Time Dimension Added

- Static Poisson MANET of density λ nodes/km^2.
- Synchronous slotted Aloha MAC: each node independently, with probability p, in each time slot $n = 1, 2, \ldots$ decides to emit one packet; $e = 1, 0$ denotes MAC decision, E$[e] = p$.

Previous Model, Time Dimension Added

- Static Poisson MANET of density λ nodes/km^2.
- Synchronous slotted Aloha MAC: each node independently, with probability p, in each time slot $n = 1, 2, \ldots$ decides to emit one packet; $e = 1, 0$ denotes MAC decision, $\mathsf{E}[e] = p$.
- I.i.d. point-to-point fading F, constant in a given time slot, may or may-not vary across times slots:
  - $\boxed{\text{slow fading}}$ (shadowing): channel conditions do not change in time,
  - $\boxed{\text{fast fading}}$: channel conditions independently re-sampled for each channel in each slot.

Previous Model, Time Dimension Added

- Static Poisson MANET of density λ nodes/km^2.
- Synchronous slotted Aloha MAC: each node independently, with probability p, in each time slot $n = 1, 2, \ldots$ decides to emit one packet; $e = 1, 0$ denotes MAC decision, $\mathsf{E}[e] = p$.
- I.i.d. point-to-point fading F, constant in a given time slot, may or may-not vary across times slots:
  - $\boxed{\text{slow fading}}$ (shadowing): channel conditions do not change in time,
  - $\boxed{\text{fast fading}}$: channel conditions independently re-sampled for each channel in each slot.
- External noise power W, may or may-not vary in time ($\boxed{\text{slow}}$ or $\boxed{\text{fast noise}}$ scenario, respectively).

Receiver Models

1. Bipolar model: each MANET node X has its dedicated receiver y (not in MANET) at a distance r km from it.

Receiver Models

1. Bipolar model: each MANET node X has its dedicated receiver y (not in MANET) at a distance r km from it.
2. Independent Poisson Nearest Receiver (IPNR) model: each transmitter selects its receiver as close by as possible in some Poisson set of potential receivers of density λ_0 nodes/km^2 (external to MANET).

Receiver Models

1. Bipolar model: each MANET node X has its dedicated receiver y (not in MANET) at a distance r km from it.
2. Independent Poisson Nearest Receiver (IPNR) model: each transmitter selects its receiver as close by as possible in some Poisson set of potential receivers of density λ_0 nodes/km^2 (external to MANET).
3. MANET Nearest Neighbor (MNN) model: each transmitter selects its receiver as close by as possible in the MANET.

SINR Capture

Transmitter X is successfully received at y at a given time slot n if the following condition

$$\mathsf{SINR} = \mathsf{SINR}(n) = \frac{F/l(|X-y|)}{W+I} \geq T\,,$$

is satisfied, where $F = F(n), W = W(n), I = I(n)$, are, respectively, $X \to y$ channel fading, external noise power and interference, and T is the SINR threshold.

SINR Capture

Transmitter $\boldsymbol{X}$ is successfully received at $\boldsymbol{y}$ at a given time slot $\boldsymbol{n}$ if the following condition

$$\text{SINR} = \text{SINR}(\boldsymbol{n}) = \frac{\boldsymbol{F}/\boldsymbol{l}(|\boldsymbol{X} - \boldsymbol{y}|)}{\boldsymbol{W} + \boldsymbol{I}} \geq \boldsymbol{T}\,,$$

is satisfied, where $\boldsymbol{F} = \boldsymbol{F}(\boldsymbol{n}), \boldsymbol{W} = \boldsymbol{W}(\boldsymbol{n}), \boldsymbol{I} = \boldsymbol{I}(\boldsymbol{n})$, are, respectively, $\boldsymbol{X} \rightarrow \boldsymbol{y}$ channel fading, external noise power and interference, and $\boldsymbol{T}$ is the SINR threshold.

$\boldsymbol{\delta} = \boldsymbol{\delta_i}(\boldsymbol{n})$ indicator of the successful transmission, i.e., that the above SINR condition holds for the MANET node $\boldsymbol{X_i}$ with its receiver $\boldsymbol{y_i}$ at time $\boldsymbol{n}$.

Local Delays

The local delay of the node $\boldsymbol{X_i}$ is the number of time slots it needs to successfully transmit a tagged packet

$$\mathbf{L}_i = \inf\{n \geq 1 : e_i(n)\delta_i(n) = 1\}\,.$$

Local Delays

The local delay of the node X_i is the number of time slots it needs to successfully transmit a tagged packet

$$\mathrm{L}_i = \inf\{n \geq 1 : e_i(n)\delta_i(n) = 1\}\,.$$

Let $\mathcal{S}$ denote all the static components of the network model (which do not vary in time n). Node locations (emitters, receivers) are in $\mathcal{S}$. Fading and noise variables are in $\mathcal{S}$ in the respective “slow” models.

Local Delays

The local delay of the node X_i is the number of time slots it needs to successfully transmit a tagged packet

$$\mathrm{L}_i = \inf\{n \geq 1 : e_i(n)\delta_i(n) = 1\}\,.$$

Let $\mathcal{S}$ denote all the static components of the network model (which do not vary in time n). Node locations (emitters, receivers) are in $\mathcal{S}$. Fading and noise variables are in $\mathcal{S}$ in the respective “slow” models.

All random elements that not in $\mathcal{S}$ vary in an i.i.d. manner in time. Thus:

> Given a realization of $\mathcal{S}$, local delays L_i are geometric random variables as the number of trials until the first success in some Bernoulli scheme with the probability of success $\pi_c(X_i, \mathcal{S}) = \mathsf{E}[e_i\delta_i \,|\, \mathcal{S}]$.

Hence,

the conditional (given all the static elements $\boldsymbol{\mathcal{S}}$ of the network) mean local delay of node $\boldsymbol{X_i}$ is equal to

$$\mathsf{E}[\mathbf{L}_i \,|\, \boldsymbol{\mathcal{S}}] = \frac{\mathbf{1}}{\boldsymbol{\pi_c(X_i, \mathcal{S})}}\,.$$

Hence,

the conditional (given all the static elements $\mathcal{S}$ of the network) mean local delay of node X_i is equal to

$$\mathsf{E}[\mathbf{L}_i \mid \mathcal{S}] = \frac{1}{\pi_c(X_i, \mathcal{S})}\,.$$

One can interpret $\pi_c(X_i, \mathcal{S})$ as the (temporal) rate of successful packet transmissions (throughput) of node X_i given all the static elements $\mathcal{S}$ of the network.

Spatial variability of local delays

Obviously, $\pi_c(X_i, \mathcal{S})$ are different for different nodes.
In other words, different MANET nodes have different throughputs and mean local delays $\mathsf{E}[\mathbf{L}_i \mid \mathcal{S}]$.

Spatial variability of local delays

Obviously, $\pi_c(X_i, \mathcal{S})$ are different for different nodes.
In other words, different MANET nodes have different throughputs and mean local delays $\mathsf{E}[\mathbf{L}_i \mid \mathcal{S}]$.

In Poisson MANET one can find nodes which have an arbitrarily small conditional temporal throughput $\pi_c(X_i, \mathcal{S})$ and thus arbitrarily large conditional mean local delay $\mathsf{E}[\mathbf{L}_i \mid \mathcal{S}] = 1/\pi_c(X_i, \mathcal{S})$.

Typical MANET node

In what follows we are interested in spatial averages of these conditional mean local delays $\mathsf{E}[\boldsymbol{L}_i \mid \boldsymbol{\mathcal{S}}]$, i.e., averages taken over a large population of MANET nodes

$$\lim_{n\to\infty} \sum_{i=1}^{n} \mathsf{E}[\boldsymbol{L}_i \mid \boldsymbol{\mathcal{S}}].$$

Typical MANET node

In what follows we are interested in spatial averages of these conditional mean local delays $\mathsf{E}[L_i \,|\, \mathcal{S}]$, i.e., averages taken over a large population of MANET nodes

$$\lim_{n\to\infty} \sum_{i=1}^{n} \mathsf{E}[L_i \,|\, \mathcal{S}].$$

Mathematically, this is equivalent to the analysis of the averaged (over Poisson pattern of nodes) local delay of the so-called $\boxed{\text{typical MANET node}}$. In Poisson case the typical node in just an extra node X_0 added to MANET, say at the origin. The spatial averaging over the pattern of nodes in this scenario is called also Palm expectation and is traditionally denoted by E^0.

Thus

$$\lim_{n\to\infty} \sum_{i=1}^{n} \mathsf{E}[L_i \,|\, \mathcal{S}] = \mathsf{E}^0[\mathsf{E}[L_0 \,|\, \mathcal{S}]] =: \mathsf{E}^0[L_0]$$

Local Delay / Throughput Paradox

Recall, $\mathsf{E}^0[L_0 \mid \mathcal{S}] = 1/\pi_c(X_0, \mathcal{S}) = 1/\mathsf{E}^0[e_0\delta_0 \mid \mathcal{S}]$ and thus

$$\mathsf{E}^0[L_0] = \mathsf{E}^0\Big[\frac{1}{\mathsf{E}^0[e_0\delta_0 \mid \mathcal{S}]}\Big]$$

Local Delay / Throughput Paradox

Recall, $\mathsf{E}^0[L_0 \,|\, \mathcal{S}] = 1/\pi_c(X_0, \mathcal{S}) = 1/\mathsf{E}^0[e_0\delta_0 \,|\, \mathcal{S}]$ and thus

$$\mathsf{E}^0[L_0] = \mathsf{E}^0\left[\frac{1}{\mathsf{E}^0[e_0\delta_0 \,|\, \mathcal{S}]}\right] \geq \frac{1}{\mathsf{E}^0[e_0\delta_0]}.$$

by Jensen's inequality.

Local Delay / Throughput Paradox

Recall, $\mathsf{E}^0[L_0 \mid \mathcal{S}] = 1/\pi_c(X_0, \mathcal{S}) = 1/\mathsf{E}^0[e_0\delta_0 \mid \mathcal{S}]$ and thus

$$\mathsf{E}^0[L_0] = \mathsf{E}^0\Big[\frac{1}{\mathsf{E}^0[e_0\delta_0 \mid \mathcal{S}]}\Big] \geq \frac{1}{\mathsf{E}^0[e_0\delta_0]}\,.$$

by Jensen's inequality.

Note that $\mathsf{E}^0[e_0\delta_0] = pp_c$, where p_c is the unconditional probability that the typical nodes successfully transmits (calculated previously, at least for the Bipolar model). Thus $\mathsf{E}^0[e_0\delta_0] = pp_c$ is the mean throughput of the typical node (average of temporal throughputs over a large population of nodes).

Local Delay / Throughput Paradox

Recall, $\mathsf{E}^0[L_0 \,|\, \mathcal{S}] = 1/\pi_c(X_0, \mathcal{S}) = 1/\mathsf{E}^0[e_0\delta_0 \,|\, \mathcal{S}]$ and thus

$$\mathsf{E}^0[L_0] = \mathsf{E}^0\Big[\frac{1}{\mathsf{E}^0[e_0\delta_0 \,|\, \mathcal{S}]}\Big] \geq \frac{1}{\mathsf{E}^0[e_0\delta_0]}\,.$$

by Jensen's inequality.

Note that $\mathsf{E}^0[e_0\delta_0] = pp_c$, where p_c is the unconditional probability that the typical nodes successfully transmits (calculated previously, at least for the Bipolar model). Thus $\mathsf{E}^0[e_0\delta_0] = pp_c$ is the mean throughput of the typical node (average of temporal throughputs over a large population of nodes).

Thus for the typical node analysis

$$\text{mean delay} \geq \frac{1}{\text{mean throughput}}$$

— a consequence of the fact that two-layer averaging.

Two extremal cases

- Fast varying network. No static elements; $\mathcal{S} = \emptyset$. Fast fading and noise. Even locations of nodes are re-sampled independently across the time slots (node mobility on the time scale of MAC! (DTN).

$$\mathsf{E}^0[L_0] = \mathsf{E}^0\left[\frac{1}{\mathsf{E}^0[e_0\delta_0 \mid \emptyset]}\right] = \frac{1}{\mathsf{E}^0[e_0\delta_0]} = \frac{1}{pp_c}.$$

In other words mean delay $= \frac{1}{\text{mean throughput}}$.

Under very mild assumptions mean throughput is non-null $\mathsf{E}^0[e_0\delta_0] > 0$ and thus mean local delay is finite $\mathsf{E}^0[L_0] < \infty \Rightarrow$ $\boxed{\text{spatial stability}}$.

- Completely static scenario. No time variability. Even MAC decisions of nodes do not change across the time slots (unrealistic!). Then

$$\mathsf{E}^{\mathbf{0}}[\boldsymbol{L_0}] = \mathsf{E}^{\mathbf{0}}\left[\frac{\mathbf{1}}{\mathsf{E}^{\mathbf{0}}[\boldsymbol{e_0\delta_0} \,|\, \boldsymbol{\mathcal{S}}]}\right] = \mathsf{E}^{\mathbf{0}}\left[\frac{\mathbf{1}}{\boldsymbol{e_0\delta_0}}\right]$$

  because the conditioning on $\boldsymbol{\mathcal{S}}$ determines MAC and SINR $\boldsymbol{e_0\delta_0}$ in this case.
  Under very mild assumptions (e.g. if $\boldsymbol{p < 1}$) $\boldsymbol{e_0\delta_0 = 0}$ with positive probability, making mean local delay of the typical node $\mathsf{E}^{\mathbf{0}}[\mathbf{L_0}] = \infty \Rightarrow$ $\boxed{\text{spatial instability}}$.
  The mean throughput of the typical node may be still positive $\mathsf{E}^{\mathbf{0}}[\boldsymbol{e_0\delta_0}] > \mathbf{0}$.

First Conclusions

Having seen the above two extremal cases, it is not difficult to understand that

the mean local delay of the typical node very much depends on

how much the time-variability

"averages out" in $\mathsf{E}[\ldots \mid \mathcal{S}]$

the spatial irregularities of the distribution of nodes

in the MANET.

First Conclusions

Having seen the above two extremal cases, it is not difficult to understand that

> the mean local delay of the typical node very much depends on
>
> how much the time-variability
>
> "averages out" in E[... | $\mathcal{S}$]
>
> the spatial irregularities of the distribution of nodes
>
> in the MANET.

In the remaining part we will give results regarding several particular receiver and space-time scenarios. The inequality mean delay $\geq 1/$mean throughput is in general strict.

Moreover, we will have in quite natural scenarios mean delay $= \infty$ while mean throughput > 0.

DETAILED ANALYSIS

Slow Noise and Fading, Bipolar Receivers

Only MAC decisions vary in time. Receives are all in fixed distance r from Poisson MANET nodes.

If $p > 0$ and the distribution of fading F and noise W is such that P$\{ WTl(r) > F \} > 0$, then there are MANET nodes which have null throughout and infinite local delay. In particular E$^{0}[\mathrm{L}_0] = \infty$.

Slow Noise and Fading, Bipolar Receivers

Only MAC decisions vary in time. Receives are all in fixed distance $\boldsymbol{r}$ from Poisson MANET nodes.

If $\boldsymbol{p} > \mathbf{0}$ and the distribution of fading $\boldsymbol{F}$ and noise $\boldsymbol{W}$ is such that P$\{\, \boldsymbol{W T l(r)} > \boldsymbol{F} \,\} > \mathbf{0}$, then there are MANET nodes which have null throughout and infinite local delay. In particular E$^{\mathbf{0}}[\mathbf{L_0}] = \infty$.

proof:

$$\begin{aligned}
\boldsymbol{\pi_c(X_0, \mathcal{S})} &= \mathsf{E}^{\mathbf{0}}[\boldsymbol{e_0 \delta_0} \,|\, \boldsymbol{\mathcal{S}}] \\
&= \mathsf{P}^{\mathbf{0}}\{\boldsymbol{e_0 F_0} \geq \boldsymbol{T l(r)(W_0 + I_0)} \,|\, \boldsymbol{\mathcal{S}}\} \\
&\leq \boldsymbol{p} \mathbf{1}(\boldsymbol{F_0^0} \geq \boldsymbol{T l(r) W}) \,.
\end{aligned}$$

The last indicator is equal to 0 with non-null probability.

Fast Fading case — Technical Lemma

Denote by $\mathcal{L}_I(\xi \mid \Phi) = \mathsf{E}[e^{-\xi I} \mid \Phi]$ the conditional Laplace transform of the interference I given Poisson pattern of emitting nodes Φ. Then

$$\mathsf{E}\Big[\frac{1}{\mathcal{L}_I(\xi \mid \Phi)}\Big] = \exp\Big\{-2\pi\alpha \int_0^\infty v\Big(1 - \frac{1}{\mathcal{L}_{eF}(\xi/l(v))}\Big)\,\mathrm{d}v\Big\}\,,$$

where $\mathcal{L}_{eF}$ is the Laplace transform of the product of the MAC indicator and Fading.

Fast Rayleigh Fading, Bipolar Receivers

The mean local delay of the typical node is equal to

$$\mathsf{E}^{0}[\mathrm{L}_0] = \frac{1}{p}\mathcal{D}_W(Tl(r)) \exp\left\{2\pi p\lambda \int_0^\infty \frac{vTl(r)}{l(v)+(1-p)Tl(r)}\,\mathrm{d}v\right\},$$

where

- $\mathcal{D}_W(s) = \mathcal{D}_W^{slow}(s) = \mathcal{L}_W(-s)$ for the slow noise case,
- $\mathcal{D}_W(s) = \mathcal{D}_W^{fast}(s) = 1/\mathcal{L}_W(s)$ for the fast noise case.

Fast Rayleigh Fading, Bipolar Receivers

The mean local delay of the typical node is equal to

$$\mathrm{E}^0[\mathrm{L}_0] = \frac{1}{p}\mathcal{D}_W(Tl(r))\exp\left\{2\pi p\lambda\int_0^\infty \frac{vTl(r)}{l(v)+(1-p)Tl(r)}\,\mathrm{d}v\right\},$$

where

- $\mathcal{D}_W(s) = \mathcal{D}_W^{slow}(s) = \mathcal{L}_W(-s)$ for the slow noise case,
- $\mathcal{D}_W(s) = \mathcal{D}_W^{fast}(s) = 1/\mathcal{L}_W(s)$ for the fast noise case.

Mean local delay of the typical node is always finite in fast noise scenario.

Fast Rayleigh Fading, Fast Noise, IPNR

(Receivers are nearest nodes in some external to MANET, Poisson set of potential receivers of density λ_0.)

$$\mathrm{E}^0[\mathrm{L}] = \frac{2\pi\lambda_0}{p}\int_0^\infty r e^{-\pi\lambda_0 r^2}\mathcal{D}_W(\mu T l(r))\,\mathcal{D}_I^{INR}(\mu T l(r))\,\mathrm{d}r$$

where

$$\mathcal{D}_I^{INR}(s) = \exp\left\{2\pi\lambda\int_0^\infty \frac{ps}{l(v)+(1-p)s}v\,\mathrm{d}v\right\}$$

and $\mathcal{D}_W(s)$ is as in Bipolar Model.

Fast Rayleigh Fading, Fast Noise, IPNR

(Receivers are nearest nodes in some external to MANET, Poisson set of potential receivers of density $\boldsymbol{\lambda_0}$.)

$$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \frac{2\pi\lambda_0}{p}\int_0^\infty r e^{-\pi\lambda_0 r^2}\mathcal{D}_W(\mu T l(r))\,\mathcal{D}_I^{INR}(\mu T l(r))\,\mathrm{d}r$$

where

$$\mathcal{D}_I^{INR}(s) = \exp\left\{2\pi\lambda\int_0^\infty \frac{ps}{l(v)+(1-p)s}v\,\mathrm{d}v\right\}$$

and $\mathcal{D}_W(s)$ is as in Bipolar Model.

Nose limited case. (When interference is perfectly canceled out.) $\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$ provided the noise $\boldsymbol{W}$ has a sufficient probability mass in the neighborhood of 0. For instance, when it has rational Laplace transform.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$ e.g. when the noise $\boldsymbol{W} > \epsilon$ is bounded away from 0.

Interference limited case. ($\boldsymbol{W = 0}$) Denote

$$\boldsymbol{\theta(p, T, \beta) = \frac{p}{(1-p)^{1-\frac{2}{\beta}}} T^{\frac{2}{\beta}} K(\beta)}\ ,$$

where $\boldsymbol{K(\beta) = \frac{2\pi\Gamma(2/\beta)\Gamma(1-2/\beta)}{\beta}}$.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$ when $\boldsymbol{p > 0}$ and $\boldsymbol{\lambda_0 \pi > \lambda\theta(p, T, \beta)}$.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$ if either $\boldsymbol{p = 0}$ or $\boldsymbol{\lambda_0 \pi < \lambda\theta(p, T, \beta)}$.

Interference limited case. ($W = 0$) Denote

$$\theta(p, T, \beta) = \frac{p}{(1-p)^{1-\frac{2}{\beta}}} T^{\frac{2}{\beta}} K(\beta)\,,$$

where $K(\beta) = \frac{2\pi\Gamma(2/\beta)\Gamma(1-2/\beta)}{\beta}$.

$\mathsf{E}^0[\mathrm{L}] < \infty$ when $p > 0$ and $\lambda_0\pi > \lambda\theta(p, T, \beta)$.

$\mathsf{E}^0[\mathrm{L}] = \infty$ if either $p = 0$ or $\lambda_0\pi < \lambda\theta(p, T, \beta)$.

- Obviously $p = 0$ makes $\mathsf{E}^0[\mathrm{L}] < \infty$.

Interference limited case. ($\boldsymbol{W = 0}$) Denote

$$\boldsymbol{\theta(p, T, \beta) = \frac{p}{(1-p)^{1-\frac{2}{\beta}}} T^{\frac{2}{\beta}} K(\beta)} \, ,$$

where $\boldsymbol{K(\beta) = \frac{2\pi\Gamma(2/\beta)\Gamma(1-2/\beta)}{\beta}}$.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$ when $\boldsymbol{p > 0}$ and $\boldsymbol{\lambda_0 \pi > \lambda\theta(p, T, \beta)}$.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$ if either $\boldsymbol{p = 0}$ or $\boldsymbol{\lambda_0 \pi < \lambda\theta(p, T, \beta)}$.

- Obviously $\boldsymbol{p = 0}$ makes $\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$.
- If $\boldsymbol{\lambda_0}$ is too small the transmitters compete for too small set of receivers making $\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$.

Interference limited case. $(W = 0)$ Denote

$$\theta(p, T, \beta) = \frac{p}{(1-p)^{1-\frac{2}{\beta}}} T^{\frac{2}{\beta}} K(\beta) \,,$$

where $K(\beta) = \frac{2\pi\Gamma(2/\beta)\Gamma(1-2/\beta)}{\beta}$.

$\mathsf{E}^0[L] < \infty$ when $p > 0$ and $\lambda_0 \pi > \lambda\theta(p, T, \beta)$.

$\mathsf{E}^0[L] = \infty$ if either $p = 0$ or $\lambda_0 \pi < \lambda\theta(p, T, \beta)$.

- Obviously $p = 0$ makes $\mathsf{E}^0[L] < \infty$.
- If λ_0 is too small the transmitters compete for too small set of receivers making $\mathsf{E}^0[L] = \infty$.
- For T and β fixed, $\mathsf{E}^0[L] < \infty$ requires that potential receivers outnumber MANET nodes by a factor which grows like $p(1-p)^{2/\beta - 1}$ when p varies.

Fast Rayleigh Fading, Fast Noise, MNN

(Receivers are nearest nodes in the MANET).

$$\mathsf{E}^{0}[\mathbf{L}] = \frac{2\pi\lambda}{p(1-p)}$$

$$\times \int_0^\infty r e^{-\pi\lambda r^2} \mathcal{D}_W(\mu T l(r))\, \mathcal{D}_I^{MNN}(r, \mu T l(r))\, \mathrm{d}r\,,$$

where

$$\mathcal{D}_I^{MNN}(r, s) = \exp\Big\{\lambda\pi \int_0^\infty \frac{ps}{l(v) + (1-p)s} v\, \mathrm{d}v$$

$$+\lambda \int_{\theta=-\frac{\pi}{2}}^{\frac{\pi}{2}} \int_{v>2r\cos\theta} \frac{ps}{l(v) + (1-p)s} v\, \mathrm{d}v\mathrm{d}\theta\Big\}$$

and $\mathcal{D}_I(s)$ is as in Bipolar Model.

We have the same type of phase transitions as for the IPNR model.

Nose limited case.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$ provided the noise W has a sufficient probability mass in the neighborhood of 0. For instance, when it has rational Laplace transform.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$ e.g. when the noise $W > \epsilon$ is bounded away from 0.

We have the same type of phase transitions as for the IPNR model.

Nose limited case.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$ provided the noise $\boldsymbol{W}$ has a sufficient probability mass in the neighborhood of 0. For instance, when it has rational Laplace transform.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$ e.g. when the noise $\boldsymbol{W} > \boldsymbol{\epsilon}$ is bounded away from 0.

Interference limited case.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] < \infty$ if $\boldsymbol{p} > \mathbf{0}$ and $\boldsymbol{\theta}(\boldsymbol{p}, \boldsymbol{T}, \boldsymbol{\beta}) < \boldsymbol{\pi}$.

$\mathsf{E}^{\mathbf{0}}[\mathbf{L}] = \infty$ if $\boldsymbol{\theta}(\boldsymbol{p}, \boldsymbol{T}, \boldsymbol{\beta}) > \mathbf{2}\boldsymbol{\pi}$.

LOCAL DELAYS
REMARKS AND CONCLUSIONS

When Spatial Instability?

The spatial instability is observed

- in Bipolar (fixed-distance) receiver model only if the noise is slow (not varying in time) and the receiver distance or SINR threshold is tuned too large.

When Spatial Instability?

The spatial instability is observed

- in Bipolar (fixed-distance) receiver model only if the noise is slow (not varying in time) and the receiver distance or SINR threshold is tuned too large.
- in IPNR and MNN (nearest neighbour receiver models) even in fast noise and fading scenario when

When Spatial Instability?

The spatial instability is observed

- in Bipolar (fixed-distance) receiver model only if the noise is slow (not varying in time) and the receiver distance or SINR threshold is tuned too large.
- in IPNR and MNN (nearest neighbour receiver models) even in fast noise and fading scenario when
  - noise does not take sufficiently often values close to 0 (e.g. for constant non-null noise)

When Spatial Instability?

The spatial instability is observed

- in Bipolar (fixed-distance) receiver model only if the noise is slow (not varying in time) and the receiver distance or SINR threshold is tuned too large.
- in IPNR and MNN (nearest neighbour receiver models) even in fast noise and fading scenario when
  - noise does not take sufficiently often values close to 0 (e.g. for constant non-null noise)
  - potential receivers do not sufficiently outnumber the emitter (MAC probability p tuned too large in MNN),

When Spatial Instability?

The spatial instability is observed

- in Bipolar (fixed-distance) receiver model only if the noise is slow (not varying in time) and the receiver distance or SINR threshold is tuned too large.
- in IPNR and MNN (nearest neighbour receiver models) even in fast noise and fading scenario when
  - noise does not take sufficiently often values close to 0 (e.g. for constant non-null noise)
  - potential receivers do not sufficiently outnumber the emitter (MAC probability p tuned too large in MNN),
  - the SINR threshold is tuned too large (in MNN model) with respect to MAC probability p and path-loss exponent β.

Delay in “RESTART” Algorithm

[Jelenkovic et-al 2007:] A file of random size B is to be transmitted over an error prone channel, with i.i.d. inter-failure times $A_1, A_2, \ldots$. If $A_i < B$ the transmission fails at the i th attempt and needs to restart (with the same B) until $A_j \geq B$. Let $N = \inf\{n \geq 1 \ s.t. \ A_n > B\}$ be the transmission delay of a tagged file.

Delay in “RESTART” Algorithm

[Jelenkovic et-al 2007:] A file of random size B is to be transmitted over an error prone channel, with i.i.d. inter-failure times $A_1, A_2, \ldots$. If $A_i < B$ the transmission fails at the i th attempt and needs to restart (with the same B) until $A_j \geq B$. Let $N = \inf\{n \geq 1 \; s.t. \; A_n > B\}$ be the transmission delay of a tagged file.

[Asmussen et al 2008:] When B has infinite support and A_n is light tailed, then N is heavy tailed including finite mean.

The physical phenomena at hand are quite different in the above RESTART algorithm and our MANET context. Nevertheless

The local delays in our MANET can be seen as instances of RESTART algorithm with variable file size replaced by spatial variability of channel conditions.

In particular, in MNN model it is the variable distance to the (nearest) receiver in conjunction with existence of big void regions in Poisson MANETS which may lead to infinite mean local delays.

The physical phenomena at hand are quite different in the above RESTART algorithm and our MANET context. Nevertheless

The local delays in our MANET can be seen as instances of RESTART algorithm with variable file size replaced by spatial variability of channel conditions.

In particular, in MNN model it is the variable distance to the (nearest) receiver in conjunction with existence of big void regions in Poisson MANETS which may lead to infinite mean local delays.

The heavy tailedness of the local delay interpreted in terms of the large fraction of MANET nodes experiencing large mean temporal local delay.

Finite Mean Local Delays and Diversity

Suggested ways of getting finite mean spatial average of the mean local delays are based on an increase of diversity:

- more variability in fading,
- more potential receivers,
- more mobility,
- more flexible (adaptive) coding schemes to break the RESTART algorithm (outage) logic.

FIRST PASSAGE PERCOLATION ON SPACE-TIME SINR RANDOM GRAPHS

OR

END-TO-END DELAYS ON LONG SOURCE-DESTINATION ROUTES

Setting

In this part we are interested in the performance of multihop routing schemes.

Setting

In this part we are interested in the performance of multihop routing schemes.

The space-time network model is as for local delays in Part II (Poisson repartition of nodes, slotted Aloha MAC). We restrict ourselves to the most favorable (from the point of view local delays) fast fading and fast noise scenario.

Setting

In this part we are interested in the performance of multihop routing schemes.

The space-time network model is as for local delays in Part II (Poisson repartition of nodes, slotted Aloha MAC). We restrict ourselves to the most favorable (from the point of view local delays) fast fading and fast noise scenario.

As before we consider SINR condition for the successful transmission.

Setting

In this part we are interested in the performance of multihop routing schemes.

The space-time network model is as for local delays in Part II (Poisson repartition of nodes, slotted Aloha MAC). We restrict ourselves to the most favorable (from the point of view local delays) fast fading and fast noise scenario.

As before we consider SINR condition for the successful transmission.

In contrast to previously considered receiver models (in particular to MANET Nearest Neighbor (MNN) receiver model), we do not prescribe any receives to emitters but consider all non-emitting at a given time nodes as potential receivers of any emitting node.

Space-time SINR Random Graph

Not-specifying particular receivers allows us to “trace” all possible paths (routs) of packets on the corresponding space-time SINR random graph.

Space-time SINR Random Graph

Not-specifying particular receivers allows us to “trace” all possible paths (routs) of packets on the corresponding space-time SINR random graph.

Nodes of this graphs are all pairs

$$(\text{a node } X_i \text{ of the network} \;\;,\;\; \text{a time slot } n).$$

Space-time SINR Random Graph

Not-specifying particular receivers allows us to “trace” all possible paths (routs) of packets on the corresponding space-time SINR random graph.

Nodes of this graphs are all pairs

$$(\text{a node } X_i \text{ of the network} \quad , \quad \text{a time slot } n).$$

Directed edges of this oriented graph connect

- all pairs $(X_i, n) \to (X_j, n+1)$ whenever X_i can successfully send packet to X_j at slot n,
- and all pairs $(X_i, n) \to (X_i, n+1)$,

i.e. all possible moves of a tagged packet from X_i at time n.

Space-time SINR Random Graph

Not-specifying particular receivers allows us to "trace" all possible paths (routs) of packets on the corresponding space-time SINR random graph.

Nodes of this graphs are all pairs

$$(\text{a node } X_i \text{ of the network} \quad , \quad \text{a time slot } n).$$

Directed edges of this oriented graph connect

- all pairs $(X_i, n) \to (X_j, n+1)$ whenever X_i can successfully send packet to X_j at slot n,
- and all pairs $(X_i, n) \to (X_i, n+1)$,

i.e. all possible moves of a tagged packet from X_i at time n.

Studying shortest paths on the above graph provide intrinsic performance limitations on all possible routing schemes.

First-Passage Percolation Problem

Our main performance characteristic is the limit of the ratio

$$\frac{\text{minimal number of hops to go from node } O \text{ to node } D}{\text{Euclidean distance } |O - D|}$$

when $|O - D| \to \infty$.
(This limit, if exists, is called “time constant” in the classical first-passage-percolation models.)

First-Passage Percolation Problem

Our main performance characteristic is the limit of the ratio

$$\frac{\text{minimal number of hops to go from node } O \text{ to node } D}{\text{Euclidean distance } |O - D|}$$

when $|O - D| \to \infty$.
(This limit, if exists, is called “time constant” in the classical first-passage-percolation models.)

The number of hops in the numerator above, called end-to-end delay (from O to D), is the sum of the local delays at all nodes visited on the shortest-time path by some tagged packet, which does not experience any queuing at nodes before being scheduled for transmission.

Two Results

1. In Poisson MANET the end-to-end delay grows faster than the distance $|O - D|$ (time constant is infinite) (principally due to large voids in the repartition of nodes).

Two Results

1. In Poisson MANET the end-to-end delay grows faster than the distance $|O - D|$ (time constant is infinite) (principally due to large voids in the repartition of nodes).
2. Adding an arbitrarily sparse, periodic infrastructure of nodes (superposing it with Poisson p.p.) makes end-to-end delay scale linearly with $|O - D|$ (time constant positive and finite).

But I’m afraid...,

But I’m afraid...,
here I do not have enough time
to precisely formulate the above results
and present their elegant proofs...

But I'm afraid...,
here I do not have enough time
to precisely formulate the above results
and present their elegant proofs...

Please consult Baccelli, B.B, Mirsadeghi (2009) "Optimal Paths on the Space-Time SINR Random Graph", arXiv:0911.3721v1:

THANK YOU