Transdimensional Modeling and Reliability Analysis of Vehicular Networks on General Street Systems

Martin Haenggi
Joint work with Jeya Pradha Jeyaraj (now at Qualcomm)

University of Notre Dame, IN, USA
ETH Zürich, Switzerland

LINCS, Paris

Mar. 10, 2022

Available at http://www.nd.edu/~mhaenggi/talks/lincs22.pdf

Overview

The two goals of this talk

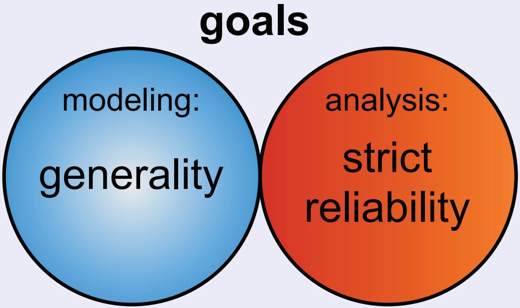

Overview

The two goals and the two approaches

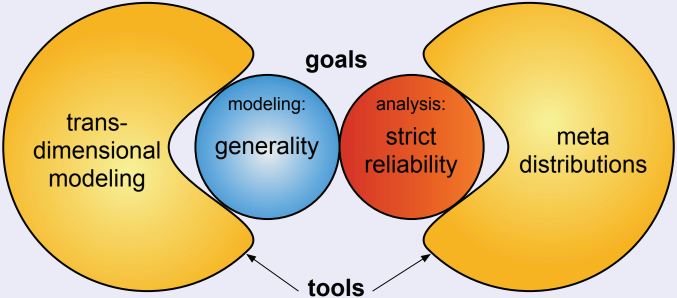

This work is based on:

- J. P. Jeyaraj and MH, "Cox Models for Vehicular Networks: SIR Performance and Equivalence", *IEEE Trans. on Wireless Comm.*, vol. 20, pp. 171-185, Jan. 2021.
- J. P. Jeyaraj, MH, A. H. Sakr, and H. Lu, "The Transdimensional Poisson Process for Vehicular Network Analysis", *IEEE Trans. on Wireless Comm.*, vol. 20, pp. 8023-8038, Dec. 2021.
- MH, "Meta Distributions—Part 1: Definition and Examples", *IEEE Comm. Letters*, vol. 25, pp. 2089-2093, July 2021.

  MH, "Meta Distributions—Part 2: Properties and Interpretations", *IEEE Comm. Letters*, vol. 25, pp. 2094-2097, July 2021.

Typical street system

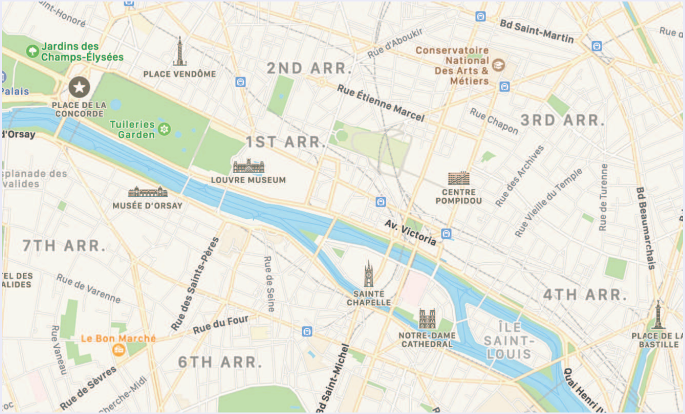

Real street systems are complicated. We need to find an abstraction that captures the pertinent features, including streets of different lengths and intersections of different order (3-way or T-junctions, 4-way, 5-way etc.).

Definition (Street system)

$\mathcal{S}$ is a stationary random measurable subset of $\mathbb{R}^2$ with $|\mathcal{S}|_2 = 0$ that contains no singletons (isolated) points a.s. By stationarity,

$$\mathbb{E}|\mathcal{S} \cap B|_1 = \tau |B|_2 \quad \text{for Borel } B.$$

τ is the mean total street length per unit area.

Example (Poisson line process (PLP))

Letting $\Phi = \{(x_i, \varphi_i)\} \subset \mathbb{R} \times [0, \pi]$ be a homogeneous Poisson point process (PPP), place lines on the plane at distances x_i from o with orientation φ_i.

Example (Poisson stick process (PSP))

At the locations of a PPP $\Phi \subset \mathbb{R}^2$, place line segments of iid half-lengths with pdf f_H and random orientation.

Vehicle locations

Vehicles form a Cox process $\mathcal{V}$ supported on $\mathcal{S}$. It is stationary with random intensity measure $\Xi(B) = \lambda|\mathcal{S} \cap B|_1$ and density $\tau\lambda$.

PLP-PPP and PSP-PPP

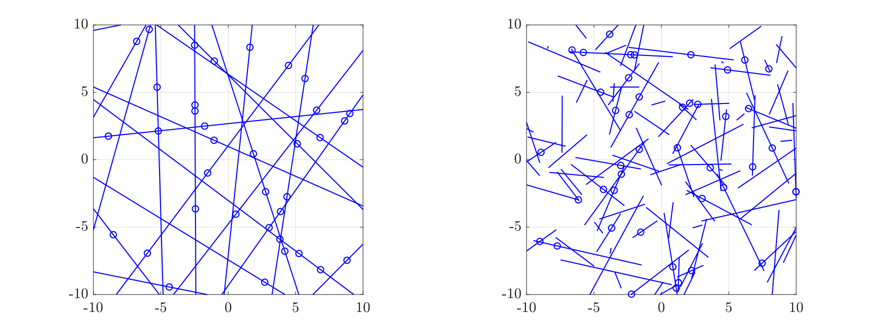

PLP-PPP and PSP-PPP.

Both models only have standard (four-way) intersections.

Street system decomposition

Let

$$\mathcal{P}_m \triangleq \{z \in \mathbb{R}^2 \colon |\mathcal{S} \cap b(z,r)|_1 \sim mr, \;\; r \to 0\}$$

be the set of points of order $m \in \mathbb{N}$ in the street system $\mathcal{S}$.
$\mathcal{P}_1$ are endpoints of streets, $\mathcal{P}_2$ interior (general) points, $\mathcal{P}_3$ T-junctions, $\mathcal{P}_4$ four-way intersections, etc.
$\{\mathcal{P}_m\}_{m\in\mathbb{N}}$ forms a partition of $\mathcal{S}$, with $\mathcal{P}_2 = \mathcal{S}$ a.e. ($\mathcal{V} \subset \mathcal{P}_2$ a.s.) and $\mathcal{P}_m$, $m \neq 2$, countable and forming a simple a stationary point process.

Street system characterization

Let

$$\mathcal{M} \triangleq \{m \in \mathbb{N} \colon \mathbb{P}(\mathcal{P}_m = \emptyset) = 0\}$$

be the index set of the non-empty components. Then we call $\mathcal{S} = \bigcup_{m\in\mathcal{M}} \mathcal{P}_m$ an $\mathcal{M}$-indexed street system.
The PLP-PPP is a $(2,4)$ street system, while the PSP-PPP is a $(1,2,4)$ street system. How about T-junctions ($m = 3$)?

Poisson lilypond model (PLM)

At the points of a PPP Φ, line segments start to grow in random directions until they hit another line segment. This leads to a $(1, 2, 3)$ street system.

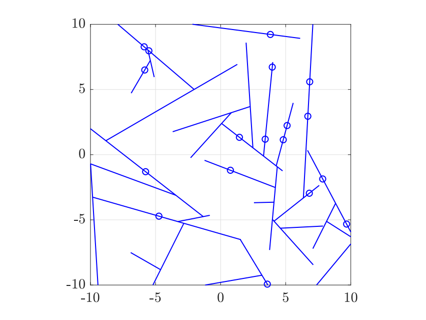

Figure: PLM-PPP

Property: In dense regions of Φ, line segments tend to be short, and vice versa. This leads to a concentration of the number of vehicles in any $B \subset \mathbb{R}^2$ (sub-Poisson distribution).

Transmission model

Vehicles transmits with probability p in each time slot. Receiving vehicles attempt to decode a message from a transmitter at fixed distance D.

SIR and reception

For a receiving vehicle at the origin, the signal-to-interference ratio (SIR) is

$$\mathsf{SIR} \triangleq \frac{gD^{-\alpha}}{\sum_{z\in\mathcal{V}} B_z g_z \|z\|^{-\alpha}},$$

where (g, g_z) are iid exponential with mean 1 and (B_z) are iid Bernoulli with mean p. We are interested in the SIR ccdf or success probability

$$\bar{F}_m(\theta) \triangleq \mathbb{P}^!_o(\mathsf{SIR} > \theta),$$

where m is the vehicle order and the Palm measure is w.r.t. the typical vehicle of order m located at the origin.

SIR distribution in PSP-PPP

In the Poisson stick Cox process (PSP-PPP) with street half-length pdf f_H,

$$\bar{F}_m(\theta) = \mathcal{L}_{I_o^m}(\theta D^\alpha)\mathcal{L}_{I_\mathrm{r}}(\theta D^\alpha),$$

where

$$\mathcal{L}_{I_o^m}(s) = \left(\int\limits_0^\infty \Big(\frac{1}{2h} \int\limits_{-h}^{h} \exp(-\lambda p s^{\delta/2} \int\limits_{(-w-h)s^{-\delta/2}}^{(-w+h)s^{-\delta/2}} \frac{1}{1+v^{\frac{2}{\delta}}} \mathrm{d}v) \mathrm{d}w \Big) \tilde{f}_H(h) \mathrm{d}h \right)^{\frac{m}{2}},$$

$$\mathcal{L}_{I_\mathrm{r}}(s) = \exp\left(-\frac{\tau}{2\pi\mathbb{E}(H)} \int\limits_0^\infty \int\limits_0^\pi \int\limits_0^{2\pi} \int\limits_0^\infty (1-\mathcal{L}_{I_a}(s))\gamma f_H(h) \mathrm{d}\gamma \mathrm{d}\phi \mathrm{d}\varphi \mathrm{d}h \right),$$

with $\mathcal{L}_{I_a}(s) = \exp\big(-\lambda p \int_{-h}^{h} (1 + (\frac{\gamma^2+u^2+2\gamma u\cos(\phi-\varphi)}{s^\delta})^{1/\delta})^{-1} \mathrm{d}u \big)$, $\delta = 2/\alpha$, and $\tilde{f}_H(h) = h f_H(h)/\mathbb{E}[H]$.

5 nested integrals—this is too complicated.

SIR distribution in PSP-PPP

In the Poisson stick Cox process (PSP-PPP) with street half-length pdf f_H,

$$\bar{F}_m(\theta) = \mathcal{L}_{I_o^m}(\theta D^\alpha)\mathcal{L}_{I_\mathrm{r}}(\theta D^\alpha),$$

where

$$\mathcal{L}_{I_o^m}(s) = \Bigg(\int_0^\infty \Big(\frac{1}{2h} \int_{-h}^{h} \exp(-\lambda p s^{\delta/2} \int_{(-w-h)s^{-\delta/2}}^{(-w+h)s^{-\delta/2}} \frac{1}{1+v^{\frac{2}{\delta}}} \mathrm{d}v) \mathrm{d}w \Big) \tilde{f}_H(h) \mathrm{d}h \Bigg)^{\frac{m}{2}},$$

$$\mathcal{L}_{I_\mathrm{r}}(s) = \exp\left(-\frac{\tau}{2\pi\mathbb{E}(H)} \int_0^\infty \int_0^\pi \int_0^{2\pi} \int_0^\infty (1 - \mathcal{L}_{I_a}(s)) \gamma f_H(h) \mathrm{d}\gamma \mathrm{d}\phi \mathrm{d}\varphi \mathrm{d}h \right),$$

with $\mathcal{L}_{I_a}(s) = \exp\big(-\lambda p \int_{-h}^{h} (1 + (\frac{\gamma^2+u^2+2\gamma u\cos(\phi-\varphi)}{s^\delta})^{1/\delta})^{-1} \mathrm{d}u \big)$, $\delta = 2/\alpha$, and $\tilde{f}_H(h) = h f_H(h)/\mathbb{E}[H]$.

5 nested integrals—this is too complicated.

Discussion

- The street geometry cannot be ignored—simply modeling vehicles as a 2D PPP is inaccurate.
- An exact analysis of the SIR distribution is possible in the case of PLP and PSP models but the results are unwieldy.
- The precise modeling of far interferers is less important than that of nearby interferers which are likely to be on the same street.

Transdimensional modeling

Model the typical vehicle's street(s) exactly (as 1D PPP(s)) and model the other vehicles as a 2D PPP.

This combination of 1D and 2D PPPs of the corresponding intensities is called a transdimensional PPP (TPPP).

TPPP of a PLP-PPP

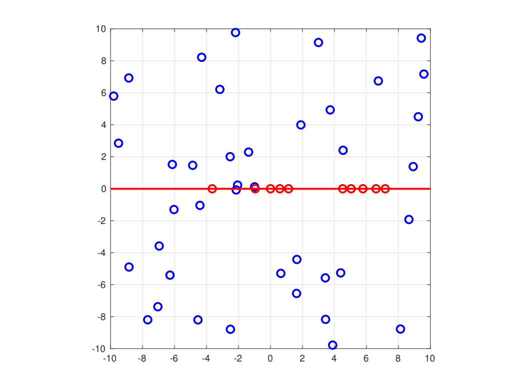

1D PPP on $(\mathbb{R}, 0)$ with density λ and 2D PPP on $\mathbb{R}^2$ with density $\lambda\tau$.

The 1D PPP is the one inherent in the Palm measure of $\{\mathrm{SIR} > \theta\}$ since the vehicle conditioned to be at the origin needs to be on a street.

Comparison of analytical expressions

SIR ccdf of the typical vehicle in the PLP-PPP

$$\bar{F}_m(\theta) = \exp(-m\lambda p D\theta^{\delta/2}\Gamma(1+\delta/2)\Gamma(1-\delta/2)) \times \exp(-2\tau \int_0^\infty (1-\mathcal{L}_{I_x}(\theta D^\alpha))\mathrm{d}x),$$

where $\mathcal{L}_{I_x}(s) = \exp\left(-\lambda p s^{\delta/2} \int_{x^2 s^{-\delta}}^\infty \frac{1}{(1+v^{1/\delta})\sqrt{v-x^2 s^{-\delta}}}\mathrm{d}v\right)$.

SIR ccdf of the typical vehicle in the corresponding TPPP

$$\bar{F}_m(\theta) = \exp(-m\lambda p D\theta^{\delta/2}\Gamma(1+\delta/2)\Gamma(1-\delta/2)) \times \exp(-\lambda p \tau \pi D^2 \theta^\delta \Gamma(1+\delta)\Gamma(1-\delta)).$$

This is closed-form and evaluated about 100,000× faster.

TPPP for PSP-PPP, general vehicle ($m = 2$), $\tau = 1$

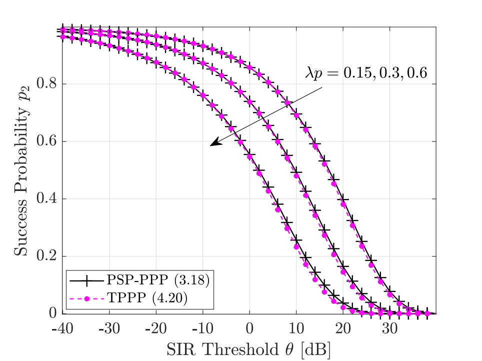

Generally, the TPPP result is asymptotically exact for both $\theta \to 0$ (1D PPP dominates) and $\theta \to \infty$ (2D PPP dominates).

TPPP of an intersection in the PSP-PPP

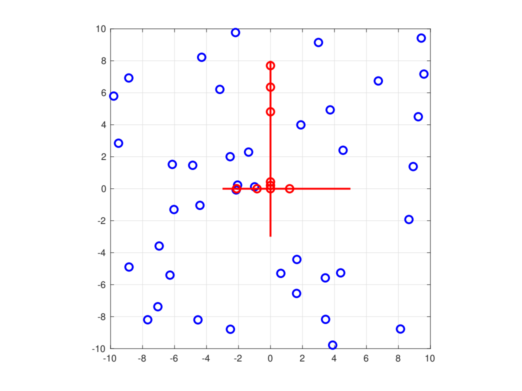

1D PPP on two line segments with density λ and 2D PPP on $\mathbb{R}^2$ with density $\lambda\tau$.

Since $m = 4$, the vehicle at o lies on two streets. Their angle is arbitrary.

TPPP for PSP-PPP, intersection vehicle ($m = 4$), $\tau = 1$

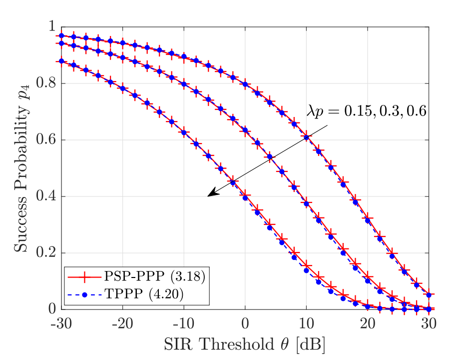

Difference PLP-PPP to TPPP for general vehicle ($m = 2$)

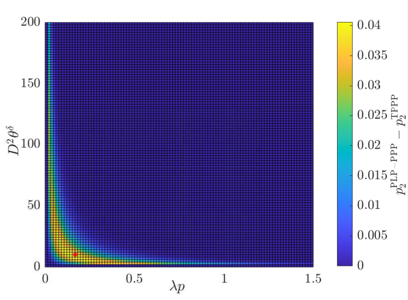

The gap is small and always positive—the TPPP approach yields a tight lower bound.

Note: Exact results with 5 nested integrals (2 over infinite range) cannot be numerically evaluated exactly in a finite amount of time. With constraints on the computational resources, the TPPP expressions may be more accurate.

The flaw of averages

Spatial averages and their interpretation

Results like

$$\mathbb{P}(\mathsf{SIR} > \theta) = 0.9$$

do not tell us anything about the fraction of links that achieve θ with probability 0.9. There is no link in the network with reliability 0.9.

They merely tell us that 90% of the links happen to achieve an SIR of θ, without any reliability information. Also, these links change in every time slot due to fading (even without macroscopic mobility).

There is no information about individual links.

Histograms and mean

Let us consider per-vehicle SIR distributions conditioned on the point process $\mathcal{V}$.

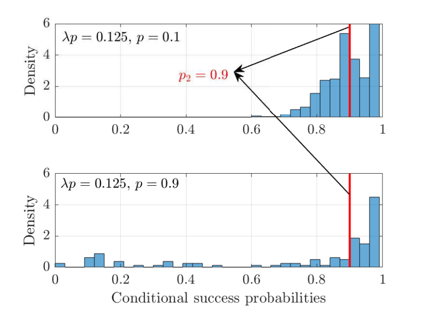

The histograms (empirical pdfs) differ vastly but $p_2 = \mathbb{P}(\mathrm{SIR} > \theta) = 0.9$ in both cases.

To guarantee a certain reliability, we need to capture the distribution of $\mathbb{P}(\mathrm{SIR} > \theta \mid \mathcal{V})$. This implies a separation of temporal and spatial randomness.

The concept of meta distributions

Let $f : \mathbb{R}^2 \mapsto \mathbb{R}$ and X, Y be random variables.

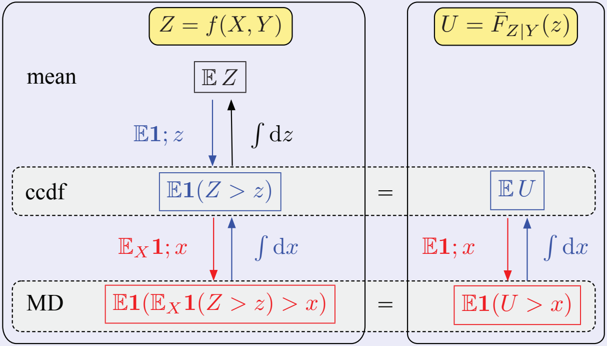

A meta distribution (MD) is the distribution of a conditional distribution. It reveals how X, Y individually affect Z. Written differently,

$$\bar{F}_{[\![Z|Y]\!]}(z,x) = \mathbb{P}(\mathbb{P}(Z > z \mid Y) > x) = \mathbb{P}(\mathbb{P}_X(Z > z) > x).$$

Toy example

MD of ratio of exponential random variables

Let X, Y be independent exponential with mean 1 and $1/\mu$, respectively. The ccdf of $Z \triangleq X/Y$ is $\bar{F}_Z(z) = \frac{\mu}{z+\mu}$ and $\mathbb{E}Z$ does not exist.
$U \triangleq \bar{F}_{Z|Y}(z) = \mathbb{E}\mathbf{1}(Z > z \mid Y) = e^{-Yz}$ and

$$\bar{F}_{[\![Z|Y]\!]}(z, x) = \mathbb{P}(U > x) = \mathbb{P}(e^{-Yz} > x) = 1 - x^{\mu/z}.$$

The ccdf of Z is retrieved by $\mathbb{E}U = \int_0^1 (1 - x^{\mu/z})\mathrm{d}x = \frac{\mu}{z+\mu}$.

This shows that when considering X and Y, the distribution of Z is just the mean $\mathbb{E}U$ while the meta distribution is the distribution of U.

Applied to vehicular networks

In vehicular networks, assuming $\mathcal{V}$ does not change during a transmission, conditioning on $\mathcal{V}$ and calculating

$$\bar{F}_{[\![\mathsf{SIR}|\mathcal{V}]\!]}(\theta, x) = \mathbb{P}^!_o(\mathbb{P}^!_o(\mathsf{SIR} > \theta \mid \mathcal{V}) > x)$$

separates time and space and reveals the per-vehicle reliability statistics.

- The MD corresponds to the fraction of vehicles that achieve $\mathsf{SIR} > \theta$ with probability at least x in each realization of $\mathcal{V}$ (in ergodic models).
- The derivative w.r.t. x corresponds to the histogram shown earlier.
- The inverse (in x) is the reliability that a certain percentile of vehicles achieve. For example, $\bar{F}^{-1}_{[\![\mathsf{SIR}|\mathcal{V}]\!]}(\theta, 0.95)$ is the reliability that 95% of the vehicles achieve.
- The standard SIR distribution is $\int_0^1 \bar{F}_{[\![\mathsf{SIR}|\mathcal{V}]\!]}(\theta, x)\mathrm{d}x$.

Toy application

Signal power in uplink V2I network

Base stations form a PPP Φ of intensity λ, and vehicles connect to the nearest one whose distance R is Rayleigh with mean $1/(2\sqrt{\lambda})$.
The signal power received is $S = h/R^2$, where h is exp. with mean 1.

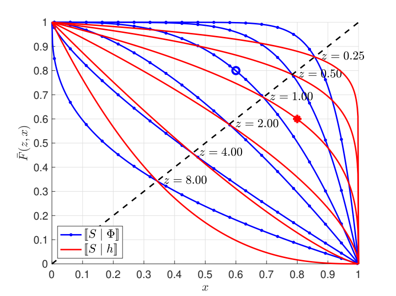

$$\bar{F}_{[\![S|\Phi]\!]}(\theta, x) = 1 - x^{\lambda\pi/\theta}$$

This is the fraction of vehicles for which $S_x > \theta$ with probability at least x, for each realization of Φ.

80% of the users achieve $S = 1$ with probability at least 0.6.

In more complete models, only the moments $M_b \triangleq \mathbb{E}[\bar{F}_{\mathsf{SIR}|\Phi}(\theta)^b]$ can be calculated.

Sidebar: Mapping moments to distributions

- Equipped with imaginary moments, we have (Gil-Pelaez),

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

  This can be tricky to evaluate. There are no closed-form solutions.
- If the moments $b \in \mathbb{N}$ are known, it is called the Hausdorff moment problem. If only a finite number of moments are known, it is truncated and has infinitely many solutions (in most cases).
- Infima and suprema can be calculated with some effort, but calculating good approximations efficiently is ongoing research.
- The beta distribution (with matching M_1 and M_2) often gives good results.
- The variance $M_2 - M_1^2$ gives basic information on how concentrated the performance is.

Moments in TPPP of PLP-PPP

For the typical vehicle of order m,

$$M_{b,m} = \exp\Big(- m\lambda p D\theta^{\delta/2}\Gamma(1+\delta/2)\Gamma(1-\delta/2)\mathcal{F}(p,\delta/2) \\ - \lambda\tau\pi D^2\theta^{\delta}\Gamma(1+\delta)\Gamma(1-\delta)\mathcal{F}(p,\delta)\Big), \quad b \in \mathbb{C},$$

where $\mathcal{F}(p,q) = pb\,{}_2F_1(1-b, 1-q; 2; p)$.

This is quasi-closed-form and quickly evaluated. The corresponding expression for the PLP-PPP is a double integral.

MDs $\bar{F}_{[\![\mathrm{SIR}|\mathcal{V}]\!]}(\theta, 0.6)$ and $\bar{F}_{[\![\mathrm{SIR}|\mathcal{V}]\!]}(\theta, 0.95)$

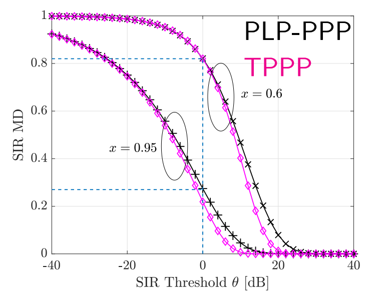

At $\theta = 1$:
$\bar{F}_{\mathrm{SIR}}(1) = 0.7$, so 70% of the links succeed.

82% of links are 60% reliable.

27% of links are 95% reliable.

Again the TPPP approach provides a tight lower bound that is asymptotically exact in the high-reliability regime.

The accuracy increases further with shadowing. As the variance $\sigma^2 \to \infty$, the TPPP result is exact.

Congestion control

Goal: Adjust the transmit probability p as a function of the vehicle density λ so that a constant reliability performance is maintained.

The standard success probability (SIR ccdf) only depends on the product λp, so choosing $p \propto 1/\lambda$ keeps the fraction of successful links constant:

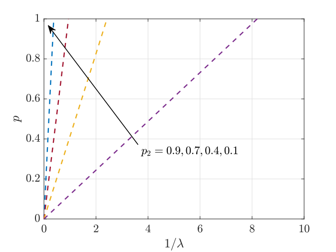

Such adaptation would result in a varying fraction of links that meet a certain reliability threshold.
Instead of $\bar{F}_{\mathsf{SIR}}(\theta)$, we should keep $\bar{F}_{[\![\mathsf{SIR}|\mathcal{V}]\!]}(\theta, x)$ constant.

MD-based congestion control

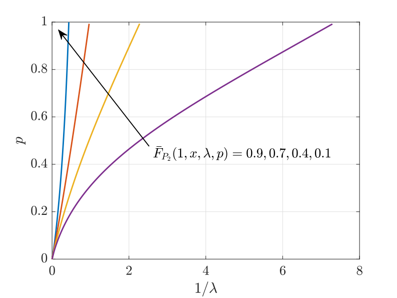

$x = 0.5$

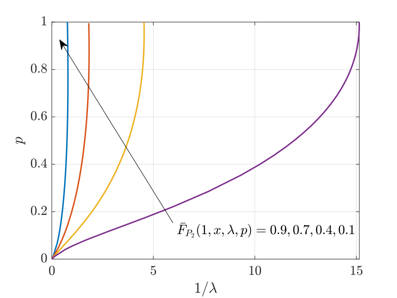

$x = 0.9$

As the reliability constraint x increases, the curves become more nonlinear. In particular, once λ exceeds a threshold, p needs to be reduced drastically with further increasing λ.

MD for broadcasting

In broadcast mode, messages are sent to several vehicles at different distances. To avoid another continuous parameter, distances are quantized into bins, say 6 bins delimited by $\{0, 50, 100, 150, 200, 250, 300\}$ m.

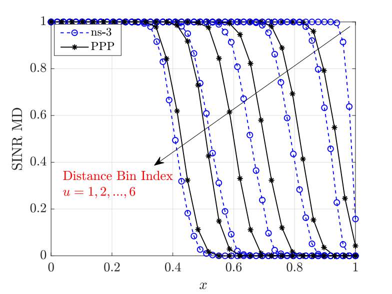

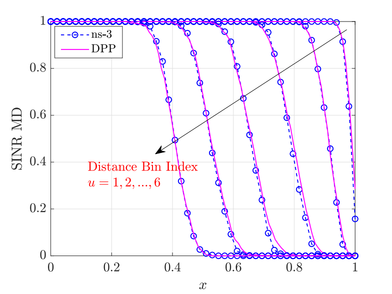

Example for binned MD. The ns3 curves show simulation results for *tLInear MEssage Rate Integrated Control* (LIMERIC) developed by Toyota. The right plot shows that the concurrently transmitting vehicles are well modeled by a determinantal point process (DPP).

Concluding remarks

- Modeling the street of the vehicle under consideration important. However, the success of the transdimensional approach reveals that modeling the streets of interfering vehicles is not critical.
- With the TPPP, general approximate results for different intersection orders can be derived (m parameter).
- The numerical evaluation of many nested integrals is never fully precise. In the high-reliability result, the TPPP-based results may be more accurate than the "exact" ones.
- The two-parameter meta distribution $\bar{F}_{[\![\mathrm{SIR}|\mathcal{V}]\!]}(\theta, x)$ addresses the inadequacy of the standard SIR ccdf $\mathbb{P}(\mathrm{SIR} > \theta)$ in quantifying the reliability of individual links. It is obtained by a Hausdorff moment transform.
- Extensions to non-Poisson deployment (non-ALOHA MACs) are possible thanks to the simplicity of the transdimensional modeling.