Palm Calculus
Part 2
Theory

JY Le Boudec

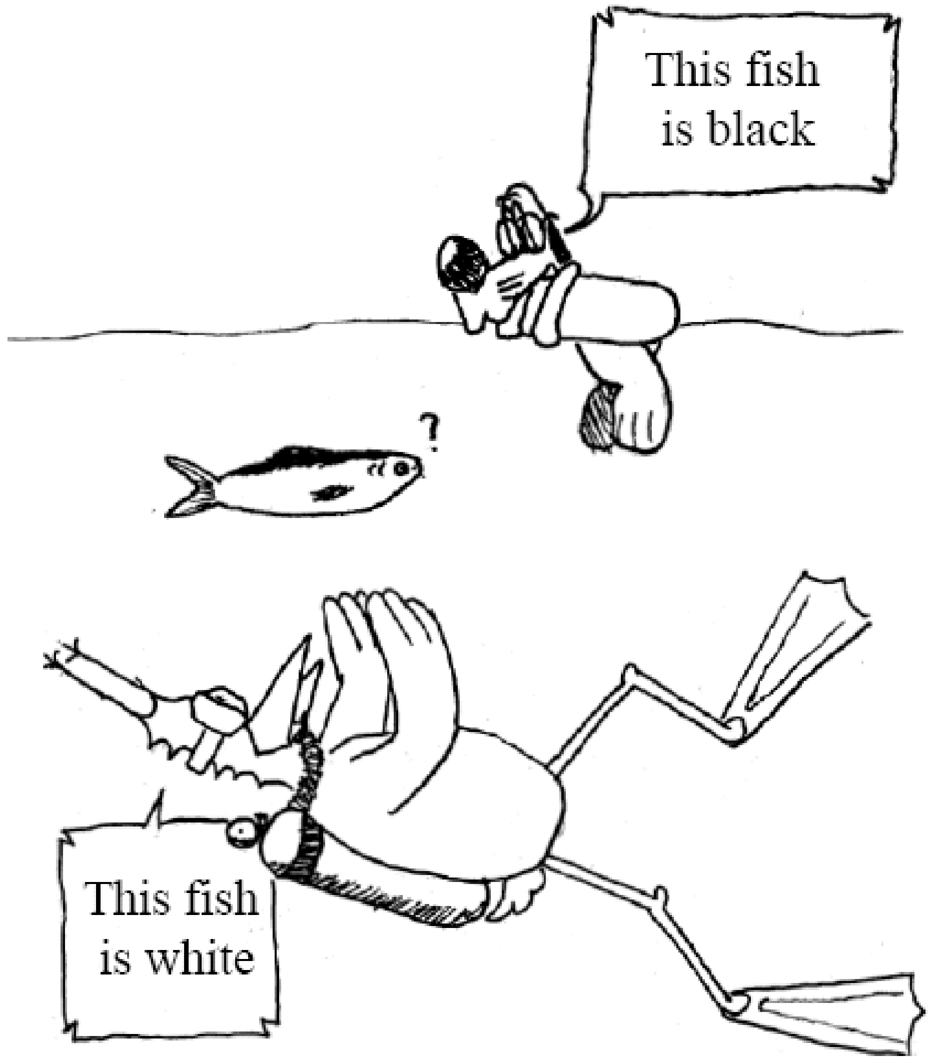

April 2019

Contents

1. Palm Calculus : Framework

A stationary process (simulation) with state S_t

Some quantity X_t measured at time t. Assume that

(S_t, X_t) is jointly stationary

I.e., S_t is in a stationary regime and X_t depends on the past, present and future state of the simulation in a way that is invariant by shift of time origin.

S_t and X_t are rigt-continuous, i.e. $X_t = X_{t^+}$ and $S_t = S_{t^+}$

Examples

- S_t= current position of mobile, speed, and next waypoint
- Jointly stationary with S_t: X_t = current speed at time t; X_t = time to be run until next waypoint
- Not jointly stationary with S_t: X_t = time at which last waypoint occurred

Arbitrary Point in Time

When X_t, S_t is jointly stationary, $E(X_t)$ is the same at all t

It represents the average seen at an arbitrary point in time

It can be shown that it is also the average seen by an external observer who observes the system at a random time, sampled from a Poisson process of any rate, independent of the simulation.

(PASTA: Poisson Arrivals See Time Averages)

Stationary Point Process

Consider some selected transitions of the simulation, occurring at times T_n.

Example: T_n= time of nth trip end

In general: given is a subset $\mathcal{F}_0 \subset \mathcal{S} \times \mathcal{S}$; we say that there is a selected transition at time t, (= an event) i.e. $t = T_n$ for some n if $(S_{t^+}, S_{t-}) \in \mathcal{F}_0$

T_n is a called a stationary point process associated to S_t

By convention, in the inversion formula:

$$\ldots < T_{-2} < T_{-1} < T_0 \le 0 < T_1 < T_2 < \cdots$$

and time $t = 0$ is the *arbitrary* point in time.

Palm Expectation

Assume: X_t, S_t are jointly stationary, T_n is a stationary point process associated with S_t

Definition : the Palm Expectation is

$$E^t(X_t) = E(X_t \mid \text{a selected transition occurs at time } t)$$

By stationarity:

$$E^t(X_t) = E^0(X_0)$$

Example:

T_n = time of nth trip end, X_t = instant speed at time t

$E^t(X_t) = E^0(X_0)$ = average speed observed at a waypoint

Take home: $E^0(\mathbf{something}) =$ average of $\mathbf{something}$, sampled at an arbitrary event time T_n

$\mathbf{E}(X_t) = \mathbf{E}(X_0)$ expresses the time average viewpoint.

$\mathbf{E}^t(X_t) = \mathbf{E}^0(X_0)$ expresses the event average viewpoint.

Example for random waypoint:

T_n = time of n^{th} trip end, X_t = instant speed at time t

$\mathbf{E}^t(X_t) = \mathbf{E}^0(X_0)$ = average speed observed at trip end

$\mathbf{E}(X_t)=\mathbf{E}(X_0)$ = average speed observed at an arbitrary point in time

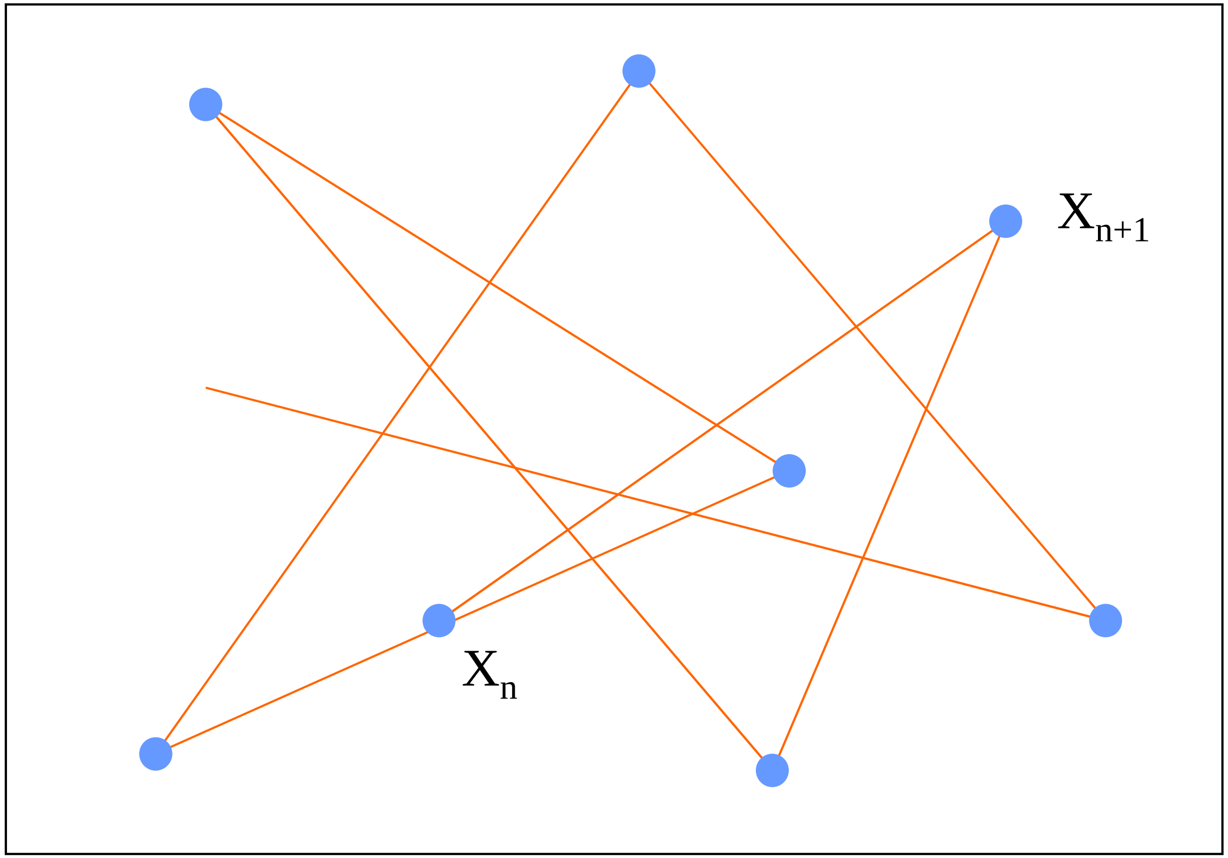

Formal Definition

In discrete time, we have an elementary conditional probability

$$\mathbb{E}^t(Y) = \mathbb{E}(Y|N(t)=1) = \frac{\mathbb{E}(YN(t))}{\mathbb{E}(N(t))} = \frac{\mathbb{E}(YN(t))}{\mathbb{P}(N(t)=1)}$$

In continuous time, the definition is a little more sophisticated

uses Radon Nikodym derivative– see lecture note for details

Also see [BaccelliBremaud87] for a formal treatment

Palm probability is defined similarly

The Palm *probability* is defined similarly, namely

$$\mathbb{P}^0(X(0) \in W) = \mathbb{P}(X(0) \in W | \text{ a point occurs at time } 0)$$

Note that $\mathbb{P}^0(T_0 = 0) = 1$, i.e., under the Palm probability, T_0 is 0 with probability 1.

Ergodic Interpretation

Assume simulation is stationary + ergodic, i.e. sample path averages converge to expectations; then we can estimate time and event averages by:

$$\mathbb{E}(X_0) = \lim_{T \to +\infty} \frac{1}{T} \sum_{s=1}^{T} X_s$$

$$\mathbb{E}^0(X_0) = \lim_{N \to +\infty} \frac{1}{N} \sum_{n=1}^{N} X_{T_n}$$

We use the same distinction for probabilities

$P^0(\text{something}) =$ proba that something happens at an arbitrary event time T_n

$P(\text{something}) =$ proba that something happens at an arbitrary time t

Intensity of a Stationary Point Process

Intensity of selected transitions: λ:= expected number of events per time unit

To estimate λ in a simulation of duration T_N with N events

$$\lambda \approx \frac{N}{T_N}$$

E.g: (Poisson process:)
events occur at times T_n such that $T_n - T_{n-1}$ ~iid Expo(λ)
(memoriless: next arrival is independent of the past)

For the Poisson process of rate λ, the intensity is also λ

E.g.: RWP times when mobiles reach a waypoint: not a Poisson process; intensity = average nb of waypoints per time unit

Palm Calculus Formula #1

Intensity Formula:

$$\frac{1}{\lambda} = \mathbb{E}^0(T_1 - T_0) = \mathbb{E}^0(T_1)$$

where by convention $T_0 \leq 0 < T_1$

Says that intensity = 1 / mean time between events

Example: Poisson process, mean time between events is the mean of Expo(λ), i.e. $\frac{1}{\lambda}$

Example: RWP: intensity is mean trip duration

The interval between 2 buses is $\sim$ $\boldsymbol{U(15, 25)}$ minutes

A. There are 2 buses in average per hour

B. There are 3 buses in average per hour

C. There are 4 buses in average per hour

D. None of the above

E. I don't know

The validity of the formula in the previous question requires that …

A. The arrival process is Poisson
B. The arrival process is stationary
C. The interarrival times are iid
D. None of the above
E. I don’t know

Palm Calculus Formula #2

Inversion Formula

$$\mathbb{E}(X_t) = \mathbb{E}(X_0) = \lambda \mathbb{E}^0 \left(\int_0^{T_1} X_s ds \right)$$

Here the quantity X_t is jointly stationary with the simulation state.

Says that the expectation of X_t, at an arbitrary point in time, is equal to $\lambda \times$ the expectation, at an arbitrary event, of the integral of X_t between two events.

Example: Joe's Waiting Time

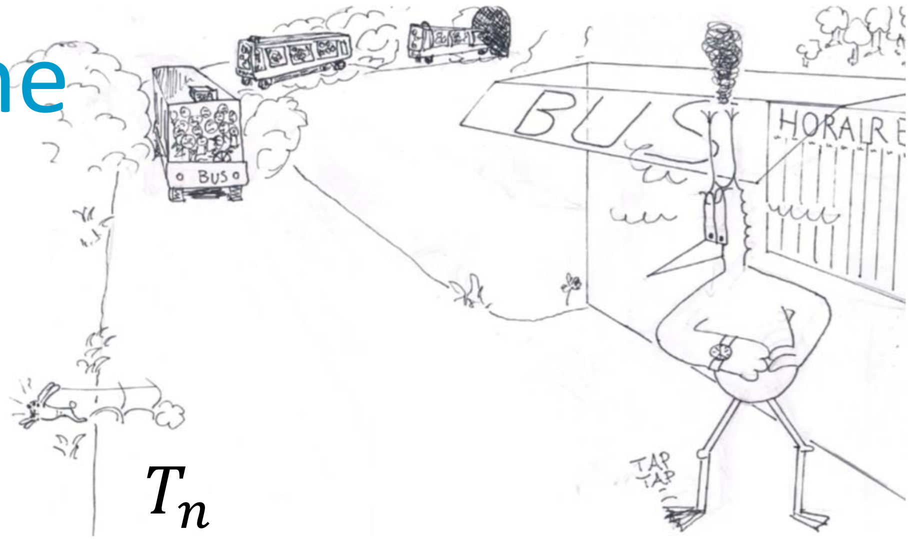

T_n: arrivals of buses

X_t = waiting time for who arrives at time t

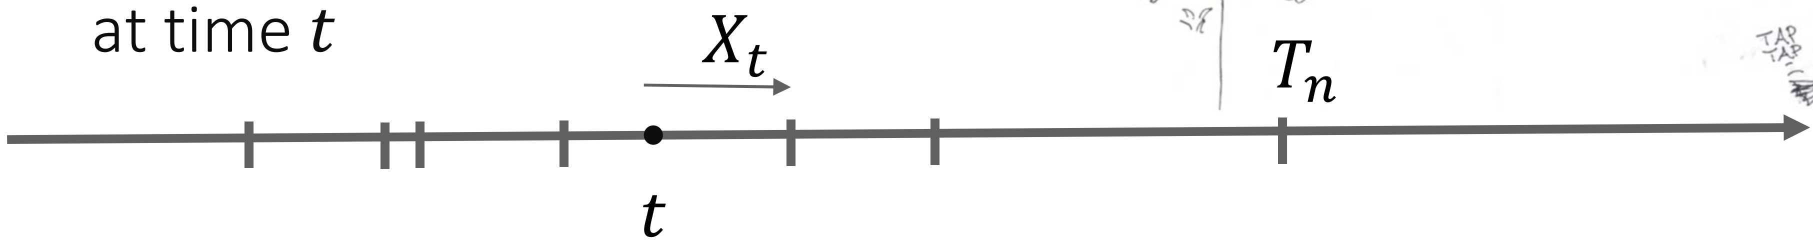

$E(X_t)$ = Joe's average waiting time = $\lambda E^0\left(\int_{0=T_0}^{T_1} X_s\, ds\right)$

For $s \in [T_0, T_1]$, $X_s = T_1 - s$ hence

$\int_{0=T_0}^{T_1} X_s\, ds = \int_{0=T_0}^{T_1} (T_1 - s) ds = \frac{1}{2}(T_1 - T_0)^2$ and

$$E(X_t) = \frac{\lambda}{2} E^0\left((T_1 - T_0)^2\right)$$

$E^0\left((T_1 - T_0)^2\right)$ is the average of the square of the interval between buses; in a simulation we estimate as $\frac{1}{N}\sum_n (T_{n+1} - T_n)^2$

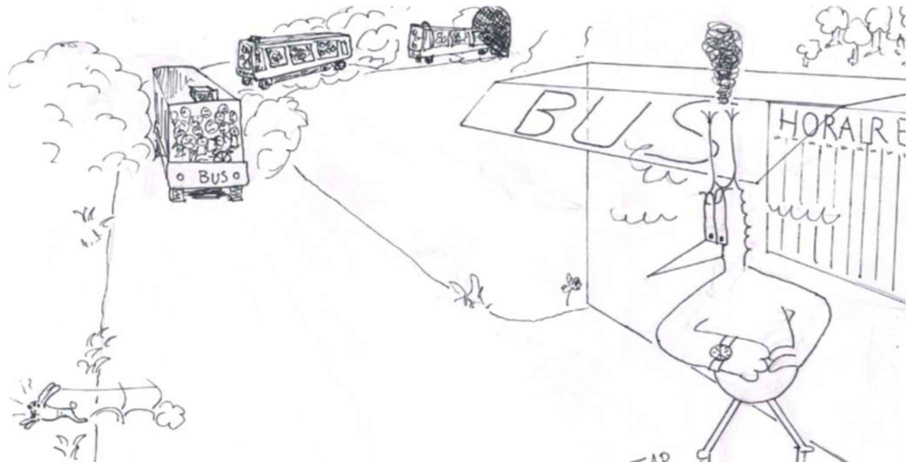

T_n: arrivals of buses

X_t = waiting time for who arrives at time t

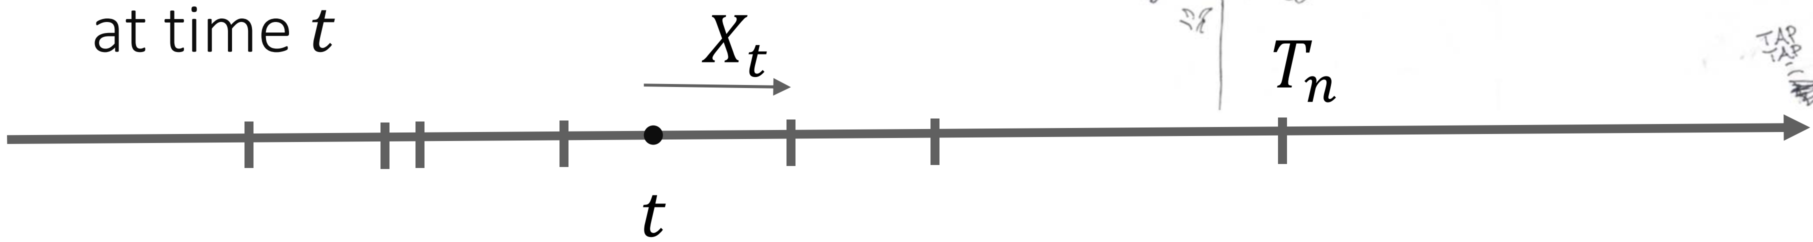

$E(X_t)$ = Joe's average waiting time $= \frac{\lambda}{2} E^0\left((T_1 - T_0)^2\right)$

Let v be the variance of the interval between buses:

$$v = E^0\left((T_1 - T_0)^2\right) - \left(E^0(T_1 - T_0)\right)^2$$

Intensity formula: $\lambda^{-1} = E^0(\,T_1 - T_0)$

Hence

$$E(X_t) = \underbrace{\frac{1}{2}\frac{1}{\lambda}}_{\text{0.5 × mean time between buses system's viewpoint}} + \underbrace{\frac{\lambda}{2} v}_{\text{penalty due to variability}}$$

Example: Gatekeeper

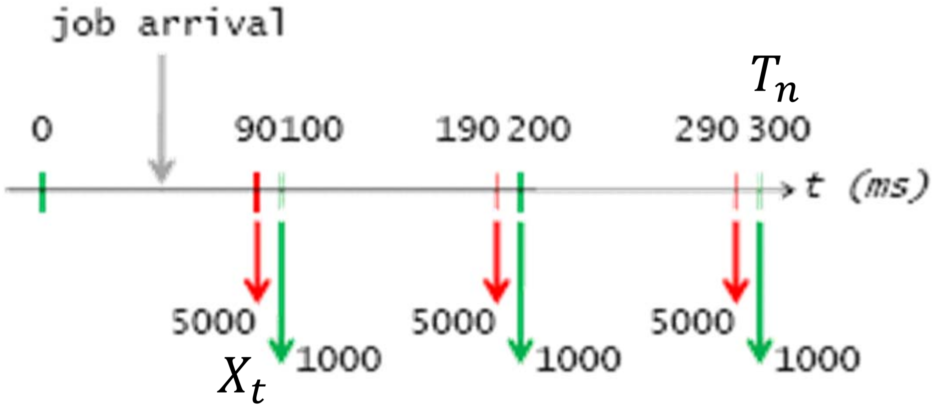

X_t = execution time of a job that would arrive at time t

T_n: wake up time of gatekeeper

$E(X_t) = E(X_0) =$ average execution time for a job that arrives at an arbitrary point in time = W_c

Inversion formula: $E(X_0) = \lambda E^0\left(\int_{T_0=0}^{T_1} X_s\, ds\right)$

$\int_{T_0=0}^{T_1} X_s\, ds = X_0(T_1 - T_0)$ because $X_s = X_0 = X_{0^+}$= execution time for any job that arrives between the two wake-up times T_0, T_1

Hence $E(X_0) = \lambda E^0\big(X_0(T_1 - T_0)\big)$

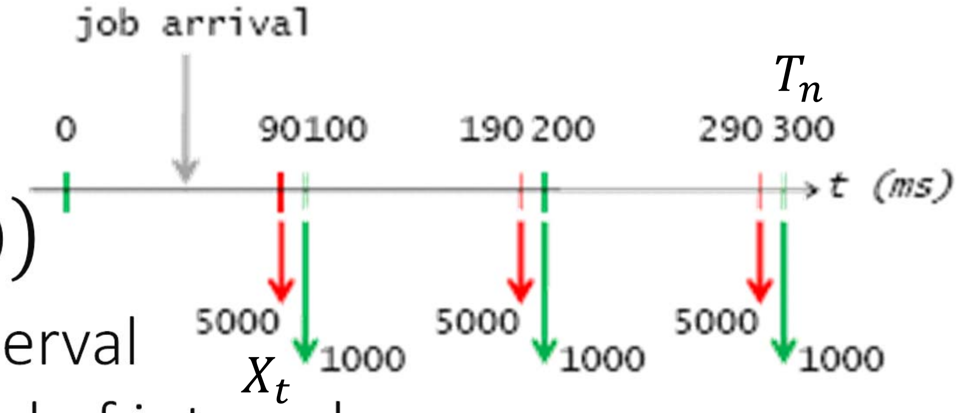

Hence $E(X_0) = \lambda E^0\big(X_0(T_1 - T_0)\big)$

Let C be covariance of wakeup interval duration and execution time at end of interval

$C = E^0(X_0(T_1 - T_0)) - E^0(X_0)E^0(T_1 - T_0)$

For example here: $E^0(X_0)$ = execution time for a job that arrives just after a wake-up time, averaged over wake-up times = W_s $=0.5 \times 5000 + 0.5 \times 1000$

Hence $E(X_0) = \lambda E^0\big(X_0(T_1 - T_0)\big) = \lambda\big(C + W_s E^0(T_1 - T_0)\big)$

Intensity formula: $\lambda^{-1} = E^0(T_1 - T_0)$

Hence $W_c = E(X_0) = \lambda C + W_s$

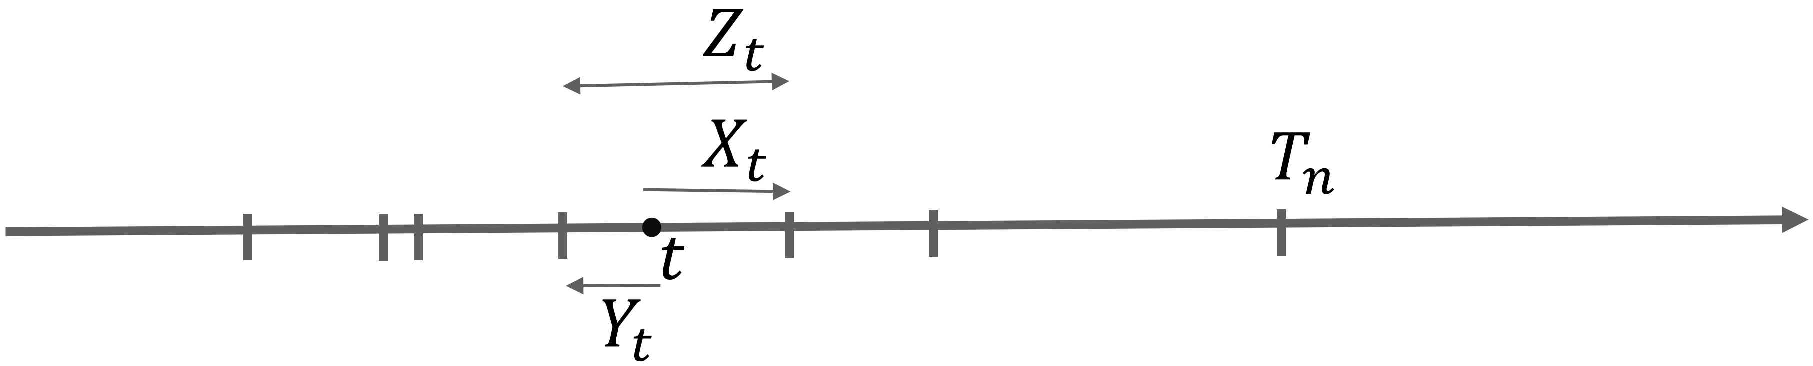

THEOREM 7.3. *Let $X(t) = T^+(t) - t$ (time until next point, also called residual time), $Y(t) = t - T^-(t)$ (time since last point), $Z(t) = T^+(t) - T^-(t)$ (duration of current interval). For any t, the distributions of $X(t)$ and $Y(t)$ are equal, with PDF:*

$$f_X(s) = f_Y(s) = \lambda \mathbb{P}^0(T_1 > s) = \lambda \int_s^{+\infty} f_T^0(u)du \tag{7.28}$$

where f_T^0 is the Palm PDF of $T_1 - T_0$ (PDF of inter-arrival times). The PDF of $Z(t)$ is

$$f_Z(s) = \lambda s f_T^0(s) \tag{7.29}$$

In particular, it follows that

$$\mathbb{E}(X(t)) = \mathbb{E}(Y(t)) = \frac{\lambda}{2}\mathbb{E}^0(T_1^2) \quad \text{in continuous time} \tag{7.30}$$

$$\mathbb{E}(X(t)) = \mathbb{E}(Y(t)) = \frac{\lambda}{2}\mathbb{E}^0(T_1(T_1+1)) \quad \text{in discrete time} \tag{7.31}$$

$$\mathbb{E}(Z(t)) = \lambda \mathbb{E}^0(T_1^2) \tag{7.32}$$

Feller's Paradox

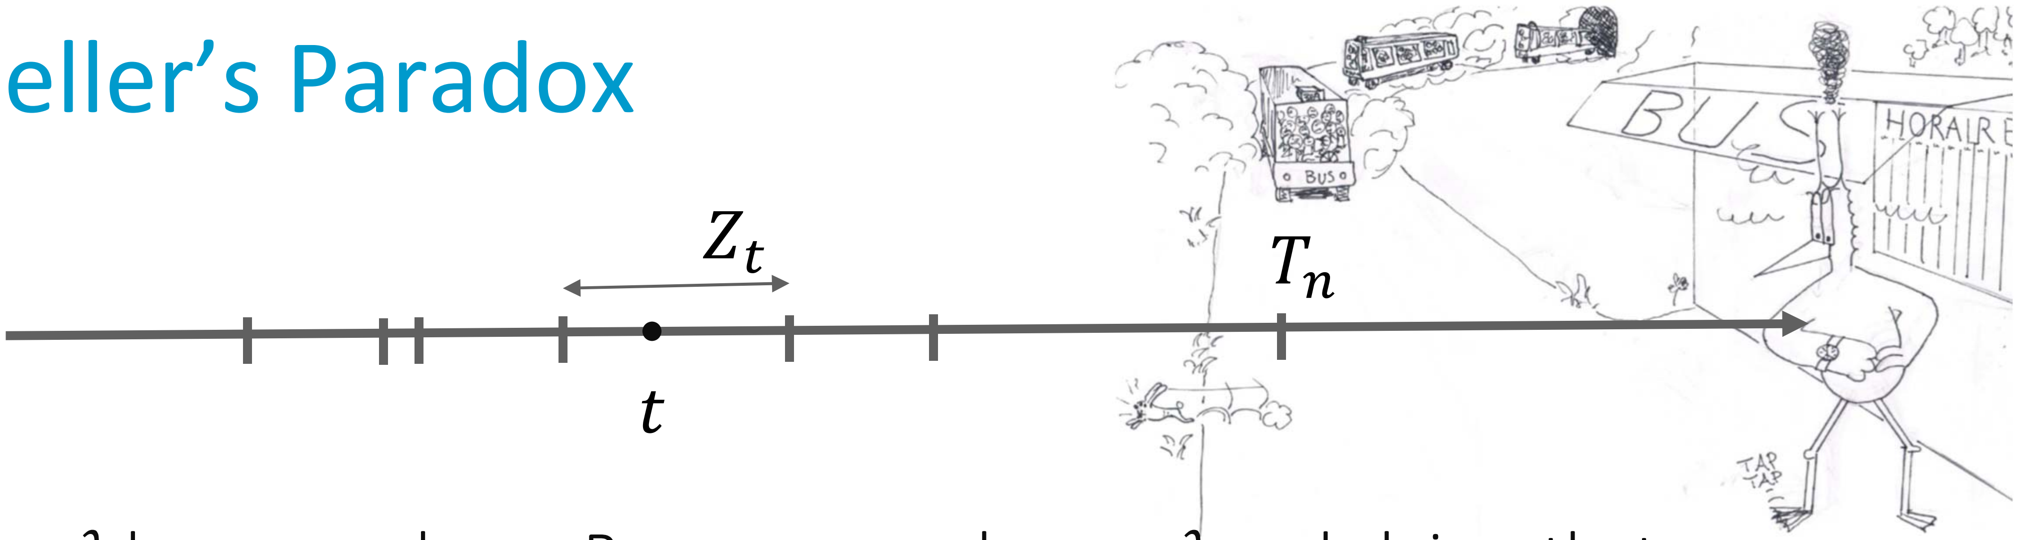

λ buses per hour. Bus company knows λ and claims that average interval between buses is $E^0(T_1 - T_0) = \frac{1}{\lambda}$

Joe arrives a bus stop and estimates Z_t (current time interval)

$$E(Z_t) = \lambda E^0\left((T_1 - T_0)^2\right) = \frac{1}{\lambda} + \lambda\, v$$

where $v = var^0(T_1 - T_0)$

Joe's estimate is always larger than inspector's (Feller's paradox)

We encountered Feller’s Paradox Already

PDFs of flow sizes

$f_F(s)$: PDF of flow size, seen by flows

$f_P(s)$: PDF of flow size, seen by packets

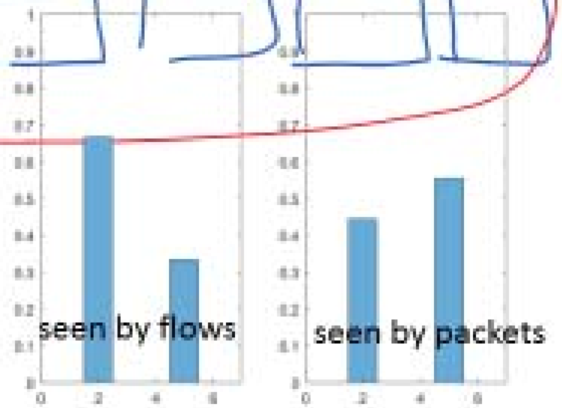

Using the same approach we obtain

$f_P(s) = \eta s \ f_F(s)$ where η is a normalization constant

21

For a Poisson process…

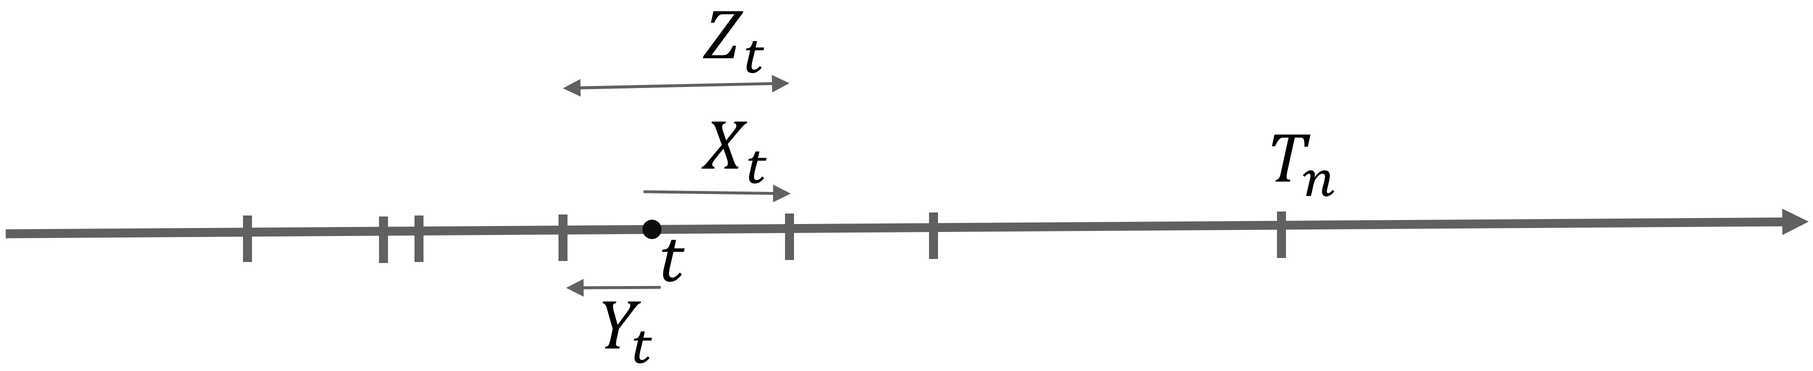

Under the time average viewpoint:

$f_X(t) = \lambda e^{-\lambda t}$ (exponential, as expected)

$f_Y(t) = \lambda e^{-\lambda t}$ (exponential as well)

$f_Z(t) = \lambda^2 t e^{-\lambda t}$ (Erlang-2, not exponential)

The duration of the time interval we are in is the sum of two independent exponential random variables

$$Z_t = X_t + Y_t$$

In average, we are in an interval of duration $\frac{2}{\lambda}$ (Feller's paradox)

A sensor senses events; the sensing interval is $\sim \boldsymbol{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^{2})$. An engineer comes and checks the current sensing interval. In average, she finds…

A. $\mu + \sigma^2$

B. $\mu(1 + \frac{\sigma^2}{\mu^2})$

C. $\mu(1 + \sigma^2)$

D. $\frac{1}{\mu}(1 + \frac{\sigma^2}{\mu})$

E. $\frac{1}{\mu}(1 + \frac{\sigma^2}{\mu^2})$

F. I don't know

A sensor senses events; the sensing interval is $\sim$ $\mathbf{expo}(\boldsymbol{\lambda})$, i.e. the event process is Poisson. An engineer comes and checks the current sensing interval. In average, she finds…

A. $\frac{1}{\lambda}$

B. $\frac{2}{\lambda}$

C. $\frac{1}{\lambda}\left(1+\frac{1}{\lambda}\right)$

D. $\frac{1}{\lambda}\left(1+\frac{1}{\lambda^2}\right)$

E. I don't know

2. Application to Simulation

Modulator-based simulation

Modulator: a sequence of states Z_n (e.g. channel state) and durations S_n stationary w.r. n

Modulated process: $Z(t), t \in [0, +\infty)$ such that $Z(t) = Z_n$ whenever $\sum_{i=1}^{n-1} S_i \leq t < \sum_{i=1}^{n-1} S_i$

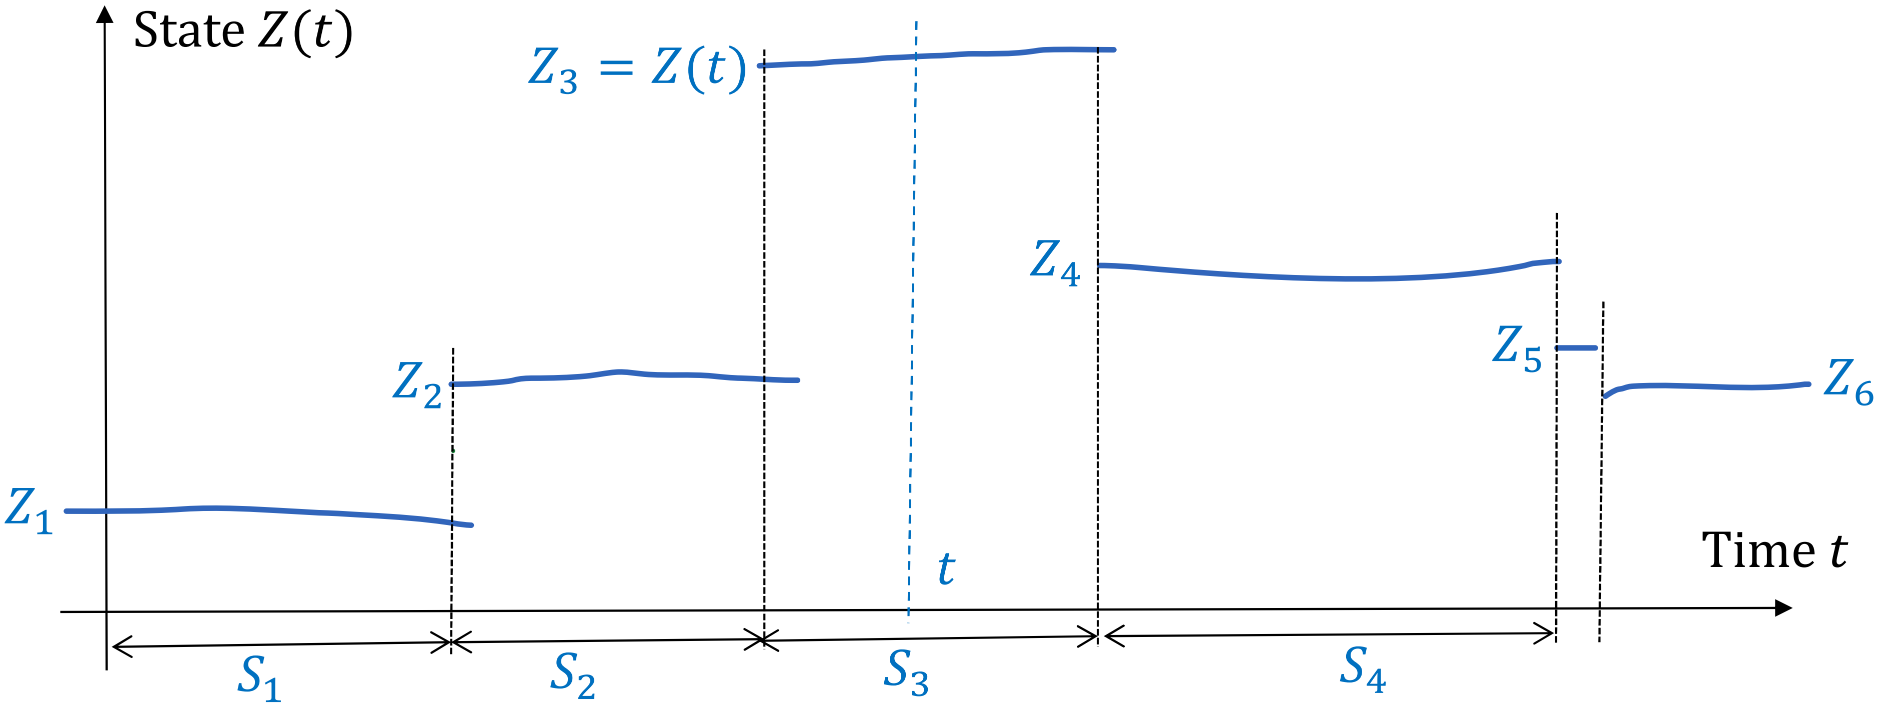

At $\boldsymbol{T_n}$, channel state is drawn at random and new state is $\boldsymbol{i}$ with proba $\boldsymbol{\pi_i^0}$. When channel state $\boldsymbol{Z(t)}$ is $= \boldsymbol{i}$, loss proba is $\boldsymbol{p_i}$ and residence time in that state is (non random) $\boldsymbol{r_i}$. What is the intensity of the point process $\boldsymbol{T_n}$?

A. $\lambda = \sum_i \pi_i^0 r_i$

B. $\lambda = \sum_i \frac{\pi_i^0}{r_i}$

C. $\lambda = \frac{1}{\sum_i \pi_i^0 r_i}$

D. None of the above

E. I don’t know

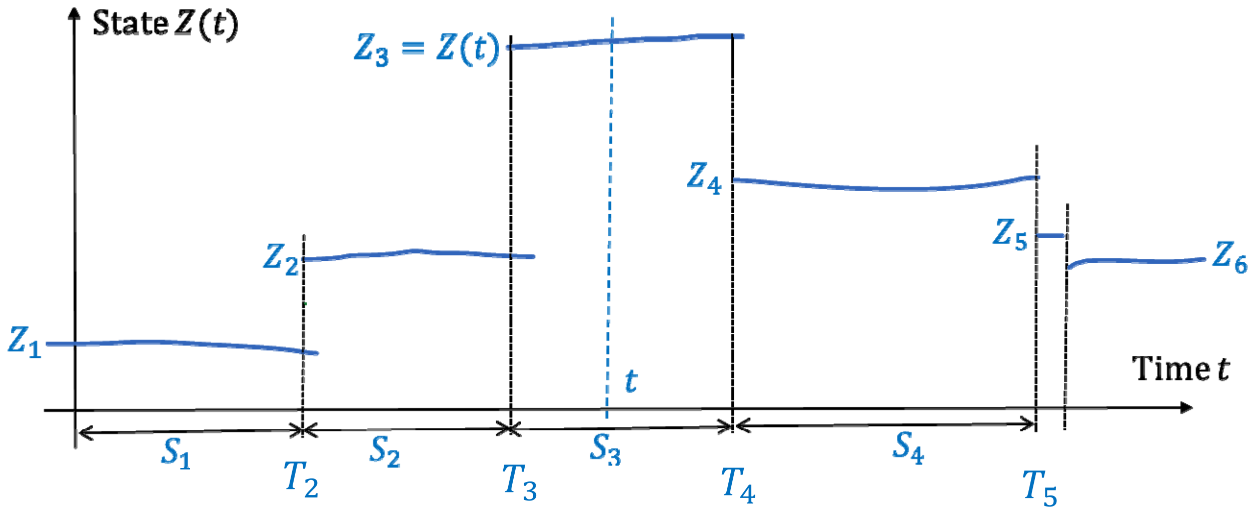

At $\boldsymbol{T_n}$, channel state is drawn at random and new state is $\boldsymbol{i}$ with proba $\boldsymbol{\pi^0(i)}$. When channel state $\boldsymbol{Z(t)}$ is $= \boldsymbol{i}$, loss proba is $\boldsymbol{p_i}$ and residence time in that state is $\boldsymbol{r_i}$. What is the loss probability $\boldsymbol{p}$ for a probe packet sent at an arbitrary point in time $\boldsymbol{t}$?

A. $\dfrac{\sum_i \pi_i^0 p_i r_i}{\sum_i \pi_i^0 r_i}$

B. $\dfrac{\sum_i \pi_i^0 r_i}{\sum_i \pi_i^0 p_i r_i}$

C. $\dfrac{\sum_i \pi_i^0 p_i r_i}{\sum_i \pi_i^0 p_i}$

D. None of the above

E. I don't know

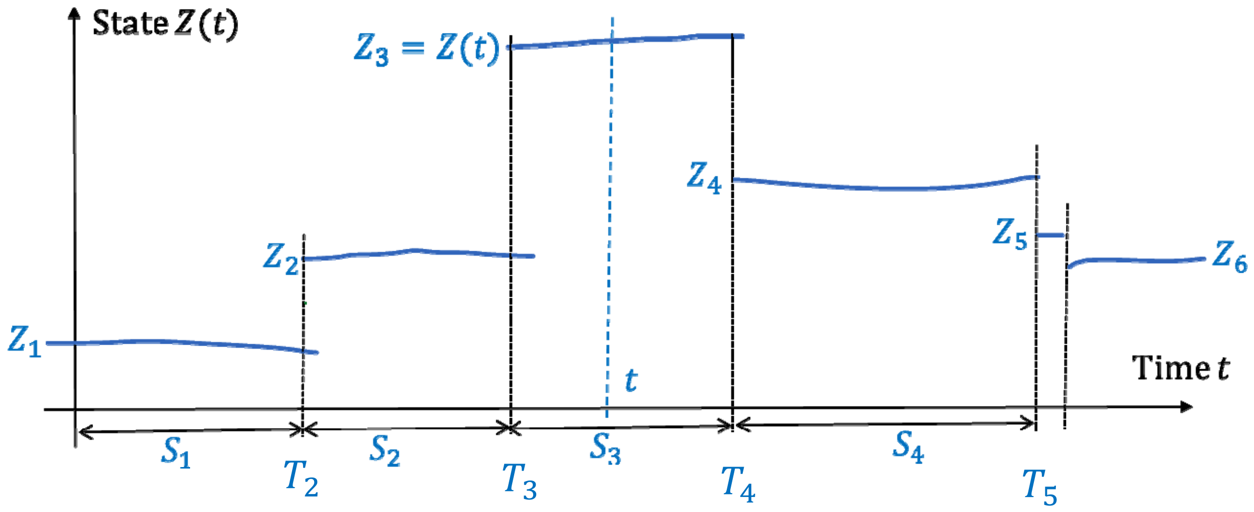

Is the previous simulation stationary ?

Seems like a superfluous question, however there is a difference in viewpoint between the epoch n and time t

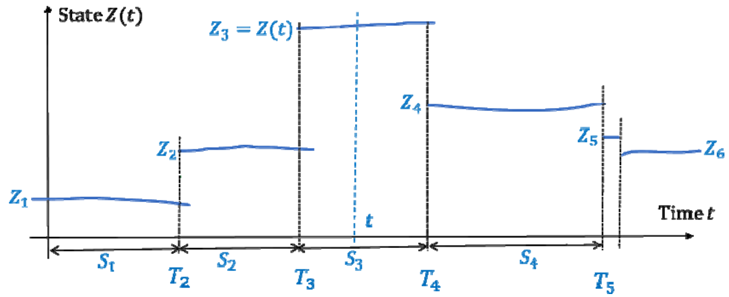

If there is a stationary regime, by the intensity formula:

$$\lambda = \frac{1}{E(S_n)}$$

So the expectation of S_n must be finite. This is also sufficient:

THEOREM 7.9. *Assume that the sequence S_n satisfies **H1** and has finite expectation. There exists a stationary process $Z(t)$ and a stationary point process T_n such that*

1. $T_{n+1} - T_n = S_n$
2. $Z_n = Z(T_n)$

RWP

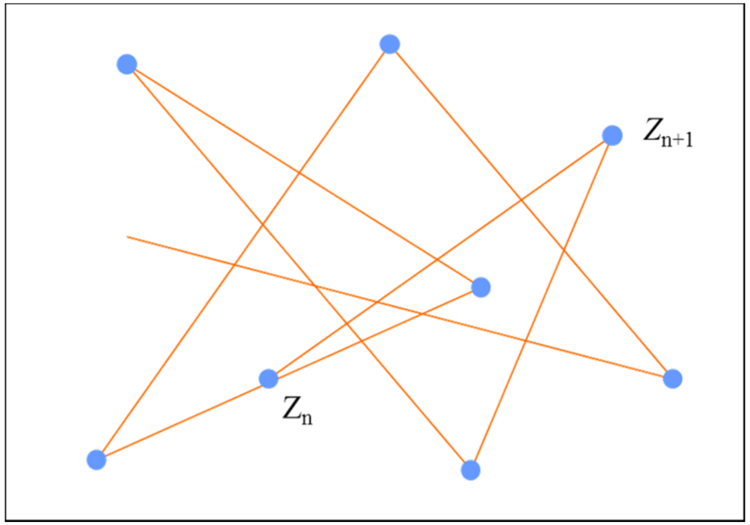

Modulator: nth trip $Z_n = (M_n, M_{n+1}, V_n)$

Modulated Process: $Z(t) =$ the trip that we are in at time t

Duration of nth trip $S_n = \frac{d(M_n, M_{n+1})}{V_n}$

Assume waypoints and speed are chosen independently (as in lab)

$$E(S_n) = E\big(d(M_n, M_{n+1})\big)\, E\left(\frac{1}{V_n}\right)$$

Assume speed is chosen uniformly between $v_{\min} \geq 0$ and $v_{\max} > 0$

$$E\left(\frac{1}{V_n}\right) = \frac{1}{v_{\max} - v_{\min}} \int_{v_{\min}}^{v_{\max}} \frac{1}{v} dv = \frac{1}{v_{\max} - v_{\min}} (\log v_{\max} - \log v_{\min})$$

There is a stationary regime $\Leftrightarrow v_{\min} > 0$

Time Average Speed, Averaged over *n* independent mobiles

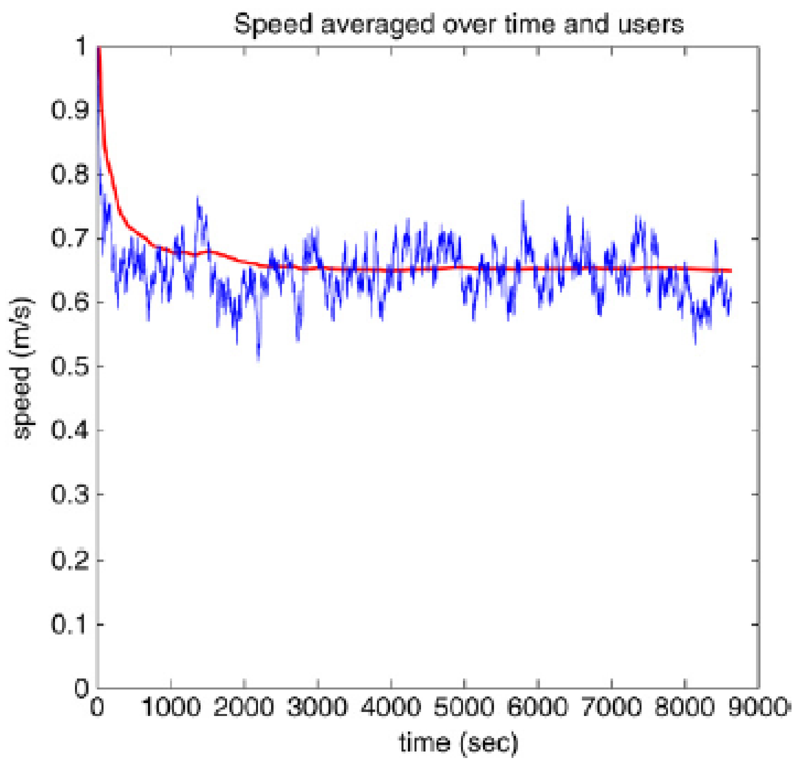

(a) $v_{\min} = 0.1$ m/s.

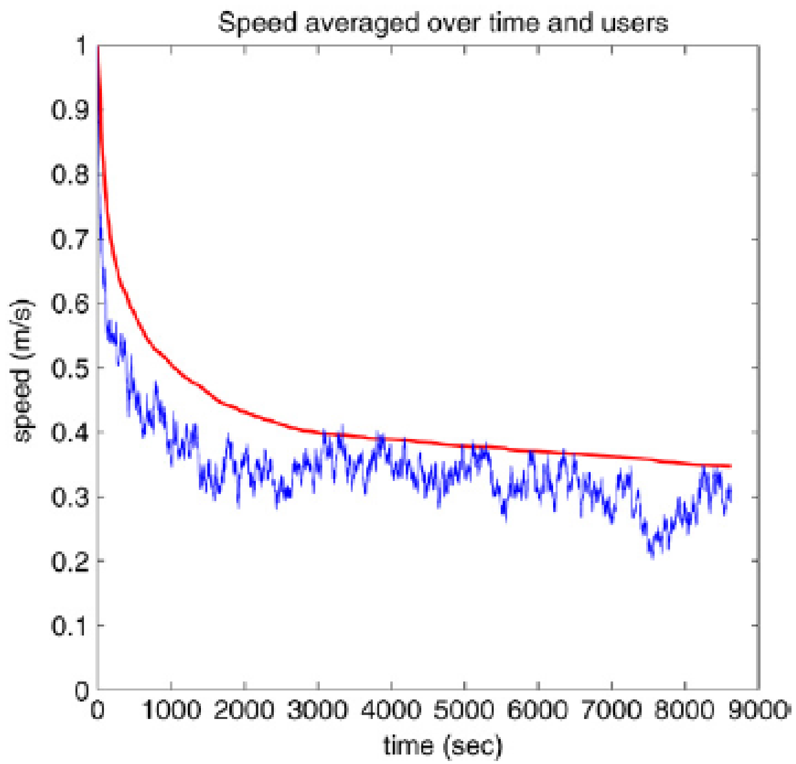

(b) $v_{\min} = 0$ m/s.

Blue line is one sample; Red line is estimate of $E(V(t))$

A Random waypoint model that has no stationary regime

When $v_{\min} > 0$ there is no stationary regime.

But this model was often used in practice "considered harmful" [YLN03]

Simulation becomes old and "freezes" – average speed $\rightarrow 0$

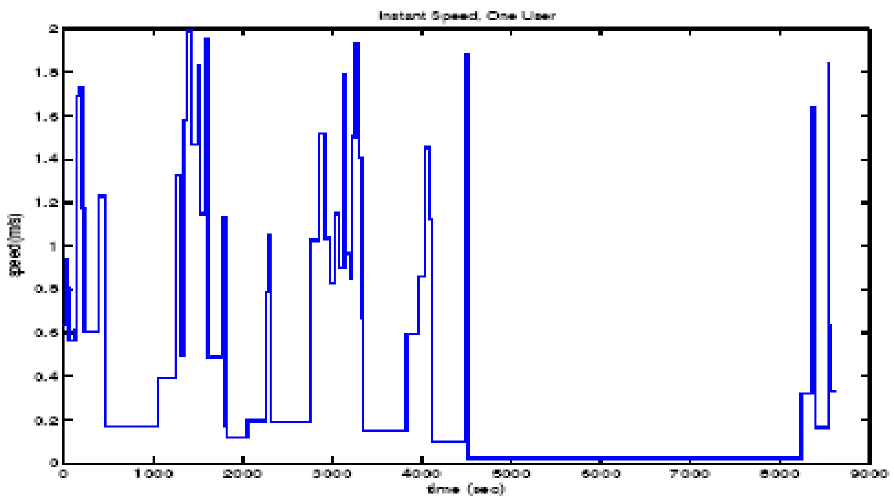

One User

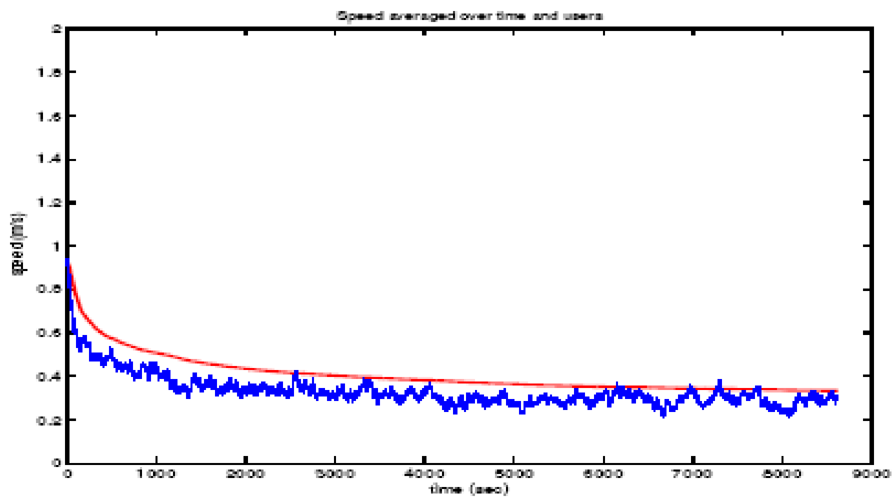

Instant Speed + Empirical speed, both averaged over users

Stationary Distribution of Speed
(For model with stationary regime)

Random Waypoint on Rectangle, without Pause:

- Speed observed at waypoints (Event average)

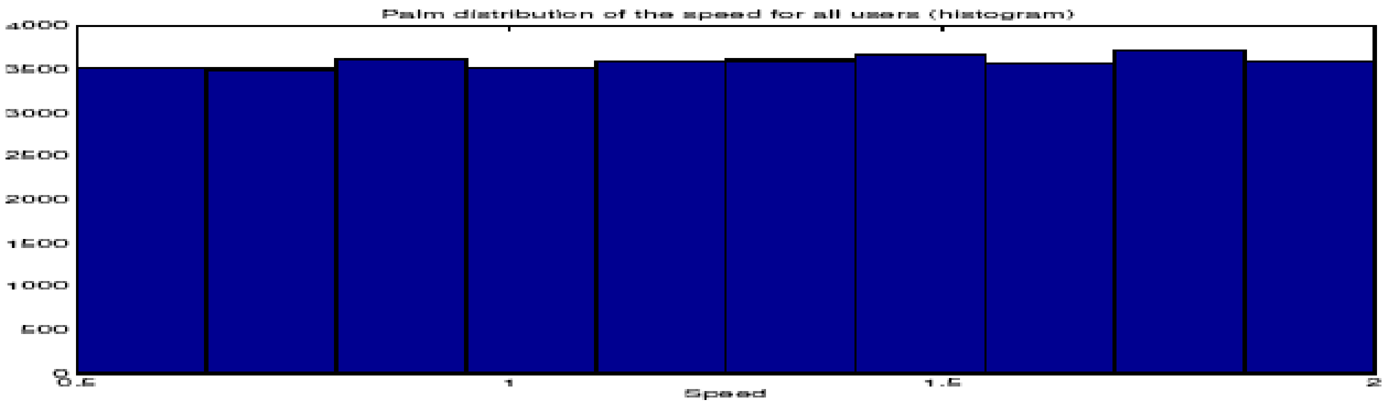

- Speed observed at an arbitrary time (Time average)

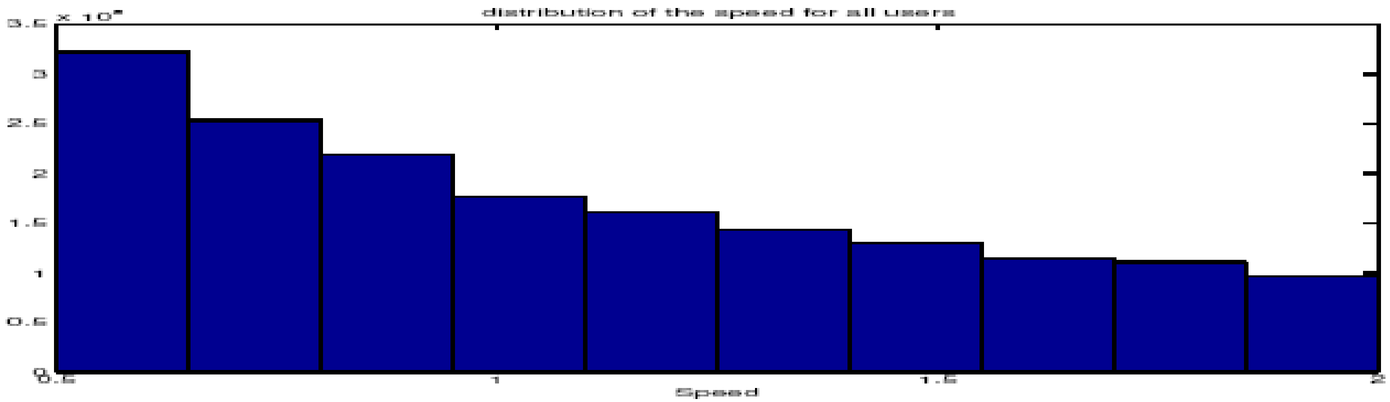

Closed Form

Assume a stationary regime exists and simulation is run long enough

Apply inversion formula and obtain distribution of instantaneous speed V(t)

$$
\begin{aligned}
\mathbb{E}\left(\phi(V(t))\right) = \lambda \mathbb{E}^0 &\left(\int_0^{T_1} \phi(V(t))\, dt\right) \\
&= \lambda \mathbb{E}^0 \left(\phi(V_0) T_1\right) \\
&= \lambda \mathbb{E}^0 \left(\phi(V_0) \frac{\|M_1 - M_0\|}{V_0}\right) \\
&= \lambda \mathbb{E}^0 \left(\|M_1 - M_0\|\right) \mathbb{E}^0 \left(\frac{\phi(V_0)}{V_0}\right) \\
&= C \int_{v_{\min}}^{v_{\max}} \frac{\phi(v)}{v} f^0_{V_0}(v) dv
\end{aligned}
$$

$$
f_{V(t)}(v) dv = \frac{C}{v} f^0_{V_0}(v) dv
$$

A (true) example: Compare impact of mobility on a protocol: Experimenter places nodes uniformly for static case, according to random waypoint for mobile case.
Finds that static is better
Find the bug !

A. Spatial distribution of nodes is different for mobile case than for static case
B. Spatial distribution of nodes is same for mobile and static cases but speed is more often small than large
C. There is no bug, mobility increases capacity
D. I don’t know

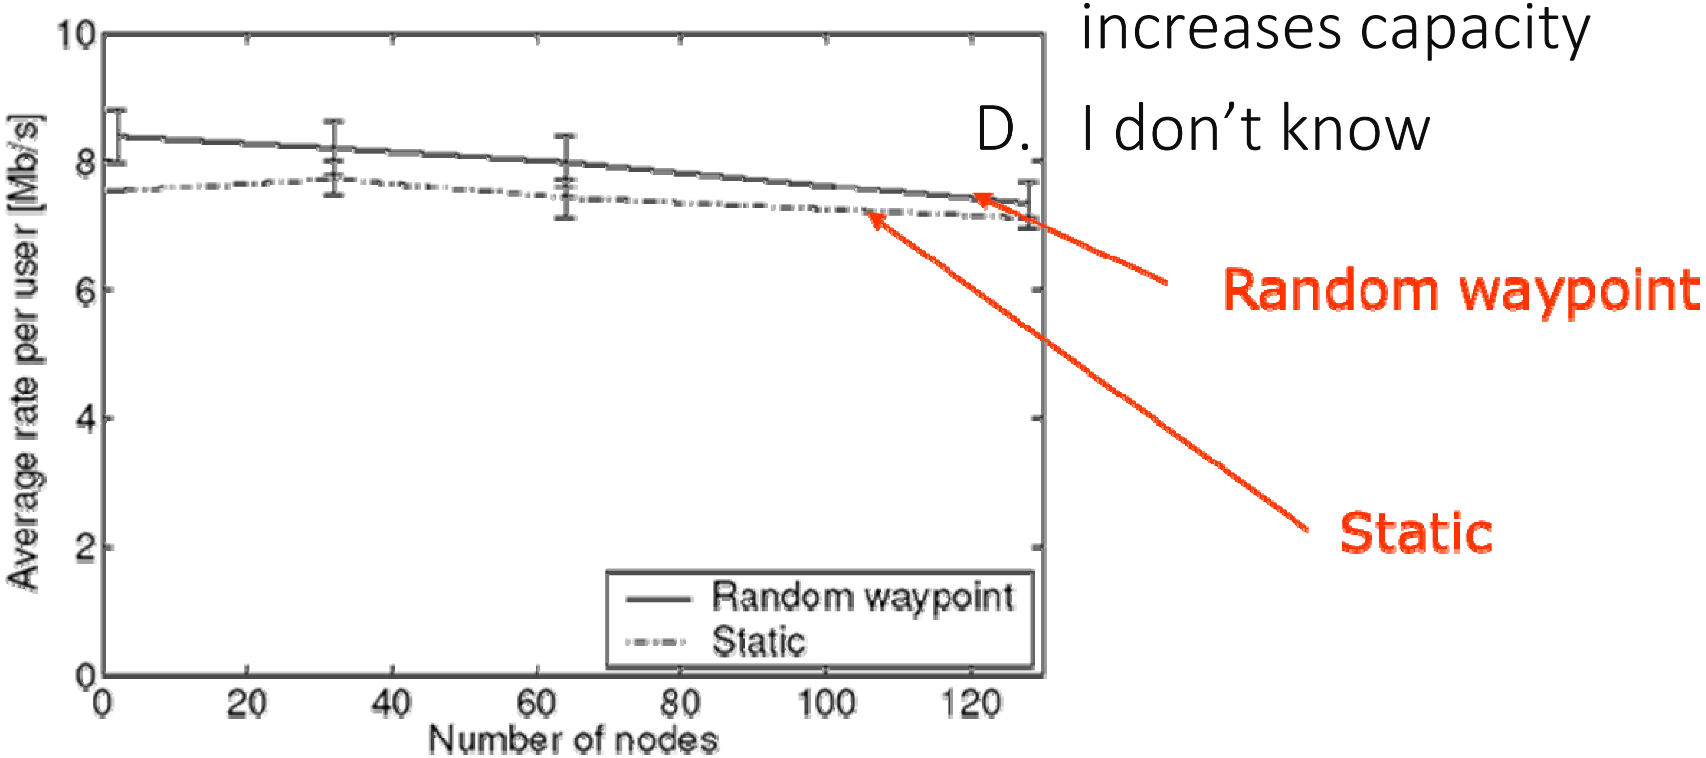

Does this model have a stationary regime ?

A mobile moves as follows
- pick a random direction uniformly in $[0, 2\pi]$
- pick a random trip duration T ~Pareto(p)
- go in this direction for duration T at constant speed ; if needed reflect at the boundary.

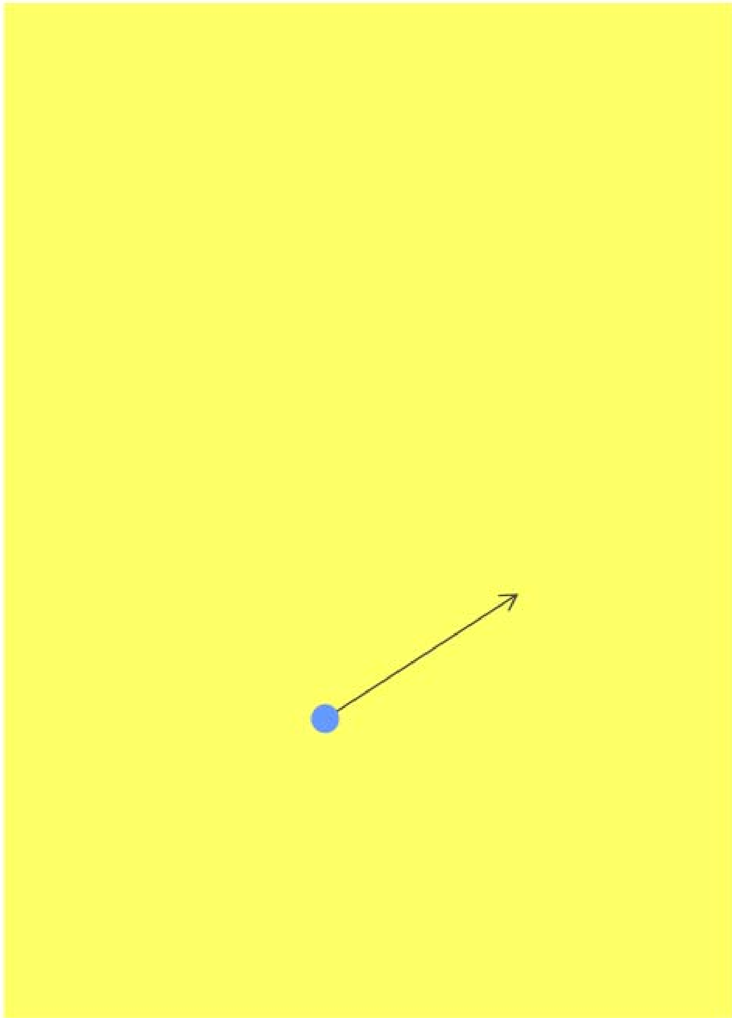

A. Yes if p>1
B. Yes if $p > 2$
C. Yes for all p
D. No
E. I don’t know

3. Perfect Simulation

Def: a simulation that starts with a sample from the stationary distribution

An alternative to removing transients

Usually difficult except for modulated models, e.g. RWP

Perfect simulation of RWP:

Sample speed V from its time stationary distribution

Sample $Prev$ and $Next$ waypoints from their joint stationary distribution

Sample M uniformly on segment $[Prev, Next]$

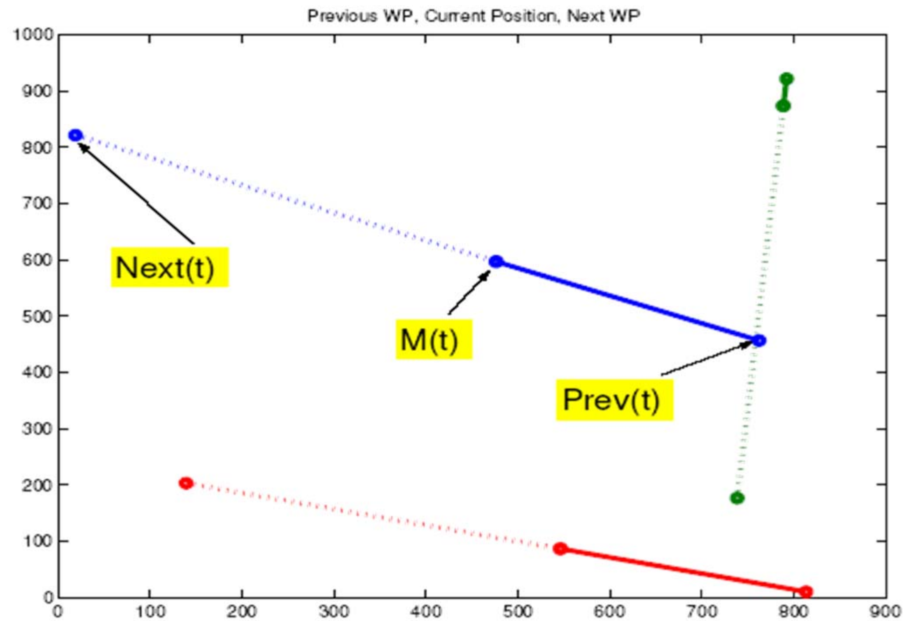

Stationary Distrib of Prev and Next

- Let $M(t)$: position at time t
- Let $Prev(t), Next(t)$: previous and next waypoints

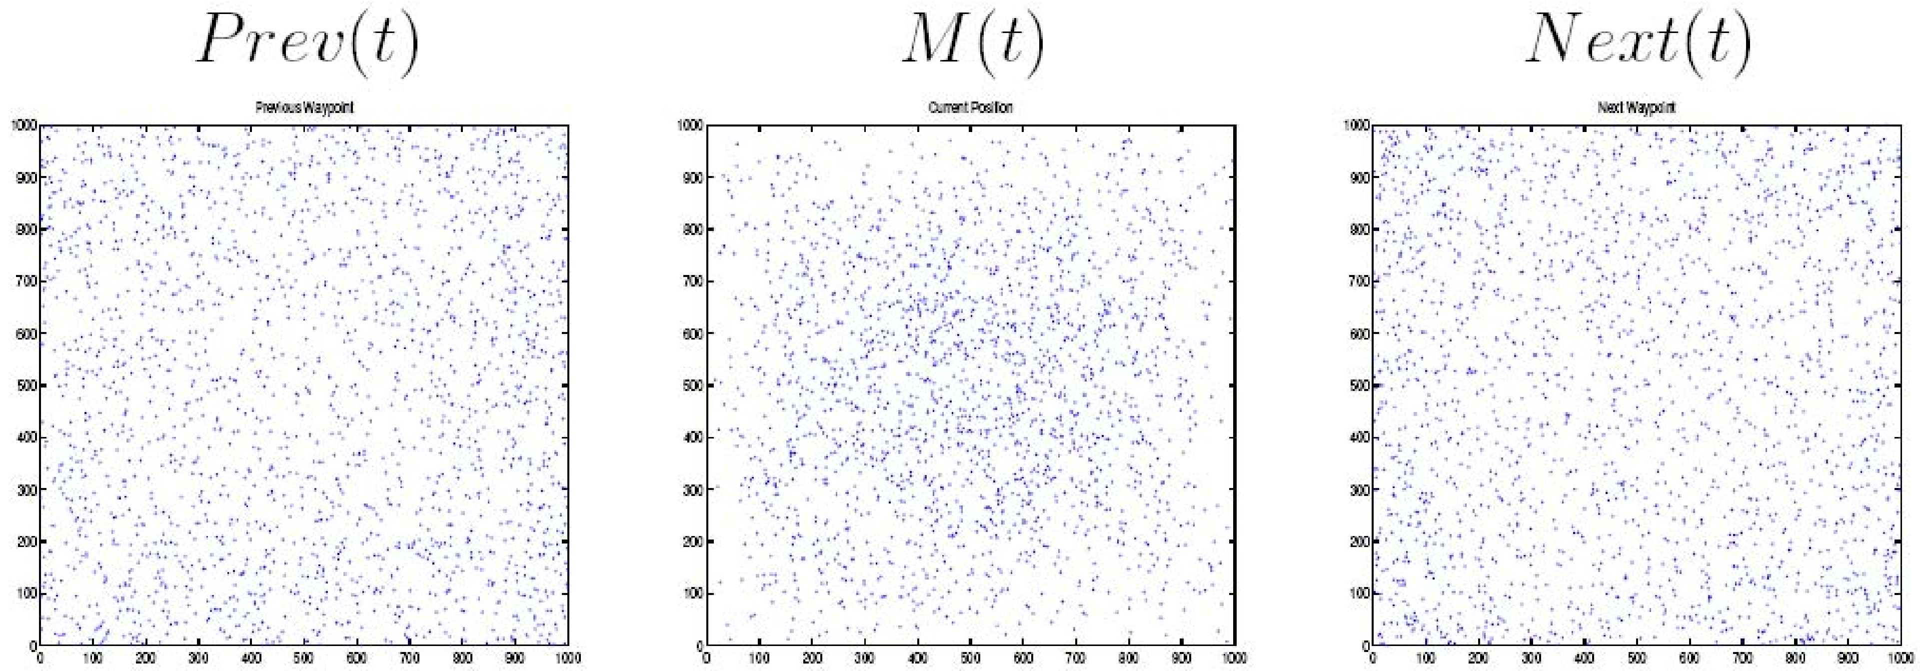

Is $\boldsymbol{M(t)}$ uniformly distributed ?

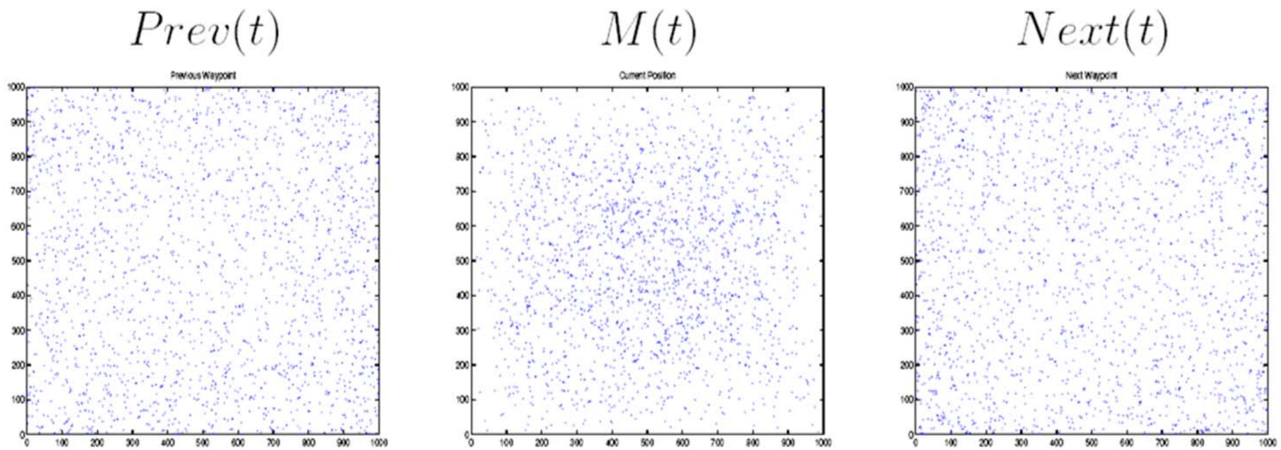

A. Yes
B. No
C. It depends on the distribution of speed
D. I don’t know

Is $\boldsymbol{Next(t)}$ uniformly distributed ?

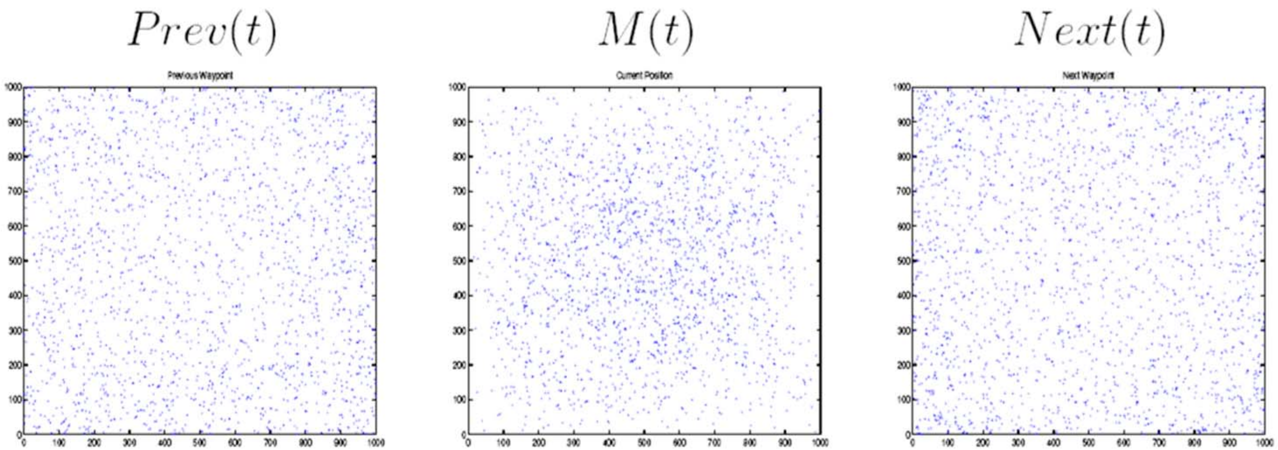

A. Yes
B. No
C. It depends on the distribution of speed
D. I don’t know

Stationary Distribution of Location Is also Obtained By Inversion Formula

- Joint distribution of $(Prev(t), M(t), Next(t))$ has a simple closed form [NavidiCamp04]:

1. $((Prev(t), Next(t))$ has density over area A

$$f_{Prev(t),Next(t)}(P, N) = K \|P - N\|$$

2. Distribution of $M(t)$ given $Prev(t) = P, Next(t) = N$ is uniform on segment $[P, N]$

$K^{-1} = \text{vol}(A)^2 \bar{\Delta}(A)$, with $\bar{\Delta}(A)$ = average distance between two points in A. For $A = [0; a] \times [0; a]$, $\bar{\Delta}(A) = 0.5214a$ [Gosh1951].

Proof. For any bounded, non-negative function ϕ:

$$\mathbb{E}(\phi(\text{Prev}(t), M(t), \text{Next}(t))) = \lambda \mathbb{E}^0 \left(\int_0^{T_1} \phi \left(M_0, M_0 + \frac{t}{T_1}(M_1 - M_0), M_1 \right) \mathrm{d}t \right).$$

By a simple change of variable in the integral, we obtain

$$\lambda \mathbb{E}^0 \left(T_1 \int_0^1 \phi(M_0, M_0 + u(M_1 - M_0), M_1) \mathrm{d}u \right).$$

Now given that there is an arrival at time 0, $T_1 = \frac{\|M_1 - M_0\|}{V_0}$ and the speed V_0 is independent of the waypoints M_0 an M_1 thus

$$= \lambda \mathbb{E}^0 \left(\frac{1}{V_0} \right) \mathbb{E}^0 \left(\|M_1 - M_0\| \int_0^1 \phi(M_0, M_0 + u(M_1 - M_0), M_1) \mathrm{d}u \right)$$

$$= K_2 \int_A \int_A \int_0^1 \phi(M_0, (1-u)M_0 + uM_1, M_1) \|M_1 - M_0\| \, \mathrm{d}u \, \mathrm{d}M_0 \, \mathrm{d}M_1$$

which shows the statement. □

No Speed Decay

Standard Simulation

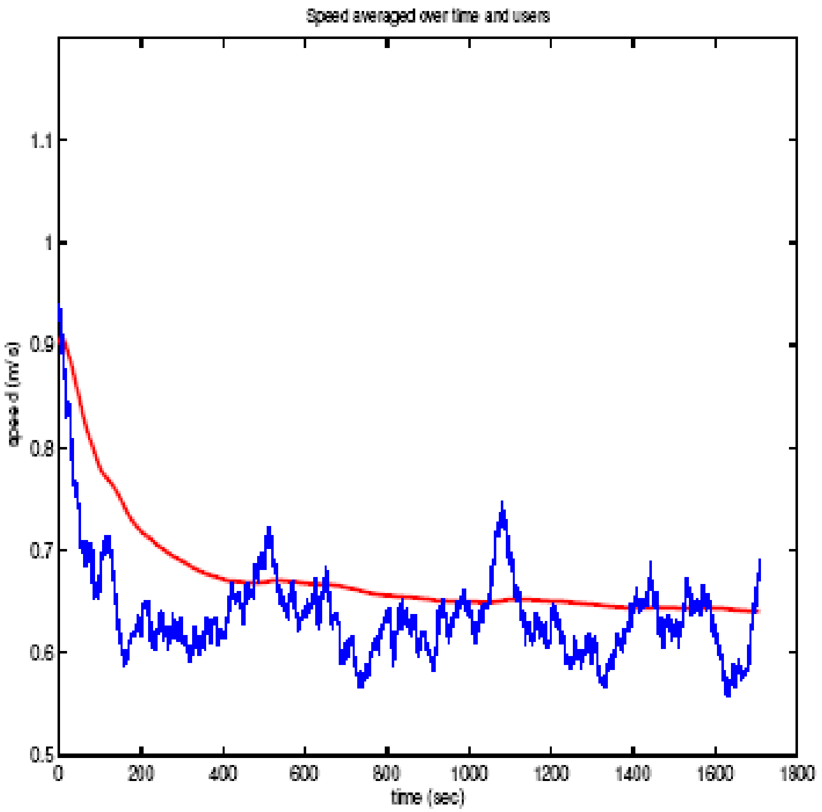

Perfect Simulation

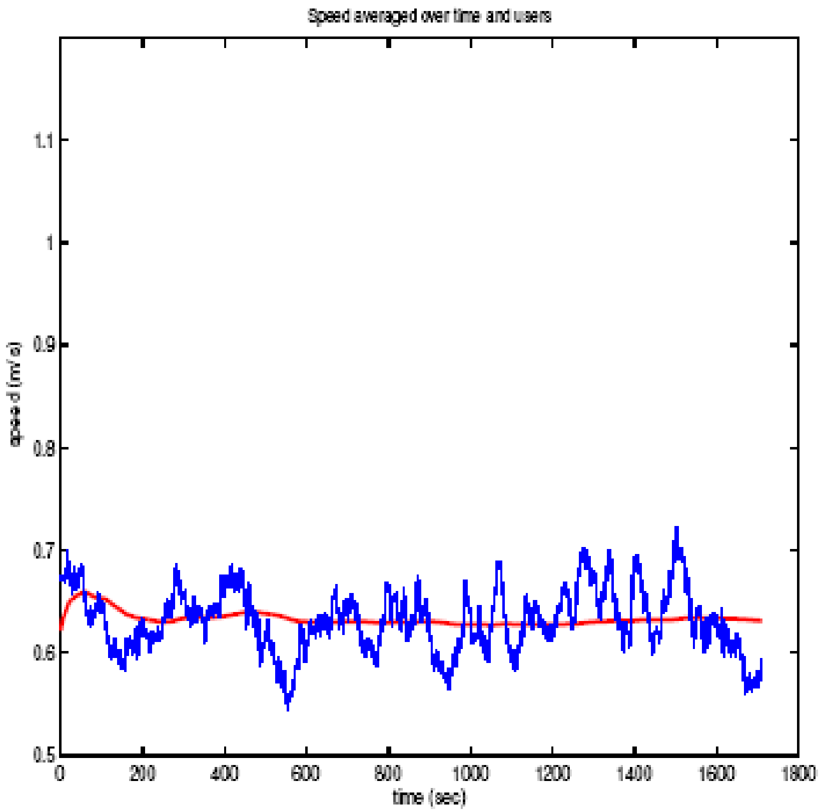

Perfect Simulation of RWP: How do you sample the speed ?

A. By rejection sampling

B. By CDF inversion

C. By an ad-hoc method

D. I don't know

$$\boxed{f_{V(t)}(v)dv = \frac{C}{v} f^0_{V_0}(v)dv}$$

Perfect Simulation of RWP: How do you sample the current segment (P, N) ?

1. $((Prev(t), Next(t))$ has density over area A

$$f_{Prev(t),Next(t)}(P, N) = K \|P - N\|$$

A. By rejection sampling

B. By CDF inversion

C. By an ad-hoc method

D. I don't know

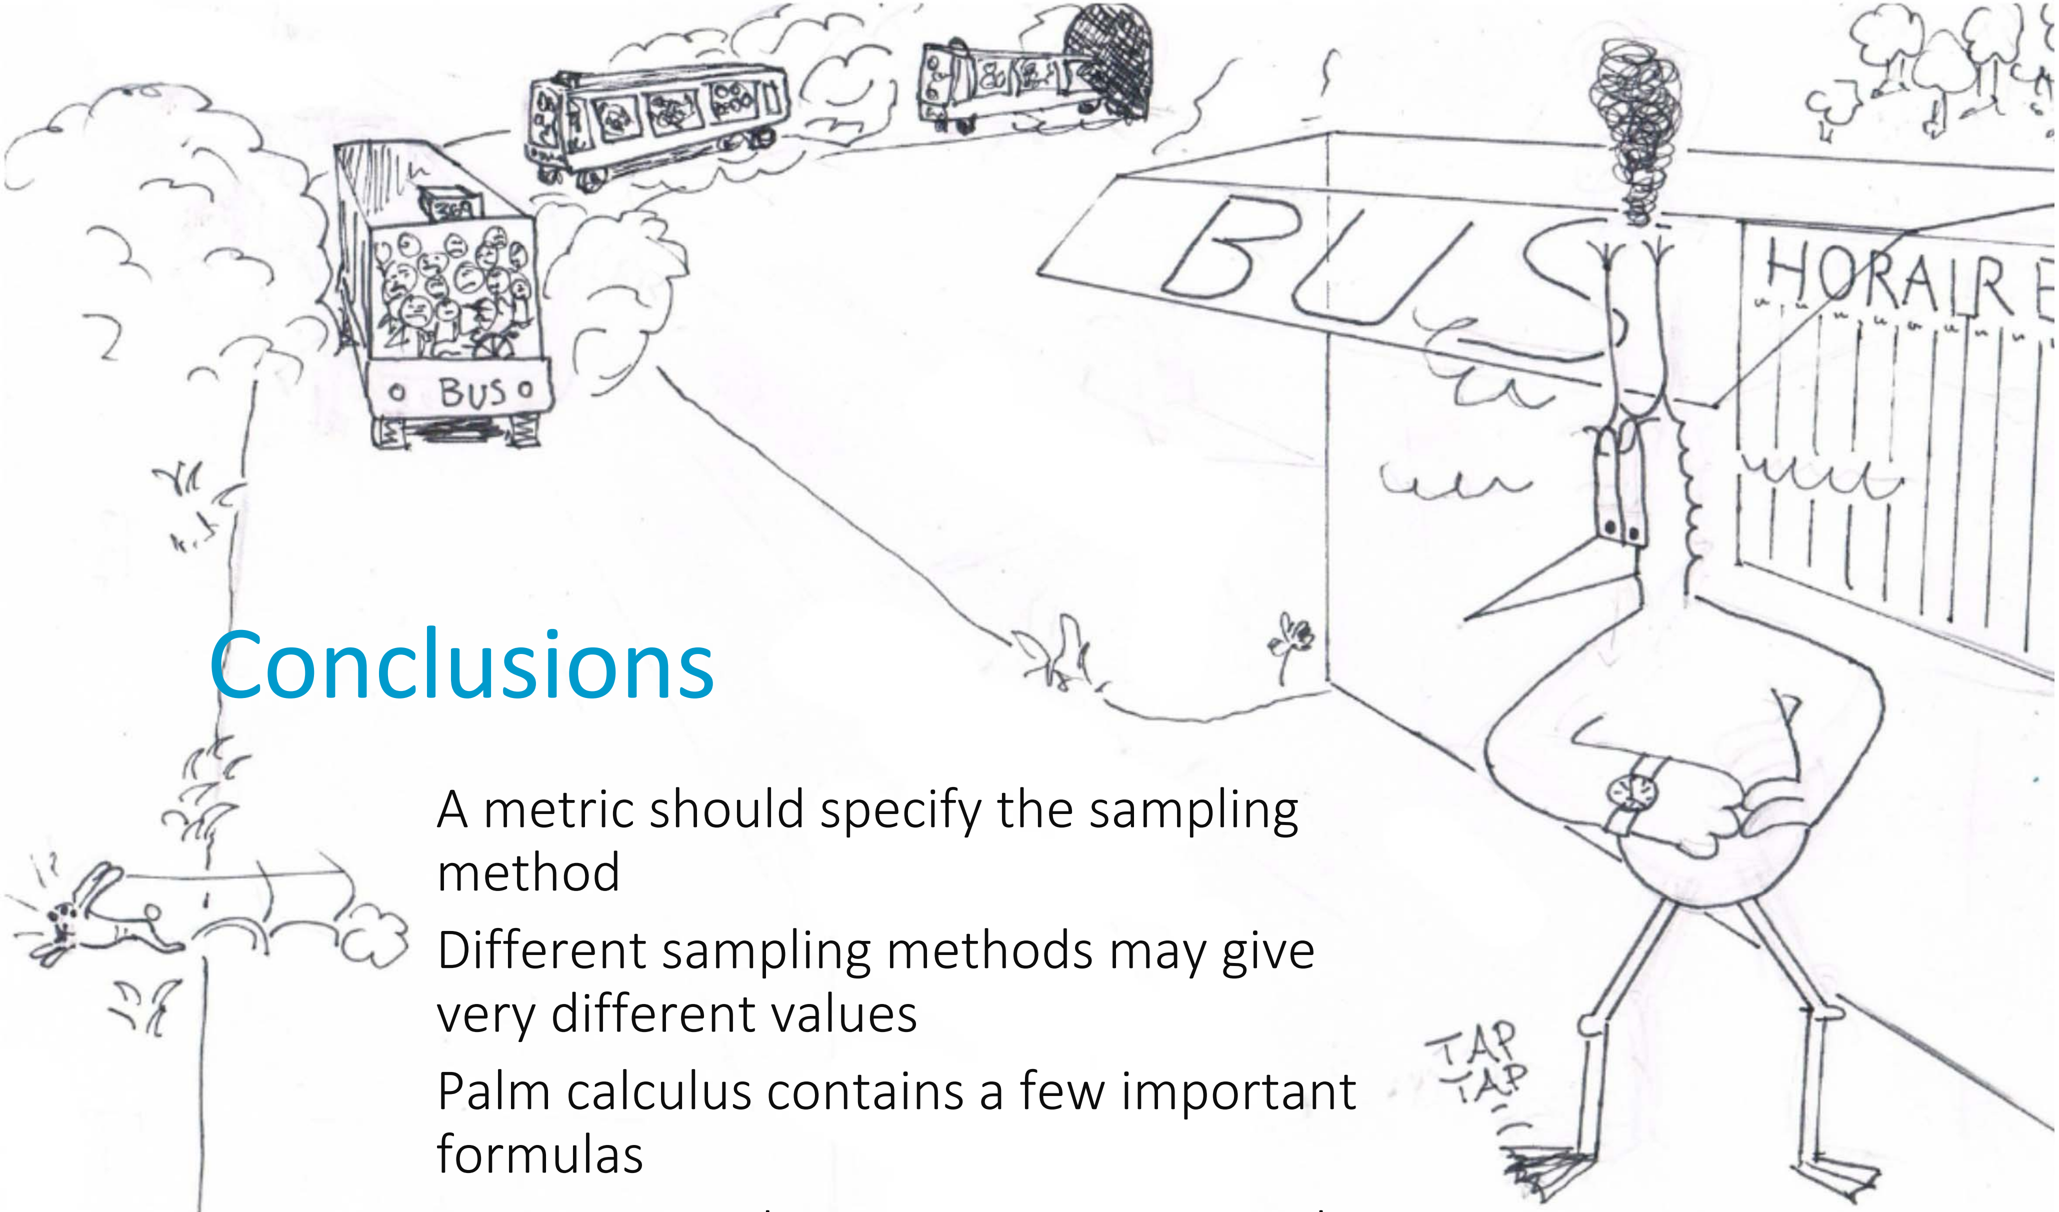

Conclusions

- A metric should specify the sampling method
- Different sampling methods may give very different values
- Palm calculus contains a few important formulas
- Freezing simulations are a pattern to be aware of