Network Engineering

Lecture 3. Continuous Time Markov Chains

Boris Bellalta

boris.bellalta@upf.edu

WN UPF

Wireless Networking

https://www.upf.edu/web/wnrg

Continuous Time Markov chains

When the system we have to model can change at any arbitrary time, and the time that the system remains in a certain state is important, we will use a CTMC to model it. Examples of stochastic processes that can be modelled with CTMCs are:

- The number of packets waiting in a queue, which depends on the time packets arrive and depart from it.
- The number of persons with active phone conversations in a cell.

Similary to DTMCs, CTMCs will be characterized by a set of states, $\mathcal{X}$, and a matrix containing the transition rates from one state to the other, $\boldsymbol{Q}$, known as infinitesimal generator or rate transition matrix.

Stationary Probability Distribution

- Two different interpretations
  - The probability that the system is in state ‘i’ when we check the system state at an arbitrary instant of time.
    - In the example: $\pi i = 2/5$ (2 of 5 events find the system in state i)
  - The fraction of time the system is in state ‘i’.
    - In the example: $\pi i = (Ti1+Ti2+Ti3)/T = 2/5$

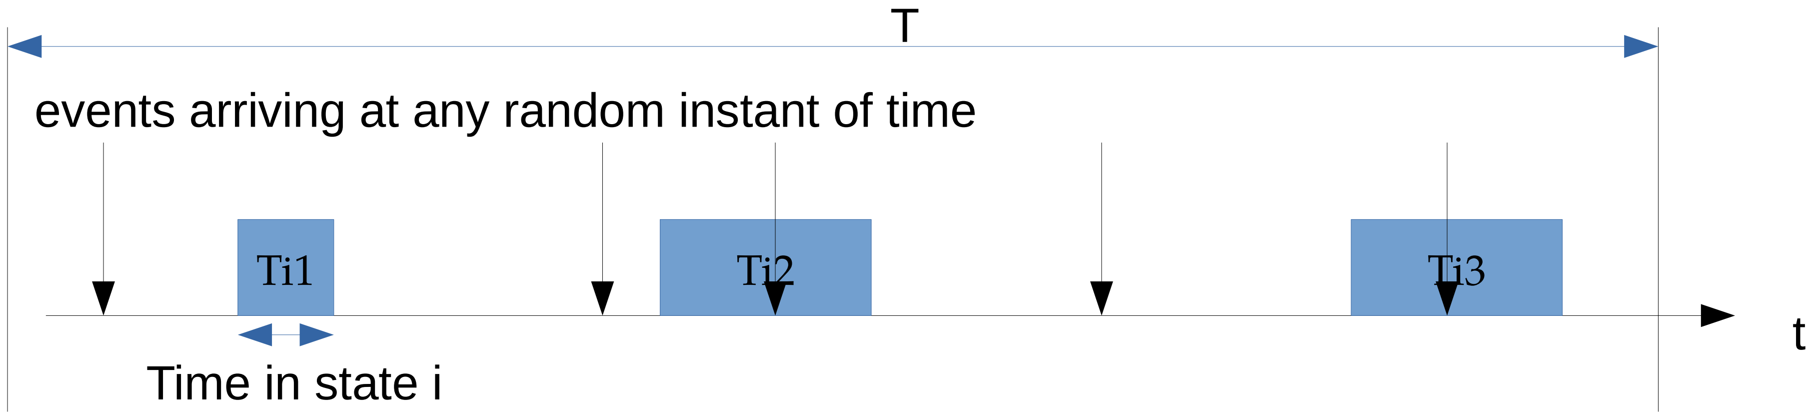

Continuous Time Markov chains

To move from a DTMC to a CTMC, we assume that the time is divided in very small time intervals of size δ, in a way that changes seem as continuous (see Figure 3.1). For instance, when we watch the Television, it seems that the images are continuous, but this is just an illusion: the images are static and change every few msecs.

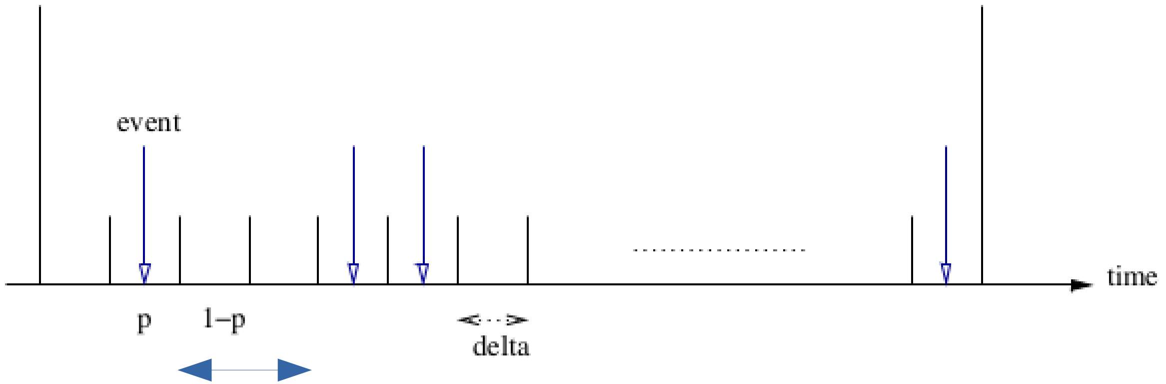

Time between two events → geometrically distributed

Continuous Time Markov chains

One of the mandatory requirements for these small intervals is that each one only can contain a single event. The probability that one period of time contains an event is $p = q\delta$, where q is the average rate (i.e. frequency) in which events happen (events / second). A second requirement is that all the periods of duration δ must have the same probability to contain or not an event, which means that the probability p must remain always constant. For example, if we have that $q = 10$ events / second, and define $\delta = 0.05$ seconds, the probability that a given period of duration δ contains an event is $p = q\delta = 0.5$.

Making it continuous

- δ->0: Binomial distribution → Poisson distribution

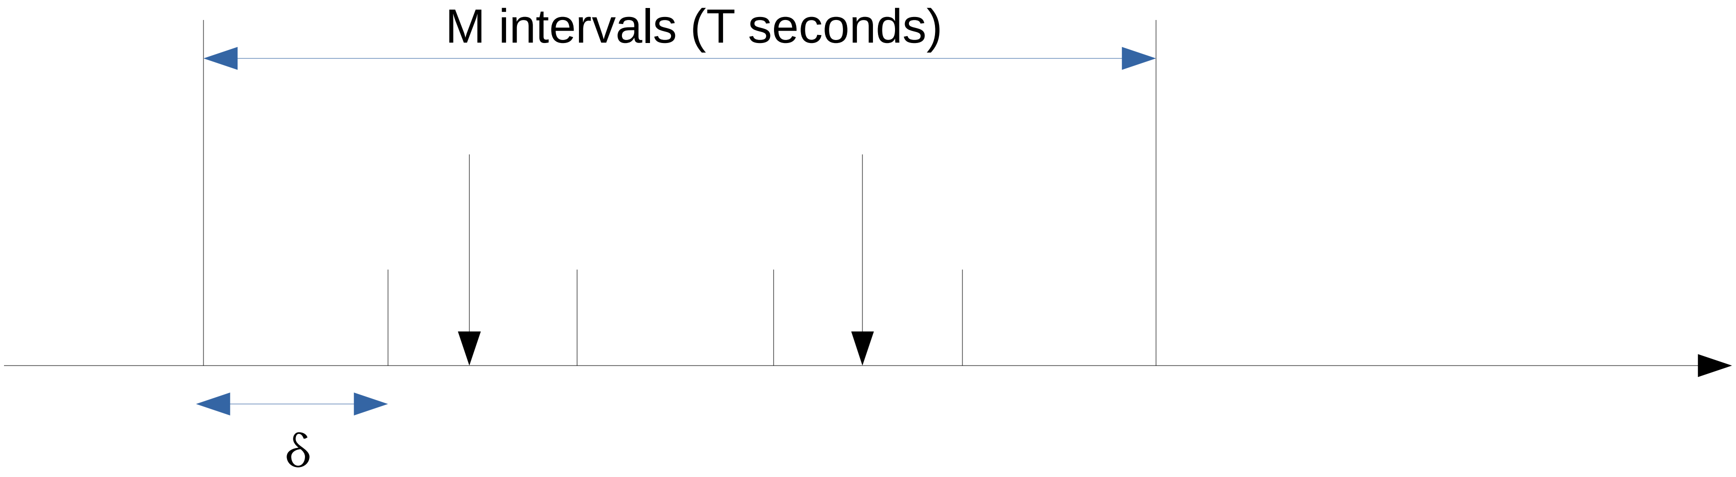

$$P(m,T) = \lim_{\delta \to 0} P(m,M) = \lim_{\delta \to 0} \binom{M}{m} p^m (1-p)^{M-m} = \frac{(qT)^m}{m!} e^{-qT}$$

Time between events must be exponential

- It results of moving from discrete to continuous
  - Time between two events (a random variable) is exponentially distributed
  - The number of events (a random variable) in a given period of time follows a Poisson distribution

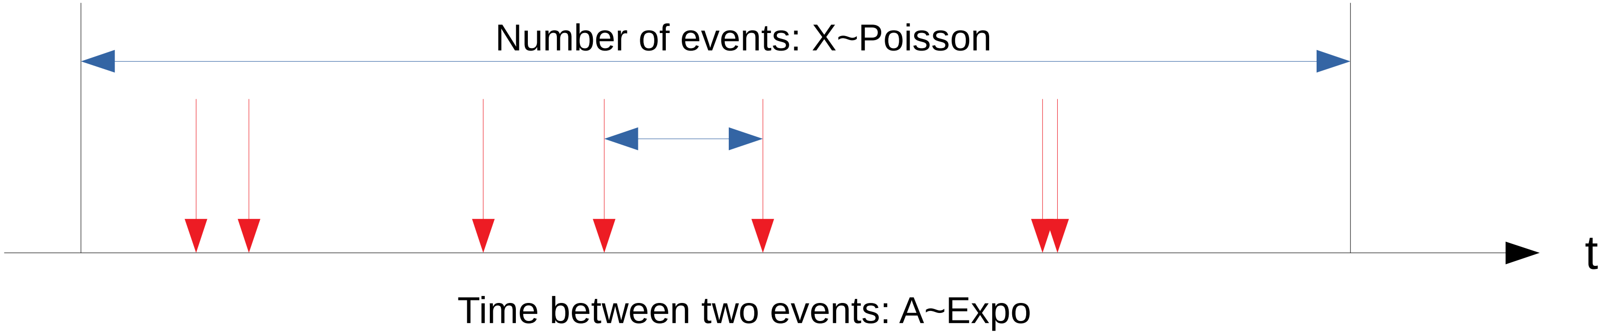

Poisson and Exponential distribution

The probability of 'm' events in 'T' follows a Poisson distribution:

$$P(m,T) = \frac{(qT)^m}{m!} e^{-qT}$$

Then, the probability of 0 events in a time τ is given by

$$P(0,\tau) = e^{-q\tau}$$

The probability that the time between two arrivals, A, is larger than τ is:

$$P(A > \tau) = e^{-q\tau}$$

Therefore, the probability A is equal or lower than τ is:

$$P(A \leq \tau) = 1 - e^{-q\tau}$$

which is the cumulative function of the exponential distribution.

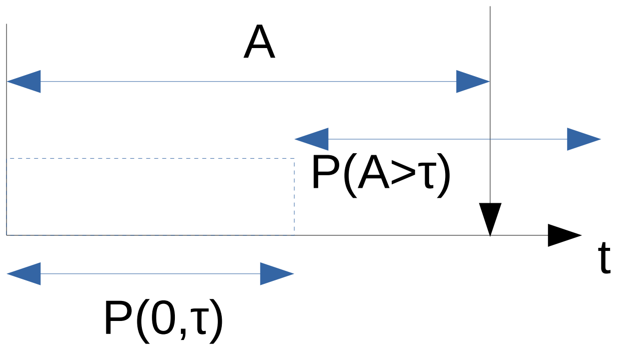

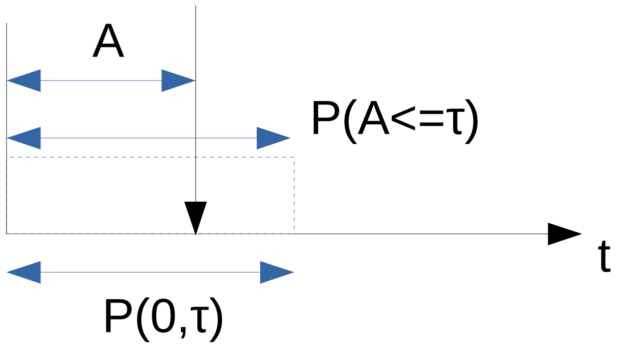

Poisson and Exponential distribution

Then, we conclude that the time between two events, A, is exponentially distributed:

$$f_A(\tau) = qe^{-q\tau} \tag{3.50}$$

Properties of the Poisson process

- **Aggregation of Poisson processes**: the aggregation of several Poisson processes results in a new Poisson process with a rate $\lambda_{\text{aggregate}}$ which is the sum of the rates λ_i of the individual Poisson processes aggregated.

$$\lambda_{\text{aggregate}} = \sum_i \lambda_i \tag{5.13}$$

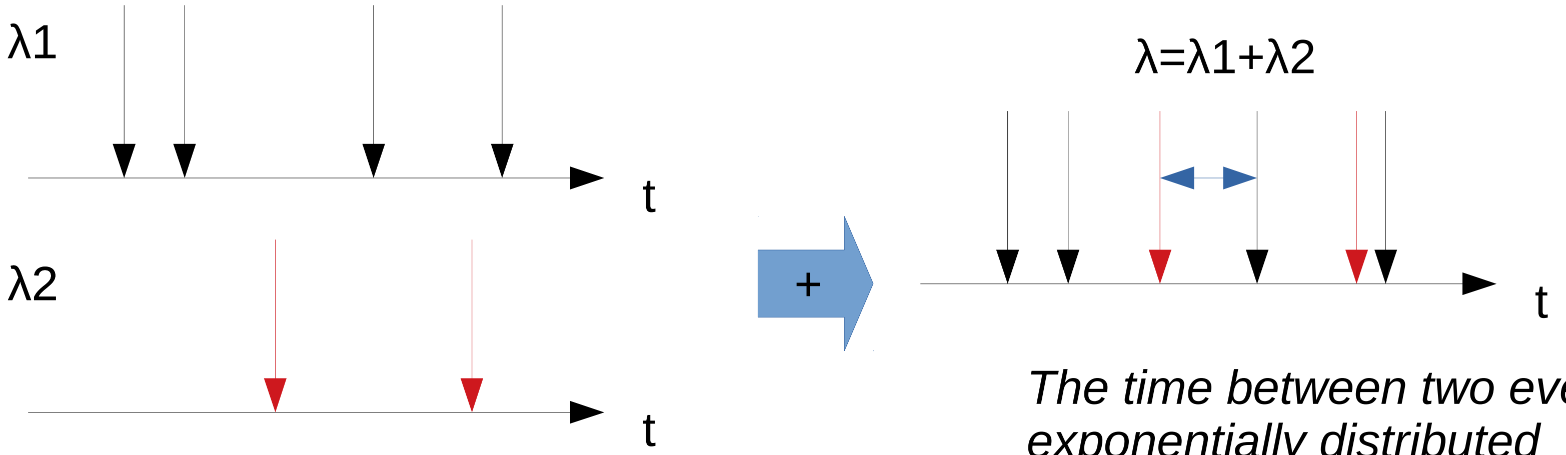

The time between two events is exponentially distributed

Properties of the Poisson process

- **Splitting a Poisson process** in several other processes selecting packets randomly with constant probability over the time causes the resulting processes to be also Poisson. For example, if we split a Poisson process in two Poisson processes, we obtain:

$$\lambda_1 = \alpha_1 \lambda_{\text{aggregate}} \tag{5.14}$$

$$\lambda_2 = \alpha_2 \lambda_{\text{aggregate}} \tag{5.15}$$

$$1 = \alpha_1 + \alpha_2 \tag{5.16}$$

$$\tag{5.17}$$

with α_1 being the probability that a packet belongs to the resulting Poisson process 1, and α_2 the opposite. Note that to obtain several Poisson processes from a single Poisson process, the assignation of a packet to the resulting process must be independent of previous decisions (i.e., stochastic).

PASTA = Poisson Arrivals see Time Averages

- When a new Poisson distributed event happens (i.e., the arrival of a packet in a buffer), it finds the system in state i with probability π_i

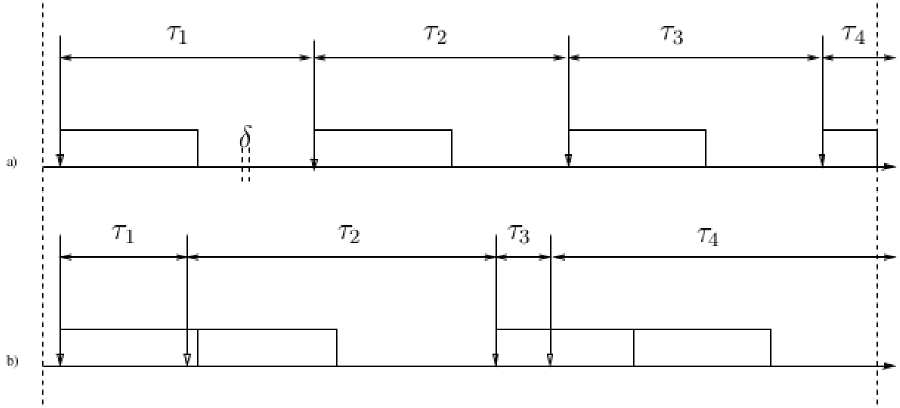

Figure 5.2: Example of the PASTA property. From the Figure, we can see that $\pi_0 = 0.5$ and $\pi_1 = 0.5$. In case a), the interarrival time is deterministic, and all packet arrivals find the system in the empty state. In case b), the interarrival time is exponentially distributed, and 2 packet arrivals find the system in state 0, and two in state 1. As we have 4 arrivals, the probability that an arrival observes the system in state i is the same as the equilibrium probability that the system is in state i (i.e. π_i)

Properties of the exponential distribution

- Memoryless property

Events must be able to arrive at any time (uniformly at random)

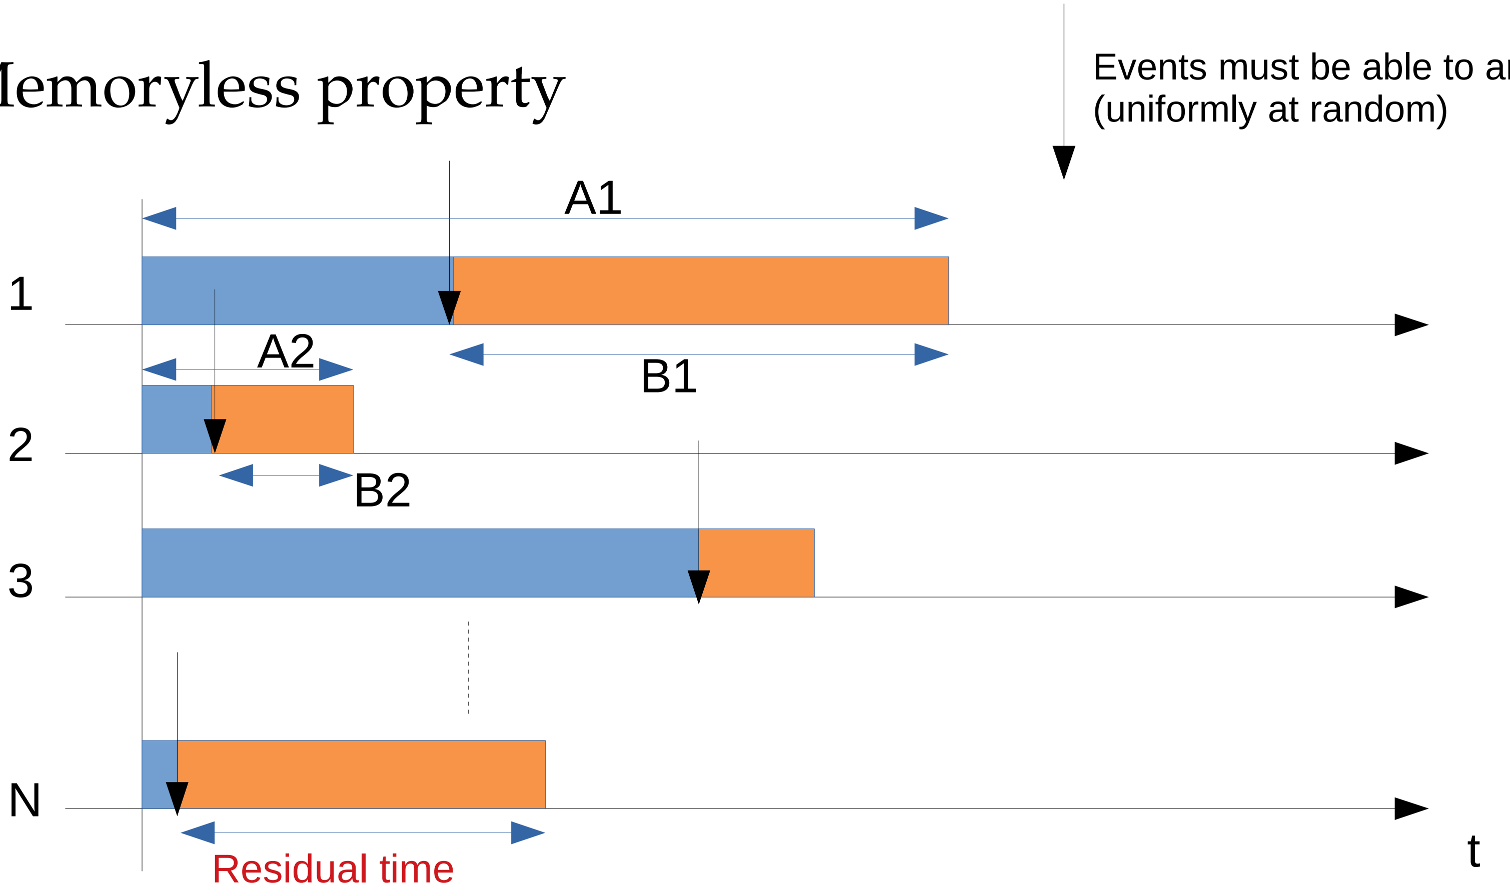

Properties of the exponential distribution

- We obtain the pdf of both data series:
  - A1,A2,A3,A4,A5,…,AN
  - B1,B2,B3,B4,B5,….,BN
- If Ai’s are exponentially distributed with parameter λ, Bi’s are also exponentially distributed with parameter λ.
- So, if the average duration of the A events is $1/\lambda$, the average duration of the B events is also $1/\lambda$.
- What does it mean? For example, I know trains arrive at the station randomly following an exponential distribution, with an average of 10 mins. Then, if I always arrive exactly when a train departs, I have to wait 10 mins in average for the next train.
  - How much time I have to wait if I arrive some time after the departure of a train to take the next one?
    - 10 mins! (in average)

Example

- Consider a network interface where packets arrive at a rate λ [packets/second] and depart at a rate μ [packets/second].
  - Events can be two types: 1) packet arrivals, 2) packet departures
  - They may happen at any arbitrary instant of time.
- The network interface has a single transmitter, and the maximum buffer size is Q=2 packets.

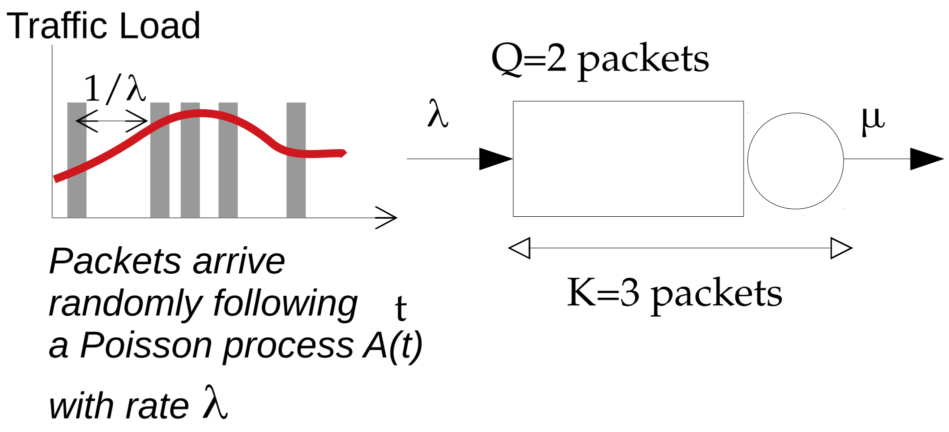

Example

- Let us assume we divide the time in very small intervals of size δ, satisfying all previous explained requirements. We have the following events:
  - The probability that in a given interval we have a packet arrival is λδ.
  - Similarly, the probability that in a given interval we have a packet departure is μδ.
  - And the probability that nothings happens is: 1-λδ-μδ.

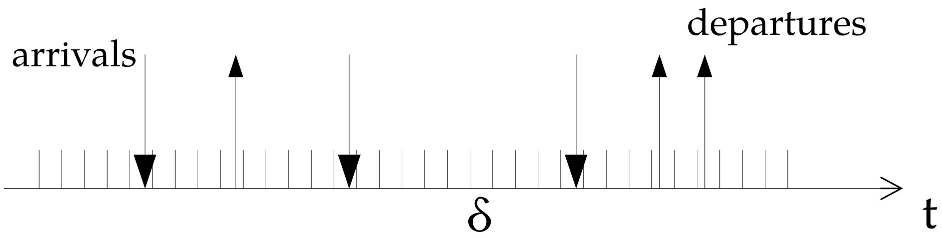

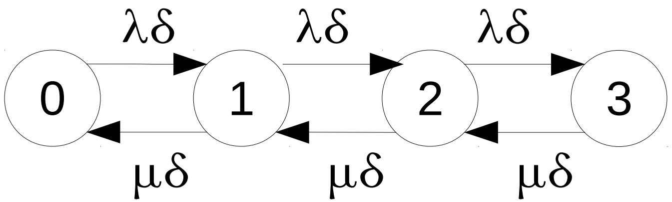

Self-transitions are omitted

CTMC: Equilibrium distribution

- Global balance equations:

$$\pi_i \sum_{\forall j \neq i} q_{i,j}\delta = \sum_{\forall j \neq i} \pi_j q_{j,i}\delta$$

$$\pi_i \sum_{\forall j \neq i} q_{i,j} = \sum_{\forall j \neq i} \pi_j q_{j,i}$$

- Local balance equations:

$$\pi_i q_{i,j} = \pi_j q_{j,i}$$

Solving a CTMC: Infinitesimal Generator Q

The infinitesimal generator is given by the following set of equations (for all i):

$$-\pi_i \sum_{\forall j \neq i} q_{i,j} + \sum_{\forall j \neq i} \pi_j q_{j,i} = 0$$

$$Q = \begin{bmatrix} -\sum_{\forall j \neq 0} q_{0,j} & q_{0,1} & q_{0,2} & \cdots \\ q_{1,0} & -\sum_{\forall j \neq 1} q_{1,j} & q_{1,2} & \cdots \\ q_{2,0} & q_{2,1} & -\sum_{\forall j \neq 2} q_{2,j} & \cdots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$

We can obtain the stationary distribution by solving

$$Q\pi = 0$$

where π is a vector representing the stationary distribution.

Example

- In a data center, a server equipped with 4 CPUs receives groups of ‘tasks’ with rate λ [tasks/second] following a Poisson process. There are two different types of groups: groups containing 1 task, and groups containing 2 tasks. The probability that given an arrival it corresponds to a group containing a single task is 0.4, and so, the probability it corresponds to a group of 2 tasks is 0.6.
  - Therefore: $\lambda 1=0.4\lambda$ and $\lambda 2=0.6\lambda$
- The time to complete a task is exponentially distributed with average value E[Ds].
  - Therefore, the rate at which tasks are completed is $\mu=1/E[Ds]$
- In case a group of tasks arrives to the server and there are not enough available CPUs, the ones that cannot be accommodated are dropped.
- Exercise: Considering the stochastic process X(t) that models the number of tasks in the system, find its stationary probability distribution.

Example

- The first thing to do is to find the state space of X(t). Since X is the number of tasks in the server, it can take the following values: 0, 1, 2, 3 and 4.
- Second, we can represent the state space and the transitions (CTMC)

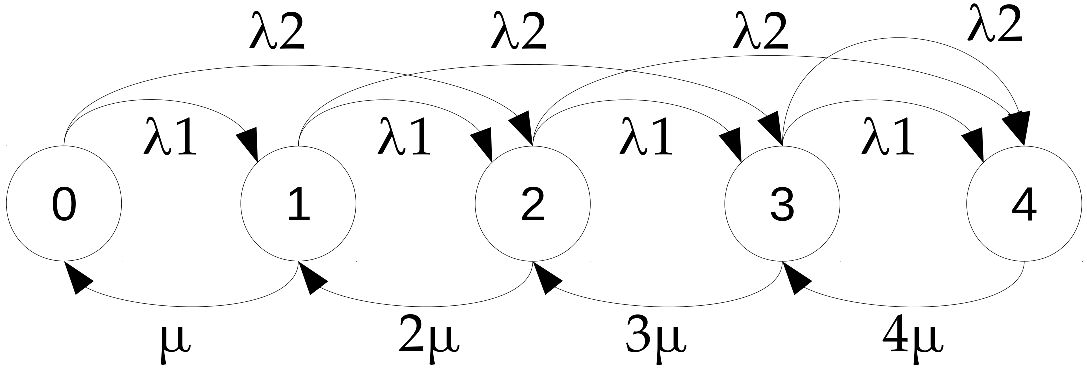

- Finally, we can write the balance equations, and solve the resulting system of equations to find the stationary probability distribution.

Example

- Note that the CTMC is not reversible, so local balance does not holds.
- Then, we must apply the global balance condition, and solve the resulting system of equations.
- Alternatively, given Q, we can obtain the stationary probability distribution as follows:
  - `std=mrdivide([zeros(1,size(Q,1)) 1],[Q ones(size(Q,1),1)]);`

Example

```
function ExampleCTMCs_Tasks()

lambda=10;
EDs=0.04;
mu = 1/EDs;

lambda1=0.4*lambda;
lambda2=0.6*lambda;

Q=[-lambda1-lambda2 lambda1 lambda2 0 0;
   mu -mu-lambda1-lambda2 lambda1 lambda2 0;
   0 2*mu -2*mu-lambda1-lambda2 lambda1 lambda2;
   0 0 3*mu -3*mu-lambda1-lambda2 lambda1+lambda2;
   0 0 0 4*mu -4*mu];

disp('Infinitesimal generator Q');
disp(Q);

std=mrdivide([zeros(1,size(Q,1)) 1],[Q ones(size(Q,1),1)]); % the left null space of Q is equivalent to solve [pi] * Q =  [0 0 ... 0 1]

disp('Stationary Probability Distribution');
disp(std);
```

```
Infinitesimal generator Q
  -10     4     6     0     0
   25   -35     4     6     0
    0    50   -60     4     6
    0     0    75   -85    10
    0     0     0   100  -100

Stationary Probability Distribution
    0.5965    0.2386    0.1193    0.0350    0.0107
```

$\pi_0=0.6;\ \pi_1=0.24;\ \pi_2=0.12;\ \pi_3=0.035;\ \pi_4=0.01$

Example

- By consider that π_i is the fraction of time the system spends in state i, we can say that:
  - The system is idle the 60 % of the time.
  - The system has the 4 CPUs working only the 10 % of the time.
  - Etc.
- Similarly, what is the probability that a new group of tasks find the system in state 2 when it arrives?
  - $\pi_2=0.24$

Modelling WIFI networks

- We can model CSMA/CA networks using CTMCs if:
  - Backoff period are continuous, and exponentially distributed.
  - The duration of the transmissions are exponentially distributed.
  - Propagation delay in the WLAN is negligible

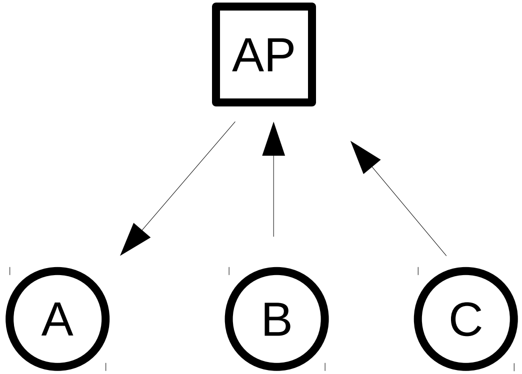

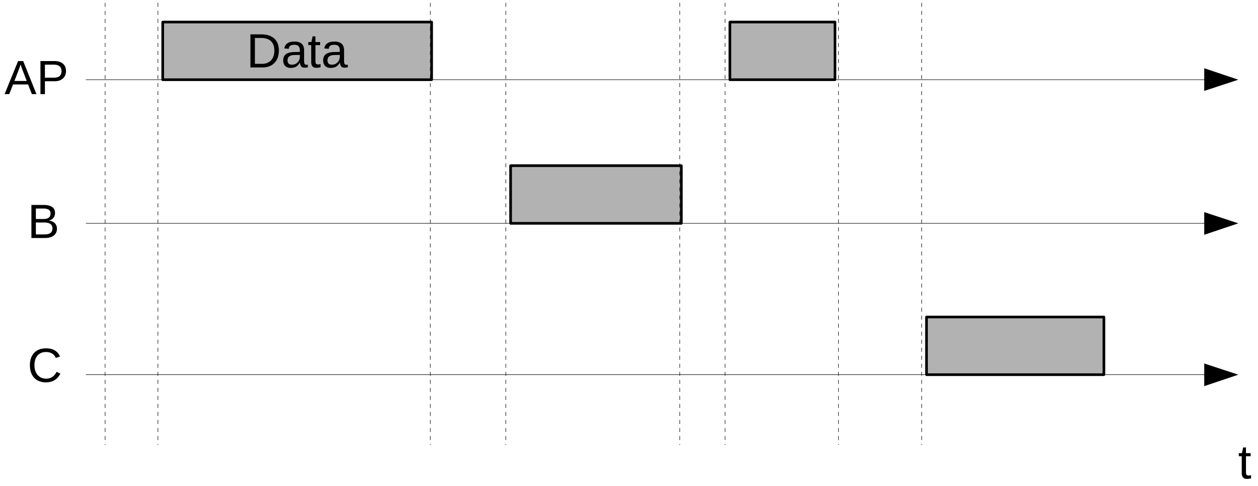

Example

- The duration of a backoff is a random variable exponentially distributed with expected value 0.01 ms, so $\lambda=1/0.01 = 1E5$ tx/second
- The duration of a transmission is a random variable exponentially distributed with expected value 0.003 s, so $\mu=1/0.003 = 333.33$ packets/sec.
- The CTMC representing the stochastic process is:

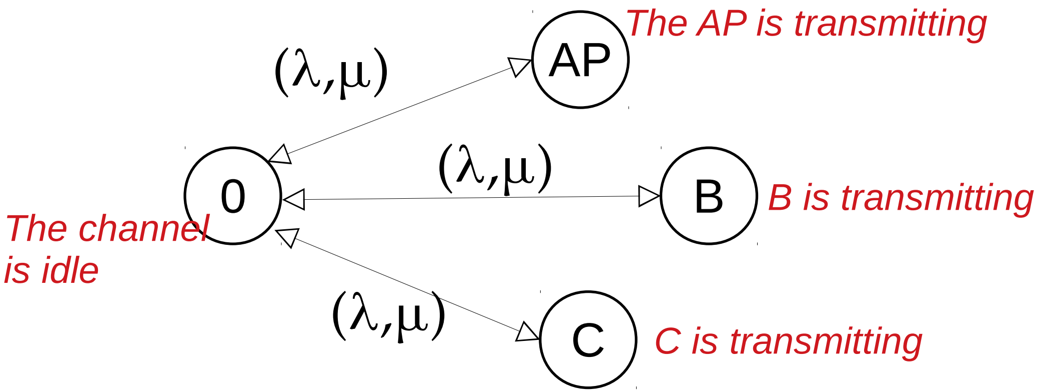

How much time (in %) the AP is transmitting?