17 Applications of the Exponential Distribution

Failure Rate and Reliability

Example 1

The length of life in years, T, of a heavily used terminal in a student computer laboratory is exponentially distributed with $\lambda = .5$ years, i.e.

$$\begin{aligned} f(t) &= .5e^{-.5t}, \quad t \geq 0, \\ &= 0, \quad \text{otherwise.} \end{aligned}$$

$\lambda = .5$ is called the **failure rate** of the terminal.

Reliability of t.

$$Rel(t) = P(T > t) = e^{-.5t}$$

For example

$$Rel(1) = P(T > 1) = e^{-.5} = 0.607.$$

the probability that the terminal will last more than 1 year.

Use *R* to calculate $P(X \leq 1)$:

```
> pexp(1, .5)
[1] 0.3934693
```

And so the reliability $P(X > 1)$:

```
> 1-pexp(1, .5)
[1] 0.6065307
```

Also

$$Rel(2) = P(T > 2) = e^{-1} = 0.368$$

In *R*

```
> 1-pexp(2, .5)
[1] 0.3678794
```

Use *R* to examine the reliability as t increases.

```
curve(exp(-.5*x), 0, 7,
  xlab = "Time in years",
  ylab = "Reliability", type = "l")
```

Reliability of a Terminal with failure rate λ =.5

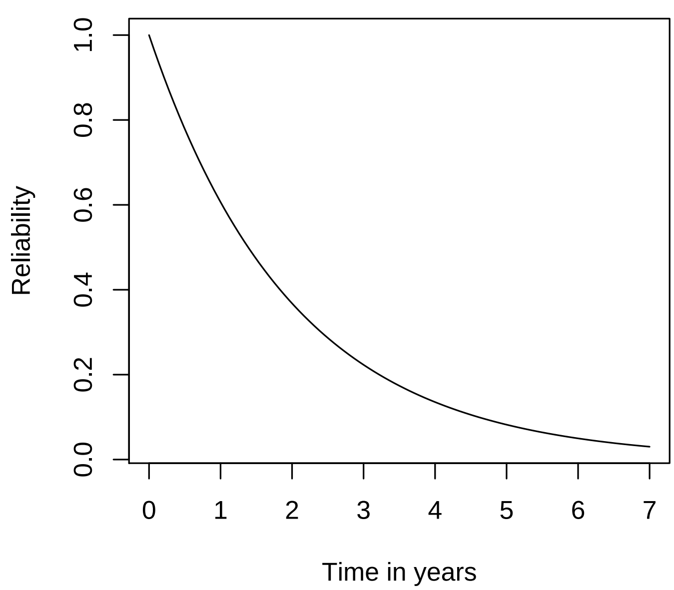

We might want to know when the reliability is just 5%:

Choose k so that

$$P(T > k) = .05$$

or equivalently, choose k so that

$$P(T \leq k) = .95$$

From *R*

```
qexp(.95, .5)
[1] 5.991465
```

95% likelihood that the machine will last less than 6 years

just a 5% chance that it will last longer than 6 years.

Rel(6) = .05.

Example:

Suppose that there are 100 terminals in the laboratory, how long will it take to have 90 of them still working, or equivalently 10% of them broken down.

We need to find k so that

$$P(T > k) = .9$$

or equivalently

$$P(T \leq k) = .1$$

In *R* :

```
qexp(.1,.5)
[1] 0.2107210
```

i.e.

$$P(T \leq .21 \text{ years}) = .1$$

a 10% chance that the terminal will last up to .21 years (3 months approx),

or a 90% chance that the lifetime will be greater than 3 months.

After just under 3 months 10 of the 100 terminals broken and about 90 still working.

To check it in *R*

```
pexp(.21, .5)
[1] 0.09967548
```

In general, if the lifetime of a machine is modeled by an exponential distribution of the form

$$\begin{aligned} f(t) &= \lambda e^{-\lambda t}, \quad t > 0 \\ &= 0, \text{ otherwise} \end{aligned}$$

then λ is the failure rate of the machine

$Rel(t) = e^{-\lambda t}$ is the reliability of the machine at time t.

Because the exponential distribution enjoys the Markov property,

$$P(T > t + s | T > t) = P(T > s)$$

i.e.

$$Rel(t + s | T > t) = Rel(s)$$

For example

$$Rel(5 | T > 2) = Rel(3)$$

which means that the probability that the terminal will last 3 years more after lasting for 2 years, is the same as the probability lasting 3 years from the start.

Breakdown is a result of some sudden failure, not wear and tear.

Example

Studies of a single-machine-tool system showed that the time the machine operates before breaking down is exponentially distributed with a mean 10 hours.

1. Determine the failure rate and the reliability.
2. Find the probability that the machine operates for at least 12 hours before breaking down.
3. If the machine has already been operating 8 hours, what is the probability that it will last another 4 hours?

$$E(T) = 10 \text{ hours}$$

Also for the exponential distribution:

$$E(T) = \frac{1}{\lambda}.$$

So $1/\lambda = 10$, giving $\lambda = 0.1$

The failure rate is 0.1

The pdf is

$$\begin{aligned} f(t) &= 0.1e^{-0.1t}, \quad t > 0 \\ &= 0 \text{ otherwise} \end{aligned}$$

Solution:

1. Failure rate = .1 hours
   Reliability: $Rel(t) = e^{-.1t}$.

2. The reliability at $T = 12$:

   $$Rel(12) = e^{-1.2} = .30$$

   In *R*

   ```
   > 1-pexp(12, .1)
   [1] 0.3011942
   ```

   i.e. just aver 30% chance that the machine will last longer than 12 hours.

3. the likelihood that the machine will last at least 12 hours given that it has already lasted 8 hours.

We seek:

$$P(T > 12 \text{ hours } |T > 8)$$

From Markov;

$$P(T > 12 \text{ hours}|T > 8 \text{ hours}) = P(T > 4) = e^{-4(0.1)} = .67032$$

In *R*

```
> 1-pexp(4, .1)
[1] 0.67032
```

$$Rel(12|T > 8) = Rel(4) = .67$$

Applications of the Exponential Distribution

Modelling Response Times

A Single Server Queue

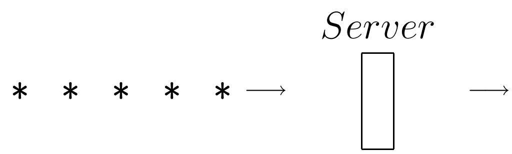

It is assumed

- jobs arrive at a single server, in accordance with a Poisson distribution.

- job is processed immediately if the queue is empty, otherwise joins the end of the queue.

- Service times are assumed to be exponentially distributed.

Called M/M/1; the two Ms refer to the arrival and service times being exponential and hence enjoy the **M**arkov or **M**emoryless property, while the 1 refers to the single server.

Arrivals

pdf

$$f(t) = \lambda e^{-\lambda t}$$

Average inter-arrival time is $1/\lambda$.

Processing time

pdf

$$f(s) = \mu e^{-\mu s}$$

cdf is

$$F(s) = P(S \leq s) = 1 - e^{-\mu s}$$

Average processing time is $1/\mu$.

Response rate depends on the arrival rate λ and the processing rate μ.

Traffic Intensity (I)

arrival rate relative to the process rate.

$$I = \frac{\lambda}{\mu}$$

- $I > 1$

  i.e. $\lambda > \mu$:

  average arrival rate exceeds the average processing rate.

- $I = 1$

  i.e. $\lambda = \mu$:

  the average arrival rate equals the average processing rate.

- $I < 1$

  i.e. $\lambda < \mu$:

  jobs are being processed faster than they arrive.

Queue Lengths

Example:

Supposing that jobs arrive at the rate of 4 per minute, use *R* to examine the queue length when the service rate is:

(a) 3.8 per minute,

(b) 4 per minute,

(c) 4.2 per minute.

You can assume that both the service times and the inter-arrival times are exponentially distributed.

Solution:

(a) service rate $\mu = 3.8 < \lambda = 4$, the arrival rate.

Therefore,
the traffic intensity $I = \lambda/\mu \approx 1.05 > 1$.

In this case we would expect the queue to increase indefinitely.

Simulating Queues in *R*

In *R*

$$rpois$$

generates random observations from a Poisson distribution.

For example:

$$rpois(10000, 4)$$

generates 10,000 Poisson $\lambda = 4$

R program for queue length

```
arrivals<- rpois(10000, 4)
service <-rpois(10000, 3.8)
queue[1] <-max(arrivals[1] - service[1], 0)
for (t in 2:10000)
queue[t] = max(queue[t-1]+arrivals[t]-service[t], 0)
plot(queue, xlab = "Time", ylab = "Queue length")
```

Queue length with traffic intensity $>$ 1.

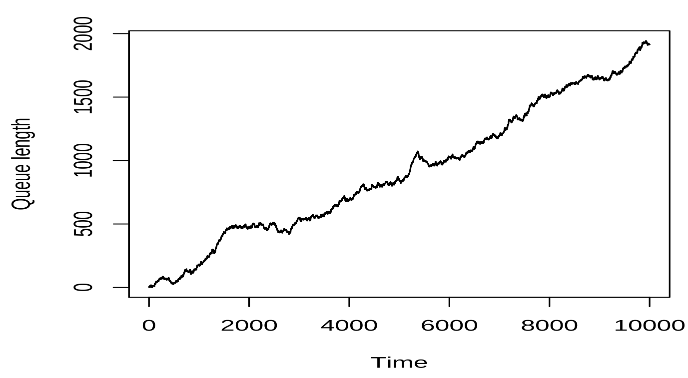

(b)**Queue length when the traffic intensity = 1**

service rate $\mu = 4$ and $\lambda = 4$ arrival rate.

Traffic intensity:

$$I = \lambda/\mu = 1.$$

Investigate with *R*

```
arrivals <- rpois(10000, 4)
service <-rpois(10000, 4)
queue[1] <-max(arrivals[1] - service[1], 0)
for (t in 2:10000)
queue[t] = max(queue[t-1]+arrivals[t]-service[t], 0)
plot(queue, xlab = "Time", ylab = "Queue length")
```

Queue pattern when the traffic intensity = 1

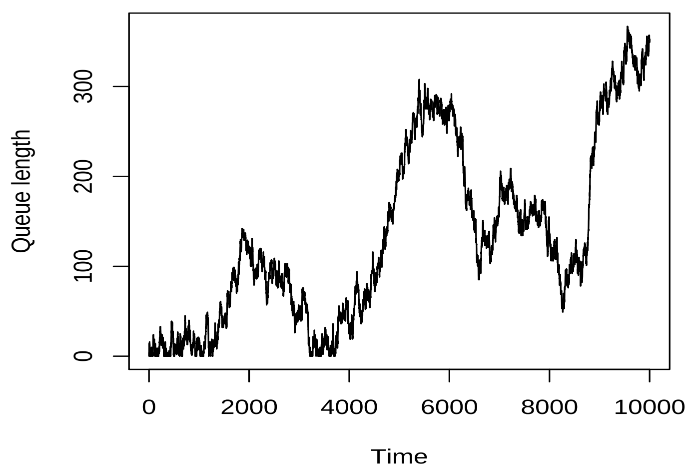

(c) **Queue length when the traffic intensity ¡ 1.**

service rate $\mu = 4.2 > \lambda = 4$ arrival rate

Traffic intensity $I = \lambda/\mu < 4/4.2 \approx 0.95 < 1$.

```
arrivals <- rpois(10000, 4)
service <-rpois(10000, 4.2)
queue[1] <-max(arrivals[1] - service[1], 0)
for (t in 2:10000)
queue[t] = max(queue[t-1]+arrivals[t]-service[t], 0)
plot(queue, xlab = "Time", ylab = "Queue length")
```

Queue pattern when the traffic intensity $I < 1$.

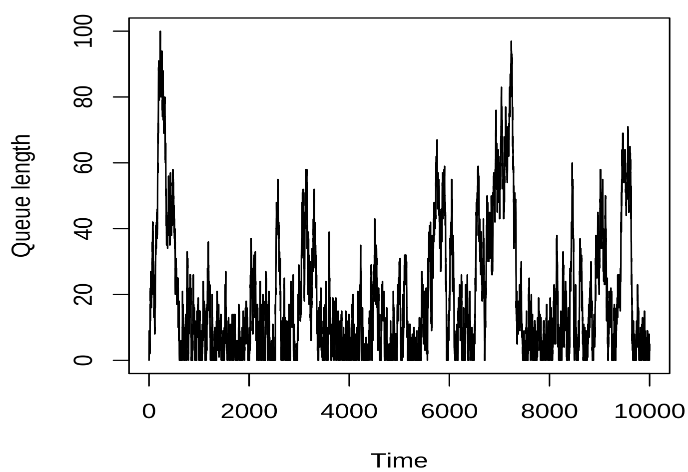

Queuing Statistics when $I < 1$

```
mean(queue)
[1] 18.4218
```

an average of 18.4218 jobs in the queue.

Average queuing time

```
mean(queue)*(1/4)
[1] 4.6
```

4.6 minutes spent in the queue on average.

Longest lengh:

```
max(queue)
[1] 103
```

Longest wait;

```
max(queue)/4
[1] 25.75
```

less than 26 minutes.

Best case

```
min(queue)
[0]
```

Queue length with arrival rate $\lambda = 4$ and processing rates $\mu = 3.8$, 4 and 4.2 jobs per minute