CSC 2232:
Topics in Computer System Performance and Reliability

Bianca Schroeder

Department of Computer Science
University of Toronto

THE EXPONENTIAL DISTRIBUTION, THE POISSON PROCESS AND THE REAL WORLD

- What are the characteristics of the exponential distribution and the Poisson process?
- Why do people like it so much?
- What does the real world look like?

A QUICK REVIEW

- The cumulative distribution function (CDF)
  - $F(x)$ is the probability that $X \leq x$
  - $P(a < X \leq b) = F(b) - F(a)$
  - $\overline{F}(x) = 1 - F(x)$ is the probability that $X > x$
- The probability density function (PDF)
  - $f(x)$ is the derivative of $F(x)$
  - $f(x)dx$ is the probability that X is in $\{x, x+dx\}$
  - $P(a < X \leq b) = \int_a^b f(x)dx$
- The mean of X is:
  - $E[X] = \int_{-\infty}^{\infty} x\, f(x)dx$

THE EXPONENTIAL DISTRIBUTION

- If X is distributed exponentially with rate λ then
  - $$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases}$$
  - $$F(x) = \begin{cases} 1 - e^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases}$$
  - $$\overline{F}(x) = \begin{cases} e^{-\lambda x} & \text{for } x \geq 0 \\ 1 & \text{for } x < 0 \end{cases}$$

THE EXPONENTIAL DISTRIBUTION

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases}$$

$$\bar{F}(x) = \begin{cases} e^{-\lambda x} & \text{for } x \geq 0 \\ 1 & \text{for } x < 0 \end{cases}$$

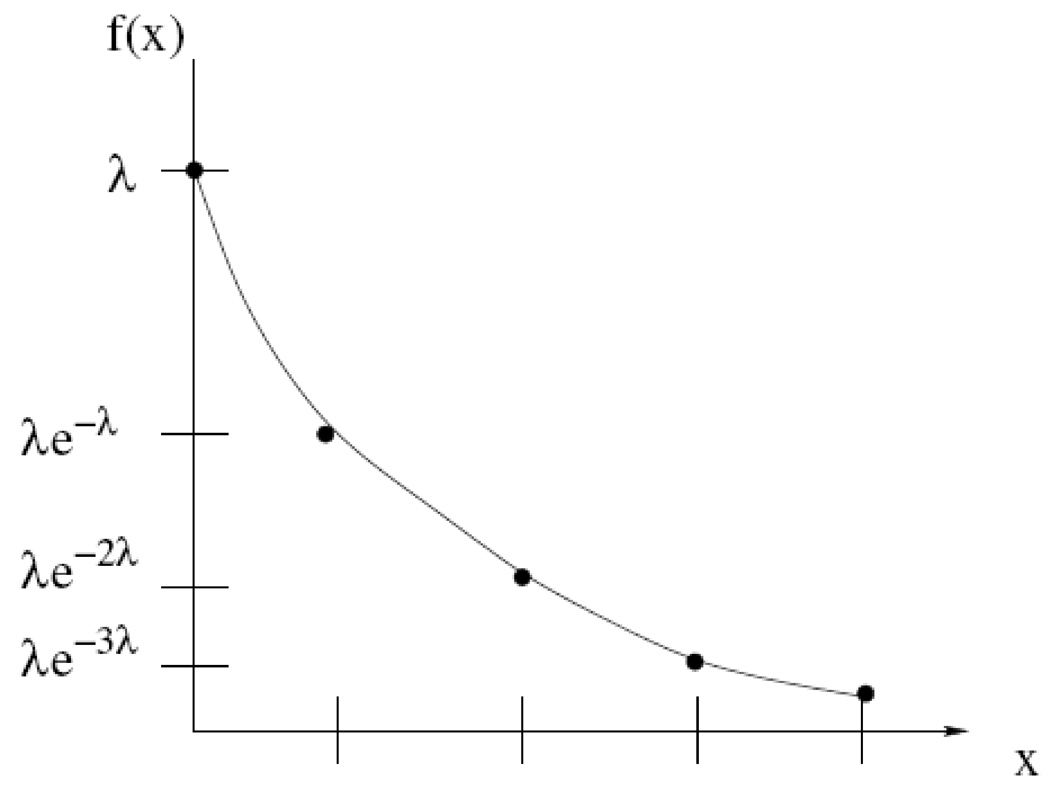

THE MEAN AND VARIANCE OF THE EXPONENTIAL DISTRIBUTION

- $f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases}$

- $E[X] = \int_{-\infty}^{\infty} x\, f(x)dx = 1/\lambda$

- $\mathrm{Var}(X) = 1/\lambda^2$ $(\mathrm{Var}(X) = E[X^2] - (E[X])^2$

- Squared coefficient of variation:
  - $C.V.^2 = \text{variance}/\text{mean}^2 = 1$

THE MINIMUM OF TWO EXPONENTIALS

- Given independent variables X1~ Exp(λ1) and X2~Exp(λ2)
- Let X = min{X1, X2}, then X~Exp(λ1+ λ2)

- Proof:
- $\Pr(X>t) = \Pr\{\min\{X1, X2\} > t\}$
- $= \Pr\{X1>t \text{ and } X2> t\}$
- $= \Pr\{X1>t\} * \Pr\{X2>t\}$
- $= e^{-\lambda 1 t} * e^{-\lambda 2t}$
- $= e^{-(\lambda 1+ \lambda 2)t}$

THE MEMORYLESS PROPERTY

- A random variable X is said to be memoryless if
- $\Pr\{X > s + t \mid X > s\} = \Pr\{X > t\}$

- Show that $X \sim \mathrm{Exp}(\lambda)$ is memoryless

  - $\Pr\{X > s + t \mid X > s\} = \Pr\{X > s+t\} / \Pr\{X > s\}$

$$= e^{-\lambda(s+t)} / e^{-\lambda s} = e^{-\lambda t}$$

$$= \Pr\{X>t\}$$

THE EXPONENTIAL DISTRIBUTION

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases} \qquad \bar{F}(x) = \begin{cases} e^{-\lambda x} & \text{for } x \geq 0 \\ 1 & \text{for } x < 0 \end{cases}$$

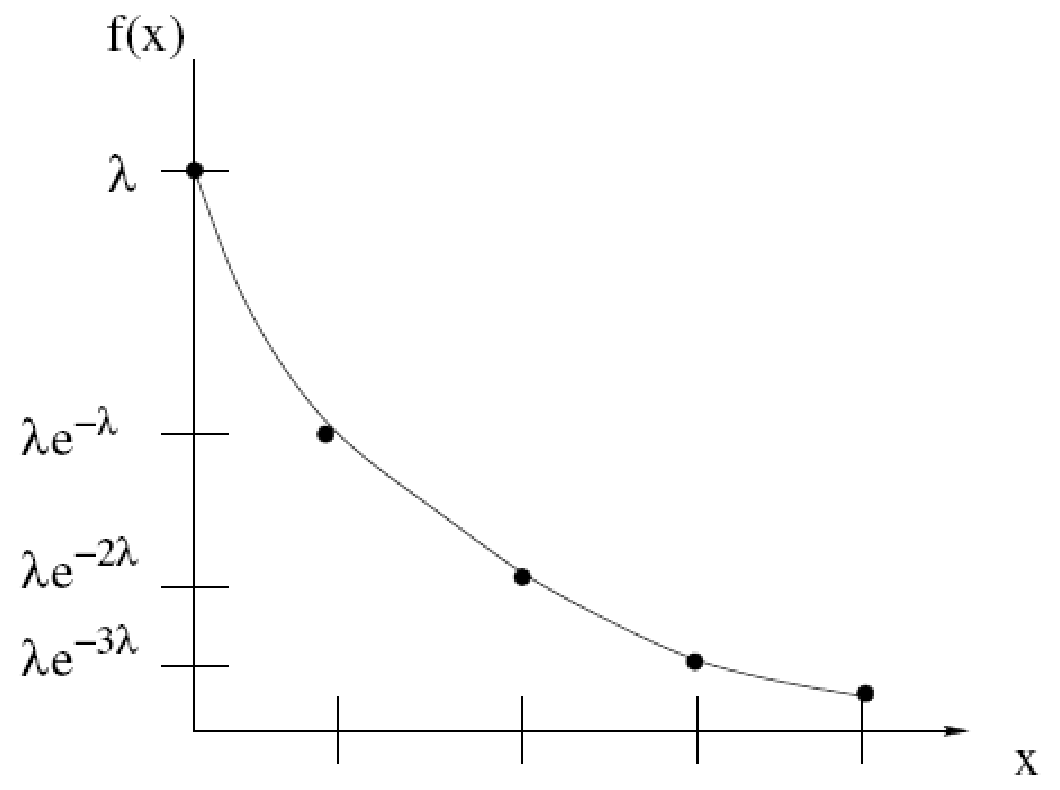

THE MEMORYLESS PROPERTY

- Is the memoryless property realistic for lifebulbs?
- Real-life examples whose lifetime can be modeled by X such that $\Pr\{X > s+t\} \mid X > s\}$ goes down as s goes up? ⇨ Increasing failure rate
- Real-life examples whose lifetime can be modeled by X such that $\Pr\{X > s+t\} \mid X > s\}$ goes up as s goes up? ⇨ Decreasing failure rate
- Failure rate $r(t) := f(t) / (1-F(t))$

EXAMPLE 1

- The time a customer spends waiting in a bank is exponentially distributed with mean 10 min.
- What is Pr[cust. waits> 5 min]
  - $\bar{F}(5min) = e^{-5/10} = e^{-0.5}$
- Pr[cust. waits > 15 min} | still in bank after 10 min]
  - $\bar{F}(5min) = e^{-5/10} = e^{-0.5}$

EXAMPLE 2

- A post office has 2 clerks. Cust. B is served by one clerk and cust. C is served by the other clerk, when A walks in. All service times are exponentially distributed.
- What is the probability that A is last to leave?
- ½

NEXT:

- The Poisson process
- But need some definitions first.

DEFINITIONS

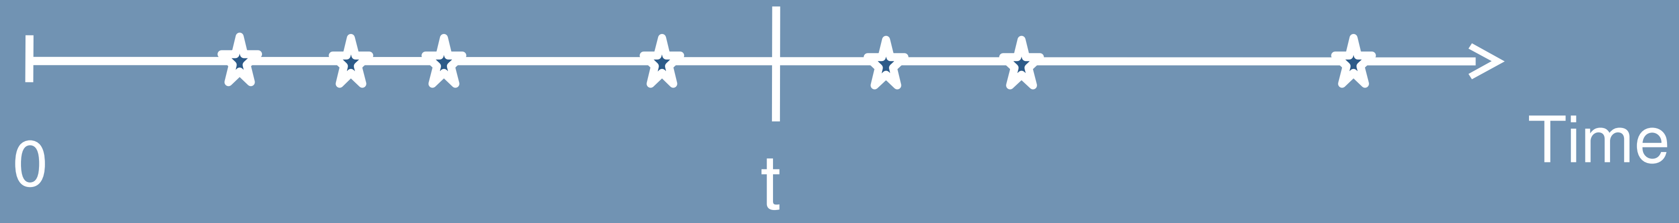

- $N(t)$:= number of events that occurred by time t

- An event sequence has independent increments if the number of events at different time intervals are independent.

- I.e. for all $t0 < t1 < t2 < \ldots. < tn$ the rand. variables

  $$N(t1)-N(t0), \quad N(t2)-N(t1), \quad \ldots\ldots\ N(tn)-N(tn-1)$$

are independent.

EXAMPLES

- Consider the following event processes:
    - People entering a store
    - Births of children
    - Goals scored by a particular soccer player?

- Do the above processes have independent increments?

DEFINITION 1 OF POISSON PROCESS

- A Poisson process having rate λ is a sequence of events such that:
  - $N(0) = 0$
  - The process has independent increments
  - The number of events in any interval of length t is Poisson distributed with mean λt, i.e.

$$\Pr\{N(t+s) - N(s)\} = n\} = e^{-\lambda t} (\lambda t)^{n} / n!$$

DEFINITION 2 OF POISSON PROCESS

- A Poisson process with rate λ is a sequence of events such that the inter-arrival times are i.i.d. exponential random variables with rate λ and N(0)=0.

MERGING POISSON PROCESSES

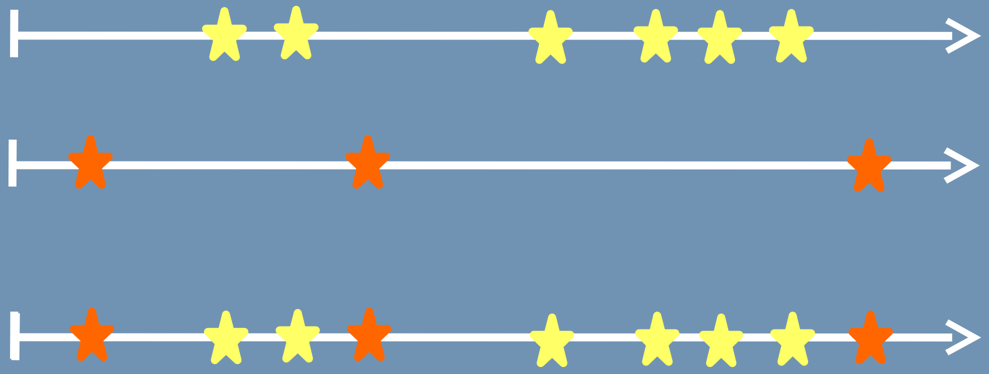

Process 1 with rate λ1

Process 2 with rate λ2

Poisson with rate (λ1+λ2)

- Theorem: The merge of process 1 and process 2 is a single Poisson process with rate (λ1 + λ2)
- Proof:
  - Process 1 has Exp(λ1) inter-arrival times
  - Process 2 has Exp(λ2) inter-arrival times
  - Time until first event has Exp(λ1+λ2)

POISSON SPLITTING

- Given Poisson process with rate λ, suppose each event is randomly classified as type A with prob. p and type B with prob 1-p.

- Then type A events form a Poisson process with rate p * λ
- And type B events form a Poisson process with rate (1-p) * λ

UNIFORMITY

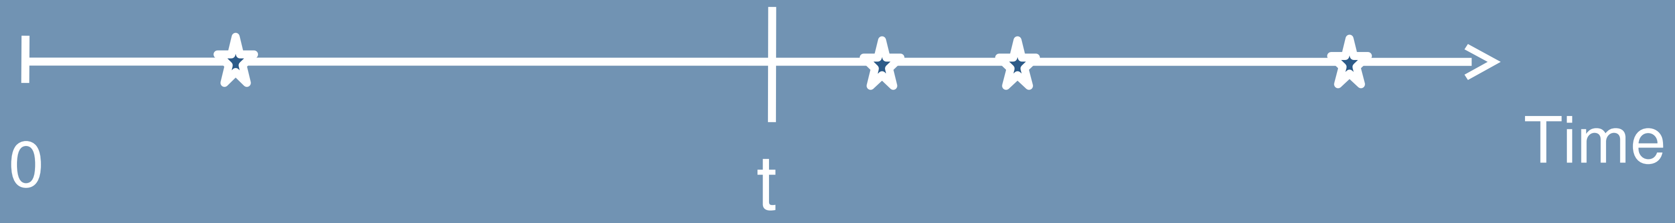

- Given that one event of a Poisson process has occurred by time t, that event is equally likely to have occurred anywhere in [0, t]

- Generalization:
- If k events of a Poisson process occur by time t, the k events are distributed independently and uniformly in [0, t].

SUMMARY OF EXPONENTIAL DISTRIBUTION & POISSON PROCESS

- The exponential distribution is memoryless
  - History doesn’t matter!
- The Poisson process is an event sequence such that inter-arrival times are iid expontial rand. vars.
  - Poisson process can be merged
  - Poisson processes can be split
  - Uniformity of events

A COMMON REAL-WORLD DISTRIBUTION

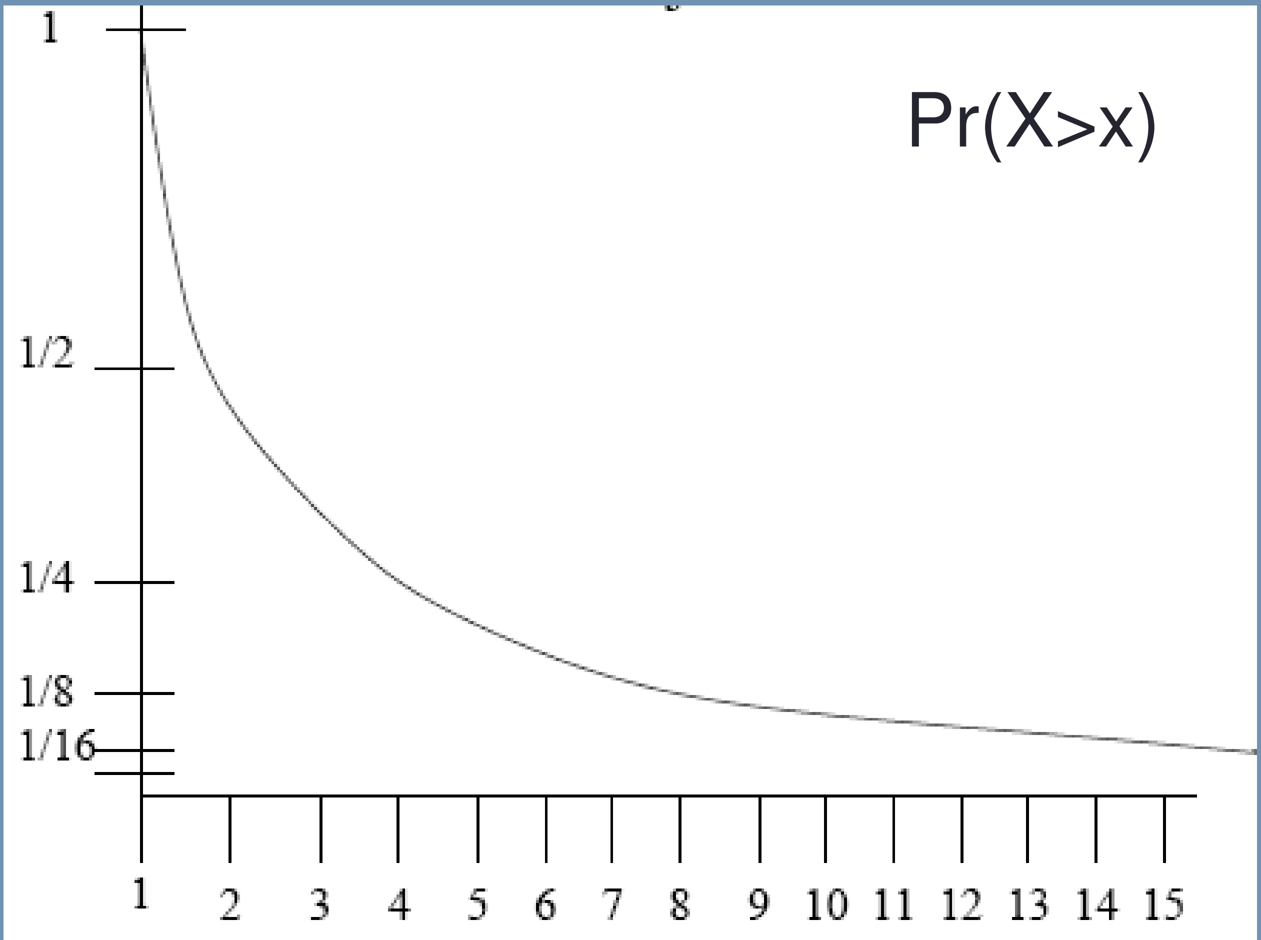

THE EXPONENTIAL DISTRIBUTION

- $$\overline{F}(x) = \begin{cases} e^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases}$$

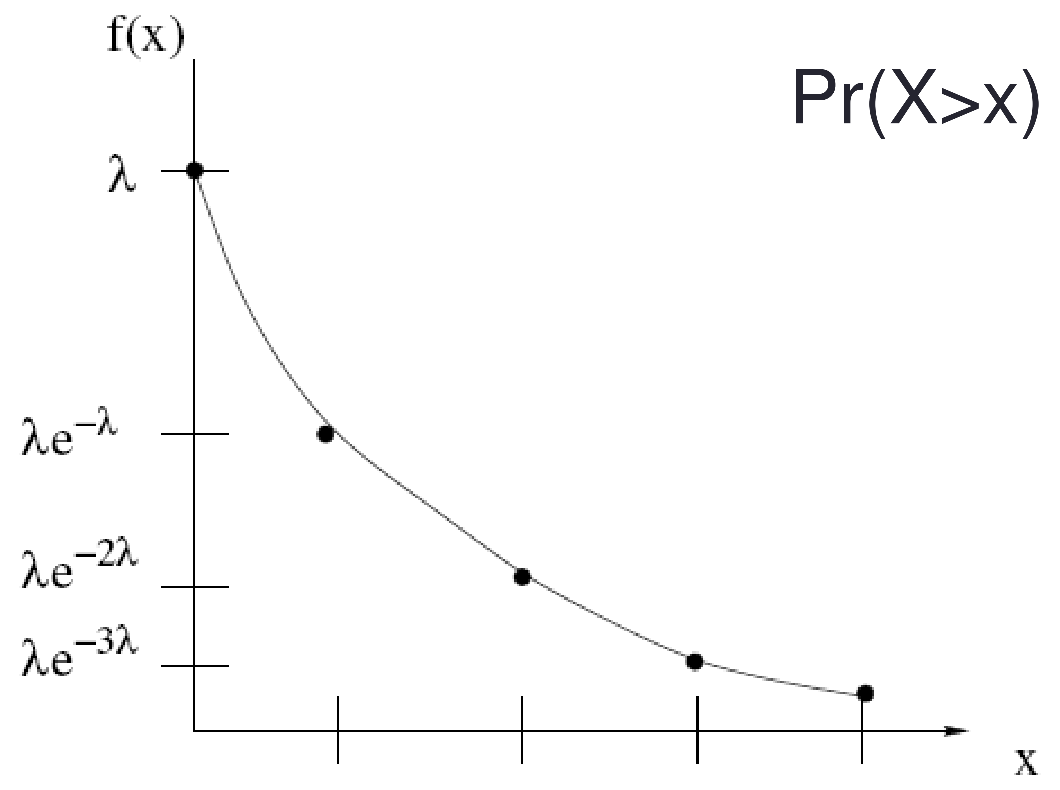

A COMMON REAL-WORLD DISTRIBUTION

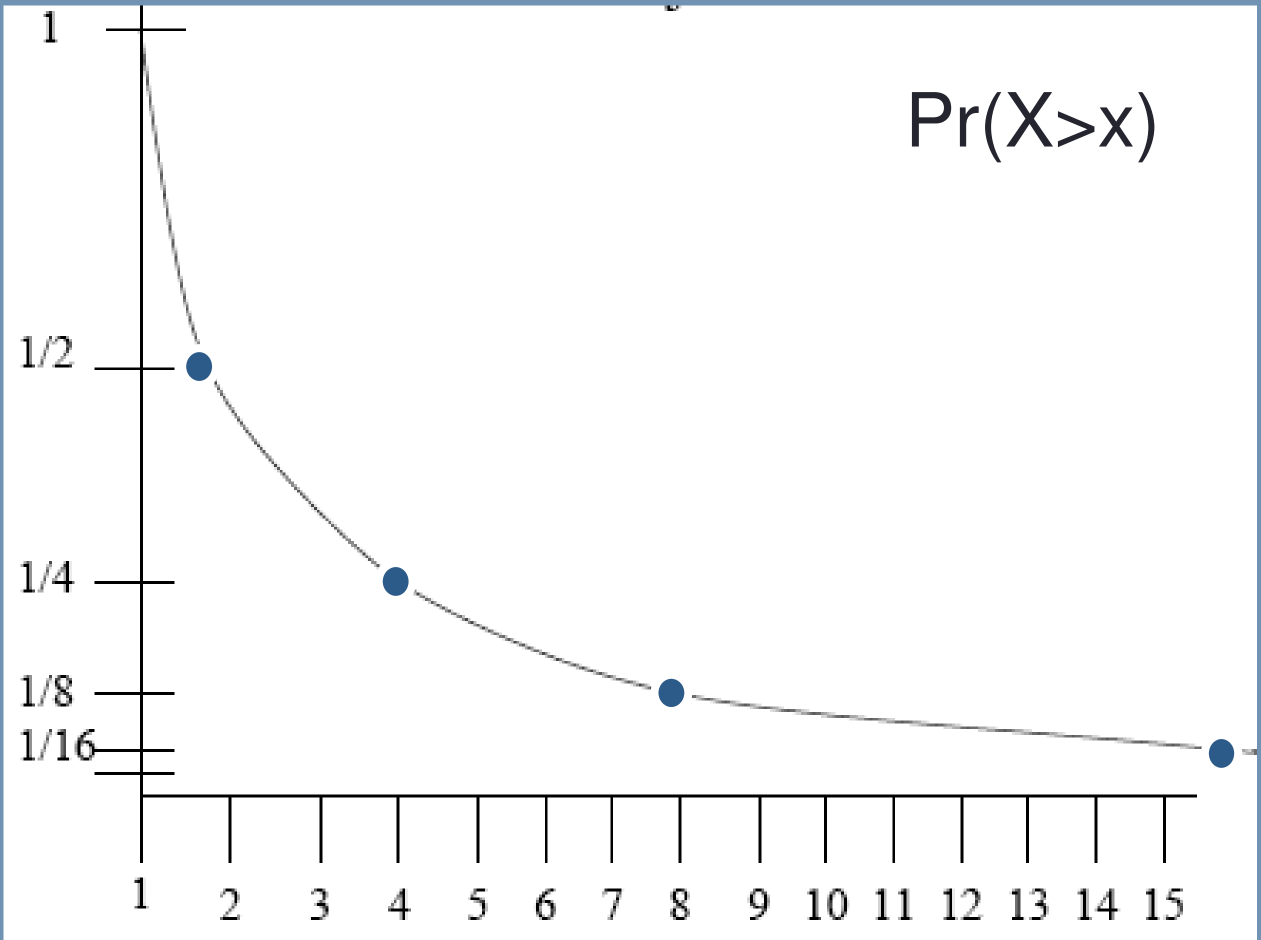

A COMMON REAL-WORLD DISTRIBUTION

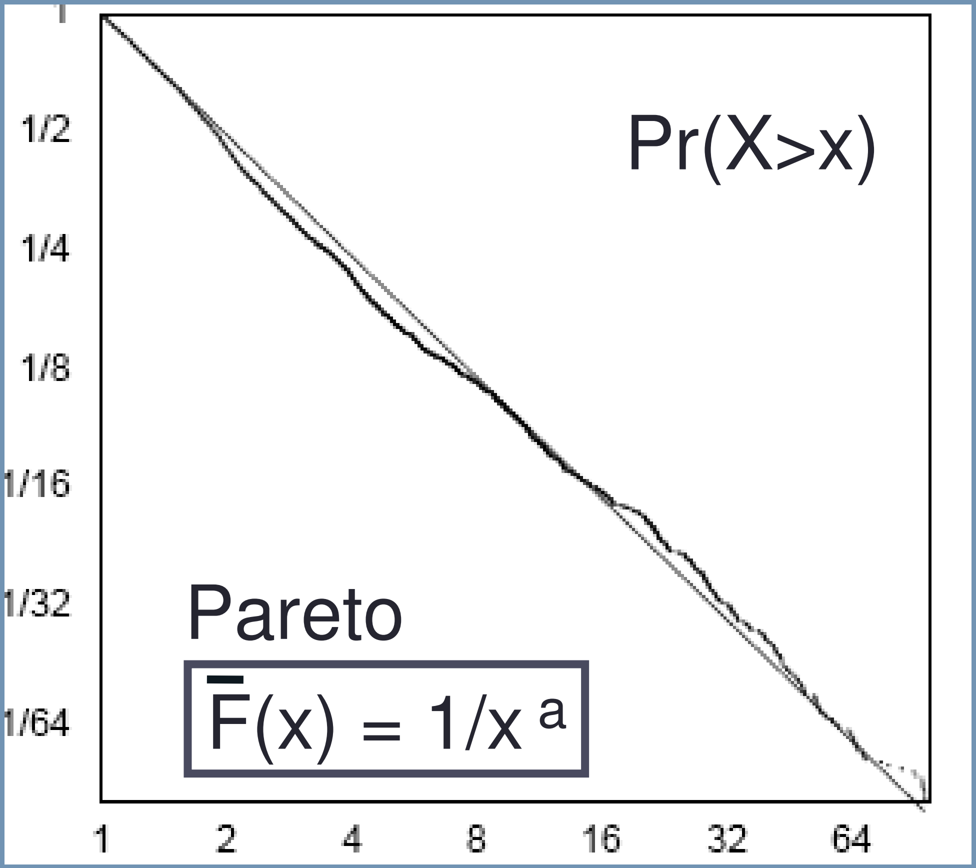

Pareto

$$\boxed{\bar{F}(x) = 1/x^{a}}$$

- a is commonly ~ 1

COMPARISON WITH EXPONENTIAL

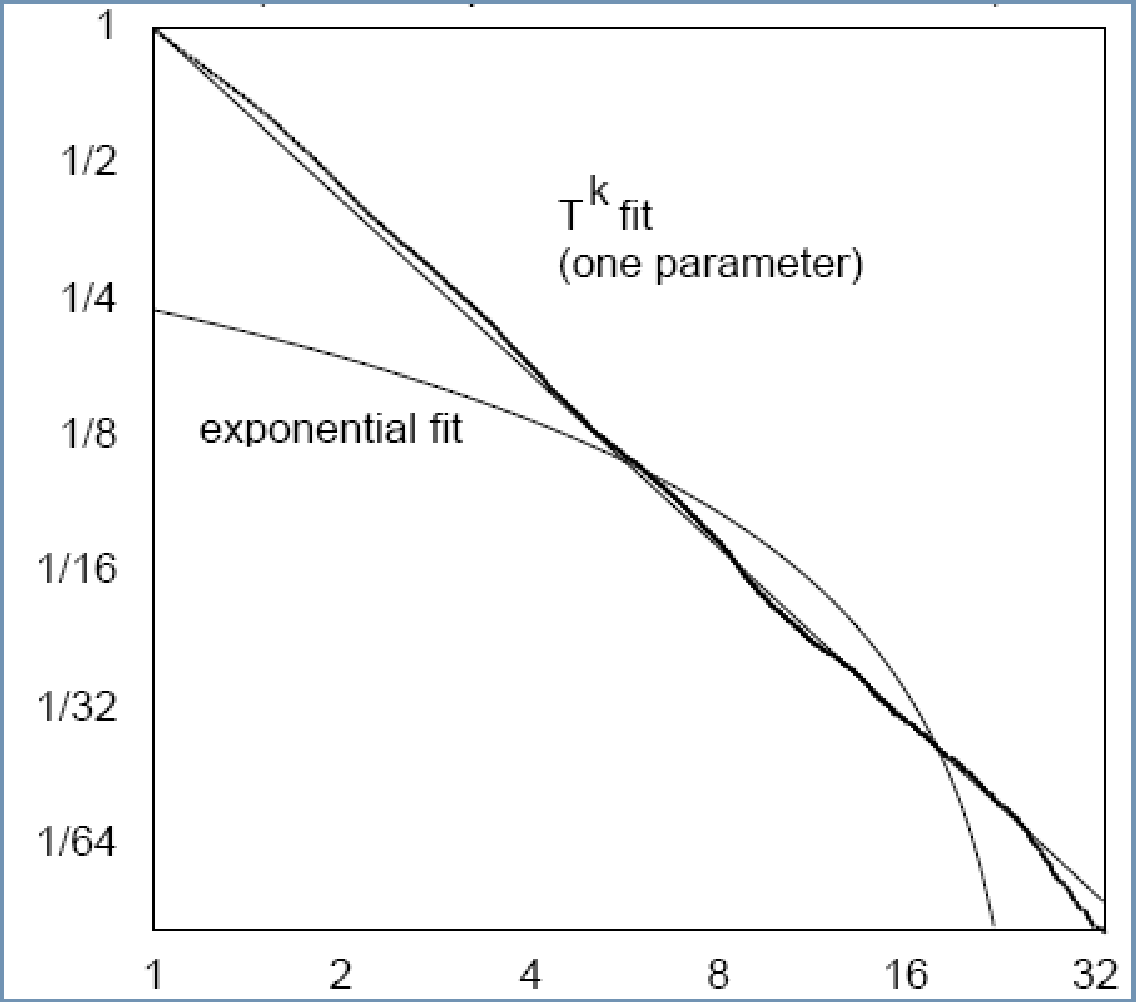

DECREASING HAZARD RATES

- *Example:* Time between failure of nodes in a cluster

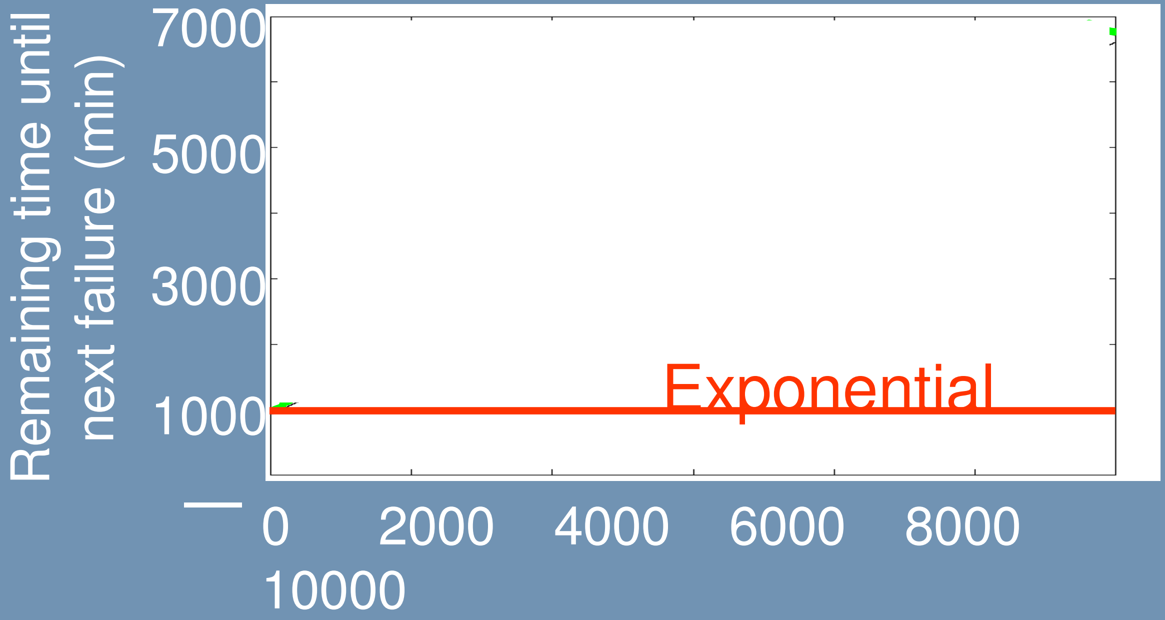

DECREASING HAZARD RATES

- *Example:* Time between failure of nodes in a cluster

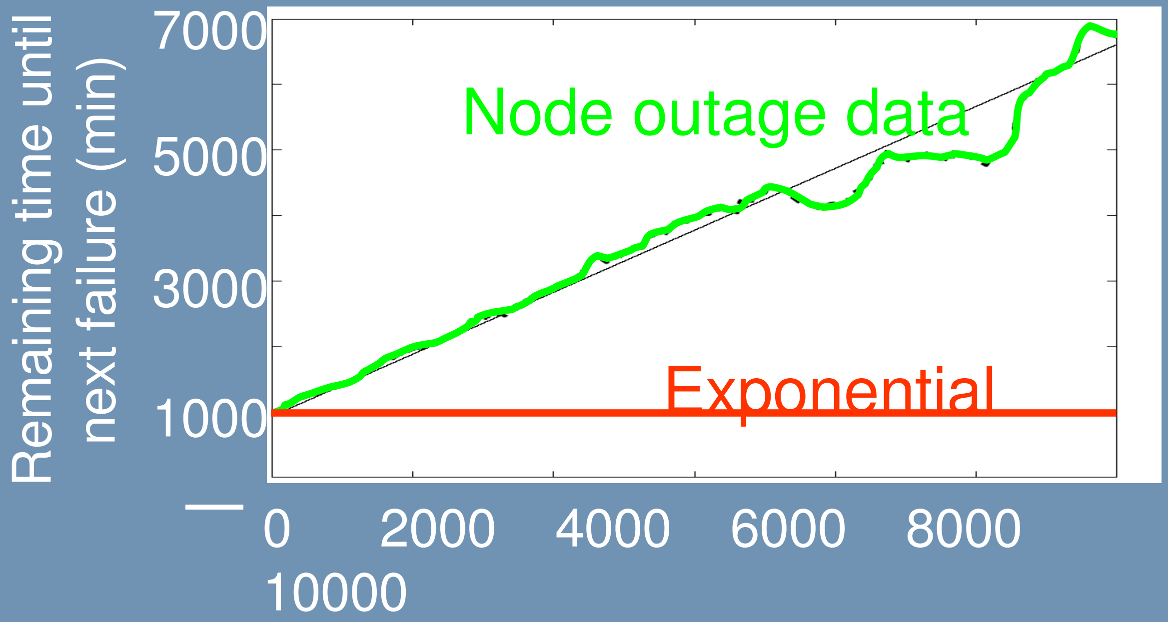

PROPERTY OF PARETO DISTRIBUTIONS

- Unbounded mean and variance (a≤1)
  - E[X] = ∞
  - Var(X) = ∞
  - => In practice, use bounded Pareto
- Decreasing failure rate
- Heavy-tail property

HEAVY-TAIL PROPERTY

- Miniscule fraction of the very largest jobs comprises half the load
  - E.g. if a=1.1 then the top 1% largest jobs make up 50% of the load
  - In comparison: in an exponential the top 1% jobs make up 5% of the load

REAL-WORLD EXAMPLES OF HEAVY-TAIL DISTRIBUTIONS

- Distribution of wealth in the population
- Number of visits to a web site
- Number of pages within a site
- Number of links to a web page
- Length of TCP flows
- File size distributions
- Job runtime distributions