Memoryless property:

- suppose time between session arrivals Z is exponentially distributed

  $\rightarrow$ note: $\Pr\{Z > y\} = \int_y^\infty be^{-bt}dt = e^{-by}$
- suppose a session has not arrived for y seconds
- what is the probability that a session will not arrive for another x seconds?

  $\rightarrow$ i.e., $\Pr\{Z > x + y \,|\, Z > y\}$?

By conditioning:

$$\Pr\{Z > x + y \,|\, Z > y\} = \frac{\Pr\{Z > x + y\}}{\Pr\{Z > y\}}$$

Hence:

$$\Pr\{Z > x + y \,|\, Z > y\} = \frac{e^{-b(x+y)}}{e^{-by}} = e^{-bx}$$

$\longrightarrow$ the past doesn't impact the future!

Of course, not surprising since exponential distribution essentially comes from independent coin tossing

$\longrightarrow$ independence over time is built in

Another interpretation/application:

$\longrightarrow$ view Z as session lifetime

... if a session has lasted for y seconds, can we predict if it will last for another x seconds?

$\longrightarrow$ no: it's just e^{-bx}

$\longrightarrow$ knowing the past doesn't help know the future

$\longrightarrow$ not good for gambling

Important empirical fact: time between session arrivals has been observed to be approximately exponentially distributed

$\longrightarrow$ e.g., TCP sessions, Web (HTTP) requests

$\longrightarrow$ refinements: additional burstiness (why?)

However, session lifetimes are not exponentially distributed!

$\longrightarrow$ tend of have "heavier" tails

$\longrightarrow$ exponential distribution: "light" tail

$\longrightarrow$ where have we seen heavier tails?

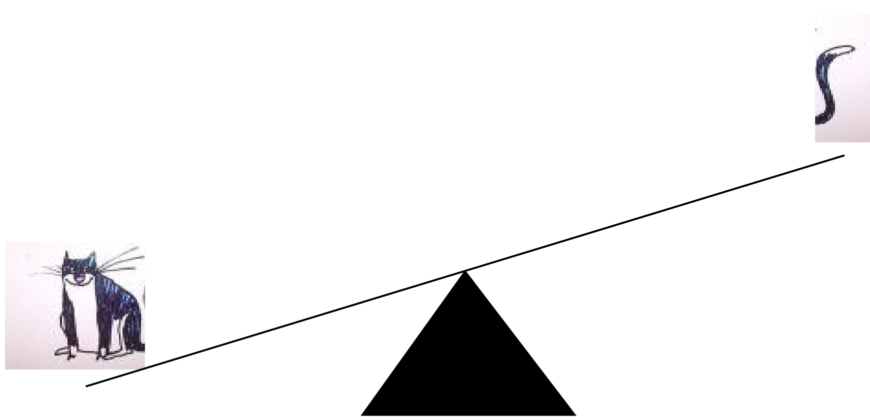

Lastly, let's count:

$\longrightarrow$ with exponential interarrivals, how many arrivals?

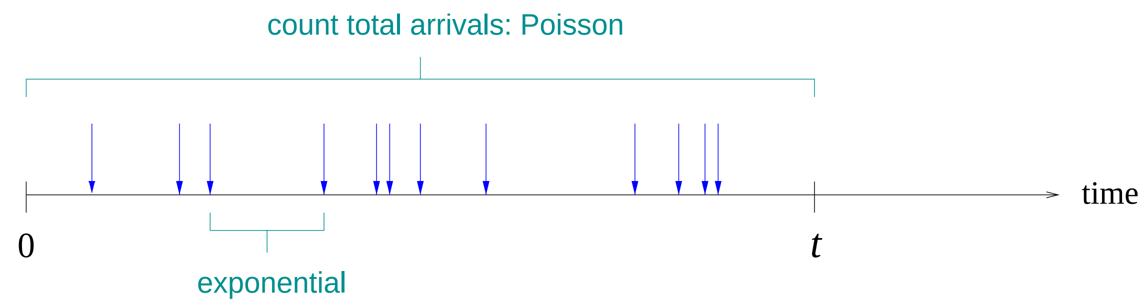

$\longrightarrow$ Poisson distribution

$\longrightarrow$ x arrivals in unit time interval: $e^{-c}c^x/x!$

$\longrightarrow$ $c = 1/b$

$\longrightarrow$ tail: $\Pr\{X > x\} < e^{-c}(ec/x)^x$

$\longrightarrow$ mean of Poisson distribution: c (so $x > c$)

$\longrightarrow$ very light

$\longrightarrow$ large deviations ("outliers") are rare

$\longrightarrow$ reincarnation of what?

Session- and Packet-Level Resource Provisioning

Viewpoint: treating packets individually is ok but ...

$\longrightarrow$ more meaningful: groups of packets

$\longrightarrow$ "packet train"

$\longrightarrow$ e.g., TCP sends window full packets

$\longrightarrow$ e.g., in multimedia frame is relevant unit

Thus:

- packet train as "session" (micro-session)

  → need to be careful about meaning of session

  → session within session within session ...

- one user engages in multiple sessions over time

  → e.g., HTTP client/server request (HTTP runs on top of TCP)

  → persistent vs. non-persistent sessions: HTTP/1.1 vs. 1.0

  → TCP connection set-up/tear-down overhead

Ex.: on/off model

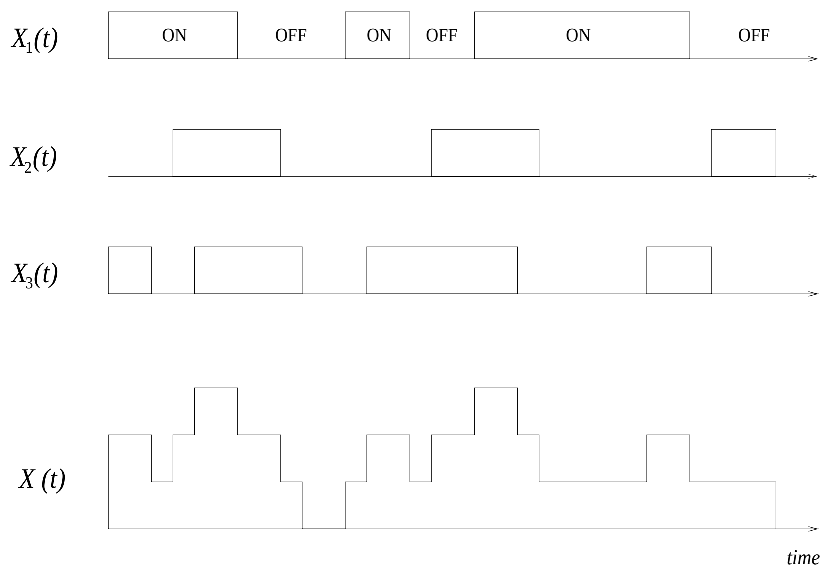

⟶ on-period: TCP file transfer

⟶ on-period length: file transfer completion time

⟶ ignore internal details within on-period: sawtooth

⟶ on-period could be VoIP session: CBR

⟶ not exactly: a user talks only 40% of the time

⟶ approximate view: ok by Amdahl's law

⟶ "don't fret about small things"

We know session arrivals are (approximately) Poisson; what about session lifetimes?

Important fact: TCP session lifetimes are heavy-tailed

$\longrightarrow$ $\Pr\{Z > x\} \approx x^{-\alpha}$

$\longrightarrow$ as opposed to: $\Pr\{Z > x\} \approx e^{-bx}$

$\longrightarrow$ exponent: $1 < \alpha < 2$ (closer to 1)

$\longrightarrow$ note: different from Internet connectivity power-law

$\longrightarrow$ much more likely session will last a long time

$\longrightarrow$ has finite mean but infinite variance

$\longrightarrow$ cat has a very fat tail ("too fat to carry")

Why would TCP session lifetimes be heavy-tailed?

$\longrightarrow$ TCP traffic makes up bulk of Internet traffic

$\longrightarrow$ greater than 80%

Important fact: TCP session lifetimes are heavy-tailed because file sizes are heavy-tailed!

⟶ empirical fact from file server (incl. Web) studies

⟶ after all, TCP mostly transports files

⟶ write simple script to tabulate file sizes on arthur

Ex.: UNIX file system study (Gordon Irlam, 1993)

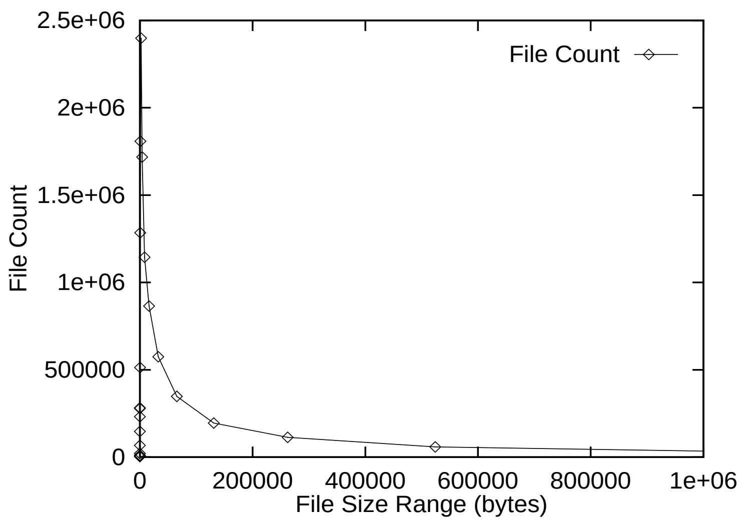

⟶ most are small, but a few are very large

How to check if files sizes are heavy-tailed?

Since $\Pr\{Z > x\} \approx x^{-\alpha}$, take logarithm on both sides:

$\longrightarrow$ $\log \Pr\{Z > x\} \approx -\alpha \log x$

$\longrightarrow$ linear function with negative slope $-\alpha$

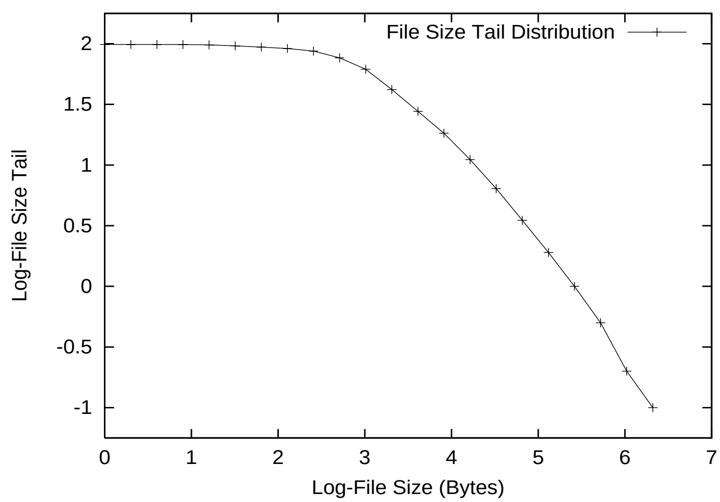

$\longrightarrow$ holds true for large x

$\longrightarrow$ what's the slope α?

$\longrightarrow$ we don't care about details of small sizes (why?)

More details: how "large" is large and "small" is small

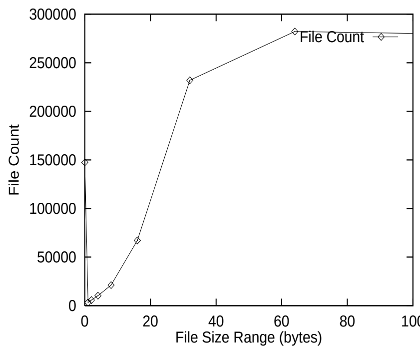

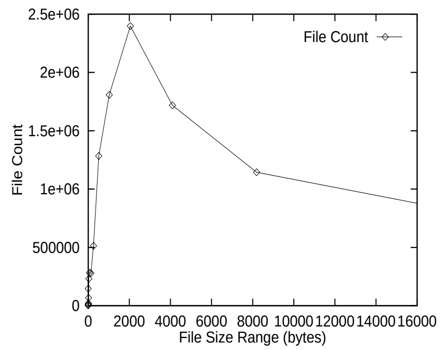

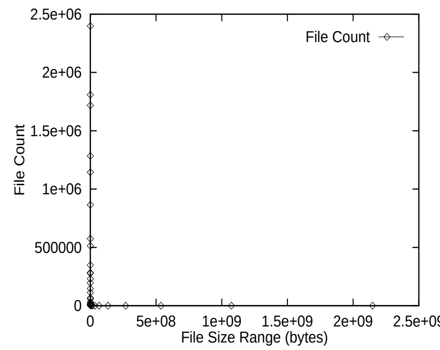

⟶ range: from 0 bytes to ∼2 GB

⟶ 90% of files are smaller than 10 KB

⟶ mean around ∼2 kB

⟶ variation in small size range

⟶ in some file systems: bimodal

Disk space consumed:

⟶ 10% of files consume 90% of space

⟶ same for bandwidth

⟶ "mice and elephants" metaphore

Consequences of heavy-tailed session lifetime:

$\longrightarrow$ for resource provisioning

$\longrightarrow$ as usual: good and bad

First, the good news!

$\longrightarrow$ what might be good?

Ex.: popular heavy-tailed distribution: Pareto

$$\Pr\{Z > x\} = \left(\frac{k}{x}\right)^{\alpha}$$

where

$0 < \alpha < 2$: shape parameter;

$k \leq x$: location parameter

mean: $E[Z] = \alpha k/(\alpha - 1)$

Good news: predictability

$\longrightarrow$ if session has lasted y sec, will last for another x sec

Compute:

$$\Pr\{Z > x + y \,|\, Z > y\} = \frac{\Pr\{Z > x + y\}}{\Pr\{Z > y\}} = \frac{(k/(x+y))^{\alpha}}{(k/y)^{\alpha}} = \left(\frac{y}{y+x}\right)^{\alpha}$$

Knowing the past allows predicting the future:

$$\Pr\{Z > x + y \,|\, Z > y\} \rightarrow 1 \quad \text{as} \quad y \rightarrow \infty$$

- by observing for longer period, can get more certainty
- average expected future duration:

  $\rightarrow E\{Z \,|\, Z > y\} = y\alpha/(\alpha - 1)$
- find a casino with heavy-tailed roulette wheel

  $\rightarrow$ no money worries!