REVIEW: M/G/1 mean queueing time

$$\mathbf{E}[T_Q] = \frac{\rho}{1-\rho} \cdot \frac{\mathbf{E}[S^2]}{2\mathbf{E}[S]}$$

$$= \frac{\rho}{1-\rho} \cdot \frac{\mathbf{E}[S]}{2} \cdot \left(C_S^2 + 1\right)$$

Question: What is the importance of load ρ?

Question: Why does C_S^2 come up? What does this have to do with the Inspection Paradox?

REVIEW: Very high queueing times

Question: If a system has low load, is its mean queueing time low?

Comparison of three server organizations from HW3

Three server organizations. Outside arrivals occur according to a Poisson proces with rate λ. Job sizes denoted by r.v. S. When a job runs on a server of speed 4x, its service time is $S/4$.

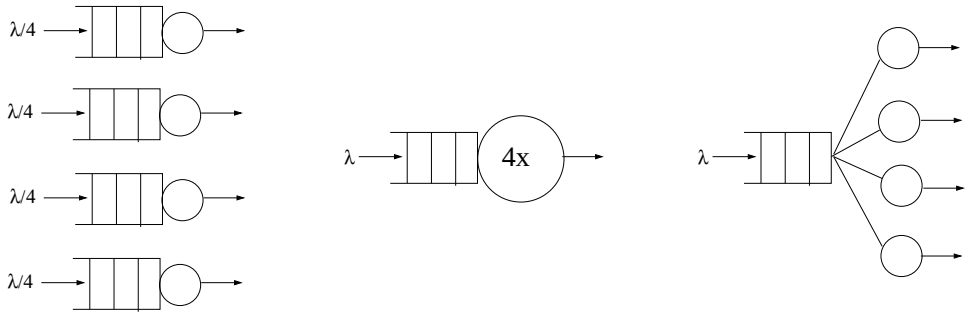

Config A Config B Config C

Question: How do A and B compare?

Question: How is the comparison between B and C affected by load ρ?

Question: How is the comparison between B and C affected by C_S^2?

So what is C_S^2 in practice?

Question: What was C_S^2 for the Exponential distribution?

Question: What was C_S^2 for the Deterministic distribution?

Question: What is C_S^2 for Uniform$(a = 1, b = 1000)$?

Question: Which do CS distributions look like?

– Distribution of job CPU usage
– Distribution of file sizes
– Distribution of IP flow times

In HW 4 you will each look at the distribution of job size in YOUR research.

For now we will hear about my exploration into this question.
Pay attention: You will follow a similar process in HW 4.

The Dark Ages of the 1990's ... my P vs. NP story

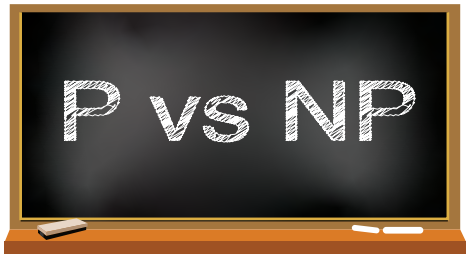

Definitions

A job's **size** refers to its total CPU requirement.
A job's **age** refers to its total CPU usage thus far.
A job's **remaining size** is its remaining CPU requirement.

Question: Which of these jobs likely has higher remaining size?

What we want to understand: $\mathbf{P}\{\text{Size} > x + a \mid \text{Size} > a\}$

Question: What's the answer if Size $\sim$ Exp(μ)?

Failure Rate – informally

$$\mathbf{P}\left\{\text{Size} > x + a \;\mid\; \text{Size} > a\right\}$$

Q: Examples?

Failure Rate more formally

Definition: For continuous X with p.d.f $f_X(t)$ and tail $\overline{F_X}(t)$, the **failure rate function $\mathbf{r_X(t)}$** for X is:

$$r_X(t) \equiv \frac{f_X(t)}{\overline{F_X}(t)}$$

- When $r_X(t)$ is strictly decreasing in t, we say: _____.
- When $r_X(t)$ is strictly increasing in t, we say: _____.

Question: What if $r_X(t)$ is constant?

To try to understand the failure rate, I measured millions of UNIX jobs ...

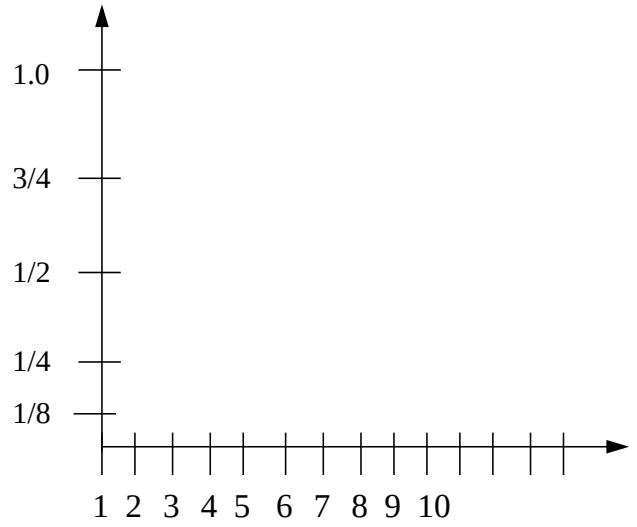

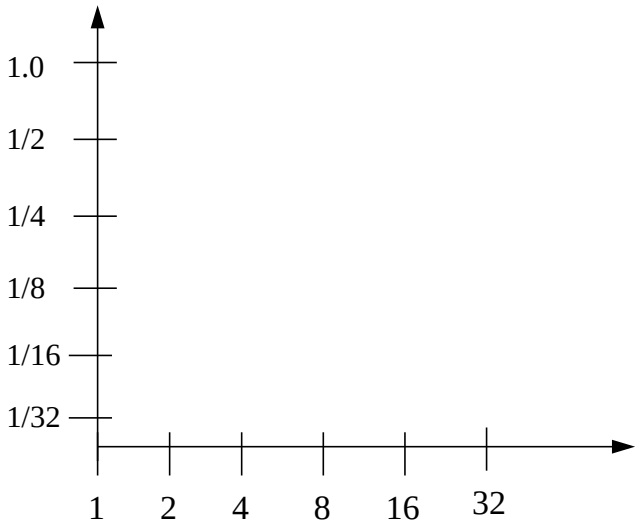

Properties of Power-law Distribution

Suppose that

$$\overline{F}_X(t) = \frac{1}{t}, \qquad t \geq 1$$

1. Is the above a valid distribution?

2. What is $\mathbf{E}[X]$?

3. What is $r_X(t)$, $t \geq 1$? What kind of failure rate is this?

4. Derive $\mathbf{P}\{X > 2a \mid X > a\}$

The Pareto(α) Distribution

Definition:

Let $0 < \alpha < 2$.

$$X \sim \text{Pareto}(\alpha) \qquad \text{if} \qquad \overline{F}(x) = x^{-\alpha}, \quad x \geq 1$$

Question: What kind of Pareto distribution did I have?

3 Properties of Pareto(α) Distribution:

- DFR

- Infinite Variance

- "Heavy-tailed property"

Q: What do the above properties tell us about migration?

Measured distribution is BoundedPareto(k, p, α)

End of the story ...

Today's Distributions – 2022

Question: How have workloads changed between 1996 and today?

See: "Borg: the next generation" by Tirmazi et al.
Proceedings of EUROSYS, 2020.

COMPUTE CONSUMPTION AT GOOGLE 2020:

- Units: _____
- Distribution: _____
- C^2 = _____
- Span in job sizes: _____
- Top 1% jobs consume ____ % of total load

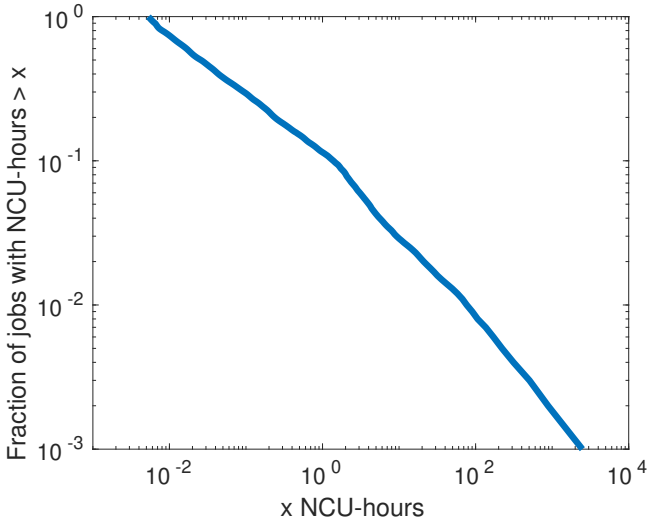

MEMORY CONSUMPTION AT GOOGLE 2020:

- Units:_____
- Distribution: _____
- $C^2 =$ _____
- Span in memory consumption: _____
- Top 1% jobs consume ____ % of total load

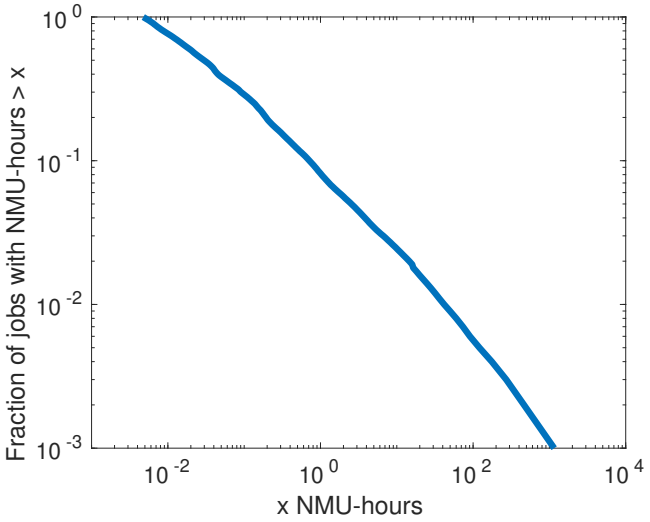

Pareto Distributions are Everywhere

- Compute consumption across jobs
- Memory consumption across jobs
- Web file sizes
- Internet node degrees
- IP flow durations
- Wireless session times
- Phone call durations
- Wealth
- Natural disasters

WHY?

SNEAK PREVIEW: What can we do when C_S^2 is high?

1. Help from multiple servers (recall what we saw in HW 3)

2. Help from scheduling (don't do FCFS)