CSC 2232:
Topics in Computer System Performance and Reliability

Bianca Schroeder

Department of Computer Science
University of Toronto

WHO AM I

WHO ARE YOU?

WHAT IS THIS CLASS ABOUT?

System performance & reliability

- Is one fast server better or two slow servers?
- How many data replicas do I need for reliability?
- How do I generate/simulate realistic workloads?
- What is the impact of workload characteristics?

⇒Methods, tools, back-of-the envelope calculations for system evaluation, simulation & measurement.

⇒Study of recent papers in the area.

FOCUS:

- Large-scale systems
- Storage systems

LOGISTICS

- **Class time: Thursday 1pm-3pm**
  - **Does anybody have serious scheduling conflicts?**
- **Office hours:**
  - By appointment
- **Class web page**
  - **www.cs.toronto.edu/~bianca/csc2232.html**

GRADING

- **30% class participation**
  - Participating in class discussions
  - Class presentation of at least one research paper
  - Possibly 1-2 homeworks
- **70% class project**
  - I will suggest possible projects
  - You can propose your own
  - Final results: Workshop quality paper
- **No exams!**

PAPER PRESENTATION

- **What I do not not want:**
  - A long laundry list of all the things the paper did
- **What I do want:**
  - A lecture style presentation of the paper (30 min)
    - Including background material your fellow classmates might need to understand it
  - A critical discussion of the paper
    - Strength & weaknesses
- **What you get:**
  - Feedback!

CLASS PROJECT

- **Must be done in a team of two**
  - Start looking for a partner now!
- **On a research project you pick**
  - Find your own or see suggestions on web page
- **Output: Workshop quality paper (10-12 pages)**
- **I will help you get there --- 5 milestones:**
  - Project proposal
  - Related work
  - Status report I
  - Status report II
  - Final report

CLASS PROJECT

- **The topic of the project must be related to the topic of the class**
- That means it should contain at least one of the following:
  - A significant modeling component
  - A significant measurement component
  - A significant data analysis component
  - A significant system evaluation component

OUTLINE

- **Administrivia** ✓
- **Motivating examples**
- **Questionnaire**

MOTIVATIONAL EXAMPLES / CASE STUDY

- **System design is often a *counter-intuitive process***
- **Do not expect to understand *everything* in the examples**
- **Don’t worry if you’re not familiar with all terminology**
- **Ask questions!**

EXAMPLE 1

- Question: Assume the arrival rate doubles. By how much do you have to increase CPU speed to keep mean response times the same?
- Does the answer change with PS scheduling instead of FCFS?

EXAMPLE 2

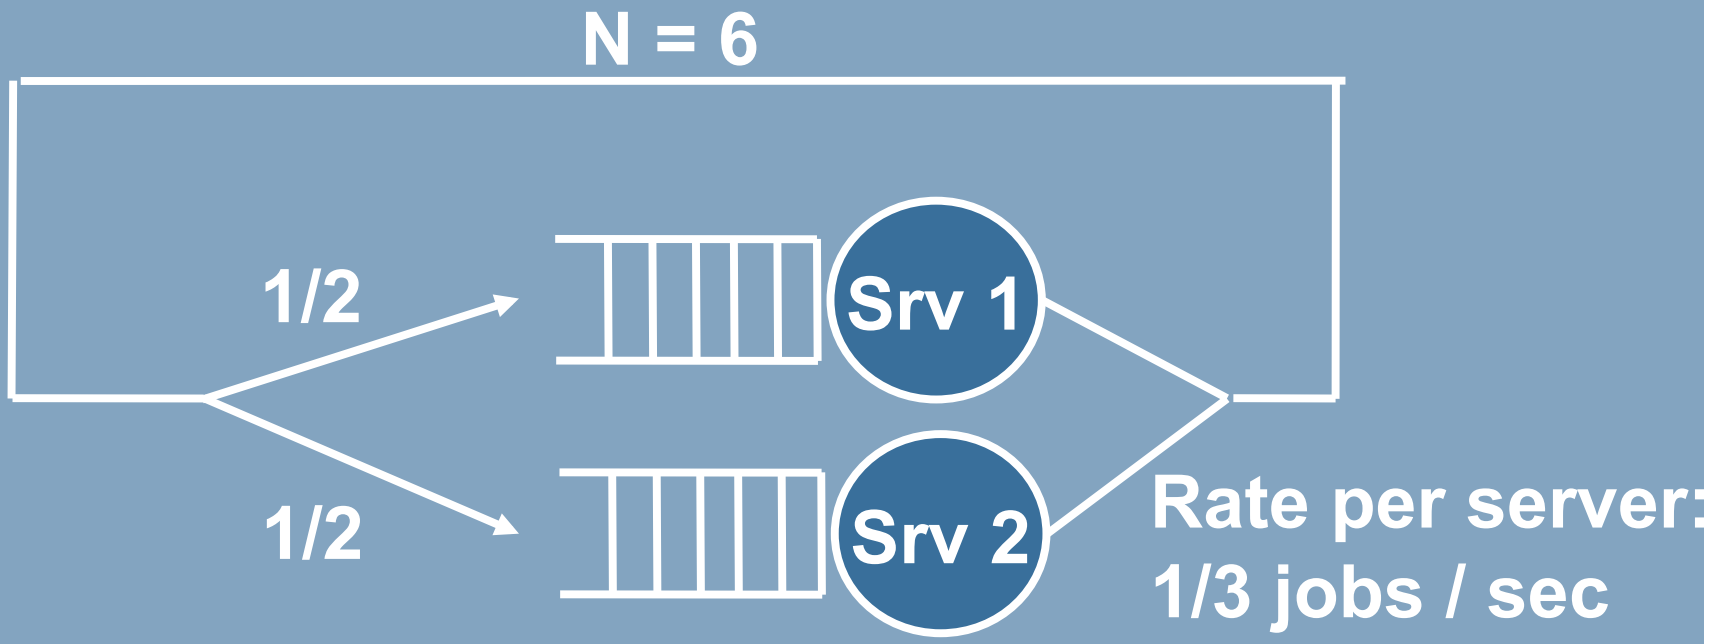

- Question: Server 1 is replaced with a server twice as fast.
- Does this change affect mean response time?
- Does this change affect throughput?

EXAMPLE 3

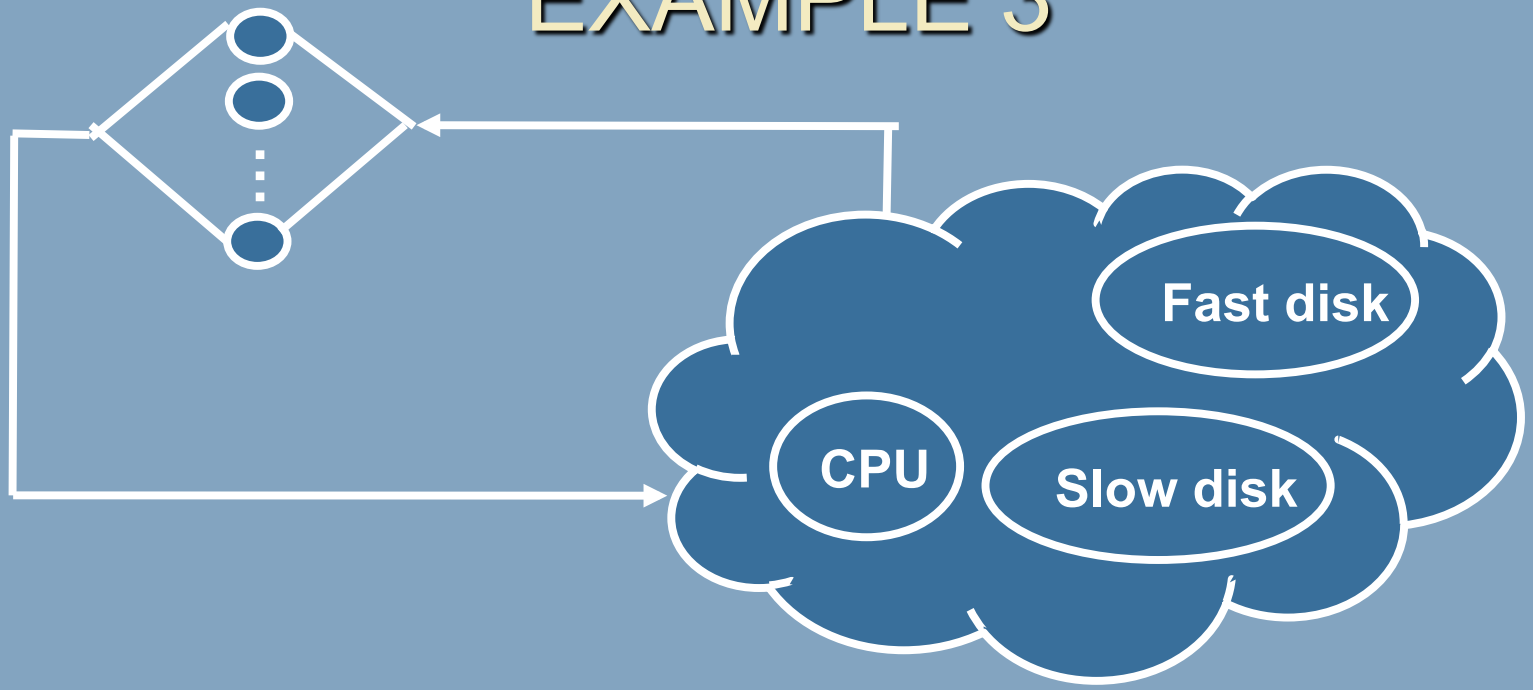

- Question: Which of the following changes is most effective in increasing throughput?
  - Replace CPU with one twice as fast?
  - Balance load between fast & slow disk?
  - Buy second fast disk?

EXAMPLE 3

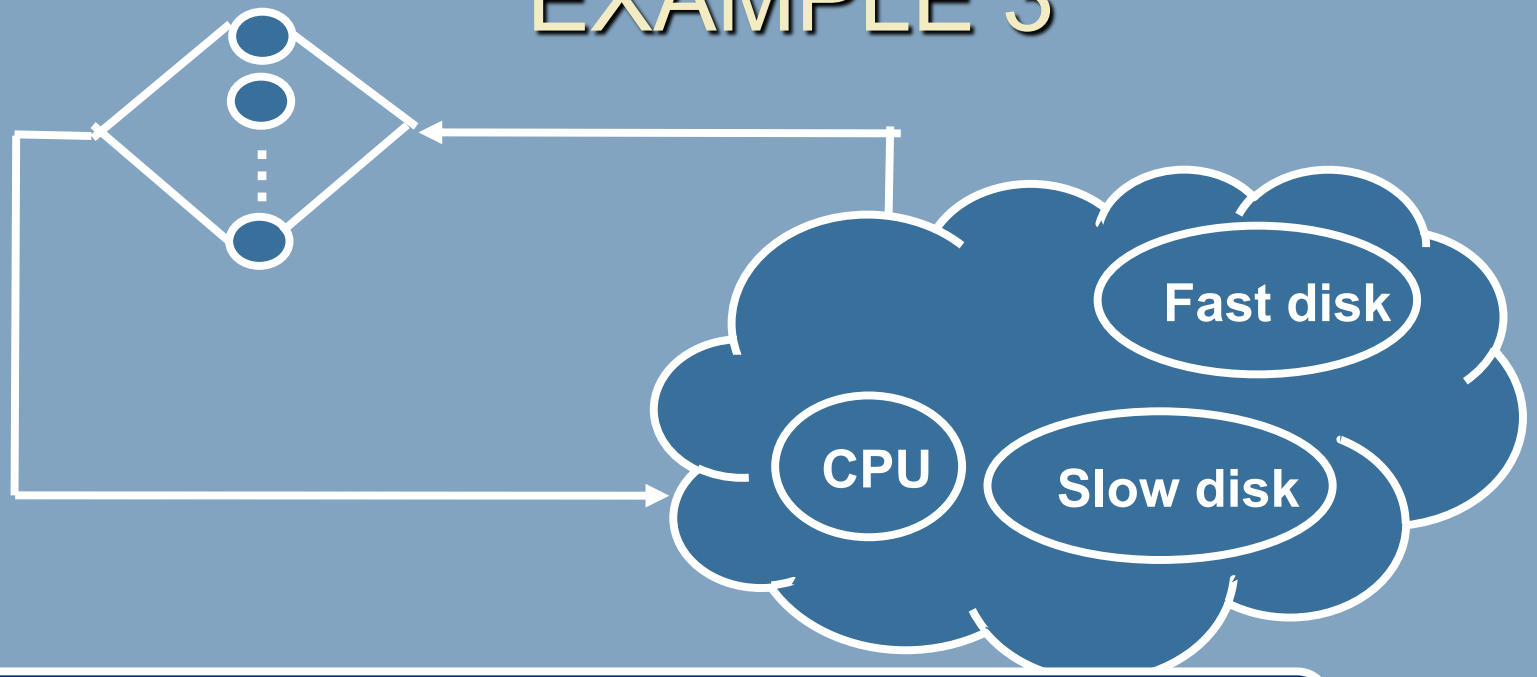

s is most

Measurements:

- **time each device is busy (utilization)**
- **number of completions per device**
- **total number completions**

?

EXAMPLE 3

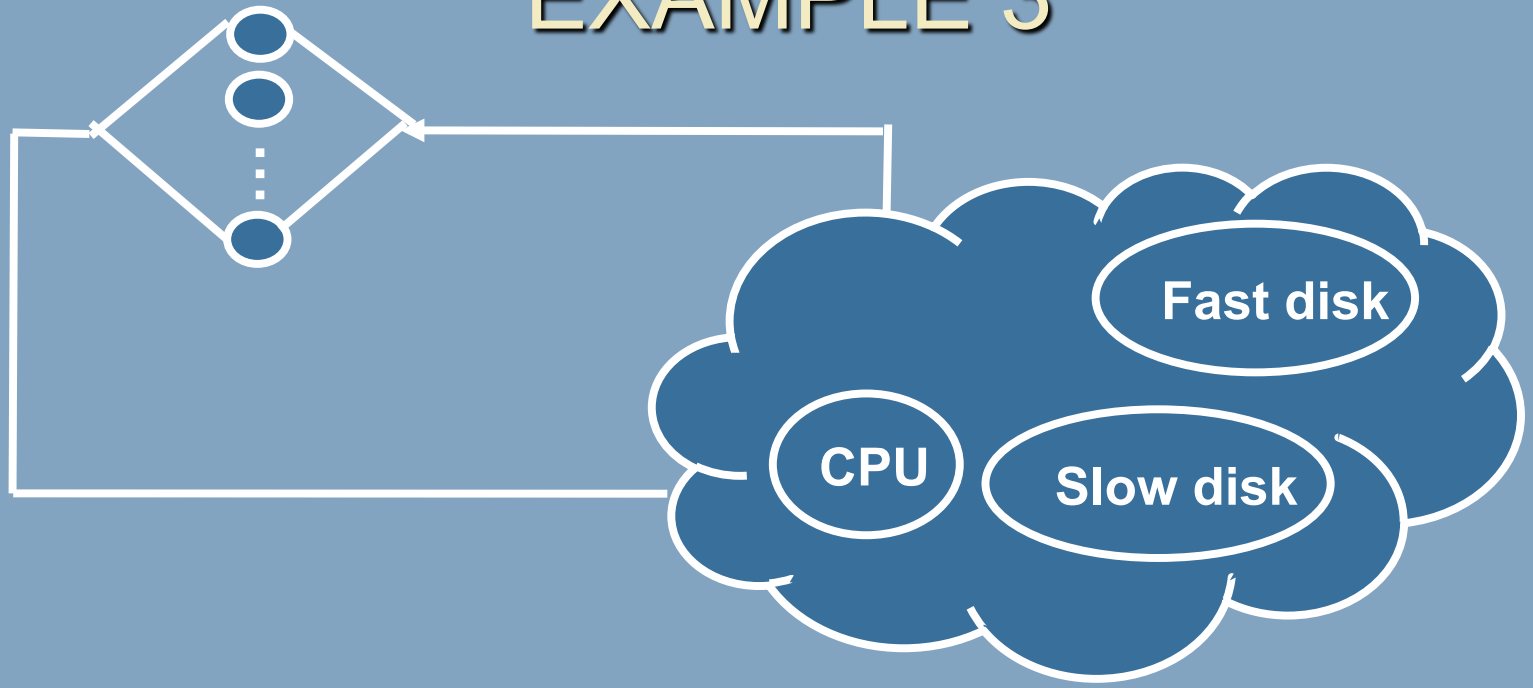

- We can solve this problem with simple back-of-the-envelope bound analysis
  - No math is necessary!
  - No assumptions on distributions
  - No knowledge of full network topology necessary

EXAMPLE 4

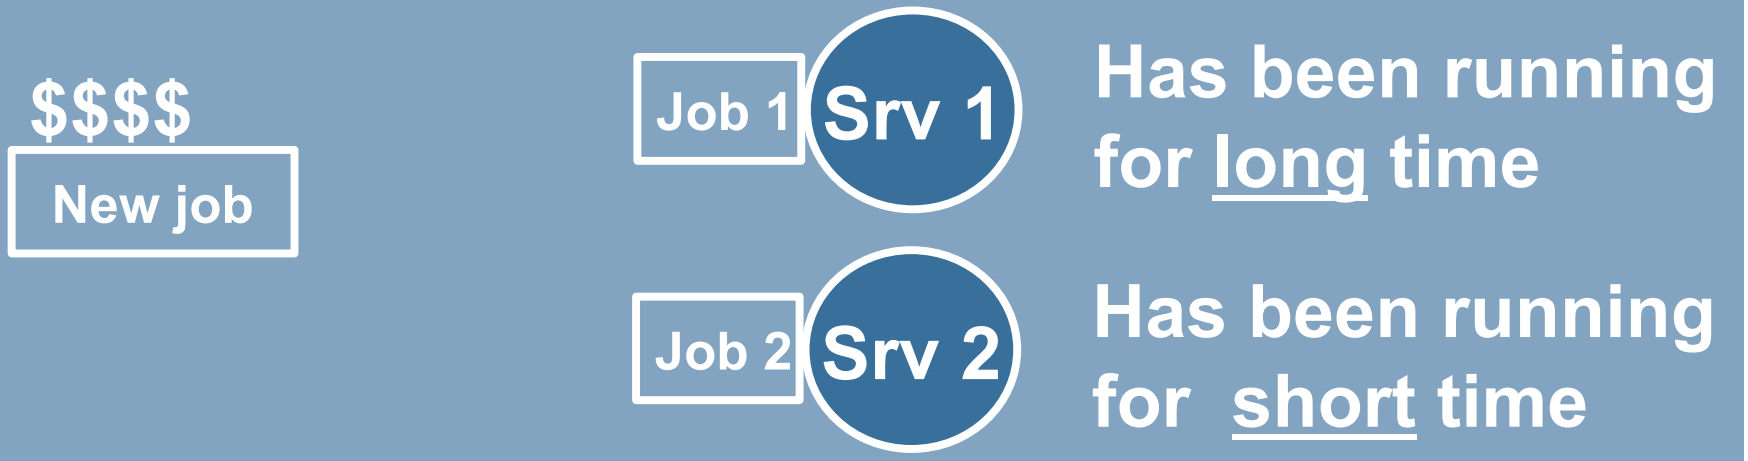

- **Question**: Which job is closer to completion, job 1 or job 2?

EXAMPLE 5

Has been running non-stop for 1 yea

Has just crashed and been fixed yesterday

- **Question**: Which server will fail first?

EXAMPLE 6

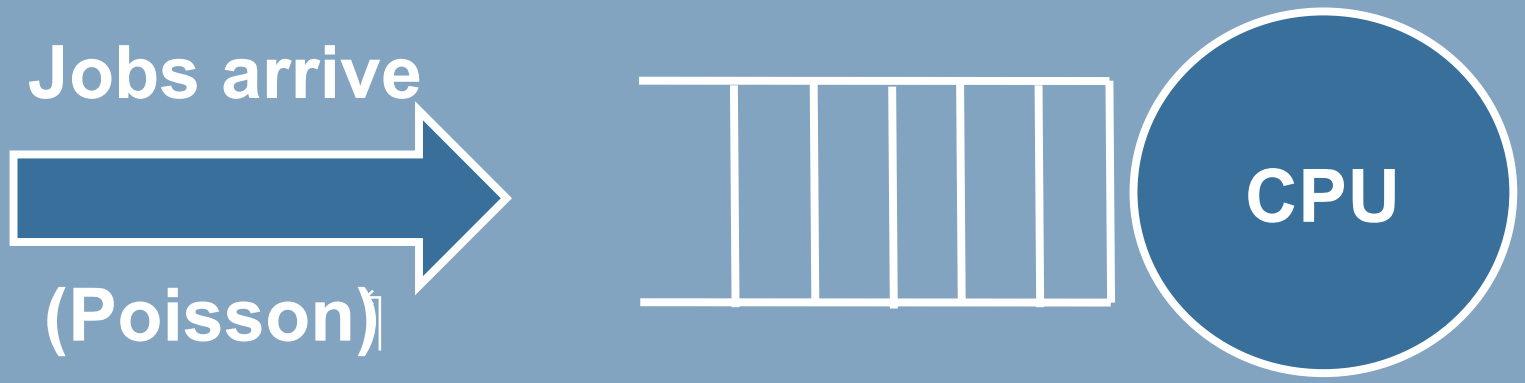

- Question: Which non-preemptive service order will result in lowest mean response time:
  - FCFS
  - LCFS = Last-Come-First-Serve
  - Random