CSC 2232: Topics in Computer System Performance and Reliability

Bianca Schroeder

Department of Computer Science
University of Toronto

SHORT ANNOUNCEMENT

- Lecture slides from last week are posted on course web page.
- Project suggestions & deadlines are posted on web page
- Reading list is posted.
  - Volunteer now! :-)
  - Need one presenter for next week …

CLASS PROJECT: MILESTONES

- Sept. 30: Project proposal
- Oct 7: Related work
- Oct 28: Status report I
- Nov 20: Status report II
- Dec 20: Final report

AGENDA:

- Two case studies from my own research
- Some project suggestions
- A few words about paper presentations

- Probably next week:
  - Queueing Terminology
  - First operational laws
  - Little's law

SOME EXAMPLES FROM MY OWN RESEARCH

SCHEDULING STATIC WEB REQUESTS

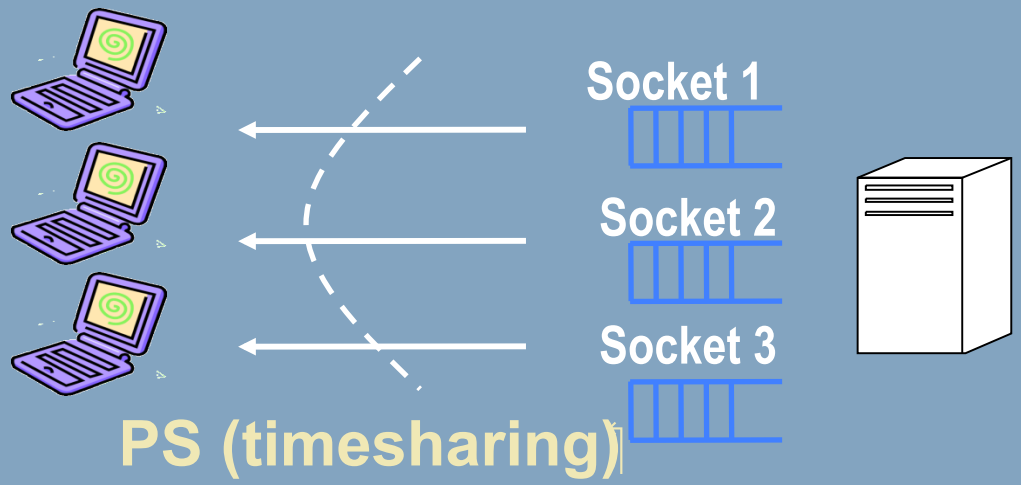

PS (timesharing)

Standard web server

Size-based scheduling for better response times.

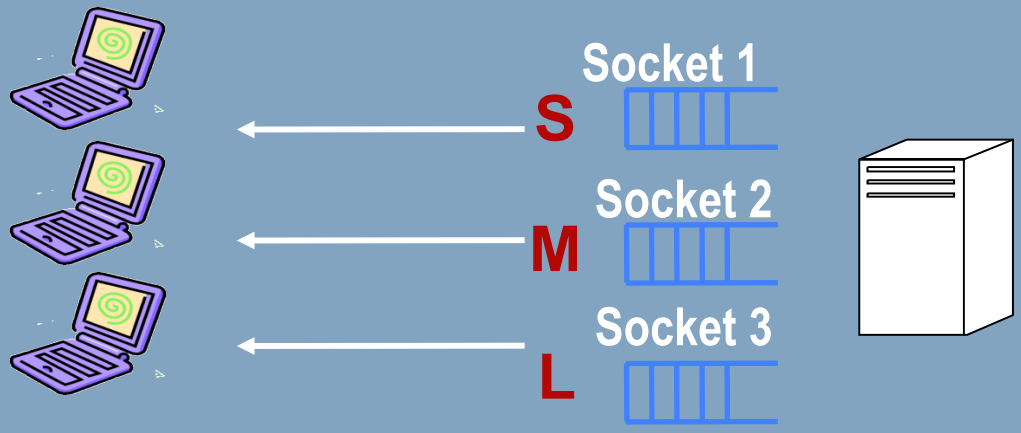

SRPT (shortest-remaining-time)

SRPT web server (kernel-level Implementation)

STATIC WEB WORKLOA
APACHE/LINUX

WHY?

- ❑ Mean file size
- ❑ File size distribution
- ❑ Access pattern
- ❑ Request rate
- ❑ CPU utilization
- ❑ Bandwidth
- ❑ Network effects

ALL THE SAME!

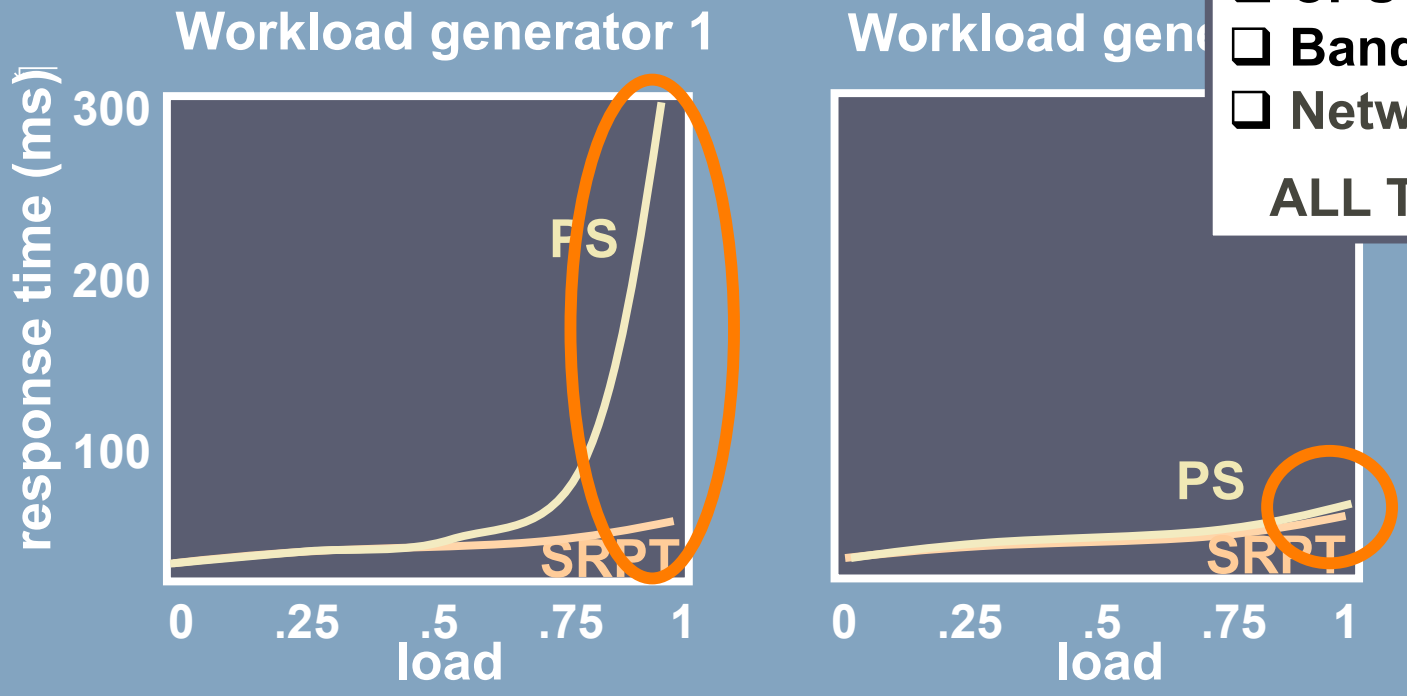

Tuning knob
- Request rate

Tuning knob
-Number of users
-Think time

DATABASE BACKEND OF E-COMMERCE SITE

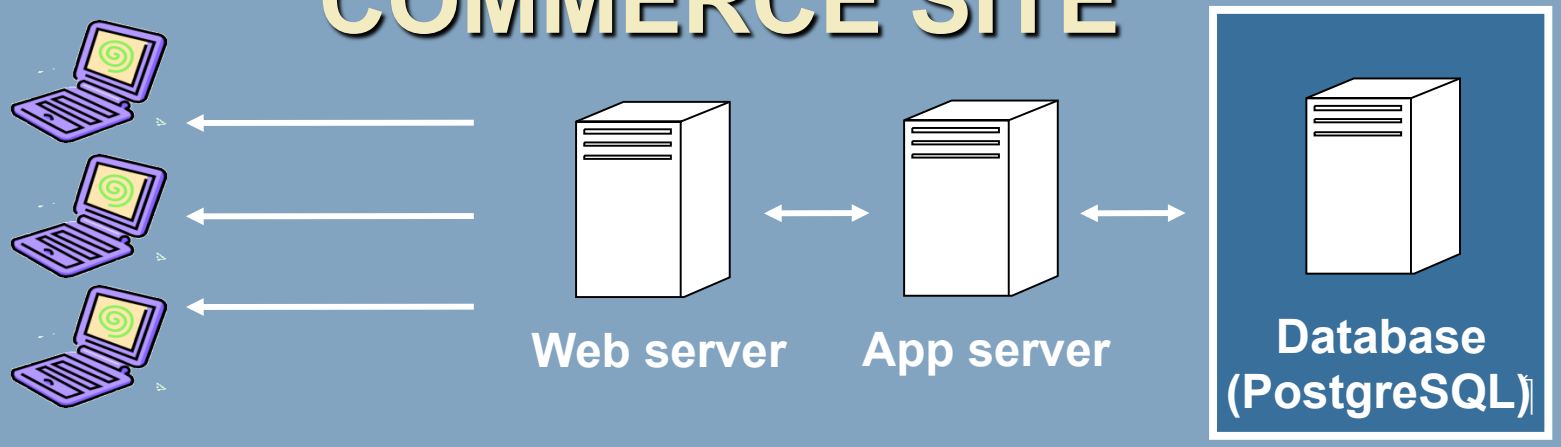

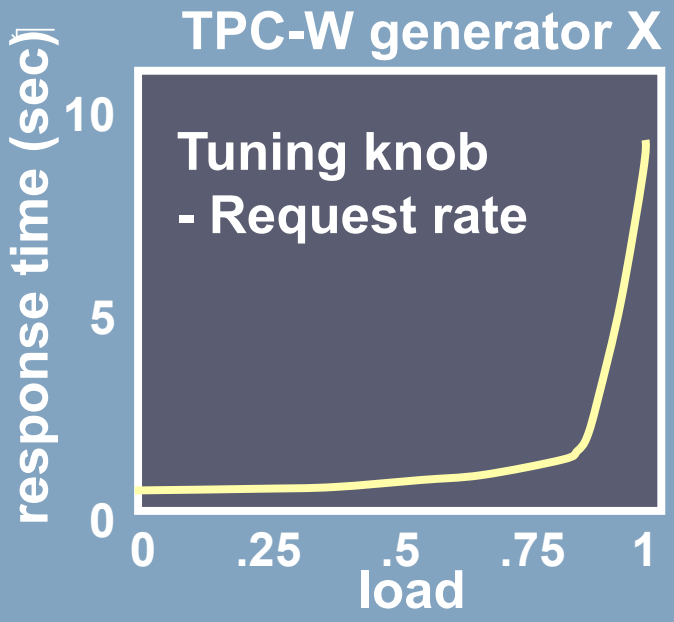

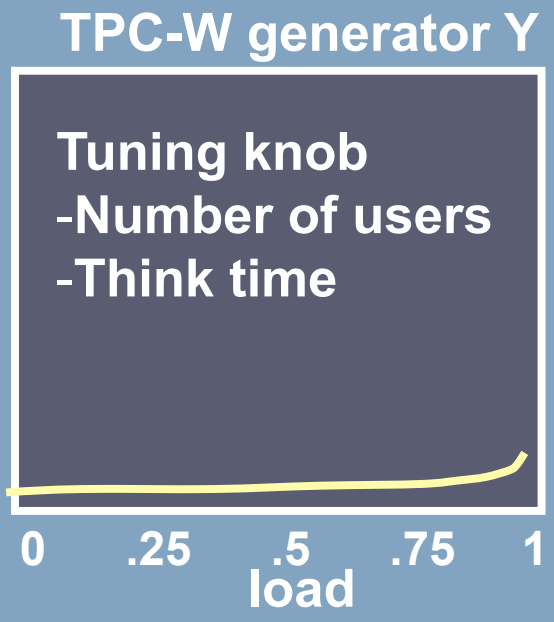

ONLINE AUCTION SITE – SIMULATION

- Based on trace from top-10 online auctioning site.

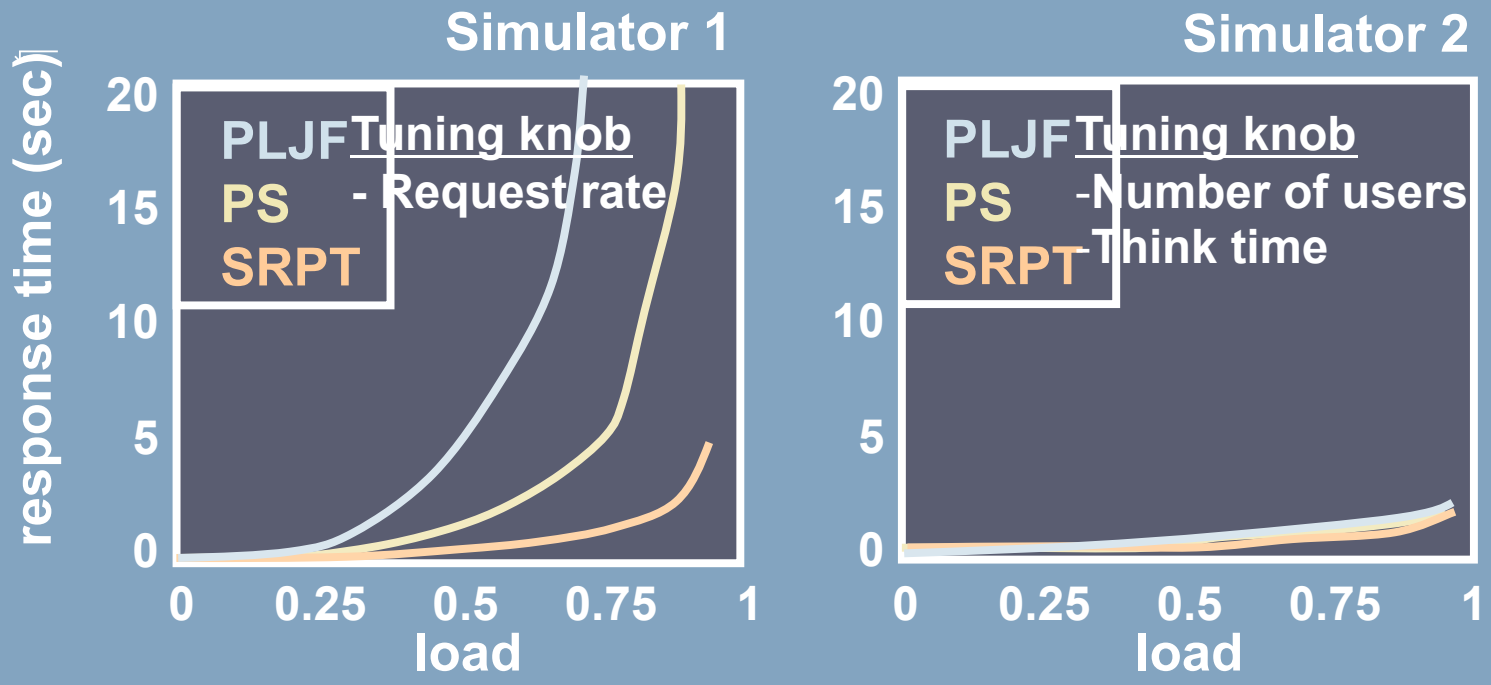

BATCH JOBS AT A SUPERCOMPUTING SITE

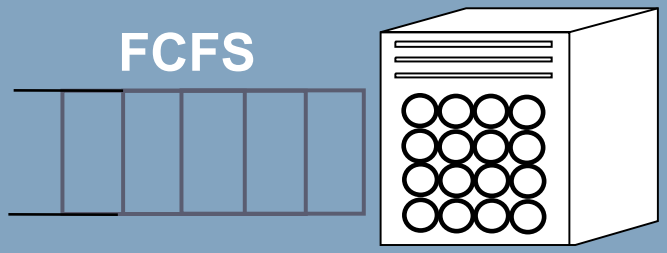

Cray J90/C90

- Simulation based on trace from Pittsburgh Supercomputing Center.

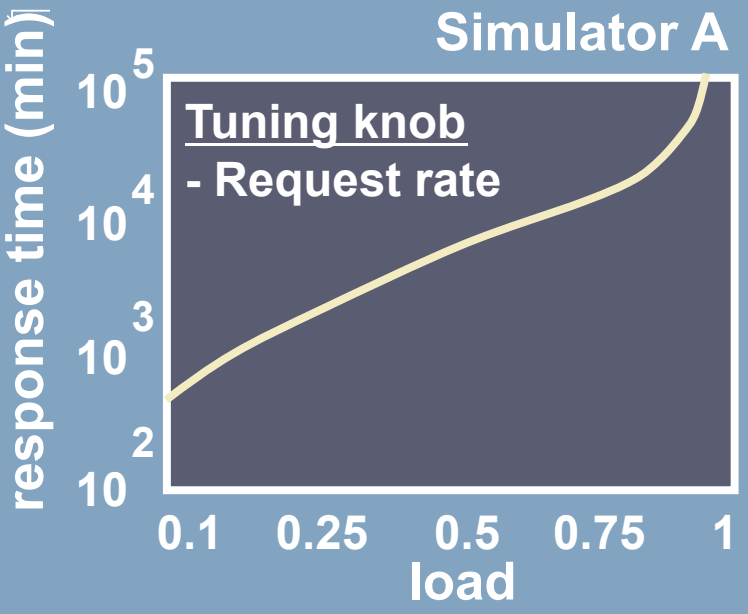

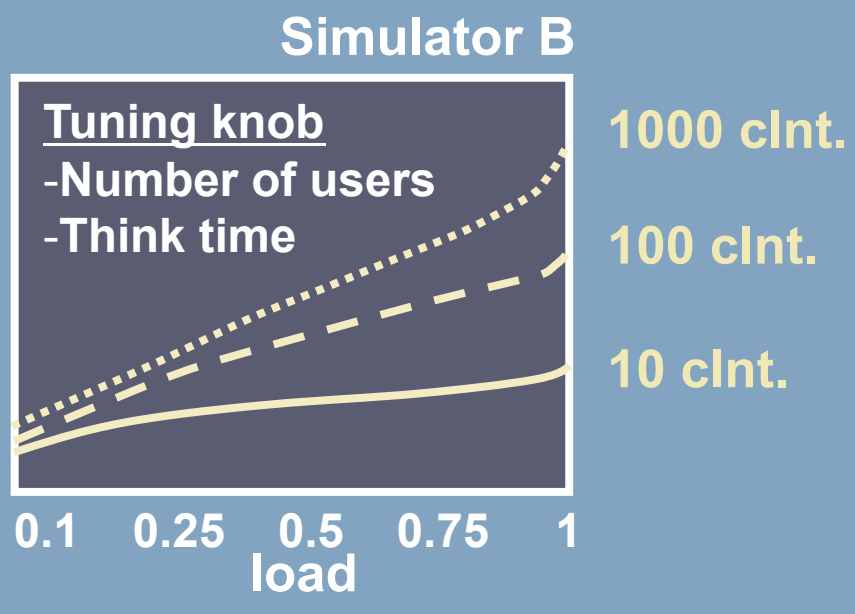

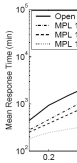

STATIC WEB

DATABASE BACKEND

ONLINE AUCTION

SUPER-COMPUTING

OPEN

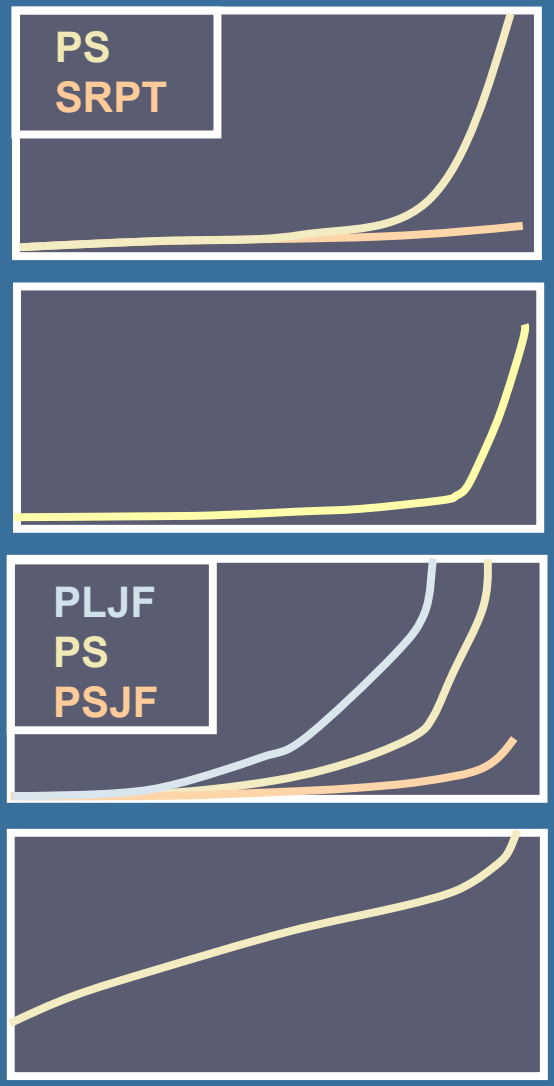

Tuning knob
- Request rate

CLOSED

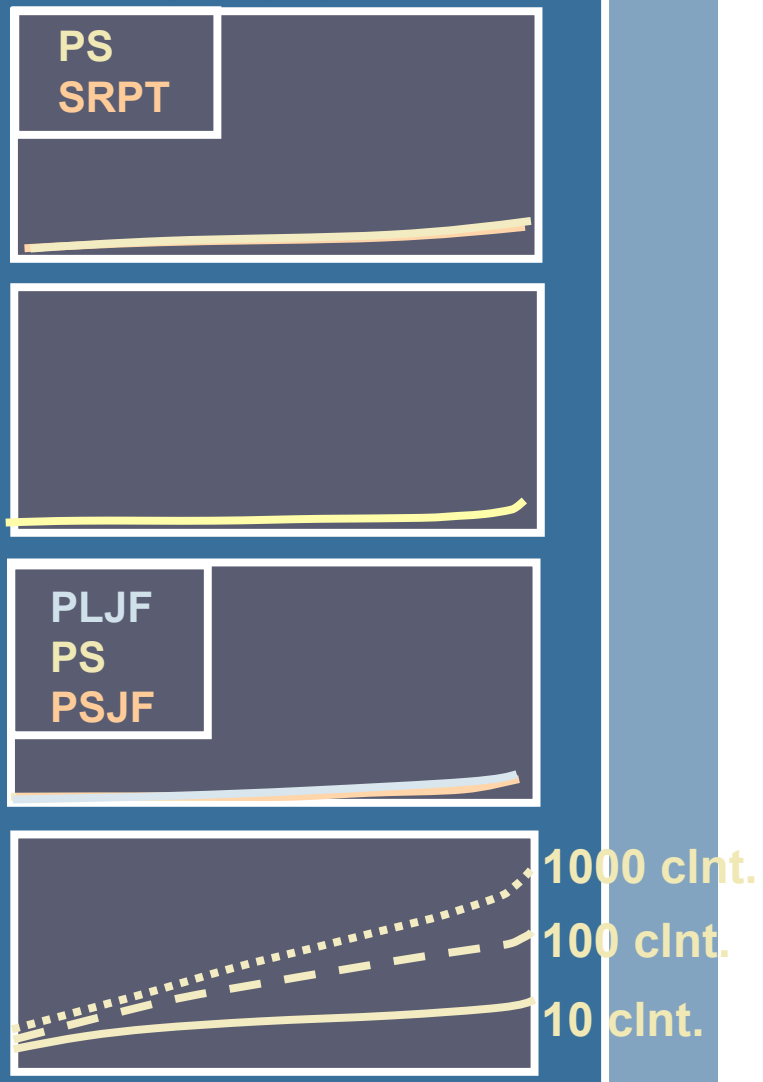

Tuning knob
-Number of users
-Think time

CLOSED SYSTEM MODEL

Model of user behavior

→ **User requests web page, receives page, reads page, clicks on new link**

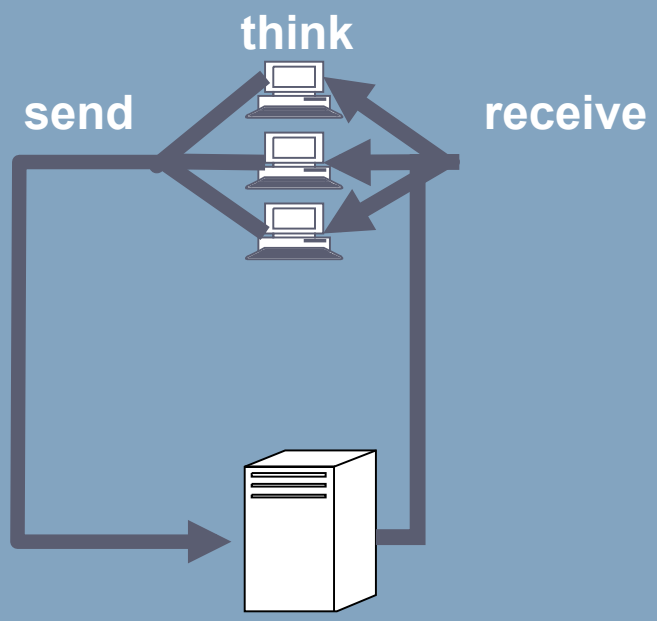

- **Fixed number of users, called the Multi-Programming-Level (MPL)**
- **Arrivals triggered by completions.**

OPEN SYSTEM MODEL

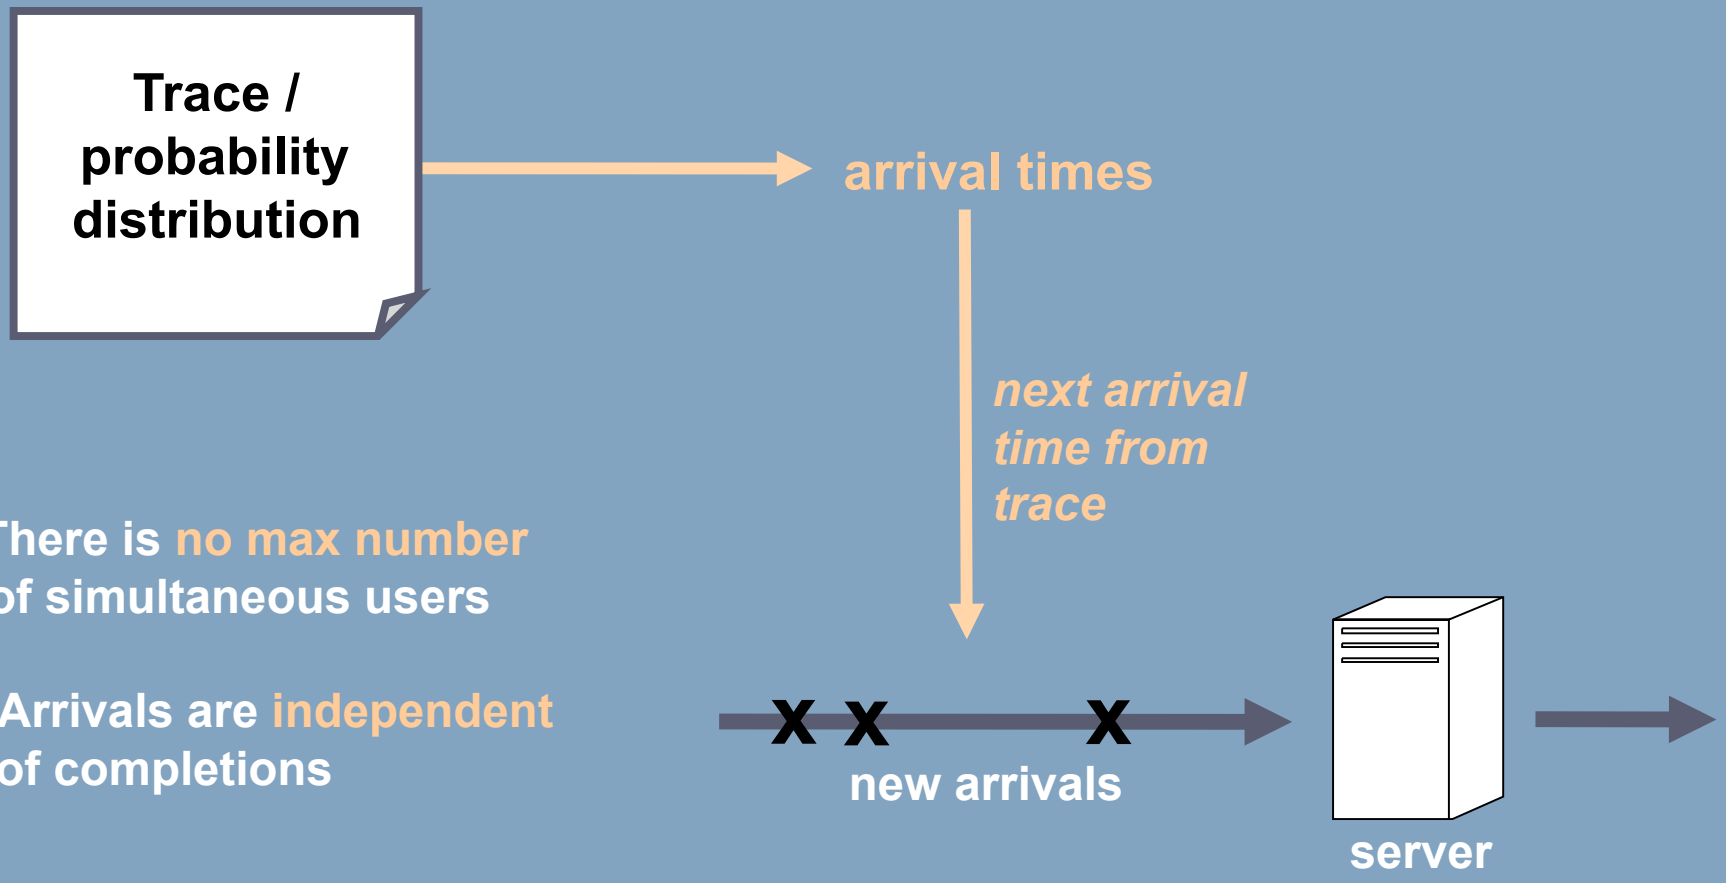

- There is no max number of simultaneous users
- Arrivals are independent of completions

WHICH MODEL DO WORKLOAD GENERATORS USE?

- G
  p
  n
- C
  n e!

WEB WORKLOAD GENERATORS

Surge
SPECWeb
TPC-W
Sclient
RUBiS
WebBench
Webjamma

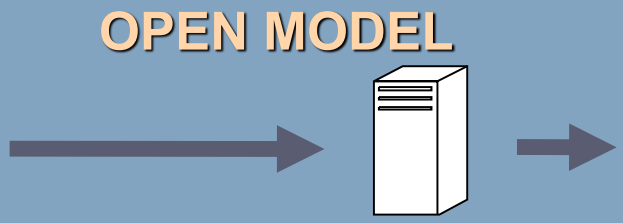

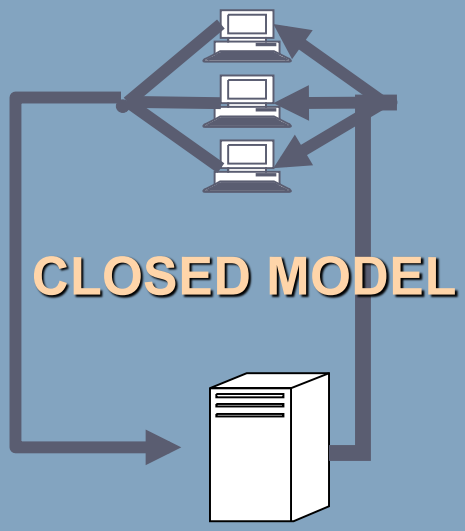

WHAT IS KNOWN IN THE LITERATURE?

- **Very little …**
- **Limited to FCFS single server queue.**
  - **Response times under open system higher than under closed [Bondi and Whitt 1986].**
  - **For MPL ->∞, closed system converges to open system [Schatte83, Schatte84].**

STILL UNANSWERED:

- What is the magnitude in difference of response times?
- What is the speed of convergence?
- How does variability (heavy tails) affect results?
- How are different scheduling disciplines affected?
- …. in practice?

PRINCIPLES FOR OPEN VS. CLOSED

- **What is the magnitude in difference of response times?**
  - **Orders of magnitude!**

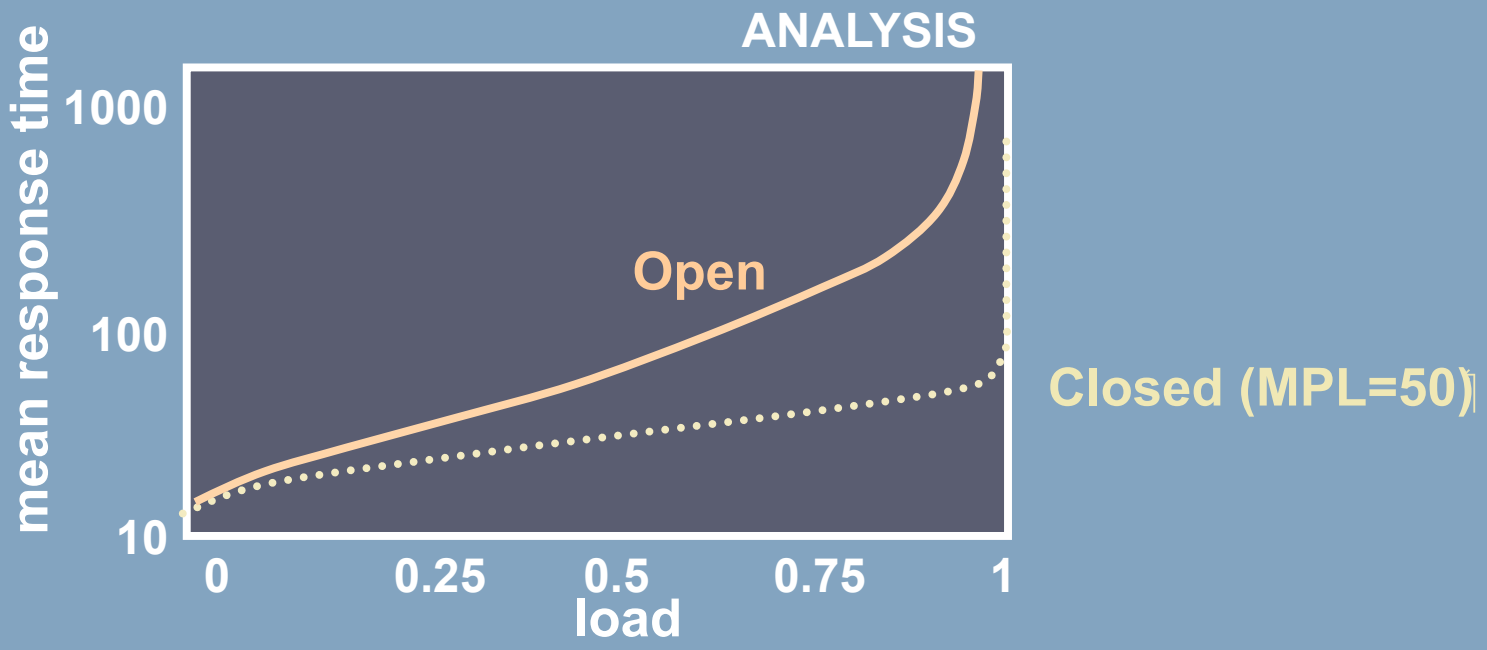

- **Why?**
  - ***Bounded number of jobs in closed system.***

PRINCIPLES FOR OPEN VS. CLOSED

- **How does variability affect open/closed response times?**
  - **Huge effect on open, limited effect on closed system.**

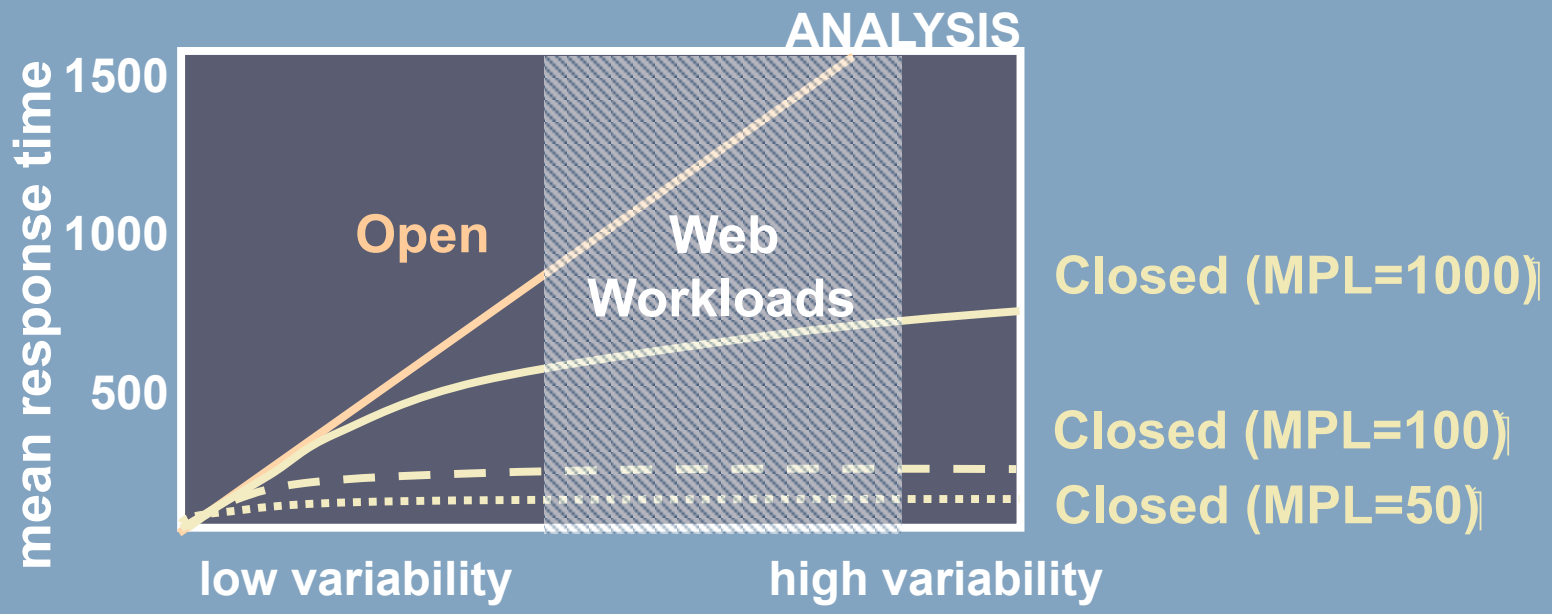

- **Why?**
  - ***Dependency between completions and arrivals in closed system reduces burstiness.***

PRINCIPLES FOR OPEN VS. CLOSED

- **Can we make closed look like open, by increasing MPL?**

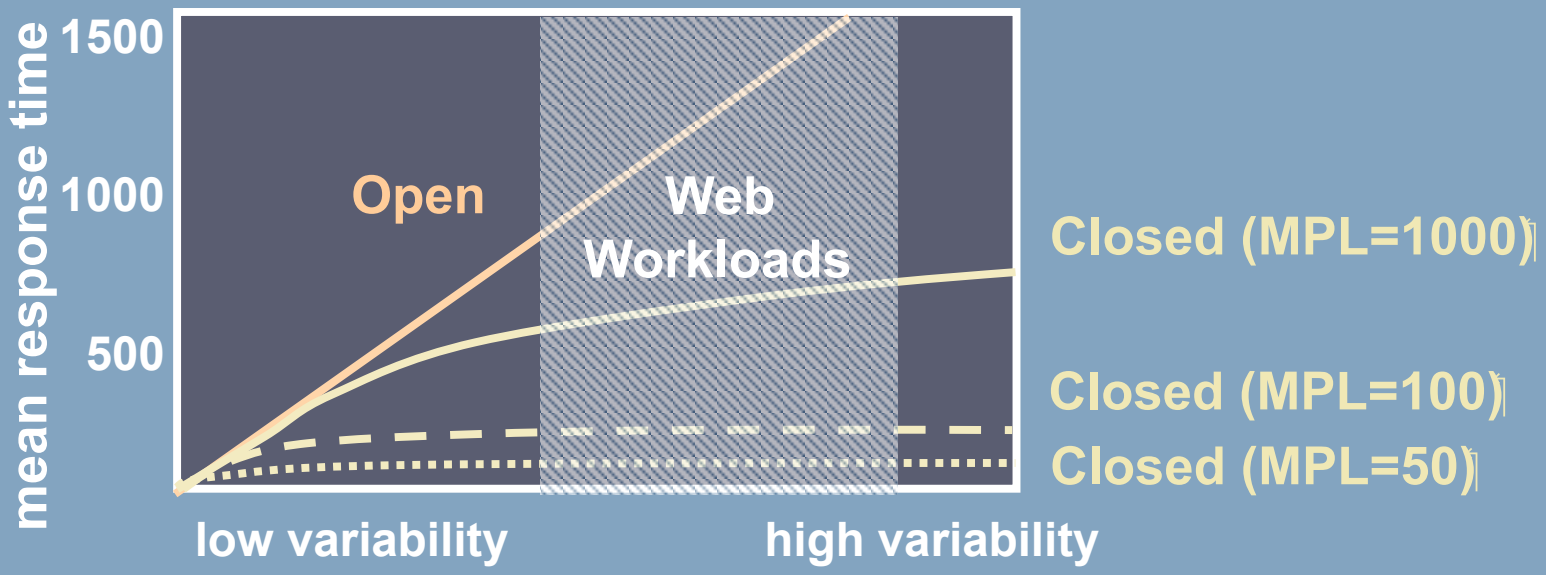

PRINCIPLES FOR OPEN VS. CLOSED

- **What is the impact of scheduling?**
  - **Huge in open system, almost none in closed system.**

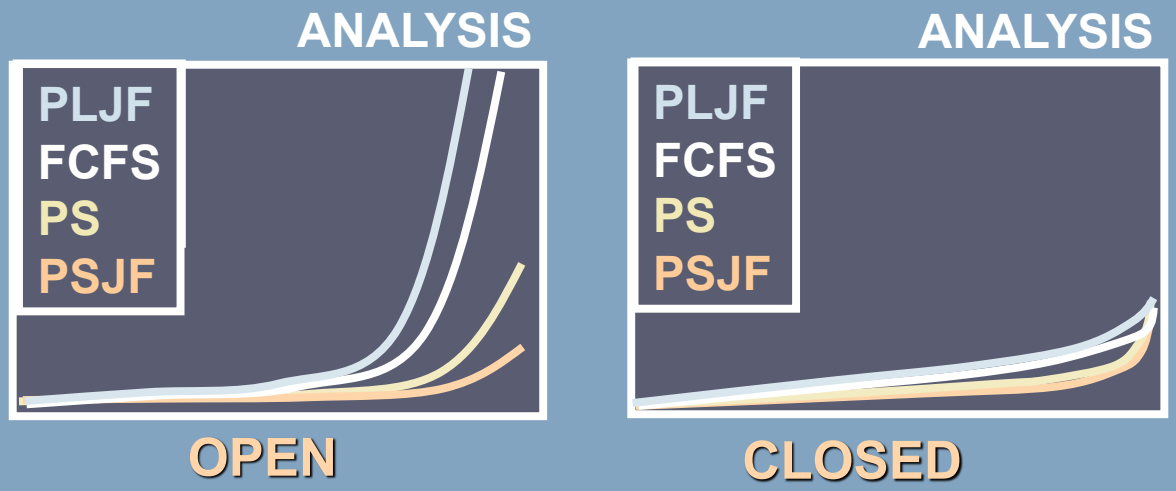

- **Why?**
  - ***Scheduling takes advantage of variability in the system.***
  - ***Closed systems reduce the effect of variability.***

HOW DO YOU CHOOSE THE RIGHT MODEL?

1. **Is there a more realistic model?**
2. **What's most representative of real systems?**

THE PARTLY-OPEN MODEL

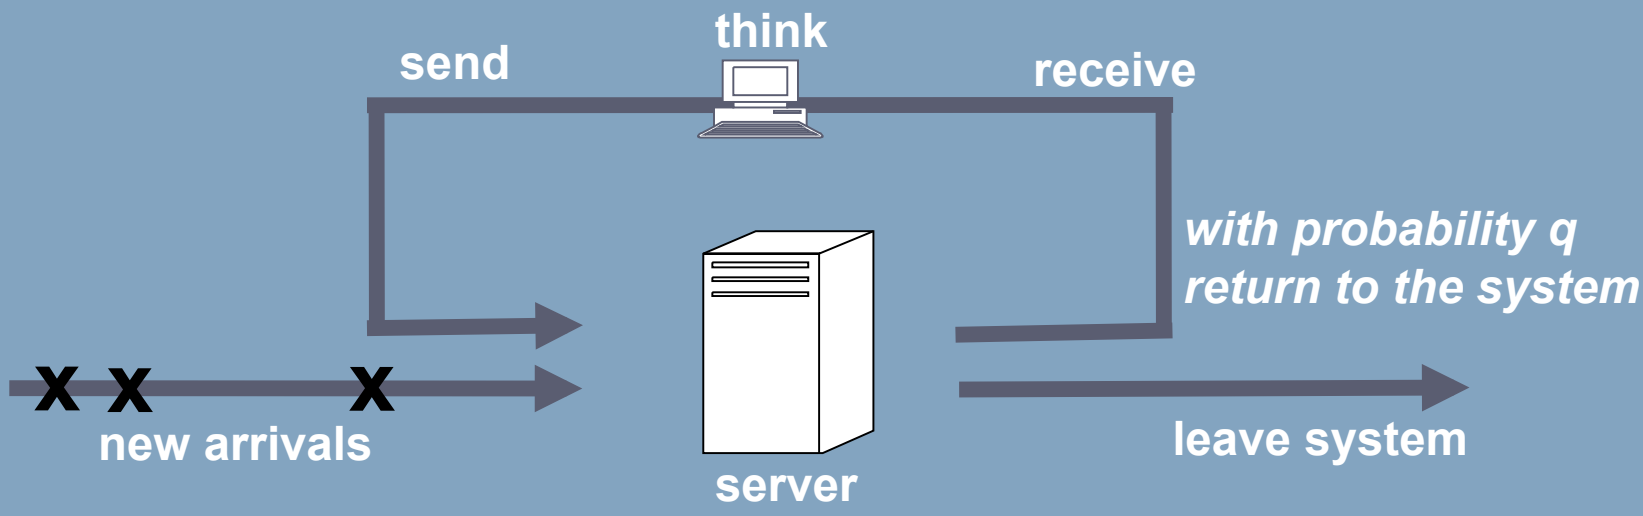

What parameters affect the load?
Does think time affect the load?
How do think times affect response times?

THE EFFECT OF THINK TIME

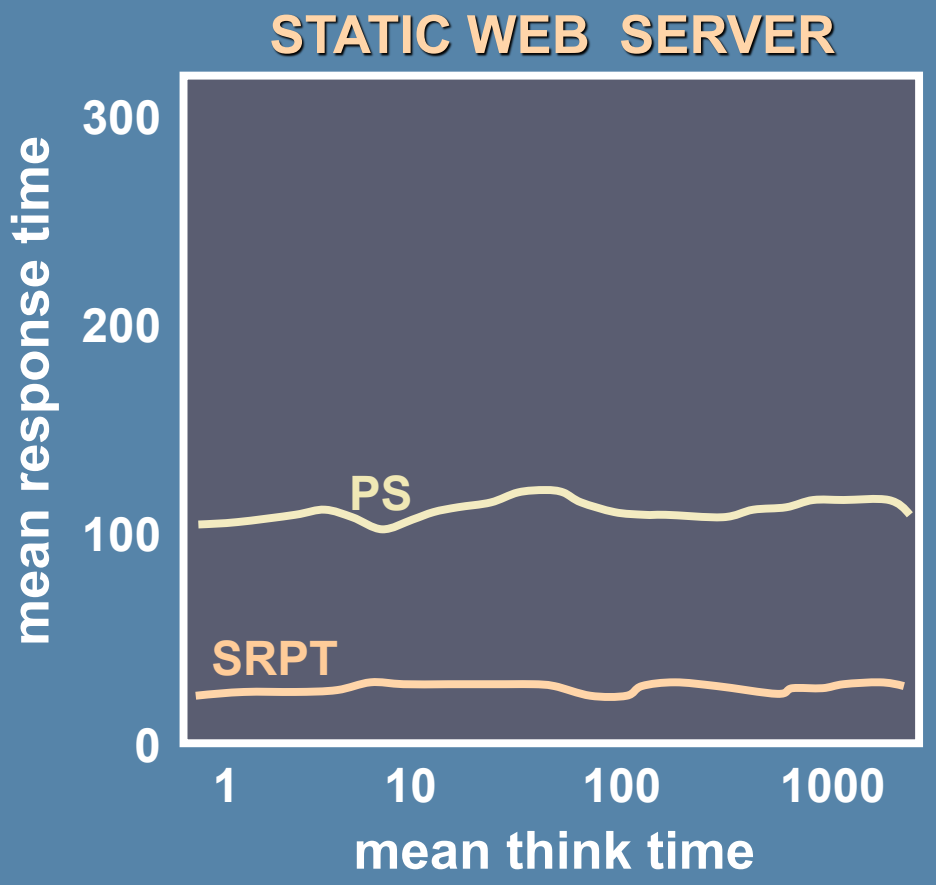

THE PARTLY-OPEN MODEL

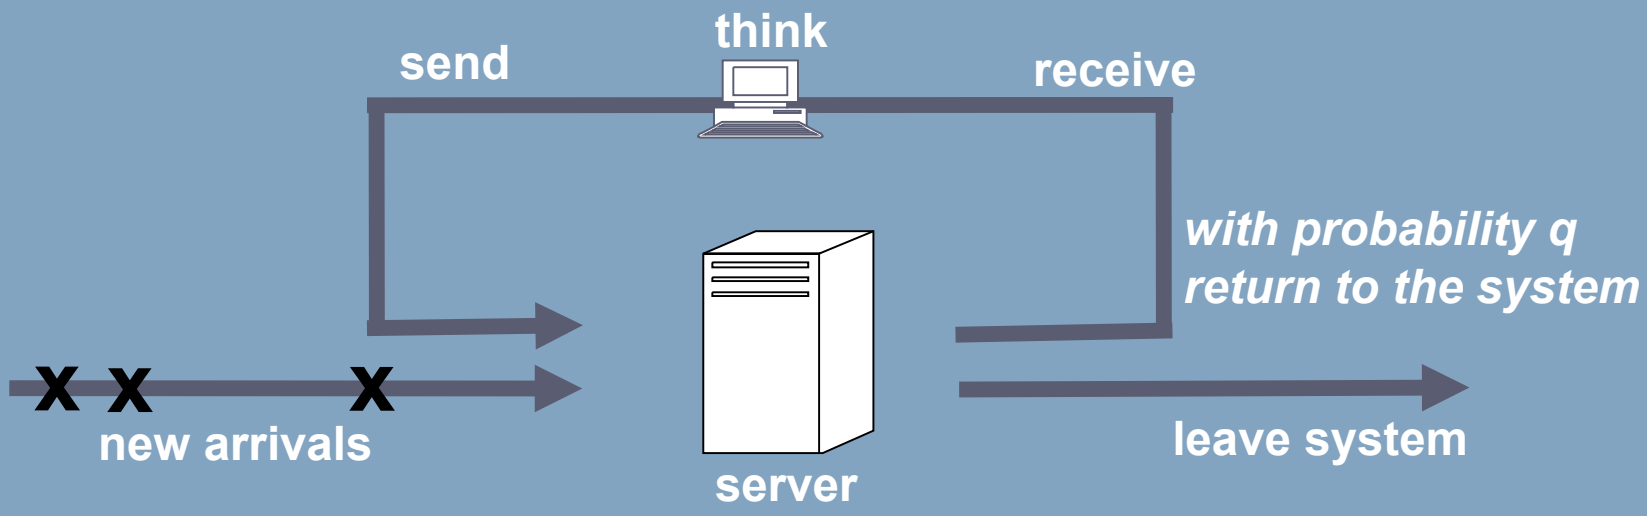

How does this model compare to Open/Closed?

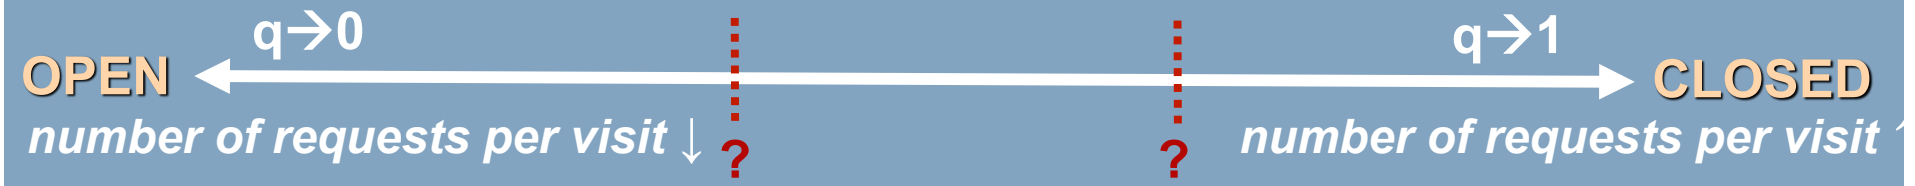

THE TRANSITION FROM OPEN → CLOSED

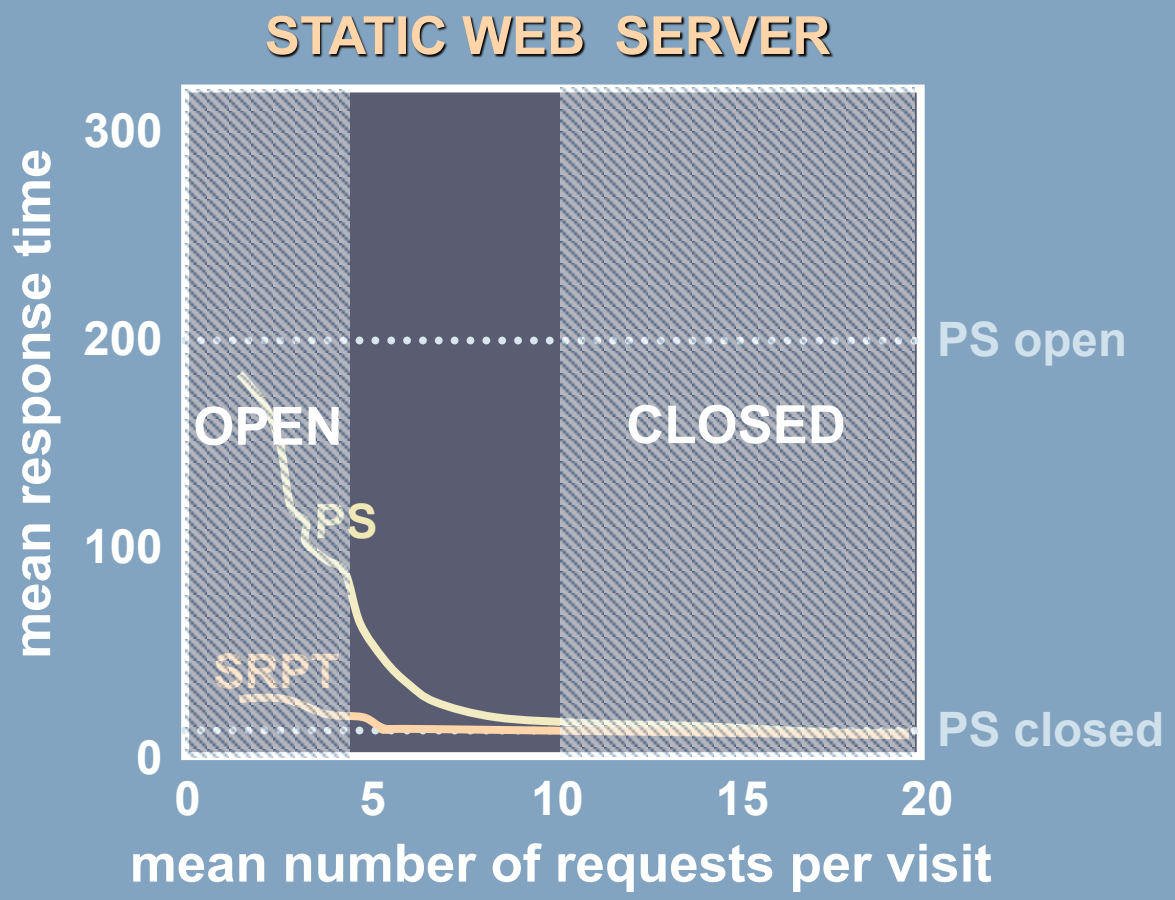

CHOOSING A SYSTEM MODEL

Real web workloads	#req. / visit
A site being "Slashdotted"	1.2
Financial service provider	1.4
CMU web server	1.8
Kasparov vs Deep Blue	2.4
Large corporate web site	2.4
Science Institute USGS	3.6
Online dept. store	5.4
Supercomp. site	6.0
World cup site	11.6
Online gaming site	12.9

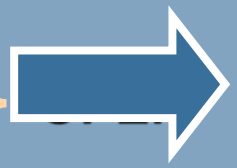

Open or Closed?

Use partly-open system to decide

PARTLY-OPEN

CLOSED

ANOTHER (VERY SHORT) EXAMPLE FROM MY OWN RESEARCH

Probability of losing data in a RAID?

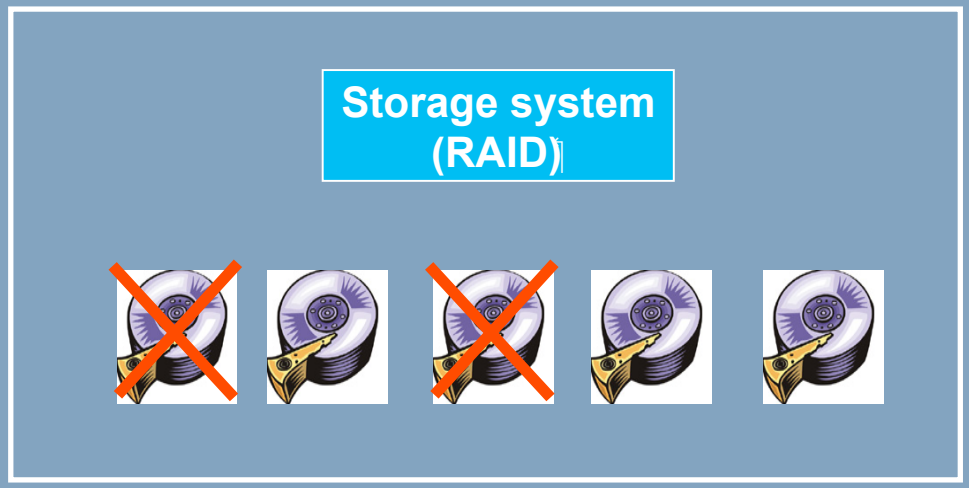

- **Depends on probability that after one drive fails, a second drive fails while reconstructing data.**

Estimating probability of data loss

- **Need probability of second failure during reconstruction**

Standard approach: Use datasheet MTTF and exponential distr.

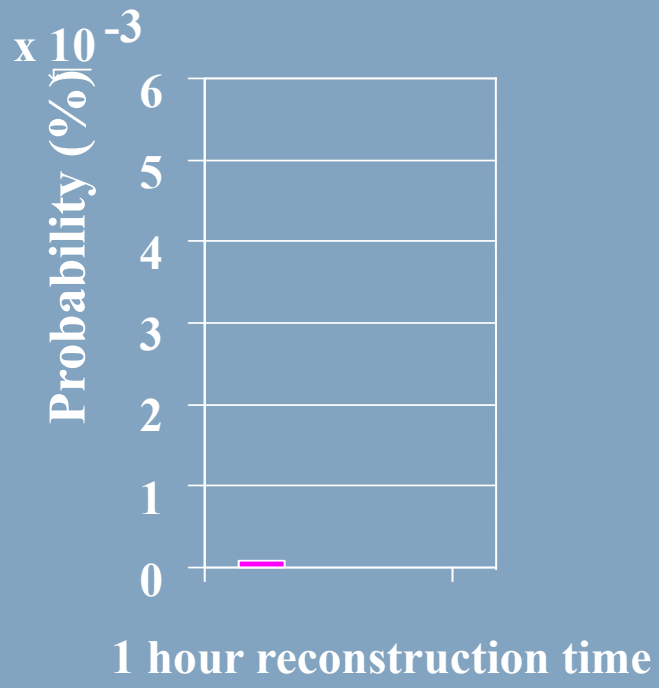

Estimating probability of data loss

- **Need probability of second failure during reconstruction**

Standard approach: Use datasheet MTTF and exponential distr.

Estimate based on data

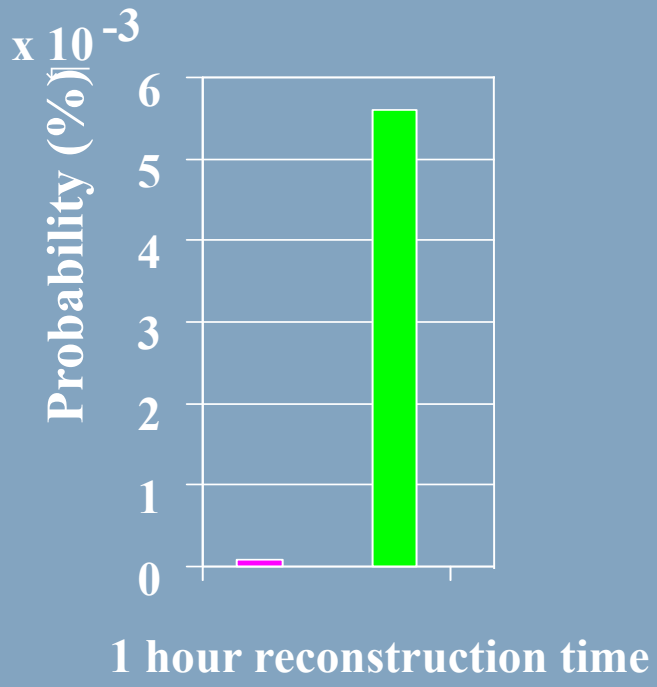

Estimating probability of data loss

- **Need probability of second failure during reconstruction**

Standard approach: Use datasheet MTTF and exponential distr.

Use measured MTTF and exponential distribution

Estimate based on data

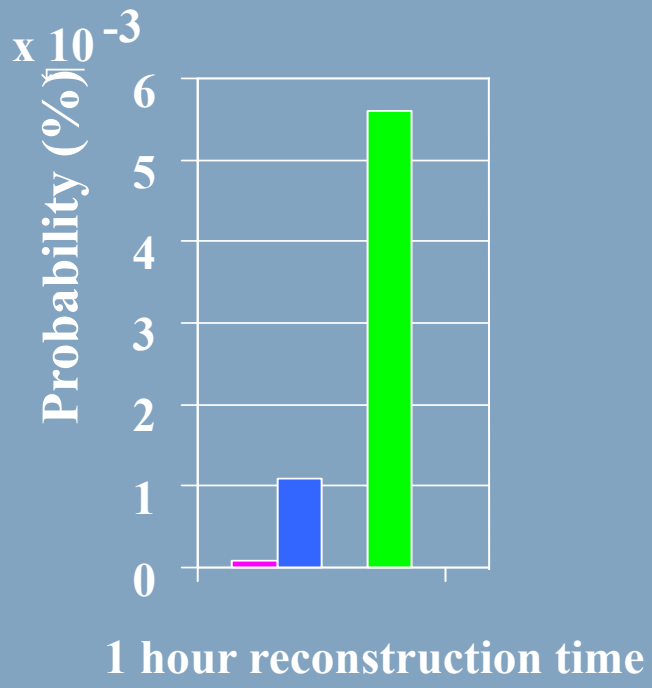

Estimating probability of data loss

- **Need probability of second failure during reconstruction**

- **Standard approach: Use datasheet MTTF and exponential distr.**
- **Use measured MTTF and exponential distribution**
- **Use measured MTTF and Weibull distribution**
- **Estimate based on data**

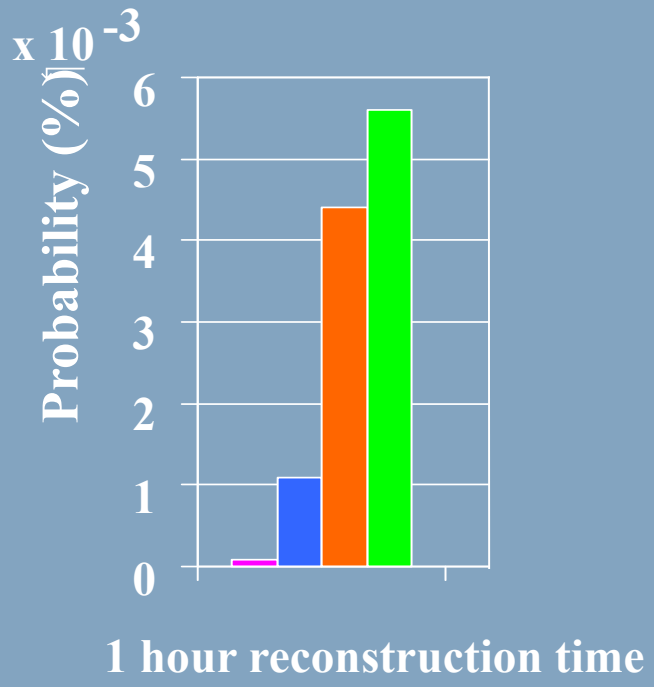

Estimating probability of data loss

- **Need probability of second failure during reconstruction**

- **Standard approach: Use datasheet MTTF and exponential distr.**
- **Use measured MTTF and exponential distribution**
- **Use measured MTTF and Weibull distribution**
- **Estimate based on data**

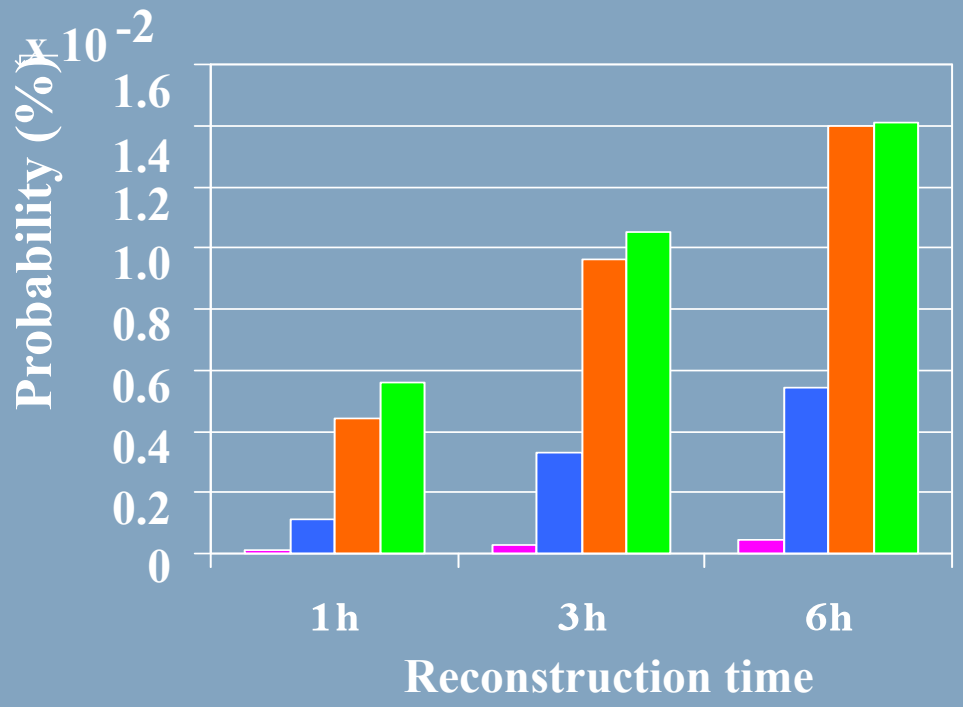

Motivating examples: conclusion

- **Intuition is not always good enough**
  - **Need back-of-the envelope calculations and analytical tools to answer questions.**
- **Workload / fault load matters hugely**
  - **Important to understand what the real world looks like!**