CS 147: Computer Systems Performance Analysis

Approaching Performance Projects

CS 147:
Computer Systems Performance Analysis
Approaching Performance Projects

Overview

Common Mistakes
- Planning Errors
- Measurement Errors
- Design Errors
- Analysis Errors
- Presentation Errors

Systematic Approach
- Pre-Planning
- Planning
- Post-Experiment

Overview
Common Mistakes
Planning Errors
Measurement Errors
Design Errors
Analysis Errors
Presentation Errors
Systematic Approach
Pre-Planning
Planning
Post-Experiment

Some Common Mistakes in Performance Evaluation

- List is long (nearly infinite)
- We'll cover the most common ones

...and how to avoid them

Some Common Mistakes in Performance Evaluation
- List is long (nearly infinite)
- We'll cover the most common ones

...and how to avoid them

No Goals

- If you don't know what you want to learn, you won't learn anything
- Hard to design good general-purpose experiments and frameworks
- So know what you want to discover
- Think before you start
- This is the most common mistake in this class!

No Goals

- If you don't know what you want to learn, you won't learn anything
- Hard to design good general-purpose experiments and frameworks
- So know what you want to discover
- Think before you start
- This is the most common mistake in this class!

Biased Goals

- Don't set out to show OUR system is better than THEIR system
  - Biases you towards using certain metrics, workloads, and techniques
    ...which may not be the right ones
- Don't let your prejudices dictate how you measure
- Instead, try to *disprove* your hypotheses
  - If you fail, that's much stronger evidence

Biased Goals

- Don't set out to show OUR system is better than THEIR system
  - Biases you towards using certain metrics, workloads, and techniques
    ...which may not be the right ones
- Don't let your prejudices dictate how you measure
- Instead, try to *disprove* your hypotheses
  - If you fail, that's much stronger evidence

Unsystematic Approach

- Avoid scattershot approaches
- Work from a plan
- Follow through on it
- Otherwise, you're likely to miss something
  - And in the end, everything will take longer

Unsystematic Approach

- Avoid scattershot approaches
- Work from a plan
- Follow through on it
- Otherwise, you're likely to miss something
  - And in the end, everything will take longer

Incorrect Performance Metrics

- If you don't measure right stuff, results won't shed much light
  - Example: instruction rates of CPUs with different architectures
  - Example: seek time on disk vs. SSD
  - Example: power consumed by mouse
- Avoid choosing metric that's easy to measure but isn't helpful
- Better to struggle to measure, but correctly capture performance

Incorrect Performance Metrics

- If you don't measure right stuff, results won't shed much light
  - Example: instruction rates of CPUs with different architectures
  - Example: seek time on disk vs. SSD
  - Example: power consumed by mouse
- Avoid choosing metric that's easy to measure but isn't helpful
- Better to struggle to measure, but correctly capture performance

Unrepresentative Workload

- If workload isn't like what normally happens, results aren't useful
- E.g., for Web browser, it's wrong to measure
  - Just text pages
  - Just pages stored on local server

Unrepresentative Workload

- If workload isn't like what normally happens, results aren't useful
- E.g., for Web browser, it's wrong to measure
  - Just text pages
  - Just pages stored on local server

Wrong Evaluation Technique

- Measurement isn’t right for every performance problem
  - E.g., issues of scaling, or testing for rare cases
- Measurement is labor-intensive
- Sometimes hard to measure peak or unusual conditions
- Decide whether to model or simulate before designing a measurement experiment

Wrong Evaluation Technique

- Measurement isn’t right for every performance problem
  - E.g., issues of scaling, or testing for rare cases
- Measurement is labor-intensive
- Sometimes hard to measure peak or unusual conditions
- Decide whether to model or simulate before designing a measurement experiment

Overlooking Important Parameters

- Try to make complete list of characteristics that affect performance
  - System
  - Workload
- Don't just guess at a couple of interesting parameters
- Despite your best efforts, you may miss one anyway
  - But the better you understand the system, the less likely you will

Overlooking Important Parameters

- Try to make complete list of characteristics that affect performance
  - System
  - Workload
- Don't just guess at a couple of interesting parameters
- Despite your best efforts, you may miss one anyway
  - But the better you understand the system, the less likely you will

Ignoring Significant Factors

- Factor: parameter you vary
- Not all parameters equally important
- More factors ⇒ experiment takes more work
  - But make sure you don't ignore significant ones
  - Give preference to those that users can vary

Ignoring Significant Factors

- Factor: parameter you vary
- Not all parameters equally important
- More factors ⇒ experiment takes more work
  - But make sure you don't ignore significant ones
  - Give preference to those that users can vary

Inappropriate Experiment Design

- Too few test runs
- Or runs with wrong parameter values
- Interacting factors can complicate proper design of experiments
  - Covered toward end of class

Inappropriate Experiment Design

- Too few test runs
- Or runs with wrong parameter values
- Interacting factors can complicate proper design of experiments
  - Covered toward end of class

Inappropriate Level of Detail

- Be sure you're investigating what's important
- Examining at too high a level may oversimplify or miss important factors
- Going too low wastes time and may cause you to miss forest for trees

Inappropriate Level of Detail

- Be sure you're investigating what's important
- Examining at too high a level may oversimplify or miss important factors
- Going too low wastes time and may cause you to miss forest for trees

No Analysis

- Raw data isn't too helpful
- Remember, final result is *analysis* that describes performance
- Preferably in compact form easily understood by others
  - Doubly important for non-technical audiences
- Common mistake in this class
  - Trying to satisfy page minimum in final report

No Analysis

- Raw data isn't too helpful
- Remember, final result is *analysis* that describes performance
- Preferably in compact form easily understood by others
  - Doubly important for non-technical audiences
- Common mistake in this class
  - Trying to satisfy page minimum in final report

Erroneous Analysis

- Often caused by misunderstanding of how to handle statistics
- Or by not understanding transient effects in experiments
- Or by careless handling of the data
- Many other possible problems in this area

Erroneous Analysis

- Often caused by misunderstanding of how to handle statistics
- Or by not understanding transient effects in experiments
- Or by careless handling of the data
- Many other possible problems in this area

No Sensitivity Analysis

- Rarely does one number or one curve truly describe a system's performance
- How different will things be if parameters are varied?
- Sensitivity analysis addresses this problem

No Sensitivity Analysis

- Rarely does one number or one curve truly describe a system's performance
- How different will things be if parameters are varied?
- Sensitivity analysis addresses this problem

Ignoring Errors in Input

- Particularly problematic when you have limited control over parameters
- If you can't measure an input parameter directly, need to understand any bias in indirect measurement

Ignoring Errors in Input

- Particularly problematic when you have limited control over parameters
- If you can't measure an input parameter directly, need to understand any bias in indirect measurement

Improper Treatment of Outliers

- Sometimes particular data points are far outside range of all others
  - Sometimes they're purely statistical
  - Sometimes indicate a true, different behavior of system
- You must determine which are which

Improper Treatment of Outliers

- Sometimes particular data points are far outside range of all others
  - Sometimes they're purely statistical
  - Sometimes indicate a true, different behavior of system
- You must determine which are which

Assuming Static Systems

- A previous experiment may be useless if workload changes

  . . . or if key system parameters change
- Don't rely blindly on old results
  - E.g. browser measurements from before the rise of HTML5

Assuming Static Systems

- A previous experiment may be useless if workload changes

  . . . or if key system parameters change
- Don't rely blindly on old results
  - E.g. browser measurements from before the rise of HTML5

Ignoring Variability

- Often, showing only mean does not paint true picture of a system
- If variability is high, analysis and data presentation must make that clear
- We'll talk quite a bit about this issue

Ignoring Variability
- Often, showing only mean does not paint true picture of a system
- If variability is high, analysis and data presentation must make that clear
- We'll talk quite a bit about this issue

Overly Complex Analysis

Occam wins!

- Choose simple explanations over complicated ones
- Choose simple experiments over complicated ones
- Choose simple explanations of phenomena
- Choose simple presentations of data

But don't go overboard the other way

Overly Complex Analysis

Occam wins!

- Choose simple explanations over complicated ones
- Choose simple experiments over complicated ones
- Choose simple explanations of phenomena
- Choose simple presentations of data

But don't go overboard the other way

Improper Presentation of Results

- Real performance experiments are run to convince others
  - Or help them make decisions
- If your results don't convince or assist, you haven't done any good
- Often, manner of presenting results makes the difference

Improper Presentation of Results

- Real performance experiments are run to convince others
  - Or help them make decisions
- If your results don't convince or assist, you haven't done any good
- Often, manner of presenting results makes the difference

Ignoring Social Aspects

- Raw graphs, tables, numbers aren't enough for many people
- Need clear explanations
- Present in terms your audience understands
- Be especially careful if going against audience's prejudices!

Ignoring Social Aspects

- Raw graphs, tables, numbers aren't enough for many people
- Need clear explanations
- Present in terms your audience understands
- Be especially careful if going against audience's prejudices!

Omitting Assumptions and Limitations

- Assumptions and limitations are usually unavoidable in real performance studies
- But:
  - Recognize them
  - Be honest about them
  - Try to understand how they affect results
  - Tell your audience about them

Omitting Assumptions and Limitations

- Assumptions and limitations are usually unavoidable in real performance studies
- But:
  - Recognize them
  - Be honest about them
  - Try to understand how they affect results
  - Tell your audience about them

A Systematic Approach To Performance Evaluation

1. State goals and define the system
2. List services and outcomes
3. Select metrics
4. List parameters
5. Select factors to study
6. Select evaluation technique(s)
7. Select workload
8. Design experiments
9. Analyze and interpret data
10. Present results

Usually, it is necessary to repeat some steps to get necessary answers

A Systematic Approach To Performance Evaluation

1. State goals and define the system
2. List services and outcomes
3. Select metrics
4. List parameters
5. Select factors to study
6. Select evaluation technique(s)
7. Select workload
8. Design experiments
9. Analyze and interpret data
10. Present results

Usually, it is necessary to repeat some steps to get necessary answers

State Goals and Define the System

- Decide what you want to find out
- Put boundaries around what you're going to consider
- This is critical step
  - Everything afterwards is built on this foundation

State Goals and Define the System

- Decide what you want to find out
- Put boundaries around what you're going to consider
- This is critical step
  - Everything afterwards is built on this foundation

List Services and Outcomes

- What services does system provide?
- What are possible results of requesting each service?
- Understanding these helps determine metrics

List Services and Outcomes

- What services does system provide?
- What are possible results of requesting each service?
- Understanding these helps determine metrics

Select Metrics

- Metrics: criteria for determining performance
- Usually related to speed, accuracy, availability
- Studying inappropriate metrics makes a performance evaluation useless

Select Metrics

- Metrics: criteria for determining performance
- Usually related to speed, accuracy, availability
- Studying inappropriate metrics makes a performance evaluation useless

List Parameters

- What system characteristics affect performance?
  - System parameters
  - Workload parameters
- Try to make list complete
- But expect you'll add to it later

List Parameters

- What system characteristics affect performance?
  - System parameters
  - Workload parameters
- Try to make list complete
- But expect you'll add to it later

Select Factors To Study

- Factors: parameters that will be varied during study
- Best to choose those most likely to be key to performance
  - Requires knowledge, experience, insight
  - Or perhaps quick pre-experiment
- Select *levels* for each factor
  - Defined as what settings you'll test
- Keep factor list and number of levels short

. . . Because it directly affects work involved

Select Factors To Study

- Factors: parameters that will be varied during study
- Best to choose those most likely to be key to performance
  - Requires knowledge, experience, insight
  - Or perhaps quick pre-experiment
- Select *levels* for each factor
  - Defined as what settings you'll test
- Keep factor list and number of levels short

. . . Because it directly affects work involved

Select Evaluation Technique

- Analytic modeling?
- Simulation?
- Measurement?
- Or perhaps some combination
- Right choice depends on many issues

Select Evaluation Technique

- Analytic modeling?
- Simulation?
- Measurement?
- Or perhaps some combination
- Right choice depends on many issues

Select Workload

- List of service requests to be presented to system
- For measurement, workload is often list of actual requests to system
  - Or special benchmark for applying load
- Must be representative of real workloads!

Select Workload

- List of service requests to be presented to system
- For measurement, workload is often list of actual requests to system
  - Or special benchmark for applying load
- Must be representative of real workloads!

Design Experiments

- How to apply chosen workload?
- How to vary factors to their various levels?
- How to measure metrics?
- All with minimal effort
- Pay attention to interaction of factors
- And remember Heisenberg

Design Experiments

- How to apply chosen workload?
- How to vary factors to their various levels?
- How to measure metrics?
- All with minimal effort
- Pay attention to interaction of factors
- And remember Heisenberg

Analyze and Interpret Data

- Data from measurement almost always has random component
- You must extract real information from the random noise
- Interpretation is key to getting value from an experiment
  - And is not at all mechanical
  - Do this well in your project

Analyze and Interpret Data

- Data from measurement almost always has random component
- You must extract real information from the random noise
- Interpretation is key to getting value from an experiment
  - And is not at all mechanical
  - Do this well in your project

Present Results

- Make them easy to understand
  - For the broadest possible audience
- Focus on most important results
- Use graphics—effectively
- Give maximum information with minimum presentation
- Make sure you explain actual implications
  - Anybody can draw a graph
  - Insight is what matters

Present Results

- Make them easy to understand
  - For the broadest possible audience
- Focus on most important results
- Use graphics—effectively
- Give maximum information with minimum presentation
- Make sure you explain actual implications
  - Anybody can draw a graph
  - Insight is what matters