

Week 1: Introduction

1. Course presentation
2. Background: Internet infrastructure
3. Types of measurements
4. Measurement platforms
- 5. Sound measurement practices**
6. Ethical issues
7. Conclusion

Goal: build trust on measurement results

- Know what to expect
- Know the measurement tool
- Know where the data comes from

Know what to expect


- Identify properties that must hold
 - E.g., RTTs $>$ speed of light
 - E.g., number of bytes in a TCP connection $<$ duration * max capacity
- When properties fail to hold
 - Incorrect assumption: improve mental model
 - Measurement error

Know the measurement tool: Study precision and accuracy

- E.g., timestamp precision: 17:21:10.154379
 - Resolution = 1 μ second?
 - Some clocks only advance every 10 millisecond

Know the measurement tool: examine outliers and spikes

- Measurement errors are often “corner cases”
 - Outlier = unusually low/high values
 - Spike = values that occur a lot
- Further analysis may identify measurement artifacts



Know the measurement tool

- Study precision and accuracy
- Examine outliers and spikes
- Monitor confounding factors
 - Monitor's CPU, memory, traffic
- Evaluate synthetic data, controlled settings
- Compare multiple methods
- Re-calibrate as needed
 - E.g., changing environments

Know where data comes from

- Log meta-data with traces
 - Any information required to fully understand measurements
 - Remember data often used for unexpected purposes
- Examples of meta-data
 - Version of measurement tool and parameters
 - When, where trace was recorded
 - Clock precision
 - Drops, missing data

Picture credits

- Slide 6: figure taken from paper: I. Cunha, F. Silveira, R. Oliveira, R. Teixeira, and C. Diot, "Uncovering Artifacts of Flow Measurement Tools", Passive and Active Measurement Conference (PAM), April 2009.