

Advanced Systems Lab

G. Alonso, D. Kossmann, T. Roscoe

Systems Group

<http://www.systems.ethz.ch>


Basic statistics

- Not a course on statistics
 - You have done that already
 - We assume familiarity with the basics
- Focus on experimental aspects
 - What and when to measure
 - Side effects and different performance patterns
 - Data distributions
 - Sampling
 - Mean, Average, Outliers, deviation, plotting
 - Confidence intervals

Accuracy vs. Precision

Accuracy = how close to the real value (often unknown)

Precision = similarity of the results of repeated experiments




When to measure

What and when to measure

- Decide on the parameters to measure:
 - Throughput, response time, latency, etc.
- Design your experiment
 - Configuration, data, load generators, instrumentation, hypothesis
- Run the experiment and start measuring:
 - When to measure (life cycle of an experiment)
 - What to measure (sampling)

Life cycle of an experiment



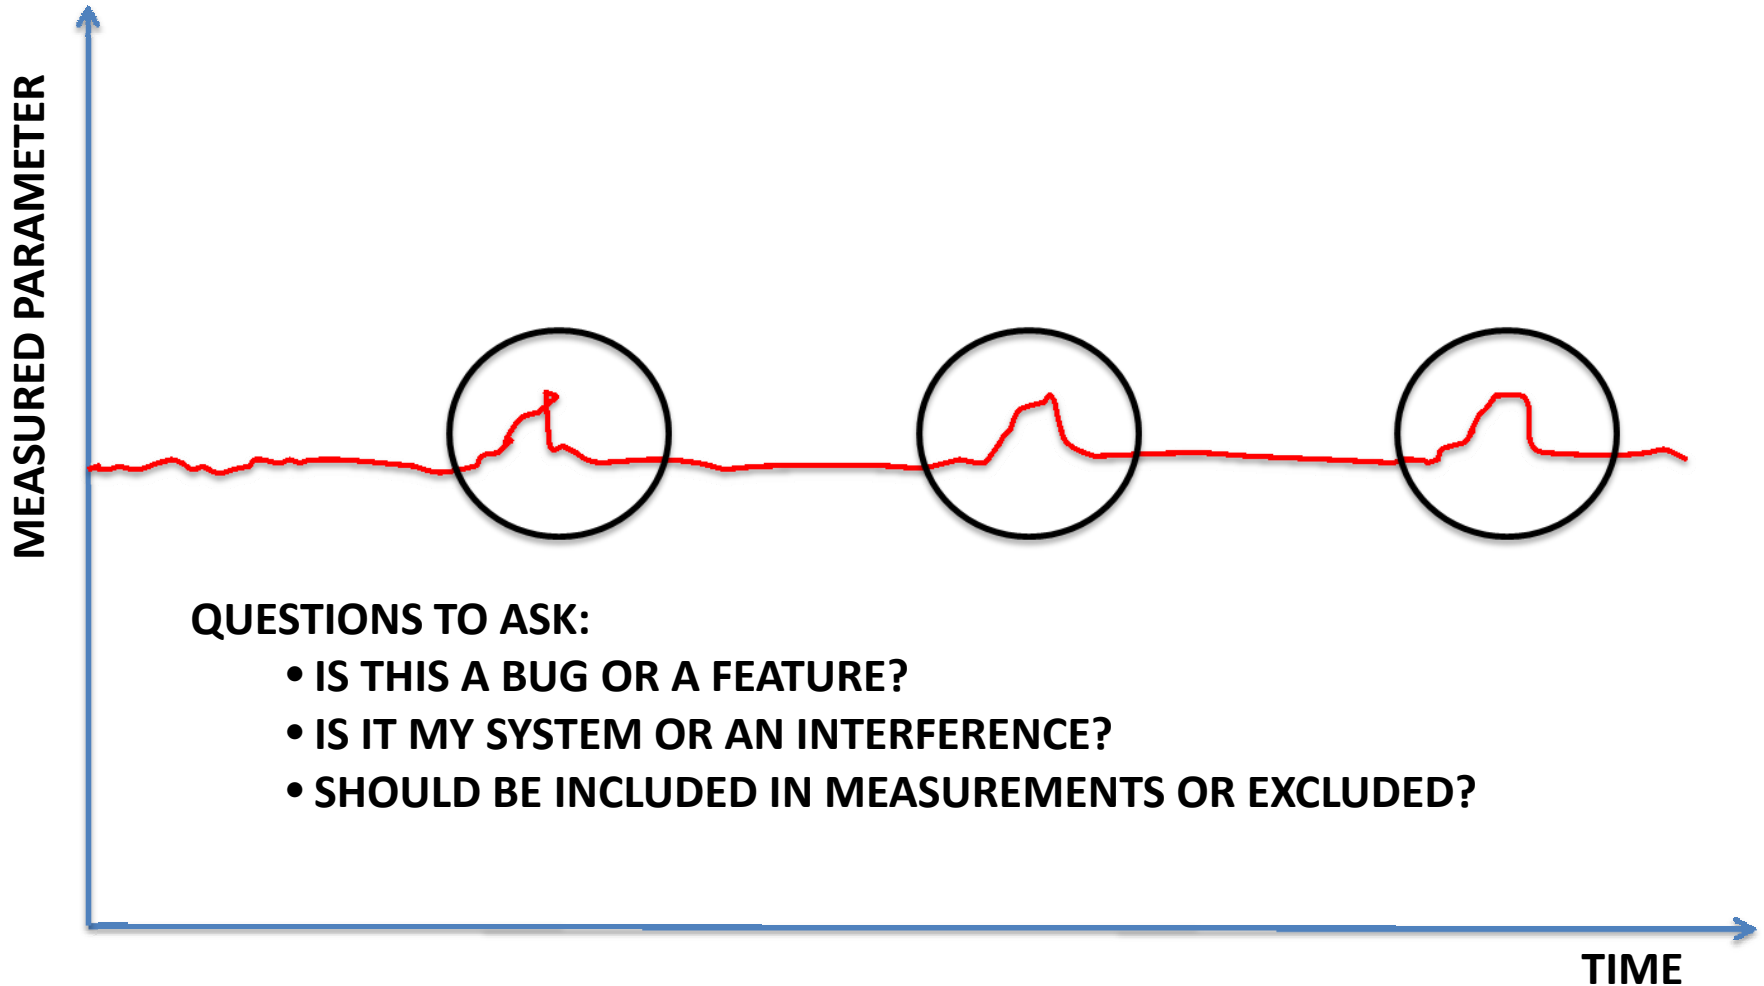
Warm up phase

- Warm up phase
 - Time until clients are all up, caches full (warm), data in main memory, etc.
 - Throughput lower than steady state throughput
 - Response time better than in steady state
- Detect by watching measured parameter changing with time
- Measure only in steady state

Cool down phase

- Cool down phase
 - Clients start finishing, resulting in less load in the system
 - Throughput is lower than in steady state
 - Response time better than in steady state
- Detect by observing when measured parameter suddenly changes behavior
- Stop measuring when clients no longer generate a steady load

Patterns to watch for - glitches

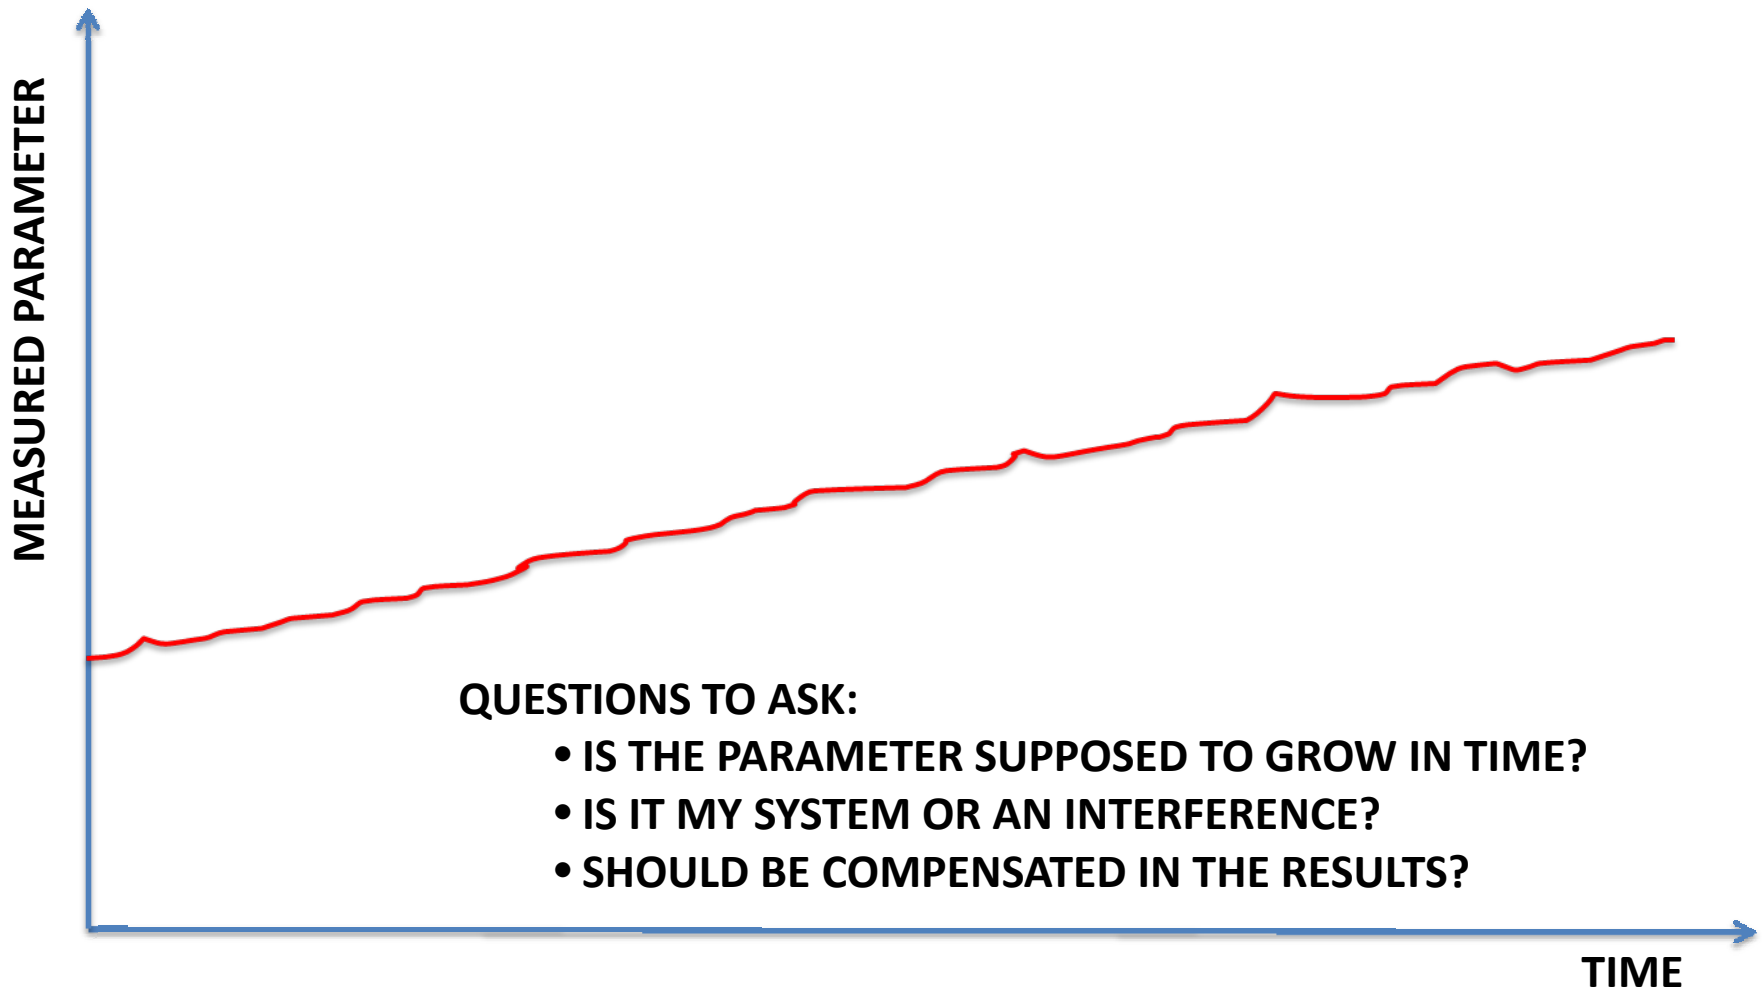


QUESTIONS TO ASK:

- IS THIS A BUG OR A FEATURE?
- IS IT MY SYSTEM OR AN INTERFERENCE?
- SHOULD BE INCLUDED IN MEASUREMENTS OR EXCLUDED?

ASSUME STEADY STATE MEASUREMENTS

Patterns to watch for - trends



ASSUME STEADY STATE MEASUREMENTS

Patterns to watch for - periodic



ASSUME STEADY STATE MEASUREMENTS

Why are these pattern relevant?

- Too few measurements and too short experiments are meaningless
 - May not capture system behavior
 - May not show pathological behavior
 - May not reflect real values
- Statistics are a way to address some of these issues by providing more information from the data and a better idea of the system behavior
 - but applying statistics to the wrong data will not help!

Data distributions

What are we measuring?





- When measuring, we are trying to estimate the value of a given parameter
- The value of the parameter is often determined by a complex combination of many effects and is typically not a constant
- Thus, the parameter we are trying to measure can be seen as a RANDOM VARIABLE
- The assumption is that this random variable has a NORMAL (GAUSSIAN) DISTRIBUTION

Central limit theorem

- Let $X_1, X_2, X_3, \dots, X_n$ be a sequence of independently and identically distributed random variables with finite values of
 - Expectation (μ)
 - Variance (σ^2)


as the sample size n increases, the distribution of the sample average of the n random variables approaches the normal distribution with a mean μ and variance σ^2/n regardless of the shape of the original distribution.

How does it work?




http://en.wikipedia.org/wiki/File:Central_limit_thm.png

Normal or Gaussian distribution



Meaning of standard deviation


- The standard deviation gives an idea of how close are the measurements to the mean
 - important in defining service guarantees
 - important to understand real behavior



Mean is 50 but standard deviation is 20, indicating that there is a large variation in the value of the parameter

Mean of a sample

- To interpret a given measurement, we need to provide complete information
 - The mean
 - The standard deviation around the mean



Mean and standard deviation

- The standard deviation defines margins around the mean:
 - 64% of the values are within $\mu \pm \sigma$
 - 95% of the values are within $\mu \pm 2\sigma$
 - 99.7% of the values are within $\mu \pm 3\sigma$
- For a real system is very important to understand what happens when the values go beyond those margins (delays, overload, thrashing, system crash, etc.)

Calculating the standard deviation

- Mean and standard deviation:

$$\mu = \frac{\sum_{i=1}^N x_i}{N} \quad \sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2}$$


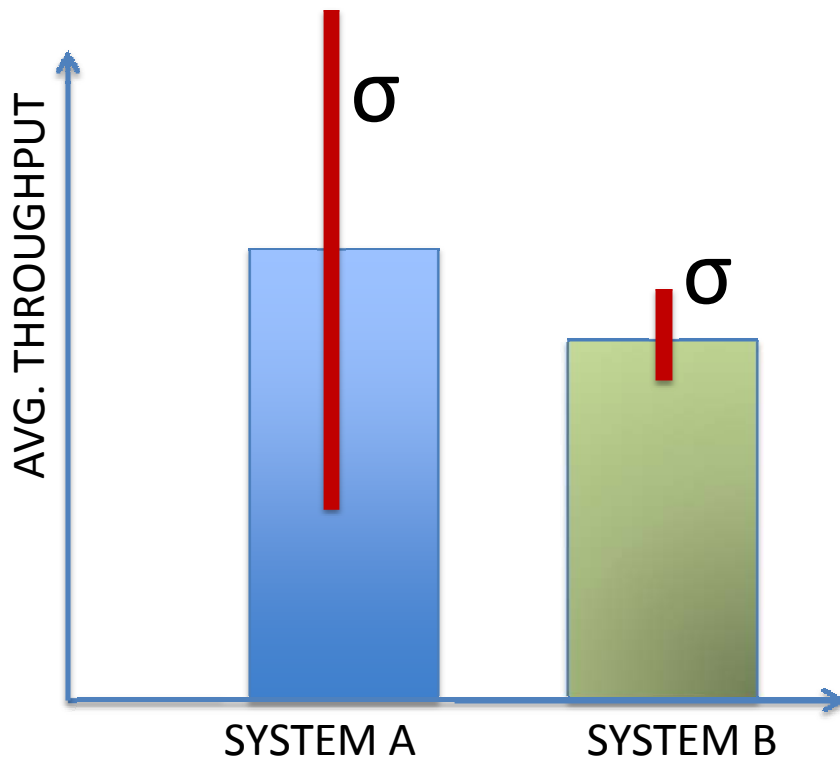
- In practice, use:

$$s = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2}$$

Comparisons

- What is better?

Deterministic behavior is often more important than good performance



In practice

- In many systems, the standard deviation is almost more important than the mean:
 - 90% of the queries need to be answered in less than X seconds
 - No web request can take longer than 5 seconds
 - Changes have to be propagated in less than 10 seconds
 - Guaranteed bandwidth higher than X 90% of the time
- Achieving determinism is often done at the cost of performance

Confidence intervals

Background

- When measuring in software system, we typically do not know neither the value of the parameter we are measuring (μ) nor its standard deviation (σ)
- Instead, we work with mean of the sample \bar{x} and the estimated standard deviation (s)
 - the result is no longer a normal distribution but a t-distribution
 - The t-distribution depends on n , the amount of samples
 - For large n , the t-distribution tends to a normal distribution

Confidence interval

- Since typically we are not measuring an absolute value (unlike in the natural sciences), the notion of confidence interval is particularly useful in computer science
- A confidence interval is a range (typically around the mean) where we can say that if we repeat the experiment 100 times, the value observed will be within the confidence interval m times (e.g., $m=95$, leading to a 95% confidence interval)

Calculation

- The confidence interval is calculated as follows:

$$CI = \bar{x} \pm t \cdot \frac{s}{\sqrt{n}}$$

Where s is the sample standard deviation, n the number of samples and t the critical value of the t-distribution

T in a table

Look up the value for the desired confidence interval and (n-1)

<i>One Sided</i>	75%	80%	85%	90%	95%	97.5%	99%	99.5%	99.75%	99.9%	99.95%
<i>Two Sided</i>	50%	60%	70%	80%	90%	95%	98%	99%	99.5%	99.8%	99.9%
1	1.000	1.376	1.963	3.078	6.314	12.71	31.82	63.66	127.3	318.3	636.6
2	0.816	1.061	1.386	1.886	2.920	4.303	6.965	9.925	14.09	22.33	31.60
3	0.765	0.978	1.250	1.638	2.353	3.182	4.541	5.841	7.453	10.21	12.92
4	0.741	0.941	1.190	1.533	2.132	2.776	3.747	4.604	5.598	7.173	8.610
5	0.727	0.920	1.156	1.476	2.015	2.571	3.365	4.032	4.773	5.893	6.869
6	0.718	0.906	1.134	1.440	1.943	2.447	3.143	3.707	4.317	5.208	5.959
7	0.711	0.896	1.119	1.415	1.895	2.365	2.998	3.499	4.029	4.785	5.408
8	0.706	0.889	1.108	1.397	1.860	2.306	2.896	3.355	3.833	4.501	5.041
9	0.703	0.883	1.100	1.383	1.833	2.262	2.821	3.250	3.690	4.297	4.781

Some observations

- For a fixed n
 - Increasing the confidence ($100\%(1-\alpha)$) implies to extend the confidence interval
- To reduce the confidence interval
 - we decrease the confidence or,
 - we increase the number of examples
- For experiments, fix a target (typically 95% confidence in a 5-10% interval around the mean) and repeat the experiments until the level of confidence is reached –if ever ...

Example

- Mean = 122
- $s = 9$
- $n = 30$
- t (two sided, 29, 95%) = 2.045
- $CI = 122 \pm (2.045 \cdot 9/30^{-2})$
- In 95 out of 100 runs, the mean will be between 119 and 125

Putting it all together

Look at all the data

- Make sure you are looking at the complete picture of the experiment and your measurements do not include side effects (warm up, cool down, repetition effects)
- Once you are sure you have identified the valid data and that it looks reasonable, then apply statistics to it

Standard deviation

- All measurements and graphs have to be accompanied by the standard deviation, otherwise they are meaningless
 - Provides an idea of the precision
 - Provides an idea of what will happen in practice
 - Provides an idea of how predictable performance is
- Repeat the experiments until you get a reasonable standard deviation (enough values are close enough to the mean)

How long to run?

- Until you reach a reasonable confidence level that the value lies within a reasonable margin of the mean
- Confidence intervals are the way to quantify how often the reported result is going to be observed in practice
 - “we repeated the experiments until we reached a 95% level confidence for an interval 5% around the mean”

Advice

- It is a good idea to run a long experiment to make sure you have seen all possible behavior of the system:
 - Glitches only every 3 hours
 - Memory leaks after 1 M transactions
- In reality, tests have to resemble how the system will be used in practice