lecture # 22 CS695

the experiments story.

motivate. formulate. design. build. deploy. measure. infer. report. repeat

(*) an empirical (measurement-based) analysis is needed for?

- establish motivation —
- correctness
- performance-cost analysis / comparisons
- ~ failure / unexpected behaviour analysis.
  under-the-hood
- ~ performance characterization

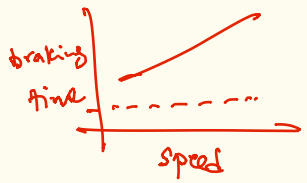

(0) meta-point / soul of all experiments

No question ⇒ No experiment!

Q ~ question (why am I doing the expt?)

A ~ accurate (only conclude what is "reasonable")

R ~ reproducible (redone & verified)

C ~ complete (cover the entire parameter space)

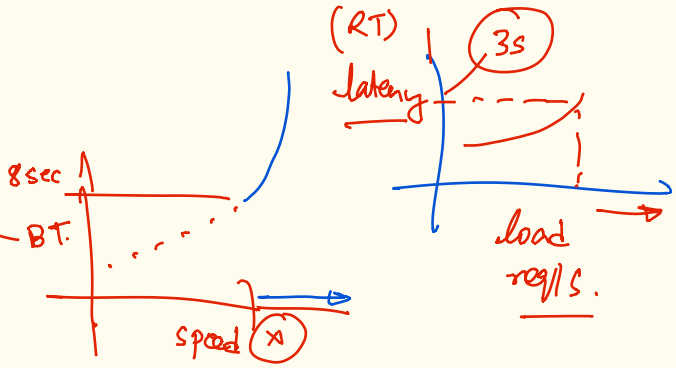

(work/demand) inputs

workloads

configuration & setup (parameters)

system

prototype

design/solution/opt.

algo./heuristic

feature/mechanism

platform

(measurements) metrics outputs

tput

latency

utilization

① Questions (type of investigation)

- correctness
- scalability (saturation behaviour)
- best case ——— worst case
- failure/root cause analysis
- comparitive analysis
- hypothesis testing — performance characterization.

① Questions
② Setup
③ Execution ← modality, types, instrumentation
④ observation & inferences.

② Setup.

~ hardware, software, workload } configurations & settings.

(a) representative — as closely should mimic "real" conditions.
- explicitly state deviations/approximations.

(b) workloads ~ needs specification
- synthetic / micro-benchmark
- trace driven replay
- ~ real world load (production systems).

(statistical) model-based load generator

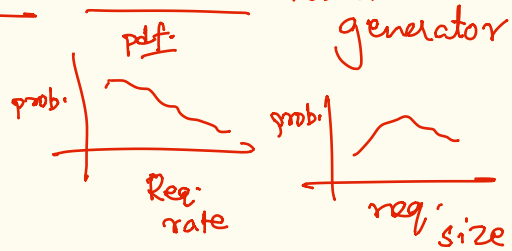

- input parameter distributions are inputs to load generator.

- file IO benchmark
  - iozone dbench
  - req. rate, file size dist².
  - r/w proportion