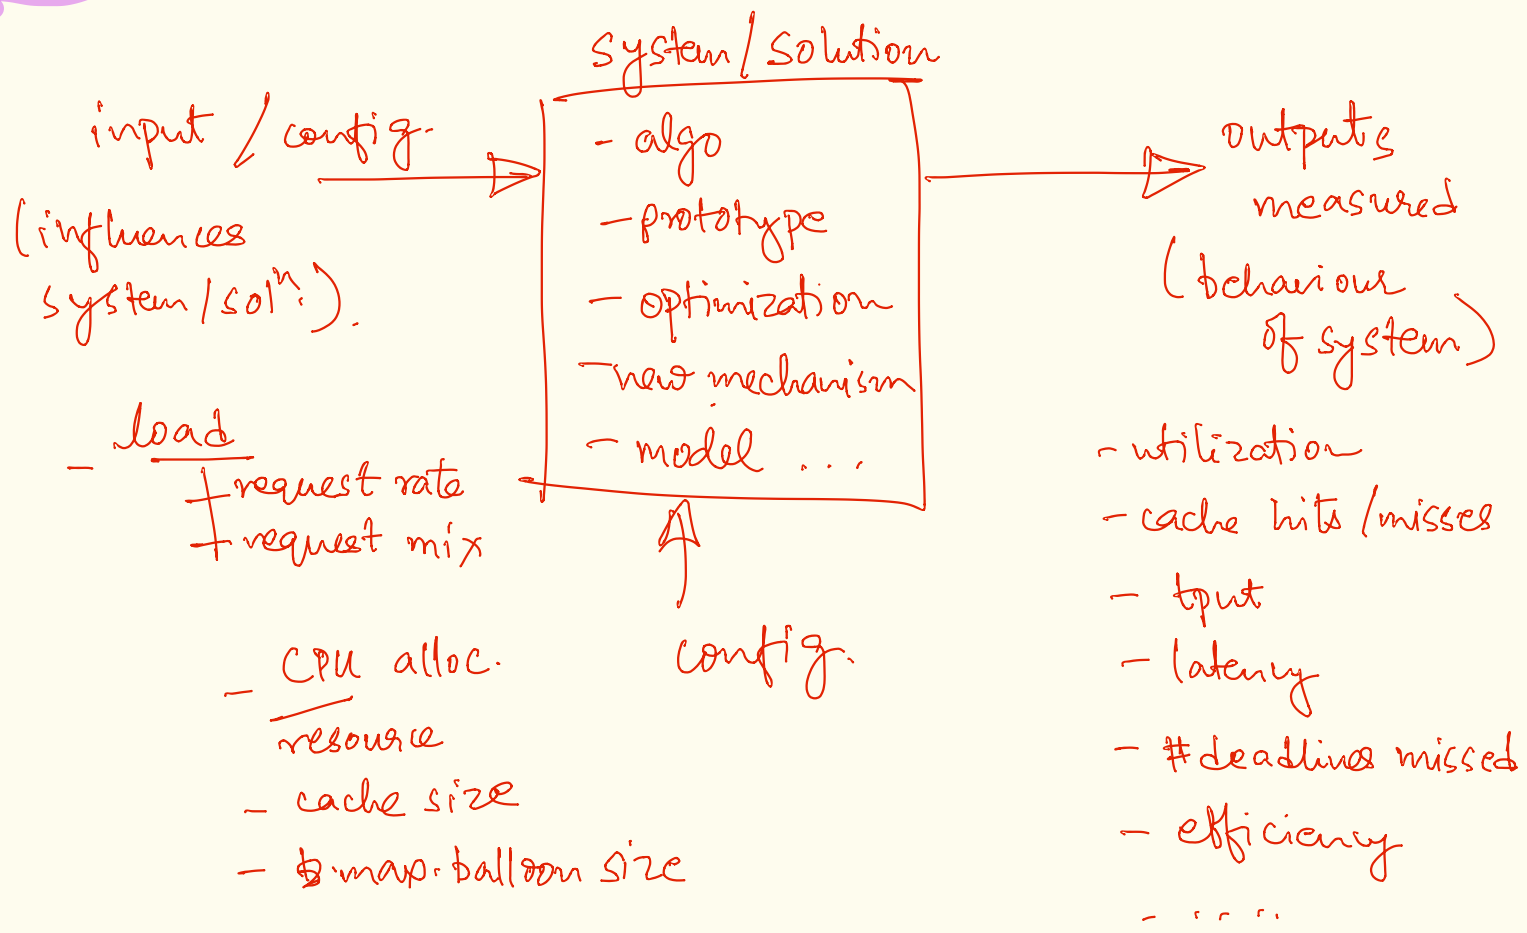



the experimentation process (design of experiments).



Q: meta-point no question, no experiment.

Q ~ question (why?)

A ~ accurate (correct inference)



R ~ reproducibility (complete specification of setup, input, config)

C ~ completeness (does it cover enough ground to answer the question).

① types of questions

- correctness — is a cache-hit always yielding an object in the cache?
- comparative analysis

↳ solⁿ.1 vs solⁿ.2 vs solⁿ...

↳ beware of comparing w/ a very bad solⁿ.


- cost vs benefit analysis
⇒ perf. vs resources / cost

- scalability / saturation behaviour

- failure analysis / performance analysis

- sensitivity analysis

- characterization




(2) setup. → hardware
software
workload.

(i) workloads

- benchmarks
synthetic workload

custom program
that generates sequence / content
of work.

- trace-driven replay



- real load

production
systems facing real requests.

(ii) modality (how to execute experiments?)

1. → ~~was~~ analytical models
2. → simulation
3. → emulation
4. → prototype
5. → production system


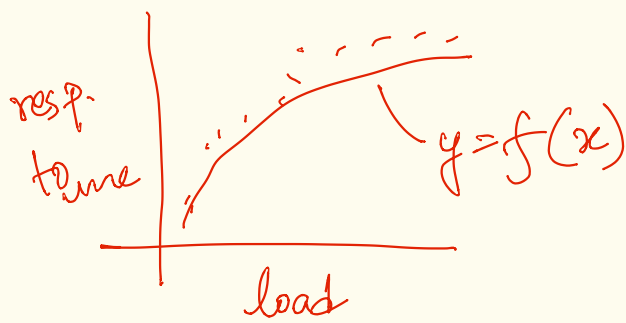
1. analytical modeling

- build a (mathematical) representation of the system/solⁿ.
and use it for expected-case behavior, what-if analysis.

- example models: queuing models, $y = f(x)$

Inputs for building models

are/can be based on real experiments/measurements.



2. simulations

- "simulate" main/solution components / action items, not all levels of functionality.

e.g: state handling — objects (IDs) in cache

procedures/sequences — scheduling / caching policy.

logic of solⁿ. / heuristic.

- when to use?

- quick fdbk / analysis of solution.

- no access to real setup / hardware

- large & complicated system.

- # config. parameters very large.


⊗ simulations are approximate!

- do not account for all system interactions.

eg: lock ordering, non-determinism in scheduling,
hw access order, hw latencies etc.

- use models

eg: on cache miss need to simulate disk
access latency



- no actual msgs or data passed or hw pkts. passed.
only events updated for these actions.

⊗ basic simulator loop.

```
while (1) {
```

```
    e = get next event (Q);
```

```
    process event (e, Q);
```

```
}
```

```
add event (e, Q) {
```

```
    t = time of event (e);
```

```
    add to queue (e, t, Q);
```

```
}
```

Q: priority queue

(based on
time of
event)

③ on cache miss event,

add a new event at time $t+d$

where 'd' is disk access latency.

& e. type is disk read completion.

③ emulation

3: - closer to real solution.

- mimic behaviour of some components of solution.

① e.g: to emulate network latency on a WAN/Internet,
add delay to delivery of each pkt. between endpoints.

~ need model to emulate add delay values!

② emulate behavior of resource / solution.

⇒ virtual disk.

④ prototype :- implement all components

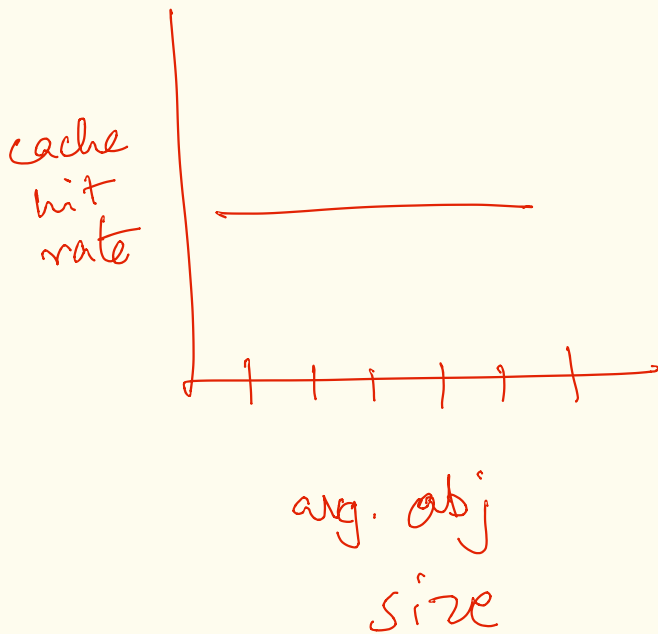
- scale of setup may be different
(compared to production system).

- larger turn around times.

(iii) how to experiment?

independent \approx dependent variables.
varied measured.

- across experiments all thing to remain same, except change in a single independent variable.



→ cache size proportionate to obj size



- cache size remains same

(i) instrumentation overheads.

- does WSS estimation slow the application down?

w/ WSS est.

w/o WSS est.

t_{put}: 1000 rps

1000 rps

- quantify overheads explicitly.

(4) observation & inferences \Rightarrow QARC

(5) representation — graphs
tables