Y.C. Tay: 3 Projects

database

target: replace TPC benchmarks

↑ bad

not application-specific

idea: empirical (application) dataset

↓ done

synthetically scale
to desired size

↓ current work

collaborative framework
for tools to tweak data

↑ from developer community

↖ application-specific

memory

target: a scientific analysis of
recency (LRU) vs popularity
in cache replacement policies

idea: use my Cache Miss Equation

$$\text{Prob(miss)} = \frac{1}{2}\left(K + \sqrt{K^2 - 4}\right)(P^* + P_0) - P_0$$

where $K = 1 + \dfrac{M^* - M_0}{M - M_0}$

↑ cache size

analyze and interpret
how recency/popularity
affects M^*, M_0, P^*, P_0

networking

target: a paradigm shift
in congestion control
from packet loss
to bandwidth-delay product

$$N_{BDP} = B_{max} R_{min}$$

↑ bottleneck bandwidth

↑ min RTT

Over 2015-17, Google shifted production traffic from TCP CUBIC (loss-based) to BBR (N_{BDP}-based) using probes to estimate B_{max} and R_{min}

idea: use queueing formula

$$X = \frac{N}{R_{min} + \frac{N}{B_{max}}}$$

regression+extrapolation (no probes)

works for any TCP version

works for video games