MichiganEngineering

Warehouse-Scale Computers
What can academics do?
A Call to Arms

Thomas F. Wenisch
The University of Michigan

EXADAPT Workshop

Can't have these under our desks…

Illustration: Bryan Christie Design

Even those that do have 'em, can't touch 'em

How can academics do data center research?

Traditional Architectural Simulators?

- 5 orders of magnitude slowdown… (e.g., GEM5)
  - × 16+ cores per chip
  - × 2 chips per box
  - × 42 boxes per rack
  - × 800+ racks per datacenter
  - Oh, need GPU models too…
  - And accurate IO…
  - And network…
  - And power/cooling systems…

== PhD awarded posthumously

How about pen & paper analysis?

- Real workloads are not Poisson [Gupta+ 2009]

Actual Google Search Leaf Interarrival Dist.

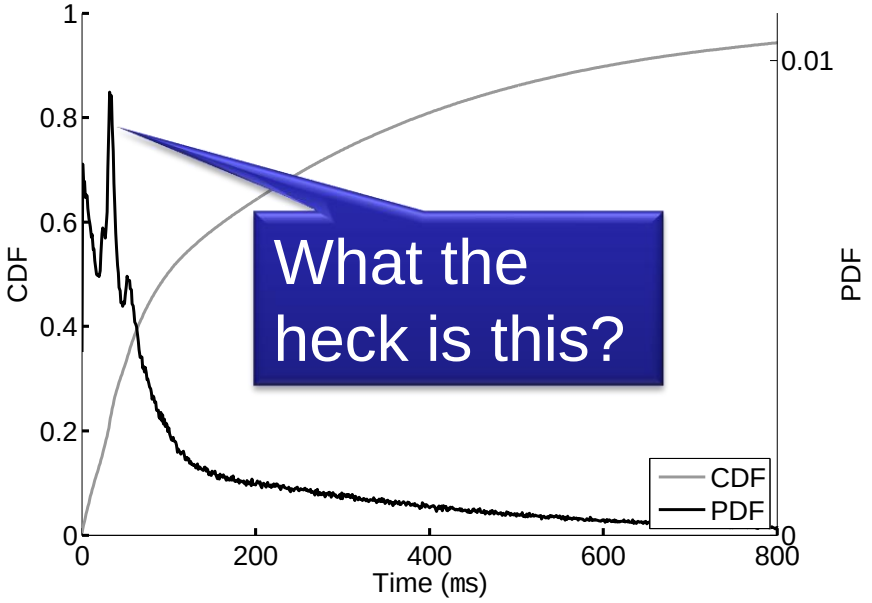

Actual Google Search Leaf Service Dist.

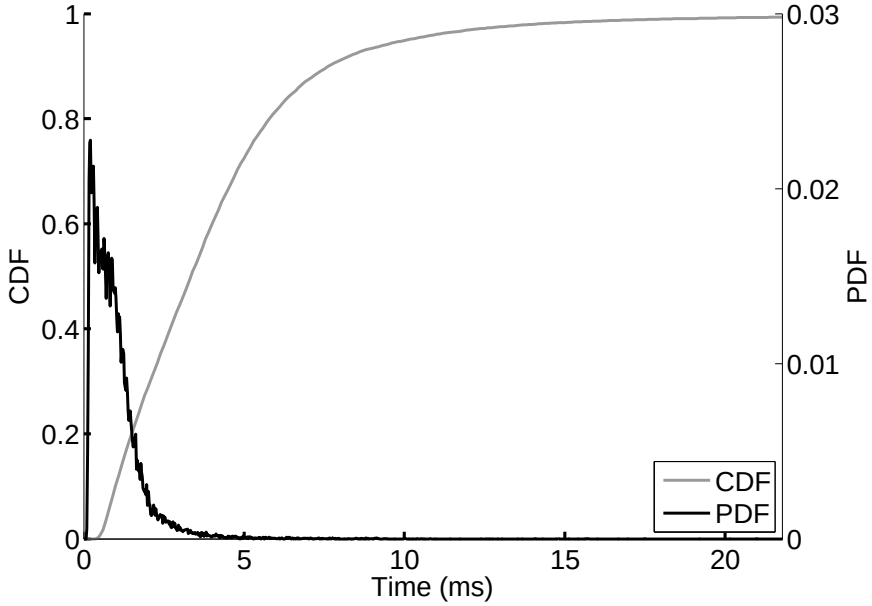

Data from [Meisner ISCA'11]

Just because the math is easy…

Doesn't mean its right

How about extrapolating from 1 node?

- OK, can work sometimes
  - E.g., memcached
  - But, accurate load-testing is really hard

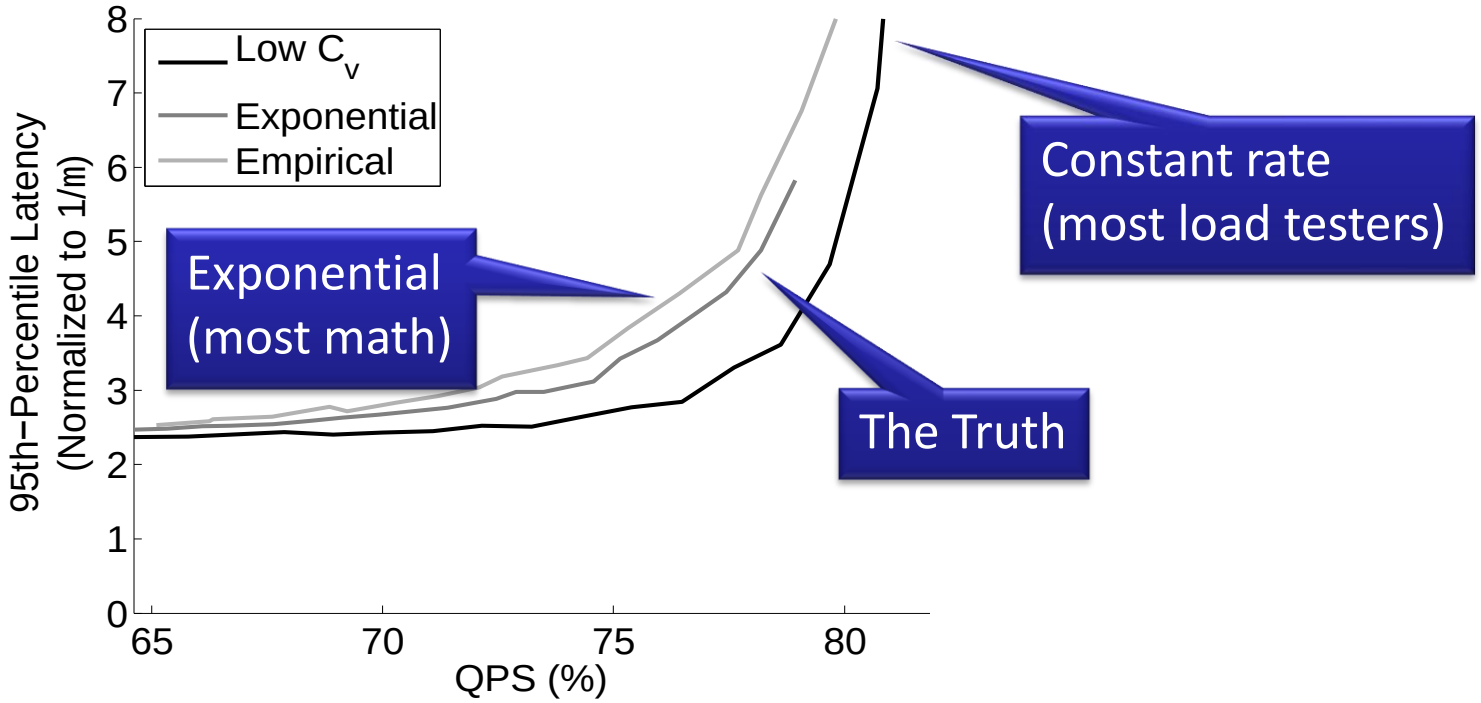

Data from [Meisner ISPASS'12]

Michigan approach: BigHouse - Raising the level of abstraction [ISPASS'12]

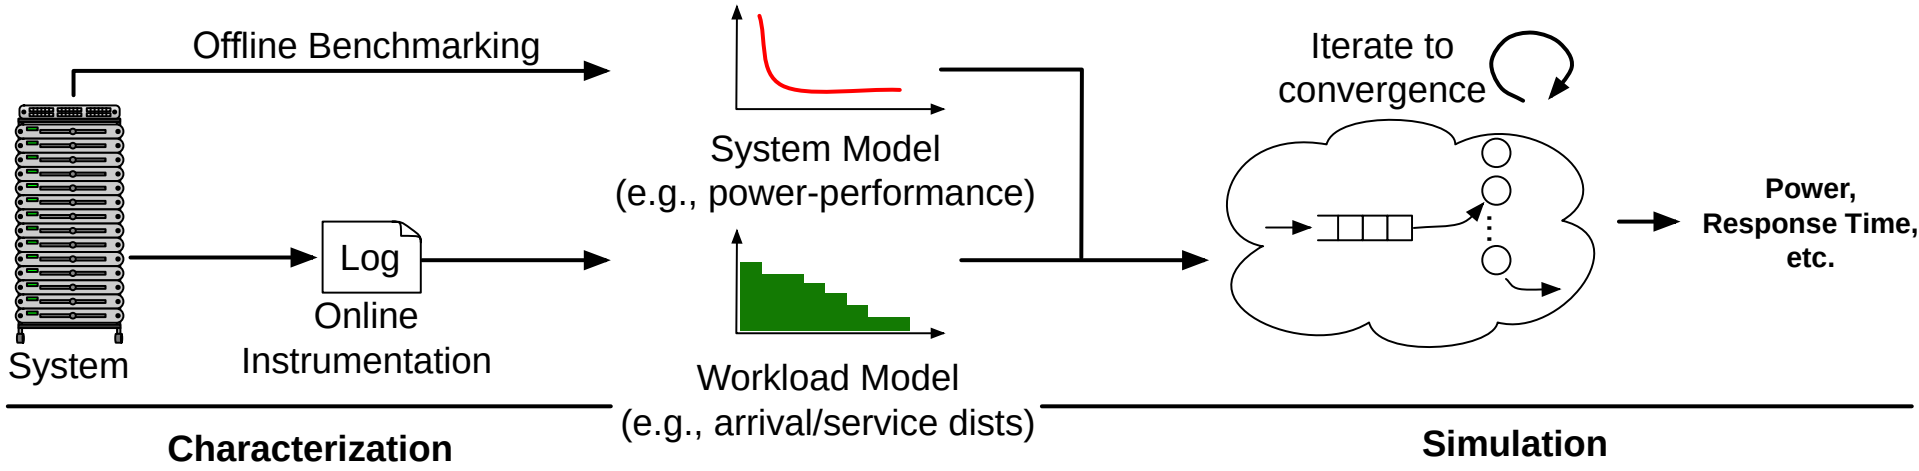

Integrates workloads, power-performance and queuing models that are otherwise analytically intractable (e.g. G/G/k)

A Call to Arms…

- Need to foster community resources
  - Traces from production facilities
  - Benchmarks of internet service workloads
  - Peer-reviewed models of app. behavior
  - Facility-scale simulation tools
  - Shared testbeds (e.g., PlanetLab)

Need industry, academia, gov't collaboration