

Modeling the Internet is fun -
but can you make a living?

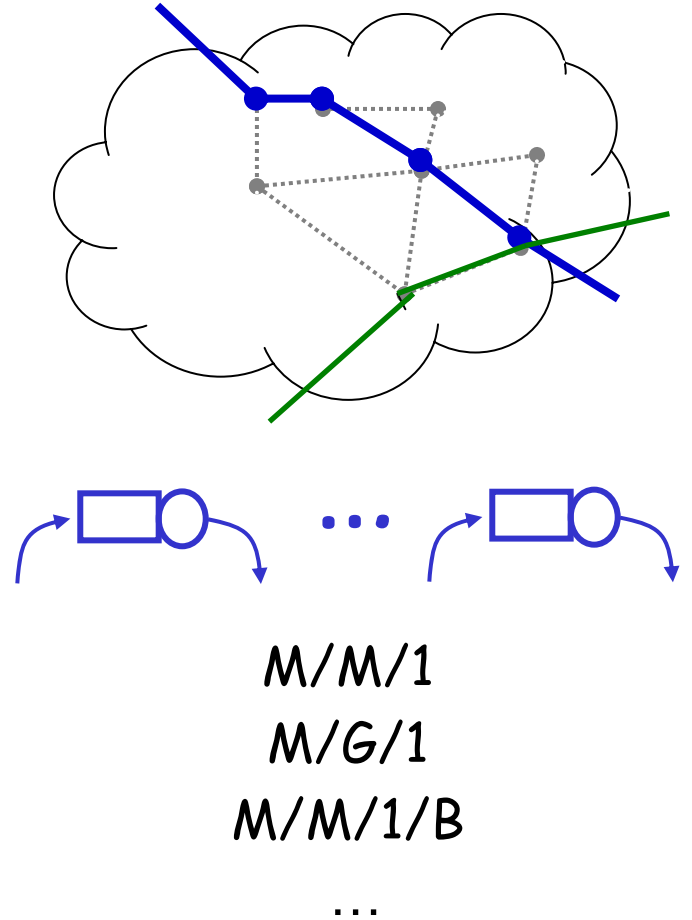
Don Towsley
Dept. of Computer Science
UMass - Amherst

Overview

- some history
- successes
 - fluids
 - Poisson
- challenge - model guided measurement
- parting thoughts

Some history

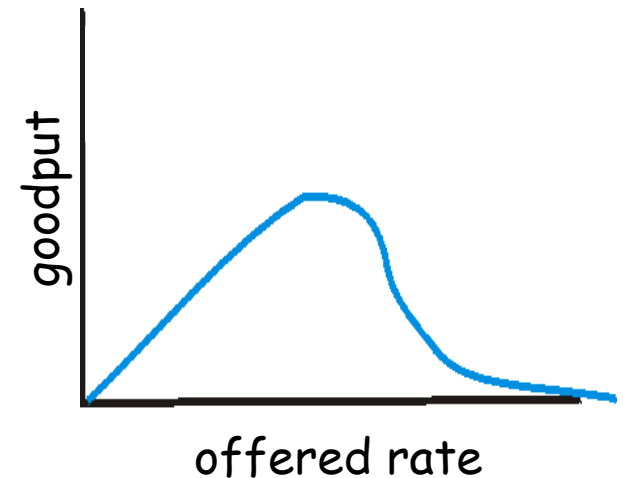
- first network model
(Kleinrock, 60s)
- independence assumption
- decomposition
- Poisson



70s - 80s age of queuing theory

- router/switch models
- flow control
- error control
- non-Poisson arrival processes

- many contributors
 - Lam, Reiser, Schweitzer



Two highlights (70s-80s)

- optimization-based routing (Gallager, 76)
- AIMD congestion control (Jacobsen, 88)

90s - diversification

- ❑ network calculus (Cruz, Parekh, Chang, LeBoudec, ...)
- ❑ effective bandwidths (Guerin, Kelly, Gibbens, Chang,...)
- ❑ long range dependence (Leland, ...)
- ❑ closed loop flow modeling (Firoiu, Padhye, Misra, ...)
- ❑ optimization-based resource allocation (Kelly,...)
- ❑ model-based measurement (Bolot, Duffield, Paxson, ...)
- ❑ wireless (Gupta, Kumar, ...)

Observations

- models serve different purposes
- performance evaluation
 - configuration
 - comparison of competing designs
- to develop understanding, algorithms, mechanisms
 - "simple as possible, but no simpler"
(Einstein)

Observations

- models serve different purposes
- performance evaluation
 - configuration
 - comparison of competing designs
- to develop understanding, algorithms, mechanisms
 - "simple as possible, but no simpler"
(Einstein)

Case studies

- congestion control and TCP
 - models to describe behavior
 - optimization-based models

- model-based measurement

Modeling TCP - an example

Modeling TCP

- background
- TCP friendliness and the IETF
- fluid and Poisson models
- optimization and control

History of TCP

- 1974: first draft of TCP/IP (Cerf, Kahn)
- 1983: ARPANET switches on TCP/IP
- 1986: congestion collapse
- 1988: congestion control for TCP
(Van Jacobson)

IETF and TCP friendliness

IETF mandated (mid 90s)

“thou shalt be TCP friendly”

operating definition was

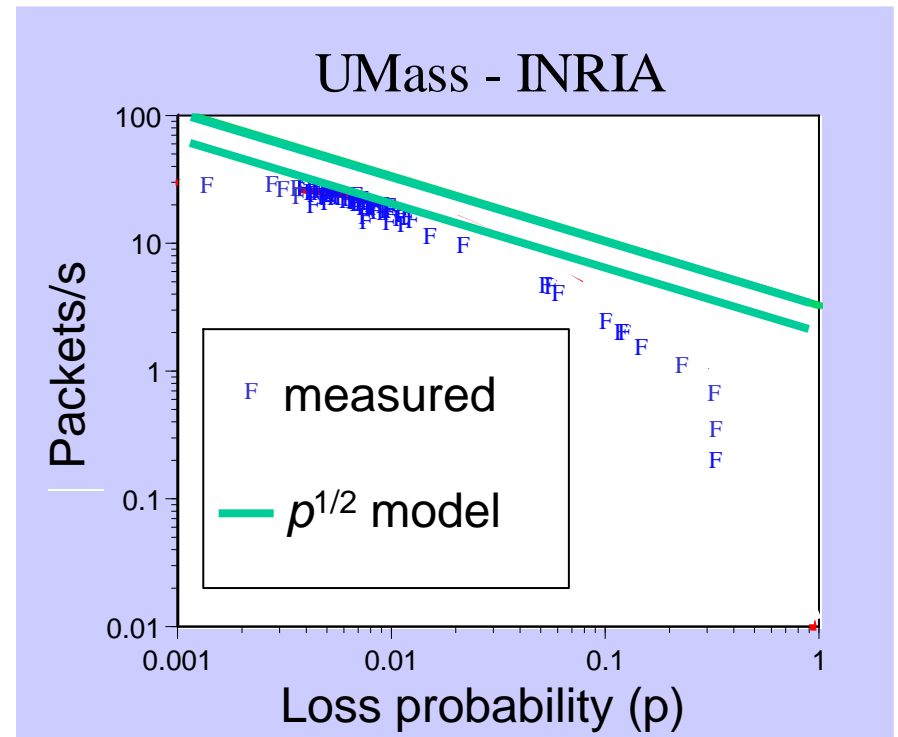
$$rate = k \times MTU / (R \times \sqrt{p})$$

p - loss probability; MTU - pkt length; R - round trip time; $k \sim 1.2-1.4$

consequence: 100s papers following party line

Theory vs measurement

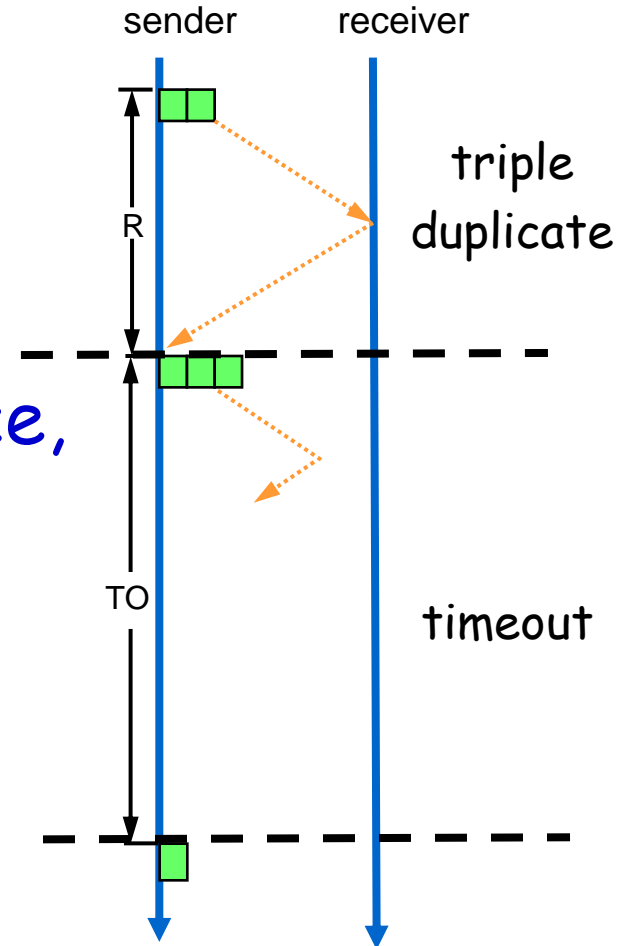
- agreement for small p
- differences for large p
- timeouts play role



PFTK formula


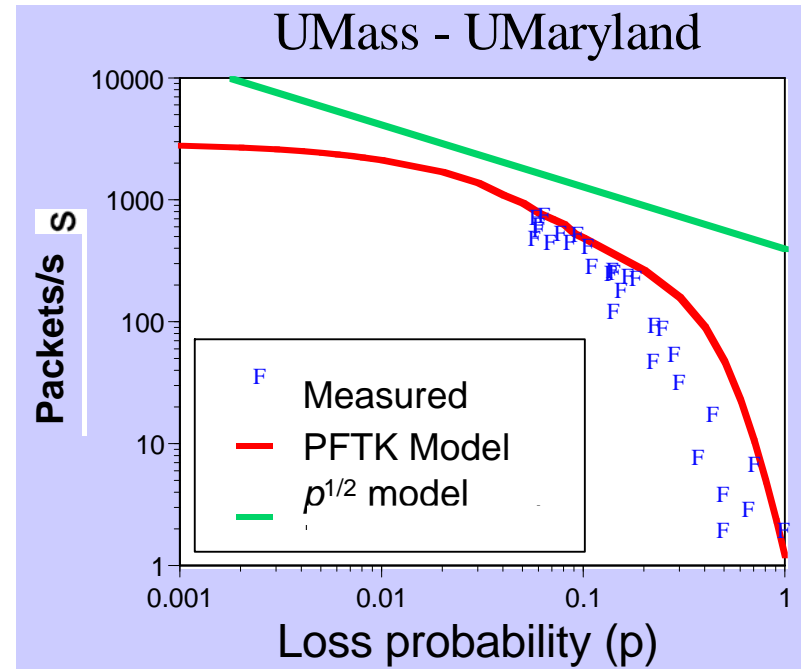
- accounts for timeouts
- drop tail loss assumption
- approximate analysis
 - independence, independence, independence

$$B(p, R) = 1/[R \sqrt{4p/3} + T_0 3(3p/2)^{1/2} p(1+32p^2)]$$



Validation

- ❑ validated against measurements
- ❑ corroborated by many studies
- ❑ insensitive to TCP version
- ❑ relatively insensitive to loss process (wireless)



Lessons

- ❑ be skeptical - don't accept party line without evidence
- ❑ follow discussion of important, timely topics -> good research

TCP friendliness redefined by IETF!

Modeling & validation go hand in hand

Fluid model (Misra)

- losses described by Poisson process

- data as fluid

- ⇒ $dW/dt = 1/R - \lambda W/2$

 - ⇒ square root p formula

- transient behavior

Fluid model - aftermath

- example of network agnostic model

- network \equiv loss

- suggests how to

- Poisson driven loss state (routers,

$$\frac{d\bar{W}_i(t)}{dt} = \frac{1}{\bar{R}_i(t)} - \frac{\bar{W}_i(t)\bar{W}_i(t-\tau_i)}{2\bar{R}_i(t-\tau_i)} P_i(t-\tau_i) \quad i=1,\dots,K$$

$$\frac{d\bar{q}_v(t)}{dt} = -\mathbf{1}\{\bar{q}_v(t) > 0\}C + \sum_{i=1}^N A_{v,i}(t)$$

$$\frac{d\bar{x}_v(t)}{dt} = \frac{\ln(1-\alpha)}{\delta} \bar{x}_v(t) - \frac{\ln(1-\alpha)}{\delta} \bar{q}_v(t)$$

set of differential equations describing flows, routers

Optimization-based Framework

Routing, congestion control: a resource optimization problem:

- allocate resources (e.g., bandwidth) to optimize some objective
- practical routing/congestion control as distributed asynchronous implementations of optimization algorithm
- foundational work
 - Gallager - routing
 - Kelly - congestion control



Model: Congestion Control

□ network: links $\{l\}$, capacities $\{c_l\}$

□ sources $S: (L(s), U_s(x_s)), s \in S$

○ $L(s)$ - links used by source s

○ $U_s(x_s)$ - utility, strictly concave increasing function of source rate x_s



Optimization Problem

$$\max_{x_s \geq 0} \sum_s U_s(x_s)$$

"system" problem

$$\text{subject to } \sum_{s \in \mathcal{S}(l)} x_s \leq c_l, \forall l \in L$$

- maximize system utility
- constraint: bandwidth used less than capacity

distributed algorithm?

Distributed algorithm

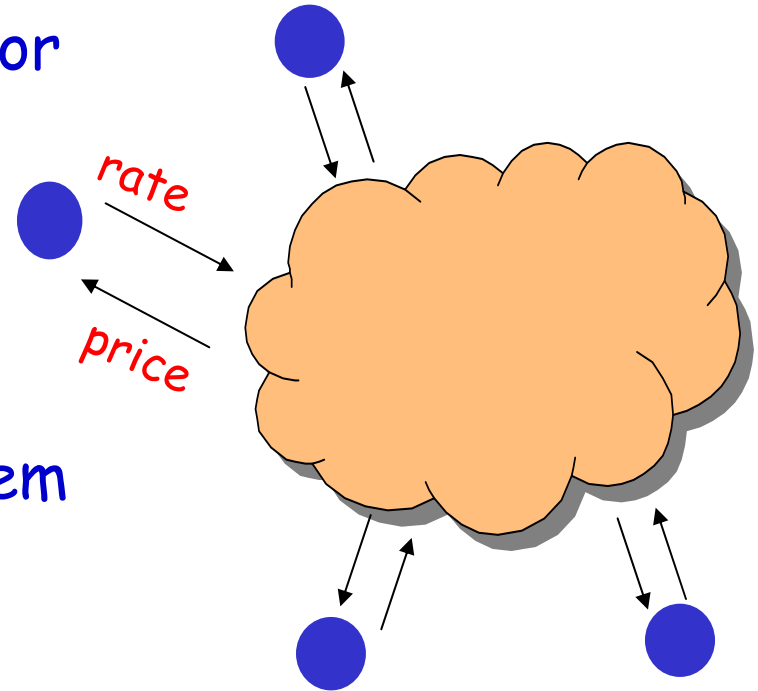
- optimization problem decouples into (Kelly):
 - greedy optimization problem for every session

$$\max U_s(x_s) - p_s x_s$$

- price p_s given by network as function of rates
- provides soln. to system problem
- for TCP, price \equiv loss prob.

what does TCP optimize?

utility function: $-1/(\text{RTT})^2 x$



Comments

- TCP is optimal
- TCP no longer a black box
- TCP SETI for large problems?

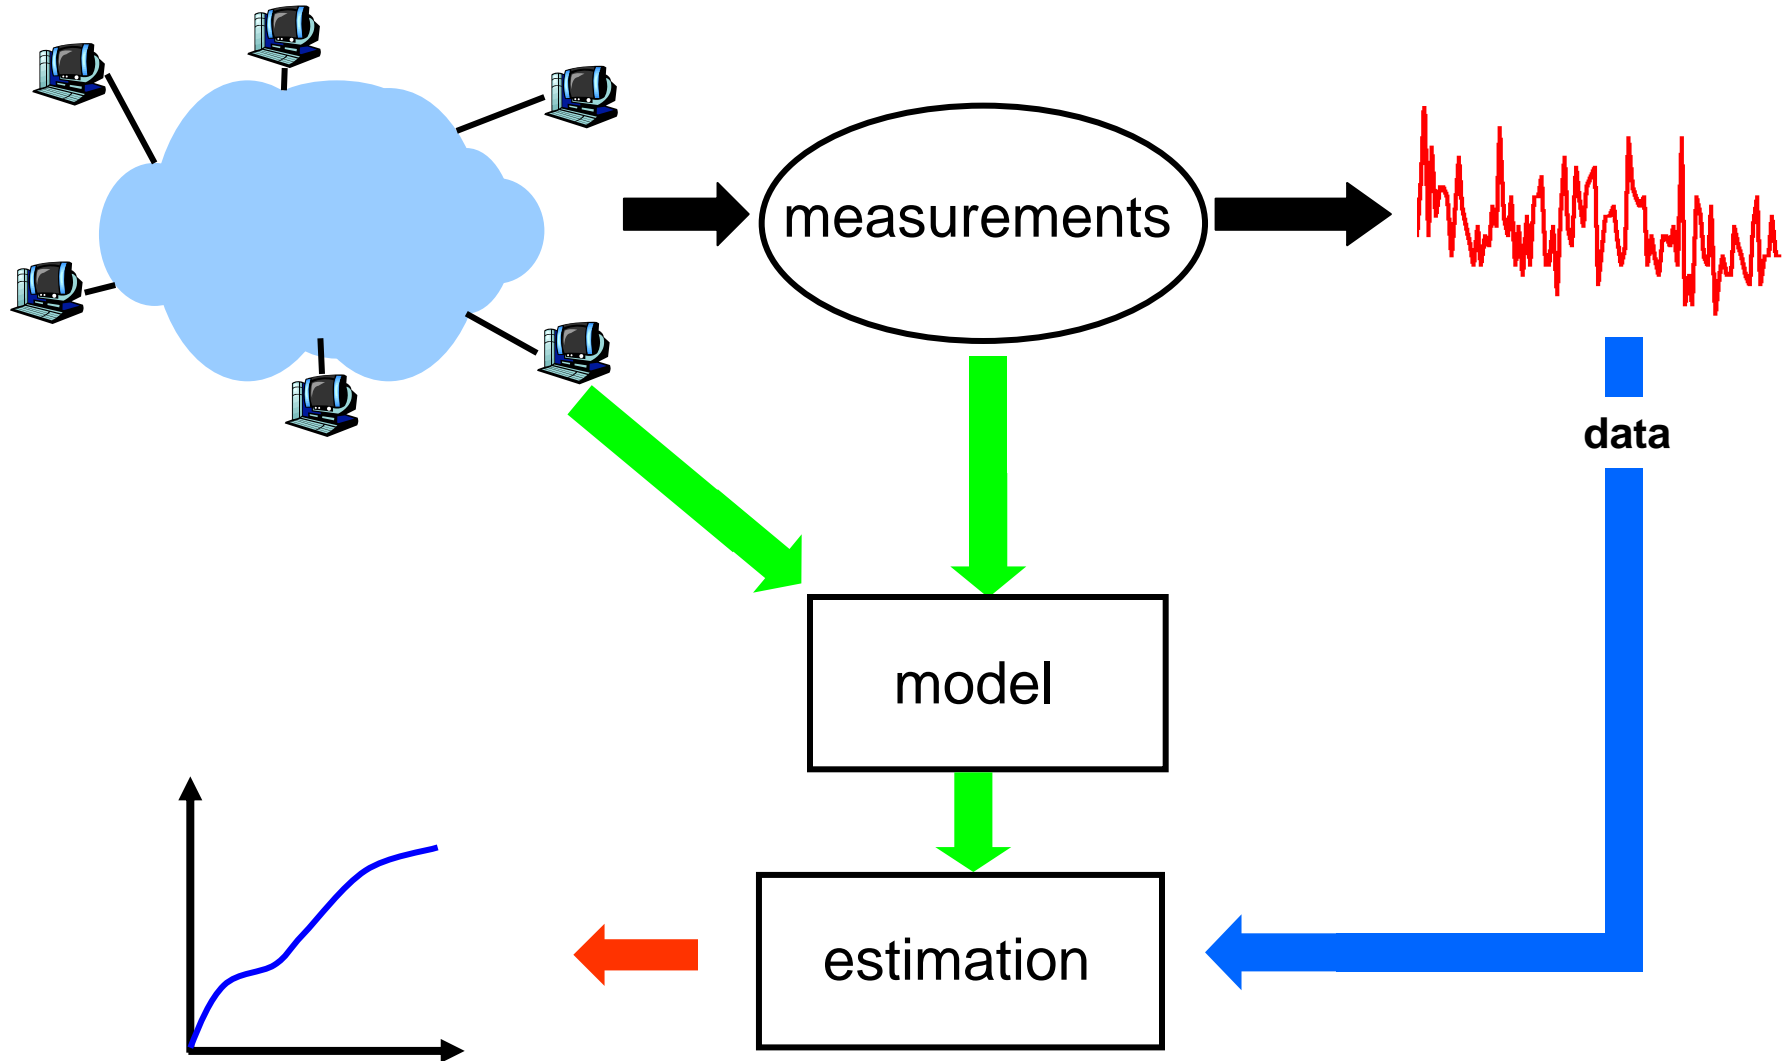
Aftermath

- new cc algorithms for high speed, small buffers
- application to other resource allocation problems - especially wireless networks
- extensions to multicast
- extensions to routing/rate control

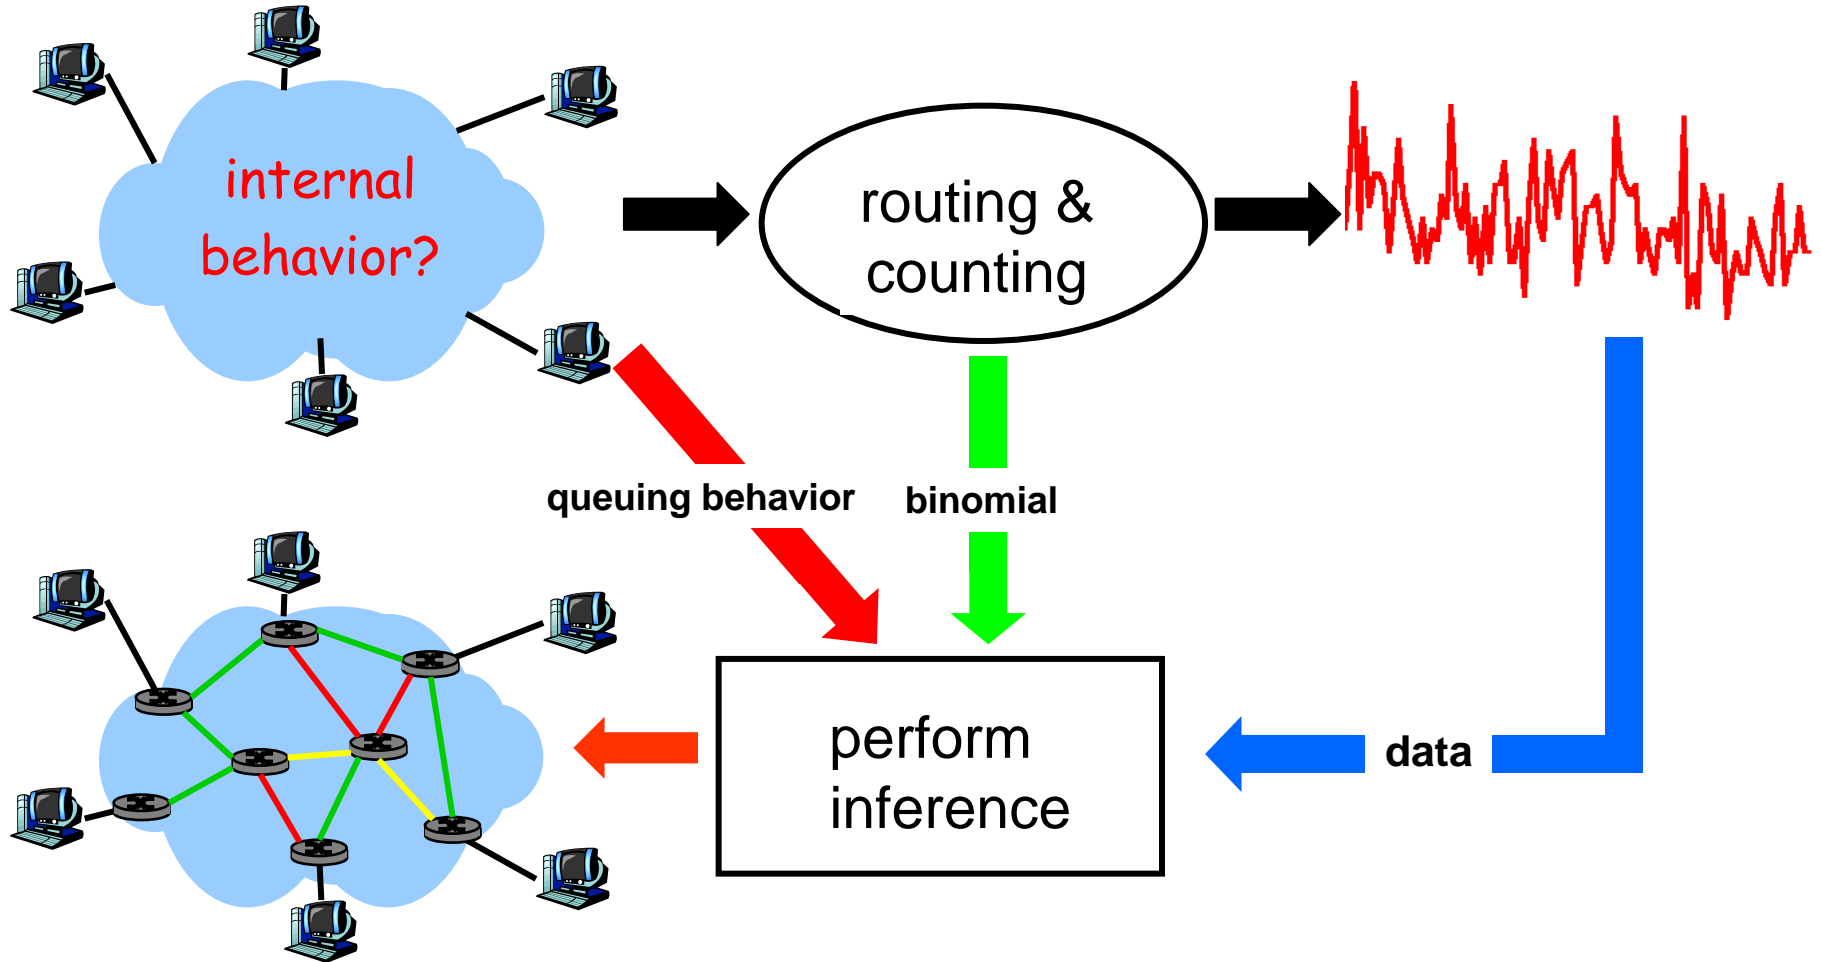
- fluid models right level of abstraction
- Poisson critical model building block

Modeling and measurements


Model-based measurement



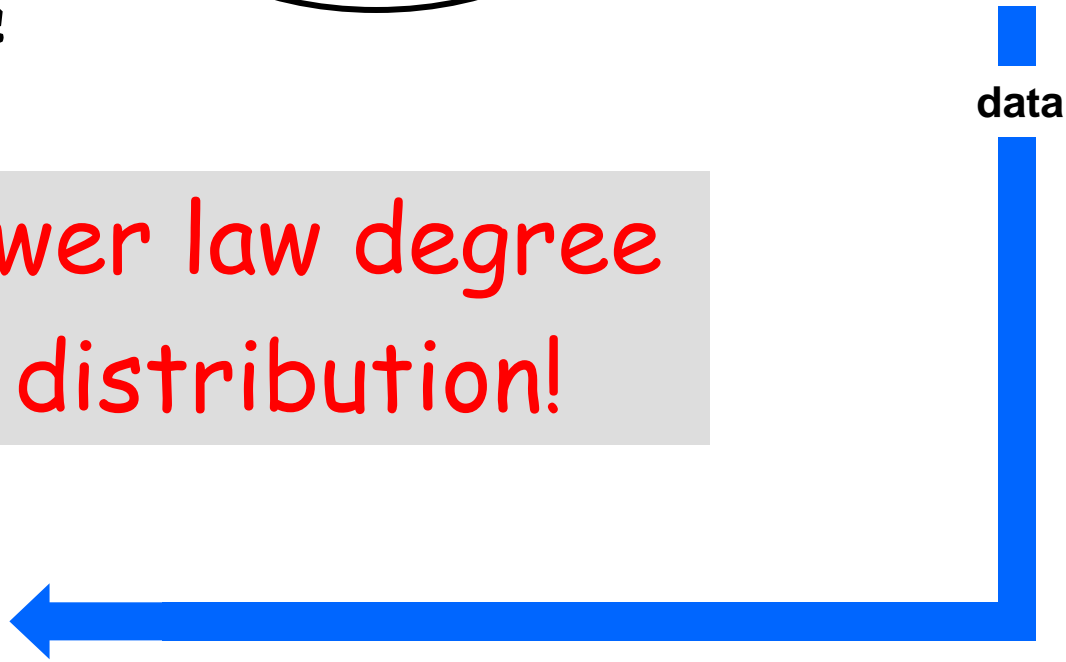
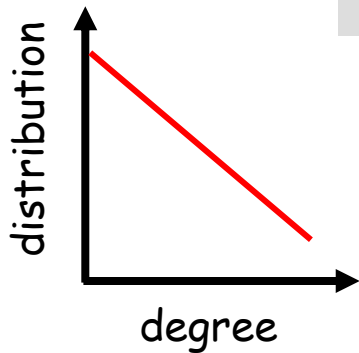
Network Tomography (MINC)



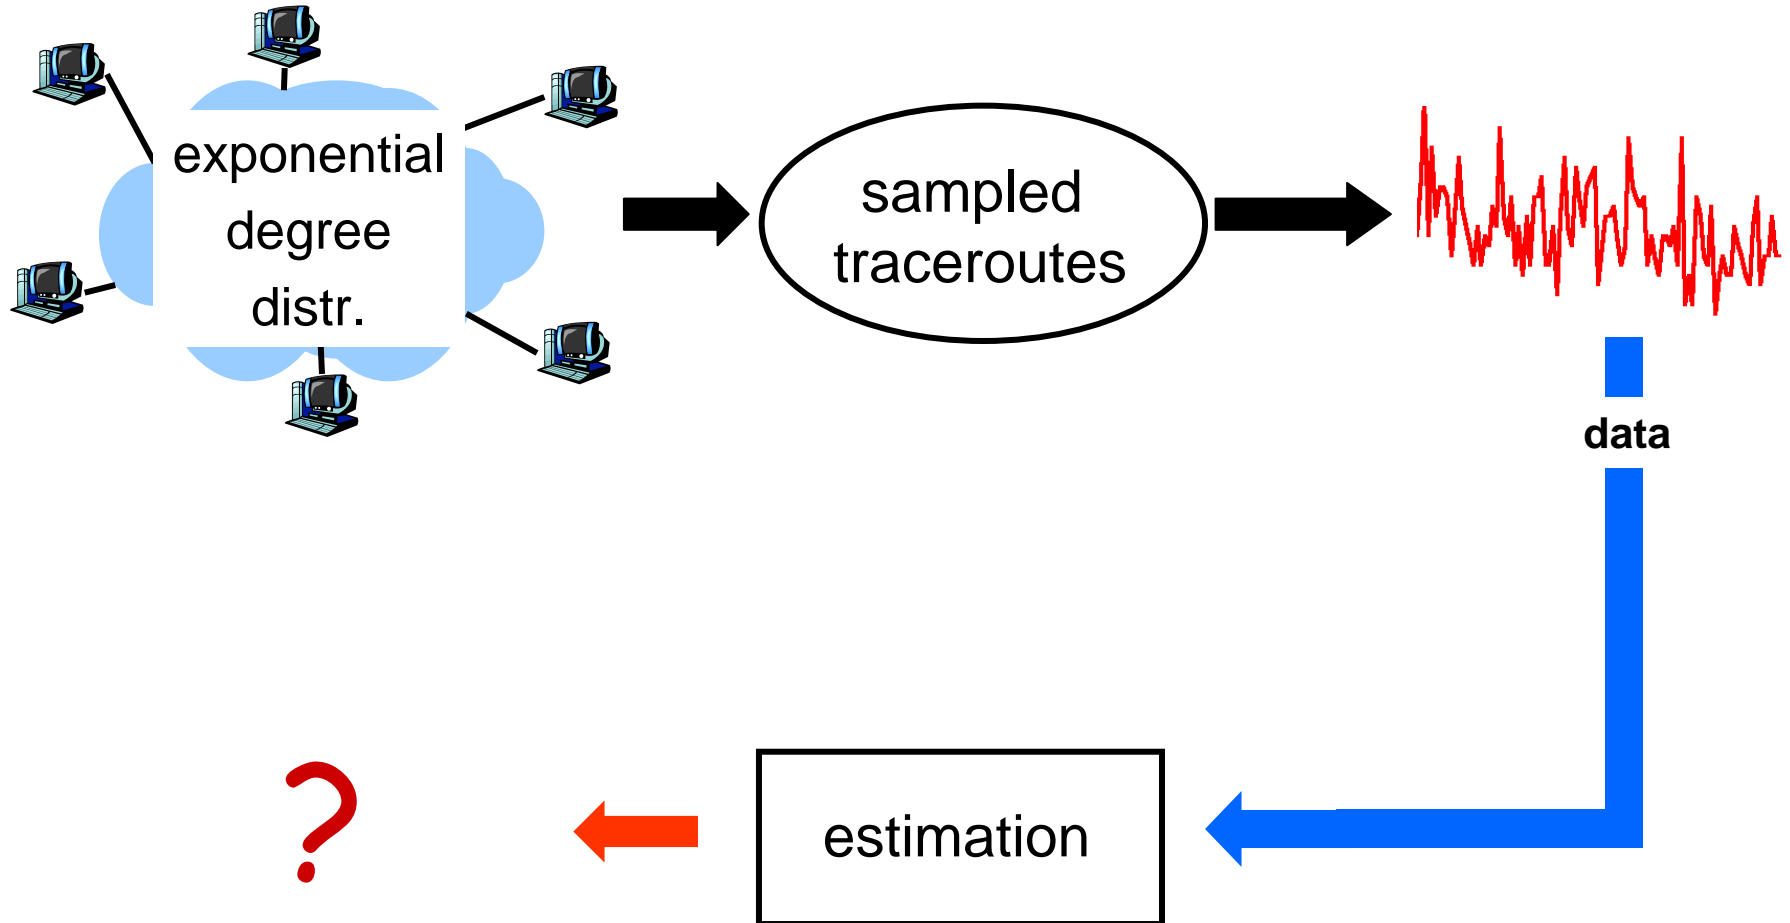
Topology inference



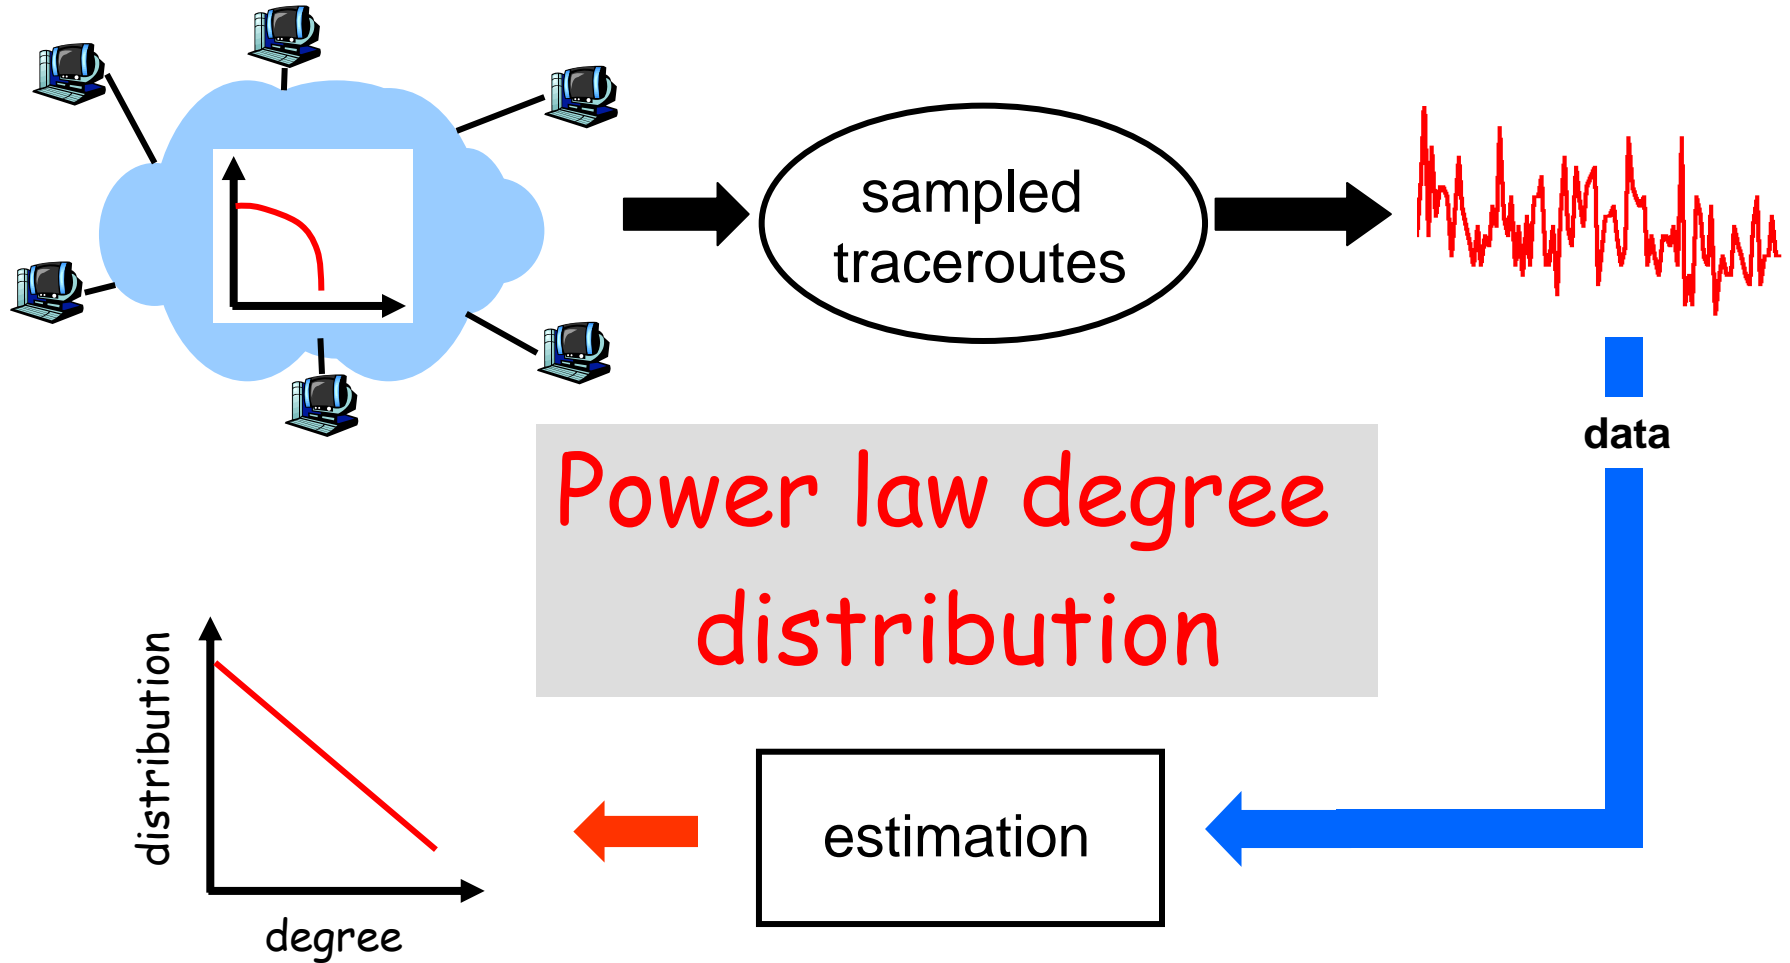
Power law degree distribution!



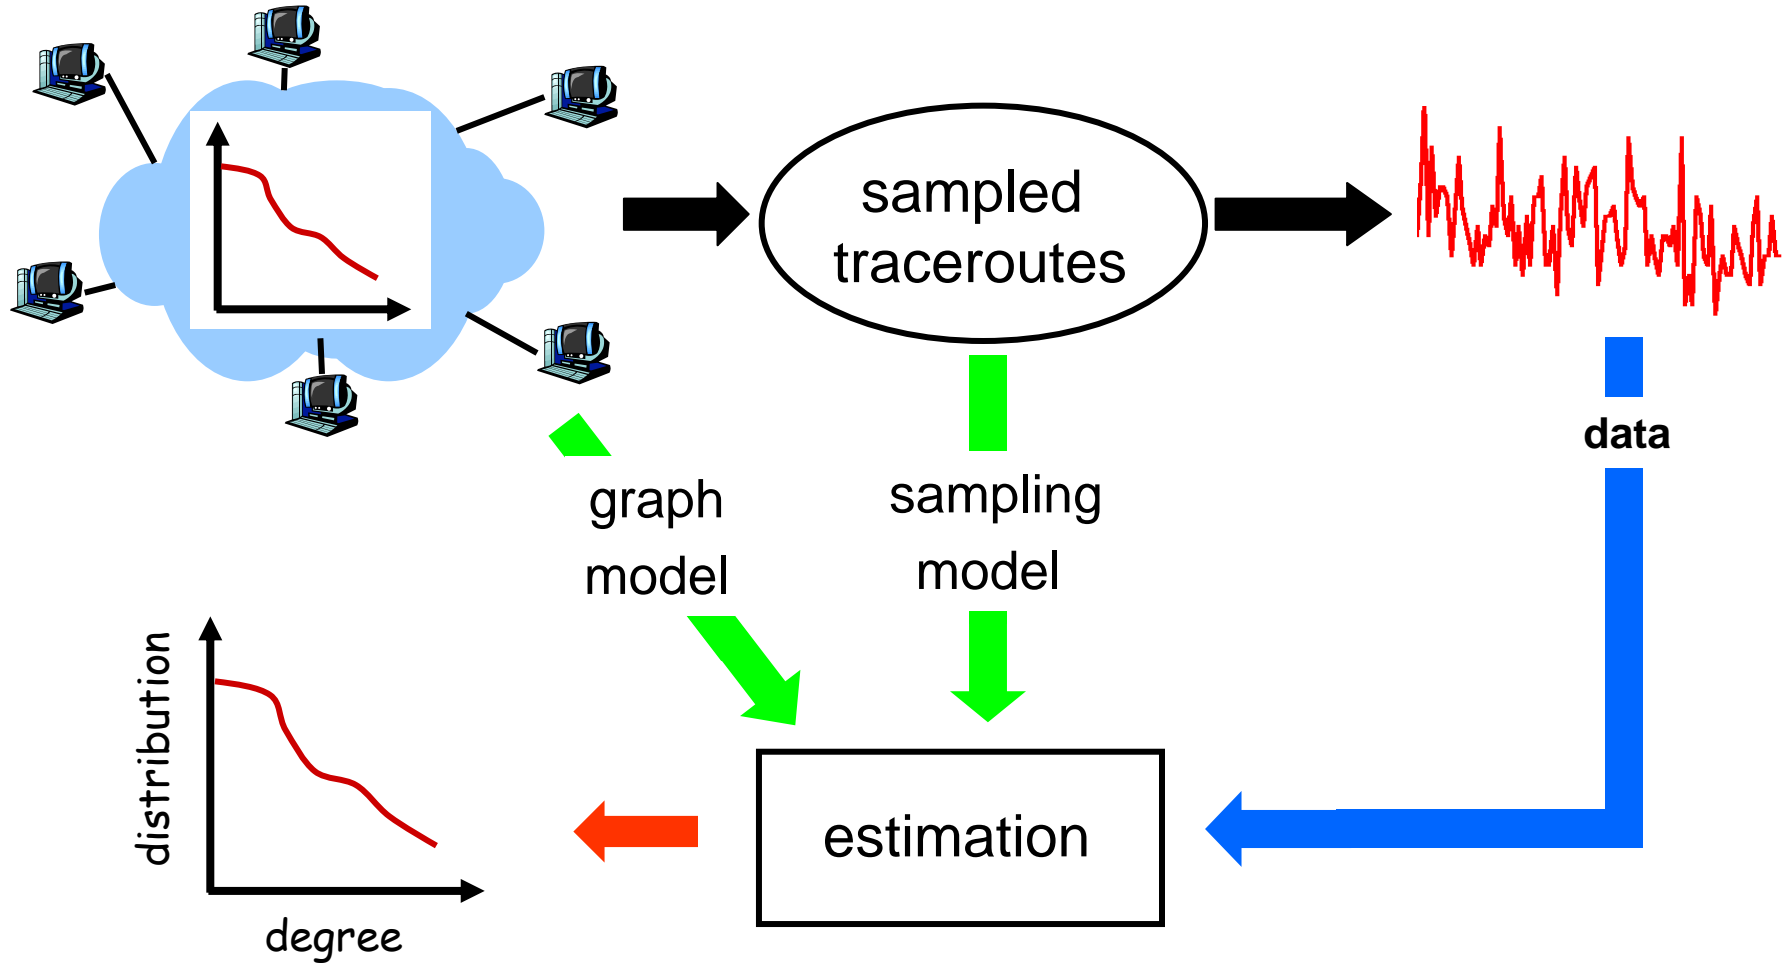
Topology inference



Topology inference



Topology inference



□ other examples

- P2P node up time distributions
- inter-contact times in DTNs

□ models can/should be used to identify bias (and often correct)

Lessons

- many reported measurement results are incorrect due to unintended biases

(not in SIGCOMM of course)

⇒ observations/conclusions may be wrong

- unintended biases introduced due to lack of model
- there **always** exists model behind measurements
- model needs to be explicit and precise

Parting thoughts: pursuing a PhD

- ❑ don't
try to develop new applications
- ❑ do
try to enable existing and new applications
- ❑ application disruption comes from outside
community
- ❑ technology disruption from within community

Parting thoughts: pursuing a PhD

- take my talk (and others) with a great deal of skepticism
- talk to other disciplines
- ear to the grapevine

... by the way, you can make a
living modeling

The end

Thanks!

Slides (will be) available at
<http://gaia.cs.umass.edu/towsley/sigcomm08.pdf>