Digging in to Ground Truth in Network Measurement

John Heidemann
University of Southern California / Information Sciences Institute
at the TMA PhD school
2017-06-19

The Internet and the Cave

(discussion)

imagine prisoners in a cave, chained to the wall

they cannot see the real w
instead only shadows o

(shadows of objec
not even of the rea

what is real?
the shadows? the objects th
the world above that inspire

what is our cave?
* is ripe atlas the cave
- n
- think about some pre-conceived idea
- cave = traceroute, shadows=output

shadows? the measurements we take
objects? the ground truth real world? the devices we do not see
the people holding the objects?
the companies (or agencies) managing (or manipulate)

Welcome to Maynooth!

The Internet and the Cave

- at one time the Internet fit on a napkin
- those days are long pa
  - many networks: >4M /24s
  - many computers: ~800M on public internet
  - many protocols
- what to do?

[Jon Postel, Dec. 1969]

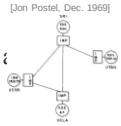

Plato's Allegory of the Cave

imagine prisoners in a cave, chained to the wall

they cannot see the real world, instead only shadows of objects

(shadows of objects, not even of the real things!)

what is real?
the shadows? the objects that cast them?
the world above that inspired them?

The Internet and the Cave

- at one time the Internet fit on a napkin
- those days are long past...
  - many networks: >4M /24s
  - may computers: ~800M on public internet
  - many protocols
- what to do?

[Jon Postel, Dec. 1969]

[map by CAIDA; data from Cheswick and Burch; 2000]

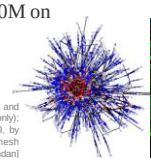

[Cable and Wireless (only); 1999, by Ramesh Govindan]

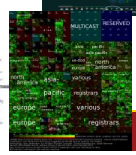

[Internet Census, USC/ISI, taken since 2006; this: 2017-02]

The Internet and the Cave

[by Markus Maurer, from Wikipedia.org]

imagine prisoners in a cave, chained to the wall

they cannot see the real w
instead only shadows of ob

(shadows of objects, not even of the real thi

what is real?
the shadows? the objects that cast the world above that inspired them

researchers are "chained" limited in what we measure

some things we can measure— but we see incomplete shadows

we should make inferences about the objects behind what we measure

we must try to imagine ideal future networks (better than we have)

Outline

- intro: Plato's cave
- **what do we want?**
- 4 case studies and 5 ground truths
- conclusions

What To Measure?

(discussion)

- how DNS resolvers are selecting?
- anomalies in traffic
- discovering structure in the address space and in routers and links that hook them up
- congestion on links IXPs
- protocol performance (QUIC vs. TCP, etc.)
- malicious queries in applications over the Internet
- deployment of new features, constraints and bugs

What To Measure?

(my take)

- topology
  - core (routers and links) and edges (hosts)
  - relationships: inter-AS relationships, AS-to-orgs
- size and capacity
  - numbers of end-systems, routers
  - amount of traffic
  - capacity of pipes and interconnection points
- traffic and applications
  - [illegible] QoE)
- reliability
  - packet loss, outages, censorship

Established Research Topics

what

- topology
  - core (routers and links) and edges (hosts)
  - relationships: inter-AS relationships, AS-to-orgs
- size and capacity
  - numbers of end-systems, routers
  - amount of traffic, capacity of pipes and interconnection points
- traffic
  - classification, trends
  - quality-of-experience
- reliability
  - packet loss, outages, censorship

how

- traceroute
- ping
- BGP peering (RouteViews)
- traffic analysis: HTTP, TCP, NTP
- (wireless stuff, also)
- platforms:
  - RIPE Atlas, CAIDA Ark, PlanetLab, private platforms
  - testbeds and emulation: DETER, Mininet private
  - from clients: Netlyzr, apps, Google ads
  - simulations: ns-2, ns-3, OpNet, custom

Defining Ground Truth

- goal: is what we measure correct?
- ground truth: **defines** what is correct
  - but sometimes it is incomplete
  - often unobtainable

but never forget that it exists; we must strive for it (there is an "outside the cave")

Elusive Ground Truth
(discussion)

- consider measuring height
  - ruler measured in cm: says h = 180cm
  - true height with ruler with infinite precision: h = 180.340cm
- is that true?
  - limitations on how accurately you can measure
  - you're taller in the morning
  - (is meter well defined)

Elusive Ground Truth
(my take)

- consider measuring height
  - ruler measured in cm: says h = 180cm
  - true height with ruler with infinite precision: h = 180.340cm
- is that true?
  - heights actually varies by around 1cm each day
  - even if true now, not true in 6 hours
- sometimes the truth varies; sometimes *no single* truth ever exists

Can we Fix Elusive Ground Truth?
(discussion)

- heights actually varies by around 1cm each day
- how to fix?
  - defining high parameters carefully
  - compute and report an average, measure multiple times
  - report error rates
  - we took height at 2:30pm

Can we Fix Elusive Ground Truth?
(my take)

- heights actually varies by around 1cm each day
- how to fix?
  - could define height more precisely
    - height must be measured at 9am
  - could define height as a range or distribution
    - 180 +/- 1cm
    - an "envelope of truth"
  - or maybe we shouldn't measure height? (it's non-stationary)
- both approaches have their place
  - range seems easier
  - WHY are you measuring?

=> "truth" is not always one value

Aside: Truth is Often an Envelope

- TCP performance as a function of loss (p) and RTT?
- bitrate = RTT^{-1} sqrt(3/(2bp)
- but there are many, different implementations
  - BSD, Linux, Windows
  - Vegas, FAST, CUBIC, BRR

where does this matter?

- validating TCP in ns-2
- TCP friendliness: contestion contorl that ties to be "like" TCP
- future TCPs (CUBIC, BBR, etc.)
- future *other* protocols (QUIC, etc.)

Outline

- intro: Plato's cave
- what do we want?
- **4 case studies and 5 ground truths**
- conclusions

Where to Get Ground Truth?
(discussion)

- DPI for traffic classification
  - modulo encryption
- SNMP to get data on congestion
- friendly network operators
- there are tradeoffs in privacy and propritariness
- testbeds
  - complete control: good: you have control, bad: you set it up, so you have know the patarameters and assumptions

Where to Get Ground Truth?
(my take)

- from the network operator
- from testbed experiments
- from simulations
- as seen in prior results

Case Study 1: Network Topology Mapping

- question: can we map ISPs, or the whole Internet?
- early work
  - "Heuristics for Internet Map Discovery" Govindan and Tangmunarunkit, INFOCOM 2000
  - "Measuring ISP Topologies with Rocketfuel", Spring, Mahajan, Wetherall, SIGCOMM 2002
  - "Macroscopic analyses of the infrastructure: measurement and visualization of Internet connectivity and performance", Huffaker, Fomenkov, Moore, Claffy, PAM 2001
- recent work
  - ITDK-2016 from CAIDA

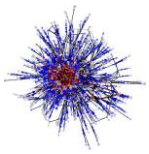

[Cable and Wireless (only); 1999, by Ramesh Govindan]

Ground Truth 1: from the Operator

- ground truth: use a few research networks
  - "Heuristics for Internet Map Discovery" Govindan and Tangmunarunkit, INFOCOM 2000
    - 2 regional networks: Los Nettos and Calren2
  - "Measuring ISP Topologies with Rocketfuel", Spring, Mahajan, Wetherall, SIGCOMM 2002
    - 3 (private) ISPs gave qualitative results
- (the Huffaker et al 2001 paper did not evaluate correctness)

Digression: Defining "Correct"
(discussion)

- ground truth: defines what is correct
- but what does "correct" mean?
- unambiguous
- something that fits the purpose of this experiment
- optimum… algorithms can prove they're the best possible
- scalable
- has high probability of being reproduced
- we can compare the

Digression: Defining "Correct"
(my take)

- ground truth: defines what is correct
- but what does "correct" mean?
- from info theory
  - precision: is what you claim always true?
  - recall: is what you claim the *complete* truth?
  - accuracy: is what you claim and reject both correct

the Confusion Matrix: Formalizing Correct

	prediction positive	prediction negative
actual positive	true positive (TP)	false negative (FN)
actual negative	false positive (FP)	true negative (TN)

recall := TP / (TP+FN)
how much do we say?

precision := TP / (TP+FP)
is what we say true?

accuracy := (TP+TN) / Population
Population := TP+FP+FP+TN

Beware papers that talk only about "correctness" without defining at what *metric* of correctness.

They often focus only on *precision* and ignore *recall*—what they say is true, but they may miss a lot (and not know it).

Back to Ground Truth from the Operator

- "Heuristics for Internet Map Discovery" Govindan and Tangmunarunkit, INFOCOM 2000
  - 2 regional networks: Los Nettos and Calren2
  - "we found all routers and all but 1 link in each network"
  - precision *and* recall are both high
    - they don't give counts, so we guessed
- "Measuring ISP Topologies with Rocketfuel" Spring, Mahajan, Wetherall, SIGCOMM 2002
  - 3 (private) ISPs gave qualitative results
  - "did we miss any pops"? "say no"
  - "did we miss any links?" "say no but we had 2 extra links
  - "how many access routers did we miss?" (a) none, (b) some, (c) extra from other ASs
  - "how many customer routers did we miss?" none will say, two say they do not know
- challenges
  - most commercial networks: topology is proprietary
  - those who share (academic nets) may not be representative
  - they don't always know the truth!

other ground truth about topology

- academic network
  - Los Nettos: LA-regional consortium
  - Calren2: California academic net
- or
  - Internet2
  - GEANT
  - pro:
    - open topologies
  - con:
    - not profit based, maybe not optimized the same way
    - Tim Griffin had a paper comparing research nets to commercial nets: "The Interdomain Connectivity of PlanetLab Nodes" PAM
- or testbeds
  - where you have to fill in all the details

Case Study 2: Edge Address Activity

- question: can we identify all active IP addresses
- early work:
  - "Census and Survey of the Visible Internet", Heidemann, Pradkin, Govindan, Papadopoulos, Bartlet, Bannister, ACM IMC 2008

[Internet Census, USC/ISI, taken since 2006; this: 2017-02]

Ground Truth 2: from *Our* Enterprise

(for Case Study: Edge Address Activity)

- ground truth from *Our* Enterprise
  - University of Southern California
- advantage
  - can talk to the operators (we know them)
  - can apply **multiple measurement methods**
    - test active probing (**ICMP**, the new method being considered)
    - against other kinds of active probing (**TCP SYNs**)
    - and against **passive traffic** analysis

Sources of Error for Edge Address Activity

- variance
  - measurement location: *doesn't matter; normal error*
  - sampling error:
    - *can predict from theory*
    - *function of probe frequency*
    - *surveys within 0.4% (with 95% confidence)*
  - births/deaths during survey: *estimate in paper*
  - probe type (ICMP vs. TCP): *ICMP consistently more complete*
- overcounting
  - routers and multi-homed hosts: *estimated at <6% in paper*
- undercounting
  - probe loss: *random due to probe order; use 1-repair process to recover single losses in survey*
  - firewalled hosts: *coming up*

⇒ warning: if error was always pro or con, can set a bound
⇒ (but no: it can go either way)

Enterprises are Not Perfect

- USC has ~89k IPv4 addresses
- management is partially decentralized
  - *no one* has complete, current status of all addrs
- current status is sensitive
  - anti-file sharing requests: who was using IP *x* and time *t*?
  - will not share DCHP information with researchers
- operator knowledge ages
  - address use changes over time; tracking is incomplete
- the **network operators don't know the ground truth**
  - **big is hard!** (even where big == one enterprise)

Evaluating at USC (Our Enterprise)

USC Survey (82k hosts)

category:		any	active
addresses probed	81,664		
non-responding	54,078		
responding any	27,586	100%	
ICMP or TCP	19,866	72%	100%
ICMP	17,054	62%	86%
TCP	14,794	54%	74%
Passive	25,706	93%	
ICMP only	656		
TCP only	1,081		
Passive only	7,720		

define **ground truth** as responds to any (TCP, ICMP, or passive traffic)

Census is *incomplete*, but can *estimate* error => recall is 62%

Evaluating at USC (Our Enterprise)

USC Survey (82k hosts)

category:		any	active
addresses probed	81,664		
non-responding	54,078		
responding any	27,586	100%	
ICMP or TCP	19,866	72%	100%
ICMP	17,054	62%	86%
TCP	14,794	54%	74%
Passive	25,706	93%	
ICMP only	656		
TCP only	1,081		
Passive only	7,720		

different **ground truth,** active probing only

Census is still *incomplete*, but can *estimate* error => recall now 86%

Advantages at *Your* Enterprise

- getting all the local traffic
- combining *passive* and *active* to get bigger view
- still not perfect
  - passive at edge misses hosts with local-only traffic
    - printers, internal telephones, etc.
  - hard to get all traffic at the edge
    - modems? internal caches? direct peering?
  - and operators don't know everything
  - and… how do we know USC is representative of the Internet as a whole?

Ground Truth 3: Random Sampling

(for Case Study: Edge Address Activity) (discussion)

- take a *random sample* of all Internet addresses
- pro:
  - could do it repeatedly
  - there is no bias
- con:
  - use of IP address space is not equally distributed
    - so many of what we pick might not be used
  - and some parts are reserved for private use
  - don't know if they're in use
    - [illegible]
    - fix: probe with TCP and ICMP

Ground Truth 3: Random Sampling

(for Case Study: Edge Address Activity) (my take)

- take a *random sample* of all Internet addresses
- pro:
  - should be unbiased (by definition)
- con:
  - what is their truth?
  - what about rare parts of the Internet?
    - 1M addresses might only get 10 servers (!), or 10 users in developing world, or…

Random Sampling for Active Addresses

random addresses (1M hosts)

category:		active
addresses probed	1,000,000	
non-responding	945,703	
responding either	54,297	100%
ICMP	40,033	74%
TCP	34,182	62%
both ICMP and TCP	19,918	
ICMP only	20,115	
TCP only	14,264	

only have *weaker* **ground truth,** active probing only

Census is still *incomplete,* but can *estimate* error
=> recall now 74%
=> **confirms prior results**

Case Study 3: IP Alias Resolution

- question: when are IP addresses in traceroutes the same device?
- early work
  - " H euristics for I nternet M ap D iscovery" Govindan and Tangmunarunkit, INFOCOM 2000
    - 2 regional networks: Los Nettos and Cabren2
  - " M easuring I S P T opologies with Rocketfuel" ,S pring,Mahajan, Wetherall, SIGCOMM 2002
    - 3 (private) ISPs gave qualitative results
- recent work:
  - " F ixing A ly Growing Pains with Velocity M Sherwood, Spring, IMC 2008
  - " P rimitives for A ctive I nternet T opology M apping Toward High-F requency C haracterization" ,B everly,B erger, Xie, IMC 2010
  - " I nternet S cale I P A lias R esolution T echniques" ,K eys,C C R 2010
  - " M apping P eering I nterconnections to F acility" , Giotsas, Smaragdakis, Huffaker, Luckie, Claffy, CoNEXT 2015 (also goes much further)

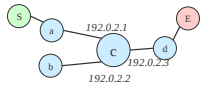

IP Alias Resolution Challenge: **c** is a multi-homed router w/192.0.2.1, .2, and .3 traceroutes from S to E could return *any* or *all* of these how to tell they are all **c**

Ground Truth 4: Prior Work

- can compare to prior published work
  - or get and run prior code
- but can compare to prior results
- challenge:
  - errors can propagate
  - " betterthanbefore" gives noclueabout" good"

Ground Truth in IP Alias Resolution

- early work
  - " H euristics for I nternet M ap D iscovery" Govindan and Tangmunarunkit, INFOCOM 2000
    - does not exlicitly validate alias resolution (!)
  - " M easuring I S P T opologies with Rocketfuel" ,S pring,Mahajan, Wetherall, SIGCOMM 2002
    - compares to prior work (Mercator) and DNS names
- recent work:
  - " F ixing A ly Growing Pains with Velocity M
    - compares to prior work: known ground truth dataset (from Mercator) and Rocketfuel
  - " P rimitives for A ctive I nternet T opology M apping Toward High-F requency C haracterization" , Beverly, Berger, Xie, IMC 2010
    - focuses on performance, not validation of results
  - " I nternet S cale I P A lias R esolution T echniques" ,K eys,C C R 2010
    - validates against datasets from 5 academic networks
  - " M apping P eering I nterconnections to F acility" , Giotsas, Smaragdakis, Huffaker, Luckie, Claffy, CoNEXT 2015 (also goes much further)
    - validates against 2 CDNs, DNS records, BGP community strings, IXPs

Case Study 4: Effects of Cable Cuts

- question: what are the effects of breaks in undersea cables on the countries they serve?
- work-in-progress (tech report)
  - " A Holistic Framework for Bridging Regional Threats to User QoE" Cai, Heidemann, Willinger, ISI-TR-687, 2013

Ex: the SeMeWe-4 cable was cut near Singapore on 2012-06-06. What is the impact on the Internet?

Challenges in Country-Level Internet Evalutaion

- specific question: how does cable outage affect YouTube in countries served by the cable?
- challenges:
  - multiple YouTube sites
  - multiple ISPs in each country
  - unknown routing, peering, ISP capacities
  - unknown other traffic on links
- *yet understanding Internet fragility is critical!*

Ground Truth 5: Modeling
(discussion)

- idea: let's model the network as best we can
- pros:
  - simplifies the problems
  - can compare your results to alternatives, based on your knowledge
- cons:
  - simplifies the problems
  - but maybe alternatives that you consider are not right or missing

Modeling What Ifs

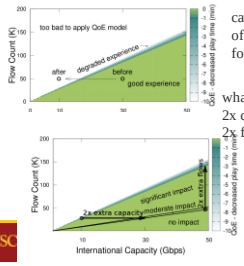

can evaluate likely outcome
of cable cut (-20Gb/s capacity)
for assumed traffic load (50k flows)

what-if:
2x capacity: always s
2x flows: right on ed

½ traffic defects
after failure?
now rest are ok!

Ground Truth 5: Modeling All Options
(my take)

- idea: let's model the network as best we can
  - look at *all possible parameters*
- pros:
  - can look at many parameters quickly
  - if all parameters give same result, have answer!
  - if *most* parameters give same result, answer is likely
  - worst case: provide *possible* outcomes, others (w/more info, or in future) can fill in
- cons:
  - can be lots of parameters!
  - each layer of model adds uncertainty
  - *not* ground truth, but all possible truths (many incorrect!)

Some Options for Ground Truth

- ask the opertors
- your enterprise
- random sampling
- prior work
- model all the things!
- (your ideas here)

Modeling Instead of Ground Truth: Cable Cuts

multiple layers of modeling; most layers are adaptive

physical -> link *-> network (policy)* *-> network (routing)* *-> transport -> app*

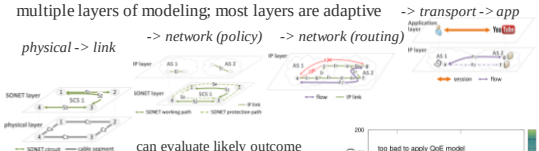

can evaluate likely outcome
of cable cut (-20Gb/s capacity)
for assumed traffic load (50k flows)

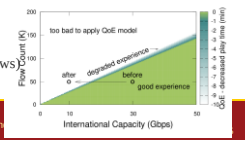

Outline

- intro: Plato's cave
- what do we want?
- 4 case studies and 5 ground truths
- **conclusions**

Plato's Allegory of the Cave

[by Markus Maurer, from Wikipedia.org]

imagine prisoners in a cave,
chained to the wall

they cannot see the real world,
instead only shadows of objects

(shadows of objects,
not even of the real things!)

what is real?
the shadows? the objects that cast them?
the world above that inspired them?

So What Is Real?
(from physical to abstract)

[by Markus Maurer, from Wikipedia.org]

the shadows?

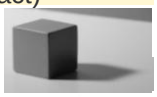

the objects that cast them?

the world above that inspired them?

an ideal world that *could* exist?

So What Is Real?
(misleading objects)

[Art by Red Hong Yi]

[by Markus Maurer, from Wikipedia.org]

the shadows?

the objects that cast them?

the world above that inspired them?

an ideal world that *could* exist?

perhaps a long time ago in a galaxy far far away...

So What Is Real?
(the truth we cannot directly see)

[Shadow picture form Quora post by Shyamala; diagram from wikipedia]

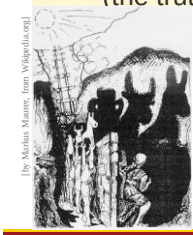

[by Markus Maurer, from Wikipedia.org]

the shadows?

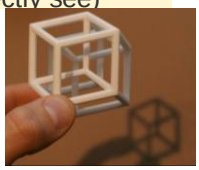

the objects that cast them?

the world above that inspired them?

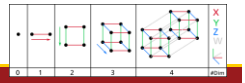

an ideal world that *could* exist?

Conclusions

- strive to search for **the truth**
  - don't stop at what you see
  - "best available data" today… can you do better tomorrow?
  - not not just what exists, but *what should be*
- use strong correctness (from info theory)
  - precision *and* recall, not just "correctness"
- be creative about ground truth
  - you can often dig it out, if you work
  - explore all possibilities if that's the best you can do