

Partial Reachability in the Internet Core: Definition, Measurement, and Application (Invited Talk)

USC/ISI Technical Report ISI-TR-742, released 2026-06-30

John Heidemann

University of Southern California

Information Sciences Institute and Thomas Lord Department of Computer Science

Marina del Rey, California, USA

johnh@isi.edu

This work is the abstract and slides of an invited keynote talk presented at ACM AINTEC 2025, the Asian Internet Engineering Conference at the Dr. Andrew L. Tan Data Science Institute, De La Salle University, Bonifacio Global City, Taguig, Philippines, on 2025-11-25.

Talk Abstract

Routing strives to connect the Internet, but failures can result in **partial reachability** in the Internet. This talk identifies partial reachability as an underappreciated area. We will show that partial reachability affects as many users as Internet outages, sometimes persisting for days or years.

We suggest three steps are critical to better understand partial reachability: We first present a **conceptual definition of the Internet core**, since, surprisingly, today there is *no clear definition of what the Internet is*. This definition is technical and neutral, based on connectivity, not authority. It enables us to understand *peninsulas* as persistent partial reachability, and *islands*, where networks are disconnected from the core.

SecoND, these definitions allow us to realize **algorithms that detect peninsulas and islands** reusing data from existing measurement systems that provide Internet-wide observations over years from thousands of observers.

Finally, we show **important practical applications of this work**. Our new analysis shows that partial reachability affects even more users than outages. We show how to improve RIPE Atlas sensitivity by removing measurement error and persistent reachability problems that otherwise that otherwise overwhelm individual events. Finally, we use our neutral definition to inform Internet policy, including questions about Internet sovereignty.

More Information

For more information, please see our technical reports “What Is The Internet? Partial Connectivity of the Internet Core” [1] and “Reasoning About Internet Connectivity” [2].

Speaker Biography

John Heidemann is the chief scientist for the Network and Cybersecurity Division of the University of Southern California/Information Sciences Institute (USC/ISI) and a research professor in the Thomas Lord Department of Computer Science at USC. At ISI he leads the ANT (Analysis of Network Traffic) Lab, developing privacy-aware methods to measure the Internet and improve network reliability, security, protocols, and critical services.

Acknowledgments

This work was done in collaboration with Guillermo Baltra, Yuri Pradkin, and Tarang Saluja.

This work is supported in part by NSF contract PIMAWAT/CNS-2319409, BRIPOD/OAC-2530698, and DARPA contract AQUARIUS.

This work is copyright ©2025 by John Heidemann and is released under a CC BY-NC-SA: Creative Commons Attribution-NonCommercial-ShareAlike license.

References

- [1] Guillermo Baltra and John Heidemann. 2023. *Detecting Partial Reachability in the Internet Core*. Technical Report arXiv:2107.11439v4. USC/Information Sciences Institute. doi:10.48550/2107.11439v4
- [2] Guillermo Baltra, Tarang Saluja, Yuri Pradkin, and John Heidemann. 2024. *Reasoning About Internet Connectivity*. Technical Report arXiv:2407.14427 [cs.NI]. arXiv. <https://ant.isi.edu/%7ejohnh/PAPERS/Baltra24b.html>

Partial Reachability in the Internet Core: Definition, Measurement, and Application

John Heidemann

joint work with Guillermo Baltra, Yuri Pradkin, Tarang Saluja

U. of Southern California / Information Sciences Institute and Thomas Lord Dept. of CS
for ACM Asian Internet Engineering Conference, AINTEC 2025

2025-11-25

This research is sponsored by NSF PIMAWAT/CNS-2319409, BRIPOD/OAC-2530698, and DARPA contract AQUARIUS/td. Prior work was supported by NSF and DHS HSAPRRPA Cyber Security Division. The views herein are those of the authors and do not necessarily represent those of DHS or the U.S. Gov't.



Copyright © 2025 by Heidemann
Release terms: CC-BY-NC 4.0 international



Partial Reachability in the Internet Core: Definition, Measurement, and Application

John Heidemann

joint work with Guillermo Baltra, Yuri Pradkin, Tarang Saluja

U. of Southern California / Information Sciences Institute and Thomas Lord Dept. of CS
for ACM Asian Internet Engineering Conference, AINTEC 2025
at the Dr. Andrew L. Tan Data Sciences Institute, Manila, Phillipines

2025-11-25

This research is sponsored by NSF PIMAWAT/CNS-2319409, BRIPOD/OAC-2530698, and DARPA contract AQUARIUS/td. Prior work was supported by NSF and DHS HSAPRRPA Cyber Security Division. The views herein are those of the authors and do not necessarily represent those of DHS or the U.S. Gov't.

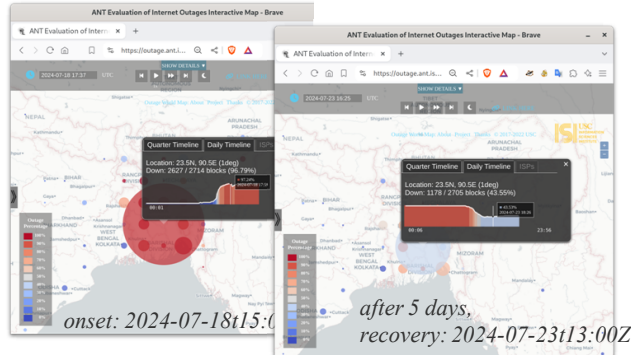


Copyright © 2025 by Heidemann
Release terms: CC-BY-NC 4.0 international



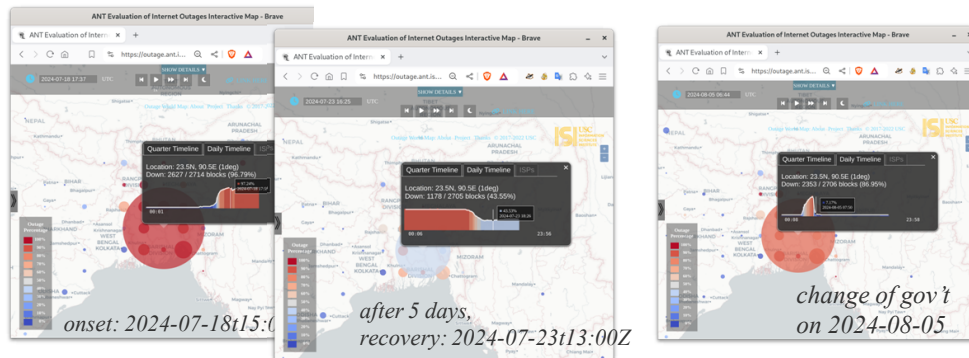
Outage detection finds things... Bangladesh's Civil Unrest, July 2024

In July 2024, a new law led to civil unrest and Bangladesh shut down their Internet *for five days*. In Aug., they did it again when the gov't resigned.



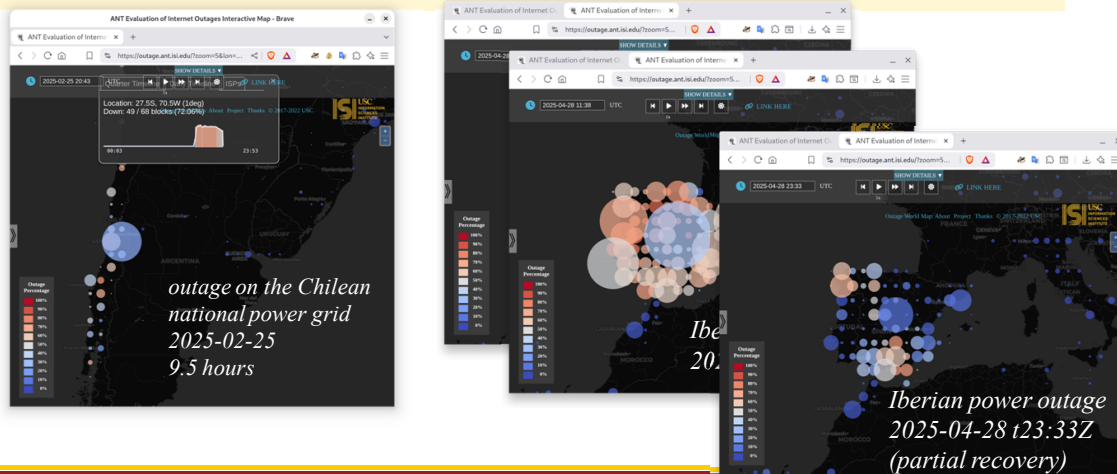
Outage detection finds things... Bangladesh's Unrest Change of Gov't, Aug. 2024

In July 2024, a new law led to civil unrest and Bangladesh shut down their Internet *for five days*. In Aug., they did it again when the gov't resigned.



Outage detection finds things...

Power Outages in Chile (2025-02) and Iberia (2025-04)



Outage Detection (Un)Certainty

- this data is from **Trinocular**
 - 6 sites: LA, Denver, Washington/DC, Tokyo, Netherlands, Greece
- tested against several other systems
- and tested against *itself*... we compare the 6 sites
 - we vote to get a **single** answer
- but is it *correct*?
 - ...but what about when our six sites *don't* agree?
 - measurement error? bugs?
 - or something else?

We Need Better Theory to Explain Partial Outages

- outage detection has struggled with partial outages for years...
 - ThunderPing (2011) *hosed* state—some pings work and some don't
 - Trinocular (2013): precision improvement small diffs are timing
 - CDN (Richter, 2018): maybe it's outages in *part* of the /24?
 - Trinocular (2019): majority voting for long disagreements
 - other systems ignore conflicting data
- common problem:
no theory explaining disagreement in reachability observations

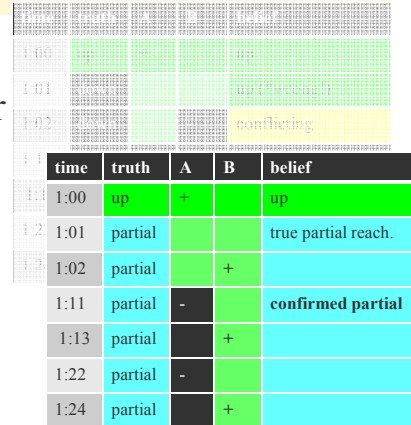
Premise 1: Partial Reachability is Real

- agreements are typical
- *some* disagreement is measurement error
 - ex: timing variation between sites (transient disagreement)

time	truth	A	B	belief
1:00	up	+		up
1:01	down			up (*wrong!)
1:02	down		-	conflicting
1:11	down	-		down (agreeing)
1:13	down		-	down
1:22	down	-		down
1:24	down		-	down

Premise 1: Partial Reachability is Real

- agreements are typical
- *some* disagreement is measurement error
 - ex: timing variation between sites
(transient disagreement)
- but **persistent disagreement is real:**
⇒ **partial reachability in the Internet**




time	truth	A	B	belief
1:00	up	+	+	up
1:01	partial			true partial reach.
1:02	partial		+	
1:11	partial	-		confirmed partial
1:13	partial		+	
1:22	partial	-		
1:24	partial		+	

Outages Given Partial Reachability

- **true outages** are disconnected from the Internet (or off)
- **partial outages are real**
 - disagreements observing reachability happen!
 - don't try to rationalize them away

If Outages are Disconnections, from What?

What is the Internet?




THE ARPANETWORK [attributed to Postel]
 DEC 1969
 4 NODES

What is the Internet, Today?

What is the Internet?

What is the Internet, today?




THE ARPANETWORK [attributed to Postel]
 the IPv4 Internet in the 2000s
 DEC 1969
 4 NODES

What is the Internet's *Edge*?

What is the Internet?

What is the Internet *edge* today?



An IPv4 census:
 all 2^{32} IPv4 addresses,
 on a Hilbert curve;
 brightness shows
 percent responding
<https://ant.isi.edu/address/browse/>

the IPv4 Internet
 in the 2000s



What is the Internet w/Cloud and NAT?


What is the Internet?

What is the Internet *edge and cloud* today?



6 billion smart-phones
 (most behind CG-NAT)

1 billion home broadband
 (many with NAT)



What is the Internet's Core?

What is the **Internet Core?**

the public, shared fabric that connects everyone (edge, cloud, you)

6 billion smart-phones (most behind CG-NAT)

1 billion home broadband (many with NAT)

Partial Reachability in the Internet Core / 2025-11-25

USCViterbi Information Sciences School of Engineering Institute

Premise 2: Define the *Internet Core*

What is the **Internet Core?**

Can it **fragment?**

Can it be **partially reachable?**

6 billion smart-phones (most behind CG-NAT)

1 billion home broadband (many with NAT)

Partial Reachability in the Internet Core / 2025-11-25

USCViterbi Information Sciences School of Engineering Institute

Prior Definitions of “Internet”

- TCP for an “internetwork” [Cerf and Khan, 1974]
- “A collection of interconnected networks is called an internet” [Postel, 1980]
 - examples were ARPAnet and X.25/X.75
- “an agreement to use an evolving set of protocols, in a globally unique address space, to enable universal data delivery” [Federal Networking Council, 1995]

Improving a Definition of “Internet”

- TCP for an “internetwork” [Cerf and Khan, 1974]
- “A collection of interconnected networks is called an internet” [Postel, 1980]
 - examples were ARPAnet and X.25/X.75
- “an agreement to use an evolving set of protocols, in a globally unique address space, to enable universal data delivery” [Federal Networking Council, 1995]

All good properties for the Internet!

But not *operationizable*.

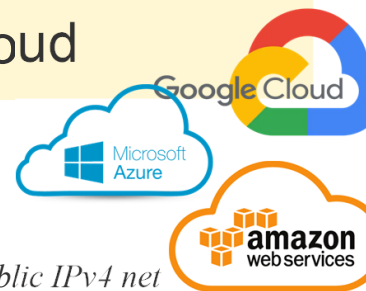
(Measurable, quantifiable—what is and is not the Internet.)

Corner Cases of “the Internet”: Bottom-Up

- 1 laptop with wifi (but not connected to anything)
 - it can run a webserver on a public IP
 - a web browser can connect to it(self)
- 2 laptops with wifi, connecting to each other
- 3 laptops, connecting over a layer-2 802.11s mesh
- 4 laptops, connecting with AODV multi-hop routing
- ...and public IPs

“Corner” Case: the Cloud

- clouds have huge private address space
 - a full 10/8, with NATs to the Internet
 - reminder: 10/8 means 10.*.*.*, or $2^{24} = 16\text{M}$ IP addresses
 - actually, clouds have *dozens* of /8s
 - someday (today?) more *more private /8s than the public IPv4 net*



“Corner” Case: Cloud and Squatting

- clouds have huge private address space
 - a full 10/8, with NATs to the Internet
 - reminder: 10/8 means 10.*.*.*, or $2^{24} = 16\text{M}$ IP addresses
 - actually, clouds have *dozens* of /8s
 - someday (today?) more *more private /8s than the public IPv4 net*
- DISA: 4 public /8s, since 1993
 - never publically routed
 - presumably active inside the U.S. DoD
 - and then routed in Jan. 2021



The Washington Post
Democracy After in Disarray

Minutes before Trump left office, millions of the Pentagon's dormant IP addresses sprang to life

After decades of not using a huge chunk of the Internet, the Pentagon has given control of millions of computer addresses to a previously unknown company in an effort to identify possible cyber vulnerabilities and threats

By Craig Timberg and Paul Sonne
April 24, 2021 at 5:19 a.m. PDT

While the world was distracted with President Donald Trump leaving office on Jan. 20, an obscure Florida company discreetly announced to the world's computer networks a startling development: It now was managing a huge unused swath of the Internet that, for several decades, had been owned by the U.S. military.

Political Threat: Countries

- a country disconnecting from the Internet
 - to prevent cheating on student exams (Iraq, Sudan)
 - because of protests and government instability (Egypt, Sudan)




Political Threat: Countries

- a country disconnecting from the Internet
 - to prevent cheating on student exams (Iraq, Sudan)
 - because of protests and government instability (Egypt, Sudan)
- a country with 24 IPv4 /8s and an aggressive national firewall



Political Threat: Countries

- a country disconnecting from the Internet
 - to prevent cheating on student exams (Iraq, Sudan)
 - because of protests and government instability (Egypt, Sudan)
- a country with 24 IPv4 /8s and an aggressive national firewall
- a country disconnection for “Internet sovereignty”
 - (or claimed to disconnect)
- multiple countries sanction and de-peering



Business Threat: ISPs Disputes

- two big ISPs with a peering dispute, won't exchange routes
- customers of those ISPs that cannot reach each other Internet

Peering Disputes Migrate to IPv6

Complaints of a "broken IPv6 Internet" at the NANOG mailing list surfaced an IPv6 peering issue between Hurricane Electric and Cogent Communications.

Rich Miller | Oct 22, 2009



Our Contributions

- **recognizing partial reachability as fundamental to network**
- **defining the Internet Core, a solution**
- **algorithms to operationalize the definition**
- **applications that show this understand helps real systems**

Details: Baltra and Heidemann, "What is the Internet? Partial Connectivity at the Internet Core", arXiv:2107.11439v3

Our Contributions

- recognizing partial reachability as fundamental to network
- **defining the Internet Core**, a solution
- algorithms to operationalize the definition
- applications that show this understand helps real systems

Details: Balra and Heidemann, "What is the Internet? Partial Connectivity at the Internet Core", arXiv:2107.11439v3

Our Definition of the Internet Core

The Internet Core is all Active IP Addresses
that can Bidirectionally Route
to more than 50% of the public, Potentially Routable IP addresses

A *conceptual* definition
(no one can instantly measure
reachability between all IPs!)



A useful limit—
a goal for
operational algorithms.



And a basis
to reason about
corner cases.



Active IPs that can Bidirectionally Route to >50% of the public, Routable IPs

- why more than 50%?
 - 50% defines **one, unambiguous component** (a majority!)
- **no central authority** or special locations
- implications:
 - ⇒ there is only *one** Internet (* but wait two slides)
 - a lower threshold allows two groups to claim “the Internet” with a plurality
 - we can end the Internet by splitting into 3 pieces, each <50%
 - no one country or organization can unilaterally claim “the Internet”

Active IPs that can Bidirectionally Route to >50% of the public, Routable IPs

why **bidirectionally route**?

- can ping both ways
- captures “universal data delivery”
- **implication: operationizable** (we can measure it!)
 - alternatives to ping TCP or HTTP, but ping is most benign

Active IPs that can Bidirectionally Route to >50% of the **public, Routable IPs**

why the active and **public, (Potentially) Routable IPs?**

- captures “globally unique address space”
- public, potentially routable IPs: necessary for universal delivery
- implications
 - we actually define *two* Internets: IPv4 and IPv6
 - private addresses are second class

Active IPs that can Bidirectionally Route to >50% of the public, Routable IPs

why the **active?**

- allocated but not-routed doesn't “count”
- IPv6 is mostly unallocated!
- implications
 - use it if you have it!

A Engineering Definition to Be Useful

- users only want working ~~e-mail~~, ~~web~~, ~~Facebook~~, phone apps
- but...
 - policy makers enact laws and rulings
 - engineers design and operate protocols and networks
 - researchers design measurement systems
- ⇒ an operational definition can provide “facts on the ground”, separate from politics or business
- users will benefit from these efforts!

Our Contributions

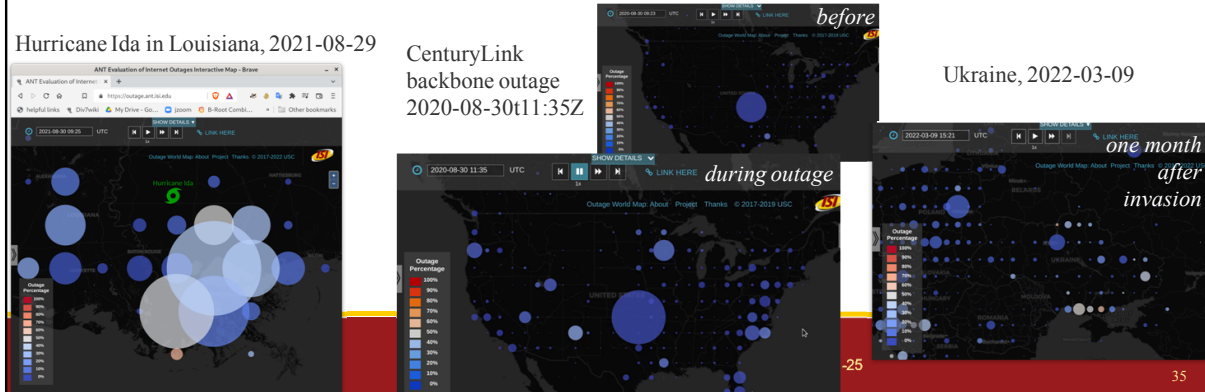
- recognizing partial reachability as fundamental to network
- defining the Internet Core, conceptually
- **algorithms to operationalize the definition**
- applications that show this understand helps real systems

Details: Baltra and Heidemann, “What is the Internet? Partial Connectivity at the Internet Core”, arXiv:2107.11439v3


Observing the (IPv4) Internet

data: <https://outage.ant.isi.edu>

- we've studied **Internet Outages** since 2013 with Trinocular
 - ping 5.1M IPv4 blocks every 11 minutes
 - from 6 vantage points (VPs): Los Angeles, Washington DC, Colorado, Tokyo, Amsterdam, Athens
- but what does it mean when VPs disagree?



Defining the Internet Clarifies Partial Connectivity





here letters are ASes,
 routes are lines (working: solid, failed: dotted)

- outage**
powered off!
- islands**
disconnected
never connected
- peninsulas**
two that cannot reach,
each other, but can
reach many


Islands

- when computers are not reachable (D & E)
 - more common than “computers off” outages (X)!
- happen anytime we have a disconnected part of the IP space

CenturyLink
backbone outage
2020-08-30t11:35Z
*computers on customer
LANs could talk to each other*



Near-total outages observed from ISI (W).
=> *islands at the observer (confirmed network problems)*



Peninsulas

- two locations cannot reach each other, but *can* reach others (B-C)
- routing transients make many short-lived peninsulas
 - see “Internet Optometry” by Bush et al, ACM IMC 2009
- sometimes persistent
 - peering disputes: like Cogent/HE peering dispute around IPv6
 - routing misconfiguration
 - firewalls


A Peninsula in Poland

On 2017-10-23,
for 3 hours from 22:02Z

one VP (W)
could reach 5 Polish ISPs

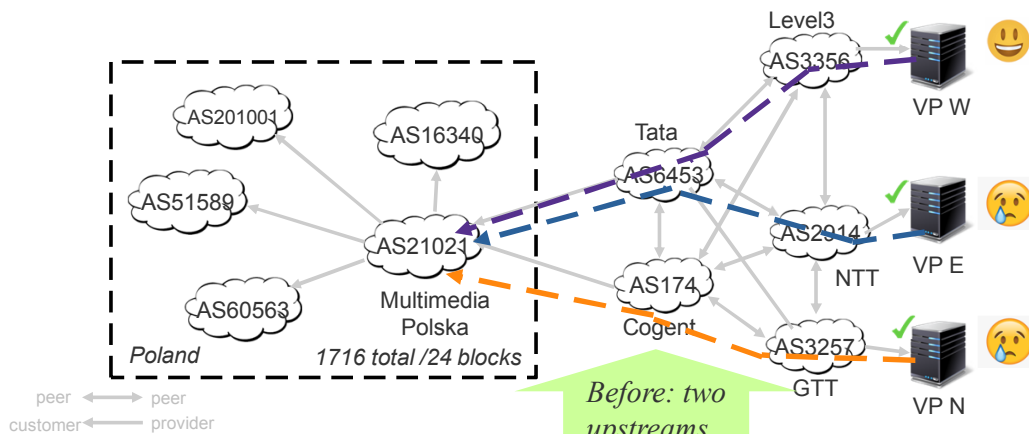
but 5 others could not.

=> A peninsula! Why?

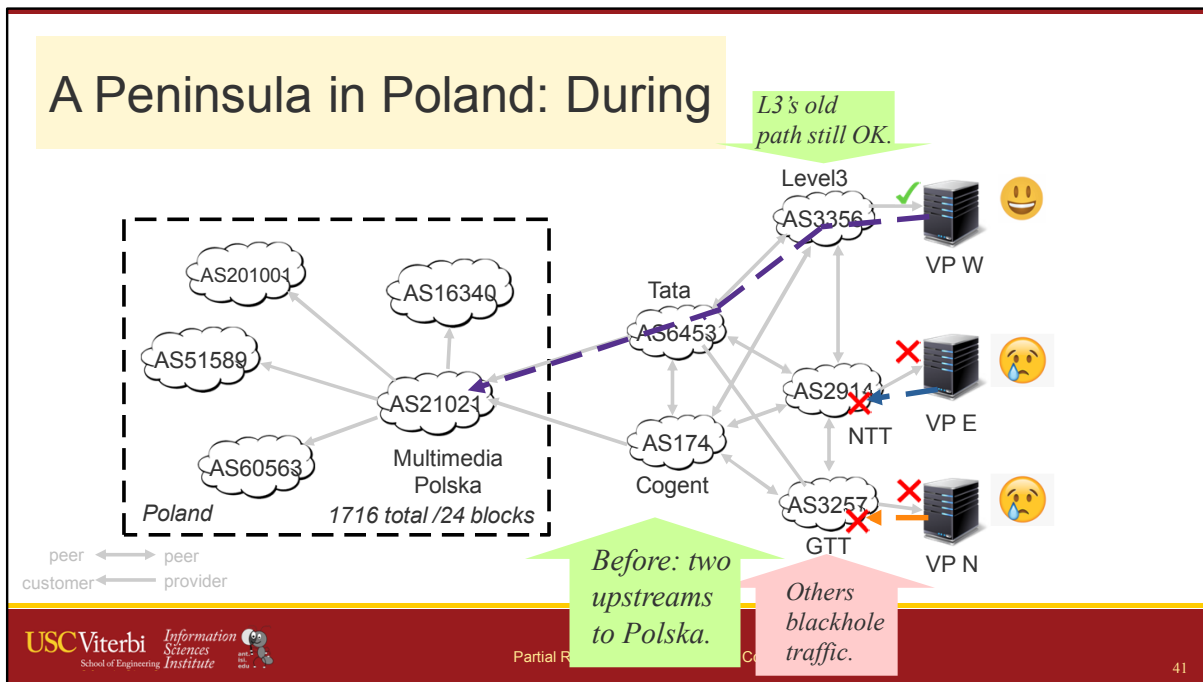


- ✓ VP W
- ✗ VP E
- ✗ VP N

A Peninsula in Poland: Before



Before: two upstreams to Polska.



Our Contributions

- recognizing partial reachability as fundamental to network
- defining the Internet Core, conceptually
- **algorithms to operationalize the definition**
 - Chiloe: an algorithm to detect islands
 - Taitao: an algorithm to detect peninsulas
- applications that show this understand helps real systems

Details: Baltra and Heidemann, "What is the Internet? Partial Connectivity at the Internet Core", arXiv:2107.11439v3

USC Viterbi Information Sciences School of Engineering Institute



Partial Reachability in the Internet Core / 2025-11-25

42

Chiloe: an Algorithm to Detect Islands


- idea: when a VP is up, but can't see most or all of the Internet
=> an island

Near-total outages observed from ISI (W).
=> *islands at the observer (confirmed network problems)*



Taitao: An Algorithm to Detect Peninsulas

- idea: probe a target network from several independent VPs
- if they disagree (some reach and others don't)
=> peninsula



Applying Taitao: Detecting Peninsulas

- data source: reanalysis of Trinocular
 - from: LA, Denver, Washington (DC), Amsterdam, Athens, Tokyo
 - to: 5M IPv4 /24s, pinging every 11 minutes
 - reanalyze 21 days starting 2017-10-10
- validate against traceroutes from CAIDA's Ark
 - 171 VPs, tracerouting every /24 once a day

Validating Taitao Peninsulas: Precision and Recall

True Positive and True Negatives
 • we see most events => high recall

False Positive (loose criteria):
 • => good precision

(strict criteria => lower precision, likely due to firewalls)

		Ark			
		Peninsula	Non Peninsula		
Taitao	Peninsula	184	251 (strict) / 40 (loose)	Precision	0.42 / 0.82
	Non Peninsula	12	1,976,701	Recall	0.94 / 0.94
				F score	0.58 / 0.88

False Negative:
 • Trinocular (all down) => firewalled

=> high recall (0.94):

we see most events

=> ok to good precision (0.42 to 0.82)

usually correct (w/ flexible comparison)

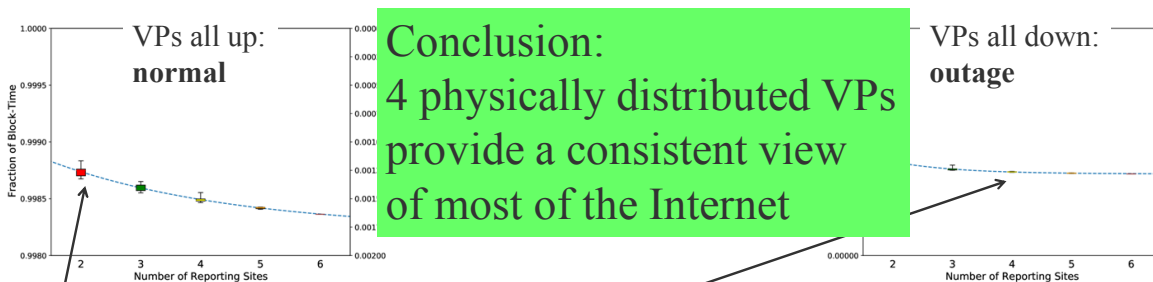
Our Contributions

- recognizing partial reachability as fundamental to network
- defining the Internet Core, conceptually
- algorithms to operationalize the definition
- applications that show **this understand helps real systems**

Details: Balra and Heidemann, "What is the Internet? Partial Connectivity at the Internet Core", arXiv:2107.11439v3

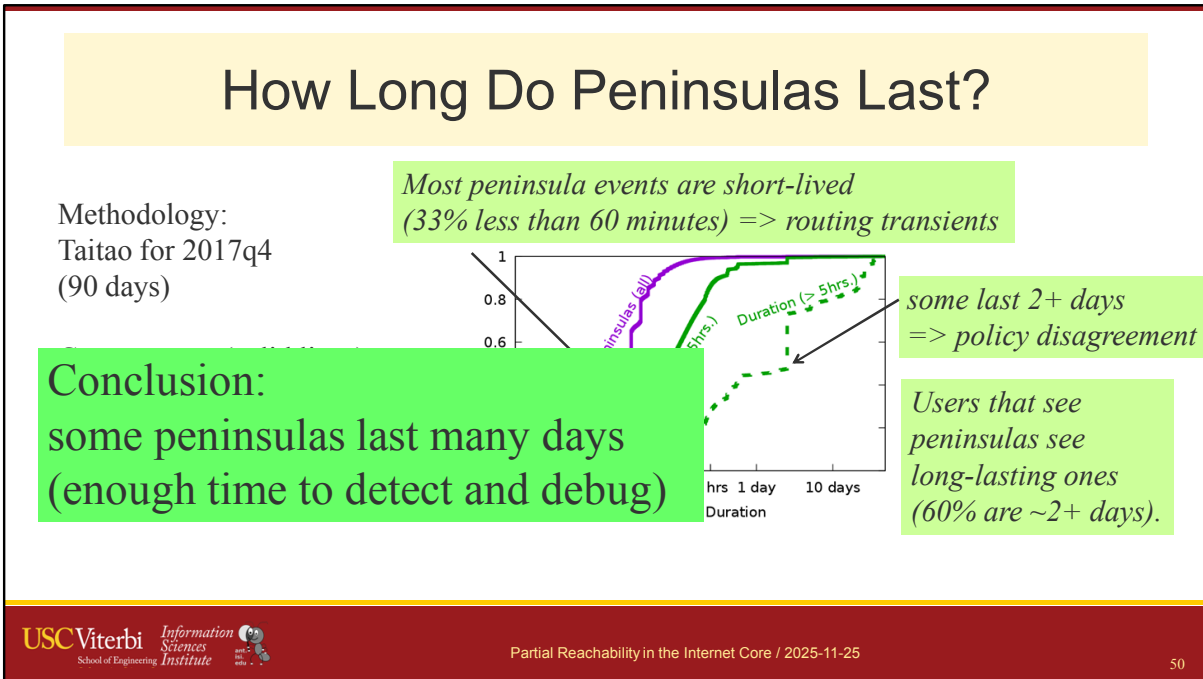
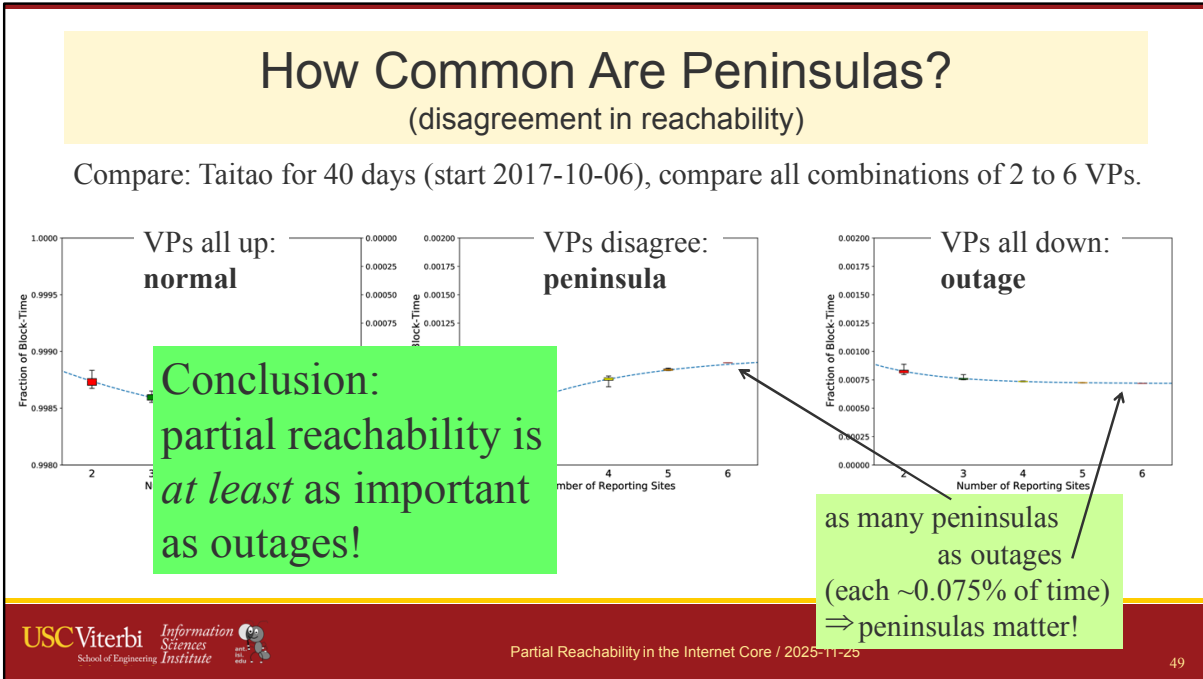
Outages from N Vantage Points

Compare: Taitao for 40 days (start 2017-10-06), compare all combinations of 2 to 6 VPs.



variance is small:
⇒ VPs are *independent*
(any 2 give the same result)

values converge
(4 VPs ≈ 6 VPs)
⇒ VPs are *enough*

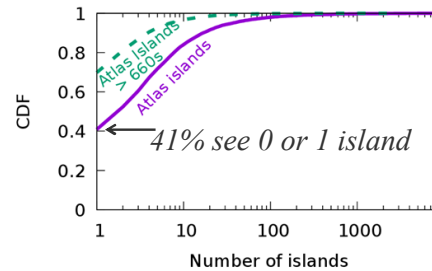


Applying Chiloe: How Common Are Islands?

islands in Trinocular (~3 y, 6 VPs, at least 660s)

site	events	events/year
W	5	1.67
C	2	0.67
J	1	0.33
G	1	0.33
E	3	1.00
N	2	0.67
all (averaged)	2.3	0.78
all (cumulative)	14	4.67

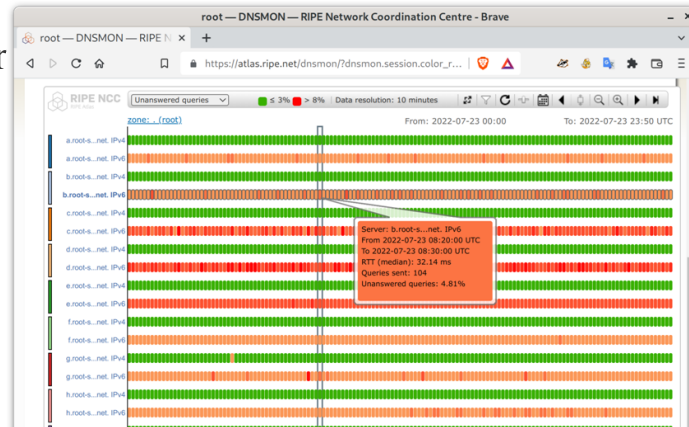
islands in RIPE Atlas
(3 months: 2021q3, ~13k VPs,
at least a minute)



=> Islands are not common, but they happen to many people.

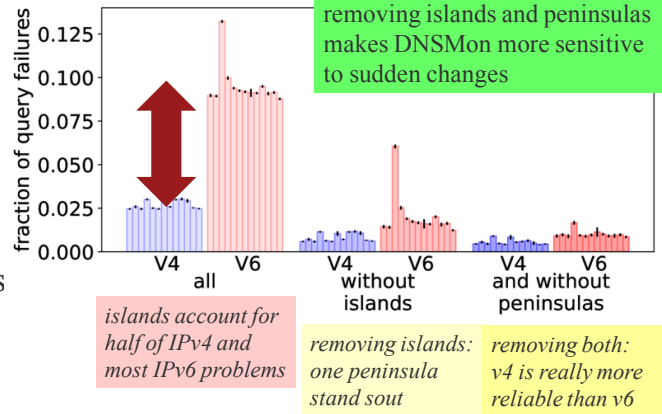
A Practical Application: Making RIPE DNSMON More Sensitive

- dnsmon.ripe.net: essential root DNS monitor
 - IPv4 and IPv6 reachability to the 13 root DNS letters
- problem: for years IPv6 reachability is bad (worse than IPv4)
- is this *real*? *fixable*?



Partial Outages and RIPE DNSMon

- the problem: DNSMon didn't really consider partial outages
 - Atlas probes were *islands*: cannot reach the outside*
 - one RSO is on a *peninsulas*: they cannot reach everywhere
- filtering these standing problems lets DNSMon see *daily change*



A Political Application: Can a Country Be Ejected?



- q: can any single country or region “take over” the Internet?
- a: no

Conclusion:
no one country or RIR can (today) control IPv4 or v6:
the Internet is a collaboration

RIR	IPv4 hosts		IPv6 /32s	
AFRINIC	121M	3.3%	9,661	3%
APNIC	892M	24.0%	88,614	27.8%
China	345M	9.3%	54,849	17.2%
ARIN	1673M	45.2%	56,172	17.6%
U.S.	1617M	43.7%	55,026	17.3%
LANIC	191M	5.2%	15,298	4.8%
RIPE NCC	826M	22.3%	148,881	46.7%
Germany	124M	3.3%	22,075	6.9%
unallocated	221	100%	318,626	100%

Conclusions

- **partial reachability is real**
- **defining the Internet** helps
Active IPs that can Bidirectionally Route
to >50% of the public, Routable IPs
- new algorithms **can measure peninsulas and islands**
- this perspective **improves**
 - **outage detection, RIPE DNSMon**
 - does this definition help clarify *your* Internet?



(more detail: our tech report:
arXiv:2107.11439 v3)