

Defining The Internet Core: Partial Connectivity and Internet Fragmentation

Guillermo Baltra and *John Heidemann*

U. of Southern California / Information Sciences Institute and CS Dept.
for the IRTF Human Rights Protocol Considerations WG


2023-07-26

This research is sponsored by NSF -2028279 (MINCEQ), CNS-2007106 (EIEIO). 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 © 2023 by Baltra and Heidemann
Release terms: CC-BY-NC 4.0 international

What is the Internet?




THE ARPA NETWORK

DEC 1969

4 NODES

[attributed to Postel]

What is the Internet, today?

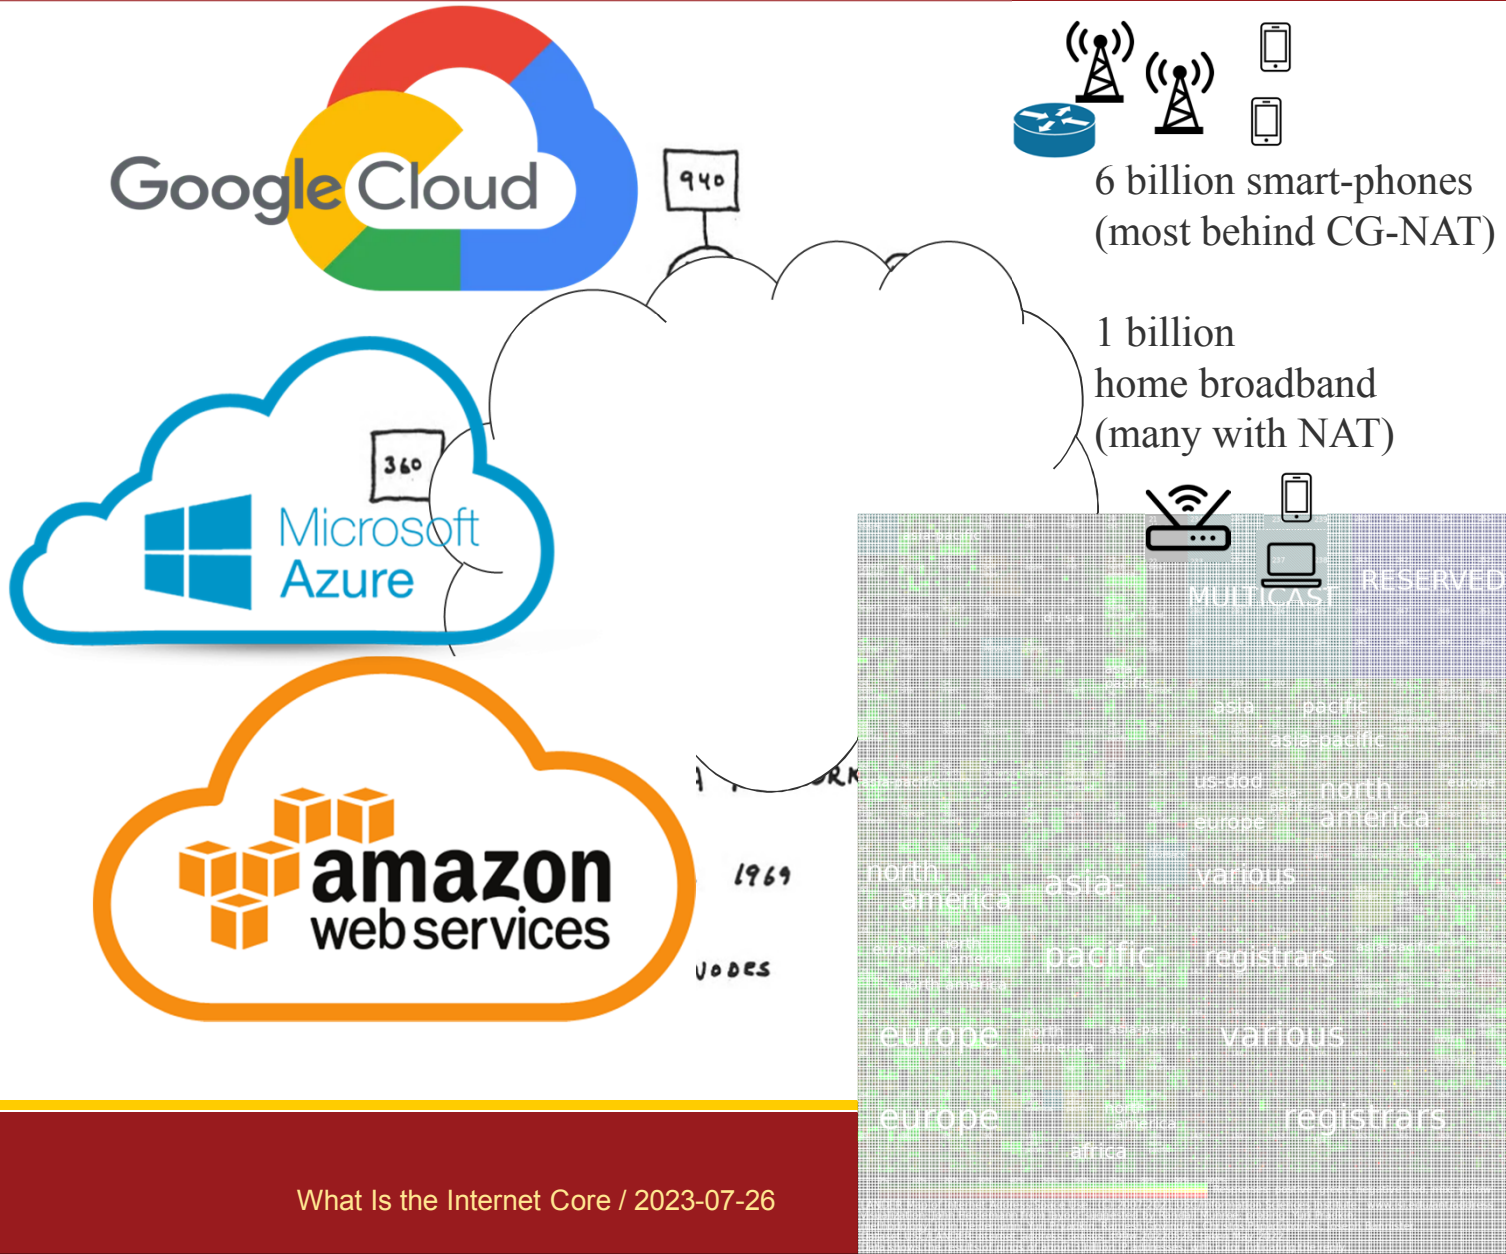


the IPv4 Internet EC 1969
in the 2000s

4 NODES

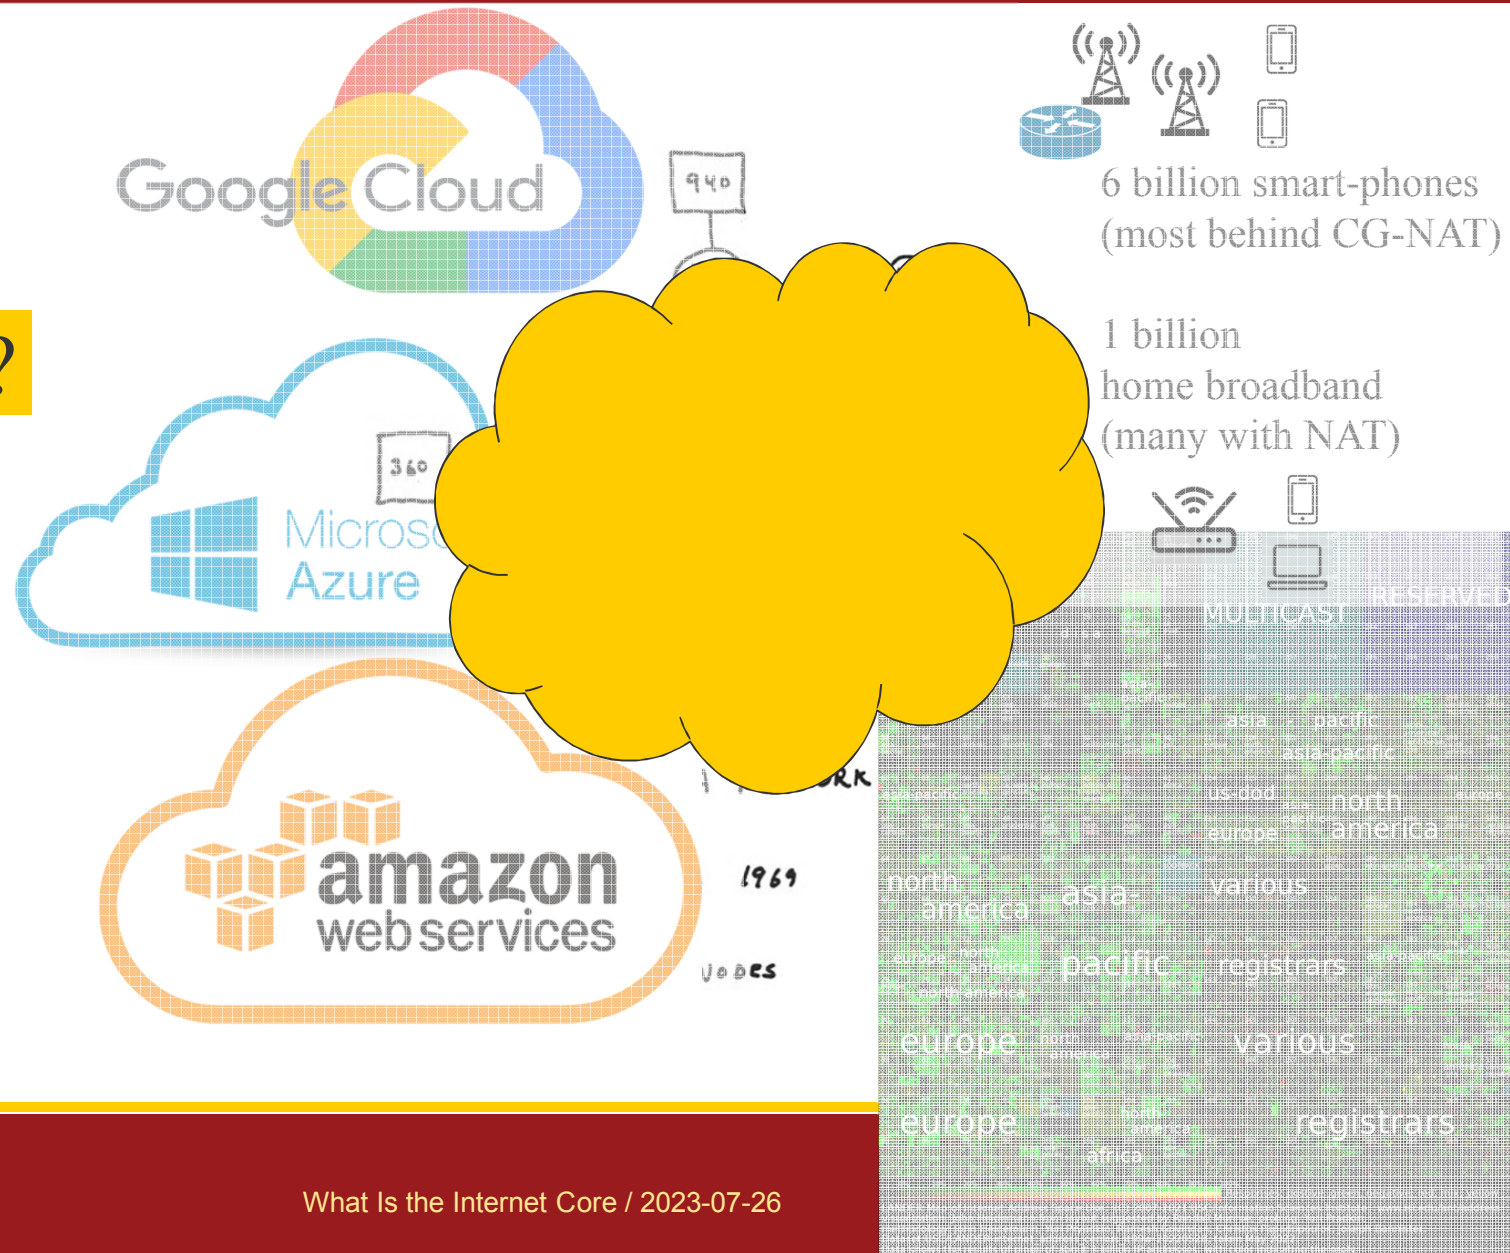
[attributed to Postel]

What is the Internet edge and cloud?



What is the Internet Core?

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




Today's Internet Stressors

- political pressure
 - Russia's "sovereign network"
 - a US or Chinese "off switch"?
 - countries have turned off their Internet



- business and technical pressures
 - persistent unreachability from peering disputes



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.

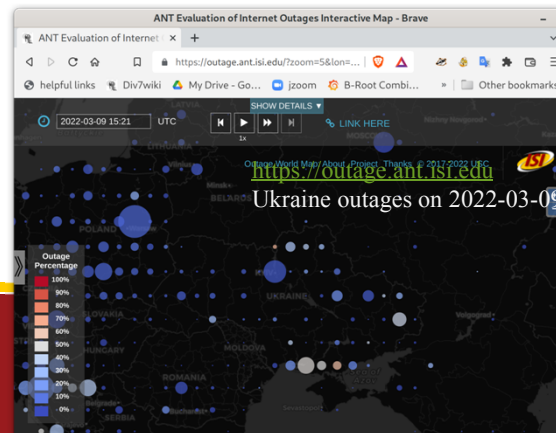
Rich Miller | Oct 22, 2009



China, Huawei propose internet protocol with a built-in killswitch

New IP sounds good in some ways, but might help authoritarian crackdowns.

- quantifying Internet reliability
 - we and others measure outages
 - but all the corner cases?



- existing stressors linger
 - NAT
 - IPv4 / IPv6 transition,
 - IP squat space and IPv4 full allocation

Stressors ... and Our Response

- political pressure
 - Russia’s “sovereign network”
 - a US or Chinese “off switch”?
 - many countries have turned the Internet off
- technical pressures
 - persistent unreachability from peering disputes
- quantifying Internet reliability—measuring outages
 - we measure outages, but all these corner cases?
- existing stressors linger
 - NAT, IPv4 / IPv6 transition, IP squat space
- a technical definition, ahead of time, to illuminate political challenges
 - ⇒ *explaining to an MP the technical limits of a legal “kill switch”*
 - ⇒ *the Internet is collaborative—no single control controls it*
- define persistent unreachability
 - ⇒ *enable measurements of reliability*
 - (so sunlight can help motivate resolution of private disputes)*
- defining outages (ISI has outage data since 2014)
 - ⇒ *must understand **partial connectivity**—it not just binary*
 - ⇒ *enable robust comparisons of ISP reliability*
- a new definition helps clarify old disputes

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

Corner Cases: 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*
- DISA: 4 public /8s, since 1993
 - never publically routed
 - presumably active inside the U.S. DoD
 - and then routed in Jan. 2021

Corner Cases: 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 disconnecting to proclaim “Internet sovereignty”
 - (or claimed to disconnect)
- multiple countries de-peering with country X due to sanctions

Our Contributions

- defining the Internet Core, conceptually
- algorithms to operationalize the definition
 - Taitao: detecting peninsulas—persistent, partial connectivity
 - Chiloe: detecting islands—fully disconnected parts
- measurements from today's Internet

Details: Baltra 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 the connected component of active, public IP addresses that can reach 50% of each other.

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.



Connected Component of Active, Public IP Addresses Reaching >50%

- why more than 50%?
 - 50% defines **one, unambiguous component** (a majority!)
- **no central authority** or special locations
- implications:
 - ⇒ there is only *one** Internet (* actually two: IPv4 and IPv6)
 - 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”

Connected Component of Active, Public IP Addresses **Reaching** >50%

reachability?

- 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

Connected Component of Active, Public IP Addresses Reaching >50%

who: the active, public IPs

why?

- captures “globally unique address space”
- public IPs: necessary for universal delivery
- active: what’s in use, not “dark space”
- implications
 - we actually define *two* Internets: IPv4 and IPv6
 - private addresses are second class

A Engineering Definition to Inform Policy

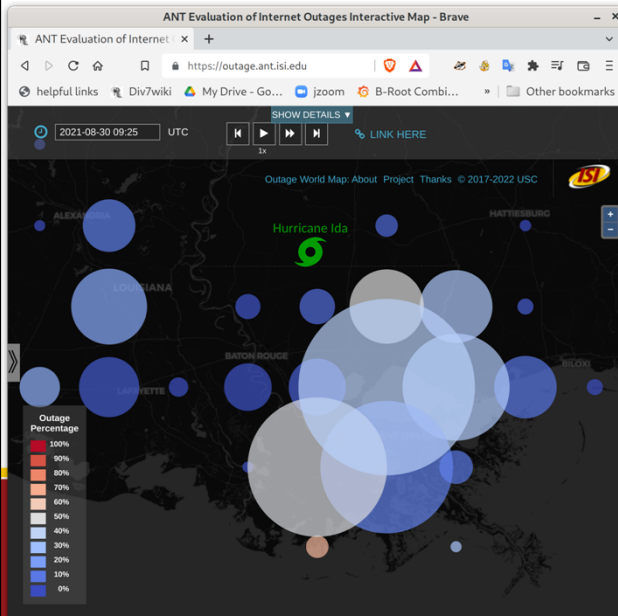
- users only want working ~~e-mail~~, ~~web~~, ~~Facebook~~, phone apps
 - but...
 - policy makers enact laws and rulings
 - engineers design and operate protocols and networks
 - researchers design measurement systems
- ⇒ an operational definition can guide them
and help promote an open Internet
- users will benefit from these efforts!

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

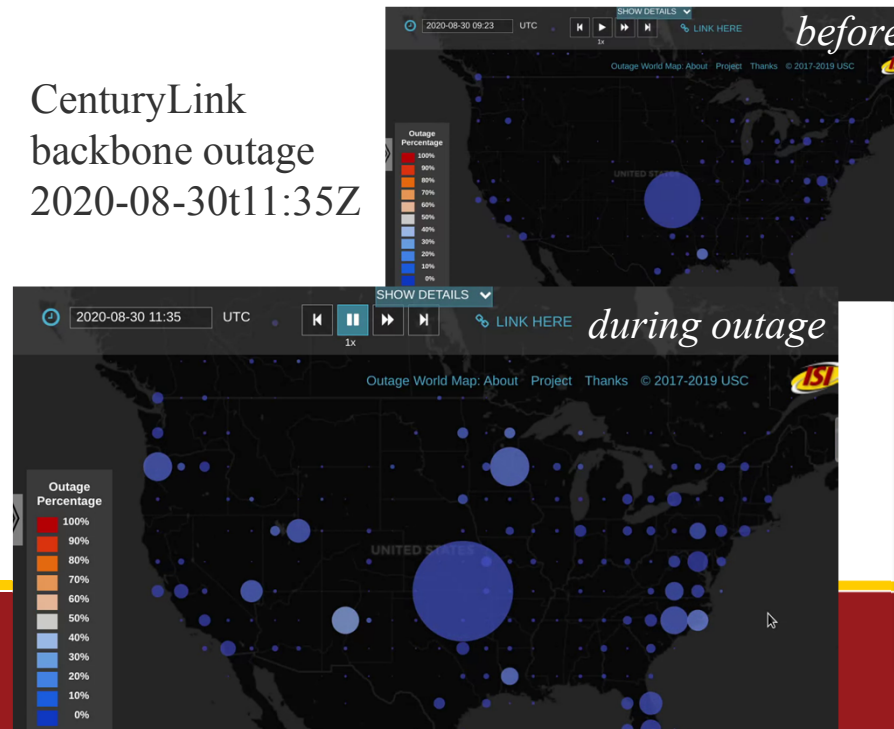
Observing the (IPv4) Internet

- 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?

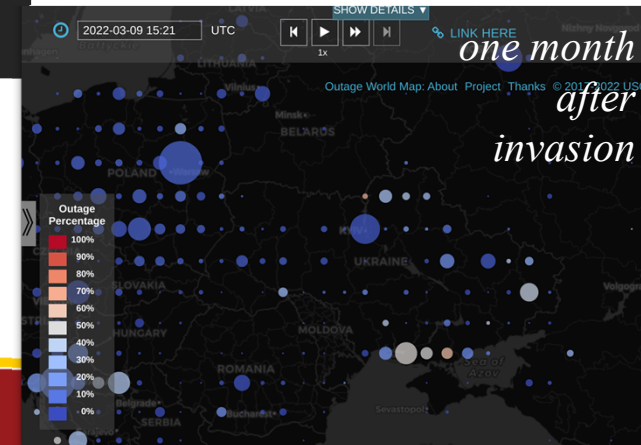
Hurricane Ida in Louisiana, 2021-08-29



CenturyLink backbone outage
2020-08-30t11:35Z



Ukraine, 2022-03-09



Defining the Internet Clarifies Partial Connectivity

when no VPs can reach the target:

outages: computers are off

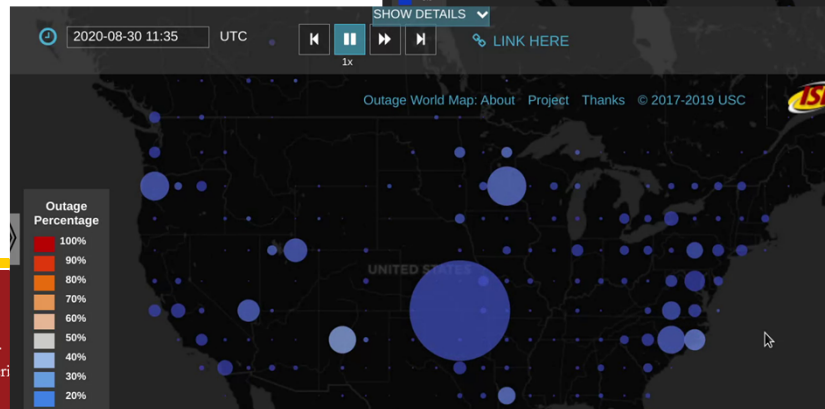
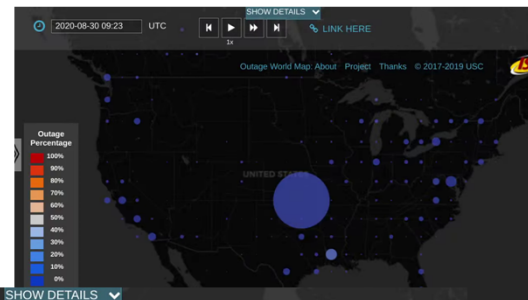
islands: computers on, but unreachable

peninsulas: when *some VPs, but not all* can reach
– partial connectivity in the Internet!

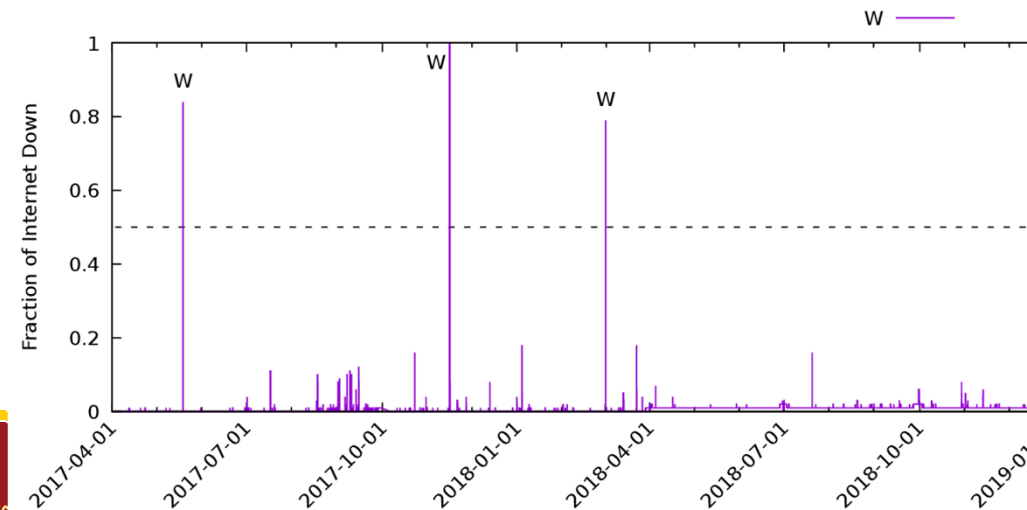
Islands

- when computers are not reachable
 - more common than “computers off” outages!
- 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

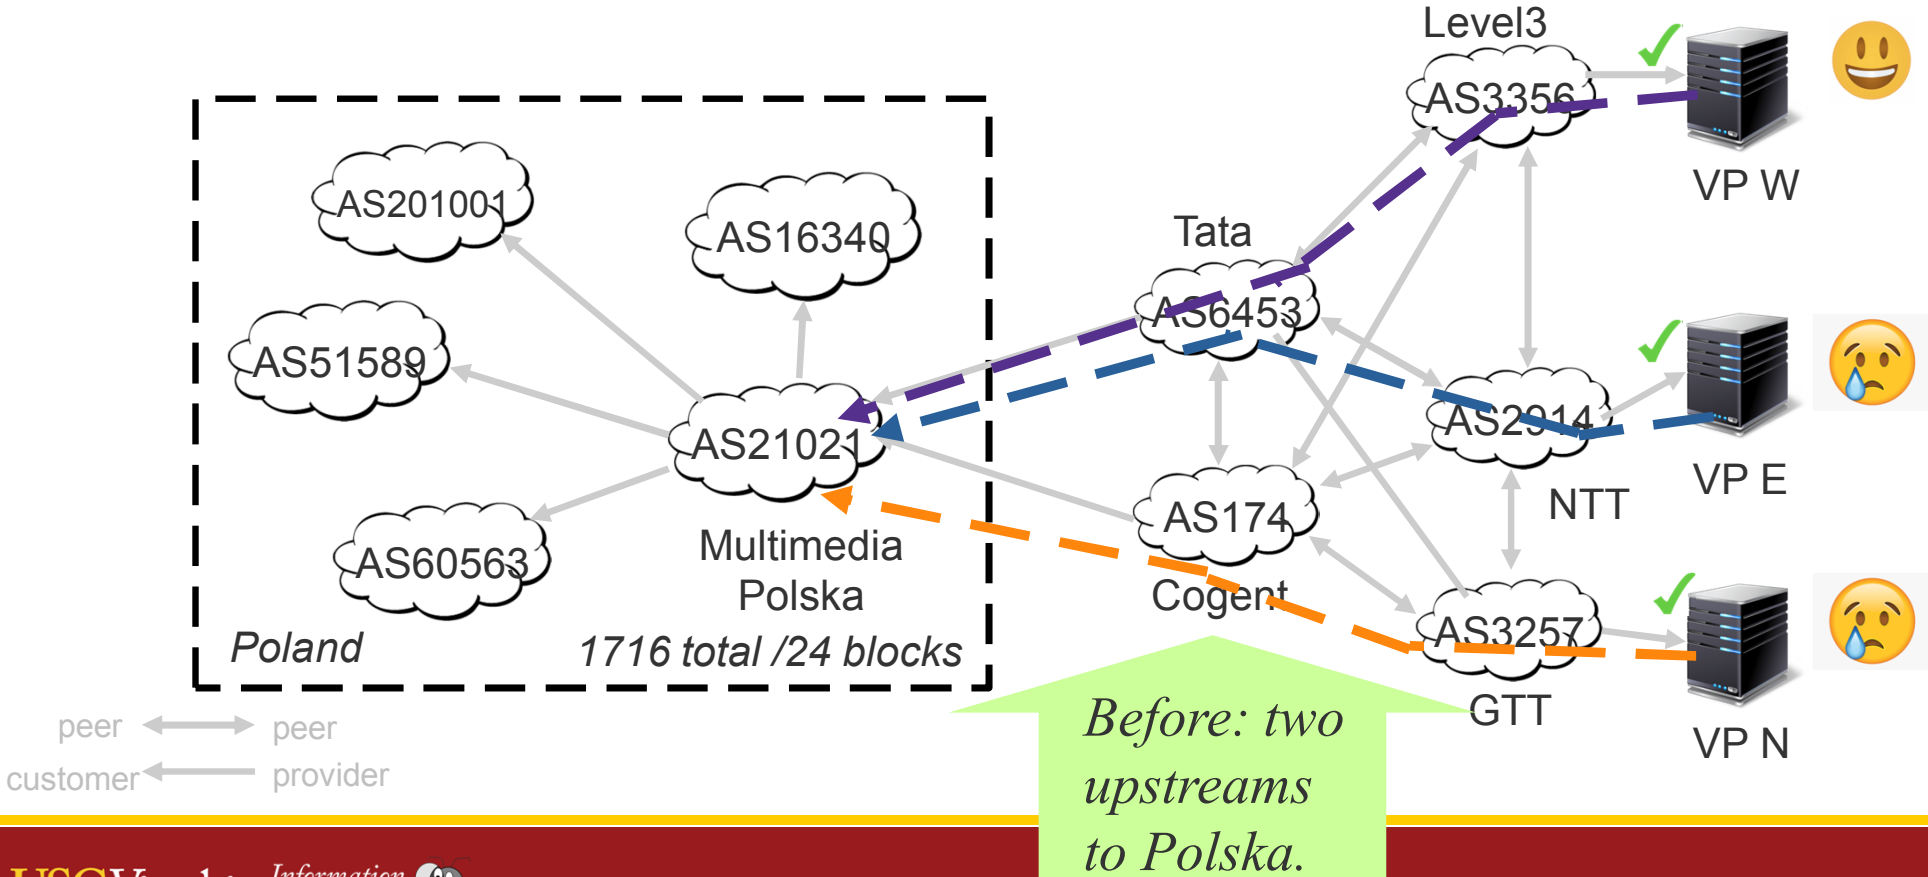
- two locations cannot reach each other, but *can* reach others
- common transient event as routes converge
 - see “Internet Optometry” by Bush et al, ACM IMC 2009
- sometimes persistent
 - peering disputes: like Cogent/HE peering dispute around IPv6
 - routing misconfiguration
 - firewalls

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?*

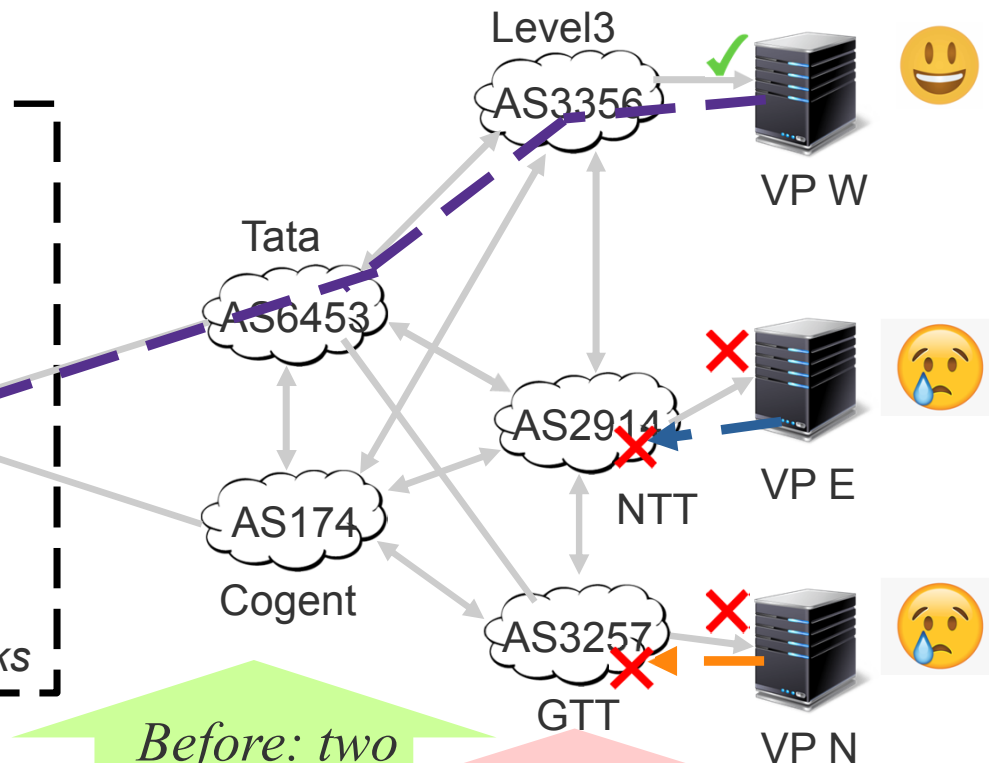



A Peninsula in Poland: Before



A Peninsula in Poland: During

L3's old path still OK.



Before: two upstreams to Polska.

Others blackhole traffic.


Contributions

- defining the Internet (a conceptual definition)
- **algorithms that operationalize the definition**
 - **Taitao: detecting peninsulas—persistent, partial connectivity**
 - **Chiloe: detecting islands—fully disconnected parts**
- data from today's Internet

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

Taitao: Detecting 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

- both systems usually agree => high recall

False Positive (loose criteria):

- Requires >1 VP down

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

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

False Negative:

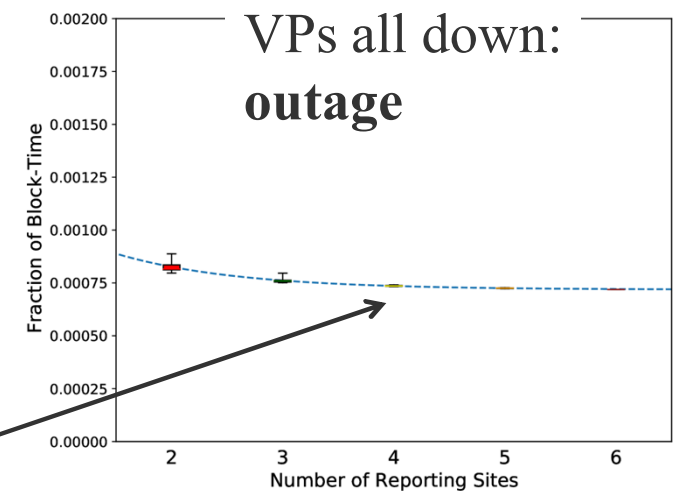
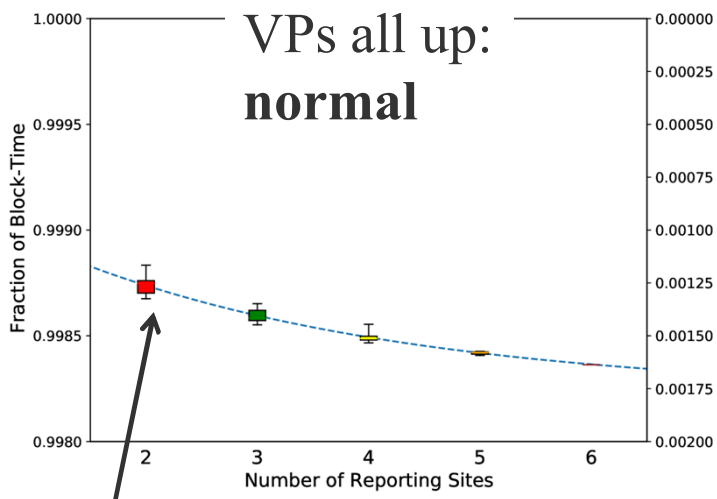
- Trinocular (all down) => firewalled

high recall => we see most peninsulas

high loose precision => most are loose peninsulas

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:
 \Rightarrow VPs are *independent*
(any 2 give the same result)

values converge
(4 VPs \approx 6 VPs)
 \Rightarrow VPs are *enough*

How Common Are Peninsulas?

(disagreement in reachability)

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



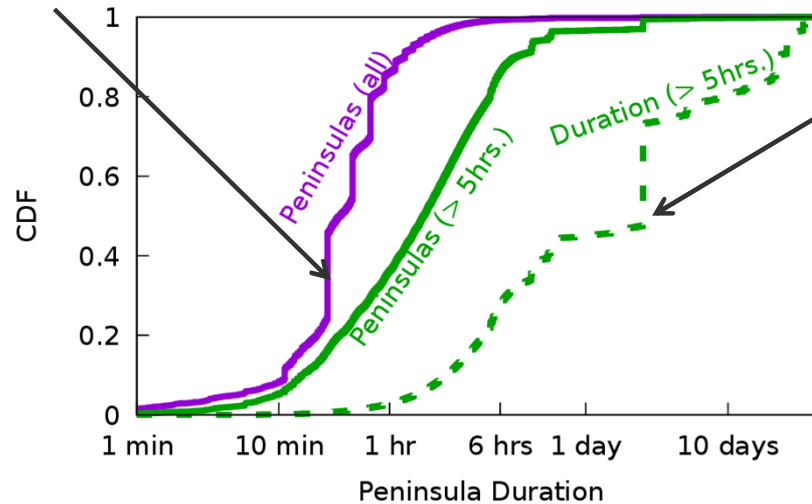
as many peninsulas
as outages
(each ~0.075% of time)
⇒ peninsulas matter!

How Long Do Peninsulas Last?

Methodology:
Taitao for 2017q4
(90 days)

Count events (solid lines)
and cumulative duration
(dashed line).

*Most peninsula events are short-lived
(33% less than 60 minutes) => routing transients*



*some last 2+ days
=> policy disagreement*

*Users that see
peninsulas see
long-lasting ones
(60% are ~2+ days).*

Secession and the Internet

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

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

Conclusions

- defining the Internet:
>50% reachability of active, public IPs
- new algorithms: Taitao for peninsulas, Chiloe for islands
- data about Internet partial connectivity
 - peninsulas are as common as outages
 - no single country can claim “the Internet”
- does this definition help clarify questions in *your* Internet?

(more detail: our tech report: [arXiv:2107.11439 v3](https://arxiv.org/abs/2107.11439))

