

# Reasoning About Internet Connectivity

Guillermo Baltra  
USC/ISI  
Marina del Rey  
California, USA  
baltra@ant.isi.edu

Tarang Saluja  
Swarthmore College  
Swarthmore  
Pennsylvania, USA  
tsaluja1@swarthmore.edu

Yuri Pradkin  
USC/ISI  
Marina del Rey  
California, USA  
yuri@isi.edu

John Heidemann  
USC/ISI and Thomas  
Lord Dept. of CS  
Marina del Rey  
California, USA  
johnh@isi.edu

## ABSTRACT

Innovation in the Internet requires a global *Internet core* to enable communication between users in ISPs and services in the cloud. Today, this Internet core is challenged by *partial reachability*: political pressure threatens fragmentation by nationality, architectural changes such as carrier-grade NAT make connectivity conditional, and operational problems and commercial disputes make reachability incomplete for months. We assert that *partial reachability is a fundamental part of the Internet core*. While other studies address partial reachability, this paper is the first to *provide a conceptual definition of the Internet core* so we can reason about reachability from principles first. Following the Internet design, our definition is guided by reachability, not authority. Its corollaries are *peninsulas*: persistent regions of partial connectivity; and *islands*: when networks are partitioned from the Internet core. We show that the concept of *peninsulas and islands can improve existing measurement systems*. In one example, they show that RIPE’s DNSmon suffers misconfiguration and persistent network problems that are important, but risk obscuring operationally important connectivity changes because they are 5× to 9.7× larger. Our evaluation also informs policy questions, showing no single country or organization can unilaterally control the Internet core.

## 1 INTRODUCTION

Today’s Internet is very different from when it was created. In 1980, Postel defined “an internet” as “a collection of interconnected networks”, such as the ARPAnet and X.25 [62]. In 1995, the Federal Networking Council defined “Internet” as (i) a global address space, (ii) supporting TCP/IP and its follow-ons, that (iii) provides services [33]. Later work added DNS [48] and IPv6. Yet today, users at home and work access the Internet indirectly through Network Address Translation (NAT) [79], or from mobile devices through Carrier-Grade NAT (CG-NAT) [68]. Many public services operate from the cloud, visible through rented or imported IP addresses, inside networks with multiple levels of virtualization [40]. Media is replicated in Content Delivery Networks (CDNs). Access is mediated by firewalls. Yet to most users, Internet services are so seamless that technology recedes and the web, Facebook, or their phone is their “Internet”.

**The Core and Its Challenges:** Today’s rich, global Internet services exist because *the Internet architecture created a single, global **Internet core** where all can communicate freely*. The Internet core has enabled 40 years of permissionless innovation [54], from e-mail and remote login to the web and streaming media, today’s Internet powers telephony and video distribution. Although today many edge devices are client-only, *continued innovation and international exchange depends on a near-completely connected Internet core*.

But today *universal reachability in Internet core is often challenged with threats of fragmentation*. Political pressure pushes to Balkanize the Internet along national borders [26, 27]. Consider Russia’s 2019 sovereign-Internet law [21, 59, 65] and national “Internet kill switches” debated in U.S. [39], the U.K., and deployed elsewhere [18, 20, 41, 77]. We suggest that *technical clarification can inform policy discussions* as threats of de-peering place the global Internet at risk (§3.2).

*Architectural pressures* from 30 years of evolution segment today’s Internet core: services are gatewayed through proprietary cloud APIs, users are usually second-class and client-only due to NAT, firewalls interrupt connectivity, and the world straddles a mix of IPv4 and IPv6. Architecture sometimes follows politics, with China’s Great Firewall managing their international communication [3, 4], and Huawei proposing “new Internet” protocols [32]. We suggest that technical methods can help us *reason about changes to Internet architecture*, both to understand the implications of partial address reachability and evaluate the maturity of IPv6.

*Operational challenges* can cause partial reachability. Peering disputes can cause long-term partial reachability [51]. Unreachability has been recognized and detected experimentally [24], and systems exist that mitigate partial reachability [2, 49, 50]. Yet an understanding how duration and breadth of partial reachability has remained elusive.

**Contributions:** Our first contribution that *an Internet core*, the global address space to which everyone interconnects, is essential to continued innovation. *Understanding partial connectivity is key to reasoning about challenges* to the Internet and innovation. We hope that recognition the importance of one, global, open Internet core will help clarify the stakes when nations assert sovereignty and architectural changes require mediated communication. We also show

that *peninsulas*, regions of partial connectivity that are sometimes long-lasting, are a real-world problem as serious as network outages [42, 63, 67, 75, 76].

Our second contribution is to offer a rigorous, conceptual definition of the Internet core as *the strongly connected component of more than 50% of active, public-IP addresses that can initiate communication with each other* (§2.1). By requiring bidirectional initiation, this definition captures the uniform, *peer-to-peer nature of the Internet core* necessary for first-class services. The 50% requirement defines *one, unique Internet core*, without central authority, historical precedent, or special locations, since multiple majorities are impossible. Unlike prior work [2, 49, 50], this *conceptual* definition avoids dependence on any specific measurement system. We have realized this definition in operational systems with two different data sources (§3.5 and elsewhere [8]). This conceptual Internet core defines an asymptote against which our current and future operational systems can compare.

Our final contribution is to *use our definition to clarify policy, architectural, and operational questions*. We bring technical light to policy choices around national networks (§3.2) and de-peering (§3.4). Our definition can help evaluate the IPv4/v6 transition and clarify operational questions in IPv4 address use, in outage detection [42, 63, 67, 75, 76], and WAN [2, 49, 50] and cloud [74] reachability optimization (§3.3). We apply our results to widely used RIPE DNSmon (§3.5). Today DNSmon shows persistent high query loss (5–8% to the DNS Root [71]), we show that most of this loss is due to misconfiguration and persistent partial connectivity. While such factors matter, they are 5× and 9.7× (IPv4 and v6) larger than other operationally important signals. Separating them therefore improves sensitivity in DNSmon [1] (§3.5).

**Artifacts and ethics:** All of the data used and created [5] in this paper is available at no cost. Our work poses no ethical concerns: we reanalyze existing data with new algorithms, and have no information about individuals. IRB review declared our work non-human subjects research (USC IRB IIR00001648).

## 2 PROBLEM: PARTIAL REACHABILITY

To understand Internet connectivity we must rigorously define *the Internet core* to which we connect, to answer the political, architectural, and operational questions from §1.

First, a definition should be both *conceptual* and *operational* [28]. Our conceptual definition in §2.1 articulates what we would *like* to observe and suggests a limit that an implementation can approach. In §3.5 we operationalize our definitions to improve understanding of DNSmon. Prior definitions are too vague to operationalize.

Second, a definition must give both sufficient *and* necessary conditions to be part of the Internet core. Prior work gave only sufficient conditions, like supporting TCP [13, 33,

62]. Our new *necessary* conditions determine when a network would *leave* the Internet core.

### 2.1 Defining the Internet Core

We define the Internet core as *the strongly-connected component of more than 50% of active, public IP addresses that can initiate communication with each other*. Computers behind NATs and cloud load-balancers are on *branches*, participating but not part of the core, often with dynamically allocated, transient public IP addresses. We believe this definition is simple, but with subtle implications. For example, it defines *two* Internet cores: one each for IPv4 and IPv6.

We build on the terms “interconnected networks”, “IP protocol”, and “global address space” from prior definitions [13, 33, 62], and their common assumption that two computers on the public Internet should be able to reach each other directly at the IP layer.

We formalize “network interconnection” [13] by considering reachability over public IP addresses: addresses  $x$  and  $y$  are interconnected if traffic from  $x$  can reach  $y$  and  $y$  can reach  $x$ . Networks are groups mutually reachable addresses.

**Why more than 50%?** We take as an axiom that there should be *one Internet core* per address space, or a reason why that no core exists. The definition must unambiguously identify “the” Internet core given conflicting claims.

Requiring a majority of active addresses ensures that there can be only one Internet core, since any two majorities must overlap. Any smaller fraction could allow two groups to make valid claims. We discuss how to identify the core in the face of conflicting claims in §3.1.

The definition of the Internet core should not require a central authority. Majority supports assessment independent of any authority, as in other distributed consensus algorithms [52, 53, 58]. Any computer to prove it is in the Internet core by reaching half of active addresses, as defined by multiple, independent, long-term evaluations [22, 44, 81]. We explicitly do not require identification of “tier-1” ISPs, an imprecise term often entangled with business concerns.

A majority defines an Internet core that can end: fragmentation occurs should the current Internet core break into three or more disconnected components where none retains a majority of active addresses. If a large enough organization, nation, or group chose to secede, or are expelled, *an* Internet core could become several no-longer internets.

**Why all and active addresses?** In each of IPv4 and IPv6 we consider all addresses equally. Public Internet addresses are global, and the Internet core intentionally designed without a hierarchy [15]. Consistent with goals for network decentralization [25], a definition should not create hierarchy, nor designate special addresses by age or importance. *Active* addresses are blocks that are reachable, defined below.

These definitions are relatively apolitical and reduce first-mover bias, discussed in §3.4. Addresses are an Internet-centric metric, unlike population or countries. Requiring activity reduces the influence of large allocated, but unused, space, such as in legacy IPv4 /8s and new IPv6 allocations.

**Reachability, Protocols and Firewalls:** End-to-end reachability avoids difficult discovery of router-level topology.


Our conceptual definition allows different definitions of reachability. Reachability can be tested by measurements with some protocol, such as ICMP echo-request (pings), or TCP or UDP queries, or by data-plane reachability with BGP. Any specific test will provide an operational realization of our conceptual definition. Particular tests will differ in how closely each approaches the conceptual ideal.

Firewalls complicate observing reachability, particularly when conditional or unidirectional. We accept that the results of specific observations may vary with different protocols or observation times; practically we see results are stable with Internet-wide measurements [8].

We have two implementations of peninsula and island detection; both use publicly-available data from existing measurement systems. One uses Trinocular [63], because of its frequent, Internet-wide ICMP echo requests (11-minutes to 5M IPv4 /24s). Prior work has shown ICMP provides the most response [9, 29, 63], and can avoid rate limiting [43], other other protocol options are possible. Our second uses RIPE Atlas because of its use in DNS (§3.5).

**Why reachability and not applications?** Users care about applications, and a user-centric view might emphasize reachability of HTTP or to Facebook rather than at the IP layer. Our second realization uses public data from RIPE Atlas, with DNS as the application, as described in §3.5. Future work may look at other, more user-centric applications. However, we suggest reachability at the IP layer is a more fundamental concept. IP has changed only twice since 1969 with IPv4 and IPv6, but dominant applications wax and wane, and applications such as e-mail extend beyond the Internet.

**Why strongly connected and bidirectional reachability?** We require bidirectional reachability (strong connectivity) to identify NAT-only computers as second class-citizens. While most computers today are behind NAT or cloud load-balancers, and NAT-ed computers are useful clients, they require protocols such as STUN [72] to rendezvous through the core, or UPnP [56] or PMP [14] to link to the core. Huge services run in the cloud by leasing public IP addresses from the cloud operator or importing their own (BYOIP). Often services use a single public IP address but employ many servers behind a load balancer [40] or IP anycast [61]. While load balancers or home routers may be on the core, and some cloud VMs use fully-reachable public addresses, devices that are not bidirectionally reachable are not part of the core.



**Figure 1: Blocks A, B and C are the connected core, with B and C peninsulas; D and E islands; X an outage.**

## 2.2 Away from the Core: Unreachability

We now use our definition of the Internet core to reason about where connectivity is incomplete: peninsulas, islands, and outages. Figure 1 shows a toy example, where long-term and current routability is shown by dotted and solid lines and white regions show current data-plane reachability. All address blocks but *E* form the core. Blocks *B* and *C* are on *peninsulas* because they do not route to each other, although data could relay through *A*. Block *X* has an *outage*; its routes are temporarily down. Blocks *D* and *E* are *islands*: *D* usually can route to the core, but not currently. *E* uses public addresses, but has never announced routes publicly.

**2.2.1 Outages.** A number of groups have examined Internet outages [42, 63, 67, 75]. These systems observe the public IPv4 Internet and identify networks that are no longer reachable—they have left the Internet. Often these systems define outages operationally (network *X* is out because none of our Vantage Points (VPs) can reach it). In this paper, we define an outage as when all computers in a block are off, perhaps due to power loss. We next define islands, when the computers are on but cannot reach the Internet core.

**2.2.2 Islands: Isolated Networks.** An *island* is a group of public IP addresses partitioned from the Internet core, but able to communicate among themselves. Operationally, outages and islands are both unreachable from an external VP, but computers in an island are on and can reach each other.

Islands occur when an organization loses all connections to the Internet core. A business with one office and one ISP becomes an island when its router’s upstream connection fails, but computers in the office can reach each other and in-office servers. An *address island* is when a computer can ping only itself. Externally, islands and outages appear identical.

**Example Islands:** Islands are common in RIPE Atlas [1] when a VP has an IPv6 address on the LAN, but lacks routes to the public IPv6 Internet. In §3.5 we show that this kind of misconfiguration accounts for 5× more IPv6 unreachability than other, more meaningful problems.

We also see islands in reanalysis of data from Trinocular outage detection [63]. Over three years, from 2017 to 2020, we saw 14 cases where one of the 6 Trinocular VPs was active and could reach its LAN, but could not reach the rest of the Internet. Network operators confirm local routing failures in several of these cases.

**2.2.3 Peninsulas: Partial Connectivity.** Link and power failures create islands, but *partial* connectivity is a more pernicious problem: when one can reach some destinations, but not others. We call a group of public IP addresses with partial connectivity to the Internet core a *peninsula*. In a geographic peninsula, the mainland may be visible over water, but reachable only with a detour; similarly, in Figure 1,  $B$  can reach  $A$ , but not  $C$ . Peninsulas occur when an upstream provider of a multi-homed network accepts traffic but drops before delivery, when Tier-1 ISPs refuse to peer, or when firewalls block traffic. Peninsula existence has long been recognized, prompting overlay networks to route around them [2, 49, 50].

**Peninsulas in IPv6:** An long-term peninsula follows from the IPv6 peering dispute between Hurricane Electric (HE) and Cogent. These ISPs decline to peer in IPv6 (IPv4 is fine), nor do they forward their IPv6 through another party. HE and Cogent customers could not reach each other in 2009 [51], and this problem persists through 2024, as we show in DNSmon (§3.5). We further confirm unreachability between HE and Cogent users in IPv6 with traceroutes from looking glasses [19, 31] (HE at 2001:470:20::2 and Cogent at 2001:550:1:a::d): neither can reach their neighbor’s server, but both reach their own.

Other IPv6 disputes are Cogent with Google [64], and Cloudflare with Hurricane Electric [34]. Disputes are often due to an inability to agree to settlement-free or paid peering.

**Peninsulas in IPv4:** We observed a peninsula lasting 3 hours starting 2017-10-23t22:02Z, where five Polish Autonomous Systems (ASes) had 1716 /24 blocks that were always reachable one Los Angeles, but not from four other VPs (as seen in public data from Trinocular [80]). Before the peninsula, these blocks received service through Multimedia Polska (MP, AS21021), via Cogent (AS174), or through Tata (AS6453). When the peninsula occurred, traffic to all blocks continued through Cogent but was blackholed; it did not shift to Tata. The successful VP could reach MP through Tata for the entire event, proving MP was connected. After 3 hours, we see a burst of 23k BGP updates and MP is again reachable from all VPs. A graph showing reachability to this peninsula is in §A.

We confirmed this peninsula with additional observations from traceroutes taken by CAIDA’s Archipelago [12] (Ark). During the event we see 94 unique Ark VPs attempted 345 traceroutes to the affected blocks. Of the 94 VPs, 21 VPs (22%) have their last responsive traceroute hop in the same AS as the target address, and 68 (73%) stopped before reaching that AS. The remaining 5 VPs were able to reach the destination AS for only some traceroutes. The large number of BGP updates suggest routing problems as a root cause.

## 3 APPLYING THE DEFINITION

### 3.1 Resolving Conflicting Claims

Our definition of the Internet core in §2.1 must resolve conflicting claims without appeal to a central authority.

We can prove the definition yields a single core (or no core). Consider a connected component with some fraction  $A$ , where  $1 > A > 0.5$ . This component *must* be larger than any other component  $X$ , as proven by contradiction: (i) assume some  $X'$  exists, such that  $X' > A$ . (ii) Since  $A > 0.5$ , then (i) implies  $X' > 0.5$ . (iii) We then must conclude that  $A + X' > 1$ , but by definition, we measure only the whole address space, so it is also required that  $A + X' \leq 1$ . Therefore  $X' < A$  and  $A$  forces a single clear component. Q.E.D.

Disagree about what addresses are in the core can be resolved by comparing evidence. Consider a simplified version of Figure 1 with three pluralities of connectivity,  $A$ ,  $B$ , and  $C$ , each representing one third of the addresses, where both  $A$  and  $B$  and  $A$  and  $C$  are strongly and directly connected, but  $B$  and  $C$  cannot directly reach each other.

In this example  $A \cup B$  and  $A \cup C$  are partially, overlapping components of strong and direct connectivity, but since  $B$  and  $C$  cannot route to each other, they may dispute the core. From our definition, *all* ( $A \cup B \cup C$ ) are in the core, but  $B$  and  $C$  are on peninsulas. Any address can reach any other from either direction (the definition of “strongly connected”), but since  $B$  and  $C$  do not exchange routes, they are partially connected peninsulas (unless one purchases transit from  $A$ ). These definitions apply if the sizes are about equal ( $|A| = |B| = |C| = 0.33$ ) or are asymmetric ( $|A| = 0.49$  and  $|B| = |C| = 0.02$ ). Since the real Internet is mostly connected, typical values are  $|A| > 0.98$  and  $|B| < 0.01$ .

Resolving competing claims require that all parties present their evidence (what addresses  $A$ ,  $B$ , and  $C$  can reach, and that they agree those addresses have the same meaning. Private addresses and address squatting (described below) are cases where addresses have different meaning.

### 3.2 Policy Applications of the Definition

We next examine how a clear definition of the Internet core can inform policy tussles [16]. Our hope is that our conceptual definition can make sometimes amorphous concepts like “Internet fragmentation” more concrete, and an operational definition can quantify impacts and identify thresholds.

**Secession and Sovereignty:** The U.S. [70], China [3, 4], and Russia [17] have all proposed unplugging from the Internet. Egypt did in 2011 [20], and several countries have during exams [23, 30, 38, 45]. When the Internet partitions, which part is still “the Internet core”? Departure of an ISP or small country do not change the Internet core much, but what if a large country, or group of countries, leave together?

Our definition resolves this question, defining the Internet core from reachability of the majority of the active, public IP addresses (§2.1). Requiring a majority uniquely provides an unambiguous, externally evaluable test for the Internet core that allows one possible answer (the partition with more than 50%). In §3.4 we discuss the corollary: the Internet core can end, turning into multiple partitions, if none retain a majority. (A plurality is insufficient.)

**Sanction:** An opposite of secession is expulsion. Economic sanctions are one method of asserting international influence, and events such as the 2022 war in Ukraine prompted several large ISPs to discontinue service to Russia [66]. De-peering does not affect reachability for ISPs that purchase transit, but Tier-1 ISPs that de-peer create peninsulas for their users. As described below in §3.4, *no single country can eject another by de-peering with it*. However, a coalition of multiple countries could de-peer and eject a country from the Internet core if they, together, control more than half of the address space.

### 3.3 Architecture and Operation Application

Defining the core also helps clarify architectural changes such as the IPv4/v6 transition and operational address reuse.

**The IPv4/v6 Transition:** We have defined two Internet cores: IPv4 and IPv6. Our definition can determine when one supersedes the other. The networks will be on par when more than half of all IPv4 hosts are dual-homed. After that point, IPv6 will supersede IPv4 when a majority of hosts on IPv6 can no longer reach IPv4. Current limits on IPv6 measurement mean evaluation here is future work. IPv6 shows the strength and limits of our definition: since IPv6 is already economically important, our definition seems irrelevant. However, it may provide sharp boundary that makes the maturity of IPv6 definitive, helping motivate late-movers.

**Repurposing Addresses:** Given full allocation of IPv4, multiple parties proposed re-purposing currently allocated or reserved IPv4 space, such 0/8 (“this” network), 127/8 (loop-back), and 240/4 (reserved) [36]. New use of these long-reserved addresses is challenged by assumptions in widely-deployed, difficult to change, existing software and hardware. Our definition demonstrates that an RFC re-assigning this space for public traffic cannot make it a truly effective part of the Internet core until implementations used by a majority of active addresses can route to it.

**IPv4 Squat Space:** IP squatting is when an organization requiring private address space beyond RFC1918 takes over allocated but currently unrouted IPv4 space [6]. Several IPv4 /8s allocated to the U.S. DoD have been used this way [68] (they were only publicly routed in 2021 [78]). By our definition, such space is not part of the Internet core without public routes, and if more than half of the Internet is squatting on it, reclamation may be challenging.

**Internet outage detection:** Outage detect systems of often reported confusing observations with a mix of positive and negative responses to active probes, such as ThunderPing’s “hosed” state [75] and observer-local problems in Trinocular [63]. Partial connectivity suggests that sometimes conflicting observations may be valid and peninsulas should be recognized legitimate occurrences for future exploration.

**Failure mitigation via routing:** Several systems have proposed using relays to route around network-level routing failures [2, 49, 50]. In addition, hypergiants operating their own backbones can select routing egress to avoid partial connectivity [74]. An understanding of partial reachability in the WAN would quantify how important such efforts are.

### 3.4 Can the Internet Core Partition?

In §3.2 we discussed secession and expulsion qualitatively. Here we ask: Does any country or group have enough addresses to secede and claim to be “the Internet core” with a majority of addresses? Alternatively, if a country were to exert control over their allocated addresses, would they become a country-sized island or peninsula? We next use our reachability definition of more than 50% to quantify control of the IP address space.

To evaluate the power of countries and Regional Internet Registries (RIRs) over the Internet core, Table 1 reports the number of active IPv4 addresses as determined by Internet censuses [44] for RIRs and selected countries. Since estimating active IPv6 addresses is an open problem, we provide allocated addresses for both v4 and v6 [46, 60]. (IPv4 has been fully allocated since 2011 [47]).

Table 1 shows that *no individual RIR or country can secede and take the Internet core*, because none controls the majority of IPv4 addresses. ARIN has the largest share with 1673M allocated (45.2%). Of countries, U.S. has the largest share of allocated IPv4 (1617M, 43.7%). Active addresses are more evenly distributed with APNIC (223M, 33%) and the U.S. (40M, 21%) the largest RIR and country.

IPv6 is also an international collaboration, since no RIR or country exceeds 50% allocation. RIPE (an RIR) is close with 46.7%, and China and the U.S. have large country allocations. With most of IPv6 unallocated, these fractions may change.

IPv4 reflects a first-mover bias, where early adopters acquired many addresses, but this factor is smaller in IPv6. Our definition’s use of active addresses also reduces this bias, since numbers of *active* IPv4 addresses is similar to allocated IPv6 addresses (legacy IPv4 addresses are less used).

### 3.5 Definitions Clarify DNSmon Sensitivity

We next show how understanding partial connectivity can improve DNSmon sensitivity. DNSmon [1] monitors the Root Server System [71] from about 10k RIPE Atlas VPs (probes) [69]. For years, DNSmon has reported IPv6 loss

RIR	IPv4 Addresses		IPv6 Addresses	
	Active	Allocated	Active	Allocated
AFRINIC	15M	2%	121M	3.3%
APNIC	223M	33%	892M	24.0%
China	112M	17%	345M	9.3%
ARIN	150M	22%	1,673M	45.2%
U.S.	140M	21%	1,617M	43.7%
LACNIC	82M	12%	191M	5.2%
RIPE NCC	206M	30%	826M	22.3%
Germany	40M	6%	124M	3.3%
Total	676M	100%	3,703M	100%

**Table 1: RIR IPv4 hosts and IPv6 /32 alloc. [46, 60].**

rates of 4-10%, 4× higher than IPv4. The DNS root is well provisioned and distributed, so why is IPv6 loss so high?


RIPE Atlas operators are aware of problems with some Atlas VPs. Some VPs support IPv6 on their LAN, but not to the global IPv6 Internet—such VPs are IPv6 islands. Atlas periodically tags and culls these VPs from DNSmon. However, our study of DNSmon for islands and peninsulas improves their results. Using concepts pioneered here (§2), we give full analysis in a workshop paper [73]; Here we add new data showing these results persist for 1 year (Figure 2).

Groups of bars in Figure 3 show query loss for each of the 13 root service identifiers, as observed from all available Atlas VPs (10,082 IPv4, and 5,173 IPv6) on 2022-07-23. (We are similar to DNSmon, but it uses only about 100 well-connected “anchors”, so our analysis is wider.) The first two groups show loss rates for IPv4 (light blue, left most) and IPv6 (light red), showing IPv4 losses around 2%, and IPv6 from 9 to 13%.


We report a VP as an island when it cannot see *any* of the 13 root identifiers over 24 hours. (This definition is stricter than our 50% definition since VPs attempt only 13 targets, not the whole Internet, and we apply it over a full day to consider only long-term trends.) The middle two groups of bars show IPv4 and IPv6 loss rates after removing VPs that are islands. Without island VPs, IPv4 loss rates drop to 0.005 from 0.01, and IPv6 to about 0.01 from 0.06. These rates represent a more meaningful estimate of DNS reliability. Users of VPs that are IPv6 islands will not expect global IPv6, and such VPs should not be used for IPv6 in DNSmon.

The third bar in each red cluster of IPv6 is an outlier: that root identifier shows 13% IPv6 loss with all VPs, and 6% loss after islands are removed. This result is explained by persistent routing disputes between Cogent (the operator of C-Root) and Hurricane Electric [57]. Omitting islands (the middle bars) makes this difference much clearer.

Finally we detect peninsulas by looking for VPs that each some but not all root servers. Peninsulas suggest persistent routing problems; they deserve attention from ISPs and root operators. The darker, rightmost two groups show loss from non-island/peninsula VPs, representing loss if routing problems were addressed. With this correction C-Root is similar to others, confirming peering disputes affect its success.



**Figure 2: Fraction of VPs observing islands and peninsulas for IPv4 and IPv6 during 2022.**



**Figure 3: Atlas queries from all available VPs to 13 Root Servers for IPv4 and IPv6 on 2022-07-23.**

This example shows how *understanding partial reachability can improve the sensitivity of existing measurement systems*. Removing islands makes it easy to identify persistent routing problems. Removing peninsulas makes transient changes (perhaps from failure, DDoS, routing) more visible. Each layer of problem is important, but by considering each separately, the interesting “signal” of routing changes (appearing in the right two groups in Figure 3), appears out from under the 5× or 9.7× times larger peninsulas and islands (the left two groups). Improved sensitivity also *shows a need to improve IPv6 provisioning*, since IPv6 loss is statistically higher than IPv4 loss (compare the right blue and red groups), even accounting for known problems. After sharing the results with root operators and RIPE Atlas, two operators adopted them in regular operation.

## 4 RELATED WORK

Prior definitions of the Internet exist at the IP-layer [13, 32, 33, 62] of their time, or the AS-level [37, 55]. IPNL proposed a core-only Internet with all users behind NAT [35]. We instead consider the IP-layer in today’s architecture to address today’s challenges (see §2).

Several systems mitigate partial outages. RON provides alternate-path routing around failures for a mesh of sites [2]. Hubble monitors multi-path reachability [49]. LIFEGUARD routes around reachability failures [50]. These systems address partial reachability; we define its scope.

Prior work studied partial reachability, showing it is a common transient occurrence during routing convergence [11]. They reproduced partial connectivity with controlled experiments; we study it in RIPE Atlas.

Active outage detection systems have encountered partial outages. ThunderPing recognizes a “hosed” state with mixed replies, but its study is future work [75]. Trinocular discards partial outages by reporting the target block “up” if any VP can reach it [63]. To the best of our knowledge, prior outage detection systems do not consistently report partial outages in the Internet core, nor do they study their extent.

We use the idea of majority to define the Internet core in the face of secession. That idea is fundamental in many algorithms for distributed consensus [52, 53, 58], with applications for example to certificate authorities [10].

Recent groups have studied the policy issues around Internet fragmentation [26, 27], but do not define it. We hope our definition can fill that need.

## 5 CONCLUSIONS

This paper affirms the importance of an Internet core for global communication, and provides a robust, operationalizable definition of that core. The definition helps identify disconnected islands and shows that partially connected peninsulas are an important challenge. The definition helps clarify what events would cause the Internet core to fragment. They also help help improve the sensitivity of operational measurement systems such as RIPE DNSmon, by distinguishing long-term partial reachability from short-term changes.

## ACKNOWLEDGMENTS

The authors would like to thank John Wroclawski, Wes Hardaker, Ramakrishna Padmanabhan, Ramesh Govindan, Eddie Kohler, and the Internet Architecture Board for their input on an early version of this paper.

The work is supported in part by the National Science Foundation, CISE Directorate, award CNS-2007106 and NSF-2028279. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

## REFERENCES

- [1] Christopher Amin, Massimo Cándela, Daniel Karrenberg, Robert Kistelevi, and Andreas Strikos. 2015. Visualization and Monitoring for the Identification and Analysis of DNS Issues. In *Proceedings of the International Conference on the Internet Monitoring and Protection*. Brussels, Belgium. [https://www.researchgate.net/profile/Massimo-Candela/publication/279516870\\_Visualization\\_and\\_Monitoring\\_for\\_the\\_Identification\\_and\\_Analysis\\_of\\_DNS\\_Issues/links/559468c808ae793d13798901/Visualization-and-Monitoring-for-the-Identification-and-Analysis-of-DNS-Issues.pdf](https://www.researchgate.net/profile/Massimo-Candela/publication/279516870_Visualization_and_Monitoring_for_the_Identification_and_Analysis_of_DNS_Issues/links/559468c808ae793d13798901/Visualization-and-Monitoring-for-the-Identification-and-Analysis-of-DNS-Issues.pdf)
- [2] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. 2001. Resilient Overlay Networks. In *Proceedings of the Symposium on Operating Systems Principles*. ACM, Chateau Lake Louise, Alberta, Canada, 131–145. <http://www-cse.ucsd.edu/sosp01/papers/andersen.pdf>
- [3] Anonymous. 2012. The collateral damage of Internet censorship by DNS injection. *ACM Computer Communication Review* 42, 3 (July 2012), 21–27. <https://doi.org/10.1145/2317307.2317311>
- [4] Anonymous. 2014. Towards a Comprehensive Picture of the Great Firewall’s DNS Censorship. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet*. USENIX, San Diego, CA, USA, 7. <https://www.usenix.org/system/files/conference/foci14/foci14-anonymous.pdf>
- [5] ANT Project. 2022. ANT IPv4 Island and Peninsula Data. [https://ant.isi.edu/datasets/ipv4\\_partial/](https://ant.isi.edu/datasets/ipv4_partial/). [https://ant.isi.edu/datasets/ipv4\\_partial/](https://ant.isi.edu/datasets/ipv4_partial/)
- [6] Cathy Aronson. 2015. To Squat Or Not To Squat? blog <https://teamarin.net/2015/11/23/to-squat-or-not-to-squat/>. <https://teamarin.net/2015/11/23/to-squat-or-not-to-squat/>
- [7] Guillermo Baltra and John Heidemann. 2020. Improving Coverage of Internet Outage Detection in Sparse Blocks. In *Proceedings of the Passive and Active Measurement Workshop*. Springer, Eugene, Oregon, USA.
- [8] Guillermo Baltra and John Heidemann. 2023. *Detecting Partial Reachability in the Internet Core*. Technical Report arXiv:2107.11439v4. USC/Information Sciences Institute. <https://doi.org/10.48550/2107.11439v4>
- [9] Genevieve Bartlett, John Heidemann, and Christos Papadopoulos. 2007. Understanding Passive and Active Service Discovery. In *Proceedings of the ACM Internet Measurement Conference*. ACM, San Diego, California, USA, 57–70. <https://doi.org/10.1145/1298306.1298314>
- [10] Henry Birge-Lee, Yixin Sun, Anne Edmundson, Jennifer Rexford, and Prateek Mittal. 2018. Bambooizing certificate authorities with BGP. In *27th USENIX Security Symposium*. USENIX, Baltimore, Maryland, USA, 833–849.
- [11] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2009. Internet optometry: assessing the broken glasses in Internet reachability. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*. ACM, Chicago, Illinois, USA, 242–253. <http://www.maennel.net/2009/imc099-bush.pdf>
- [12] CAIDA. 2007. Archipelago (Ark) Measurement Infrastructure. website <https://www.caida.org/projects/ark/>. <https://www.caida.org/projects/ark/>
- [13] Vint Cerf and Robert Kahn. 1974. A Protocol for Packet Network Interconnection. *IEEE Transactions on Communications* COM-22, 5 (May 1974), 637–648. <http://sysnet.ucsd.edu/classes/cse222/wi03/papers/cerf-tcp-toc74.pdf>
- [14] S. Cheshire and M. Krochmal. 2013. *NAT Port Mapping Protocol (NAT-PMP)*. RFC 6886. Internet Request For Comments. <https://doi.org/10.17487/RFC6886>
- [15] David D. Clark. 1988. The Design Philosophy of the DARPA Internet Protocols. In *Proceedings of the 1988 Symposium on Communications Architectures and Protocols*. ACM, 106–114.
- [16] David D. Clark, John Wroclawski, Karen Sollins, and Robert Braden. 2002. Tussle in Cyberspace: Defining Tomorrow’s Internet. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Pittsburgh, PA, USA, 347–356. <http://www.acm.org/sigcomm/sigcomm2002/papers/tussle.pdf>
- [17] CNBC. 2019. Russia just brought in a law to try to disconnect its Internet from the rest of the world. <https://www.cnbc.com/2019/11/01/russia-controversial-sovereign-internet-law-goes-into-force.html>
- [18] N. Coca. 2018. China’s Xinjiang surveillance is the dystopian future nobody wants. *Engadget* (Feb. 22 2018). <https://www.engadget.com/2018-02-22-china-xinjiang-surveillance-tech-spread.html>
- [19] Cogent. 2021. Looking Glass. <https://cogentco.com/en/looking-glass>.
- [20] James Cowie. 2011. Egypt Leaves the Internet. Renesys Blog <http://www.renesys.com/blog/2011/01/egypt-leaves-the-internet.shtml>. <http://www.renesys.com/blog/2011/01/egypt-leaves-the-internet.shtml>
- [21] RBC daily. 2021. Russia, tested the Runet when disconnected from the Global Network. website [https://www.rbc.ru/technology\\_and\\_media/21/07/2021/60f8134c9a79476f5de1d739](https://www.rbc.ru/technology_and_media/21/07/2021/60f8134c9a79476f5de1d739). [https://www.rbc.ru/technology\\_and\\_media/21/07/2021/60f8134c9a79476f5de1d739](https://www.rbc.ru/technology_and_media/21/07/2021/60f8134c9a79476f5de1d739)
- [22] Alberto Dainotti, Karyn Benson, Alistair King, kc claffy, Eduard Glatz, Xenofontas Dimitropoulos, Philipp Richter, Alessandro Finamore, and Alex C. Snoeren. 2016. Lost in Space: Improving Inference of IPv4 Address Space Utilization. *IEEE Journal of Selected Areas in Communication* 34, 6 (April 2016), 1862–1876. <https://doi.org/10.1109/JSAC.2016.2559218>
- [23] Dhaka Tribune Desk. 2018. Internet services to be suspended across the country. *Dhaka Tribune* (Feb. 11 2018). <http://www.dhakatribune.com/regulation/2018/02/11/internet-services-suspended-throughout-country/>

- [24] Amogh Dhamdhere, David D. Clark, Alexander Gamero-Garrido, Matthew Luckie, Ricky K. P. Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex C. Snoeren, and kc claffy. 2018. Inferring Persistent Interdomain Congestion. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Budapest, Hungary, 1–15. <https://doi.org/10.1145/3230543.3230549>
- [25] DINRG. 2021. Decentralized Internet Infrastructure Research Group. <https://irtf.org/dinrg>.
- [26] William J. Drake, Vinton G. Cerf, and Wolfgang Kleinwächter. 2016. *Internet Fragmentation: An Overview*. Technical Report. World Economic Forum. [https://www3.weforum.org/docs/WEF\\_FII\\_Internet\\_Fragmentation\\_An\\_Overview\\_2016.pdf](https://www3.weforum.org/docs/WEF_FII_Internet_Fragmentation_An_Overview_2016.pdf)
- [27] William J. Drake (moderator). 2022. Internet Fragmentation, Reconsidered. CITI Seminar on Global Digital Governance at IETF 115. <https://www8.gsb.columbia.edu/citi/GlobalDigitalGovernance>
- [28] Peter K. Dunn. 2021. Scientific Research Methods. <https://bookdown.org/pkaldunn/Book/>.
- [29] Zakir Durumeric, Michael Bailey, and J Alex Halderman. 2014. An Internet-wide view of Internet-wide scanning. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*. USENIX, San Diego, California, USA, 65–78. <https://jhalderm.com/pub/papers/scanning-sec14.pdf>
- [30] Economist Editors. 2018. Why some countries are turning off the internet on exam days. *The Economist* (July 5 2018). <https://www.economist.com/middle-east-and-africa/2018/07/05/why-some-countries-are-turning-off-the-internet-on-exam-days> (Appeared in the Middle East and Africa print edition).
- [31] Hurricane Electric. 2021. Looking Glass. <http://lg.he.net/>.
- [32] Engadget. 2020. China, Huawei propose internet protocol with a built-in killswitch. <https://www.engadget.com/2020-03-30-china-huawei-new-ip-proposal.html>.
- [33] Federal Networking Council (FNC). 1995. Definition of “Internet”. [https://www.nitrd.gov/historical/fnc/internet\\_res.pdf](https://www.nitrd.gov/historical/fnc/internet_res.pdf).
- [34] HE forums. 2017. Cloudflare Blocked on Free Tunnels now? <https://forums.he.net/index.php?topic=3805.0>.
- [35] Paul Francis and Ramakrishnan Gummadi. 2001. IPNL: A NAT-extended internet architecture. In *Proceedings of the ACM SIGCOMM Conference*. ACM, San Diego, CA, USA, 69–80. <https://doi.org/10.1145/964723.383065>
- [36] V. Fuller, E. Lear, and D. Meyer. 2008. Reclassifying 240/4 as usable unicast address space. (March 2008). <https://datatracker.ietf.org/doc/html/draft-fuller-240space-02> Work in progress (Internet draft draft-fuller-240space-02.txt).
- [37] Lixin Gao. 2001. On Inferring Autonomous System Relationships in the Internet. *ACM/IEEE Transactions on Networking* 9, 6 (Dec. 2001), 733–745. <https://doi.org/10.1109/90.974527>
- [38] Samuel Gibbs. 1996. Iraq shuts down the Internet to stop pupils cheating in exams. *The Guardian* (18 May 1996). <https://www.theguardian.com/technology/2016/may/18/iraq-shuts-down-internet-to-stop-pupils-cheating-in-exams>
- [39] GovTrack.us. 2020. Unplug the Internet Kill Switch Act would eliminate a 1942 law that could let the president shut down the internet. <https://govtrackinsider.com/unplug-the-internet-kill-switch-act-would-eliminate-a-1942-law-that-could-let-the-president-shut-78326f0ef66c>. <https://govtrackinsider.com/unplug-the-internet-kill-switch-act-would-eliminate-a-1942-law-that-could-let-the-president-shut-78326f0ef66c>
- [40] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, and Parveen Pat. 2009. VL2: A Scalable and Flexible Data Center Network. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Barcelona, Spain, 51–62. <http://ccr.sigcomm.org/online/files/p51.pdf>
- [41] James Griffiths. 2019. Democratic Republic of Congo internet shutdown shows how Chinese censorship tactics are spreading. *CNN* (Jan. 2 2019). <https://edition.cnn.com/2019/01/02/africa/congo-internet-shutdown-china-intl/index.html>
- [42] Andreas Guillot, Romain Fontugne, Philipp Winter, Pascal Merindol, Alistair King, Alberto Dainotti, and Cristel Pelsser. 2019. Chocolate: Outage Detection for Internet Background Radiation. In *Proceedings of the IFIP International Workshop on Traffic Monitoring and Analysis*. IFIP, Paris, France, 8. <https://clarinet.u-strasbg.fr/~pelsser/publications/Guillot-chocolate-TMA2019.pdf>
- [43] Hang Guo and John Heidemann. 2018. Detecting ICMP Rate Limiting in the Internet. In *Proceedings of the Passive and Active Measurement Workshop*. Springer, Berlin, Germany, to appear.
- [44] John Heidemann, Yuri Pradkin, Ramesh Govindan, Christos Papadopoulos, Genevieve Bartlett, and Joseph Bannister. 2008. Census and Survey of the Visible Internet. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Vouliagmeni, Greece, 169–182. <https://doi.org/10.1145/1452520.1452542>
- [45] Jon Henley. 2018. Algeria blocks internet to prevent students cheating during exams. *The Guardian* (22 June 2018). <https://www.theguardian.com/world/2018/jun/21/algeria-shuts-internet-prevent-cheating-school-exams>
- [46] IANA. 2021. IPv6 RIR Allocation Data. <https://www.iana.org/numbers/allocations/>.
- [47] ICANN. 2011. *Available Pool of Unallocated IPv4 Internet Addresses Now Completely Emptied*. Announcement. ICANN. <https://itp.cdn.icann.org/en/files/announcements/release-03feb11-en.pdf>
- [48] Internet Architecture Board. 2000. *IAB Technical Comment on the Unique DNS Root*. RFC 2826. Internet Request For Comments. <https://www.rfc-editor.org/rfc/rfc2826>
- [49] Ethan Katz-Bassett, Harsha V Madhyastha, John P John, Arvind Krishnamurthy, David Wetherall, and Thomas E Anderson. 2008. Studying Black Holes in the Internet with Hubble. In *Proceedings of the USENIX Conference on Networked Systems Design and Implementation*. ACM, San Francisco, CA, 247–262.
- [50] Ethan Katz-Bassett, Colin Scott, David R. Choffnes, Ítalo Cunha, Vytautas Valancius, Nick Feamster, Harsha V. Madhyastha, Tom Anderson, and Arvind Krishnamurthy. 2012. LIFEGUARD: Practical Repair of Persistent Route Failures. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Helsinki, Finland, 395–406. <https://doi.org/10.1145/2377677.2377756>
- [51] DataCenter Knowledge. 2009. Peering Disputes Migrate to IPv6. <https://www.datacenterknowledge.com/archives/2009/10/22/peering-disputes-migrate-to-ipv6>.
- [52] Leslie Lamport. 1998. The Part-Time Parliament. *ACM Transactions on Computer Systems* 16, 2 (May 1998), 133–169. <https://doi.org/10.1145/279227.279229>
- [53] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems* 4, 3 (July 1982), 382–401.
- [54] Mark A. Lemley and Lawrence Lessig. 2001. The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era. *UCLA Law Review* 48 (April 2001), 925–972. <https://law.stanford.edu/publications/the-end-of-end-to-end-preserving-the-architecture-of-the-internet-in-the-broadband-era/>
- [55] Matthew Luckie, Bradley Huffaker, Dhamdhere, Vasileios Giotsas, and kc claffy. 2013. AS Relationships, Customer Cones, and Validation. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Barcelona, Spain, 243–256. <http://conferences.sigcomm.org/imc/2013/papers/imc039-luckieAemb.pdf>
- [56] Brent A. Miller, Toby Nixon, Charlie Tai, and Mark D. Wood. 2001. Home Networking with Universal Plug and Play. *IEEE Communications*




- Magazine* 39, 12 (Dec. 2001), 104–109. <https://doi.org/10.1109/35.968819>
- [57] Rich Miller. 2009. Peering Disputes Migrate to IPv6. website <https://www.datacenterknowledge.com/archives/2009/10/22/peering-disputes-migrate-to-ipv6>. <https://www.datacenterknowledge.com/archives/2009/10/22/peering-disputes-migrate-to-ipv6>
- [58] Satoshi Nakamoto. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. Released publically <http://bitcoin.org/bitcoin.pdf>. <http://bitcoin.org/bitcoin.pdf>
- [59] BBC News. 2019. Russia internet: Law introducing new controls comes into force. website <https://www.bbc.com/news/world-europe-50259597>. <https://www.bbc.com/news/world-europe-50259597>
- [60] NRO. 2021. IPv4 Address Space Registry. <https://www.nro.net/about/rirs/statistics/>.
- [61] C. Partridge, T. Mendez, and W. Milliken. 1993. *Host Anycasting Service*. RFC 1546. Internet Request For Comments. <https://www.rfc-editor.org/rfc/rfc1546.txt>
- [62] Jonathan B. Postel. 1980. Internetwork Protocol Approaches. *IEEE Trans. Comput.* 28, 4 (April 1980), 604–611. <https://doi.org/10.1109/TCOM.1980.1094705>
- [63] Lin Quan, John Heidemann, and Yuri Pradkin. 2013. Trinocular: Understanding Internet Reliability Through Adaptive Probing. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Hong Kong, China, 255–266. <https://doi.org/10.1145/2486001.2486017>
- [64] Dan Rayburn. 2016. Google Blocking IPv6 Adoption With Cogent, Impacting Transit Customers. <https://seekingalpha.com/article/3948876-google-blocking-ipv6-adoption-cogent-impacting-transit-customers>.
- [65] Reuters. 2021. Russia disconnected from internet in tests as it bolsters security. website <https://www.reuters.com/technology/russia-disconnected-global-internet-tests-rbc-daily-2021-07-22/>. <https://www.reuters.com/technology/russia-disconnected-global-internet-tests-rbc-daily-2021-07-22/>
- [66] Reuters. 2022. website <https://www.reuters.com/technology/us-firm-cogent-cutting-internet-service-russia-2022-03-04/>. <https://www.reuters.com/technology/us-firm-cogent-cutting-internet-service-russia-2022-03-04/>
- [67] Philipp Richter, Ramakrishna Padmanabhan, Neil Spring, Arthur Berger, and David Clark. 2018. Advancing the Art of Internet Edge Outage Detection. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Boston, Massachusetts, USA, 350–363. <https://doi.org/10.1145/3278532.3278563>
- [68] Philipp Richter, Florian Wohlfart, Narseo Vallina-Rodriguez, Mark Allman, Randy Bush, Anja Feldmann, Christian Kreibich, Nicholas Weaver, and Vern Paxson. 2016. A Multi-perspective Analysis of Carrier-Grade NAT Deployment. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Santa Monica, CA, USA. <https://doi.org/10.1145/2987443.2987474>
- [69] RIPE NCC Staff. 2015. RIPE Atlas: A Global Internet Measurement Network. *The Internet Protocol Journal* 18, 3 (Sept. 2015), 2–26.
- [70] Sen. John D. Rockefeller. 2009. Cybersecurity Act of 2010. <https://www.congress.gov/bill/111th-congress/senate-bill/773>.
- [71] Root Operators. 2016. <http://www.root-servers.org>.
- [72] J. Rosenberg, J. Weinberger, C. Huitema, and R. Mahy. 2003. *STUN—Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)*. RFC 3489. Internet Request For Comments. <ftp://ftp.rfc-editor.org/in-notes/rfc3489.txt>
- [73] Tarang Saluja, John Heidemann, and Yuri Pradkin. 2022. Differences in Monitoring the DNS Root Over IPv4 and IPv6. In *Proceedings of the National Symposium for NSF REU Research in Data Science, Systems, and Security*. IEEE, Portland, OR, USA, to appear.
- [74] Brandon Schlinker, Hyojeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V. Madhyastha, Italo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. 2017. Engineering Egress with Edge Fabric: Steering Oceans of Content to the World. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Los Angeles, CA, USA, 418–431. <https://doi.org/10.1145/3098822.3098853>
- [75] Aaron Schulman and Neil Spring. 2011. Pingin’ in the Rain. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Berlin, Germany, 19–25. <https://doi.org/10.1145/2068816.2068819>
- [76] Anant Shah, Romain Fontugne, Emile Aben, Cristel Pelsser, and Randy Bush. 2017. Disco: Fast, Good, and Cheap Outage Detection. In *Proceedings of the IEEE International Conference on Traffic Monitoring and Analysis*. Springer, Dublin, Ireland, 1–9. <https://doi.org/10.23919/TMA.2017.8002902>
- [77] Berhan Taye and Sage Cheng. 2019. Report: the state of internet shutdowns. blog <https://www.accessnow.org/the-state-of-internet-shutdowns-in-2018/>. <https://www.accessnow.org/the-state-of-internet-shutdowns-in-2018/>
- [78] Craig Timberg and Paul Sonne. 2021. Minutes before Trump left office, millions of the Pentagon’s dormant IP addresses sprang to life. *The Washington Post* (Apr. 24 2021). <https://www.washingtonpost.com/technology/2021/04/24/pentagon-internet-address-mystery/>
- [79] Paul F. Tsuchiya and Tony Eng. 1993. Extending the IP Internet Through Address Reuse. *ACM Computer Communication Review* 23, 1 (Jan. 1993), 16–33. <http://www.cs.cornell.edu/People/francis/tsuchiya93extending.pdf>
- [80] USC/LANDER Project. 2017. Internet Outage Measurements. IMPACT ID USC-LANDER/LANDER:internet\_outage\_adaptive\_a30all-20171006 at [https://ant.isi.edu/datasets/internet\\_outages/](https://ant.isi.edu/datasets/internet_outages/).
- [81] Sebastian Zander, Lachlan L. H. Andrew, and Grenville Armitage. 2014. Capturing Ghosts: Predicting the Used IPv4 Space by Inferring Unobserved Addresses. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Vancouver, BC, Canada, 319–332. <https://doi.org/10.1145/2663716.2663718>

## A EXAMPLES

In §2.2 we define islands and peninsulas as two cases of partial connectivity in the Internet. Here we give two real-world examples of islands and peninsulas that we discovered in Trinocular data.

For our example island (Figure 4) and peninsula (Figure 5), we show a graph that counts the number of active IP addresses that are reachable from the 6 Trinocular observers. Each line represents the best estimate of the current number of active addresses from each observer. Most of the lines generally overlap and show a V-shaped dip during the island, but one (e, the green line), stands out as fairly stable over this period. Because we probe only a few addresses per round, the estimate of active addresses updates slowly, and lags the true value after an abrupt change in reachability.

While *visually* these graphs look the same, with one VP able to reach the destination block while all others fail, we distinguish the island from the peninsulas with additional information. For the example island, the destination block also hosts VP E, so we know with confidence E can reach itself but cannot reach the rest of the Internet, making it an island. For the example peninsula, all VPs are external to the



**Figure 4: Estimates of an island in Trinocular data starting 2017-06-03t23:06Z and lasting just longer than one hour.**

unreachable block. We confirm that it is a peninsula using BGP and traceroutes. We can therefore distinguish them as island and peninsula both from the Trinocular observations, and to confirm that claim with external data sources.


### A.1 An Example Island

We defined island in §2.2.2, and looked for the systematically in Trinocular data (as described elsewhere [8]).

Figure 4 shows an example island we discovered. This island begins at 2017-06-03t23:06Z and lasts about 64 minutes.

For this hour, VP E and this network are part of an island, cut off from the rest of the internet and other VPs. Because VP “E” is inside the island, it always sees 13 (or 14) active IP addresses in its top green line Figure 4. By contrast, the other 5 sites see a steady-state of 10 VPs, dropping to 0 during the island, as the gradually re-scan and fail to reach previously active IP addresses. Other VPs rediscover all 10 addresses over the next two hours after the island ends. (VP E sees 3 more IP active addresses that the other VPs for this block presumably because those targets have firewall rules that only allow replies to sources originating from the same block.)

The network being scanned here was the same network block hosting VP E, and we confirmed that this network was disconnected from network operators. During the hour-long island VP E had 5 Trinocular rounds to scan the whole Internet, and it concluded that about 80% of the Internet was unreachable. It actually had failures to all of the Internet, but it dismisses brief unreachability to 20% of blocks due to a conservative choice from the FBS algorithm [7]. (This algorithm requires that outages in sparsely active blocks are only confirmed after all probed addresses in the block respond



**Figure 5: Estimates of reachable addresses during a peninsula found in Trinocular data starting 2017-10-23t22:02Z and lasting about 3 hours.**

negatively. Incremental scanning all addresses in about 20% of blocks takes longer than one hour in this dataset.)

### A.2 An Example Peninsula

We defined peninsulas in §2.2.3, and discussed an example that occurred in 2017-10-23 in Poland.

This peninsula was discovered by algorithms we developed (described elsewhere [8]). Figure 5 shows our best estimates of the number of responsive addresses from each of our 6 observers. We believe this block has 78 active addresses (based on about 3 years of history), and we scan a mean of 2.3 addresses each 11-minute round.

We believe the peninsula begins at 2017-10-23t22:02:24, the time the first observer (n) has no successful queries, and ends 2h46m later when another observer (e) is successful. The w observer is successful for the whole period, as shown by its orange line staying fixed at 27 during the peninsula, confirming that its queries sent reach a responsive address every 660 s. By contrast, the other five observers have no positive responses during the peninsula, so their active address count drops to zero by midnight. After the peninsula ends, queries the five sites are again successful and the estimate of active addresses climbs slow back to 34 just after t06:00.

To explain why the observed number of active addresses lags true reachability, recall that, in each round, Trinocular probes the minimum number of addresses to confirm block reachability. After the peninsula recovers, this means one address each round confirms reachability, explaining the the 33 rounds (5.5 h) it takes to return to 34 active addresses. (Each site actually probes two addresses in the round just after recovery.) Negative information is acquired more quickly, because each site requires 3 to 5 negative responses to confirm

unreachability during the peninsula. Although Trinocular is willing to send up to 15 queries per round, this block usually full and stable (typically 34 of 78 addresses responding, giving an expected response around around  $A = 0.44$ ), so it can conclude unreachability from fewer queries.

This example is representative of other peninsulas we have seen.