

Understanding Partial Reachability in the Internet Core

Guillermo Baltra 

USC/ISI, Marina del Rey, California, USA

Tarang Saluja 

Swarthmore College, Swarthmore, Pennsylvania, USA

Yuri Pradkin 

USC/ISI, Marina del Rey, California, USA

John Heidemann  

USC/ISI, Marina del Rey, California, USA

Abstract

Routing strives to connect all the Internet, but compete: political pressure threatens routing fragmentation; architectural changes such as private clouds, carrier-grade NAT, and firewalls make connectivity conditional; and commercial disputes create partial reachability for days or years. This paper suggests *persistent, partial reachability is fundamental to the Internet* and an underexplored problem. We first *derive a conceptual definition of the Internet core* based on connectivity, not authority. We identify *peninsulas*: persistent, partial connectivity; and *islands*: when computers are partitioned from the Internet core. Second, we develop algorithms to observe each across the Internet, and apply them to two existing measurement systems: Trinocular, where 6 locations observe 5M networks frequently, and RIPE Atlas, where 13k locations scan the DNS roots frequently. Cross-validation shows our findings are stable over *three years of data*, and consistent with as few as 3 geographically-distributed observers. We validate peninsulas and islands against CAIDA Ark, showing good recall (0.94) and bounding precision between 0.42 and 0.82. Finally, our work has broad practical impact: we show that *peninsulas are more common than Internet outages*. Factoring out peninsulas and islands as noise can *improve existing measurement systems*; their “noise” is 5× to 9.7× larger than the operational events in RIPE’s DNSmon. We show that most peninsula events are routing transients (45%), but most peninsula-time (90%) is due to a few (7%) long-lived events. Our work helps inform Internet policy and governance, with our neutral definition showing no single country or organization can unilaterally control the Internet core.

2012 ACM Subject Classification Networks → Network measurement; Networks → Public Internet; Networks → Network economics

Keywords and phrases Internet, Internet reliability, Network outages, Active measurements

Digital Object Identifier [10.4230/OASICS.NINeS.2026.4](https://doi.org/10.4230/OASICS.NINeS.2026.4)

Funding The work is supported in part by the National Science Foundation, CISE Directorate, award CNS-2007106 NSF-2028279, OAC-2530698, and DARPA Contract HR001124C0403. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF or DARPA. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

Acknowledgements 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. We would also like to thank the NINeS chairs and steering committee for building a venue that is receptive for this kind of paper, and to our shepherd, Michael Schapira, and the reviewers for their input.



© Guillermo Baltra, Tarang Saluja, Yuri Pradkin, John Heidemann;
licensed under Creative Commons License CC-BY 4.0

1st New Ideas in Networked Systems (NINeS 2026).

Editors: Katerina J. Argyraki and Aurojit Panda; Article No. 4; pp. 4:1–4:31



Open Access Series in Informatics

OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The Internet was created to allow disparate networks to communicate [18, 73, 20], making *network partition* its nemesis. Routing is designed to heal partitions, so that “communication must continue despite loss of networks or gateways” [20]. Yet the reality of partitions prompts leadership-election algorithms such as Paxos [60].

Worse than complete network partition is *long-lived partial reachability*. Although transient reachability problems are well known (for example, [100]), and human errors occur [64], *policy choices* can cause persistent partial connectivity. Economic differences result in peering disputes [67, 76, 42]; while political choices can limit access [80], or emphasize sovereignty [71, 26, 79]. Research [2, 57, 58] and production [89, 102] work around persistent unreachability.

Challenges: But today *universal reachability in the Internet core is often challenged*: Political pressure may Balkanize the Internet along national borders. Examples include Russia’s 2019 sovereign-Internet law [71, 26, 79] and national “Internet kill switches” that are debated in U.S. [46] and the U.K., and deployed elsewhere [25, 23, 48, 93]. These pressures prompted policy discussions about fragmentation [33, 34]. We suggest that *technical methods can help inform policy discussions* and that threats such as de-peering place the global Internet at risk. We will show that no single country can unilaterally control the Internet core today (§6.2), and that de-peering *can* fragment the Internet core into pieces (§6.1).

Architecturally, 25 years of evolution have segmented the Internet core: many services live in clouds; users are usually second-class clients due to Network Address Translation (NAT); firewalls interrupt connectivity; and Internet has both IPv4 and IPv6. Politics can influence architecture, with China’s Great Firewall [4, 5], and a proposed “new Internet” [39]. We suggest that technical methods help us *reason about changes to Internet architecture*, to understand implications of partial reachability and evaluate IPv6 deployment.

Operationally, even when ISP peering is mature, disputes can cause long-term partial unreachability [67]. Such unreachability detected experimentally [31], and systems built to mitigate partial reachability [2, 57, 58]. We show several operational uses of our work. We show that *accounting for partial reachability can make existing measurement systems more sensitive*. By applying these results to widely used RIPE DNSmon (§6.3), we show that its observations of persistent high query loss (5–8% to the DNS Root [85]) are mostly measurement error and persistent partial connectivity. These factors are $5\times$ and $9.7\times$ (IPv4 and v6) larger than operationally important signals. Our analysis also helps resolve uncertainty in Internet outage detection (§6.2), clarifying “corner cases” due to conflicting observations [90, 75, 91, 81, 49]. We show partial reachability is a common cause, and it occurs at least as often as complete outages (§5.1). Finally, our work helps quantify the applicability of route-failure mitigation [2, 57, 58], and of cloud egress selection [89].

Contributions: Our first contribution is to *recognize that partial reachability is a fundamental part of the Internet*, and addressing it requires a *rigorous definition of what is the Internet’s core* (§2). In 1982, the Internet was 83 hosts [92] globally reachable with TCP/IP [73]. 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 [41]. Later work added DNS [56] and IPv6. But today’s Internet is much changed: Both users on PCs and the majority of users on mobile devices access the Internet indirectly through NAT [96] and Carrier-Grade NAT (CG-NAT) [82]. Many public services operate from the cloud, visible through rented or imported IP addresses, backed by network virtualization [47]. Media is replicated in Content Delivery Networks (CDNs). Access is mediated by firewalls.

data source	num. VPs	measurement		
		freq.	targets	duration
Trinocular [75]	6 ^a	11 min.	5M /24s	4 years
RIPE Atlas [83]	12,086 ^b	5 min.	13 RSOs	3 years
CAIDA Ark [14]	171 ^c	24 hrs.	all IPv4	selected
Routeviews [65]	55 ^d	1 hour	all IPv4	selected

a: In 2017 and 2019. b: On 2024-01-30. c: On 2017-12-01. d: In 2024-01.

■ **Table 1** Types of data sources used in this paper.

Yet users find Internet services so seamless that technology recedes and the web, Facebook, and phone apps are their “Internet”.

We define *Internet core* as the strongly-connected component of more than 50% of active, public IP addresses that can bidirectionally route to each other (§2.1). This definition has several unique characteristics. First, captures the uniform, *peer-to-peer nature of the Internet core* necessary for first-class services. Second, it defines *one, unique* Internet core by requiring reachability of more than 50%—there can be only one since multiple majorities are impossible. Finally, unlike prior work, this *conceptual* definition avoids dependence on any specific measurement system, nor does it depend on historical precedent, special locations, or central authorities. Although an operational measurements will reflect observation error, the conceptual Internet core defines an asymptote against which our current and future measurements can compare, unlike prior definitions from specific systems [2, 57, 58].

Our second contribution is to use this definition to identify two classes of persistent unreachability (§2.3), and develop algorithms to quantify each (§3). We define *peninsulas* as when a network sees persistent, partial connectivity to part of Internet core. We present the *Taitao* algorithm to detect peninsulas that often result from peering disputes or long-term firewalls. We define *islands* as when one or more computers are partitioned from the main Internet core as detected by *Chiloe*, our second algorithm.

We apply these algorithms to data from two operational systems (Table 1): Trinocular, with frequent measurements of 5M networks from six Vantage Points (VPs) [75], and RIPE Atlas, with frequent measurements of the DNS root [85] from 13k VPs [83]. By applying new algorithms to existing, publicly available, multi-year data we are able to provide longitudinal analysis with some results covering more than three years. These two systems demonstrate our approach works on active probes covering millions of networks (although from few observers) and also from more than 13k VPs (although probing only limited destinations), strongly suggesting the results generalize, since no practical system can cover the $O(n^2)$ cost of all destinations from all sources.

In addition varying VPs and destinations across the design space, we evaluate the accuracy of our systems with rigorous measurements (§4). We quantify the independence of the Trinocular sites (§4.3) with cross-validation. Our analysis shows that combinations of any three independent VPs provide a result that is statistically indistinguishable from the asymptote §5.1. We show our results are stable over more than three years with samples from Trinocular (§4.2) and continuous results from RIPE Atlas (§6.3). Finally, we validate both algorithms against a third measurement system, CAIDA Archipelago, where 171 VPs scan millions of networks, daily [13]. Although comparing very different systems is challenging, these results provide strong bounds on accuracy (§4.1), with very good recall (0.94) and reasonable precision (lower and upper bounds from 0.42 and to 0.82).

Our final contribution *uses these algorithms to address current operational questions*. We

show that partial reachability is a *pervasive problem* today, meriting attention. We prove that peninsulas occur *more* often than outages, as subject of wide attention [90, 29, 75, 91, 28, 99]. We bring technical light to policy choices around national networks (§6.2) and de-peering (§6.1). We improve sensitivity of RIPE Atlas’ DNSmon [1] (§6.3), resolve corner cases in outage detection (§6.2), and quantify opportunities for route detouring (§5.1).

These contributions range from a theoretical definition, to experimental measurements, and their practical application. Each depends on the other—the definition enables the algorithms, which are then applied to show utility.

Artifacts and ethics: Data used (Table 1) and created [7] in this paper is available at no cost. Our work poses no ethical concerns (§A) by not identifying individuals and avoiding additional traffic by reanalysis with new algorithms. IRB review says it is non-human subjects research (USC IRB IIR00001648).

2 Problem: Partial Reachability

Understanding partial reachability requires a rigorous definition of *what* is being reached. We next define *the Internet core* to which we connect, to answer the political, architectural, and operational questions from §1.

We suggest a definition must be both *conceptual* and *operational* [35]. Our conceptual definition (§2.1) articulates what the Internet *is* and *is not*. it provides a goal which our implementation (§3) approximates, and we apply it improve real-world, operational systems (§6.3). Prior definitions [18, 73, 41] 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 properties the core must have (sufficient conditions, like supporting TCP). We add *necessary* conditions to define when networks *leave* the Internet core (§6.1).

2.1 The Internet: A Conceptual Definition

We define the Internet core as *all active IP addresses that can Bidirectionally Route to more than 50% of the public, Potentially Reachable Internet*. We define these key terms next, and expand their motivation and implications later (§2.2).

Two addresses are *Bidirectionally Routable* when each can initiate a connection to the other. In our realization we measure connectivity with either ICMP echo-request or with DNS queries and replies, considering alternatives in §2.2.

The *Potentially Reachable Internet* is all IP addresses in a graph-theoretic strongly-connected component, with graph edges defined by Bidirectional Routability. This definition means any node in the set can reach any other, either directly or perhaps through one or more hops.

2.2 Motivation for *This* Definition

We define the potentially reachable Internet via observation, so it depends only on testable, shared information, and not a central authority such as ICANN. Defining the Potentially Reachable Internet as active addresses also implies that the vast parts of unallocated IPv6 do not change our conclusions.

Why both bidirectional routability and potential reachability? *Bidirectional Routability* is connectivity in the networking sense, so each address must have a routing table entry that covers the other, and there must be some BGP-level reachability between them.

Graph-Theoretic Reachability shows transitive connectivity, even when disputes mean some pairs cannot reach each other.

Bidirectional Routability is required to capture the idea of IP routing from prior definitions [18, 73, 41], where all hosts should be able to communicate directly. It excludes private, NAT'ed addresses [78], which, although useful clients, require rendezvous protocols (STUN [86], UPnP [66], or PMP [19]) to partially link to the core, and also non-public cloud addresses hidden behind load balancers [47]. However, cloud VMs with fully-reachable public addresses are part of the core, including cloud-hosted services using public IP addresses from the cloud operator or their own (BYOIP).

Graph-Theoretic Reachability is required to define what “100%” is, so we guarantee one (or no) Internets by looking for a non-overlapping majority, even in the face of conflicting claims (§B). The combination of terms help us resolve such conflicts as different peninsulas sharing a common Internet core (although perhaps requiring relay through a third party).

Why more than 50%? We take as an axiom that there should be *one Internet core* per address space (IPv4 and IPv6), or a reason why that Internet core no longer exists. Thus we require a definition to unambiguously identify “the” Internet core given conflicting claims; any larger value is excessive, and anything smaller would allow multiple viable claims. (In practice, Figure 8 we see 98.5–99.5% agreement on the core, so values at the 50% threshold are unlikely.)

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 §B.

The definition of the Internet core should not require a central authority. “Majority” supports assessment independent of any authority. Any computer to prove it is in the Internet core by reaching half of active addresses, as defined by multiple, independent, long-term evaluations [51, 103, 27]. It also avoids identification of “tier-1” ISPs, an imprecise term determined only by private business agreements.

Finally, 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 or group chose to secede, or are expelled, the Internet core could become several no-longer internets (§6.1).


Why all addresses? In each of IPv4 and IPv6 we consider all addresses equally. Public Internet addresses are global, and the Internet core was intentionally designed without a hierarchy [20]. Consistent with decentralization trends [32], a definition should not create hierarchy, nor designate special addresses by age or importance.

These definitions are relatively apolitical and reduce first-mover bias, discussed in §6.1. 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 may be measured by protocols such as ICMP echo-request (pings), DNS or HTTP queries, or by data-plane reachability with BGP. Any specific test will provide an operational realization of our conceptual definition. (Measurement must tolerate transient failures, perhaps with multiple targets (Trinocular) or retransmissions (Atlas).) §5.1 examines how well using ICMP-based measures converge, and §6.3 shows DNS stability over years.

Firewalls complicate observing reachability and can make it conditional. We accept that



■ **Figure 1** *A, B and C are the connected core, B and C peninsulas, D and E islands, X is out.*

the results of specific observations may vary with different protocols or observation times; experiments show results are stable (§5.1). Measurement allows us to evaluate policy-driven unreachability (see Appendix G.2 in [9]).

We have two implementations of peninsula and island detection; both use publicly-available data from existing measurement systems. One uses Trinocular [75], because of its frequent, Internet-wide ICMP echo requests (11-minutes to 5M IPv4 /24s). Prior work has shown ICMP provides the most response [10, 75, 36], and can avoid rate limiting [50], other other protocol options are possible. Our second uses RIPE Atlas because of its use in DNS (§6.3).

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 §6.3. Many large outages are failures of applications such as DNS [74]; their study would require a different evaluator than IP reachability. 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 applications wax and wane, and some (like e-mail) extend beyond the Internet.

2.3 Cases of Partial Reachability

We use our definition of the Internet core to consider three types of partial reachability, shown in Figure 1. Here long-term and current routability are 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.

2.3.1 Outages

A number of groups have examined Internet outages [90, 75, 81, 49]. 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 since none of our 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.

2.3.2 Islands: Isolated Networks

An *island* is a group of public IP addresses partitioned from the Internet core, but still able to communicate among themselves. Operationally, outages (*X* in Figure 1) and islands (like

D and E) are both unreachable from external VPs and appear identical, but computers in an island are on and can reach each other.

Islands occur when an organization is no longer connected to the Internet core. A business with one ISP becomes an island when its router upstream connection fails, even though computers in the business can reach each other. An *address island* is when a computer can reach only itself.

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 §6.3 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 [75]. 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. We provide one example in Appendix E.1 of [9].

2.3.3 Peninsulas: Partial Connectivity

Link and power failures create islands, *peninsulas* are *partial* connectivity, when a group of public IP addresses can reach some destinations, but not others. (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 it cannot deliver or forward, when Tier-1 ISPs refuse to peer, or when firewalls block traffic. Experimental overlay networks route around peninsulas [2, 57, 58].

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 [67], and this problem persists through 2025, as we show in DNSmon (§6.3). We further confirm unreachability between HE and Cogent users in IPv6 with traceroutes from looking glasses [38, 24] (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 include Cogent and Google [76], and Cloudflare and Hurricane Electric [42]. Disputes can arise from 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 [98]). 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. We provide additional details in Appendix E.2 of [9].

3 Detecting Partial Connectivity

We now introduce the *Taitao* algorithm to detects peninsulas, and *Chiloe*, islands (names from Patagonian geography).

3.1 Taitao: a Peninsula Detector

Peninsulas occur when portions of the Internet core are reachable from some locations and not others. They can be seen by two VPs disagreeing on reachability.

Detecting peninsulas presents three challenges. Without VPs everywhere, when all VPs are on the same “side” of a peninsula (A and C in Figure 1), their reachability agrees even though VPs may disagree (like B). Second, asynchronous observations test reachability at different times: observations in Trinocular spread over 11 minutes, and in Atlas, 5 minutes. Observations at times before and after a network change will disagree, but both are true—a difference due to weak synchronization, and not a peninsula. Third, connectivity problems near the observer (or when an observer is an island) should not reflect on the intended destination.

We identify peninsulas by detecting disagreements in block state by comparing successful VP observations that occur at about the same time. Since probing rounds occur asynchronously, we compare measurements within the measurement system’s window (11 or 5 minutes for Trinocular and Atlas). This approach sees peninsulas lasting longer than one window duration, but may miss briefer ones, or when VPs are not on “both sides”.

Formally, $O_{i,b}$ is the set of observers with valid observations about block b at round i . We look for disagreements in $O_{i,b}$, defining $O_{i,b}^{up} \subset O_{i,b}$ as the set of observers that measure block b as up at round i . We detect a peninsula when:

$$0 < |O_{i,b}^{up}| < |O_{i,b}| \quad (1)$$

When only one VP reaches a block, we must classify it as a peninsula or an island, as described next.

3.2 Chiloe: an Island Detector

According §2.3.2, islands occur when the Internet core is partitioned, and the component with fewer than half the active addresses is the island. Typical islands are much smaller.

We can find islands by looking for networks that are only reachable from less than half of the Internet core. However, to classify such networks as an island and not merely a peninsula, we need to show that it is partitioned, which requires global knowledge. In addition, if islands are partitioned from all VPs, we cannot tell an island, with active but disconnected computers, from an outage, where they are off.

For these reasons, we must look for islands that include VPs in their partition. Because we know the VP is active and scanning we can determine how much of the Internet core is in its partition, ruling out an outage. We also can confirm the Internet core is not reachable, to rule out a peninsula.

Formally, we say that B is the set of blocks in the Internet core. $B_{i,o}^{up} \subseteq B$ are blocks reachable from observer o at round i , while $B_{i,o}^{dn} \subseteq B$ is its complement. We detect that observer o is in an island when it thinks half or more of the observable Internet core is down:

$$0 \leq |B_{i,o}^{up}| < |B_{i,o}^{dn}| \quad (2)$$

This method is independent of measurement systems, but is limited to detecting islands that contain VPs, so *any deployment will certainly undercount islands*. We evaluate islands in Trinocular and Atlas (§5.5), confirming more VPs see more islands, but that *nearly all reported islands are correct*.

Finally, because observations are not instantaneous, we must avoid confusing short-lived islands with long-lived peninsulas. For islands lasting longer than 11-minutes, we also require $|B_{i,o}^{up}| \rightarrow 0$. With $|B_{i,o}^{up}| = 0$, it is an address island.

Trinocular	Sites Up	Ark		
		Conflicting	All Down	All Up
Conflicting	1	20	6	15
	2	13	5	11
	3	13	1	5
	4	26	4	19
	5	83	13	201
Agree	0	6	97	6
	6	491,120	90	1,485,394

Table 2 Trinocular and Ark agreement table. Dataset A30, 2017q4.

Taitao		Ark	
		Peninsula	Non Peninsula
Peninsula		184	251 (<i>strict</i>) 40 (<i>loose</i>)
Non Peninsula		12	1,976,701

Table 3 Taitao confusion matrix. Dataset: A30, 2017q4.

3.3 Deployment with Existing Systems

We have deployed our algorithms as extensions to two systems: Trinocular and RIPE Atlas. In both cases, each system provides data to us via existing APIs and we then apply Taitao and Chiloe and share results back. Processing time for both is modest, with DNSmon running in minutes and Trinocular taking less time than Trinocular outage detection.

For DNSmon, we provide daily outages and peninsulas since 2022-01-01 on a public website [88]. We have also discussed these results with RIPE and the root operators; RIPE currently identifies islands manually, and one root operator is using our results to guide operations. We provide 3.5 years Trinocular analysis at our website [6], and are working with Trinocular operators to operationalize our algorithms.

4 Validating our approach

We next validate our algorithms with three data sources.

4.1 Can Taitao Detect Peninsulas?

We compare Taitao detections from 6 VPs to independent observations taken from more than 100 VPs in CAIDA’s Ark [14]. This comparison is challenging, because both Taitao and Ark are imperfect operational systems that differ in probing frequency, targets, and method. Neither defines perfect ground truth, but agreement suggests likely truth.

We believe this complexity is warranted because Ark provides a more diverse perspective (with 171 locations), if we can account for its much sparser frequency. Ark traceroutes also allow us to assess *where* peninsulas begin. We expect to see a strong correlation between Taitao peninsulas and Ark observations. (We considered RIPE Atlas as another external dataset, but its coverage is sparse, while Ark covers all /24s.)

Identifying comparable blocks: We study 21 days of Ark observations from 2017-10-10 to -31. Ark covers all networks with two strategies. With team probing in 2017, a 40 VP “team” traceroutes to all routed /24 about once per day. For prefix probing, about 35 VPs each traceroute to .1 addresses of all routed /24s every day. We use both types of data: the three Ark teams and all available prefix probing VPs. We group results by /24 blocks, considering /24s instead of ASes to be sensitive to intra-AS peninsulas.

Ark differs from Taitao’s Trinocular input in three ways: the target is a random address or the .1 address in each block; it uses traceroute, not ping; and it probes blocks daily, not every 11 minutes. Sometimes these differences cause Ark traceroutes to fail when a simple ping succeeds. First, Trinocular’s targets respond more often because it uses a curated hitlist [40] while Ark does not. Second, Ark’s traceroutes can terminate due to path *loops* or *gaps* in the path, (in addition to succeeding or reporting target unreachable). We do

not consider results with gaps, so problems on the path do not bias results for endpoints reachable by direct pings.

To correct for differences in target addresses, we must avoid misinterpreting a block as unreachable when the block is online but Ark’s target address is not, we discard traces sent to never-active addresses (those not observed in 3 years of complete IPv4 scans), and blocks for which Ark did not get a single successful response. Since dynamic addressing [72] means Ark often fails with an unreachable last hop, we see conflicting observations in Ark, implying false peninsulas. We therefore trust Ark confirmation of outages and full reachability, but question Ark-only peninsulas.

To correct for Ark’s less frequent probing, we compare *long-lived* Trinocular down-events (5 hours or more). Ark measurements are infrequent (once every 24 hours) compared to Trinocular’s 11-minute reports, so short Trinocular events are often unobserved by Ark. (Since outage durations are heavy-tailed, 5 h gives Ark some time to confirm without discarding too many events.) To confirm agreements or conflicting reports from Ark, we require at least 3 Ark observations within the peninsula’s span of time. Varying these parameters is potential future work; with small quantitative changes likely, but changes to overall bounds unlikely.

We filter out blocks with frequent transient changes or signs of network-level filtering, as prior work [75, 91, 81]. We define the “reliable” blocks suitable for comparison as those responsive for at least 85% of the quarter from each of the 6 Trinocular VPs. (This threshold avoids diurnal blocks or blocks with long outages; values of 90% or less have similar results.) We also discard flaky blocks whose responses are frequently inconsistent across VPs. (We consider more than 10 combinations of VP as frequently inconsistent.) For the 21 days, we find 4M unique Trinocular /24 blocks, and 11M Ark /24 blocks, making 2M blocks in both available for study.

Results: Table 3 shows outcomes, treating Taitao as prediction and Ark as truth, with details in Table 2. Dark green indicates true positives (TP): when (a) either both Taitao and Ark show mixed results, both indicating a peninsula, or when (b) Taitao indicates a peninsula (1 to 5 sites up but at least one down), Ark shows all-down during the event and up before and after. We treat Ark in case (b) as positive because the infrequency of Ark probing (one probe per team every 24 hours) means we cannot guarantee VPs in the peninsula will probe responsive targets in time. Since peninsulas are not common, so too are true positives, but we see 184 TPs.

We show *true negatives* as light green and neither bold nor italic. In almost all of these cases (1.4M) both Taitao and Ark reach the block, agreeing. The vast majority of these are an artifact of our use of Ark as “ground truth”, when it is not designed to accurately measure partitions. The challenge of an Ark claim of peninsula is that about 5/6ths of Ark probes fail in the last hop because it probes a single random address (see [75] figure 6). As a result, while positive Ark results support non-partitions, negative Ark results are most likely a missed target and not an unreachable block; we expand on this analysis in Appendix F.1 of [9]. We therefore treat this second most-common result (491k cases) as a true negative. For the same reason, we include the small number (97) of cases where both Ark and Taitao report all-down, assuming Ark terminates at an empty address. We include in this category the 90 events where Ark is all-down and Trinocular is all-up. We attribute Ark’s failure to reach its targets to infrequent probing.

We mark *false negatives* as red and bold. For these few cases (only 12), all Trinocular VPs are down, but Ark reports all or some responding. We believe these cases indicate blocks that have chosen to drop Trinocular traffic.

Finally, yellow italics shows when Taitao’s peninsulas are *false positives*, since all Ark

Trinocular		Chiloe	
		Island	Peninsula
	Block Island	2	0
	Addr Island	19	8
	Peninsula	2	566

(a) Chiloe confusion matrix

Sites	Events	Per Year
W	5	1.67
C	11	3.67
J	1	0.33
G	1	0.33
E	3	1.00
N	2	0.67
All (norm.)	23	7.67 (1.28)

(b) Detected islands

■ **Table 4** (a) Chiloe confusion matrix, events between 2017-01-04 and 2020-03-31, datasets A28 through A39. (b) Islands detected from 2017q2 to 2020q1.

probes reached the target block. This case occurs when either traffic from some Trinocular VPs is filtered, or all Ark VPs are “inside” the peninsula. Light yellow (strict) shows all the 251 cases that Taitao detects. For most of these cases (201), five Trinocular VPs responding and one does not, suggesting network problems are near one of the Trinocular VPs (since five of six independent VPs have working paths). Discarding these cases we get 40 (orange); still conservative but a *looser* estimate.

The strict scenario sees precision 0.42, recall 0.94, and F_1 score 0.58, and in the loose scenario, precision improves to 0.82 and F_1 score to 0.88. We consider these results a strong lower bound on the size of problem, and confirmation that the peninsulas detected by Taitao are correct.

Of course custom measurement could align with our analysis and should close this bound, but the need to build in long-term, existing data, motivates these early, rough bounds. We expect future work to tighten these bounds.

4.2 Can Chiloe Detect Islands?

Chiloe (§3.2) detects islands when a VP within the island can reach less than half the rest of the world.

Trinocular: To validate Chiloe’s correctness, we compare when a single VP believes to be in an island, against what the rest of the world believes about that VP. We begin with Trinocular, where we have strong evidence for a few VPs, then we summarize Atlas with 13k VPs.

Islands are unreachable, like D in Figure 1. We measure blocks, so if any address in block D can reach another, it is an island. If no external VPs can reach D ’s block, Chiloe confirms an island, but some VP reaching D ’s block implies a peninsula. In §4.3 we show that Trinocular VPs are independent, and therefore no two VPs live within the same island. We believe this definition is the best possible ground truth, since perfect classification requires instant, global knowledge and cannot be measured in practice.

We take 3 years worth of data from all six Trinocular VPs. From Trinocular’s pacing, we analyze 11-minute bins.

In Table 4a we show that Chiloe detects 23 islands across three years. In 2 of these events, the block is unreachable from other VPs, confirming the island with our validation methodology. Manual inspection confirms that the remaining 19 events are islands too, but at the address level—the VP was unable to reach anything but did not lose power, and other addresses in its block were reachable from VPs at other locations. These observations suggest a VP-specific problem making it an island. Finally, for 2 events, the prober’s block was reachable during the event by every site including the prober itself which suggests partial connectivity (a peninsula), and therefore a false positive.

	C	J	G	E	N
W	0.017	0.031	0.019	0.035	0.020
C		0.077	0.143	0.067	0.049
J			0.044	0.036	0.046
G				0.050	0.100
E					0.058

■ **Table 5** Similarities all VPs.
Dataset: A30, 2017q4.

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

■ **Table 6** RIR IPv4 hosts and IPv6 /32 allocation [53, 54].

In the 566 non-island events (true negatives), a single VP cannot reach more than 5% but less than 50% of the Internet core. In each of these cases, one or more other VPs were able to reach the affected VP’s block, showing they were not an island (although perhaps a peninsula). The table omits the frequent events when less than 5% of the network is unavailable from the VP, although they too are true negatives.

Bold red shows 8 false negatives. These are events that last about 2 Trinocular rounds or less (22 min), often not enough time for Trinocular to change its belief on block state.

Atlas: With 13k VPs, RIPE Atlas provides a broader view of islands. We find 188 (v4) and 388 (v6) Atlas VPs are islands (§6.3), accounting for *the majority of DNS unreachable events*. RIPE operators confirmed these are often misconfigurations.

Operators: Beyond this quantitative comparison, we discussed islands with Trinocular and RIPE Atlas operators. They confirm our examples and trends (Figure 7).

4.3 Are the Sites Independent?

Our evaluation assumes VPs do not share common network paths. VPs improve path diversity by network diversity and physical distance, particularly with today’s “flatter” Internet [59]. We next quantify and validate this assumption.

We measure similarity of observations between pairs of VPs. We examine only cases where one of the pair disagrees with some other VP, since when all agree, we have no new information. If the pair agrees with each other, but not with the majority, the pair shows similarity. If they disagree with each other, they are dissimilar. We quantify similarity S_P for a pair of sites P as $S_P = (P_1 + P_0)/(P_1 + P_0 + D_*)$, where P_s indicates the pair agrees on the network having state s of up (1) or down (0) and disagrees with the others, and for D_* , the pair disagrees with each other. S_P ranges from 1, where the pair always agrees, to 0, where they always disagree.

Table 5 shows similarities for each pair of the 6 Trinocular VPs (as half of the symmetric matrix). No two sites have a similarity more than 0.14, and most pairs are under 0.08. This result shows that no two sites are particularly correlated.

4.4 Stability Across Time

We confirm our results are not time-dependent by repeating key results in multiple years, including operational result from 2022 to 2025 (Figure 7 in §6.3), and confirm all results with multiple sources and dates (see Appendix F.2 of [9]). We expect these results to apply today since partial reachability has persisted since 2001 [2], with some events lasting years [42], as our results document (Figure 7). We use older data in some examples to avoid limitations

of measurement deployments. During 2017q4, Trinocular had six active VPs and Ark had three teams, providing strong statements from many perspectives. Trinocular had fewer VPs in 2019 and early 2020, and Ark has fewer teams today, but 2020 gives quantitatively similar results (see Appendix F.2 of [9]). §5.4 uses 2020q3 data because Ark observed a very large number of loops in 2017q4.

4.5 Varying Parameters and Geography

Our algorithms are influenced by the parameters in our data sources, including how often and where they probe, where they are placed, and how many VPs they employ, and how much data we analyze. We vary *all of these parameters* across our datasets (see Table 1), but the requirement for Internet-wide data spanning months and years means we depend on existing deployed infrastructure. Systematically varying VP frequency and location is challenging future work.

We believe these diverse data sources *confirm our results apply over a range of geographic locations*. We study locations quantitatively in §4.3) and confirm stable results with Atlas across 3k ASes and 12k locations in §6.3. Thus, while we certainly greatly *undercount* the absolute numbers of peninsulas and islands observed from Trinocular’s 6 locations (§5), Atlas confirms these trends apply with 12k VPs.

IPv6: Given data, our algorithms apply to both IPv4 and IPv6. We provide results for both v4 and v6 with RIPE Atlas and DNSmon (§6.3), and for Internet-wide v4 with Trinocular. Internet-wide IPv6 results depend on v6 outage detection, an area of active and future research.

5 Internet Islands and Peninsulas

We now examine islands and peninsulas in the Internet core.


5.1 How Common Are Peninsulas?

We estimate how often peninsulas occur in the Internet core in three ways. First, we directly measure the visibility of peninsulas by summing the duration of peninsulas as seen from six VPs. Second, we confirm the accuracy of this estimate by evaluating its convergence as we vary the number of VPs—more VPs show more peninsula-time, but a result that converges suggests it is approaching the limit. Third, we compare peninsula-time to outage-time, showing that, in the limit, observers see both for about the same duration. Outages correspond to service downtime [101], and are a recognized problem in academia and industry. Our results show that *peninsulas are as common as outages*, suggesting peninsulas are an important new problem deserving attention.

Peninsula-time: We estimate the duration an observer can see a peninsula by considering three types of events: *all up*, *all down*, and *disagreement* between six VPs. Disagreement, the last case, suggests a peninsula, while agreement (all up or down), suggests no problem or an outage. We compute peninsula-time by summing the time each target /24 has disagreeing observations from Trinocular VPs.

We have computed peninsula-time by evaluating Taitao over Trinocular data for 2017q4 [97]. Figure 2 shows the distribution of peninsulas measured as a fraction of block-time for an increasing number of sites. We consider all possible combinations of the six sites.

First we examine the data with all 6 VPs (the rightmost points). We see that peninsulas (the middle, disagreement graph) are visible about 0.00075 of the time. This data suggests



■ **Figure 2** Distribution of block-time fraction: all-down (left), disagreement (center), and all-up (right), events ≥ 1 hour. Data: 3.7M blocks, 2017-10-06 to -11-16, A30.

peninsulas regularly occur, appearing at least 0.075% of the time. Fortunately, large peninsulas are rare from many locations—our 6 VPs almost always see the same targets.

Convergence: While more VPs provide a better view of the Internet core’s overall state, but the *global fraction* of affected networks will show diminishing returns after major problems are found. That is previously inferred outages (all unreachable) should have been peninsulas, with partial reachability. All-down (left) decreases from an average of 0.00082 with 2 VPs to 0.00074 for 6 VPs. All-up (right) goes down a relative 47% from 0.9988 to 0.9984, while disagreements (center) increase from 0.0029 to 0.00045. Outages (left) converge after 3 sites, as shown by the fitted curve and decreasing variance. Peninsulas and all-up converge more slowly. We conclude that *a few, independent sites (3 or 4) converge on a good estimate of the fraction of true islands and peninsulas.*

We support this claim by comparing all non-overlapping combinations of 3 sites. If all combinations are equivalent, then a fourth site will not add new information. Six VPs yield 10 possible sets of 3 sites; we examine those combinations for each of 21 quarters, from 2017q2 to 2020q1. When we compare the one-sample Student *t*-test to evaluate if the difference of each pair of combinations of those 21 quarters is greater than zero, none of the combinations are rejected at confidence level 99.75%, suggesting that any combination of three sites is statistically equivalent and confirm our claim that a few sites are sufficient for estimation.

Relative impact: Finally, comparing outages (the left graph) with peninsulas (the middle graph), we see both occur about the same fraction of time (around 0.00075). This comparison shows that *peninsulas are about as common as outages*, suggesting they deserve more attention.

Generalizing: We confirm that each of these results holds in a subsequent year in Appendix F.2 of [9], suggesting the result is not unique to this quarter. While we reach a slightly different limit (in that case, peninsulas and outages appear about in 0.002 of data), we still see good convergence after 4 VPs.

While this data demonstrates convergence on the *rate* of peninsulas and islands, we confirm the rate and show a larger absolute *number* of peninsulas with DNSmon’s 12k VPs.

5.2 How Long Do Peninsulas Last?

Peninsulas have multiple root causes: some are short-lived routing misconfigurations while others reflect long-term disagreements in routing policy. In this section we determine the distribution of peninsulas in terms of their duration to determine the prevalence of persistent peninsulas. We will show that there are millions of brief peninsulas, likely due to transient routing problems, but that 90% of peninsula-time is in long-lived events (5 h or more, following §4.1).

We use Taitao to see peninsula duration for all detected in 2017q4: some 23.6M peninsulas

affecting 3.8M unique blocks. If instead we look at *long-lived* peninsulas (at least 5 h), we see 4.5M peninsulas in 338k unique blocks.

Figure 4 examines peninsula duration in three ways: a cumulative distribution (CDF) counting all peninsula events (left, solid, purple line), the CDF of the number of peninsulas for VP-down events longer than 5 hours (middle, solid green line), and the cumulative size of peninsulas for VP down events longer than 5 hours (right, green dashes).

We see that there are many very brief peninsulas (purple line): about 65% last only 20–60 minutes (~ 2 –6 observations). With two or more observations, these events are not just one-off measurement loss. These results suggest that while the Internet core is robust, there are many small connectivity glitches (7.8M events). Events that are two rounds (20 minutes) or shorter may be due to transient BGP blackholes [12].

The number of day-long or multi-day peninsulas is small, only 1.7M events (2%, the purple line). However, about 57% of all peninsula-time is in such longer-lived events (the right, dashed line), and 20% of time is in events lasting 10 days or more, even when longer than 5 hours events are less numerous (compare the middle, green line to the left, purple line). Day-long events persist long enough for human network operators to respond, and events lasting longer than a week suggest potential policy disputes and *intentional* unreachability. Together, these long-lived events suggest that there is benefit to identifying non-transient peninsulas and addressing the underlying routing problem.

5.3 What Sizes Are Peninsulas?

When network issues cause connectivity problems like peninsulas, the *size* of those problems may vary, from country-size (see Appendix G.2 in [9]), to AS-size, and also for routable prefixes or fractions of prefixes. We next examine peninsula sizes.

We begin with Taitao peninsula detection at a /24 block level. We match peninsulas across blocks within the same prefix by start time and duration, both measured in one hour timebins. This match implies that the Trinocular VPs observing the blocks as up are also the same.


We compare peninsulas to routable prefixes from Routeviews [65], using longest prefix matches with /24 blocks.

Routable prefixes consist of many blocks, some of which may not be measurable. We therefore define the *peninsula-prefix fraction* for each routed prefix as fraction of blocks in the peninsula that are Trinocular-measurable blocks. To reduce noise provided by single block peninsulas, we only consider peninsulas covering 2 or more blocks in a prefix.

Figure 3a shows the number of peninsulas for different prefix lengths and the fraction of the prefix affected by the peninsula as a heat-map, where we group them into bins.

We see that about 10% of peninsulas are likely due to routing problems or policies, since 40k peninsulas affect the whole routable prefix. However, a third of peninsulas (101k, at the bottom of the plot) affect only a very small fraction of the prefix. These low prefix-fraction peninsulas suggest that they happen *inside* an ISP and are not due to interdomain routing.


Finally, we show that *long-lived peninsulas are likely due to routing or policy choices*. Figure 3b shows the same data source, but weighted by fraction of time each peninsula contributes to the total peninsula time during 2017q4. Here the larger fraction of weight are peninsulas covering full routable prefixes—20% of all peninsula time during the quarter (see left margin).



■ **Figure 3** Peninsulas measured with per-site down events longer than 5 hours. Dataset A30, 2017q4.

Sites Up	Target AS		Target Prefix	
	At	Before	At	Before
0	21,765	32,489	1,775	52,479
1	587	1,197	113	1,671
2	2,981	4,199	316	6,864
3	12,709	11,802	2,454	22,057
4	117,377	62,881	31,211	149,047
5	101,516	53,649	27,298	127,867
1-5	235,170	133,728	61,392	307,506
6	967,888	812,430	238,182	1,542,136

■ **Table 7** Halt location of failed traceroutes for peninsulas longer than 5 hours. Dataset A41, 2020q3.



■ **Figure 4** Cumulative peninsulas and peninsula duration. Dataset A30, 2017q4.


5.4 Where Do Peninsulas Occur?

Firewalls, link failures, and routing problems cause peninsulas on the Internet, and can occur at AS boundaries or inside an AS. We next show that *many peninsulas occur at AS boundaries, consistent with policies as a cause* for long-lived events. (Short-lived events at AS boundaries may be routing transients or operator error that is quickly corrected.)

To detect *where* the Internet breaks into peninsulas, we look at traceroutes that failed to reach their target address, either due to a loop or an ICMP unreachable message. Then, we examine if the traceroute halts *at* the target AS and target prefix, or *before* the target AS and prefix.

For our experiment we run Taitao to detect peninsulas at target blocks over Trinocular VPs, we use Ark’s traceroutes [15] to find last IP address before halt, and we get target and halting ASNs and prefixes using RouteViews.

In Table 7 we show how many traces halt *at* or *before* the target network. The center, gray rows show peninsulas (disagreement between VPs) with their total sum in bold. For all peninsulas (the bold row), more traceroutes halt at or inside the target AS (235k vs. 134k, the left columns), but they more often terminate before reaching the target prefix (308k vs. 61k, the right columns). (While traceroutes are imperfect, these large differences (2× or more) suggest a robust qualitative conclusion.) This difference suggests policy is implemented at or inside ASes, but not at routable prefixes. By contrast, outages (agreement with 0 sites up) more often terminate before reaching the target AS. Because peninsulas are more often at or in an AS, while outages occur in many places, it suggests that long-lived peninsulas are policy choices consistent with public operator reports [67, 62, 3, 77, 94, 17].



■ **Figure 6** CDF of islands detected by Chiloe for data from Trinocular (3 years, Datasets A28-A39) and Atlas (2021q3).

5.5 How Common Are Islands?

Multiple groups have shown that there are many network outages in the Internet [90, 75, 91, 81, 49]. We have described (§2) two kinds of outages: full outages where all computers at a site are down (perhaps due to a loss of power), and islands, where the site is cut off from the Internet core, but computers at the site can talk between themselves. We next use Chiloe to determine how often islands occur. We study islands in two systems with 6 VPs for 3 years and 13k VPs for 3 months.

Trinocular: We first consider three years of Trinocular data (Table 1), from 2017-04-01 to 2020-04-01. We run Chiloe across each VP for this period.

Table 4b shows the number of islands per VP over this period. Over the 3 years, all six VPs see from 1 to 5 islands. In addition, we report as islands some cases even though not the *entire* Internet core is unreachable. This apparent discrepancy from our definition reflects the limitations of our necessarily non-instantaneous measurement of the Internet. We expect such cases, and perhaps other 12 non-islands where 20% to 50% is inaccessible, are *short-lived* true islands, that are incompletely measured because the island recovers before we complete an 11 minute-long evaluation of all 5M networks for a full Internet scan (see §C.2 for details).

RIPE Atlas: For broader coverage we next consider RIPE Atlas’ 13k VPs for all of 2021q3 [69]. While Atlas does not scan the whole Internet core, they do scan most root DNS servers every 240s. Chiloe would like to observe the whole Internet core, and while Trinocular scans 5M /24s, it does so with only 6 VPs. To use RIPE Atlas’ VPs, we approximate a full scan with probes to 12 of the DNS root server systems (G-Root was unavailable in 2021q3). Although far fewer than 5M networks, these targets provide a very sparse sample of usually independent destinations since each is independently operated. Thus we have complementary datasets with sparse VPs and dense probing, and many VPs but sparse probing. In other words, to get many VP locations we relax our conceptual definition by decreasing our target list.

Figure 5a shows the CDF of the number of islands detected per RIPE Atlas VP during 2021q3. During this period, 55% of VPs observed one or no islands (the solid line). We compare to Trinocular with only events longer than 660s (the dashed line). We see that 60% of VPs have no islands; 19%, one; with 21% seeing more. The annualized rate of the stable VPs that see 2 or fewer islands is 1.75 islands per year (a lower bound, since we exclude less stable VPs), compared to 1.28 for Trinocular (Table 4b). We see islands are more common in Atlas, perhaps because it includes many VPs in homes.

We conclude that islands *do* happen, but rarely, and occur at irregular times. This finding is consistent with importance of the Internet at the locations where we run VPs.

5.6 How Long Do Islands Last?

Islands causes range from brief connectivity loss to long-term policy differences, so we next evaluate island duration.

We compare the distributions of island durations observed from RIPE Atlas (the left line) and Trinocular (right) in [Figure 5b](#). Since Atlas’ frequent polling means it detects islands lasting seconds, while Trinocular sees only islands of 660 s or longer, we split out Atlas events lasting at least 660 s (middle line). All measurements follow a similar S-shaped curve, but for Trinocular, the curve is truncated at 660 s. With only 6 VPs, Trinocular sees far fewer events (23 in 3 years compared to 235k in a yearly quarter with Atlas), so the Trinocular data is quantized. In both cases, about 70% of islands are between 1000 and 6000 s. This graph shows that Trinocular’s curve is similar in shape to Atlas-660 s, but about 2× longer. All Trinocular observers are in datacenters, while Atlas devices are often at homes, so this difference may indicate that datacenter islands are rarer, but harder to resolve.

5.7 What Sizes Are Islands?

5.7.1 Island Size via Traceroute

First we evaluate island sizes, comparing traceroutes before and during an island. We use traceroutes from RIPE Atlas VPs sent to 12 root DNS servers for 2021q3 [70]. [Figure 5c](#) shows the distribution of number of traceroute hops reaching target (green), and *not* reaching their target (purple), for VPs in islands (§5.5).

Most islands are small, with 70% at 0 or 1 hop. We believe huge islands (10 or more hops) are likely false positives.

5.7.2 Country-sized Islands

We have some evidence of country-sized islands: In 2017q3, on 8 occasions it appears that most or all of China stopped responding to external pings (visualized in [Figure 10](#) in §C.1). We found no problem reports on network operator mailing lists, so we believe these outages were ICMP-specific and likely did not affect web traffic. Since there were no public reports, we assume the millions of computers inside China continued to operate, suggesting that China was briefly a country-wide *ICMP-island*. Such large examples have not re-occurred.

6 Applying These Tools

6.1 Can the Internet Core Partition?

In §6.2 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 6](#) reports the number of active IPv4 addresses as determined by Internet censuses [51] for RIRs and selected countries. Since estimating active IPv6 addresses is an open problem, we provide allocated addresses for both v4 and v6 [53, 54]. (IPv4 has been fully allocated since 2011 [55]).

[Table 6](#) 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 surpasses a 50% allocation for control. RIPE (an RIR) is close with 46.7%, and China and the U.S. have large allocations; with most v6 unallocated, this balance 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).

6.2 Other Applications of the Definition

We next examine how a clear definition of the Internet core can inform policy tussles [21]. 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. [84], China [4, 5], and Russia [22] have all proposed unplugging from the Internet. Egypt did in 2011 [25], and several countries have during exams [45, 30, 52, 37]. 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 (§2.1) resolves this question, since requiring a majority defines an Internet core that can end (§6.1) if multiple partitions leave none with a majority.

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 [80]. 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 §6.1, *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.

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 (loopback), and 240/4 (reserved) [43]. 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 [8]. Several IPv4 /8s allocated to the U.S. DoD have been used this way [82] (they were only publicly routed in 2021 [95]). 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.

The IPv4/v6 Transition: We have defined two Internet cores: IPv4 and IPv6. Our definition can determine when one supersedes the other. After more than half of all IPv4 hosts are dual-homed, 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, and show the strength and limits of our definition: since IPv6 is already economically




Figure 7 Fraction of VPs observing islands and peninsulas for IPv4 and IPv6, 2022–2025.




Figure 8 Atlas queries from all available VPs to 13 Root Servers for IPv4 and IPv6 on 2022-07-23.

important, a definition seems unnecessary. But providing a sharp threshold that makes the maturity of IPv6 definitive may help motivate late-movers.

Outage Detection: Prior outage detection systems have struggled with conflicting observations [90, 75, 91, 81, 49]. We instead recognize such cases as peninsulas in a normal Internet, not measurement error. (We expand in §6.4.)

6.3 Improving DNSmon Sensitivity

DNSmon [1] monitors the Root Server System [85], with each RIPE Atlas VP measuring its anycast-determined neighbor [83]. For years, DNSmon has often reported IPv6 loss rates of 4-10%. Since the DNS root is well provisioned and anycast, we expect minimal or no loss.


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 and §3), we give full analysis in a workshop paper [87]; Here we add new data showing these results persist for 3 years (Figure 7).

Groups of bars in Figure 8 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 apply Chiloe to these VPs, detecting as islands those VPs that cannot see *any* of the 13 root identifiers over 24 hours. (This definition is stricter than regular Chiloe because these VPs attempt only 13 targets, 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 188 v4 and 388 v6 VPs that are islands. Without islands, v4 loss drops to 0.005 from 0.01, and v6 to 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 [67]. Omitting islands (the middle bars) makes this difference much clearer.

Applying Taitao to detect peninsulas, we find 14 to 57 v4 peninsulas and 266 (Cogent) and 19 to 49 (others) v6 peninsulas. Peninsulas suggest persistent routing problems meriting attention from ISPs and root operators. The darker, right two groups show loss remaining (after removal of islands and peninsulas), representing *underlying events worth root operator*



■ **Figure 9** Ark traceroutes sent to targets under partial outages (2017-10-10 to -31). Dataset A30.

attention. These bars show all letters see similar events rates, *after* we remove persistent problems.

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 these problems can be interesting, but considering each separately, the interesting “signal” of routing changes (appearing in the right two groups in Figure 8), is hidden under the $5\times$ or $9.7\times$ 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.

6.4 Outages Given Partial Reachability

We next re-evaluate reports from existing outage detection systems, considering how to resolve conflicting information in light of our new algorithms. We compare findings to external information in traceroutes from CAIDA Ark.

Figure 9 compares Trinocular with 21 days of Ark topology data, from 2017-10-10 to -31 from all 3 probing teams. For each Trinocular outage we classify the Ark result as success or three types of failure: unreachable, loop, or gap.

Trinocular’s 6-site-up case suggests a working network, and we consider this case as typical. However, we see that about 25% of Ark traceroutes are “gap”, where several hops fail to reply. We also see about 2% of traceroutes are unreachable (after we discard traceroutes to never reachable addresses). Ark probes a random address in each block; many addresses are non-responsive, explaining these.

With 1 to 11 sites up, Trinocular is reporting disagreement. We see that the number of Ark success cases (the green, lower portion of each bar) falls roughly linearly with the number of successful observers. This consistency suggests that Trinocular and Ark are seeing similar behavior, and that there is partial reachability—these events with only partial Trinocular positive results are peninsulas.

Since 5 sites give the same results as all 6, single-VP failures likely represent problems local to that VP. This data suggests that all-but-one voting will track true outages.

With only partial reachability, with 1 to 4 VPs (of 6), we see likely peninsulas. These cases confirm that partial connectivity is common: while there are 1M traceroutes sent to outages where no VP can see the target (the number of events is shown on the 0 bar), there are 1.6M traceroutes sent to partial outages (bars 1 to 5), and 850k traceroutes sent to definite peninsulas (bars 1 to 4). This result is consistent with the convergence we see in Figure 2.

7 Related Work

Prior definitions of the Internet exist at the IP-layer [18, 73, 41, 39] of their time, or the AS-level [44, 63]. We consider the IP-layer, and seek to address today’s challenges (see §2).

Cannon explored legal definitions of the Internet [16], recognizing limitations of early definitions and need to be application-independent. Like us, he considers connectivity and addressing important, but he questions if a firm legal definition is possible. While we do not comment legalities, we suggest our technical definition may address his questions.

Several systems mitigate partial outages. RON provides alternate-path routing around failures for a mesh of sites [2]. Hubble monitors in real-time reachability problems when working physical paths exist [57]. LIFEGUARD, remediates route failures by rerouting traffic using BGP to select a working path [58]. While addressing the problem of partial outages, these systems do not quantify their duration or scope.

Prior work studied partial reachability, showing it is a common transient occurrence during routing convergence [12]. They reproduced partial connectivity with controlled experiments; we study it from Internet-wide VPs.

Internet scanners have examined bias by location [51], more recently looking for policy-based filtering [99]. We measure policies with our country specific algorithm, and we extend those ideas to defining the Internet core.

Active outage detection systems have encountered partial outages. Thunderping’s “hosed” state recognizes mixed replies, but its study is future work [90]. Trinocular discards partial outages by reporting the target block “up” if any VP can reach it [75]. Disco identifies partial connectivity as future work [91]. None of these systems consistently report partial outages in the Internet core, nor 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 [61, 60, 68], for example, with applications to certificate authorities [11].

Recent work considered policies about Internet fragmentation [33, 34], but do not define it—a need we hope to meet.

8 Conclusions

Our new definition of the Internet core leads to new algorithms: Taitao, to find peninsulas of partial connectivity, and Chiloe, to find islands. We validate these algorithms and show partial reachability is as common as simple outages. They have important applications about Internet sovereignty and to improve outage and DNSmon measurement systems.

References

- 1 Christopher Amin, Massimo Cándela, Daniel Karrenberg, Robert Kisteleki, and Andreas Strikos. 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, June 2015.
- 2 David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the Symposium on Operating Systems Principles*, pages 131–145, Chateau Lake Louise, Alberta, Canada, October 2001. ACM.
- 3 Nate Anderson. Peering problems: digging into the Comcast/Level 3 grudge match. *Ars Technica*, Dec. 9 2010. URL: <https://arstechnica.com/tech-policy/2010/12/comcastlevel3/>.

- 4 Anonymous. The collateral damage of Internet censorship by DNS injection. *ACM Computer Communication Review*, 42(3):21–27, July 2012. doi:10.1145/2317307.2317311.
- 5 Anonymous. 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*, San Diego, CA, USA, August 2014. USENIX.
- 6 ANT Project. Ant internet islands and peninsula datasets. https://ant.isi.edu/datasets/ipv4_partial/, January 2017. URL: https://ant.isi.edu/datasets/ipv4_partial/.
- 7 ANT Project. ANT IPv4 island and peninsula data. https://ant.isi.edu/datasets/ipv4_partial/, November 2022.
- 8 Cathy Aronson. To squat or not to squat? blog <https://teamarin.net/2015/11/23/to-squat-or-not-to-squat/>, November 2015.
- 9 Guillermo Baltra, Tarang Saluja, Yuri Pradkin, and John Heidemann. Understanding partial reachability in the internet core (extended). Technical Report arXiv:2601.12196, arXiv, January 2026. The technical report includes additional appendices.
- 10 Genevieve Bartlett, John Heidemann, and Christos Papadopoulos. Understanding passive and active service discovery. In *Proceedings of the ACM Internet Measurement Conference*, pages 57–70, San Diego, California, USA, October 2007. ACM. doi:10.1145/1298306.1298314.
- 11 Henry Birge-Lee, Yixin Sun, Anne Edmundson, Jennifer Rexford, and Prateek Mittal. Bamboozling certificate authorities with BGP. In *27th USENIX Security Symposium*, pages 833–849, Baltimore, Maryland, USA, 2018. USENIX.
- 12 Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Internet optometry: assessing the broken glasses in internet reachability. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*, pages 242–253, Chicago, Illinois, USA, November 2009. ACM. URL: <http://www.maennel.net/2009/imc099-bush.pdf>.
- 13 CAIDA. Archipelago (Ark) measurement infrastructure. website <https://www.caida.org/projects/ark/>, 2007.
- 14 CAIDA. The CAIDA UCSD IPv4 routed /24 topology dataset - 2017-10-10 to -31. https://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml, 2017.
- 15 CAIDA. The CAIDA UCSD IPv4 routed /24 topology dataset - 2020-09-01 to -31. https://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml, 2020.
- 16 Robert Cannon. Will the real Internet please stand up: An attorney’s quest to define the internet. In *Proceedings of the TPRC, the Research Conference on Communication, Information and Internet Policy*. TPRC, March 2002. Also in: *Rethinking Rights and Regulations*. doi:10.7551/mitpress/5932.003.0007.
- 17 Ben Cartwright-Cox. Cogent-tata peering dispute? Nanog mailing list, May 17 2024. URL: <https://mailman.nanog.org/pipermail/nanog/2024-May/225651.html>.
- 18 Vint Cerf and Robert Kahn. A protocol for packet network interconnection. *IEEE Transactions on Communications*, COM-22(5):637–648, May 1974.
- 19 S. Cheshire and M. Krochmal. NAT port mapping protocol (NAT-PMP). RFC 6886, Internet Request For Comments, April 2013. doi:10.17487/RFC6886.
- 20 David D. Clark. The design philosophy of the DARPA Internet protocols. In *Proceedings of the 1988 Symposium on Communications Architectures and Protocols*, pages 106–114. ACM, August 1988.
- 21 David D. Clark, John Wroclawski, Karen Sollins, and Robert Braden. Tussle in cyberspace: Defining tomorrow’s internet. In *Proceedings of the ACM SIGCOMM Conference*, pages 347–356, Pittsburgh, PA, USA, August 2002. ACM.
- 22 CNBC. 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>, 11 2019.
- 23 N. Coca. China’s xjiang surveillance is the dystopian future nobody wants. *Engadget*, Feb. 22 2018. URL: <https://www.engadget.com/2018-02-22-china-xjiang-surveillance-tech-spread.html>.

- 24 Cogent. Looking glass. <https://cogentco.com/en/looking-glass>, 05 2021.
- 25 James Cowie. Egypt leaves the Internet. Renesys Blog <http://www.renesys.com/blog/2011/01/egypt-leaves-the-internet.shtml>, January 2011.
- 26 RBC daily. Russia, tested the Runet when disconnected from the global network. website https://www.rbc.ru/technology_and_media/21/07/2021/60f8134c9a79476f5de1d739, July 2021.
- 27 Alberto Dainotti, Karyn Benson, Alistair King, kc claffy, Eduard Glatz, Xenofontas Dimitropoulos, Philipp Richter, Alessandro Finamore, and Alex C. Snoeren. Lost in space: Improving inference of IPv4 address space utilization. *IEEE Journal of Selected Areas in Communication*, 34(6):1862–1876, April 2016. doi:10.1109/JSAC.2016.2559218.
- 28 Alberto Dainotti, kc claffy, Alistair King, Vasco Asturiano, Karyn Benson, Marina Fomenkov, Brad Huffaker, Young Hyun, Ken Keys, Ryan Koga, Alex Ma, Chiara Orsini, and Josh Polterock. IODA: Internet outage detection & analysis. Talk at CAIDA Active Internet Measurement Workshop (AIMS), March 2017. URL: http://www.caida.org/publications/presentations/2017/ioda_aims/ioda_aims.pdf.
- 29 Alberto Dainotti, Claudio Squarcella, Emile Aben, Marco Chiesa, Kimberly C. Claffy, Michele Russo, and Antonio Pescapé. Analysis of country-wide Internet outages caused by censorship. In *Proceedings of the ACM Internet Measurement Conference*, pages 1–18, Berlin, Germany, November 2011. ACM. doi:10.1145/2068816.2068818.
- 30 Dhaka Tribune Desk. Internet services to be suspended across the country. *Dhaka Tribune*, Feb. 11 2018. URL: <http://www.dhakatribune.com/regulation/2018/02/11/internet-services-suspended-throughout-country/>.
- 31 Amogh Dhamdhare, David D. Clark, Alexander Gamero-Garrido, Matthew Luckie, Ricky K. P. Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex C. Snoeren, and kc claffy. Inferring persistent interdomain congestion. In *Proceedings of the ACM SIGCOMM Conference*, pages 1–15, Budapest, Hungary, August 2018. ACM. doi:10.1145/3230543.3230549.
- 32 DINRG. Decentralized Internet Infrastructure Research Group. <https://irtf.org/dinrg>, 05 2021.
- 33 William J. Drake, Vinton G. Cerf, and Wolfgang Kleinwächter. Internet fragmentation: An overview. Technical report, World Economic Forum, January 2016. URL: https://www3.weforum.org/docs/WEF_FII_Internet_Fragmentation_An_Overview_2016.pdf.
- 34 William J. Drake (moderator). Internet fragmentation, reconsidered. CITI Seminar on Global Digital Governance at IETF 115, October 2022. URL: <https://www8.gsb.columbia.edu/citi/GlobalDigitalGovernance>.
- 35 Peter K. Dunn. Scientific research methods. <https://bookdown.org/pkaldunn/Book/>, 05 2021.
- 36 Zakir Durumeric, Michael Bailey, and J Alex Halderman. An internet-wide view of internet-wide scanning. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, pages 65–78, San Diego, California, USA, August 2014. USENIX.
- 37 Economist Editors. Why some countries are turning off the internet on exam days. *The Economist*, July 5 2018. (Appeared in the Middle East and Africa print edition). URL: <https://www.economist.com/middle-east-and-africa/2018/07/05/why-some-countries-are-turning-off-the-internet-on-exam-days>.
- 38 Hurricane Electric. Looking glass. <http://lg.he.net/>, May 2021.
- 39 Engadget. China, Huawei propose internet protocol with a built-in killswitch. <https://www.engadget.com/2020-03-30-china-huawei-new-ip-proposal.html>, 2020.
- 40 Xun Fan and John Heidemann. Selecting representative IP addresses for Internet topology studies. In *Proceedings of the ACM Internet Measurement Conference*, pages 411–423, Melbourne, Australia, November 2010. ACM. doi:10.1145/1879141.1879195.
- 41 Federal Networking Council (FNC). Definition of “Internet”. https://www.nitrd.gov/historical/fnc/internet_res.pdf, 1995.
- 42 HE forums. Cloudflare blocked on free tunnels now? <https://forums.he.net/index.php?topic=3805.0>, 12 2017.

- 43 V. Fuller, E. Lear, and D. Meyer. Reclassifying 240/4 as usable unicast address space. Work in progress (Internet draft draft-fuller-240space-02.txt), March 2008. URL: <https://datatracker.ietf.org/doc/html/draft-fuller-240space-02>.
- 44 Lixin Gao. On inferring autonomous system relationships in the Internet. *ACM/IEEE Transactions on Networking*, 9(6):733–745, December 2001. doi:10.1109/90.974527.
- 45 Samuel Gibbs. Iraq shuts down the Internet to stop pupils cheating in exams. *The Guardian*, 18 May 1996. URL: <https://www.theguardian.com/technology/2016/may/18/iraq-shuts-down-internet-to-stop-pupils-cheating-in-exams>.
- 46 GovTrack.us. 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-78326> November 2020.
- 47 Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, and Parveen Pat. VL2: A scalable and flexible data center network. In *Proceedings of the ACM SIGCOMM Conference*, pages 51–62, Barcelona, Spain, August 2009. ACM.
- 48 James Griffiths. Democratic Republic of Congo internet shutdown shows how chinese censorship tactics are spreading. *CNN*, Jan. 2 2019. URL: <https://edition.cnn.com/2019/01/02/africa/congo-internet-shutdown-china-intl/index.html>.
- 49 Andreas Guillot, Romain Fontugne, Philipp Winter, Pascal Merindol, Alistair King, Alberto Dainotti, and Cristel Pelsser. Chocolate: Outage detection for internet background radiation. In *Proceedings of the IFIP International Workshop on Traffic Monitoring and Analysis*, Paris, France, June 2019. IFIP.
- 50 Hang Guo and John Heidemann. Detecting ICMP rate limiting in the Internet. In *Proceedings of the Passive and Active Measurement Workshop*, page to appear, Berlin, Germany, March 2018. Springer.
- 51 John Heidemann, Yuri Pradkin, Ramesh Govindan, Christos Papadopoulos, Genevieve Bartlett, and Joseph Bannister. Census and survey of the visible Internet. In *Proceedings of the ACM Internet Measurement Conference*, pages 169–182, Vouliagmeni, Greece, October 2008. ACM. doi:10.1145/1452520.1452542.
- 52 Jon Henley. Algeria blocks internet to prevent students cheating during exams. *The Guardian*, 22 June 2018. URL: <https://www.theguardian.com/world/2018/jun/21/algeria-shuts-internet-prevent-cheating-school-exams>.
- 53 IANA. IPv4 address space registry. <https://www.nro.net/about/rirs/statistics/>, 05 2021.
- 54 IANA. IPv6 RIR allocation data. <https://www.iana.org/numbers/allocations/>, 01 2021.
- 55 ICANN. Available pool of unallocated IPv4 internet addresses now completely emptied. Announcement, ICANN, February 2011. URL: <https://itp.cdn.icann.org/en/files/announcements/release-03feb11-en.pdf>.
- 56 Internet Architecture Board. IAB technical comment on the unique DNS root. RFC 2826, Internet Request For Comments, May 2000. URL: <https://www.rfc-editor.org/rfc/rfc2826>.
- 57 Ethan Katz-Bassett, Harsha V Madhyastha, John P John, Arvind Krishnamurthy, David Wetherall, and Thomas E Anderson. Studying black holes in the internet with hubble. In *Proceedings of the USENIX Conference on Networked Systems Design and Implementation*, pages 247–262, San Francisco, CA, 2008. ACM.
- 58 Ethan Katz-Bassett, Colin Scott, David R. Choffnes, Ítalo Cunha, Vytas Valancius, Nick Feamster, Harsha V. Madhyastha, Tom Anderson, and Arvind Krishnamurthy. LIFEGUARD: Practical repair of persistent route failures. In *Proceedings of the ACM SIGCOMM Conference*, pages 395–406, Helsinki, Finland, August 2012. ACM. doi:10.1145/2377677.2377756.

- 59 Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. Internet inter-domain traffic. In *Proceedings of the ACM SIGCOMM Conference*, pages 75–86, New Delhi, India, August 2010. ACM. doi:10.1145/1851182.1851194.
- 60 Leslie Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, May 1998. doi:http://dx.doi.org/10.1145/279227.279229.
- 61 Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.
- 62 Mike Leber. Re: Ipv6 internet broken, cogent/telia/hurricane not peering. NANOG mailing list, October 2009. URL: <https://mailman.nanog.org/pipermail/nanog/2009-October/014017.html>.
- 63 Matthew Luckie, Bradley Huffaker, Dhamdhere, Vasileios Giotsas, and kc claffy. AS relationships, customer cones, and validation. In *Proceedings of the ACM Internet Measurement Conference*, pages 243–256, Barcelona, Spain, October 2013. ACM.
- 64 Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP misconfiguration. In *Proceedings of the ACM SIGCOMM Conference*, pages 3–16, Pittsburgh, Pennsylvania, USA, August 2002. ACM. doi:10.1145/633025.633027.
- 65 D. Meyer. University of Oregon Routeviews. <http://www.routeviews.org>, 2018.
- 66 Brent A. Miller, Toby Nixon, Charlie Tai, and Mark D. Wood. Home networking with universal plug and play. *IEEE Communications Magazine*, 39(12):104–109, December 2001. doi:10.1109/35.968819.
- 67 Rich Miller. Peering disputes migrate to IPv6. website <https://www.datacenterknowledge.com/archives/2009/10/22/peering-disputes-migrate-to-ipv6>, Oct. 22 2009.
- 68 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Released publicly <http://bitcoin.org/bitcoin.pdf>, March 2009.
- 69 RIPE NCC. RIPE atlas IP echo measurements in IPv4. [https://atlas.ripe.net/measurements/\[1001,1004,1005,1006,1008,1009,1010,1011,1012,1013,1014,1015,1016\]/](https://atlas.ripe.net/measurements/[1001,1004,1005,1006,1008,1009,1010,1011,1012,1013,1014,1015,1016]/), 2021q3.
- 70 RIPE NCC. RIPE atlas IP traceroute measurements in IPv4. [https://atlas.ripe.net/measurements/\[5001,5004,5005,5006,5008,5009,5010,5011,5012,5013,5014,5015,5016\]/](https://atlas.ripe.net/measurements/[5001,5004,5005,5006,5008,5009,5010,5011,5012,5013,5014,5015,5016]/), 2021q3.
- 71 BBC News. Russia internet: Law introducing new controls comes into force. website <https://www.bbc.com/news/world-europe-50259597>, March 2019.
- 72 Ramakrishna Padmanabhan, Amogh Dhamdhere, Emile Aben, kc claffy, and Neil Spring. Reasons dynamic addresses change. In *Proceedings of the ACM Internet Measurement Conference*, pages 183–198, Santa Monica, CA, USA, November 2016. ACM. doi:10.1145/2987443.2987461.
- 73 Jonathan B. Postel. Internetwork protocol approaches. *IEEE Transactions on Computers*, 28(4):604–611, April 1980. doi:10.1109/TCOM.1980.1094705.
- 74 Matthew Prince. Cloudflare outage on November 18, 2025. Cloudflare blog <https://blog.cloudflare.com/18-november-2025-outage/>, November 2025.
- 75 Lin Quan, John Heidemann, and Yuri Pradkin. Trinocular: Understanding Internet reliability through adaptive probing. In *Proceedings of the ACM SIGCOMM Conference*, pages 255–266, Hong Kong, China, August 2013. ACM. doi:10.1145/2486001.2486017.
- 76 Dan Rayburn. Google blocking IPv6 adoption with Cogent, impacting transit customers. <https://seekingalpha.com/article/3948876-google-blocking-ipv6-adoption-cogent-impacting-transit-customers>, 03 2016.
- 77 Dan Rayburn. Google blocking IPv6 adoption with Cogent, impacting transit customers. web page <https://seekingalpha.com/article/3948876-google-blocking-ipv6-adoption-cogent-impacting-transit-customers>, March 2016. URL: <https://seekingalpha.com/article/3948876-google-blocking-ipv6-adoption-cogent-impacting-transit-customers>.

- 78 Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. Address allocation for private internets. RFC 1918, Internet Request For Comments, February 1996.
- 79 Reuters. 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/>, July 2021. URL: <https://www.reuters.com/technology/russia-disconnected-global-internet-tests-rbc-daily-2021-07-22/>.
- 80 Reuters. website <https://www.reuters.com/technology/us-firm-cogent-cutting-internet-service-russia-2022-03-04/>, July 2022. URL: <https://www.reuters.com/technology/us-firm-cogent-cutting-internet-service-russia-2022-03-04/>.
- 81 Philipp Richter, Ramakrishna Padmanabhan, Neil Spring, Arthur Berger, and David Clark. Advancing the art of Internet edge outage detection. In *Proceedings of the ACM Internet Measurement Conference*, pages 350–363, Boston, Massachusetts, USA, October 2018. ACM. doi:10.1145/3278532.3278563.
- 82 Philipp Richter, Florian Wohlfart, Narseo Vallina-Rodriguez, Mark Allman, Randy Bush, Anja Feldmann, Christian Kreibich, Nicholas Weaver, and Vern Paxson. A multi-perspective analysis of carrier-grade NAT deployment. In *Proceedings of the ACM Internet Measurement Conference*, Santa Monica, CA, USA, November 2016. ACM. doi:10.1145/2987443.2987474.
- 83 RIPE NCC Staff. RIPE Atlas: A global Internet measurement network. *The Internet Protocol Journal*, 18(3):2–26, September 2015.
- 84 Sen. John D. Rockefeller. Cybersecurity act of 2010. <https://www.congress.gov/bill/111th-congress/senate-bill/773>, 2009.
- 85 Root Operators. <http://www.root-servers.org>, April 2016.
- 86 J. Rosenberg, J. Weinberger, C. Huitema, and R. Mahy. STUN—simple traversal of user datagram protocol (UDP) through network address translators (NATs). RFC 3489, Internet Request For Comments, December 2003.
- 87 Tarang Saluja, John Heidemann, and Yuri Pradkin. 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*, page to appear, Portland, OR, USA, December 2022. IEEE.
- 88 Tarang Saulja and Yuri Pradkin. RIPE Atlas islands and peninsulas. https://ant.isi.edu/ripe_atlas_islands/, September 2022.
- 89 Brandon Schlinker, Hyojeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V. Madhyastha, Italo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. Engineering egress with Edge Fabric: Steering oceans of content to the world. In *Proceedings of the ACM SIGCOMM Conference*, pages 418–431, Los Angeles, CA, USA, August 2017. ACM. doi:10.1145/3098822.3098853.
- 90 Aaron Schulman and Neil Spring. Pingin’ in the rain. In *Proceedings of the ACM Internet Measurement Conference*, pages 19–25, Berlin, Germany, November 2011. ACM. doi:10.1145/2068816.2068819.
- 91 Anant Shah, Romain Fontugne, Emile Aben, Cristel Pelsser, and Randy Bush. Disco: Fast, good, and cheap outage detection. In *Proceedings of the IEEE International Conference on Traffic Monitoring and Analysis*, pages 1–9, Dublin, Ireland, June 2017. Springer. doi:10.23919/TMA.2017.8002902.
- 92 D. Smallberg. Who talks TCP? RFC 832, Internet Request For Comments, December 1982.
- 93 Berhan Taye and Sage Cheng. Report: the state of internet shutdowns. blog <https://www.accessnow.org/the-state-of-internet-shutdowns-in-2018/>, 8 July 2019.
- 94 ThinkBroadband. NTT/Cogent peering dispute increasing latency for some routes. <https://www.thinkbroadband.com/news/9896-ntt-cogent-peering-dispute-increasing-latency-for-some-routes>, Feb. 16 2024.

- 95 Craig Timberg and Paul Sonne. Minutes before Trump left office, millions of the Pentagon’s dormant IP addresses sprang to life. *The Washington Post*, Apr. 24 2021. URL: <https://www.washingtonpost.com/technology/2021/04/24/pentagon-internet-address-mystery/>.
- 96 Paul F. Tsuchiya and Tony Eng. Extending the IP Internet through address reuse. *ACM Computer Communication Review*, 23(1):16–33, January 1993.
- 97 USC/LANDER Project. Internet outage measurements. listed on web page <https://ant.isi.edu/datasets/outage/>, October 2014.
- 98 USC/LANDER Project. Internet outage measurements. IMPACT ID USC-LANDER/LANDER:internet_outage_adaptive_a30all-20171006 at https://ant.isi.edu/datasets/internet_outages/, October 2017.
- 99 Gerry Wan, Liz Izhikevich, David Adrian, Katsunari Yoshioka, Ralph Holz, Christian Rossow, and Zakir Durumeric. On the origin of scanning: The impact of location on internet-wide scans. In *Proceedings of the ACM Internet Measurement Conference*, pages 662–679, Pittsburgh, PA, USA, October 2020. ACM. doi:10.1145/3419394.3424214.
- 100 Feng Wang, Zhuoqing Morley Mao, Jia Wang, Lixin Gao, and Randy Bush. A measurement study on the impact of routing events on end-to-end Internet path performance. In *Proceedings of the ACM SIGCOMM Conference*, pages 375–386, Pisa, Italy, August 2006. ACM. doi:10.1145/1159913.1159956.
- 101 Samuel Woodhams and Simon Migliano. The global cost of internet shutdowns in 2020. <https://www.top10vpn.com/cost-of-internet-shutdowns/>, 01 2021.
- 102 Kok-Kiong Yap, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holliman, Gary Baldus, Marcus Hines, Taeun Kim, Ashok Narayanan, Ankur Jain, Victor Lin, Colin Rice, Brian Rogan, Arjun Singh, Bert Tanaka, Manish Verma, Puneet Sood, Mukarram Tariq, Matt Tierney, Dzevad Trumic, Vytautas Valancius, Calvin Ying, Mahesh Kallahalla, Bikash Koley, and Amin Vahdat. Taking the edge off with Espresso: Scale, reliability and programmability for global Internet peering. In *Proceedings of the ACM SIGCOMM Conference*, pages 432–445, Los Angeles, CA, USA, August 2017. ACM. doi:10.1145/3098822.3098854.
- 103 Sebastian Zander, Lachlan L. H. Andrew, and Grenville Armitage. Capturing ghosts: Predicting the used IPv4 space by inferring unobserved addresses. In *Proceedings of the ACM Internet Measurement Conference*, pages 319–332, Vancouver, BC, Canada, November 2014. ACM. doi:10.1145/2663716.2663718.

A Discussion of Research Ethics

Our work poses no ethical concerns as described in §1. We elaborate here.

First, we collect no additional data, but instead reanalyze data from several existing sources (see Appendix D.1 of [9]). Our work therefore poses no additional load on the Internet, nor any new risk from data collection.

Our analysis poses no risk to individuals because our subject is network topology and connectivity. There is a slight risk to individuals in that we examine responsiveness of individual IP addresses. With external information, IP addresses can sometimes be traced to individuals, particularly when combined with external data sources like DHCP logs. We avoid this risk in three ways. First, we do not have DHCP logs for any networks (and in fact, most are unavailable outside of specific ISPs). Second, we commit, as research policy, to not combine IP addresses with external data sources that might de-anonymize them to individuals. Finally, except for analysis of specific cases as part of validation, all of our analysis is done in bulk over the whole dataset.

We do observe data about organizations such as ISPs, and about the geolocation of blocks of IP addresses. Because we do not map IP addresses to individuals, this analysis poses no individual privacy risk.

Finally, we suggest that while our work poses minimal privacy risks to individuals, to also provides substantial benefit to the community and to individuals. For reasons given in the introduction it is important to improve network reliability and understand how networks fail. Our work contributes to that goal.

Our work was reviewed by the Institutional Review Board at our university and because it poses no risk to individual privacy, it was identified as non-human subjects research (USC IRB IIR00001648).

B Proof of Majority Enforcing One or No Internet

Our definition in §2.1 is complete, and Bitcoin provides an example of majority forcing consensus. However, here we provide a proof and discuss scenarios that, at first glance, may appear challenging.


Our conceptual definition is “the strongly-connected component of more than 50% of active, public IP addresses that can initiate communication with each other”, is chosen to ensure there can be only one Internet in each address space (IPv4 and IPv6). We next prove this definition yields one result, both with and without peninsulas.

The reasoning for this choice in §2.1 is straightforward: if a connected component has some fraction A , where $1 > A > 0.5$, then this component *must* be larger than any other component B . One can prove this by contradiction: (i) assume some B' exists, such that $B' > A$. (ii) Since $A > 0.5$, then (i) implies $B' > 0.5$. (iii) We then must conclude that $A + B' > 1$, but by definition, we measure only the whole address space, so it is also required that $A + B' \leq 1$. Therefore $B' < A$ and A forces a single clear component. Q.E.D.

Resolving competing “cores”: This definition handles cases with multiple overlapping but incompletely communicating groups. If members of those groups can reach half the active addresses, they are part of the Internet even if some are on peninsulas relative to each other. Consider a simplified version of Figure 1 with only three pluralities of connectivity, A , B , and C , each representing one third of the addresses, where A and B are strongly and directly connected, and A and C are strongly and directly connected, but B and C cannot directly reach each other. (Recall that strong connections in graph theory means bi-directional connectivity, but it does not require *direct* and allows connections through multiple hops.) In this example, B and C can reach each other, but only through A , so they are strongly connected but not directly connected. Our Internet core requires strong connections, but if it required direct connections, it would become a clique (a fully connected graph), forbidding peninsulas.

In this example there are two, partially overlapping, large, components that are both strongly and directly connected: $A \cup B$ and $A \cup C$. Here *all* ($A \cup B \cup C$) are part of the Internet, because any address can directly reach more than half of the active addresses: address $b \in B$ can reach $A \cup B$, $c \in C$ can reach $B \cup C$, and $a \in A$ can reach anyone. While all addresses are in one Internet, B and C are on peninsulas. The example in Figure 1 is similar to this thought experiment. In practice, we know that peninsulas occur in less than 1% of block-time (§5.1), so typically $A \geq 0.98$, with other components $B, C < 0.01$, quite different from this theoretical case where $A = B = C = 0.33$, or an asymmetric case where $A = 0.49$ and $B = C = 0.02$. However, the definition applies whenever $A \cup B \cup C > 0.5$.

Resolving disagreements with incomplete knowledge: In the above discussion we apply our conceptual definition assuming an omniscience view of connectivity. All parties must agree that A directly reaches both B and C , but B and C can reach each other only indirectly through A . An omniscient observer must recognize they are all part of the same



■ **Figure 10** Unreachable blocks over time. Large spikes are unreachability to Chinese-allocated IPv4 addresses. Dataset: A29, 2017q3.

core, in spite of the peninsula.

In practice, no real-world system will have omniscient knowledge of connectivity. However, this scenario works even with incomplete knowledge. Imagine observers only in B and C both might assert they are “the” core, since both can observe direct, strong connectivity to more than half of the active, public addresses.


When faced with seemingly conflicting claims of what the core is, all parties must share their observations with each other to make their case. In this case, B and C will recognize they are both reporting A as part of their core, and that A overlaps—they must therefore recognize the reachable core is $A \cup B \cup C$, even though they cannot directly reach each other.

This seeming disagreement highlights the requirement that B and C recognize that the A they each measure is the same A . This requirement is met by our definition of what a public, global address space is—we assume some authority allocated addresses. In today’s Internet, this authority is IANA. Note that IANA is not saying who is in our out of the Internet, but only who is responsible for a given fraction of the address space.

If all parties cannot agree on a shared address space, then our definition cannot be used. For example, if one party asserts the entire 0/0 IPv4 address space is theirs to reallocate, then one cannot use address to resolve disputes. Fortunately, address assignment has historically been coordinated to avoid overlaps. (One exception is DISA’s 4/8 prefixes. These were clearly allocated to DISA, but lack of global routing prompted multiple organizations to squat on them, using them as additional private address space. Fortunately this variance is not a practical problem for several reasons: Since 2021 DISA has announced routes for these blocks on the public Internet. Their actual allocation has never been disputed. And even if they were disputed, this 4/256ths of the address space is not enough to change control of a majority.)

C Additional Results about Islands

We define islands and give examples in §2.3.2. Here we supplement those results with examples of country-sided islands (§5.7.2). We also show the raw data we use to justify our choice of 50% unreachability to define islands in Trinocular (§C.2).



■ **Figure 11** Islands detected across 3 years using six VPs. Datasets A28-A39.

C.1 Visualizing Potential 2017q3 Islands

In §5.7.2 we discuss evidence for country-sized islands. In 2017q3, on 8 occasions it appears that most or all of China stopped responding to external pings. Figure 10 shows the number of /24 blocks that were down over time, each spike more than 200k /24s, between two to eight hours long.

C.2 Longitudinal View Of Islands

We first consider three years of Trinocular data (described in Appendix D.1 of [9]), from 2017-04-01 to 2020-04-01. Figure 11 shows the fraction of the Internet that is reachable as a dotted line at the 50% threshold that Chiloe uses to detect an island (§3.2). We run Chiloe across each VP for this period.