

Ebb and Flow: Implications of ISP Address Dynamics

Guillermo Baltra^{1,2}, Xiao Song^{1,2}, and John Heidemann^{1,2}

¹ University of Southern California, Los Angeles CA 90089, USA

² Information Sciences Institute, Marina del Rey CA 90292, USA

{baltra,songxiao,johnh}@isi.edu

Abstract. *Address dynamics* are changes in IP address occupation as users come and go, ISPs renumber them for privacy or for routing maintenance. Address dynamics affect address reputation services, IP geolocation, network measurement, and outage detection, with implications of Internet governance, e-commerce, and science. While prior work has identified diurnal trends in address use, we show the effectiveness of Multi-Seasonal-Trend using Loess (MSTL) decomposition to identify both daily and weekly trends. We use ISP-wide dynamics to develop IAS, a new algorithm that is the first to automatically detect ISP maintenance events that move users in the address space. We show that 20% of such events result in /24 IPv4 address blocks that become unused for days or more, and correcting nearly 41k false outages per quarter. Our analysis provides a new understanding about ISP address use: while only about 2.8% of ASes (1,730) are diurnal, some diurnal ASes show more than 20% changes each day. It also shows greater fragmentation in IPv4 address use compared to IPv6.

1 Introduction

Millions of devices connect to the Internet everyday, but some come and go. Many ISPs *dynamically assign* devices to public IP addresses. While some users have IP addresses that are stable for weeks, ISPs often reassign users for many reasons: to promote privacy, to prevent servers on “home” networks, and to shift users away from routers scheduled for maintenance. IP policies vary: some renumber users every day [33, 17, 23, 20, 28], some show large diurnal changes [26].

Understanding ISP address dynamics is important in Internet policy, network measurement, and security. In Internet policy, ISPs need to make business-critical decisions that include purchasing carrier-grade NAT equipment versus acquiring more address space, or evaluating the costs of carefully reusing limited IPv4 space versus transitioning to IPv6. Regulators like national or Regional Internet Registries (RIRs) must consider address dynamics when crafting policies about transferring limited IPv4 address space and tracking IPv4 and IPv6 routing table sizes. For network measurement and security, dynamics affect services like IP address reputation [24, 7], IP geolocation [15], and generating IPv4 [8] and

IPv6 hitlists [9, 11, 18, 5, 10, 34]. Stable addresses also simplify attack targeting, traffic fingerprinting, and have implications for privacy and anonymity.

The topic of address dynamics has been explored previously. Some have shown the stagnation of the total number of active IPv4 addresses and identified address block activity patterns [28]. Others have tracked address changes for a subset of addresses and analyzed lease durations [20, 21], studied diurnal patterns where blocks stay active during the day, but remain inactive during the night [26], and built a statistical model from few ISPs to provide address churn estimation [17]. While important, all prior work focuses on behavior inside specific address blocks, not ISP-wide.

Address dynamics also affect the accuracy of outage detection. Diurnal changes can misinterpret nighttime quiet as outages [26] and must be considered in passive data [12]. CDN-based outage detection showed that ISP-level user movement, often correlated with scheduled router maintenance, can produce false outages [27]. Although CDN-measurements were more robust to ISP-level events than prior work, they did not quantify how often maintenance events happen, nor suggest how to address this problem in other systems. Finally, recent work showed that tracking changes in address use can detect changes in human behavior, such as shifts to working from home [31]. The goal of this paper is to explore how these related outcomes—diurnal effects, maintenance detection, and tracking overall user changes—all benefit from an improved understanding of address dynamics. We then use this understanding to build more accurate models of address activity, improving these existing services.

The primary contribution of this paper is to develop ISP Availability Sensing (IAS, §3.4), a new algorithm that identifies maintenance events in the ISP, allowing us to recognize that apparent outages are actually users being reassigned. This algorithm uses ISP Diurnal Detrending (IDD, §3.3), to separate daily and weekly patterns from underlying trends and residual, both of which are important in detection algorithms. Our second contribution is to validate IAS using data from ISPs with known maintenance patterns and data from RIPE Atlas (§4.1). Although such events have been previously identified, IAS and this validation are the first to automate their detection. Our final contribution is to use IAS and IDD for three new observations: we quantify how many ISPs are diurnal (§5.4), how many maintenance events occur (§5.1), and how IPv6 shows more consistent address usage than IPv4 (§5.2).

All of the data used and created in this paper is available at no cost [32]. While our primary datasets date from 2017 to 2020, when data from 6 Trinocular sites was available, we confirm key results hold in 2022 (§5.4). Evidence of diurnal behavior in networks has been shown over nearly a decade of observations (from [26] to [31]), so we expect our algorithms and our qualitative observations to apply today. Our work poses no ethical concerns (§A): although we use data on individual IP addresses, we have no way of associating them with individuals. Our work was IRB reviewed and identified as non-human subjects research (USC IRB IIR00001648).

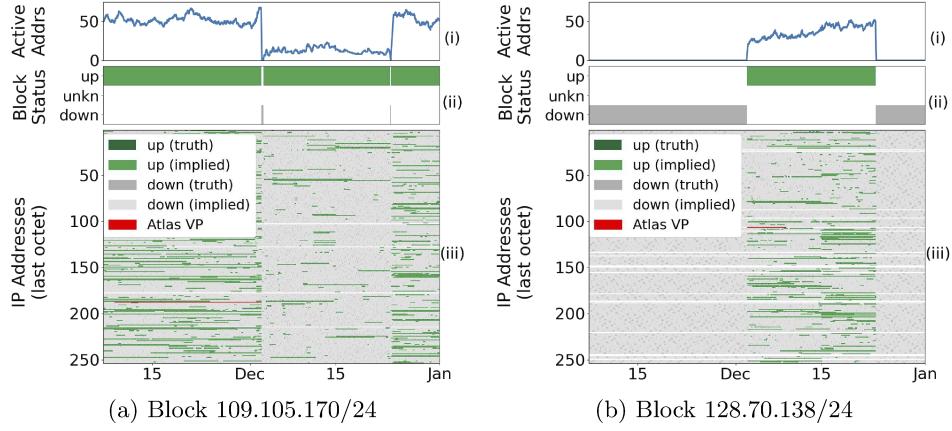


Fig. 1: Sample Atlas VP (Id 1001049) shifting between different /24 blocks during maintenance events, 2020q4. Datasets: Trinocular A42N, Atlas MsmId 1010.

2 Implications of Address Dynamics

This paper examines two challenges in how addresses are managed: maintenance events and diurnal networks. Both cause problems to outage detection systems because they cause individual /24 blocks to become vacant and so appear unreachable, resulting in a *false outage*. An outage is incorrect because users are receiving service elsewhere (due to maintenance or other address reassignment), or are sleeping (diurnal changes). We focus on blocks going fully vacant, since outage detection already accommodates blocks that partially reduce use [26].

Our premise is that a *whole-ISP* viewpoint is necessary to address these challenges. In fact, initially we hoped we could fix algorithms that examined only individual /24s. However, because ISPs can shift addresses from one to many /24s, we believe that a combination of block-level and ISP-level examination is necessary to provide outcomes robust to now known challenges.

ISP Maintenance Events: Figure 1 shows two /24 address blocks, with green dots showing ping responses (the bottom area marked (iii)), block status (area (ii) shown up, unknown, or down), and a count of number of active addresses (the line graph in area (i)). We have a RIPE Atlas VP [30] that moved from the left block to the right during three weeks in December—a maintenance event that left the left block mostly idle. This event is difficult for an external outage detection system to handle: gaps on December 3 and 24 are outages in the left block, and the much lower utilization for the period is hard to track. Yet someone watching the *whole Autonomous System* (AS) would realize users shifted addresses temporarily. In §3.1 we show how we build an AS-level view, and in §3.4 how we can avoid false outages from these type of events, while in §5.3 we show how often these events occur.

Diurnal ASes: Other ASes have large changes in user populations over the course of a day. In these *diurnal ASes*, dynamic users (possibly on mobile devices, or connecting during the workday) disconnect at night, leaving their

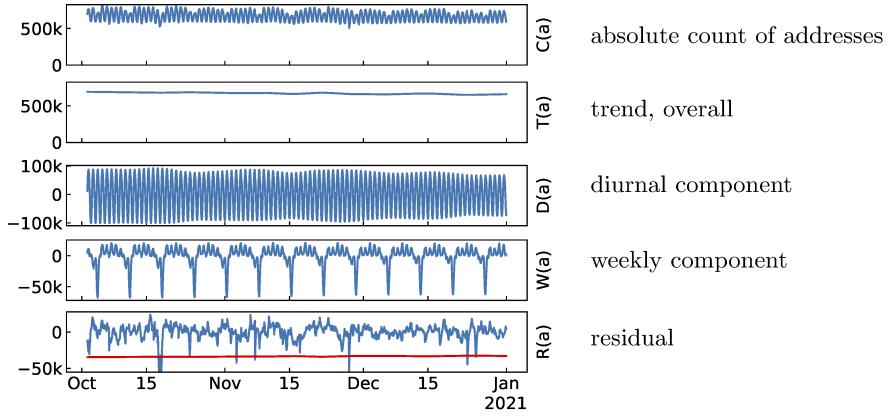


Fig. 2: MSTL decomposition of AS9829 during 2020q4. Dataset: A42.

addresses unused. Although many users in North America and Europe have always-on home routers with device changes hidden behind NAT, many users in Asia, South America, and Africa disconnect at night [26].

The top graph in Figure 2 shows the number of active IP addresses in AS9829 (Bharat Sanchar Nigam Limited) over three months, based on measurements updated every 11 minutes (§3.1). With an average of 600k active addresses (the second graph), this major national Indian ISP has *many* active users. But the timeseries in the top shows daily changes in the number of active addresses of $\pm 20\%$, more than 100k users! In addition, close examination shows activity drops for two days after every five—evidence of weekends. These trends show that this AS has *strong diurnal and weekly trends*.

Such large shifts cause problems for outage detection systems, because losing 100k users every night vacates some /24 blocks. (§B shows a specific block with daily outages.) Tracking outages across ASes with this much daily change motivates diurnal AS detection (§3.2) and detrending (§3.3). We show the importance of detrending and tracking ISP-wide behavior in §4.2.

3 Methodology

We next describe how we track AS-wide address usage, identify diurnal ASes and trends, and detect maintenance events.

3.1 AS-wide Address Accumulation

Since ASes move users around in their address space, we find that *the number of active addresses across the AS* helps characterize the current population. The first step in our AS-wide algorithms is to track active addresses. Determining active addresses in an AS is difficult because some ASes have millions of addresses, too many to monitor instantly—we reanalyze existing data to get an approximate snapshot of global state.

Our input is from Trinocular [32], since it is publicly available and it has years of address-specific ping responses covering much of the Internet (4M to 5M /24

blocks). Trinocular sends between 1 and 16 ICMP echo-requests to each block every 11 minutes, each to a different address. Addresses rotate in a fixed order unique to each block, so a single Trinocular site will scan all planned addresses in 48 hours or less.

We accumulate individual observations from incremental Trinocular scans to approximate current state following prior work [1, 31] and validation [31, 3]. Combining results from all six Trinocular sites cuts worst-case latency to eight hours (each site scans independently with different and varying phases). We update estimates incrementally each 11-minute round, so even this worst case usually tracks diurnal changes. We apply 1-loss repair to recover from measurement loss [25]. For efficiency, we aggregate results by the hour.

We add AS information from Routeviws [16] and combine reports for all addresses in each AS. The result is $C_i(a)$, a timeseries *counting addresses* for each AS a at time i .

We currently treat each AS as independent. Although most large ISPs employ multiple ASes (for example, one each for Asian, American, and European operations), in §5.2 we show that renumbering usually occurs within the same AS, so this simplification does not change our primary results.

3.2 Diurnal ISP Detection

Given $C(a)$, address counts for an AS (§3.1), our next step is to identify ASes with a strong diurnal component. Following prior block-level diurnal analysis [26], we take the Fast-Fourier Transform (FFT) of this timeseries, giving a set of coefficients showing the strength and phase at all frequencies. We then label that AS as diurnal if the energy in the frequency corresponding to a 24-hour period is the largest of all other (non-zero frequency) components.

3.3 ISP Diurnal Detrending (IDD)

Since we know some ASes are strongly diurnal, we next *decompose* $C(a)$ to extract long-term trends, cyclic components, and any residual changes. Each component is useful to identify usual events.

We apply MSTL [4] to extract *four* components, one for diurnal (daily) behavior, one for weekly patterns, along with trend and residual components. We find some networks have both diurnal and weekly patterns, while others are only diurnal. We decompose $C(a)$ in four components: trend (T), diurnal (D), weekly (W), and residual (R) components.

Figure 2 shows trend decomposition of AS9829 during 2020q4 (three months). The top graph shows the AS-wide timeseries $C(a)$, ranging from 500k to 800k active addresses. The next graph down $T(a)$ shows the long-term trend. We can see that this AS has a static user population over this quarter.

The third and fourth graphs show $D(a)$ and $W(a)$, how much regular change there is each day and week. The strong diurnal pattern that we first identified at the 24 h frequency in the FFT (§3.2) shows up in $D(a)$ with swings that range across 30% of responsive addresses ($\pm 100k$). The weekly component ($W(a)$) shows a weekend drop of about 50k addresses. Diurnal and weekly trends are both visible in $C(a)$, but easier to quantify after decomposition.

The residual in the final row, $R(a)$, isolates any remaining changes. We use this residual when detecting address dynamics in §5.1.

3.4 ISP Availability Sensing (IAS)

The ISP Availability Sensing algorithm (IAS) recognizes maintenance events by comparing a global count of active users at AS-level against local changes in portions of the network. Our insight is that the AS-wide count of active users *remains stable* during maintenance, even though specific parts of the network *lose* and *add* users. This stability distinguishes user movement from outages.

Detecting AS-wide Address Stability We first show that the AS' active addresses are roughly stable.

We define Δ_i as the relative fraction of change of active addresses across an AS at time interval i , using the residual and trend decomposition from §3.3: $\Delta_i = R_i(a)/T_i(a)$.

When $\Delta_i = 0$ there is no change in number of active users. Of course we expect some accident changes ($\Delta_i \neq 0$) as individual hosts come and go in real networks, or due to loss of probing packets or replies. These small changes appear in the residual, $R(a)$, in Figure 2. Finally, we identify large changes ($\Delta_i \geq 0.05$) as outages, while smaller changes ($0 < \Delta_i < 0.05$) are more typical jitter. We select a threshold large enough to avoid noise, but not so large as to miss outages. As future work we hope to optimize this value by training against external observations.

IAS assumes complete knowledge of each AS' address space. However, ASes bring new address space on-line to serve new customers, such space may not immediately appear in $C(a)$. We evaluate how frequently we miss users due to unmonitored address space in §4.3.

Detecting network changes IAS' second requirement is the presence of some blocks changing. We enforce this requirement by identifying the number of blocks that change state (become or cease being reachable) from outage detection. We currently require $\delta = 4$ blocks to change state, and study this choice in §4.2. Larger values reduce the number of events (Figure 3). We select $\delta = 5$ because it is just past the knee of the curve, in a plateau $5 \leq \delta \leq 9$ (Table 1). Exploration of an adaptive threshold to account for larger ASes is future work.

4 Validation

We validate IAS against external sources, and then examine design choices.

4.1 IAS Detecting Known ISP Maintenance?

We next evaluate IAS' ability to detect maintenance events using ground truth from an ISP. Our first source of ground truth are ISPs that have public maintenance windows; we use this case study to validate. Lumen Technologies (AS209, at the time called CenturyLink) announces that midnight to 6 a.m. local time [14] is a public maintenance window, and they report specific events [6]. We identify Lumen address blocks from 18 peers in Routeviews [16].

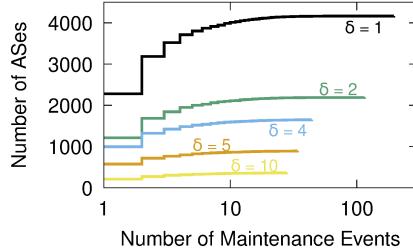


Fig. 3: Cumulative number of ASes with x or fewer maintenance events (2020q4).

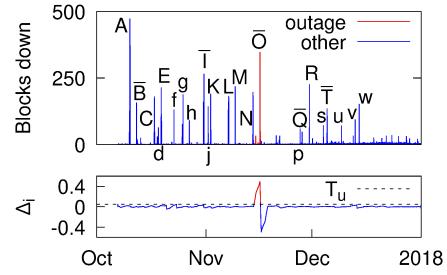


Fig. 4: Down events (top) from six observers (top) and Δ_i (bottom) from Los Angeles. Lumen AS209. Dataset: A30, 2017q4.

Figure 4 shows data for 2017q4. The top graph counts the number of blocks that are down over time, from all Trinocular data, but without IAS.

IAS finds 23 events in this merged data, each involving 35 or more blocks. Of the 23 events, more than half (13, indicated with capital letters) are in Lumen’s published maintenance window, suggesting they are maintenance events. Five of these (with lines over the letters) are documented events and datacenters from the service log on their website. (We do not examine regional aspects of maintenance, since our work focused on ISP-wide dynamics.) IAS identifies *all* these events as maintenance, except for event (\bar{O}) . The best-available ground truth (service logs and maintenance window), suggests that these events are true positives.

We believe event (\bar{O}) on 2017-11-16 is a true outage. We show Δ_i , the changes in the number of active addresses from one observer, in the bottom graph. The 40% drop in active addresses, followed by the recovery indicates an outage. IAS classifies event (\bar{O}) as an outage, not a maintenance event. This event is unusual, in that it affected Los Angeles VP mostly (20,211 blocks and 8.5 hours, not shown in graph) than from other sites (where it was 348 blocks and 2 hours).

2024-01-23 It shows that the IAS will correctly pass through large outages (a true negative).

4.2 Validating IAS and IDD from RIPE Atlas

RIPE Atlas VPs live in edge networks and report their current IP address, providing ground truth for ISP maintenance.

Atlas as Ground Truth We take Atlas VPs built-in measurement 1010 [29] as our known addresses. We aggregate VP address changes using 4-minute timebins, since new address reports are provided every 4 minutes. We omit address changes where Atlas VP failed to reach a root DNS server to rule out address changes due to outages [20]. Finally, to rule out individual VP changes, we require four VPs in the same AS to move at about the same time to declare a maintenance event. During 2020q4, we count 164 events by this criteria.

		blocks that change state (δ_i)											all			
		0	1	2	3	(1-3)	4	5	6	7	8	9	≥ 10	(≥ 4)		
Addr. drop	without	$\Delta_i \leq 5\%$	42	15	9	4	(28)	11	4	7	8	6	5	42	(83)	153
	IDD	$\Delta_i > 5\%$	1	0	0	0	(0)	0	0	0	0	0	0	10	(10)	11
	with	$\Delta_i \leq 5\%$	43	15	9	4	(28)	11	4	7	8	6	5	52	(93)	164
	IDD	$\Delta_i > 5\%$	0	0	0	0	(0)	0	0	0	0	0	0	0	(0)	0

Table 1: Atlas VP address change events (≥ 4 VPs) compared against IAS detection thresholds, 2020q4.

Validating IAS We first validate IAS with IDD, considering its two requirements: stable AS-level addresses ($\Delta_i \leq 5\%$) and four or more blocks that move ($\delta_i \geq 4$). We show our results in [Table 1](#).

We first look at the bottom two rows of [Table 1](#), labeled “with IDD”. The fourth row shows no blocks change when $\Delta_i > 5\%$ because of the large threshold. The third row shows 164 blocks where VPs move, of which 93 are found in IAS and occur as part of a large movement (right, in green), while 43 move by themselves (left, gray), and 28 move with 1 to 3 others (center, yellow). We consider the 93 to be IAS successes. All will be found and recognized as maintenance events.

The 43 in gray represent independent movements that are not large maintenance events, but may be routers at home rebooting to a new address. These are not found by IAS, but are not necessarily maintenance events, so we count them neither as true nor false positives.

Finally, the 28 marked yellow are likely maintenance events that IAS misses as being too small. These are false negatives.

Not having all negative cases prevents us from computing recall and precision, but we can show a True Positive rate of 0.77 ($93/(28 + 93)$). We conclude that IAS works reasonably well, although there is room for improvement.

Validating IDD To validate the importance of IDD, we turn it off and compare the results in the first two rows of [Table 1](#) with the bottom two rows. IDD helps filter out diurnal changes, making large shifts more common: compare the 10 cases with $\Delta_i > 5\%$ without IDD to zero cases with it. We also see that it helps IAS: the TPR is 0.75 without IDD ($83/(28 + 83)$) compared to 0.77 with IDD. We conclude that accounting for diurnal changes helps.

4.3 Does Unmonitored Space Harm IAS?

Measurement systems do not track the complete address space, discarding some segments due to low response rate, or lack of any historic responses [\[2\]](#). Users reassigned to unmonitored space implies that IAS may erroneously infer outages due to drops in the total active address count, IAS false negatives.

For 2020q4, we find that the majority of reassessments (51%) occur within monitored addresses, and that most addresses (84%) stay in the same category. IAS is not impeded by incomplete measurement (see [§C](#)).

	IPv4		IPv6	
Total VPs	12,855	100.0%	6,319	100.0%
Do not change IP	8,501	66.1%	4,730	74.9%
Change IP	4,354	33.9%	1,589	25.1%
Do not change routable prefix	973	7.6%	1,182	18.7%
Change routable prefix	3,381	26.3%	407	6.4%
Do not change AS	2,411	18.8%	75	1.1%
Change AS	970	7.5%	332	5.3%

Table 2: Active RIPE Atlas Vantage Points during 2020q4

4.4 Choice of Spatial Granularity

We next consider what spatial granularity to use when tracking address dynamics. Our goal is that IAS can identify when users move. To do so, we must assess how “far” a user moves: do they stay in the same routable prefix, or within the same AS, or move between ASes. We compare address movement (a baseline) against how often a device moves within a routable prefix or an AS.

Table 2 shows how often 12,855 RIPE Atlas VPs change their address, routable prefix, or AS in 2020q4, for both IPv4 and IPv6. We see that the majority of devices are stable, with 66.1% (v4) and 74.9% (v6) never changing address and 7.6% and 18.7% staying in the same routable prefix. Of the remaining that move, some change only once, but many change frequently, perhaps because they are in ISPs that renumber their users regularly. We conclude that most devices are very stable, but a few move frequently,

Surprisingly, we find about 7.5% (970 v4) and 5.3% (332 v6) change AS. As changes are very rare, with a few (2%) changing once, perhaps because a user changed their home ISP. The remaining 3% change frequently, perhaps because they are mobile and regularly move between home and work.

We conclude that AS granularity is almost always suitable to capture most movement and so IAS’ use of ASes is correct.

5 Evaluation

We now study ISP address dynamics across the Internet with IDD and IAS. We evaluate the addressing efficiency, improvements to outage detection, quantify diurnalness, and compare IPv4 and IPv6 management practices.

5.1 Quantifying ISP Address Dynamics

Several groups have looked at different aspects of address dynamics [26, 17, 28, 20, 27, 21, 31]. While prior work identified ISP maintenance as a type of network disruption [28], they did not quantify how often such events occur. We examine maintenance and diurnal events over a quarter using IAS.

We use IAS to identify maintenance events across all 63k ASes active in 2020q4. Figure 3 shows the cumulative distribution of number of maintenance events for 6.5% of ASes with at least one event in this period. We compare results for different detection thresholds, that is, the number of Trinocular blocks going

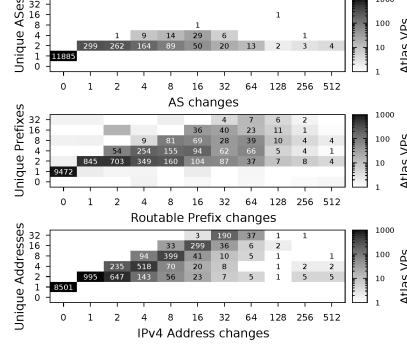


Fig. 5: IPv4 changes by AS (top), routable prefix (center), and address (bottom), for Atlas VPs with at least one change, 2020q4.

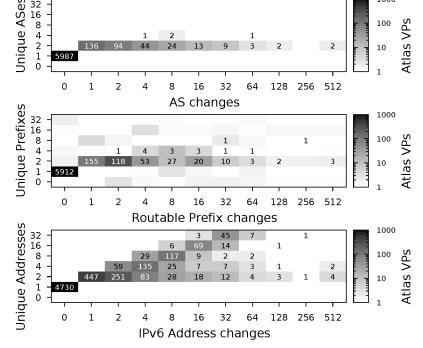


Fig. 6: IPv6 changes by AS (top), routable prefix (center), and address (bottom), for Atlas VP with at least one change, 2020q4.

up or down during the same timebin, with an AS-level responsive address count remaining stable (less than 5% drops).

With a threshold of one changed block (the minimum), IAS detects at least one event in 2k ASes. The number of ASes decreases to only 210 ASes with our strictest threshold ($\delta = 10$).

We also see that some ASes regularly move users around, seeing 100 maintenance events in these 90 days. One example of these frequent-maintenance ISPs are ISPs that renumber users every 24 hours.

The area under each curve corresponds to the number of maintenance events and diurnal events that occurred during this quarter. For our default preferred threshold of four blocks, we see 2.5k events in 2020q4. At a block level, these events cause 41k false outages.

5.2 Address Space Refactoring

Address management is a business-critical decision for ISPs. Limited IPv4 addresses require careful management and reuse, while IPv6 transition requires updating current practices. Each of these choices incurs costs. To provide ground truth for the kind of AS-level changes observed by our IDD and IAS algorithms, we next examine address churn in both IPv4 and IPv6 inferred from RIPE Atlas.

We take Atlas VP IPv4 and IPv6 address changes during 2020q4, and perform longest prefix match for these addresses using Routeviews RIB archives to obtain routable prefixes and ASes for the quarter. We count the number of times each VP changes address, routable prefix and AS, and the times unique addresses are assigned. [Figure 5](#) and [Figure 6](#) show IPv4 and IPv6 aggregates as heatmaps. This analysis of RIPE Atlas data provides new ground truth about the number of maintenance events we aim to detect with IAS.

For IPv4 *address* changes ([Figure 5](#), bottom) we observe that most VPs do not change address during the quarter, but those that do change, often are

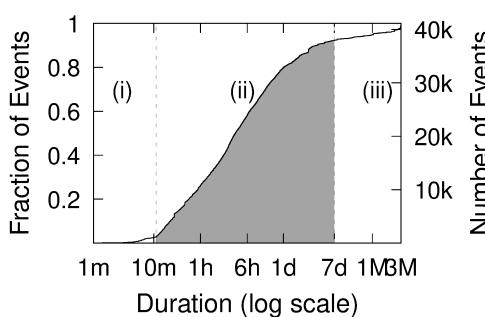


Fig. 7: CDF of unresponsive duration of blocks before or after an IAS-detected maintenance event, 2020q4.

	IPv4	
ASes	62,310	100.0%
Diurnal	1,730	2.8%
Non-diurnal	60,580	97.2%
Routable prefixes	606,187	100.0%
Diurnal	30,029	5.0%
Non-diurnal	576,158	95.0%
/24 blocks	5,124,967	100.0%
Diurnal	111,908	2.2%
Non-diurnal	5,013,059	97.8%

Table 3: Number of diurnal networks at different granularities. 2022q4.

assigned a new address, as the dark heatmap diagonal shows. Prefixes (middle) are mostly reused, although many address changes involve a routable prefix change, too. Finally, addresses almost always stay in the same AS (top).

On the other hand, in IPv6 we observe that address changes generally occur within the same prefix. Our data confirms that IPv4 exhaustion and fragmentation make address management more challenging, while IPv6 uses fewer prefixes more efficiently because addressing aligns better with routing. While prior work has mentioned these trends [20, 21], we look at both protocols together. Quantifying behavior at granularities other than routable prefixes is future work.

5.3 How Often Does IAS Repair False Outages?

Analysis of CDN traffic [27] showed that incorrect block-level outages are often due to users being assigned to different IP addresses. Measurements from end-user devices show users changing addresses, but external outage detection systems cannot distinguish the now-vacant old address block from a network problem. After users are reassigned, their old address blocks remain empty for minutes or months, and external outage detection systems (like Trinocular or Thunderping) incorrectly interpret this absence as an outage.

Figure 7 shows the distribution of durations of the time a block stays unresponsive before or after the 41k address reassessments that are detected by IAS during 2020q4 (§5.1). All of these events are false outages that IAS repairs. To understand root causes we identify three regions of unresponsive duration.

Durations less than 11 minutes (the bottom 2% on the left, about 800 unshaded events) are less than the scanning frequency of our data source (Trinocular). Such short outages occur in blocks with large numbers of scanned addresses where only a few are in use (blocks with large $|A(b)|$ and small A , with terms from [25]). These false outages of shorter-than-probing-interval events is natural due to the nature of active probing at regular intervals; IAS suggests these are measurement “noise”.

The majority of events (the center, shaded region, 88% or about 36k events) are blocks that are inactive between 11 minutes and one day. These false outages are typically due to diurnal address assignment, when customers are regularly

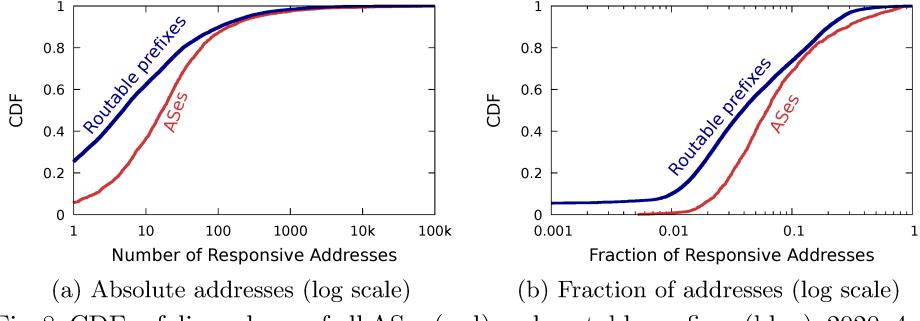


Fig. 8: CDFs of diurnalness of all ASes (red) and routable prefixes (blue); 2020q4.

reassigned, and part of the assigned blocks sometimes appear to have no activity. IAS detects and corrects these false outages because it knows the AS-wide activity is constant.

Finally, about 10% or 4k blocks are empty for more than a week. Trinocular already classified these blocks as “gone dark”, inferring that a long period of no responses cannot be a transient outage, but must be ISP renumbering. Our AS-wide analysis with IAS confirms that this policy is correct.

5.4 How Many ASes Are Diurnal?

Diurnal networks are important to assess allocation policies in Internet governance and to avoid false outages in outage detection. While diurnal behavior has previously studied at the */24* block granularity [26, 31], we next examine diurnalness in the larger groupings of routable prefixes and ASes.

Here we use address accumulation data for 2022q4 from Trinocular, following §3.1, grouped by prefixes and ASes from Routeviews [19] on 2022-10-01. We assess diurnalness as described in §3.2 .

In Table 3 we show how many diurnal networks are detected at different granularities. While only 2.8% of ASes are diurnal (1,730), 112k blocks are diurnal. Recent work has used these blocks to help understand human activity [31].

5.5 How Much of a Diurnal AS is Diurnal?

Although we can identify networks as diurnal as we described in §3.2, in many ISPs, only *part* of the AS is diurnal, while part is more static.

Here we examine what fraction of an AS’ address space is diurnal. Prior work has examined individual blocks, but our decomposition allows us to examine “diurnalness” for the AS as a whole. We judge AS-level diurnalness by the size of the daily change in addresses. From our MSTL decomposition (§3.3), that is $(P_{95}(D(a)) - P_5(D(a)))/P_{95}(C(a))$, where P_n is the n -percentile of the given timeseries over the quarter.

Figure 8 shows how diurnal ASes (red) and routable prefixes (blue) are by numbers (Figure 8a) and fraction (Figure 8b) of responsive addresses.

First, we see that *most networks are not very diurnal*: activity in 85% of ASes change by 100 addresses or fewer each day, accounting for only 20% of their

address space. This stability is typical of ISPs with customers using always-on home gateways. Stable address usage is why IAS can detect maintenance events.

When comparing routable prefixes to ASes, we see that routable prefixes are more often mostly diurnal (comparing the two lines in [Figure 8b](#)). Although most prefixes are fairly stable (69% change by only 10% of their active addresses), some (about 20%) have a very large daily swing (15% of addresses or more). Finally, because routable prefixes are necessarily smaller than ASes, they see a smaller absolute size of diurnal change (compare the lines in [Figure 8a](#)).

This trend suggests that routable prefixes are a useful size to study diurnality, and it supports suggestion for its study in [§5.4](#).

6 Related Work

Other work has considered maintenance events in relation to outages. Richter et al. used internal information from clients to demonstrate that address reassignment cause false outages, defining *disruptions* to include both true and false outages [27]. Guillot et al. proposed Chocolatine to detect outages at AS level and geographical areas using Internet background radiation [12]. These works do not show how to differentiate true outages from maintenance events, nor perform a quantitative analysis of maintenance events. Recently Padmanabhan et al. showed that some events span parts of multiple blocks [22], a result consistent with our goal of studying whole ISPs.

Other groups have studied address changes and usage. Some have examined the duration hosts keep the same address [13, 20, 21], estimated Internet-wide address churn [17], and address utilization [28]. However, these techniques either do not scale to the entire address space, are estimations, or require CDN access, while we do Internet-wide, third-party detection and identify ISP renumbering.

Address counting was first used in outage detection with CDN data [27], then darknet analysis [12]. We previously described address accumulation with Trinocular [1, 31]. We use this signal to detect ISP maintenance.

Previous work has considered seasonal patterns. Quan et al. detected diurnal patterns using FFT at block level [26]; Chocolatine used SARIMA in ISP-wide detection to factor out seasonal trends [12]. Unlike prior work, we show the importance of multi-seasonal trends to account for both daily and weekend effects, and show that we can distinguish maintenance events from outages at the /24-granularity.

7 Conclusions

AS-wide diurnal changes and maintenance are part of our Internet ecosystem, yet they challenge outage detection systems. Our new IAS algorithms, with IDD, can often recover from such dynamics. We showed that IAS is effective and can correct 41k false outages per quarter.

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

Bibliography

- [1] Guillermo Baltra and John Heidemann. 2019. *Improving the Optics of Active Outage Detection (extended)*. Technical Report ISI-TR-733. USC/Information Sciences Institute.
- [2] Guillermo Baltra and John Heidemann. 2020. Improving Coverage of Internet Outage Detection in Sparse Blocks. In *Proceedings of the Passive and Active Measurement Conference*. Springer, Eugene, Oregon, USA.
- [3] Guillermo P. Baltra. 2023. *Improving network reliability using a formal definition of the Internet core*. Ph.D. Dissertation.
- [4] Kasun Bandara, Rob J Hyndman, and Christoph Bergmeir. 2021. MSTL: A Seasonal-Trend Decomposition Algorithm for Time Series with Multiple Seasonal Patterns. <https://doi.org/10.48550/ARXIV.2107.13462>
- [5] Robert Beverly, Ramakrishnan Durairajan, David Plonka, and Justin P. Rohrer. 2018. In the IP of the Beholder: Strategies for Active IPv6 Topology Discovery. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Boston, Massachusetts, USA, 308–321. <https://doi.org/10.1145/3278532.3278559>
- [6] CenturyLink. 2019. Event History. <https://status.ctl.io/history>.
- [7] Fail2ban. 2023. <https://github.com/fail2ban/fail2ban>.
- [8] Xun Fan and John Heidemann. 2010. Selecting Representative IP Addresses for Internet Topology Studies. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Melbourne, Australia, 411–423. <https://doi.org/10.1145/1879141.1879195>
- [9] Paweł Foremski, David Plonka, and Arthur Berger. 2016. Entropy/IP: Uncovering Structure in IPv6 Addresses. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Santa Monica, CA, USA, 167–181. <https://doi.org/10.1145/2987443.2987445>
- [10] Oliver Gasser, Quirin Scheitle, Paweł Foremski, Qasim Lone, Maciej Korczyński, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Boston, Mass., USA, to appear. <https://doi.org/10.1145/3278532.3278564>
- [11] Oliver Gasser, Quirin Scheitle, Sebastian Gebhard, and Georg Carle. 2016. Scanning the IPv6 Internet: Towards a Comprehensive Hitlist. In *Proceedings of the IFIP International Workshop on Traffic Monitoring and Analysis*. IFIP, Louvain La Neuve, Belgium. <http://tma.ifip.org/2016/papers/tma2016-final51.pdf>
- [12] Andreas Guillot, Romain Fontugne, Philipp Winter, Pascal Merindol, Alistair King, Alberto Dainotti, and Cristel Pelsser. 2019. Chocolatine: Outage detection for internet background radiation. In *2019 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–8.
- [13] John Heidemann, Yuri Pradkin, Ramesh Govindan, Christos Papadopoulos, Genevieve Bartlett, and Joseph Bannister. 2008. Census and Survey

of the Visible Internet. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Vouliagmeni, Greece, 169–182. <https://doi.org/10.1145/1452520.1452542>

- [14] Lumen Technologies. 2023. Lumen Master Service Agreement. <https://www.lumen.com/en-us/about/legal/business-customer-terms-conditions.html> <https://www.lumen.com/en-us/about/legal/business-customer-terms-conditions.html>
- [15] MaxMind. 2017. GeoIP Geolocation Products. <http://www.maxmind.com/en/city>.
- [16] D. Meyer. 2018. University of Oregon Routeviews. <http://www.routeviews.org>.
- [17] Giovane CM Moura, Carlos Ganán, Qasim Lone, Payam Poursaied, Hadi Asghari, and Michel van Eeten. 2015. How dynamic is the isps address space? towards internet-wide dhcp churn estimation. In *2015 IFIP Networking Conference (IFIP Networking)*. IEEE, 1–9.
- [18] Austin Murdock, Frank Li, Paul Bramsen, Zakir Durumeric, and Vern Paxson. 2017. Target Generation for Internet-wide IPv6 Scanning. In *Proceedings of the ACM Internet Measurement Conference*. ACM, San Diego, CA, USA, 242–253. <https://doi.org/10.1145/3131365.3131405>
- [19] University of Oregon. 2020. Route Views Archive Project. <http://archive.routeviews.org/bgpdata/2020.10/RIBS/rib.20201001.0000.bz2>.
- [20] Ramakrishna Padmanabhan, Amogh Dhamdhere, Emile Aben, kc claffy, and Neil Spring. 2016. Reasons Dynamic Addresses Change. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Santa Monica, CA, USA, 183–198. <https://doi.org/10.1145/2987443.2987461>
- [21] Ramakrishna Padmanabhan, John P. Rula, Philipp Richter, Stephen D. Strowes, and Alberto Dainotti. 2020. DynamIPs: Analyzing address assignment practices in IPv4 and IPv6. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*. ACM, Barcelona, Spain, 55–70. <https://doi.org/10.1145/3386367.3431314>
- [22] Ramakrishna Padmanabhan, Aaron Schulman, Alberto Dainotti, Dave Levin, and Neil Spring. 2019. How to find correlated Internet failures. In *Proceedings of the Passive and Active Measurement Conference*. Springer, Puerto Varas, Chile. <http://cseweb.ucsd.edu/~schulman/docs/pam19-corrfail.pdf>
- [23] David Plonka and Arthur Berger. 2015. Temporal and Spatial Classification of Active IPv6 Addresses. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Tokyo, Japan, 509–522. <https://doi.org/10.1145/2815675.2815678>
- [24] The Spamhaus Project. 2023. <https://www.spamhaus.org>.
- [25] Lin Quan, John Heidemann, and Yuri Pradkin. 2013. Trinocular: Understanding Internet Reliability Through Adaptive Probing. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Hong Kong, China, 255–266. <https://doi.org/10.1145/2486001.2486017>

- [26] Lin Quan, John Heidemann, and Yuri Pradkin. 2014. When the Internet Sleeps: Correlating Diurnal Networks With External Factors. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Vancouver, BC, Canada, 87–100. <https://doi.org/10.1145/2663716.2663721>
- [27] Philipp Richter, Ramakrishna Padmanabhan, Neil Spring, Arthur Berger, and David Clark. 2018. Advancing the Art of Internet Edge Outage Detection. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Boston, Massachusetts, USA, 350–363. <https://doi.org/10.1145/3278532.3278563>
- [28] Philipp Richter, Georgios Smaragdakis, David Plonka, and Arthur Berger. 2016. Beyond counting: new perspectives on the active IPv4 address space. In *Proceedings of the 2016 Internet Measurement Conference*. 135–149.
- [29] RIPE Atlas. 2023. IPv4 B-root ping. <https://atlas.ripe.net/measurements/1010/>.
- [30] RIPE NCC Staff. 2015. RIPE Atlas: A Global Internet Measurement Network. *The Internet Protocol Journal* 18, 3 (Sept. 2015), 2–26. <http://ipj.dreamhosters.com/wp-content/uploads/2015/10/ipj18.3.pdf>
- [31] Xiao Song, Guillermo Baltra, and John Heidemann. 2023. Inferring Changes in Daily Human Activity from Internet Response. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Montreal, QC, Canada, to appear. <https://doi.org/10.1145/3618257.3624796>
- [32] USC/ISI ANT project. 2017. <https://ant.isi.edu/datasets/all.html>. (Updated quarterly).
- [33] Yinglian Xie, Fang Yu, Kannan Achan, Eliot Gillum, Moises Goldszmidt, and Ted Wobber. 2007. How Dynamic are IP Addresses?. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Kyoto, Japan, 301–312. <https://doi.org/10.1145/1282380.1282415>
- [34] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. Rusty clusters?: dusting an IPv6 research foundation. In *Proceedings of the 22nd ACM Internet Measurement Conference*. ACM, France, 395–409. <https://doi.org/10.1145/3517745.3561440>

A Research Ethics

Our work poses no ethical concerns for several reasons.

First, we collect no additional data, but instead reanalyze data from existing sources. Our work therefore poses no additional risk in 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

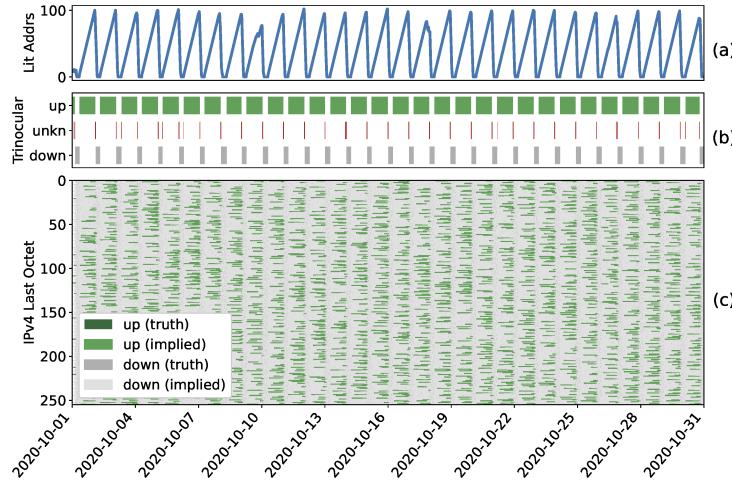


Fig. 9: Diurnal blocks in AS9829 observed from Trinocular, Los Angeles, October 2020. Dataset A42.

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 A Sample Block with Diurnal Behavior

The block in Figure 9 from AS9829 shows how one /24 occupancy varies over the course of a day. Green dots show active addresses, and gray non-response. This address block is 50% full every day at its peak, but empty every night. This trend can be seen in the count of active address (the top graph). It causes daily outage events in Trinocular, as shown in the middle graph, showing *up* most of the day but *down* every night. Blocks that look like this are common in this AS, and they show the need for our IDD algorithm (§3.3).

C Does Unmonitored Space Harm IAS?

Measurement systems do not track the complete address space, as some segments are discarded due to low response rate, as well as addresses that historically have not responded [2]. Users reassigned to unmonitored space implies that IAS may erroneously infer outages due to drops in the total active address count, IAS

FROM / TO	active	non-trackable	inactive		
active	66,892	51%	3,487	3%	5,086
non-trackable	3,392	3%	30,101	22%	1,251
inactive	4,915	4%	1,303	1%	14,602

Table 4: Atlas VP address changes in Trinocular (un)monitored address space.

false negatives. To evaluate if unmonitored space interferes with IAS, we count the number of times known VPs move to and from our underlying measurement system’s unmonitored address space. We expect most of VPs to move within monitored addresses, as unmonitored space has been historically unresponsive implying low usage.

Trinocular strives to probe as much as it can (the *active* addresses), Trinocular excludes addresses for two reasons, *inactive* addresses used to reply to pings but have not in two years, and *non-trackable* blocks have less than three responsive addresses.

As with §4.2, we use RIPE Atlas VPs as ground truth, since they track their current IP addresses. Table 4 counts how many addresses Atlas VPs have in each of the three Trinocular categories (active, inactive, non-trackable).

As expected, the majority of reassessments (51%) occur within monitored addresses (the top, left, green cell). In addition, most addresses (84%) stay in the same category (the diagonal).

A few addresses (7% in the yellow, left column) become active as they move in to measurable space, and about an equal number move out (the 7% in the red, top row). Finally, a surprisingly large 35% are never tracked (the gray region). Since the IAS goal is identify steady or changing addresses, never tracked blocks do not matter. The number that becomes and cease to be active is small (7% each) and about equal in size, so they should not skew IAS. We therefore conclude IAS is not impeded by incomplete measurement.