

A First Look at IPv4 Transfer Markets

Ioana Livadariu, Ahmed Elmokashfi
Simula Research Laboratory
{ioana, ahmed}@simula.no

Amogh Dhamdhere, Kc Claffy
CAIDA / UC San Diego
{amogh, kc}@caida.org

ABSTRACT

In February 2011 the Internet Assigned Numbers Authority (IANA) exhausted its free pool of IPv4 addresses, and the regional registries (RIRs) have started to run out of IPv4 addresses as well. As RIRs have started rationing allocations, *IPv4 transfer markets* have emerged as a new mechanism to acquire IPv4 addresses. Barring a few high-profile exceptions [30], IPv4 transfers have largely flown under the radar. In this work, we use the lists of transfers published by three RIRs to characterize the transfer market — the types of players involved, the sizes and characteristics of transferred address blocks, and the visibility of transferred address blocks in the routing table before and after the transfer. Next, we take first steps toward detecting address transfers using BGP data from the Routeviews and RIPE repositories from 2004-2013. We identify reasons why legitimate changes in prefix origin could be mistakenly inferred to be transfers, and implement a series of 10 filters that remove 86% of candidate transfers. Our results indicate that BGP-based detection of transfers is prone to false positives due to significant noise in BGP data, while some transfers remain undetectable as they involve non-BGP speakers. We describe some additional data sources and analysis techniques that may help reveal an opaque market for IPv4 address block transfers.

Categories and Subject Descriptors

C.2.2 [COMPUTER-COMMUNICATION NETWORKS]: Network Protocols—*Routing Protocols*

1. INTRODUCTION

On February 3, 2011 the Internet Assigned Numbers Authority (IANA) allocated the remaining five /8 blocks to the Regional Internet Registries (RIRs) [21]. APNIC started allocating addresses from its last /8 in April 2011 [5]; RIPE followed suit in September 2012 [35]. The other RIRs are expected to run out within a few years [21]. RIRs have drafted stricter allocation policies as a response to this impending address exhaustion [8, 32]. Meanwhile, the replacement protocol IPv6 has been lagging in deployment, even if recent signs are positive [19]. The slow transition to IPv6, and the possibility of a dual-stacked world for the foreseeable

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT'13, December 9–12, 2013, Santa Barbara, California, USA.
Copyright 2013 ACM 978-1-4503-2101-3/13/12 ...\$15.00.
<http://dx.doi.org/10.1145/2535372.2535416>.

future, means that there is a continual need for IPv4 addresses. Recognizing this need, the RIRs have allowed intra-registry transfers of IPv4 addresses via a market mechanism, starting with RIPE NCC in December 2008 [28, 33] and followed by ARIN [4] and APNIC [7] in mid-2009 and early 2010, respectively. Inter-registry transfers are currently only authorized between ARIN and APNIC [11]. Recently, ARIN, RIPE NCC and APNIC have made lists of completed transfers available to the public in an attempt to provide more transparency into the address transfer process.

IPv4 address transfer markets are a subject of heated debate: they can extend the usable life of IPv4, but they could also delay the adoption of IPv6 or halt it altogether [16], cause further fragmentation of the address space and larger IPv4 routing tables, or generate destabilizing speculation and/or hoarding behavior. It is not clear that address space owners, especially holders of legacy space, will adhere to RIR transfer policies; even now address blocks may be changing hands without the knowledge of the RIRs. We believe an empirically grounded characterization of address transfer activity will inform the on-going debate on the relative benefits and harms of IPv4 address space markets.

We use RIR-published lists of address transfers to characterize the observable transfer market – the types of entities buying and selling addresses, sizes of transferred address blocks, and what happens to prefixes after they are transferred. We find that 75% of transferred addresses come from legacy allocations, 85% of transferred address blocks appear in the routing table after the transfer, and transferred prefixes are generally lightly utilized. These results suggest that thus far markets are facilitating a healthy redistribution of IPv4 address space. Second, we experiment with a method to detect address transfers using BGP routing data. A major challenge is that prefixes may change origin ASes for reasons other than transfers, e.g., movements internally within an organization, transient prefix hijacks, and traffic engineering. We devise a set of 10 *filters* to remove apparent transfers that may be due to such reasons. Our approach is useful in reducing the number of candidate transfers from an initial count of several thousands per month to a few hundreds. Even after this filtering, however, 99.5% of the apparent transfers do not appear in the lists published by RIRs. It is not possible to confirm or eliminate these as real transfers without soliciting validation from the involved parties. There are undoubtedly false positives due to the noisiness of BGP data, and our results indicate that it is necessary to augment BGP-based detection with other measurement techniques that rely on the data-plane.

2. DATASETS

We briefly describe the data we use in our analysis of published transfers (Section 3) and in our method for inferring transfers from BGP routing data (Section 4).

Lists of transferred address blocks: Three registries have published lists of IP address transfers in their region: ARIN [10], APNIC [6], and RIPE NCC [34]. ARIN publishes this information dating back to October 2009, and APNIC from November 2010 onwards. RIPE NCC has started publishing transfers recently, starting October 2012. The ARIN list also specifies inter-RIR transfers between ARIN and APNIC.

BGP data: We collected historical BGP data from the two major public repositories at RouteViews [3] and RIPE [36]. Our data spans 9 years from 2004 to 2012. We use it to build prefix-AS maps for use in transfer inference (Section 4).

AS relationships: We use CAIDA’s AS relationship classification algorithm [27] to infer business relationships between ASes. The AS-rank algorithm classifies AS links into two types, customer-provider or settlement-free peer.

AS classification: We use an AS classification scheme from our previous work [18] to classify ASes into four categories: Enterprise Customer (EC), Small Transit Provider (STP), Large Transit Provider (LTP), and Content/Access/Hosting Provider (CAHP). Our previous work provides details regarding the validation and robustness of our classification, which achieves an accuracy of 85% [18].

WHOIS data: We use delegation files published by RIRs [9] to find the allocation date associated with given address blocks. We use CAIDA’s AS-to-organization dataset [23] to group together ASes belonging to the same organization. This dataset is built from WHOIS data, and extends the method proposed in previous work on AS-to-organization mapping [14].

IP census data: In Section 3 we study address space utilization by networks that either buy or sell IPv4 addresses. To measure the utilization of IP address space, we use data from ISI’s July 2012 Internet Census [22], which consists of a list of probed and responding IP addresses from a scan of the entire IPv4 Internet.

3. OBSERVABLE TRANSFER MARKET

In this section, we use the lists of transferred address blocks published by ARIN, RIPE NCC and APNIC to characterize the publicly documented transfer market. Our goal is to look into the types of organizations that are buying and selling address space, the nature of address blocks being traded, and the visibility of transferred address blocks in the routing table.

General statistics of market transfers

ARIN reports 111 completed transfers from October 2009 to March 2013. 26 of these transfers involve sub-prefixes of 198.32.0.0/16, a block reserved by ARIN for Internet Exchange points (IXPs). These transfers turn out to be triggered by the re-numbering of the DNS "L" root server [17]. RIPE NCC reports 31 transfers from October 2012 to March 2013, and APNIC reports 277 transfers between November 2010 and March 2013. APNIC’s list includes non-market transfers, e.g., those due to merger and acquisition (M&A). We manually inspect and remove 45 transfers that either occurred within the same organization, or due to M&A activities. This leaves 232 transfers, which may still include non-market transfers, i.e., we consider it an upper bound on the number of market transfers in the APNIC region. Eight authorized transfers from the ARIN to APNIC regions are reported between October 2012 and January 2013 [10].

Figure 1 shows the number of transfers per month reported by each RIR. ARIN transfers peak on two occasions, in December 2009 and August 2011. The first peak corresponds to 13 transfers from the IXP block mentioned above. Since we are interested in market transfers, we ignore these transfers, leaving 85 market




Figure 1: The number of transfers reported by each RIR over time. The number of transfers per month in the APNIC region shows an increasing trend.

Table 1: Transferred blocks classification

Class	Before	After
A-I	Never routed	unrouted
A-II	Unrouted for the last two years (routed before that)	unrouted
B	routed	unrouted
C-I	Never routed	routed
C-II	unrouted for the last two years (routed before that)	routed
D	routed	routed

transfers. The second peak represents Microsoft’s purchase of 38 address blocks from Nortel [30]. Over half (45) of the transferred blocks in ARIN used to be owned by Nortel. While there are only on average 32 transfers per year in the ARIN region, APNIC and RIPE had 83 and 62 reported transfers, respectively. More transfer activity in the APNIC and RIPE regions is consistent with address scarcity in those regions, which are already allocating addresses from their last /8, unlike ARIN [21].

Counting the number of transferred addresses, rather than the number of transactions, reveals a different picture. The number of transferred addresses is higher in the ARIN region than APNIC or RIPE, mostly due to the Nortel transfer. In terms of size of transferred blocks, ARIN exhibits larger transfers, with its first quartile at a /16; the same quartile is /19.5 and /20 for the RIPE and APNIC regions, respectively. The median size of a transferred prefix is /21 in the ARIN and RIPE regions, while it is /23 for the APNIC regions. These differences are also consistent with relatively more address scarcity in the APNIC and RIPE regions.

Where are addresses going?

Understanding the geographies of the transfer market gives insight into its spread and into the role played by brokerages such as the IPv4 Market Group [2] in bringing together sellers and buyers. Most of the transferred addresses move between organizations that are registered (according to WHOIS data) in different countries. 53% and 66% of all transferred addresses fall into this category in APNIC and RIPE, respectively.¹ Buyers come from 13 countries in APNIC and 12 in RIPE, but the numbers are dominated by a few countries. Buyers from Japan and Hong Kong dominate the APNIC transfers (60% of addresses), while in the RIPE region buyers from Great Britain and France dominate (52% of addresses). The prevalence of inter-country transfers suggests that brokers may play a valuable role in bringing together buyers and sellers.

Are transferred blocks used immediately?

While the lists of transfers published by RIRs show which address blocks are being transferred, they do not say whether the transferred

¹We do not measure these statistics for transfers within the ARIN region, since all buyers and sellers are from the U.S. or Canada.

Table 2: Transferred address blocks from each category in each RIR and overall. The numbers in parentheses are percentages of all prefixes from that RIR.

RIR	A-I	A-II	B	C-I	C-II	D
APNIC	12 (5)	5 (2)	20 (9)	36 (16)	35 (15)	124 (53)
ARIN	4 (5)	4 (5)	5 (6)	26 (31)	10 (11)	36 (42)
RIPE	0 (0)	0 (0)	3 (10)	2 (6)	7 (23)	19 (61)
ARIN-APNIC	1 (12)	0 (0)	2 (25)	2 (25)	0 (0)	3 (38)
Total	17 (5)	9 (2)	30 (8)	66 (19)	52 (15)	182 (51)

prefixes appear in the routing table. To understand the life-cycle of transferred blocks, we divide them into six classes (described in Table 1) based on their presence in the routing table before and after the transfer. To determine if a prefix was ever advertised in the past, we use historical routing tables collected from the Routeviews and RIPE NCC repositories from 1998 onwards. Table 2 shows the number of prefixes in each class for each RIR. Most transferred blocks (85% of transferred prefixes from all RIRs combined) are currently advertised, i.e., classes C-I, C-II, and D. 34% of all transferred blocks were either advertised for the first time after their transfer, or reappeared following an absence of at least two years, i.e., Classes C-I and C-II. This observation indicates that address space is transferred to organizations that actually use it. The percentage of transferred prefixes that are currently unrouted (classes A-I, A-II, and B) is 15%. Investigating these prefixes shows that the bulk of them are either allocated to national registries such as JPNIC, their transfers took place recently in the past six months, or they are part of a buyout of prefixes of which some are already advertised. For example, 8 of the currently unrouted prefixes in ARIN were bought along with other blocks from Nortel by Microsoft, Salesforce, and Vodafone.

A plausible explanation for the unrouted blocks is that networks are not ready to deploy acquired addresses immediately. We measured the latency between address acquisition and advertisement for all registries, and find that it takes on average 1.1, 2.4, and 6.7 months to re-advertise an acquired prefix in RIPE, APNIC, and ARIN respectively. The difference in latency indicates that ASes in the RIPE and APNIC regions are buying addresses to meet more immediate needs than in the ARIN region. A more pessimistic interpretation of unrouted blocks is that some organizations are acquiring addresses to *hoard* them, anticipating an increase in their value. Only 15% of transferred prefixes remain unadvertised so far, so such pessimism may be premature. It is, however, a strong motivation for tracking advertisement latency of transferred addresses.

Do buyers need the addresses more than sellers?

We used the IP census data to measure the utilization of transferred prefixes before they were transferred, i.e., the fraction of a prefix that the probing was able to reach.² Since we have only a single snapshot of ISI’s IP census data from July 2012 [22], we examine only subsequent transfers, and find that utilization of transferred prefixes (before their transfer) was significantly lower than other blocks. On average, the utilization of transferred blocks was 0.8%, 3.9% and 2.9% for ARIN, APNIC, and RIPE, respectively, compared to an average of 9.4% for all other routed prefixes.

We also found that buyer ASes utilize their *other* (i.e., non-transferred) address space more than sellers: the median utilization of address space already owned by buyers was about 5%, 8%, and

19% in the ARIN, APNIC, and RIPE regions, respectively; corresponding utilization of addresses already owned by sellers was 0.9%, 2.5%, and 5.3%, respectively. These differences imply that ASes are offering underutilized space for transfer, while buyers seem driven by an actual need for more addresses.

What kind of ASes are involved in transfers?

We next investigate the types of ASes that are involved in transfers, using our AS classification listed in Section 2. We note that the scheme classifies ASNs, and does not include organizations that do not have AS numbers. Such *non-speakers* also appear as buyers or sellers in the RIR transfer lists. We find that acquisitions and sales are dominated by networks at the edge of the AS ecosystem, namely CAHPs, ECs, and non-BGP speakers. All purchases in RIPE, 75% of transfers in APNIC, and 87% of the transfers in ARIN are by such edge networks. We also investigate the allocation dates of the AS buyers, to develop a better understanding of their motives. The number of BGP-speaking buyers (those with AS numbers) is 63 (APNIC), 33 (ARIN), and 21 (RIPE). The number of non-BGP speaking buyers (organizations with no AS numbers who have their addresses advertised by their providers) is 10 and 2 in APNIC and ARIN, respectively. Of the BGP-speaking buyers, 13 ASes in APNIC, 7 in ARIN, and 10 in RIPE were allocated after January 2011, indicating that many buyers are newly joining edge ASes. Interestingly, a previous study of IPv6 deployment found that edge ASes are lagging in IPv6 deployment [19].

Are mostly legacy prefixes transferred?

In the years predating RIRs, IANA directly allocated addresses to requesting networks. Such *legacy allocations* account for about 40% of the total IPv4 address space [24]. The allocation of legacy addresses was not tightly coupled with need, e.g., entire /8 blocks were allocated to universities that required (and utilized) much less space. We find that legacy addresses account for 75% of all transferred IP addresses, disproportionately higher than their contribution to the entire IPv4 address space (40%). This high fraction of transferred legacy allocations suggests a healthy redistribution process. Mueller et al. [29] also found a similarly large number of legacy allocations involved in transfers published by ARIN and APNIC between 2009 and 2012.

How much money has been transferred so far?

According to the list of published transfers, 11515904 IP addresses have changed hands so far. The Nortel and Borders sales in 2011 were reported to be at the rate of \$11.25 and \$12 per IP address, respectively. Prices for other transactions have not been published, but assuming \$11.25 per address, roughly \$130 million has been exchanged for addresses so far. Mueller et al. [29] estimated that \$60 million had changed hands based on ARIN and APNIC transfers from 2009 to mid-2012. Our larger estimate covers transfers until March 2013 and also includes transfers in the RIPE region. Mueller et al. also discarded transfers that involved national registries, e.g., JPNIC in the APNIC region, which we do not. As mentioned earlier, the APNIC transfer lists possibly include non-market transfers. Excluding APNIC transfers reduces our estimate to \$91 million.

Do markets affect IPv6 adoption?

To explore the relationship between IPv4 transfer markets and IPv6 adoption, we check whether ASes that are receiving IPv4 address transfers (as recorded by the RIRs) are also originating IPv6 prefixes in publicly available BGP data. We found 33 recipient ASes in the ARIN region, of which 18 (55%) originated IPv6 prefixes in April 2013. In the RIPE region, 10 out of 21 recipient ASes originated IPv6 prefixes, while the corresponding number for APNIC was 33 out of 63 (52%). Previous work found that the fraction of

²We recognize the problem with defining *utilization* as “ICMP responding” given unknown filtering behavior, but we will use this terminology in this section.

ASes from the IPv4 graph that are also present in IPv6 was less than 20% for each RIR [19]. Nonetheless, the fact that about half of the ASes receiving IPv4 transfers have not deployed IPv6 does indicate that some organizations view the IPv4 transfer market as mechanism to avoid deploying IPv6 immediately.

4. INFERRING TRANSFERS IN BGP DATA

In this section we propose and evaluate a technique to automatically infer transfers of IP blocks using only information from routing tables. We infer *candidate transfers* by observing changes in the origin AS of prefixes seen in BGP routing tables over time, assuming the origin AS of a prefix is its “owner”. We build our prefix-AS mappings from BGP data, assigning the origin AS of a prefix as its “owner”. We construct one prefix-AS mapping on each of the first 7 days of every month, combining the largest routing table from the RIPE collector rrc04 with the largest from the Routeviews collector routeviews2 on that day. This gives us 7 prefix-AS mappings for that month. We then use *majority filtering* to retain only those prefix-AS mappings that were seen in at least 4 of the 7 daily snapshots, and seen from monitors that observed more than 90% of the routed prefixes. This process yields a single prefix-AS mapping for each month, and a total of 1014215 mappings over nine years. Majority filtering ensures that our prefix-AS mapping in each snapshot is robust to transient events in the routing table.

We consider a *full transfer* to have occurred when we observe the origin AS change for an entire address block at a single point in time. In a *partial transfer*, we observe a change in the origin AS for a part of an advertised address block. Consider two prefixes p_a and p_b , where the address space of p_a completely overlaps the address space of p_b . Assume that p_a is advertised by AS A from time t_1 until t_2 , and p_b by AS B from t_3 until t_4 , where the time periods do not overlap. In this case, we identify a partial transfer of p_b at time t_3 from seller A to buyer B. The reverse situation could also occur, where the buyer acquires multiple smaller address blocks and starts advertising the total acquired space.

This basic approach can mis-infer transfer activity based on origin AS changes due to entirely different reasons, e.g., traffic engineering, mergers and acquisitions, prefix hijacks. To remove such non-transfers from the list of candidate transfers, we implement the following set of filters, and apply them to candidate transfers in the following order.

Transient filter: BGP misconfigurations or prefix hijacks can cause prefixes to temporarily change origin. To disqualify such cases as candidate transfers, we consider only prefixes that have been advertised by the same AS for at least four months before and after being transferred.

RIR filter: We filter out candidate transfers in which one of the involved ASes is assigned to an Internet Registry, according to WHOIS data.

Delegation filter: Provider aggregatable (PA) space refers to address space that is assigned by a service provider to its customers in such a way that routing information for many customers can be aggregated. In contrast, Provider Independent (PI) address space is assigned to end-users. To avoid false positives due to the delegation of PA address space by transit providers to their customers, we use the method proposed by Cittadini et al. [15] to classify prefixes into four categories: top, lonely, deaggregated and delegated. A prefix p_t is considered to be a top prefix if other prefixes exist in the routing table for which p_t is a covering prefix. A prefix p_l is classified as a lonely prefix if no covering or covered prefixes appear in the routing table for p_l . A prefix p_{da} is a deaggregated prefix if it is covered by another prefix p'_t and both prefixes are advertised by the same AS. A prefix p_{de} is a delegated prefix if it is covered by

Table 3: Filtering the initial set of 173,105 candidate transfers. The filters are listed in the order in which they were applied.

Filter	Removed transfers	%
Transient	43,851	25.33%
RIR	217	0.12%
Delegation	25664	14.82%
Map-to-org	22284	12.87%
Set	47	0.02%
AS Death	5712	3.29%
AS Born	18217	10.52%
Allocation	441	0.25%
Transit Provider	26418	15.26%
Customer-Provider	2420	1.39%

another prefix p'_t but the two prefixes are advertised by two different ASes. The delegation filter removes a candidate transfer if the prefix is classified as delegated.

Map-to-Organization filter: If a candidate transfer involves two ASes belonging to the same organization, then the movement could simply reflect a change of internal policy rather than an actual transfer. If a prefix moves between two ASes belonging to the same organization (according to our AS-to-organization map [23]), then we remove this prefix from the candidate transfers.

Set filter: The origin of a prefix can be an AS set, indicating that the origin is an unordered sequence of aggregated ASes. For prefixes originated by AS sets that appear in our snapshots, we filter candidate transfers where the prefixes moves between two ASes in the same set.

AS Death, AS Born filters: The AS Death filter is designed to remove non-market transfers due to mergers and acquisitions. We consider an AS to be involved in a merger or acquisition if it sells parts of its address space and disappears from the routing table at some time after 3 months from the date of the transfer. The AS Born filter is designed to remove candidate transfers where organizations that previously held Provider Aggregatable (PA) space acquire AS numbers and begin to advertise the space themselves. We filter candidate transfers in which the recipient ASes appeared in the routing table for less than twelve months before the transfer. **Allocation filter:** We use delegation files provided by the five RIRs [9] to remove prefixes that appear in BGP routing tables before they are allocated (“bogon” prefixes).

Transit Provider filter: An organization that owns a prefix p (but no AS number) may switch upstream providers from A to B, which appears as a transfer of prefix p from A to B. Transit ISPs are unlikely to transfer IP addresses since they are a critical input to the transit business, so this filter removes transfers for which we have classified the seller as a (small or large) transit provider.

Customer-Provider filter: Provider Aggregatable (PA) address space that is sub-allocated by providers to their customers or returned by customers to providers could appear in our candidate list of transfers. The customer-provider filter uses AS-relationship files to filter the candidate transfer of prefix p from AS A to AS B if the following conditions are satisfied: A is the provider of B at time t , and the provider continues to advertise blocks adjacent to p after t .

5. EVALUATION RESULTS

Table 3 shows the number of candidate transfers removed by each filter when applied to our initial list of 173,105 candidate transfers. The transient, delegation, map-to-Organization filters removed most (53%) candidate transfers. After applying all ten filters, the final set of candidate transfers is almost an order of magni-




Figure 2: Inferred initial and final transfers. Our filtering methodology reduces the set of candidate transfers by an order of magnitude.

tude smaller. Figure 2 plots the number of initial candidate transfers and the number of transfers inferred after filtering. Our detection stops in December 2012 because of the four-month limit required by our transient filter.

We evaluate the BGP-based detection method by comparing the inferred transfers with the ground truth transfer transactions published by RIRs. Our method depends on prefix announcements, and thus can only detect transfers involving prefixes that were routed before and after the transfer, i.e., class D prefixes as defined in Section 3. Table 4 presents the outcome of this evaluation. Some transfers are not detectable using our method, because the involved prefixes do not change origins, e.g., they move between non-BGP customers of the same AS. Our method detects most detectable transfers, except those in the APNIC region, where non-market transfers muddle the data. Such non-market transfers are due to mergers and acquisitions that do not result in any change in the origin AS of the prefix, and are thus false negatives with our methodology. Four of the 7 false negatives in ARIN are filtered by our "Transit Provider Filter". One of them was owned by Norand Co., a non-BGP speaker, and advertised by AT&T, which appears as its owner. The remaining three are filtered by the *AS Born Filter*. The false negative in RIPE is caused by the seller continuing to advertise (intentionally or erroneously) parts of the transferred block after the transfer, and is thus not a false negative in the true sense.

We have presented our experiences with detecting transfers using publicly available BGP data. While the BGP-based method does detect transfers, 99.5% of the candidate transfers it infers are not in the lists published by RIRs. Also, some published transfers are undetectable due to the presence of non-BGP speakers. IP address blocks owned by non-BGP speaking organizations are advertised by the organization's upstream providers. When a non-BGP speaker switches providers, its address space appears in the BGP data as transferred between the involved providers – these are false positives of the BGP-based method. For detecting candidate transfers that occur on the black or grey market, false positives are tolerable, because the goal is to reduce the huge number of apparent transfers in BGP data to just the suspicious cases. We can then analyze the suspicious cases in detail if they are small in number. False negatives and undetectables are more problematic, as they mean that some transfers may pass undetected unless they are reported to the RIRs. We are planning to extend our method to reduce the number of false positives and undetectables caused by PA space. One approach we are considering is the limited use of WHOIS data for identifying PA space and non-BGP speaking organizations. While additional filters and the use of some WHOIS data may improve the detection accuracy, our results indicate that we may need to look at

Table 4: Comparison of inferred candidate transfers with transfer lists published by RIRs.

RIR	Undetectable	False negatives	Detected
APNIC	56	23	25
ARIN	8	7	17
RIPE NCC	4	1	3

other data sources to augment BGP-based detection. We briefly describe some possible approaches in Section 7.

6. RELATED WORK

Closely related to our work is that of Mueller et al. [29], who studied the IPv4 transfer market from 2009 to mid-2012 using WHOIS data, transfer lists published by ARIN and APNIC, and bankruptcy court orders reporting completed transfers. They found that 89% of transferred addresses were from legacy allocations; we find that 75% of transferred addresses in our extended dataset come from legacy allocations. In addition, we have added an analysis of the transfer listings from RIPE NCC, which did not publish transfers until October 2012. Further, Mueller et al. [29] did not look into whether transferred prefixes show up in the routing table, the time lag between transfer and advertisement, the utilization of transferred prefixes and ASes involved in transfers, and the relation to IPv6 adoption.

Several research publications, press articles and presentations at research and operational venues have debated the pros and cons of IPv4 address transfer [12, 13, 20, 25, 26], a debate we briefly summarized in Section 1. Osterweil et al. [31] argued that RIRs should deregulate the sale of IP addresses altogether, and implement resource certification as proof of ownership of IP blocks. We are not aware of other empirical analysis of transfers that have already occurred, or methods to detect address transfers that may be occurring without the knowledge of RIRs.

7. DISCUSSION AND FUTURE WORK

We analyzed address transfers that have already been completed (and documented by RIRs) finding that 75% of transferred addresses involve legacy allocations, involve underutilized prefixes, and are generally advertised via BGP into the global routing table shortly after the transfers take place. These indicators suggest that address markets, at least what we can see of them, are serving their intended purpose of facilitating a redistribution of underutilized addresses to entities that need them. Next, we inferred a number of *apparent transfers* using BGP data, and developed some techniques to filter out false candidate transfers based on other AS behavior that might share BGP-observable characteristics of transfer behavior. While some of these inferred transfers correspond to transfers published by RIRs, we infer many more that are not in RIR transfer listings. Some of these are surely false positives in our methodology, but others could be the result of transfers happening on an underground market, bypassing the RIRs. There is no way to confirm these one way or the other without soliciting ground truth from the parties involved in these transfers.

The number of transferred prefixes (reported in the lists published by RIRs) is still small (a total of 382 transferred prefixes as compared to more than 400K routed prefixes), and we have discerned no significant effects on global routing tables, such as increased fragmentation of address blocks or unusual growth in routing table sizes. But we have observed an increasing trend in the number of transfers published by RIRs (2, 109, and 184 in 2010, 2011, and 2012, respectively), indicating that the popular-

ity of these markets is increasing. This popularity will be amplified by the fact that address transfer markets have now been legitimized by RIR community policies; these communities are now formulating policies for inter-RIR address transfers, and even discussing policies for legitimizing transfers that are not needs-based.

There is also speculation that availability of IP addresses on the market will inflate their value, as organizations with underutilized address space stand to profit significantly from selling it. Organizations could also circumvent RIRs and purchase addresses on the black market, leading to less transparency into the ownership of IP numbering resources. Although we believe that the stability of address markets is best ensured with transparent disclosure of transactions, such disclosure is not guaranteed in the current RIR policy architecture. It is thus worth pursuing inference methods that can improve transparency in address transfer markets without formal registration requirements, ideally using publicly available data so others can verify inferences.

We plan to refine our methodology, publish our list of inferred transfers, and solicit ground truth from parties involved in these inferred transfers so we can improve the accuracy of our inference methodology and shed additional light on the IPv4 address transfer market, including its impact on the size and stability of the BGP routing system. Our experience with BGP-based detection methods leads us to believe that other approaches such as DNS and data plane measurements may be needed to discern the signal from the noise. We have experimented with using reverse DNS mappings (and looking for differences in these mappings before and after the apparent transfer of a prefix) to identify transferred prefixes, particularly transfers that involve non-BGP speakers. DNS mappings are unlikely to change when a non-BGP speaker switches providers, while they should change for legitimate transfers. We have also experimented with using data plane measurements of RTTs from fixed vantage points (CAIDA's Archipelago infrastructure [1]) toward IP addresses in prefixes we suspect as being transferred. Both these techniques showed promise in reducing false positives and undetectables due to non-BGP speakers. They are the focus of our current work. It is important to note that while BGP data is available historically from the Routeviews and RIPE RIS repositories, this is not the case for data plane measurements. We believe it is worth developing and deploying these measurements on an ongoing basis, to assist in detecting grey-market transfers in the future.

Acknowledgments

We thank our shepherd, Christian Kreibich, and the anonymous reviewers for their constructive comments. Amogh Dhamdhere and Kc Claffy were supported by the National Science Foundation (NSF) under grant CNS-1111449. Ahmed Elmokashfi and Ioana Livadariu are supported by the Norwegian Research Council, grant number 209954/S10.

8. REFERENCES

- [1] Archipelago Measurement Infrastructure. <http://www.caida.org/projects/ark/>.
- [2] IPv4 Market Group. <http://ipv4marketgroup.com/home/>.
- [3] University of Oregon Route Views Project. <http://www.routeviews.org/>.
- [4] ARIN Number Resource Policy Manual, 2013. <https://www.arin.net/policy/nrpm.html>.
- [5] APNIC. APNIC IPv4 Address Pool Reaches Final /8. <http://www.apnic.net/publications/news/2011/final-8>.
- [6] APNIC. APNIC IPv4 Transfers. <http://ftp.apnic.net/transfers/apnic/>.
- [7] APNIC. APNIC Transfer, Merger, Acquisition, and Takeover Policy. <http://www.apnic.net/policy/transfer-policy>.
- [8] APNIC. Policies for IPv4 Address Space Management in the Asia Pacific Region. <http://www.apnic.net/policy/add-manage-policy#9.10>.
- [9] ARIN. Address Allocation and Assignment Reports. <ftp://ftp.arin.net/pub/stats/>.
- [10] ARIN. Inter-RIR and Specified Transfers of Internet Number Resources. <https://www.arin.net/knowledge/statistics/transfers.html>.
- [11] ARIN. Inter-RIR transfers. https://www.arin.net/resources/request/transfers_8_4.html.
- [12] I. V. Beijnum. Trading IPv4 Addresses Will End in Tears. *Ars Technica*, 2011.
- [13] I. V. Beijnum. IPv4 Address Transfer Markets are Forming where we Least Expected. *Ars Technica*, 2012.
- [14] X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger. Towards an AS-to-Organization Map. In *Proc. ACM SIGCOMM IMC*, 2010.
- [15] L. Cittadini, W. Mühlbauer, S. Uhlig, R. Bush, P. Francois, and O. Maennel. Evolution of Internet Address Space Deaggregation: Myths and Reality. *IEEE JSAC*, Aug 2010.
- [16] K. Claffy. Exhausted IPv4 Address Architectures. CAIDA Blog, 2010.
- [17] K. Davies. IANA Update for CCTLD Registries, 2008.
- [18] A. Dhamdhere and C. Dovrolis. Twelve Years in the Evolution of the Internet Ecosystem. *IEEE/ACM Transactions on Networking*, 19(5), 2011.
- [19] A. Dhamdhere, M. Luckie, B. Huffaker, A. Elmokashfi, K. Claffy, and E. Aben. Measuring the Deployment of IPv6: Topology, Routing, and Performance. In *Proc. ACM SIGCOMM IMC*, 2012.
- [20] A. Dul. Economics of IPv4 markets on IPv6 deployment. NANOG 53, 2011.
- [21] G. Huston. IPv4 Address Report, 2009. <http://www.potaroo.net/tools/ipv4/index.html>.
- [22] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister. Census and Survey of the Visible Internet. In *Proc. ACM SIGCOMM IMC*, 2008.
- [23] B. Huffaker, K. Keys, M. Fomenkov, and K. Claffy. AS-to-Organization Dataset. <http://www.caida.org/research/topology/as2org>.
- [24] IANA. IANA IPv4 Address Space Registry. <http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml>.
- [25] T. B. Lee. The Case for a Free Market in IPv4 Addresses. *Ars Technica*, 2011.
- [26] W. Lehr, T. Vest, and E. Lear. Running on empty: the challenge of managing Internet addresses. In *Proc. TPRC*, 2008.
- [27] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and K. Claffy. AS Relationships, Customer Cones and Validation. In *Proc. ACM SIGCOMM IMC*, 2013.
- [28] M. Mueller. Europe Pioneers IPv4 Address Transfer Markets. Internet Governance Project Blog, 2008.
- [29] M. Mueller, B. Kuerbis, and H. Asghari. Dimensioning the Elephant: An Empirical Analysis of the IPv4 Number Market. In *Proc. TPRC*, 2012.
- [30] Network World. Microsoft Pays Nortel \$7.5 million for IPv4 Addresses, 2011.
- [31] E. Osterweil, S. Amante, D. Massey, and D. McPherson. The Great IPv4 Land Grab: Resource Certification for the IPv4 Grey Market. In *Proc. HotNets-X*, 2011.
- [32] RIPE NCC. Allocations from the Last /8. <http://www.ripe.net/ripe/policies/proposals/2010-02>.
- [33] RIPE NCC. IP Address Allocation and Assignment Policies for the RIPE NCC Service Region. <http://www.ripe.net/ripe/docs/ripe-553>.
- [34] RIPE NCC. IPv4 Transfer Statistics. <http://www.ripe.net/lir-services/resource-management/ipv4-transfers/table-of-transfers>.
- [35] RIPE NCC. RIPE NCC Begins to Allocate IPv4 Address Space From the Last /8. <http://www.ripe.net/internet-coordination/news/announcements/ripe-ncc-begins-to-allocate-ipv4-address-space-from-the-last-8>.
- [36] RIPE NCC. Routing Information Service (RIS), 2008. <http://www.ripe.net/ris/>.