Behind IP Prefix Overlaps in the BGP Routing Table

Quentin Jacquemart[1], Guillaume Urvoy-Keller[1], and Ernst Biersack[2]

[1] Univ. Nice Sophia Antipolis, CNRS, I3S, UMR 7271, 06900 Sophia Antipolis
(firstname.lastname@unice.fr)
[2] Swimming pool and relaxation area
erbi@e-biersack.eu

Abstract. The IP space has been divided and assigned as a set of IP prefixes. Due to the longest prefix match forwarding rule, a single assigned IP prefix can be further divided into multiple distinct IP spaces; resulting in a BGP routing table that contains over half a million distinct, but overlapping entries. Another side-effect of this forwarding rule is that any anomalous announcement can result in a denial of service for the prefix owner. It is thus essential to describe and clarify the use of these overlapping prefixes. In order to do this, we use Internet Routing Registries (IRR) databases as semantic data to group IP prefixes into *families of prefixes* that are owned by the same organization. We use BGP data in order to populate these families with prefixes that are announced on the Internet. We introduce several metrics which enable us to study how these families behave. With these metrics, we detail how organisations prefer to subdivide their IP space, underlining global trends in IP space management. We show that there is a large amount of information in the IRR that appears to be actively maintained by a number of ISPs.

1 Introduction

The IP space has been divided into a set of IP prefixes that are assigned to organizations by RIRs (Regional Internet Registries). These organizations can choose to further divide the IP prefixes they were assigned in smaller IP spaces that they can use as independent networks. This is possible because packets are routed according to the longest prefix match rule. In other words, any traffic will always be forwarded to the smallest IP space (i.e. the most specific prefix) containing the destination IP address. This can be useful in order to do traffic engineering, e.g. to make sure off-site servers are reachable from the global Internet. At the same time, the recent attack against Spamhaus demonstrated that announcing more specific prefixes is an effective DoS (Denial of Service) attack [12]. Even in the case of misconfigurations, large-scale repercussions can be disastrous [9].

The BGP (Border Gateway Protocol) routing table currently contains over half a million entries. With so many entries, it is improbable that there is no overlap among them. As a result, it is essential to describe and understand the uses of overlapping prefixes from a BGP point of view. A way of doing this would be to create pairs of overlapping prefixes, and then compare them together. As an example, let us consider the three overlapping prefixes a/8, a.b/16, and a.b.c/24. How relevant is the study of the three pairs of these prefixes? If the organization to which the /8 has been assigned is

an ISP, the /16 prefix might have been sold to one of its customer; and it is solely this customer who decided to create the /24 subnet. Hence, the comparison between the /8 and the /24 is not meaningful. Conversely, if the organization behind the /8 prefix is not an ISP, the /16 prefix is likely not a sub-allocation, but the result of network engineering.

Therefore, by trying to simply compare pairs of overlapping prefixes, we overlook assignment policies. Namely, IP blocks are assigned by RIRs to organizations. These organizations then use their IP space as they see fit. An ISP, for example, will most likely sell a part of its IP space to customers, who, in turn, will use the (sub) IP space as they see fit. As a result, simply comparing any pair composed of any overlapping prefix disregards the fact that different entities may administer the prefixes. In order to overcome this problem, we use the prefix assignment information included in IRR (Internet Routing Registry) databases in order to cluster overlapping BGP prefixes into *families of prefixes*. Prefixes inside these families are then guaranteed to be under the control of a single organization. Consequently, their comparison can be done without ambiguity.

In this paper, we present a method to group BGP prefixes into a set of prefix families with the help of the contents of the IRR databases. These families are composed of two types of prefixes: children prefixes, which are BGP announcements that are *not* included as-is in the IRR databases; and family fathers, which *are* included in the IRR databases. We define a set of metrics to analyse the behaviour of these families that shed light on how an organization sub-divides its own IP space into smaller networks for its own use. We look at a few real-world examples of families, and show that the behaviour inside groups of families of tier-1 ISPs, tier-3 ISPs, and private corporations, is comparable. At the same time, we investigate the distributions of prefixes inside BGP and inside the IRR databases, and offer possible reasons for their large size difference.

2 Data Sources

2.1 IRR Databases

We were able to secure access to the IRR databases of the five RIRs: AfriNIC, ARIN, APNIC, LACNIC, and RIPE. These databases contain information directly provided by network operators, on a voluntary basis, about their routing policies and announcements. They are composed of different objects that represent, among other things, people, IP address allocation, and AS numbers. We extract information from the `inetnum` objects, which contain "details of an allocation or assignment of IPv4 address space" [1].

Table 1 details the number of entries that we extracted from each database on August 1st, 2014. Because some RIRs include information about special-use IP space (e.g. the private IP space) for user friendliness, we discard 0.07% of overall entries. Finally, we obtain 8,364,909 distinct IP prefixes from the IRR databases.

The accuracy of IRR databases is widely debated among the community. For example, [10] underlines the inconsistencies among the distributed content of the database, as well as the varying level of accuracy depending on the considered database and object. However, by comparing the origin AS inside the IRR with the one in BGP, [7] shows

RIR	Parsed prefixes
AfriNIC	72,516
APNIC	1,432,154
ARIN	2,696,539
LACNIC	322,828
RIPE	3,846,706
Total	8,370,743
Filtered	8,364,909

Table 1. Number of CIDR IP prefixes extracted from IRR databases per RIR on August 1st, 2014

Name	Meaning
announced father	prefix is in IRR and BGP
unannounced father	prefix only in IRR
child	prefix only in BGP
announced family	at least the father or one child prefix is seen in BGP
nbr. children	nbr. of internal assignments (i.e. dividing IP space)
nbr. subfamilies	nbr. of external assignments (i.e. delegating IP space)
overlap ratio	fraction of father's IP space used by either children, or subfamilies

Table 2. Quick guide to the metrics defined for the analysis

that around 90% of autonomous systems register at least a subset of their BGP prefixes in the IRR database, making the information it contains valuable. In the end, even though IRR information needs to be considered with a grain of salt, we demonstrated in [13] that it provides a unique insight into BGP ground truth information.

2.2 BGP Data

Our source of BGP data is RIPE RIS [8]. We parse binary files that contain a dump of the BGP messages exchanged between the RIPE collector router and its BGP peers. We focus on `update` messages, that contain prefix announcements and withdrawals, as well as the AS path to the prefix. The AS path is an attribute that contains the list of ASNs (Autonomous System Numbers) which need to be crossed before reaching the destination. The last number in this list is known as the origin AS, i.e. the AS in which the prefix resides.

We process BGP `update` messages according to RFC4271. Namely, we maintain an adjacency table for each of our peers. A prefix is reachable if at least one of our peers has announced it; and is not reachable once every peer that had announced it has withdrawn it. In this way, we are able to build our own BGP routing table, which is dynamically updated as BGP messages flow between routers.

We selected RIPE's Amsterdam collector (rrc00) as our data feeder. It is the best-connected RIPE collector, with over 40 geographically diversified peers. The selected time window for the analysis was the whole month of August 2014, where we counted 629,595 distinct IP prefixes.

3 Methodology

3.1 Definitions

In Section 1, we stated why simply comparing overlapping prefixes together does not produce meaningful results. Instead, we use a combination of semantic data that we extract from the IRR database, and routing information that we get from BGP. In this section, we present how we group these elements into *families of prefixes* that are composed of a family *father*, of *children*, and of *subfamilies*.

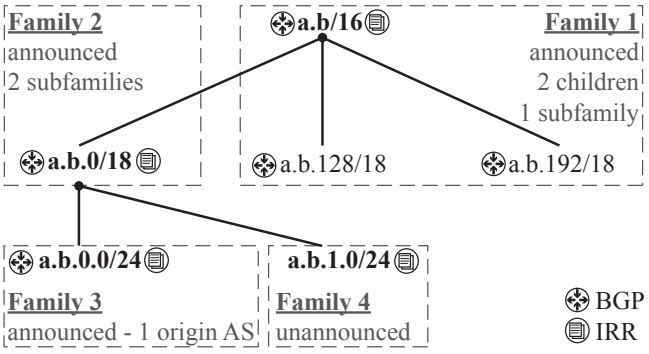

Fig. 1. Example of constitution of families and subfamilies

Each prefix included in the IRR database is *always* the **father of a family**. Consequently, we have as many distinct families as the number of filtered IRR prefixes (see Table 1). Because most of these prefixes overlap, some family fathers completely include some other family fathers. This situation leads to subfamilies. A **subfamily** is a family whose father is completely included in the IP space generated by another family's father.

For example, there are 4 distinct families in Figure 1, because the prefixes a.b/16, a.b.0/18, a.b.0.0/24, and a.b.1.0/24 are included in the IRR database. Incidentally, these 4 prefixes are the fathers of their families. However, some of these fathers overlap. As a result, in Figure 1, Family 2 is a subfamily of Family 1, because the father of Family 2 is more specific than the father of Family 1. Similarly, Family 3 and Family 4 are subfamilies of Family 2. However, neither Family 3 nor Family 4 is a subfamily of Family 1 because Family 2 "hides" them from Family 1. This accounts for the fact that a.b.0/18 has been delegated to another entity (because it has an IRR entry). In other words, the organization responsible for Family 2 is the one in charge to further subdivide this IP space.

Once the families have been put together, we populate them with BGP data. A prefix seen in BGP is either a family father, or a prefix more specific than a family father. In the first case, there is nothing left to do: a father already belongs to the family it defines. In the second case, the prefix is added to the family as a child prefix. A **child** is a prefix

seen in BGP that is more specific than the family father, but not declared in the IRR database as having been assigned to another entity. The child is consequently managed by the organization linked with the IRR record of its family father.

Continuing with the example depicted in Figure 1, three family fathers are announced in BGP: a.b/16, a.b.0/18, and a.b.0.0/24. Moreover, two non-IRR prefixes (a.b.128/18 and a.b.192/18) are also announced. Since they are both more specific than a.b/16, they are added as children in Family 1.

To summarize, we use the prefixes in the IRR database as a binding link between an organization in the real-world, and one or several BGP prefixes. An IRR prefix induces a family, which contains a certain number of children (BGP prefixes).

3.2 Metrics

In this section, we present the metrics that will be used in Section 4 to analyze prefix families.

The **number of children in a family** indicates the number of assignments that have been done internally in this family. In other words, this is the number of distinct IP zones that exist in this family, each possibly leading to different locations, but which should all be under the authority of the same organization. We put this number in relation with the **number of *aggregated* children in a family**, which is the number of prefixes resulting from an aggregation process on the children prefixes. Both sets of prefixes generate the exact same IP space, but the aggregated set does so with the minimal number of prefixes. Consequently, a difference in the number of children and the number of aggregated children indicates that internal assignments were done with contiguous IP blocks. For example, in Figure 1, Family 1 has 2 children: a.b.128/18 and a.b.192/18. These prefixes define IP addresses that are contiguous, and they are aggregated as a.b.128/17. Thus, Family 1 has only 1 aggregated child.

The **number of subfamilies** in a family indicates the number of prefixes that have been delegated to other entities. This number is a constant in our method, because it results from the contents of the IRR database. We put this number in relation with the **number of *announced* subfamilies**, i.e. the number of subfamilies that were actually announced in BGP. We consider that **a family is announced** at time t if either the family father or one of the family child is announced in BGP at t. As an example, Figure 1 depicts Family 2, which has 2 subfamilies. However, since a.b.1.0/24 is not announced in BGP (and has no child), it is marked as unannounced. Consequently, Family 2 only has 1 announced subfamily. Please note that we use the term "unannounced" to refer to the fact that a prefix is not seen as-is from the BGP control plane; it does not imply that the prefix is not used, or that no host is connected using an IP address included in this prefix. The prefix can be routed by a less specific announcement (resulting, for example, from a route aggregation).

The **children overlap ratio** is the ratio of the number of IP addresses available to family children divided by the number of IP addresses available to the family father. In the same fashion, the **subfamily overlap ratio** is the ratio of the number of IP addresses available to the *announced* subfamilies divided by the number of IP addresses available to the family father. For example, the children overlap ratio for Family 1 of Figure 1 is 0.5, and the subfamily overlap ratio is 0.25.

Because the information contained in this section is quite dense, Table 2 provides a summary of the metrics that have been defined; and which can be used as a quick reference guide while going through the results presented in Section 4.

4 Results

4.1 BGP vs IRR Database

In this section, we briefly compare the prefixes inside the IRR database and the prefixes announced in BGP.

In Section 2, we saw that we parsed over $8 \cdot 10^6$ IP prefixes from the IRR, and just a little less than 630k from the BGP control plane. When we compared the distribution of the number of prefixes in both sources according to their mask length, we saw that the number of prefixes in both sources was comparable for /24's and larger prefixes. For smaller prefixes (i.e. prefixes with a mask length $>$ /24), there is at least a factor 100 of difference in the number of entries. In other words, only 1% of IRR prefixes were seen as-is from the BGP control plane, meaning that only 1% of families whose father is more specific than a /24 prefix are *announced*. This phenomenon can be explained by two reasons. First, these prefixes have a long mask, and BGP good practices indicate that prefixes longer than /24s should not be propagated [5], and confirming previous experiments on that topic [3]. Second, IRR database entries are not restricted to BGP users. Any assignment of IP blocks, for example by an ISP, is a potential entry in the IRR database, even though the ISP and its customer are not connected via BGP (but, for example, via DSL or cable). This also explains the high number of /32 entries in the IRR database (i.e. single IP addresses): these may be dedicated servers, and an entry in the IRR provides the rightful technical contact information. For unannounced families, there is a difference between the owner of the IRR prefix and the (BGP) manager of the prefix. The manager of the prefix is generally the ISP of the owner, the one that makes sure that the network is adequately connected. The owner of the prefix is the organization actually hosting machines on the IP addresses within the prefix, which is what the IRR entry specifies. For example, one of Eurecom's prefix is 193.55.113.0/24, which is announced from our provider, Renater, as an aggregated /15 prefix. However, the `inetnum` object for the prefix points to Eurecom, even though it is maintained by Renater.

We now focus on the relative size of children and subfamilies in a family. Figure 2 plots the distribution of the mask length of children prefixes according to the mask length of the family father. The x axis represents the mask length of the family father, and the y axis represents the mask length of the child. The plot data is the histogram of the distribution: the thicker the line is at a coordinate, the more prefixes there are of this size. As we can see, the bulk of the distribution is around children with a mask length of 24, regardless of the father. The fact that the distribution of children prefixes does not depend on the size of the father is surprising. Indeed, one would expect larger families to divide their IP space into bigger zones. The sparsity of available IPv4 addresses could explain this observation since RIRs and, consequently, ISPs prefer to distribute smaller blocks.

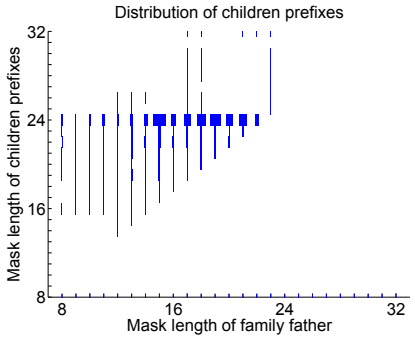

Fig. 2. Distribution of children prefix mask length depending on family father mask length

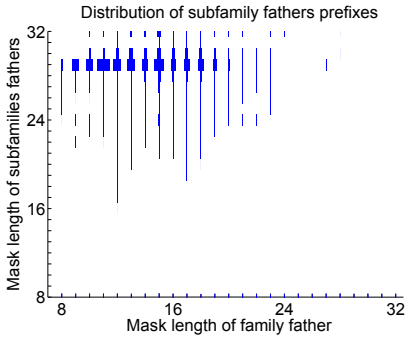

Fig. 3. Distribution of subfamily fathers prefix mask length depending on family father mask length

Figure 3 plots the distribution of the mask length of subfamily fathers prefixes according to the mask of the family father. Here, the bulk of the distribution is around /29, regardless of the mask length of the family father. This raises the question of why these assignments appear to be so popular. In our view, a /29 prefix contains 6 usable IP addresses, which, in today's Internet, is just enough for a small-to-medium size corporation: a couple of IP addresses for publicly accessible servers, plus a couple more for NAT gateways. As tier-3 ISPs typically offer Internet access to a number of SMEs, this could naturally result in a predominance of /29 assignments.

Finally, we have 194,465 families announced in BGP. This amounts to only 2.32% of the total number of families from the IRR database. The results that we present now apply only to those announced prefixes; nothing else can be said about the other ones strictly from a BGP point of view. Moreover, the figures in the remainder of this section *always* plot the *time-weighted average* of the specified metrics. As a result, plots of discrete metrics show continuous values. For example, a plot showing 0.1 child could mean that there was a single child, but 10% of the time.

4.2 Children and Subfamilies

In this section we study the number of children and subfamilies a family has. We look at the IP space occupied by children and subfamilies; and we look at the correlation between children and subfamilies.

Figure 4 plots the number of children per family, and the number of aggregated children per family. It shows that only around 25% of families have, on average one or more children, while the probability of having a large number of children decreases very rapidly. Comparing with the number of aggregated children, we see that in 16% of cases, the families have one aggregated child. This means that, for 16% of families, the IP space dedicated to children is contiguous. This indicates that prefix owners prefer to assign contiguous IP space in order to avoid fragmentation (which may lead to more complex and error-prone network configurations).

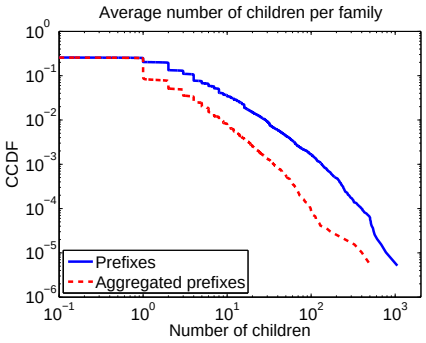

Fig. 4. Number of children per family

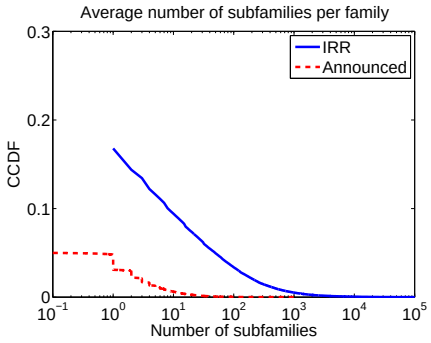

Fig. 5. Number of subfamilies per family

Figure 5 plots the number of subfamilies per family. The plot indicates that, in the IRR database, only 17% of families have at least one subfamily. This can be explained by the rather large number of very-specific assignments (masks $\geq$ /29) in the IRR database: these prefixes are directly allocated to end networks, not to networks providers. On the other hand, only 6% of *announced* families have at least one subfamily, with 1% of them having less than one subfamily on average.

Figures 4 and 5 underline that the vast majority of announced families have neither children nor subfamilies. Table 3 shows the proportion of announced families, according to having children or subfamilies. 73% of families do not have children or subfamilies. In other words, for 73% of the announced families, the prefixes announced in BGP match the prefixes that were assigned, as shown in the IRR database. No further sub-allocation was done by the end-user, either internally (i.e. using child prefixes), or externally (i.e. subfamilies).

Furthermore, there is no correlation between the number of children and the number of subfamilies. The Pearson correlation coefficient, as well as the Spearman correlation coefficient have values between 0.14 and 0.25, depending on if we include or not families without any child or subfamily. In other words, a lot of children implies neither few, and neither a lot of subfamilies; and vice versa. We further study these two dimensions (number of children, and number of subfamilies) in Section 4.3, where we present case studies.

Because we study the relationships between overlapping prefixes, we must limit our analysis to the 27% of families that *do* have children, subfamilies, or both (see Table 3). Consequently, the results presented in the remainder of this section only apply to these 27% of (announced) families.

We now focus on the fraction of IP space of the family father that was allocated to children or to subfamilies. Figure 6 plots the children overlap ratio and the family overlap ratio. It shows that there is no overlap by subfamilies for about 80% families. This is because, as indicated in Table 3, in most cases, there is no subfamily when there are children. In contrast, children can occupy a much larger fraction of the family father IP space, up to 100% in 45% of the cases.

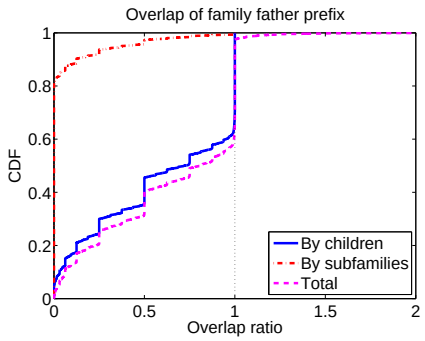

Fig. 6. Prefix overlap within a family

Child	Subfamily	Count	%
N	N	141,883	72.96%
N	Y	1,930	0.99%
Y	N	42,734	21.96%
Y	Y	7,918	4.07%
Announced families		194,465	100%

Table 3. Announced families

Figure 6 also plots the sum of both of these ratios for the families. Interestingly, this ratio exceeds 1 for a few cases. Effectively, this means that, for about 3% of families, children prefixes and subfamily prefixes overlap the family father more than once. Consequently, they also overlap each other. An example from the real world for this situation is the following one. The IRR database lists five prefixes: 5.102.0.0/19, 5.102.{0,8,16,24}.0/21. All these prefixes are also announced in BGP, plus two more: 5.102.{0,16}.0/20. As a result, the /19 family has four subfamilies that fully overlap the family father, and two children, which also fully overlap the family father. All prefixes are originated by a single AS, and belong to the same organization (a tier-3 ISP). It is worth noting that the time-weighted average values of these metrics were over 0.9 in both cases, indicating that this configuration was not transient.

4.3 Real-World Case Studies

We consider in this section a few real-world cases to illustrate the typical relationship that can exist between the business of a company and the breakdown of its prefixes into subfamilies and children. We pick 21 companies – listed in Table 4 – that can be classified into three categories: tier-1, tier-3 and private corporations. The classification is approximate because companies acting as tier-1 providers can also run a tier-3 business at the same time, i.e. directly connecting end-users/small companies to the Internet. This is, for instance, the case of AT&T and Deutsche Telekom.

When looking at ISPs, and regardless of their size (i.e. tier-1 or tier-3), we observe a trend of having a large number of subfamilies and a comparatively smaller number of children. The sheer number of subfamilies suggests that ISPs routinely insert information about prefix delegation in the IRR database. This is in line with expectations: ISPs typically offer Internet access to other companies, and thus assign a set of IPs to its clients. Doing so, the ISPs choose to push this information into the IRR database, because it can be used for administrative purposes. We also observe again the trend that only a small fraction of these families are announced in BGP. This is because ISPs mostly provide Internet connectivity local businesses or home users, that would reap no benefit from the complexity and overload of running a BGP router.

Business type	Name	#fathers	#children	#subfamilies
Tier-1 ISP	AT&T	64	363	582,863
	Cogent	39	87	1,416
	DeutscheTlkm	26	5	58,055
	NTT	152	466	2,744
	TeliaSonera	9	0	247
Tier-3 ISP	Belgacom	15	0	3,710
	Comcast	66	119	14,945
	Free	15	8	3,864
	Rogers	36	187	23,778
	Tele2	29	4	2,852
Private Corp.	Amazon	18	1	15
	Apple	2	196	1
	BBC	2	2	61
	DHL	2	21	0
	eBay	5	1	0
	HSBC	5	6	0
	Microsoft	40	86	3
	OVH	43	9	27,489
	Philips	8	0	0
	Sony	3	2	0
	Yandex	49	18	2,191

Table 4. Real-world case studies

For private corporations, the number of children is much higher than the number of subfamilies. We attribute this to corporations considering internal network policies as private information, thus not wanting to reveal additional company information it in the IRR database (e.g. branch office location). We see two noticeable exceptions: Yandex and OVH. Yandex operates the largest search engine in Russia, along with a number of additional services (cloud storage, etc). The reason for the large number of subfamilies might be due to Yandex pushing up information concerning client companies (e.g. in the case of Web hosting service) in the IRR. The case of OVH is easier to diagnose: OVH offers PaaS and IaaS services, and reports in the IRR database the set of addresses assigned to each clients, just like an ISP would do.

5 Related Work

Previous work in this area can be divided into two categories: works that analyze the BGP routing table growth; and works that aim at validating BGP routing announcements using IRR data.

The evolution of the BGP routing table has been studied many times, most famously by [6], which reports on the growth of the routing table size from the mid 1990's to today. The analysis also includes AS number usage, average AS path length, and other typical BGP aspects. Other papers, such as [2], investigate the reasons behind this

growth, and classify the prefixes inside the routing table depending on the reason for which they need to be announced. The methodology used by [4] to study the evolution of aggregation practices over time may bear some similarity to ours, but differs in several key aspects. Most notably, [4] provides limited prefix grouping methodology, where we make active use of the semantic information found in the IRR databases in order to group prefixes into families that are owned by the same organization. We consequently consider assignments made at the edge of the network by tier-2/tier-3 ISPs. We better illustrate and explain the relationships between the overlapping prefixes inside these families, whereas [4] focuses more on the dynamics of the BGP announcement and their consequences on BGP router processes. As a result, our methods are not directly comparable, even though the BGP-sides of the analyses exhibit similar global trends.

Validation of routing data based on IRR databases entries has been attempted to make the BGP infrastructure more robust. For example, [11] used IRR data to build a tool that informs network administrators of an anomaly that should be further investigated. More recently, [7] studied the validity of the association between a prefix and its origin AS in the IRR. The overall conclusion of this type of work is that the quality of the data inside IRR databases is highly dependent on the RIR. However, it also appears that more recent studies suggest that the IRR provides information that can be used in order to improve the security level of BGP.

6 Conclusion and Future Work

In this paper, we detailed how we use assignment data from the IRR database as semantic anchor points in order to cluster prefixes from the BGP routing table into families, inside of which we can non-ambiguously study the overlap among these prefixes.

We showed that the IRR database contains many times more prefixes than the BGP routing table. This is particularly true for prefixes with a mask length longer than 24. At the same time, we found that only 2.32% of the families induced by these IRR entries were effectively seen from BGP. We attribute this difference to the fact that IRR entries are not restricted to BGP players, but can exist due to any IP assignment. For example, there are single IP addresses (i.e. /32 prefixes) with an IRR entry for administrative reasons.

We showed that 74% of the announced families do not have children. This means that, for these families, only the prefix that was assigned is announced in BGP, which does not lead to (additional) routing table entries. It is also in accordance with the standard BGP good-practice of always announcing the assigned prefix. For about 15% of all families – but about half of the families with children or subfamilies – this practice is not met, which means that these families are part of the *dormant* IP space, which appears to be more vulnerable to malicious prefix hijacking attacks, as demonstrated by [14].

A key take-away from our study is that a joint analysis of BGP and the IRR database sheds light on the way the IRR is used, and also enables to uncover different types of business practices. For instance ISPs (large, or small) are more likely to register their customer in the IRR database, leading to a greater number of subfamilies than children. Clients of ISPs being, most of the time, relatively small, the most popular flavour

of subfamily is a prefix of mask length 29, which constitutes enough addresses for a small business. In our view, this implies that ISPs devote a lot of energy to populate and maintain their IRR entries. We argue that this is a proof that ISPs find the information in the IRR valuable. Consequently, even though IRR information is not perfect, it *cannot* be dismissed as entierly stale, inaccurate, and/or bogus. It provides a unique (administrative-level) insight into IP networks, and can help better understand a number of routing phenomenons, as we demonstrated in [13].

We see different possible ways of extending the scope of this work. First, we will study the AS-level relationship between father prefixes, their children, and between families and their subfamilies. Second, we would like to further study unannounced families. A first clue to know how much of that information is stale would be perform IP-level measurements, such as traceroutes, in order to see how the IP-level topology for addresses within the unannounced family differs from the topology inside the announced family. A complementary result from this experiment would be further validating the prefix/organization mappings that are available in the IRR.

References

1. APNIC: Using Whois: Quick Beginners Guide. http://www.apnic.net/apnic-info/whois_search/using-whois/guide
2. Bu, T. et al.: On Characterizing BGP Routing Table Growth. Comput. Netw. (May 2004)
3. Bush, R., Hiebert, J., Maennel, O., Roughan, M., Uhlig, S.: Testing the reachability of (new) address space. In: Proceedings of the 2007 SIGCOMM Workshop on Internet Network Management. pp. 236–241. INM '07 (2007)
4. Cittadini, L., Muhlbauer, W., Uhlig, S., Bush, R., Francois, P., Maennel, O.: Evolution of internet address space deaggregation: Myths and reality. IEEE J.Sel. A. Commun. 28(8), 1238–1249 (Oct 2010)
5. Hu, X., Mao, Z.: Accurate real-time identification of ip prefix hijacking. In: IEEE Symposium on Security and Privacy (May 2007)
6. Huston, G.: BGP Reports. http://bgp.potaroo.net/
7. Khan, A., Kim, H., Kwon, T., Choi, Y.: A comparative study on IP prefixes and their origin ases in BGP and the IRR. Computer Communication Review pp. 16–24 (2013)
8. RIPE NCC: Routing Information Service. http://www.ripe.net/ris/
9. RIPE NCC: YouTube Hijacking: A RIPE NCC RIS case study. http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study (March 2008)
10. Siganos, G., Faloutsos, M.: Analyzing bgp policies: methodology and tool. In: INFOCOM 2004. vol. 3, pp. 1640–1651 (March 2004)
11. Siganos, G., Faloutsos, M.: Neighborhood watch for internet routing: Can we improve the robustness of internet routing today? In: IEEE INFOCOM (2007)
12. Toonk, A.: Looking at the spamhaus DDOS from a BGP perspective. http://www.bgpmon.net/looking-at-the-spamhouse-ddos-from-a-bgp-perspective/ (March 2013)
13. Vervier, P.A., Jacquemart, Q., Schlamp, J., Thonnard, O., Carle, G., Urvoy Keller, G., Biersack, E., Dacier, M.: Malicious BGP hijacks: appearances can be deceiving. In: ICC CISS 2014, IEEE International Conference on Communications. Sydney, Australia (06 2014)
14. Vervier, P.A., Thonnard, O., Dacier, M.: Mind your blocks: On the stealthiness of malicious BGP hijacks. In: NDSS 2015, Network and Distributed System Security Symposium, 8-11 February 2015, San Diego, California, USA. San Diego, CA, USA (02 2015)