A Square Law Revisited

Brian E. Carpenter
The University of Auckland
brian@cs.auckland.ac.nz

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

An earlier study observed that until 2008, the size of the BGP4 system for IPv4 appeared to have grown approximately in proportion to the square root of the host count of the globally addressable Internet. This article revisits this study by including IPv4 data until 2020 and adding IPv6 data. The results indicate that BGP4 for IPv4 is continuing to scale steadily even as IPv4 approaches its end of life, and that it is working as it should for IPv6, except for a slight concern that the number of announced routes is trending upwards faster as time goes on.

CCS CONCEPTS

• **Networks** → **Network architectures**; **Network dynamics**; **Public Internet**;

KEYWORDS

Internet topology, BGP, inter-domain routing

1 INTRODUCTION

More than a decade ago, a previous article [1] examined data on the size of the Internet and on deployment of the BGP4 inter-domain routing protocol [18]. The article also gave a general introduction to the role of BGP4 and its growth, which is not repeated here. The results indicated a square law relationship between the number of unique Autonomous Systems (AS) in the BGP4 system and a lower-bound estimate of the total number of active and addressable Internet hosts. To be precise, Figure 3 of [1] showed that over a number of years, the the number of ASes appeared to remain in a linear relationship with the square root of the number of hosts. If we let A be the AS count, and H be the host count, then the data up to 2008 fitted the relationship $A = 1.35\sqrt{H} - 4615$, where the standard error of the coefficient was 2.2%.

The article [1] hypothesised that this relationship resulted from the overall "shape" of the Internet being a star of stars, with the host count being proportional to its circumference. Consistently with that hypothesis, a significant factor at the time was that the large majority (86 %) of ASes were purely originators of routes, and most of the rest originated some routes as well as providing transit [17], [16].

Another observation was that the number of BGP4 routes per AS was approximately 9 over the years 2001-2008. This was interpreted to indicate the success of CIDR [8] in achieving a reasonable level of route aggregation.

Much has changed since 2008. Firstly, the above observations related entirely to IPv4, since IPv6 deployment was very small during the years concerned. Secondly, exhaustion of the IPv4 address pool was still well in the future. Thirdly, network address translation (NAT) was largely confined to the very edges of the network. Fourthly, Internet access from mobile telephones was still relatively new and not universal. For these reasons, this article revisits the data to see what has happened.

2 PRIMARY DATA

As in [1], the primary data source for BGP4 and AS statistics are observations collected over many years at APNIC's research and development network by Geoff Huston [9]. For the number of IPv4 hosts, the same proxy is used as before, namely the number of IPv4 reverse lookup (IN-ADDR.ARPA) domain names collected in the regular surveys carried out by ISC Internet Domain Survey [13]. The validity of this measurement is discussed at more length in [1]. It is worth recalling that this number counts the number of active and globally routeable IPv4 addresses that have associated DNS information; thus multiple systems behind a single NAT will count as one. It does not claim to be a count of the number of hosts that have access to the Internet via NATs or other screening mechanisms, it does not include hosts whose connectivity is internal to a corporate or private network, and it certainly does not indicate the number of human users. What it gives is an estimate of the number of IPv4 addresses than can effectively be reached across the Internet, which is a lower bound on the number of IPv4 hosts that can be reached; as a reminder of its origin, this paper refers to it as the "domain count." Other interpretations of the size of the Internet may give much higher numbers (for example the report[14] estimates 4.5 billion users), but they are not numbers that have relevance to the evolution of BGP4.

There is no automatic relationship between the number of BGP4 routes announced and the number of active IPv4 addresses, because BGP4 routes lead to IPv4 address prefixes of various sizes which are used with varying efficiency, i.e., not all possible addresses within each prefix are active. Readers interested in the overall behaviour of the BGP4 system as such should consult [10] and many other articles on the same web site. Here, as in [1], we focus on the *observed* relationship between the BGP4 system and the number of visible IP addresses.

Unfortunately the ISC survey was ended in January 2019, and it never covered IPv6. The observation date chosen is the middle of July each year (except for 2003, when the ISC survey did not take place in July).

The APNIC data series for IPv6 started in 2004 but is otherwise directly comparable with the data series for IPv4. To estimate the host count, the solution adopted is to make use of the data series on IPv6 utilisation published by Google Inc. [5, 12], which started in 2008. This gives the percentage of Google users who *actually* access Google services via IPv6 on any given date. Since, until now, the vast majority of IPv6 end-users were also IPv4 users (by dual-stack or an equivalent solution), we estimate the total number of active IPv6 hosts as the ISC IPv4 domain count for a given date multiplied by the IPv6 percentage observed by Google on that date. The observation date chosen is the middle of July each year. As explained in [5], the Google IPv6 percentage is not dependent on the actual IP addresses involved, but only on the IP version number used by a random sample of Google users. This is fortunate, because it makes the estimated IPv6 host count independent of whether those hosts are using stable or temporary IPv6 addresses (also known as "privacy" addresses). If anything, we obtain an underestimate, because domestic gateways that have one IPv4 address may support several dual-stack users. This difference makes any direct comparison of the IPv4 and IPv6 trends of limited value.

Clearly this method is far from perfect and will be no use in the future, because the ISC survey has ended and because the number of IPv6-only end users will increase. However, it serves the purposes of this article, which covers a time span when essentially all IPv6 access was from dual-stack users.

2.1 IPv4 Data

The data for IPv4 are shown in Table 1. Due to a discontinuity in the data, the domain count for 2016-07 is a linear interpolation between 2016-01 and 2017-01. The drop in the domain count in 2018 is assumed to be real, and most likely a reflection of the increase in carrier-grade NAT (CGN) usage by Internet Service Providers (ISPs), since a single CGN hides far more end-users than an ordinary domestic NAT. CGN deployment has been significant since at least 2015 [19]. The driver for this was of course the ongoing depletion of the supply of IPv4 addresses, which indeed almost ran out in several regions in the same period [2, 7]. Where one IPv4 address would previously have been used by a single user or a small group behind a domestic or small office NAT, now quite large numbers of users are hidden behind a single address. The observations in [19] suggest that at least 64 subscribers per IPv4 address is common practice, and much higher numbers are possible.

The domain counts for 2019 and 2020 are arbitrarily set to the 2018 value, since, with the end of the ISC survey, there is no scientific justification for any other method of extrapolation. In any case, because of IPv4 address exhaustion, it certain that the actual count did not increase. (Clearly there is no implication that the actual number of end-users fell.)

Unfortunately no extrapolation to 2021 is scientifically reasonable, so evaluating any impact of the COVID-19 pandemic is impossible.

2.2 IPv6 Data

The data for IPv6 are shown in Table 2. As explained above, the IPv6 host counts were estimated as the IPv4 domain count multiplied by the fraction of IPv6 Google users in mid-July. The numbers for 2019 and 2020 are therefore based on frozen IPv4 numbers, and no 2021 data are considered. Since the IPv6 fraction observed by Google has increased every year, the estimated IPv6 number increases even as the IPv4 number reaches saturation.

3 GRAPHICAL PRESENTATION

The following figures show two different aspects of the above data sets; the earlier data were presented similarly in [1].

3.1 IPv4 Graphs

The relationship between the BGP4 table size (i.e. the number of IPv4 routes advertised) and the number of active ASes observed remains straightforward, as it was in [1]. To a reasonable approximation it is linear (see Fig. 1) and continued exactly as might have been predicted in 2008. The visual continuity does hide a change: in 2008 the slope was computed as 8.06 (using gnuplot's least-squares fit); today it is 11.98. Nevertheless, the smooth curve does imply that there has been no major change in the extent to which ASes are announcing aggregated BGP4 routes. Consistently, other work [4] concludes that while de-aggregation of prefixes in support of traffic engineering certainly occurs, its impact on the BGP4 table size and churn "has not changed for the worse in recent years."

Table 1: IPv4 BGP4 data and domain counts

Year, Month	BGP4 size	Unique AS count	Ratio	Domain count
1994-07	18468			3864000
1995-07	27717			8200000
1996-07	36851			16729000
1997-07	46948	2473	19.0	26053000
1998-07	52199	3695	14.1	36739000
1999-07	62318	5287	11.8	56218000
2000-07	83921	7942	10.6	93047785
2001-07	103095	11283	9.1	125888197
2002-07	111940	13283	8.4	162128493
2003-01	118231	14355	8.2	171638297
2004-07	138726	17498	7.9	285139107
2005-07	163442	20001	8.2	353284187
2006-07	189700	22569	8.4	439286364
2007-07	228856	25762	8.9	489774269
2008-07	273992	28811	9.5	570937778
2009-07	296627	31827	9.32	681064561
2010-07	328905	34812	9.45	768913036
2011-07	366456	38267	9.58	849869781
2012-07	419241	41626	10.07	908585739
2013-07	470159	44635	10.53	996230757
2014-07	509336	47613	10.70	1028544414
2015-07	560654	51127	10.97	1033836245
2016-07	619860	54530	11.37	1055713573
2017-07	675216	58038	11.63	1074971748
2018-07	726322	61442	11.82	1015787389
2019-07	787434	65005	12.11	1015787389
2020-07	824214	68775	11.98	1015787389

Table 2: IPv6 BGP4 data and estimated domain counts

Year, Month	BGP4 size	Unique AS count	Ratio	Host count (est.)
2004-07	626	443	1.41	
2005-07	762	552	1.38	
2006-07	708	599	1.18	
2007-07	878	746	1.18	
2008-07	1422	1062	1.34	799313
2009-07	1951	1495	1.31	1362129
2010-07	3122	2296	1.36	1691609
2011-07	6750	4379	1.54	2549609
2012-07	9820	5844	1.68	6360100
2013-07	13651	7262	1.88	13947231
2014-07	18479	8542	2.16	38056143
2015-07	23438	9917	2.36	72368537
2016-07	31059	11920	2.61	116128493
2017-07	41479	13664	3.04	182745197
2018-07	54464	15596	3.49	223473225
2019-07	72037	17567	4.10	274262595
2020-07	91093	20068	4.54	314894091

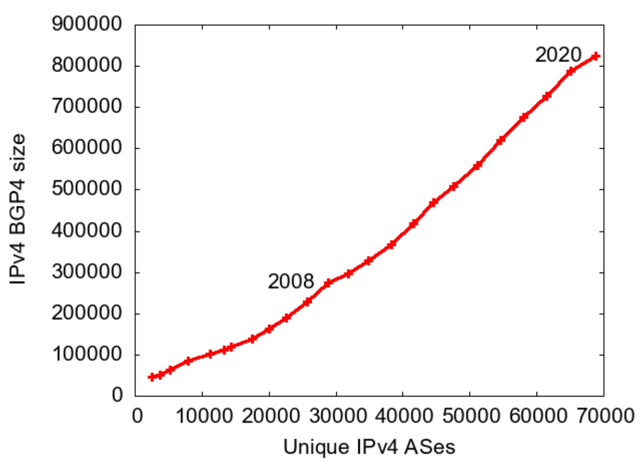

Figure 1: IPv4 routes vs AS count (slope = 11.98)

The other relationship between the AS count and the *square root* of the domain count was strikingly linear in [1] and continued so until 2012, after which it changed dramatically (see Fig. 2). Until 2012, the data remained consistent with the square law observation in [1]. However, as noted in Section 2.1, the growth in the domain count slowed down quite abruptly from 2013 and actually reversed itself in 2018, and this effect is of course exagerrated in the graph by using the square root. Whatever aspect of the BGP4 topology held constant from 1994 to 2012 has changed substantially since then. In particular the "star of stars" hypothesis of [1] clearly no longer applies today.

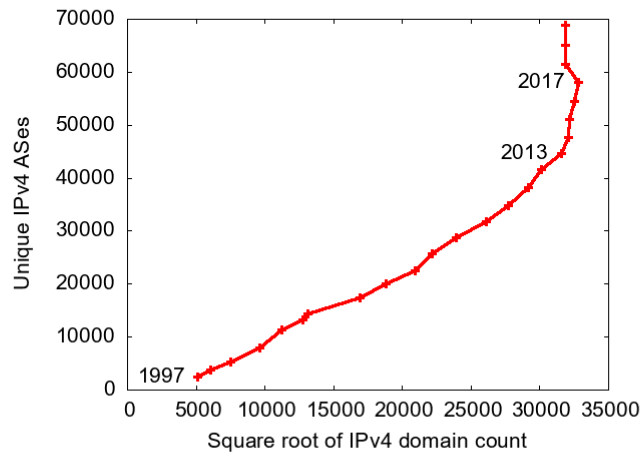

Figure 2: IPv4 AS count vs square root of domain count

3.2 IPv6 Graphs

The relationship between the BGP4 table size and the number of active ASes observed (see Fig. 3) is less linear than in the case of IPv4, presumably because the IPv6 network is still young and is evolving as more IPv6 sites with their own prefixes appear. Strikingly, however, its slope is 4.03, a manifestation that IPv6 route aggregation is in some sense more successful to date than for IPv4. The various Regional Internet Registries that allocate prefixes to operators all have policies that strongly recommend that these prefixes are *not* de-aggregated. However, the slope is clearly increasing, i.e. aggregation appears to be getting worse as IPv6 deployment goes on. This is consistent with the observation that the number of more specific BGP4 routes for IPv6 is now at 60% of the total forwarding table, whereas for IPv4 it is steady at 51%[11].

The relationship between the AS count and the *square root* of the estimated host count is shown in Fig. 4 and is harder to interpret. It seems that for several years (as in the early years for IPv4) the relationship was linear, with a slope of about 1.34, almost identical to the IPv4 slope reported in [1]. The "star of stars" hypothesis could perhaps be revived for those years. However, visibly the slope of the curve is increasing in recent years, so any extrapolation to the future would be implausible.

4 DISCUSSION AND CONCLUSION

The most notable feature of the new data presented here, compared to [1], is the most recent section of the graph in Fig. 2. As the exhaustion of IPv4 addresses took hold, and as CGN became prevalent, the previous approximate square law relationship between the number of BGP4 autonomous systems and the size of the addressable IPv4 Internet, which had continued from 1997 until 2013, suddenly and almost literally hit the wall. This is completely explained by CGN deployment, since it means that from the addressability viewpoint, the "edge" of the Internet is pulled back from the subscriber's site to the CGN, which is inside the provider's infrastructure. This applies both to traditional fixed-line subscribers and to mobile subscribers.

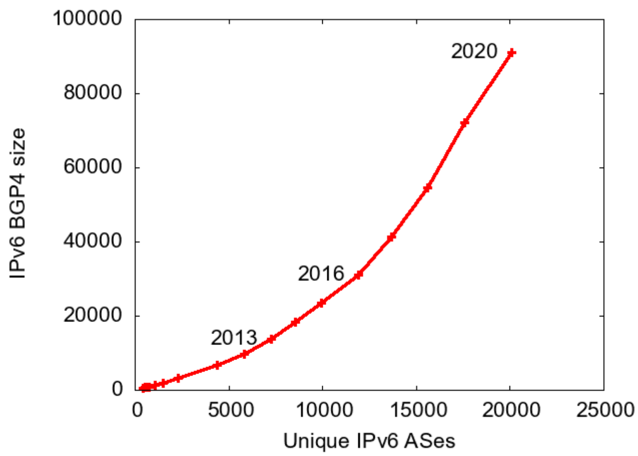

Figure 3: IPv6 routes vs AS count (slope = 4.03)

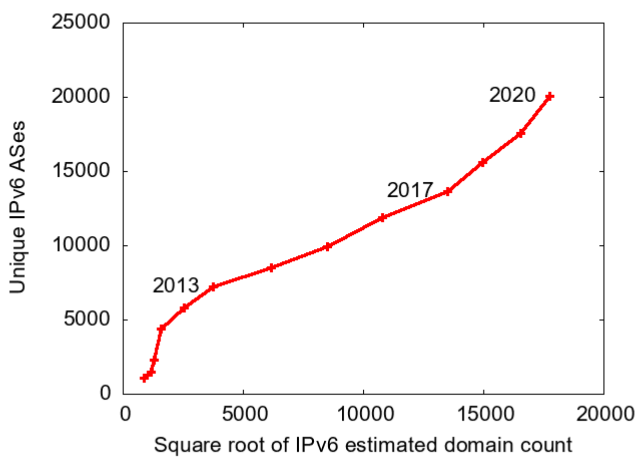

Figure 4: IPv6 AS count vs square root of estimated host count

Even if the subscriber is a mobile telephone using IPv4 as a secondary service over IPv6, for example with the 464XLAT technique [15], the effect is the same as CGN, since many users will share a single IPv4 address. In these cases, a given subscriber may even be dynamically allocated ports from different IPv4 addresses. The notion of "one subscriber, one address" has more or less broken down. We can clearly see that during the years since 2014 until the end of the ISC domain survey, growth of the addressable edge of the IPv4 Internet essentially stopped. (This does not of course imply that growth in the number of IPv4 hosts inside home and enterprise networks stopped, but that growth is invisible to the metrics used in this article.)

As a side note, the "wall" turned out to be at approximately one billion IPv4 addresses, whereas the theoretical size of the address space is 2^{32}, or approximately four billion. This is explicable as follows. Firstly, some of the address space is reserved for special purposes such as multicast addressing. Secondly, many addresses are used within large enterprises and are never active on the Internet as such. These addresses are hidden, but not necessarily wasted. Thirdly, not all subnets are full, in the sense that in a subnet with n bits available, often there are fewer than 2^n hosts on the subnet, and the remaining addresses cannot be used. In terms of the logarithmic host density ratio for address assignment [6], we can calculate that the final HD-ratio of the visible IPv4 Internet is approximately $log(10^9)/log(4 \times 10^9) = 0.94$. This is well above the limit that according to [6] corresponds to "a high pain level, at which operators are ready to make drastic decisions." While the scientific value of the HD-ratio is debatable, switching to IPv6 is indeed a drastic decision for network operators.

It is nevertheless reassuring for the stability of the IPv4 BGP4 system that the smooth relationship between the number of routes advertised in BGP4 and the number of active autonomous systems was not affected by this rather dramatic change at the edge (Fig. 1).

As for IPv6, Fig. 3 is of slight concern, as it shows a tendency for the number of routes to increase faster than the number of ASes as time goes on, especially since 2016. The policies of various Regional Internet Registries (e.g. [3]) allow the assignment of 48-bit "provider independent" (PI) prefixes directly to user sites, in parallel with the assignment of shorter prefixes to ISPs. PI prefixes *might* be aggregated to some extent by upstream ISPs, but this is not guaranteed. The graph suggests that the fraction of unaggregated routes has increased since 2016, as more and more operators have enabled IPv6. As noted in Section 3.2, the fraction of more specific IPv6 routes is higher than for IPv4. While this is not an immediate concern, it is a reminder that BGP4 is not infinitely scalable.

Finally, Fig. 4 suggests that at the moment, the BGP4 topology for IPv6 is following the same trend as for IPv4 in earlier years, although extrapolation of this curve is not justified.

In conclusion, the data show that BGP4 for IPv4 is continuing to scale predictably even as IPv4 approaches its end of life, and that it is working as it should for IPv6 except for a slight concern that the number of announced routes per autonomous system is trending upwards faster as time goes on.

ACKNOWLEDGMENTS

Helpful comments on an early version of this article were received from Nevil Brownlee, Geoff Huston, George Michaelson and Jordi Palet Martinez.

REFERENCES

[1] Brian E Carpenter. 2009. Observed Relationships between Size Measures of the Internet. *ACM SIGCOMM Computer Communication Review.* 39, 2 (2009), 6–12.

[2] Japan Network Information Center. 2017. IPv4 address pool ran out in the Asia Pacific region. https://www.nic.ad.jp/en/ip/ipv4pool/

[3] RIPE Network Coordination Centre. 2020. IPv6 Address Allocation and Assignment Policy. https://www.ripe.net/publications/docs/ripe-738

[4] Luca Cittadini, Wolfgang Mühlbauer, Steve Uhlig, Randy Bush, Pierre François, and Olaf Maennel. 2010. Evolution of Internet Address Space Deaggregation: Myths and Reality. *IEEE Journal on Selected Areas in Communications* 28, 8 (2010), 1238–1249. https://doi.org/10.1109/JSAC.2010.101002

[5] Lorenzo Colitti, Steinar H. Gunderson, Erik Kline, and Tiziana Refice. 2010. Evaluating IPv6 adoption in the Internet. In *PAM 2010.* http://www.pam2010.ethz.ch/papers/full-length/15.pdf

[6] Alain Durand and Christian Huitema. 2001. The Host-Density Ratio for Address Assignment Efficiency: An update on the H ratio. *Internet RFC 3194* (2001). https://doi.org/10.17487/RFC3194

[7] American Registry for Internet Numbers. 2015. ARIN IPv4 Free Pool Reaches Zero. https://www.arin.net/vault/announcements/2015/20150924.html

[8] V. Fuller, T. Li, J. Yu, and K. Varadhan. 1993. Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy. *Internet RFC 1519* (1993). https://doi.org/10.17487/RFC1519
[9] Geoff Huston. [n.d.]. BGP Routing Table Analysis Reports. http://www.potaroo.net/
[10] Geoff Huston. 2021. BGP in 2020 - The BGP Table. https://www.potaroo.net/ispcol/2021-01/bgp2020.html
[11] Geoff Huston. 2021. Private Communication.
[12] Google Inc. [n.d.]. IPv6 Statistics. https://www.google.com/intl/en/ipv6/statistics.html
[13] Internet Systems Consortium. [n.d.]. Internet Domain Survey. https://www.isc.org/solutions/survey
[14] Simon Kemp. 2020. Digital 2020: Global Digital Overview. https://datareportal.com/reports/digital-2020-global-digital-overview
[15] Masataka Mawatari, Masanobu Kawashima, and Cameron Byrne. 2013. 464XLAT: Combination of Stateful and Stateless Translation. *Internet RFC 6877* (2013). https://doi.org/10.17487/RFC6877
[16] Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. 2008. In Search of the elusive Ground Truth: The Internet's AS-level Connectivity Structure. In *SIGMETRICS'08*. 217–228.
[17] Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. 2007. Observing the Evolution of Internet AS Topology. In *Proceedings of the ACM SIGCOMM 2007 Conference*. 313–324.
[18] Yakov Rekhter and Tony Li. 1995. A border gateway protocol 4 (BGP-4). *Internet RFC 1654* (1995). https://doi.org/10.17487/RFC1654
[19] Philipp Richter, Florian Wohlfart, Narseo Vallina-Rodriguez, Mark Allman, Randy Bush, Anja Feldmann, Christian Kreibich, Nicholas Weaver, and Vern Paxson. 2016. A Multi-Perspective Analysis of Carrier-Grade NAT Deployment. In *Proceedings of the 2016 Internet Measurement Conference* (Santa Monica, California, USA) *(IMC '16)*. Association for Computing Machinery, New York, NY, USA, 215–229. https://doi.org/10.1145/2987443.2987474