

Comparing Community Networks to the Internet: an empirical Study of BGP Behaviour

Bart Braem, Johan Bergs, Chris Blondia
iMinds - University of Antwerp - MOSAIC Research Group
{bart.braem, johan.bergs, chris.blondia}@uantwerpen.be

Abstract—Community networks are a new, emerging form of organising local network access and access to the public Internet where people themselves operate and build their networks. However, little research has been performed on the stability of the resulting bottom-up infrastructure. In this work, the BGP behaviour of community networks is empirically studied by evaluating dumps of BGP messages from the AWMN and Guifi community networks and comparing the results to data from the public Internet. The preliminary results show that community networks are significantly less stable, although the public Internet shows a significantly higher ratio of update messages over withdrawal messages.

Keywords—Community Networks, BGP, Internet

I. INTRODUCTION

All over the world community networks are arising as a solution to provide Internet connectivity to areas where commercial Internet is not economically or practically feasible. Sometimes community networks are started *just for fun*, because network engineers or hackers want to experiment with running their own network without relying on the limitations imposed by ISP services. Built and managed by a community of users rather than by a commercial third party, community networks frequently rely on wireless technology, because of the flexibility of using (unlicensed) wireless spectrum when deploying organically. Some networks are even rolling out fiber themselves, leading to very high speed backbone links which in the past traditionally were reserved to more expensive commercial offerings[1].

Because of the DIY approach in community networks, a robust routing protocol is required to cope with frequent outages and topology changes[2]. Moreover, the routing has to be highly distributed and scalable, because of the internal organisation of community networks. Some networks use ad-hoc network routing protocols such as OLSR for this[3], other use novel mesh firmwares like LibreMesh[4], while the largest community networks AWMN and Guifi simply use the Border Gateway Protocol (BGP)[5] with in most cases one BGP AS number corresponding to one node. It is the performance of these BGP networks this work wants to evaluate, as an initial assessment of the stability of community networks from a routing protocol perspective.

Extensive research has been performed on the stability of BGP in the public Internet, especially related to the stability of the BGP routing tables[6], [7]. This work starts from BGP dumps rather than from the routing table itself. Other researchers take an analytical view on BGP stability[8], while we use an empirical approach. To the best of our knowledge,

this is the first study of community networks stability. While research in this field has been growing steadily, the overall stability of this kind of networks is still unknown. This is also the major contribution of our work, to present initial stability measurements of community networks.

The remainder of this work is organized as follows. In section II we introduce the most important concepts of BGP, followed by an overview of the measurement setup in section III. Section IV presents the analysis results and section V concludes this work.

II. THE BORDER GATEWAY PROTOCOL (BGP)

BGP is a well studied exterior gateway routing protocol, already at version 4[9]. At a high level, the protocol exchanges routing information between different Autonomous Systems (AS's) to make path decisions based on path length, peering policies and rules. BGP is best known for being the routing protocol on the Internet, connecting different entities around the world to each other, while adhering to peering agreements. Although the protocol comes with its own set of challenges and security issues[10], its distributed nature, well-known behaviour and broad support in even simple routing solutions led to a broad adoption in community networks.

The protocol has an extensive list of features and configuration possibilities[11]. In what follows we describe the most relevant functional components, focusing on network stability. The two most important messages for this analysis are updates (also called announcements) and withdrawals. While the former serves as a message to inform neighbours and eventually the entire network about newly available routes, the latter signals message deletion from the network. Notice that in a relatively stable network the amount of withdrawals is supposed to be smaller than the number of updates, which indicates a network with possibly multiple parallel routes and no permanent path deletion.

III. MEASUREMENT SETUP

To study the behaviour of BGP in community networks, a BGP monitor has been installed in the AWMN and Guifi networks. The resulting data sets are publicly available¹. Because of connectivity issues some data sets are empty. This work will focus on the data sets from November 2013 in the case of AWMN and June 2014 in the case of Guifi. Efforts to include more recent data are ongoing, however considering publicly

¹Datasets available at <http://opendata-awmn.confine-project.eu/> and <http://opendata-guifi.confine-project.eu/>.

available growth charts of Guifi² we believe the behaviour will be comparable as the network growth seems to stabilize. Unfortunately no growth data from AWMN is available. We consider this an important point of attention for future work.

The data sets contain the daily update/withdrawal and Routing Information Base (RIB) messages, dumped at intervals of 15 minutes. As AWMN is federated with other community networks over the FEDERICA European federation network³, BGP traffic from other networks is expected to appear in the dumps. However, active BGP filtering at the peering links is in place, reducing the number of entries from other community networks to a negligible amount.

We limit ourselves to an analysis of the update and withdrawal messages, as they give a direct indication of the dynamic nature of the networks. In section IV-C, we give an overview of how these update and withdrawal messages affect the overall size of the RIB table.

For analysis of the data the update messages were parsed using the RIPE NCC libBGPDump tool⁴. The dumps were first uncompressed and then inserted in a MySQL database for further analysis. The following columns were taken into account: message timestamp, message type (announcement or withdrawal), originating IP and AS, AS path and announced subnet. In total 2.681.777 rows were inserted in the MySQL database for AWMN, for Guifi this amounted to 4.396.650 rows. These rows were then queried with SQL queries to generate the plots in what follows.

To compare the behaviour of BGP in community networks to the behaviour of BGP in the public Internet, we used the BGP data from Potaroo.net[12], a website which maintains a list of analyses of Internet BGP behaviour.

IV. ANALYSIS

Based on the measurement data gathered with the measurement setup described in the previous section, we first perform an analysis of community networks separately. Then we compare the observed behaviour against public Internet observations.

A. BGP in AWMN

To provide context for the analysis of the stability of the community networks under study, figure 1 gives an overview of messages observed in AWMN during a single day, in this case November 10, 2013. It becomes clear immediately that community networks show strong variations in network stability over time, with peaks in this figure around 1PM and 6PM (UTC). As the number of announcements is much higher than the number of withdrawals, clearly multiple paths between AS'es are present, as confirmed by public topology information of AWMN and Guifi. Although this indicates higher resilience in case of failures of single links, it does imply a higher routing message load.

To assess the influence of individual community network nodes on the overall network stability, for this same day in

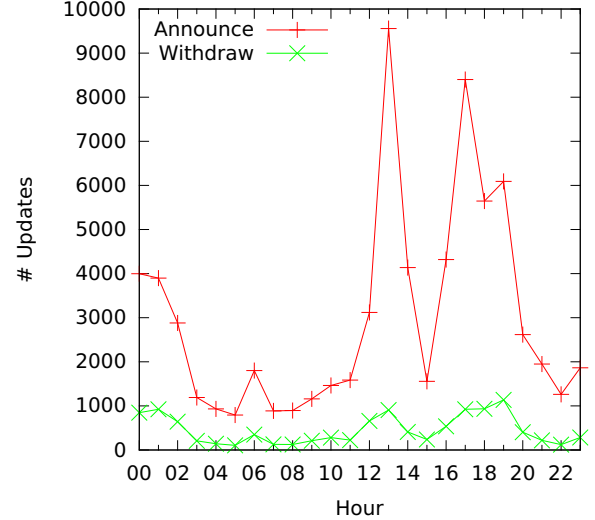


Fig. 1. Updates and withdrawals in AWMN for November 10, 2013.




Fig. 2. Cumulative updates and withdrawals in AWMN for November 10, 2013.

AWMN figure 2 shows the cumulative updates and withdrawals, relative to the AS percentiles. Clearly less than 50% of all AS numbers is responsible for the largest number of updates and withdrawals, with a strong exponential distribution of the announcements and withdrawals. From this we conclude that a number of weak points in the networks are responsible for most routing traffic and as a consequence most instability. This is illustrated by the fact that the top 5% of nodes each generate more than 1000 announcements and 200 withdrawals in a single day.

After zooming out, figure 3 gives the total number of announcements and withdrawals over a period of one month (November 2013) in AWMN, grouped per day. This figure illustrates the unstable, time-independent and largely unpredictable network behaviour per day.

²Growth chart or *corba de creixement* available at <https://guifi.net/guifi/menu/stats/nodes>

³See <http://www.fp7-federica.eu/>

⁴See <https://bitbucket.org/ripencc/bgpdump/wiki/Home>

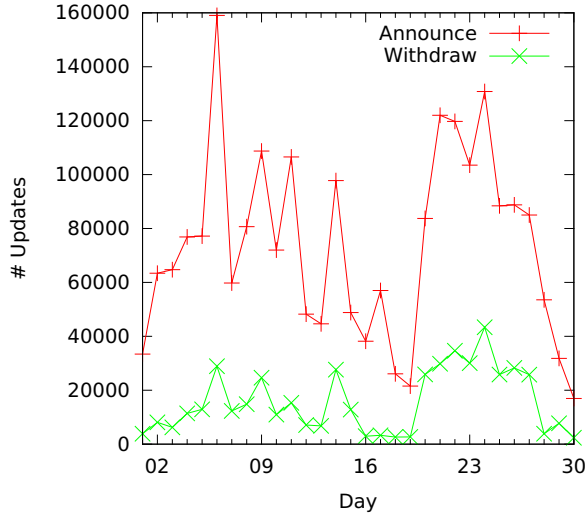


Fig. 3. Total updates and withdrawals per day in AWMN for November 2013




Fig. 5. Maximum number of updates and withdrawals per day in AWMN for November 2013




Fig. 4. Median number of updates and withdrawals per day in AWMN for November 2013

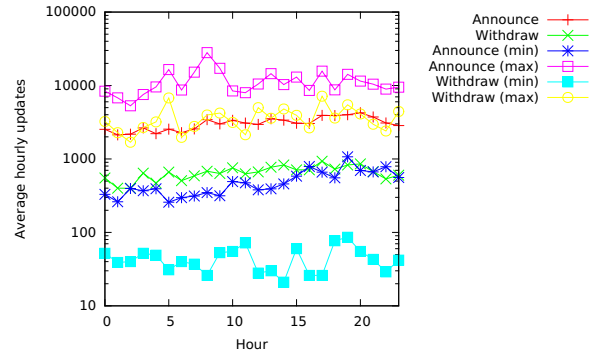


Fig. 6. Monthly average number of updates and withdrawals per hour in AWMN for November 2013

For the same range figures 4 and 5 give the median number of updates per AS and maximum number of updates per AS respectively, showing how the distribution of the number of updates is not uniform. A small part of all AS'es are generating a relatively large amount of announcements, indicating that only parts of the AWMN network are unstable.

When considering the monthly averages per hour of AWMN as illustrated by figure 6, a weak diurnal pattern can be observed. During working day hours, from 6AM until about 8PM, more announcements are generated and the number of withdrawals is slightly higher. This indicates a larger degree of network instability or at least of changes to the network topology. In this case we believe this is the result from the very nature of community networks, where members maintain the network during daytime. Especially the high number of

updates in the late afternoon and the evening is not expected in commercial Internet offerings, where most operations happen during low-traffic periods. However, the pattern is quite weak, possibly smoothed out by the stronger, time-independent variations in the network.

B. BGP in Guifi

The other community network that had its BGP traffic monitored is Guifi in Spain, which is slightly larger than AWMN. Again, the data from one month, in this case June 2014, is analysed. The same graphs are presented, except for the introductory figures 1 and 2.

Comparing the maximum number of announcements and withdrawals per day between AWMN (figure 3) and Guifi (figure 7), it is clear that both network behave similarly: the Guifi network also shows an unpredictable, unstable maximum amount of updates per day. Interestingly, the number of withdrawals is significantly higher in Guifi than in AWMN.

When the statistics per AS are compared, some interesting differences can be observed. In the Guifi network, the maxi-

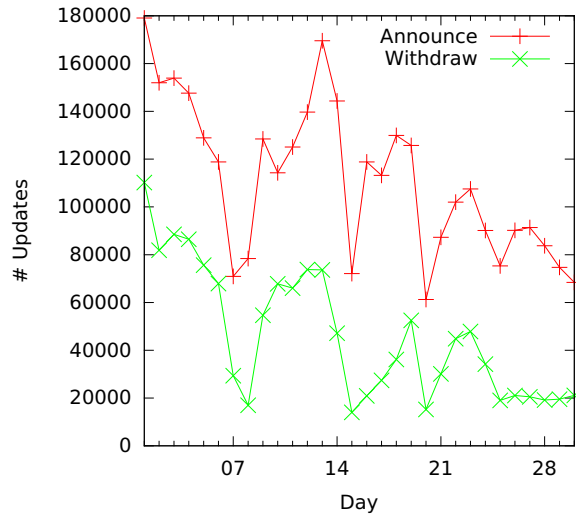


Fig. 7. Total updates and withdrawals per day in Guifi for June 2014

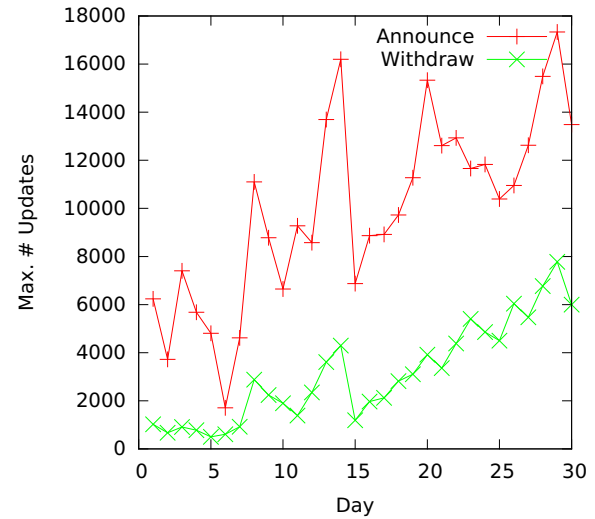


Fig. 9. Maximum number of updates and withdrawals per day in Guifi for June 2014

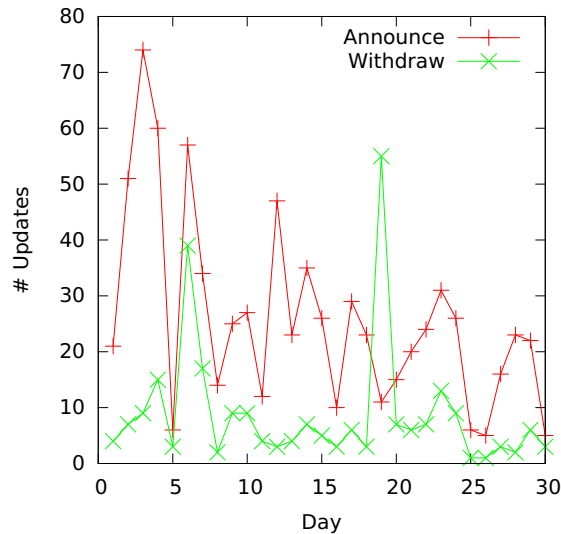


Fig. 8. Median number of updates and withdrawals per day in Guifi for June 2014

mum number of updates per AS per day, as shown in figure 9, is higher than in AWMN (figure 5). In both AWMN and Guifi, some high spikes can be observed, but on average, the values are higher in Guifi. This is certainly the case when the number of withdrawals are compared. On the other hand, if the median number of updates per AS per day are compared (in figures 4 and 8), the values in Guifi are approximately 100 times lower. This is an indication that, generally speaking, the Guifi network is mostly very stable, but does contain unstable sections that generate a disproportionately large amount of announcements and withdrawals. Finally, when the hourly values averaged over a month are compared, as shown in figures 6 and 10, no significant differences can be observed. Both networks follow a similar, weak diurnal pattern.

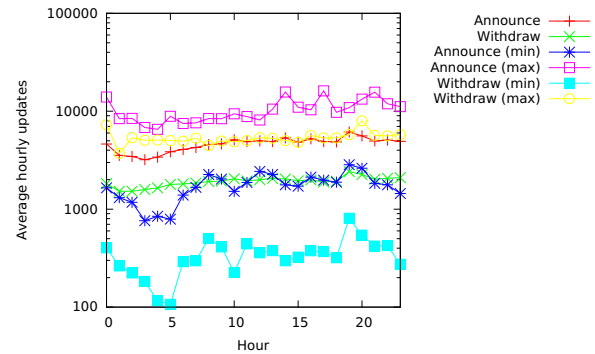


Fig. 10. Monthly average number of updates and withdrawals per hour in Guifi for November 2013

C. RIB

The number of announcements and withdrawals does not give a direct indication on whether the network is growing or shrinking: a large number of updates and withdrawals might just indicate that a certain, possibly small, part of the network is behaving erratically and continuously disconnects and reconnects to the rest of the network. In order to verify the stability of the overall size of the network, the size of the RIB table in the AWMN network in the observed month, November 2013, is plotted in figure 11. In this graph, an entry represents a RIB table size snapshot, taken with an interval of 2 hours.

From the graph, it can be seen that the size of the RIB is almost constant, varying a little in time. The large dip present in the graph at the 6th and 7th of November corresponds to the large spike in the same period in figure 3. This indicates that around this period, a (portion of) the AWMN network had some stability problems and the affected nodes were effectively no longer present in the RIB table. It is interesting to note that the spikes in figure 3 near the end of the month do not translate in large changes in the RIB table size. This is probably due




Fig. 11. AWMN RIB table size per day for November 2013.

to the fact that the increased number of announcements is compensated for by a simultaneous increase in the number of withdrawals.

Similar data on the RIB table size for the Guifi network were unfortunately not available for the observed month (June 2014).

D. Comparing Community networks and the Internet

To analyse the similarities and differences between community networks and autonomous systems in the public Internet, we considered the data from Potaroo.

Figure 12 shows hourly average update prefix rate per second, as measured by Potaroo for the public Internet for a week in May 2015. For comparison, based on the data from AWMN and Guifi, we derived similar data shown in figure 13 and figure 14 respectively. For the public Internet the average withdrawal rate is between 0.1 and 1 per second, for Guifi it is around 1 per second and for AWMN this same rate shows more variation and is closer to 0.1 on average. For the update rate we see that the Internet shows a rate of 1 to 10 per second, while for Guifi this rate is closer to 1. Again, AWMN shows larger variation.

To consider these numbers relative to the number of involved AS'es, first we have to estimate the number of autonomous systems in all three cases. The number of autonomous systems on the Internet as counted by Potaroo is about 80.000⁵. For AWMN no clear numbers are available. The AWMN WIND Node Database⁶ counts about 12000 active nodes and about 1200 backbone nodes. Although the difference between the two types is hard to derive and this data is known to be outdated, it is certain that the total number of AS numbers in AWMN will be much smaller than this in the Internet. A safe bet seems about 8000 nodes, or a factor ten less than in the public Internet. From the aforementioned Guifi growth statistics we can derive that about 25000 nodes were present in the network during the measured period.

With these numbers, a direct comparison between the rates reported in figure 12 for the public Internet, figure 13 for AWMN and figure 14 for Guifi becomes possible. This shows that relative to the network size and based on the updates, on




Fig. 12. Hourly average update prefix rate per second for public Internet




Fig. 13. Hourly average update prefix rate per second for AWMN

average AWMN has a stability which is comparable to this of the public Internet. Guifi is even a bit more stable. However, when looking at the withdrawal rates, these are comparable on average for the public Internet and AWMN (for a factor 10 less nodes) and even higher for Guifi (for a factor 3 less nodes). We consider this an important indication for higher instability in community networks.

When considering the peak prefix update rate per second as depicted in figures 15, 16 and 17 for respectively the public Internet, AWMN and Guifi, the withdrawal rates are surprisingly similar for the public Internet and AWMN, for Guifi the averages are higher. On the public Internet the peak prefix update rate is significantly larger, which is unexpected given the total number of autonomous systems. When comparing the overall behaviour, the peak rate is similar. We currently do not have a full explanation for this, we can only guess that this is




Fig. 14. Hourly average update prefix rate per second for Guifi

⁵See <http://www.potaroo.net/tools/asn32/>

⁶<http://wind.awmn.net/>




Fig. 15. Peak prefix update rate per second in the public Internet




Fig. 16. Peak prefix update rate per second in AWMN

caused by the consideration of only peak rates.

V. CONCLUSIONS AND FUTURE WORK

Community networks are formed by individuals connecting to an existing network, by pointing one or multiple antennas on their roofs. There is a similarity to starting an ISP where you begin peering with other (transit) providers, however at a different scale and cost. More interesting for this paper, in both cases you will start with a single AS number in which you announce your subnet over BGP.

From the initial data on only the community networks we can conclude that a strong variation is present, with a small number of nodes causing instability in the networks. A weak diurnal pattern in the data can be observed.

When comparing the community networks data to the




Fig. 17. Peak prefix update rate per second in Guifi

public Internet however, it is clear that in community networks the update rates and withdrawal rates per node are higher. The ratio of the update rates over the withdrawal rates is larger for the public Internet, due to the larger number of routes.

In general, we can conclude that to a certain degree community networks behave as the public Internet and could well be the Future Internet, with a higher degree of instability. This forms an important challenge to tackle with future research.

For this study of BGP behaviour, we plan to expand this analysis beyond the data from AWMN and Guifi, considering similar data from the Wireless Belgium community network. Moreover, we want to extend this analysis to more data from AWMN and Guifi itself, for more recent months and with more data on e.g. the network size. Finally, we believe comparing these results to data from OLSR-based community networks can help form a more general image about community network stability.

ACKNOWLEDGMENT

This work is supported by the CONFINE Integrated Project 288535. The authors would like to thank the Guifi and AWMN community networks for access to the BGP dumps.

REFERENCES

- [1] B. Braem, C. Blondia, C. Barz, H. Rogge, F. Freitag, L. Navarro, J. Bonicioli, S. Papathanasiou, P. Eschrich, R. Baig Viñas *et al.*, "A case for research with and on community networks," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 3, pp. 68–73, 2013.
- [2] D. Vega, L. Cerda-Alabern, L. Navarro, and R. Meseguer, "Topology patterns of a community network: Guifi. net," in *Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2012 IEEE 8th International Conference on. IEEE, 2012, pp. 612–619.
- [3] P. Millan, C. Molina, E. Medina, D. Vega, R. Meseguer, B. Braem, and C. Blondia, "Tracking and predicting link quality in wireless community networks," in *Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2014 IEEE 10th International Conference on. IEEE, 2014, pp. 239–244.
- [4] J. Crowcroft, A. Wolisz, A. Sathiseelan, J. Crowcroft, A. Wolisz, and A. Sathiseelan, "Towards an affordable internet access for everyone: The quest for enabling universal service commitment (dagstuhl seminar 14471)," *Dagstuhl Reports*, vol. 4, pp. 78–1377.
- [5] J. Avonts, B. Braem, and C. Blondia, "A questionnaire based examination of community networks," in *Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2013 IEEE 9th International Conference on. IEEE, 2013, pp. 8–15.
- [6] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "Bgp routing stability of popular destinations," in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*. ACM, 2002, pp. 197–202.
- [7] G. Huston, "Analyzing the internet's bgp routing table," *The Internet Protocol Journal*, vol. 4, no. 1, pp. 2–15, 2001.
- [8] R. Sami, M. Schapira, and A. Zohar, "Searching for stability in interdomain routing," in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 549–557.
- [9] Y. Rekhter and T. Li, "A border gateway protocol 4 (bgp-4)," 1995.
- [10] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding bgp mis-configuration," in *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4. ACM, 2002, pp. 3–16.
- [11] R. White, D. McPherson, and S. Sangli, *Practical Bgp*. Addison Wesley Longman Publishing Co., Inc., 2004.
- [12] G. Huston, "Potaroo.net," 2009.