Understanding the Internet Backbone in the Circumpolar North with Looking Glass

Moro Bamber Mary Kollander Asia Brooks Jedidiah Pollard Pradeeban Kathiravelu
University of Alaska Anchorage, Anchorage, AK, 99508, USA
{mwbamber, mgkollander, abrooks14, mjpollard, pkathiravelu}@alaska.edu

***Abstract*—Understanding how the Internet's major Autonomous Systems (ASes) interconnect and interact can provide exciting insights into resolving the inequity in Internet access, providing better connectivity to everyone. Looking Glass servers function as a real-time source of routing and information on Border Gateway Protocol (BGP), the exterior gateway protocol of the Internet. They provide a snapshot of the performance of these ASes.**

This paper presents a historical timeline and background of the Internet ecosystem and then looks into the backbone routing of the Internet through Looking Glass servers. We developed an easy-to-use Python toolkit to understand the Internet backbone and use our toolkit in the context of the circumpolar north, specifically Alaska and Canadian Territories. We discuss our findings and observations from two experiments. The first looks at popular websites and the ASes they exist in, and the second is an in-depth look at the BGP routing to the University of Alaska AS. The toolkit and the findings of our experiments show how Internet connectivity is across various ASes.

***Index Terms*—Border Gateway Protocol (BGP), Looking Glass**

I. Introduction

Understanding how the data is transferred across the Internet is crucial for Internet performance, especially in the remote regions of the Internet, such as the Circumpolar North. Large organizations on the Internet, such as the Internet Service Providers (ISPs), form the Internet Backbone. These organizations are vast, often globe-spanning networks of Autonomous Systems (ASes). NTT, Lumen, and AT&T are such entities, often titled Tier 1 ISPs [1]. The Internet uses Border Gateway Protocol (BGP) [2] to enable data transfer across multiple ASes [3]. To determine the best route to a prefix, BGP has several attributes describing its best path. Chief of these attributes is local preference, a policy decision made by the system administrator. After local preference, BGP uses the shortest AS-PATH, or fewer hops, to reach the destination. BGP allows ISPs to configure their network to their liking. Problems can arise like a prefix route being described as going through a rival ISP, and while this could be a viable route, an ISP might not want to send traffic through that ISP when they can choose a more favorable route [4].

ISPs transfer messages between each other at Internet Exchange Points (IXPs) [5] that are on neutral ground. IXPs are hubs within the global network connectivity. They are physical locations where ISPs and other ASes directly exchange Internet traffic. IXPs do not prefer one Tier 1 ISP over another; they simply facilitate cooperation. Due to the amount of traffic they carry, IXPs use the fastest routers on the Internet. The primary purpose of IXPs is to reduce the need for lengthy data transmission routes through transit providers, therefore creating better efficiency and reliability of the global Internet. Several large IXPs, such as the Equinix Exchange, London Internet Exchange, and NetIX, have emerged as key players in major Internet hubs. Each serves as a focal point for exchanging large amounts of internet traffic. Non-profit organizations or private companies typically facilitate the establishment of IXPs. However, governments and large industries also play a role in encouraging the development of new IXPs.

To understand the Internet performance of a region, one must understand the Internet's architecture and the region's unique challenges under consideration. Internet peering is the foundation of the modern landscape of the Internet. Therefore, understanding internet peering is essential for those interested in computer networking. Developed nations usually have faster Internet with widespread access. Distant countries sadly have high latency and low throughput due to them having to route Internet traffic through major IXPs that tend to be geographically distant from these countries. On the other hand, large regions such as Alaska, Canadian Territories, and Greenland are often overlooked in Internet performance studies as they are grouped with the developed countries with low latency. However, their Arctic landscape and low population density have caused them to have poor Internet access, a unique case that deserves detailed studies in the local and international context.

This paper aims to address this identified gap in the research by finding connections from historical developments, developing a framework that learns the Internet backbone, and thus providing a better understanding of the Internet backbone in the circumpolar north. Our research aims to better understand BGP backbone routing through Looking Glass [6] servers and Autonomous System (AS) peering, focusing on Alaska and Canadian Territories. To understand the internet backbone, one must understand BGP, and to see BGP in action, one can use Looking Glass. A server running Looking Glass software is a real-time routing and BGP information source. ISPs use Looking Glass to know how different servers are reached. We

explore the internet backbone and BGP with two experiments. We use the term upstream to refer to the immediate AS before the destination AS in a BGP prefix. The main contributions of this paper are:

- (C_1) A comprehensive analysis of Internet peering, the voluntary interconnection of diverse networks among large and small Internet Service Providers (ISPs).
- (C_2) An exploration of the historical evolution of the Internet, diving into the opportunities and challenges that paved the way for the impactful invention and adoption of Internet peering.
- (C_3) Understanding the technical intricacies of Internet peering and the economic models driving strategies employed by ISPs.
- (C_4) *BGP python AS Lookup*[1], an open-source tool to analyze the Internet backbone, using Lumen's Looking Glass.
- (C_5) A detailed evaluation of popular websites as test subjects to see what AS it was a part of and what is upstream of this AS.
- (C_6) Studying the Circumpolar North networks, using the University of Alaska AS as a sample, and cross-referencing the BGP routes it advertises against what we already know about its upstream.

This paper presents our *BGP python AS Lookup* framework in detail. Section II provides a historical background and details the current status of the Internet ecosystem in general. Section III presents our approach to understanding the Internet backbone with Looking Glass. Section IV shows the framework in action and evaluates its performance at measuring the Internet behavior in the Circumpolar North. Finally, Section V concludes the paper with a discussion of our findings.

II. Background

Internet peering is a fundamental part of global network infrastructure, facilitating data exchange between ISPs and other ASes of the Internet. It allows networks to exchange traffic without relying on third parties [7], with this data exchange between interconnected networks occurring at IXPs. It enhances the efficiency of the Internet by reducing latency and lowering costs associated with data transmission. When ISPs and networks peer, direct connections enable traffic exchange more efficiently than when routed through multiple intermediaries. This improves the speed of data transfer and creates a resilient internet infrastructure. Furthermore, Internet peering promotes a decentralized and interconnected Internet ecosystem where entities can cooperate effectively.

A. Historic Perspective

ARPANet [8] was the first operational, wide-area network developed by and for the Department of Defense organization Advanced Research Projects Agency, ARPA, now more commonly known as DARPA. The network was built to facilitate the transmission of messages and data between research colleagues and government workers in disparate locations. It is the foundation of the Internet as we know it. The first message ever sent on ARPANet was conducted between an SSD computer at UCLA and another computer at Stanford. ARPANet was declared operational in 1971.

While impressive, ARPANet's scope was limited to just some select research facilities, campus laboratories, and government installations included as nodes in the network, bound together by a standard directive in researching defense-oriented studies. Institutions outside this initiative saw the benefits of such a communications network. Still, they did not want to be bound by the scope and directive of ARPANet necessarily. These various pressures and incentives drove the push for the creation of the Computer Science Network, or CSNet [9]. In January 1981, the NSF awarded a $5 million to fund to expand a network across these networks [9]. In the years to come, CSNet would expand from only three sites in 1981 to as many as 180 institutions, not just limited to the United States but with connections to computer science departments across several developed nations.

Architects aimed to create a simple model that would necessitate traffic flow between ARPANet and CSNet without heavy administrative policies complicating or obfuscating the process. This was officially the start of peering [7]. With the deployment of CSNet and early peering efforts made between this new network and the existing ARPANet, the NSF sought to foster even more interconnection between research facilities across the United States. The NSF had in 1985 funded the development of five supercomputing centers on five different campuses: Princeton, Cornell, University of Pittsburgh, U.I. Urbana-Champaign, and U.C. San Diego, with a sixth separate site at the National Center for Atmospheric Research [10]. They developed NSFNet [10], a backbone network to interconnect these six NSF-funded regional networks across the U.S. in 1986.

By the early 1990s, the NSF determined that the Internet should be autonomous, meaning it could operate independently from any funding or support from the government [11]. Meanwhile, as government partners in industry, research, and academia thrived using the NSFNet, commercial groups took notice of the benefits such a network afforded them and began to create commercial ISPs independently. Because they could not freely exchange traffic and peer with the NSFNet due to the NSFNet's overriding directives and restrictions, these businesses coalesced funding to create in 1991 the first commercial IXP, named Commercial Internet eXchange (CIX), with interconnection on a settlement-free peering basis [12].

NSF crafted an NSFNET Transition Plan that became a reality in 1994. NSFNet was partitioned into Network Service Providers (NSPs). Network Access Points (NAPs) were

[1]https://github.com/UnderYourSpell/BGP_python_AS_Lookup.

established to route traffic between these NSPs nationwide at critical nodes and locations. Regional networks, such as Sprint, PacBell, AADS, and eventually MAE-East, were designated priority NAPs. These regional networks were then encouraged to pay for internet access from these new NSPs with NAP access. Lastly, a special function called the Routing Arbiter was created to collect and propagate routing information across the NAPs [7]. The NSFNet was formally retired in April 1995 [13]. During this transition, industry experts and corporate representatives began to obfuscate any information coming out of their networks. Introducing competition meant these members were incentivized to hide information from industry peers. This created additional hurdles to identifying, diagnosing, or rectifying developing problems in this new Internet.

With packet loss becoming a notable issue at these NAPs, large ISPs migrated to only having private point-to-point connections with other large ISPs to avoid peering free with smaller ISPs among the NAPs [7]. However, the point-to-point connection model was expensive to prevent congestion. They were also inefficient, delivered 18 months late while the Internet traffic doubled each year [7]. Ultimately, driven by the economic perspectives, major ISPs decided on an Internet-exchange-based model to replace the point-to-point circuits. The major US-based ISPs enrolled into distinctly carrier-neutral IXPs across the U.S., which would then dominate the peering in the U.S. and worldwide [7].

B. Current Landscape

Large ISPs and ASes impact various aspects of the Internet ecosystem. Creating a decentralized network of direct connections between the two, Internet peering removes the need to rely on singular pathways. Not only does this make the Internet more resistant to disruptions, but it also helps maintain a continuous flow of data. Internet Peering also plays a vital role in enhancing global network redundancy. Since data can flow through multiple routes, there is a decreased risk of failure caused by a single point. However, Internet peering can influence internet connectivity depending on geographic location. Large ISPs and ASes engaged in peering may influence the establishment of IXPs in targeted locations. This means that governments and large industries may collaborate with ISPs to create new IXPs, strengthening the local internet infrastructure while contributing to the interconnected global network at the same time.

Internet peering allows the networks to hand off traffic between each other's customers without paying a third party to carry traffic across the Internet for them. There are 7.7 exabytes of traffic daily on the Internet, equivalent to 1 billion gigabytes. The Internet traffic consists of the data transmitted between users and servers or between servers. This can include web traffic, search engine queries, and email traffic. Companies opt to form peering agreements that allow them to retain control of routing paths and improve performance. Internet companies have assured security and performance as the economy has shifted to the Internet. Three approaches have been used primarily: (i) settlement-free peering, (ii) paid peering, and (iii) transit.

Settlement-free peering is an agreement where the traffic each network sends to the other is roughly equal. This is performed when both networks agree not to charge each other for the exchange in traffic. When engaging in Settlement-Free Peering, ISPs have requirements and expectations all parties must meet before traffic can be exchanged [14]. Large ISPs require a minimum of 6-8 locations from a predetermined list. The ratio of incoming traffic required is around 2:1; this is considered a roughly equal exchange for large ISPs. An example of this is from Netflix. In the past, Netflix's Open Connect Program [15] was a settlement-free arrangement where access providers agreed to house Netflix servers in their data centers or connect to Netflix at carrier-neutral IXPs. This agreement allows their customers direct access to Netflix's content library, bypassing third-party involvement.

Paid peering is done where there are imbalances in traffic, meaning a company has more leverage. A paid agreement can be established, consisting of a network paying the other for carrying its traffic. The payment amount is determined based on the volume and nature of the traffic exchanged. Companies with a worldwide presence usually engage in paid peering agreements with smaller ASes. Peering arrangements can be bilateral or multilateral. Networks may choose to peer to reduce the cost of transit services, especially if they have significant traffic to exchange with other networks [16]. However, if a network is smaller or serves less densely populated regions, it may rely on transit providers to access the broader Internet. The size and influence of the networks involved play a role in the dynamics. There are regulatory bodies, such as the Federal Communications Commission (FCC), European Commission, and Autorité de Régulation des Communications Électroniques, des Postes et de la Distribution de la Presse (ARCEP, in France), that monitor and shape peering agreements and interconnections. Many countries have set guidelines that impact peering agreements [17].

All these historical design decisions and the fast-evolving Internet landscape impact any region's current Internet performance. In the following sections, we will use the Circumpolar North as the sample location for our analysis due to its unique nature.

III. Approach

Several frameworks that can monitor the Internet performance and interconnections in real-time and historically are available. These frameworks allow us to study how local ISPs interconnect to access content and connect users to the broader Internet. Knowing the history of the Internet and the intentions of the ASes helps understand the local Internet performance better.

We evaluate the Internet in the Circumpolar North using several tools as part of our methodology. Looking Glass servers

are beneficial for network administrators to see if the prefixes they are advertising are reachable. We also use the website bgp.tools to find what BGP prefixes specific IP addresses were advertising and information about ASes. Despite this array of tools, no framework supports displaying particular details on an AS to be put into a database for analysis. Therefore, we developed *BGP python AS Lookup*, leveraging data from the BGPView API[2]. *BGP python AS Lookup* features a simple user interface (UI) that allows users to enter an ASN or an IPv4 address, as shown by Figure 1. If an IP is entered, the ASN that the IPv4 address is in is displayed. The user can then enter that ASN to see a multitude of information about the AS, including the upstream peers.

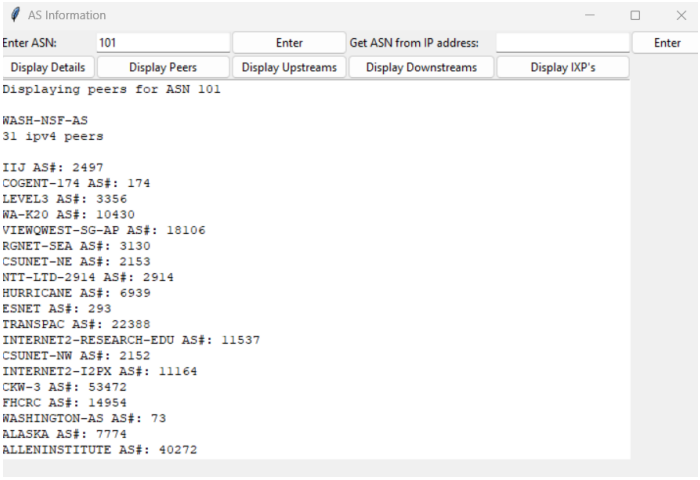

(a) Display of peers.

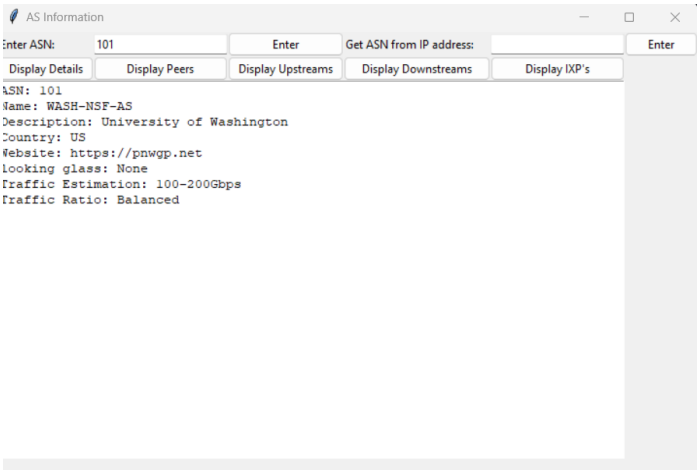

(b) AS information.

Fig. 1: *BGP python AS Lookup* UI.

A repository found on Cisco Developer[3] inspired the development of *BGP python AS Lookup*. The app works by appending the entered Autonomous System Number (ASN) [18] to a link to the API, which then responds with JSON [19] data consisting of the necessary information about the BGP paths of the AS. It then parses through the JSON data for the specific information each section of information may need for the given AS. The information is stored in object variables so the app can store more than one AS's data. It uses string concatenation to display the result in plaintext format. We develop *BGP python AS Lookup* in Python and use TKinter, a built-in Python graphics library [20], to build its user interface. This application proved worthwhile as it expedites the data-gathering process to understand the network performance in Circumpolar North and elsewhere throughout our evaluations.

[2]https://bgpview.docs.apiary.io/.

[3]https://github.com/pyjoepy06/bgp_python_path_lookup.

IV. Evaluation

We evaluate our *BGP python AS Lookup* framework for its functionality. We use *BGP python AS Lookup* to find i) the ASes of 18 websites, including Facebook, GitHub, and the University of Alaska (UA), using their IPv4 addresses and ii) the upstream ASes of these systems. Figure 2 compiles the peers into a chart detailing how many times an AS occurred upstream of the ASes we looked at. The sample size was small to prevent a significant saturation of ASes unique to a particular AS from showing up.

id	website	IP Address	managed by	AS#	Peers	Upstream
1	microsoft.com	20.76.201.171	Microsoft	8075	42	7
2	alaska.edu	137.229.114.150	U of Alaska	7774	2	2
3	weather.gov	104.117.232.18	Akami	16625	14	12
4	twitter.com	104.244.42.65	Twitter Inc	13414	24	6
5	georgefm.co.nz	3.163.189.14	Amazon	16509	61	20
6	australia.gov.au	13.107.213.70	Microsoft	8075	42	7
7	last.fm	34.96.123.111	Google Cloud	396982	2	1
8	fis-ski.com	151.101.2.217	Fastly.com	54113	59	36
9	GitHub.com	192.30.255.113	GitHub Inc.	36459	11	4
10	bgp.tools	185.230.223.150	Ben Cartwright Cox	206924	11	4
11	play.max.com/	13.224.14.120	Amazon	16509	61	20
12	www.netflix.com	44.242.60.85	Amazon	16509	61	20
13	wolframalpha.com	140.177.8.192	Wolfram AS	11106	2	2
14	www.bundeskanzleramt.gv.at	85.158.224.156	Bundesrechenzentrum GmbH	8692	3	1
15	p4.org/ecosystem/	138.68.30.104	DigitalOcean, LLC	14061	29	9
16	getpocket.com	3.163.189.94	Amazon	16509	61	20
17	linkedin.com	13.107.42.14	Microsoft	8068	1	1
18	facebook.com	157.240.3.35	facebook	32934	51	14

Fig. 2: Results from website upstream AS collection.

Our observations highlight what ASes are most commonly upstream of the AS that popular websites advertise their prefix through. The most common upstream peer for all websites is NTT (Nippon Telegraph and Telephone) AS2914 with 11 occurrences; NTT is Japan's fifth largest publicly traded company. Tied at ten occurrences, each is LEVEL3 AS 3356 and TWELVE99 AS 1299.AS3356 is Lumen, formerly Level 3 Parent LLC, ranked one by cone size (number of direct or indirect customers) worldwide. AS1299 is Arelion Sweden AB,

ranked two in cone size worldwide (AS Rank Caida). The other more frequently occurring players are a who's who of massive ISPs, all with impressive peering capabilities.

Using the website bgp.tools, we look up BGP sessions for the UA network AS7774. We wanted to see the BGP sessions for UA - Alaska-AS 7774 - and cross-reference them with different Looking Glass server traceroutes. By comparing the actual BGP session to traceroutes from different ASes, the route advertising done by BGP becomes clearer. For example, based on the upstreams of Alaska AS (AS7774 bgp.tools), we know that a route can pass through AS101 or AS7782 before getting to the UA system. The path before AS101 and AS7782 is determined by what is upstream of either of those two ASes. From looking at what is upstream of both these ASes, we can understand the BGP routes AS7774 advertises. AS7782, or Alaska Communications, has upstreams AS6461 Zayo Bandwidth and AS6969 Hurricane Electric. These two ASes are some of the largest in the world, but AS101's upstreams are also Tier 1 ISPs. Figure 3 presents the results visually.

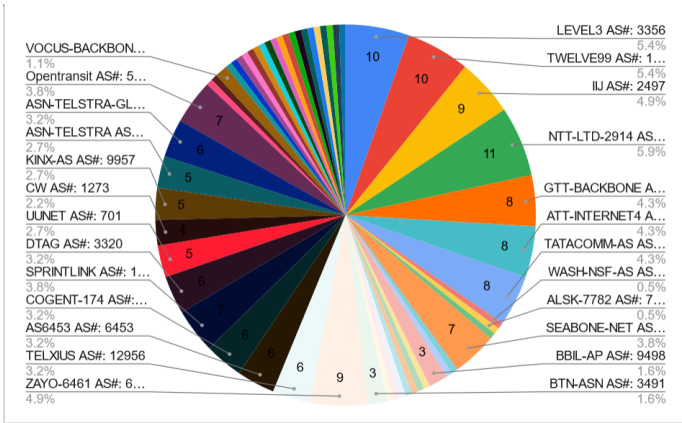

Fig. 3: Visual representation of the results.

Example of Looking Glass of Lumen AS3365 to Alaska AS7774:

```
Tracing route to 137.229.114.150
1 ae2.3602.ear3.Seattle1.level3.net (4.69.203.169) 0
    ms 0.0.0.0 (0.0.0.0) * *
2 4.2.138.138 (4.2.138.138) 0ms 0ms 2ms
3 ae17--201.icar-sttl1-3.infra.pnw-gigapop.net
    (209.124.188.216) 1ms 5ms 1ms
4 209.124.188.217 (209.124.188.217) 40ms 40ms 40ms
5 swf-mx480-1.ne.alaska.edu (137.229.255.233) 40ms 40
    ms 40ms
```

An example BGP announcement from the prefix in AS 7774:

```
Type: BGP
BGP.as\_path: 8426 3356 101 101 101 101 101 7774
BGP.community: (101,11103) (101,11303) (101,20000)
    (101,24100) [AS3356: North America] (3356,22)
    (3356,100) [AS3356: Customer route] [AS3356: USA
    ] (3356,903) [AS3356: SEA1 - Seattle] (8975,515)
     (8975,599) [AS34307: Prepend 1 time to all
    peers]
```

Hurricane Electric's (HE) Looking Glass shows that before the route gets to the UA AS, it passes through Alaska Communications AS 7782. The best route for getting from HE AS6939 to AS7774 is via 7782. This is because HE is upstream of Alaska Communications but not upstream of the University of Washington. Using BGP.tools' Super Looking Glass (Figure 4), we can see the BGP routes advertised by the IP 137.229.0.0/16 in AS 7774. Most BGP routes are advertised with AS6939 and AS7782. Of the 1253 paths advertised and shown on BGP.tools' Super Looking Glass for querying BGP sessions, only 41 hops from AS101 to AS7774.

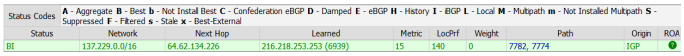

Fig. 4: HE Looking Glass.

BGP python AS Lookup, in conjunction with these freely available tools, allowed us to evaluate and study the network path on the Internet to reach out to the ASes in Alaska. Such tools help the network operators and enthusiasts to understand and appreciate network interconnections.

V. Discussion

Backbone routing through Looking Glass servers and AS peering is a vast topic to study. Further research is needed to uncover aspects of the internet backbone that should be studied more in the context of Alaska and Northern Canada. One must understand how ASes peer with each other to understand the Internet. To know how ASes peer with each other, one must understand how BGP allows an AS to advertise routes to IP prefixes on their system. The design decisions of the Internet have been made for historical reasons, as we discussed earlier. We explored the basics and evolution of Internet peering and how it has shaped the present-day economic models in charge of network data exchange. We gained an understanding of BGP and AS peering, what that looks like, and how it happens. Then, we devised experiments to identify trends in this aspect of the Internet.

We found that many direct upstreams of popular websites and the AS they exist in are Tier 1 ISPs. We also observed that peering between an AS and others can be affected by who may be directly upstream. For the case of Alaska AS7774, AS7782 had 1212/1253 of advertised routes for a prefix in AS7774. This is because of Alaska Communications' upstreams HE AS6939 and Tier 1 ISP Zayo AS6461. The HE upstream is critical as HE advertises an astonishing amount of BGP routes, and ranked by bgp.tools as the top 1 for known peers globally. The Internet backbone and BGP are critical to our modern lives, although the end users overlook them since they do not directly interact with domestic customers.

We executed our initial experiments from Alaska and then further expanded the scope to include the other regions of the Circumpolar North, consisting of the Arctic and Sub-Arctic. Northern Canada has some striking similarities and

considerable differences from Alaska. The Canadian territories are vast expanses of wilderness with little infrastructure connecting them to the rest of the country. Connecting to the World Wide Web is often difficult for these regions. For populated areas such as Whitehorse/YT and Yellowknife/NW, Internet connection options are controlled by just one ISP, Northwestel. Northwestel is downstream of Bell Canada, one of the two major ISPs in Canada. Peers with Northwestel include the governments of both the Yukon and Northwest Territories. A submarine cable owned by Quintillion lands at Iqaluit, the capital of Nunavut, where Northwestel and Bell provide internet access. Interestingly, a peer of Northwestel is a company called OneWeb AS800. OneWeb is a program by Eutelsat that brings the Internet to remote regions via satellites. Most Northern Canadian and rural Alaskan communities rely on satellite internet, and that option will only improve in the coming years.

Using the *BGP python AS Lookup* helped us gain insights into certain exciting behaviors of the ASes, specifically the tier-1 ISPs. For example, HE AS6939 advertises 180k BGP routes on most internet exchanges, and the paths it advertises are preferred over transit. So, when using sites that collect data from RIPE RIS and Routeviews - the two major BGP route collectors - one must understand that there is a bias toward HE BGP routes. By doing this, HE hides transit paths from the BGP route collector (NANOG).

Understanding BGP prefixes will help understand the users of a website and the Internet backbone as a whole. For example, our research shows that AWS or Amazon AS16509 originate many BGP prefixes. This was also noticed by the administrator of bgp.tools because it was taxing their ability to fetch real-time data. This is because AWS is getting a head start on prefix hijacks. We specifically looked into Alaska and Canada as part of this research. We use popular websites as test subjects to see what AS it was a part of and what is upstream of this AS. We also examined the University of Alaska AS and cross-referenced the BGP routes it advertised against what we already knew about its upstreams. As a future work, we aim to expand the development of *BGP python AS Lookup* and look into other regions of the Circumpolar North.

Acknowledgment: This work is supported by the Elizabeth Tower Endowment for Canadian Studies.

REFERENCES

[1] M. Winther, "Tier 1 isps: What they are and why they are important," *IDC White Paper*, pp. 1–13, 2006.

[2] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol (s-bgp)," *IEEE Journal on Selected areas in Communications*, vol. 18, no. 4, pp. 582–592, 2000.

[3] I. Van Beijnum, *BGP: Building reliable networks with the Border Gateway Protocol*. " O'Reilly Media, Inc.", 2002.

[4] J. W. Stewart III, *BGP4: inter-domain routing in the Internet*. Addison-Wesley Longman Publishing Co., Inc., 1998.

[5] J. C. Cardona Restrepo and R. Stanojevic, "A history of an internet exchange point," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 2, pp. 58–64, 2012.

[6] J. Welch, "Through the looking glass: Classifying anomalous bgp communities," Ph.D. dissertation, Monterey, CA; Naval Postgraduate School, 2020.

[7] W. B. Norton, *The 2014 Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2014.

[8] M. Hauben, "History of arpanet," *Site de l'Instituto Superior de Engenharia do Porto*, vol. 17, pp. 1–20, 2007.

[9] D. Comer, "The computer science research network csnet: a history and status report," *Communications of the ACM*, vol. 26, no. 10, pp. 747–753, 1983.

[10] D. L. Mills and H.-W. Braun, "The nsfnet backbone network," in *Proceedings of the ACM workshop on Frontiers in computer communications technology*, 1987, pp. 191–196.

[11] V. G. Cerf, "Thoughts on the national research and education network," Tech. Rep., 1990.

[12] S. Estrada, "Commercialization and the commercial internet exchange: How the cix can help further the commercialization of the internet," *Internet Research*, vol. 2, no. 3, pp. 24–28, 1992.

[13] S. R. Harris and E. Gerich, "Retiring the nsfnet backbone service: Chronicling the end of an era," *ConneXions*, vol. 10, no. 4, pp. 2–11, 1996.

[14] A. Nikkhah and S. Jordan, "Requirements of settlement-free peering policies," in *GLOBECOM 2022-2022 IEEE Global Communications Conference*. IEEE, 2022, pp. 3617–3622.

[15] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig, "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the netflix cdn," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 28–34, 2018.

[16] X. Wang, Y. Xu, and R. T. Ma, "Paid peering, settlement-free peering, or both?" *IEEE/ACM Transactions on Networking*, vol. 29, no. 2, pp. 585–594, 2021.

[17] A. Lodhi, A. Dhamdhere, C. Dovrolis *et al.*, "Analysis of peering strategy adoption by transit providers in the internet." in *INFOCOM Workshops*, 2012, p. 177.

[18] G. Huston, "Exploring autonomous system numbers," *The Internet Protocol Journal*, vol. 9, no. 1, pp. 2–23, 2006.

[19] F. Pezoa, J. L. Reutter, F. Suarez, M. Ugarte, and D. Vrgoč, "Foundations of json schema," in *Proceedings of the 25th international conference on World Wide Web*, 2016, pp. 263–273.

[20] F. Lundh, "An introduction to tkinter," *URL: www. pythonware. com/library/tkinter/introduction/index. htm*, 1999.