2020 12th International Conference on Cyber Conflict
20/20 Vision: The Next Decade
T. Jančárková, L. Lindström,
M. Signoretti, I. Tolga, G. Visky (Eds.)
2020 © NATO CCDCOE Publications, Tallinn

Measuring the Fragmentation of the Internet: The Case of the Border Gateway Protocol (BGP) During the Ukrainian Crisis

Frédérick Douzet
Professor
GEODE
University Paris 8
Saint-Denis, France
douzet@univ-paris8.fr

Louis Pétiniaud
PhD Candidate
GEODE
University Paris 8
Saint-Denis, France
l.petiniaud@gmail.com

Loqman Salamatian
GEODE
University Paris 8
Saint-Denis, France
salamatianloqman@gmail.com

Kevin Limonier
Associate Professor
GEODE
University Paris 8
Saint-Denis, France
Klimonier02@univ-paris8.fr

Kavé Salamatian
Professor
GEODE
University of Savoy
Annecy, France
kave.salamatian@gmail.com

Thibaut Alchus
GEODE
University Paris 8
Saint-Denis, France
thibaut.alchus@gmail.com

Abstract: This paper presents the results of a year-long research project conducted by GEODE (geode.science), a multidisciplinary team made up of geographers, computer scientists and area specialists.

We developed a new methodology for mapping cyberspace in its lower layers (infrastructures and routing protocols) in order to measure and represent the level of fragmentation of the Internet in areas of geopolitical tensions using the Border Gateway Protocol (BGP). Our hypothesis was that BGP could be used for geopolitical reasons in the context of a large-scale crisis, leading to a further fragmentation of the Internet. We focused on the Ukrainian crisis.

BGP is a core protocol of cyberspace that connects the tens of thousands of autonomous systems (ASes) that compose the Internet. Based on a 35-year-old technology, this protocol is easy to manipulate to re-route Internet traffic or even to cut off entire regions (BGP hijacks). Our results show actions on BGP implemented right after the 2014 Maidan Revolution, when Russian forces took control of the Crimean Peninsula and started to back separatist forces in Eastern Ukraine. In both cases, Russian authorities and separatist forces modified BGP routes in order to divert the local Internet traffic from continental Ukraine – drawing a kind of "digital frontline" consistent with the military one. The study of Donbass and of the Crimean Peninsula leads to important methodological findings to (1) define and map digital borders at the routing level; (2) analyze the strategies of actors conducting actions via BGP; (3) categorize these strategies, from traffic re-routing to cutting-off entire regions for intelligence or military purposes; and (4) anticipate future uses for BGP manipulations by identifying strategic bottlenecks within the network.

Keywords: *cyberspace, Ukraine, BGP, Russia, Crimea, Donbass, autonomous systems*

1. INTRODUCTION

On December 23, 2019, Russia claimed to have successfully tested disconnecting its network from the global Internet in an attempt to run a domestic alternative. Months earlier, the country had announced that it considered briefly unplugging itself from the Internet to test its cyber defense. This took place as the law n°608767-7 on the creation of a "sovereign Internet" came into force in November,[1] requiring technical alterations to provide Russia with the ability to control the Internet access points at its borders and to continue operating its domestic network in the event that it was disconnected from the global Internet.

These initiatives demonstrate the depth of Russia's strategic reflection on the structure of its connectivity and on the geopolitical importance of data routing. They are part

1 Federal Law n°608767-7 "On information, information technologies and information defense," https://sozd.duma.gov.ru/bill/608767-7.

of a larger strategy developed by Russia to secure sovereign control over what the authorities perceive as their national network, a geopolitical representation best captured by the term RuNet, widely adopted in Russia, and embodied by national platforms like Yandex or Vkontakte, to designate the post-Soviet linguistic, ethnic and cultural subspace of the web. The RuNet has since been used by Russian authorities to promote the representation of a sovereign cyberspace (Limonier 2018).

This strategy is not unprecedented. In November 2019, Iran accomplished just that when it cut off most traffic from the global Internet while operating its domestic network fully. The architecture of connectivity had been purposely redesigned to allow selective censorship of international traffic by connecting Iran's network to the outside with only three operators controlled by the government, thus creating a huge domestic intranet (Salamatian et al. 2019).

These initiatives have triggered concerns inside the Internet governance community about the increasing fragmentation of cyberspace and the risks it poses for its security and stability, not to mention online freedom and human rights. The question we ask in this paper is: "How can we measure and represent the fragmentation of cyberspace?" This paper presents the results of a year-long research project conducted by GEODE (geode.science), a multidisciplinary team composed of geographers, computer scientists and area specialists. We have developed a new methodology to map cyberspace in its lower layers (infrastructures and routing protocols) in order to measure and represent the fragmentation of the Internet in areas of geopolitical tensions using the Border Gateway Protocol (BGP).

Efforts to map cyberspace have focused on the physical infrastructure of the Internet, which is composed of cables, servers and other physical equipment that are grounded in physical territory and which can easily be mapped with the traditional tools of political and physical geography (Dodge and Kitchin 2001; Musiani et al. 2016). Other efforts have also attempted to capture the overall data traffic (Faravelon, Frénot, and Grumbach 2016). In the 2010s, much attention was given to the informational layer of cyberspace in the wake of jihadist propaganda and manipulations of information during democratic elections, leading to innovative cartographies of social networks and of the modes of content propagation (Howard et al. 2018; Limonier 2017). The strategic dimension of the BGP architecture and data routing, however, has been given much less attention in the scientific literature.

Jesse Sowell illustrates the importance of the lack of a top-down central governance model and the emergence of bottom-up governance models for groups of network operators – Internet exchange (IXP) groups – and other actors (Sowell 2012). This is made possible through the use of the BGP by autonomous systems (ASes) to establish

connections and exchange information between each other. BGP determines the routes data take and has been leveraged in the past by stakeholders to route traffic through specific paths and control the flow of information (Feamster and Ramachandra 2006). A relative flattening of the Internet structure has been observed, resulting from the emergence of major content providers like Netflix, Google, Amazon, Akamai, etc., along with major cable providers such as Angola Cables,[2] Me-We-Se,[3] and even Google, which owns 8.5% of Submarine Cables Worldwide (Zimmer 2018), that maintain a large part of the Internet traffic inside their networks (Wong 2016); however, BGP still has a primary role especially at the international level. It has also been manipulated by countries in order to block access to some content, to exclude some users from the Internet, to hijack traffic from other countries, or attack other countries' infrastructures. Many studies have focused on the inherent fragilities of a routing system designed in 1989 (Vervier, Thonnard, and Dacier 2015; Butler et al. 2010). Additionally, several articles have explored the BGP strategies of several nation-states (Edmundson et al. 2018; Wählslich et al. 2012).

Our hypothesis was that BGP could be manipulated for geopolitical reasons in the context of a large-scale crisis, leading to a further fragmentation of the Internet. We decided to focus on the Ukrainian crisis for several reasons.

First, the Ukrainian crisis presents a unique example of recent and direct military, economic, identity and diplomatic confrontation with Russia in the context of a major territorial conflict in Europe. Crimea and the two self-proclaimed republics of Donetsk and Luhansk in East Ukraine are spatial entities with disputed sovereignty sitting at the intersection of territorial and digital rivalries of power. In that sense, Crimea in particular can be perceived as a laboratory for Russia's strategies of appropriation.

Second, the anarchic development of the Internet in Russia and Ukraine has led to an abundance of ASes in both states, which provides larger sets of data with a greater level of precision. Finally, Russia has recently been testing methods to develop sovereign control of its network – particularly its physical infrastructure – through the re-nationalization of data networks, such as the obligation made in 2015 to maintain the data of Russian citizens in the country (Limonier 2018). But at the same time, Russia enjoys a very rich network with multiple external connections and its actors have been nurtured in the libertarian culture of Internet pioneers (Ermoshina and Musiani 2017). More recently, Russian authorities have asserted a need to organize the RuNet single-handedly.[4]

2 Angola Cables have emerged as an important actor in maritime Internet cables by providing direct links from Africa to South America.

3 Major maritime cables between Europe, the Middle-East and Asia.

4 "Совбез России поручил создать «независимый интернет» для стран БРИКС RBC," November 28, 2017, accessed March 9, 2020, https://www.rbc.ru/technology_and_media/28/11/2017/5a1c1db99a794783ba546aca.

This paper offers an overview of the topology of the Ukrainian network and its level of complexity in 2019. Then, through a longitudinal analysis of BGP data since 2013, it demonstrates the marginalization of Donbass and the appropriation of Crimea in cyberspace and raises the question of strategies of control these disputed territories have been subjected to, thus revealing the success and limits of Russia's venture for sovereign control in cyberspace.

2. METHODOLOGY

A. What is an Autonomous System?

The Internet is a network of networks characterized by its lack of centrality. It results from the interconnection of approximately 92,000 nodes (as of August 2019) called autonomous systems. An autonomous system (AS) is itself a network that manages its internal routing, distributes IP addresses to its customers and defines its access policies. Data transiting through the Internet from one point of the world to another usually crosses several independent ASes (6 on average) (Leguay et al. 2005).

Autonomous systems vary greatly in size and importance. A basic taxonomy divides them into three categories – Tier 1, 2 and 3 – which form an arborescent and partly hierarchical network structure. The most common types of Tier 1 ASes are intercontinental backbone carriers – such as Level 3 or Telia – or large national Internet Service Providers (ISP) – such as AT&T (United States), Orange (France), Rostelecom (Russia). Tier 2 ASes are generally medium-sized providers operating on regional or local scales. Tier 3 ASes (or "stub domains") are smaller networks run by a single company or university.

AS numbers, along with the blocks of IP addresses (or "prefixes"[5]) they manage, are allocated by the five Regional Internet Registries (RIR), themselves answering to the Internet Corporation for Assigned Names and Numbers (ICANN), one of the most important regulatory bodies of the Internet today. The administrator – either private or public – of each autonomous system determines a routing policy for its AS, which involves deciding which ASes to establish connections with and the behavior of its external routers when receiving data to be forwarded.

B. Why is BGP Political?

Notably, the security aspect of BGP routing in cases of traffic hijack, i.e., a redirecting of the traffic through malicious network nodes, has already stirred awareness of the political dimensions of routing.

However, BGP is political in ways that have not been investigated as much.

5 Set of several contiguous IP addresses that an AS can then assign to its users or customers.

First, an AS administrator wishing to connect to the global Internet has to establish relationships with other autonomous systems already connected to the network. The relationship can be of two types: a customer-to-provider (i.e., commercial) relationship, with an Internet Service Provider for instance; or a "peering," where two ASes estimate that they share approximately the same amount of traffic and set up a non-monetary relationship that allows their customers to exchange traffic. Moriano et al. analyzed the economic dimension of routing decisions (Moriano, Achar, and Camp 2016). Despite these choices being economic in nature, they also bear a political dimension.

Second, AS administrators implement routing algorithms that decide which path the packets of data will take to reach a destination, depending on commercial or security criteria, as well as on geopolitical considerations. When an AS receives information about a new possible path to reach a specific IP address, it chooses whether to change the path according to its preferences or keep the existing one. These routing policies integrate the basic rules of BGP along with the preferences set by AS administrators to create complex algorithms (Van Beijnam 2002).

Third, ASes contribute to the production of territories (Painter 2010). Through their routing policies, they define the paths and therefore the shapes of cyberspace. They are also implemented on physical territories and play a crucial role in providing places and people with Internet access and services, contributing to the development of territories. This is particularly true of remote places that rely on a limited number of ASes in order to access the global Internet, thus creating a digital territory defined by the topology of a network dependent on a few specific ASes. At the local level, the structure of ASes is critical to the resilience of the network (Chiu et al. 2015) and can result from spatial power strategies of various actors, a form of topological power (Allen 2011). The interconnection between states' ASes helps us understand how some countries might exert influence on others through connectivity and what relationships of dependency may exist.

Finally, BGP has been conceived of without security in mind and is very easy to manipulate for malicious or strategic purposes, such as espionage, censorship, disconnection, traffic hijack or the obfuscation of cyber attacks (Butler et al. 2010). Bearing the risk of observed BGP hijacks that could have resulted in threats of large-scale data exfiltration, Benton and Camp (2016) have proposed using BGP filters to ensure that packets are not being routed through problematic jurisdictions.

The strategic dimension of BGP deserves empirical studies. But mapping BGP data is a tremendous challenge because of the highly dynamic nature of this system. Routers can fail or restart. External connections between autonomous systems change at a very

fast pace and are announced through constant updates. For instance, an AS managed by the Russian company Vimpelcom (AS 8402) was found to have generated over 95,000 updates in seven days, which is not an extreme number. In addition, an autonomous system can change its information any time: the AS number can be reallocated, or the administrator can change its physical address and relocate to a different country. Relationships between ASes and routing policies are, for the most part, confidential and one challenge in measuring the Internet is to develop inference techniques to guess the policies of network operators and their relationships.

Despite these caveats, we were able to collect and process data to infer and map the topology of the Ukrainian network and its evolutions. Our approach is fundamentally interdisciplinary and involves research and fieldwork by regional specialists in geopolitics combined with the methodologies of computer science and mathematics.

C. What Data Did We Collect and Use?

Not all peering and customer-to-provider relationships are announced publicly. Our cartography is therefore mostly based on inference data, as opposed to data collected directly from operators. The AS relation graphs we infer are known to be incomplete. In particular, BGP path filtering policies do not expose less-preferred paths that would be chosen if the preferred announced paths were not available (Gregori et al. 2012). For this reason, we need to cross and combine our data with other sources (such as active measurements and IXP membership datasets) in order to obtain a consistent view of the network that can be mapped. Yet even this limited and incomplete view of the full AS graph is enough to monitor major changes to the Internet structure in Ukraine.

We have developed a BGP observatory that generates, every minute, a snapshot of a real-time AS graph that contains approximately 89,000 nodes and 220,000 links obtained by processing up to 30 BGP flows – announcing possible paths through a series of ASes – coming from different routers across the network. We have used the largest source of publicly available BGP routing data in 2019, RouteViews,[6] and the RIPE Routing Information Service (RIS),[7] which aggregates BGP messages from BGP monitors at cooperating ASes. These snapshots allow for a continuous monitoring of the logical layer of cyberspace at the AS level. We have collected more than ten terabytes of snapshots of AS graphs for a period of over three years. The AS graphs are inferred using path updates advertised by the routers running BGP to update neighboring routing tables (Roughan et al. 2011; Salamatian, Kaafar, and Salamatian 2018).

6 "Routeviews", accessed March 9, 2020, http://www.routeviews.org/routeviews.

7 "Routing Information Service – RIS," RIPE, accessed March 9, 2020, https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris.

We used several datasets and statistical methods:

1. We used graphs from the BGP observatory to represent the connectivity between the individual network operators (AS). Using that observatory, we augmented the BGP announcements by adding relevant information like 1) the name associated with the AS, 2) the country where the AS was registered, 3) the number of IP address prefixes announced by the AS and 4) the number of times a connection has appeared on the routing table.

2. We used the Potaroo blog to get statistics about the number of prefixes and ASes associated with each country year after year.[8]

3. We gathered the *AS relationships* (Dimitropoulos et al. 2007) inferred by the Caida Research Center at University of California San Diego, which indicate the underlying economic forces that drive the evolution of the Internet topology and its hierarchy.

4. In addition, we collected latency data using the Atlas network provided by RIPE, which allows any Internet user to install a probe on their server that can then be used by any other user to launch precise measures of connectivity.

Based on this data, our ambition was to study the topology of the networks and its consistence with the evolution of the topography of the country in a context of large-scale geopolitical crisis. The topological approach is highly valuable for approaching the reticular space of non-contiguous, enclave or exclave territories and the strategies of actors to reach this territory across space (Painter 2010; Latour 1987, 2005). By focusing on the crucial aspects of connectivity, such as data transits and network properties, the topological approach helps mobilize relevant concepts, such as accessibility, inclusion, borders, disjunction, continuity, intersection, connection and nodality (degree to which a node is the point of convergence between different routes) (Severo and Venturini 2016). In a nutshell, the topological approach is "first and foremost a reduction of complexity in the name of representing more complexity" (Piper 2013).

D. Limitations of Our Methodology

The BGP view is well-known to be incomplete. In particular, peer-to-peer (p2p) links are known to be harder to observe than customer-to-providers (c2p) links (Gao 2001; Ager et al. 2012; Cohen and Raz 2006). A contribution of this paper is to show that even this incomplete view provides valuable geopolitical insights. Moreover, c2p links reflect real economic strains and are therefore better indicators of the power relationships that shape the topology of the network.

8 "BGP Routing Table Analysis Reports", Houston G. Blog, accessed March 9, 2020, https://bgp.potaroo.net/.

Another shortcoming of BGP analysis is usually caused by the incompleteness of available information on AS owners stored in the Whois registry, as well as the unreliability of IP geolocation databases at the regional and local levels (Poese et al. 2011). In this work, we compensate for these shortcomings through a qualitative analysis, based on OSINT (Open Source Intelligence). We use various sources of information to find geographical data on the most important actors of Ukraine's connectivity and their relationships to policymakers.

In addition, our analysis is based on the routes available for data traffic and not on the quantification of the actual volume of traffic that circulates through these routes, as we cannot access this level of granularity in BGP data. However, we are able to evaluate the importance of a link through the number of announced BGP paths, and the number of BGP prefixes that cross it. Although these values do not precisely give the amount of traffic, it allows us to understand how central a link is for the overall routing. Moreover, we consider all the ASes to be nodes, despite their diversity (governmental, private, universities, geographically bounded to cities, etc.).

Last but not least, we need to acknowledge BGP's intense fungibility and lack of fixed relationships: on average, more than 5,000 route changes happen every second in the whole Internet. Most of them result from operational constraints (like a router rebooting), but some of them are also caused by relationship changes between ASes. This is the reason why we track data overtime in order to be able to provide longitudinal studies and avoid over-interpreting isolated instances of routing changes. Nevertheless, BGP is a highly dynamic environment and no cartography could possibly pretend to be fully accurate and exhaustive.

3. STRUCTURE OF CONNECTIVITY IN UKRAINE

A. A Rich and Diverse Network

We first looked at the architecture of Ukraine's ASes and the way they are connected to the rest of the world. Our first graph (Figure 1) represents, as of June 26, 2019, all Ukrainian ASes and their immediate neighbors, meaning ASes that have a direct relationship with at least one Ukrainian AS. Each node represents an AS, and each link a relationship (commercial or peering). For clarity, we eliminated from the graph ASes with fewer than five neighbors and provided the name of significant ASes only.[9] The nodes are colored according to the country the ASes are registered in.[10] Although this information is often reliable, it can hide part of the reality. Large ASes that operate

9 The names of the ASes (including quotation marks, numbers and capital letters) are based on the RIPE database. They are the official names of the autonomous systems. As such, the names are based on the decision of the administration of each AS, and do not always match the name of their parent company.

10 "List of country codes and RIRs," RIPE, accessed March 9, 2020, https://www.ripe.net/participate/member-support/list-of-members/list-of-country-codes-and-rirs.

internationally are likely to change their country of registration for political reasons, as we will see below, hence the need for qualitative research for graph analysis.

FIGURE 1. REPRESENTATION OF UKRAINIAN AUTONOMOUS SYSTEMS AND THEIR DIRECT NEIGHBORS, JUNE 2019

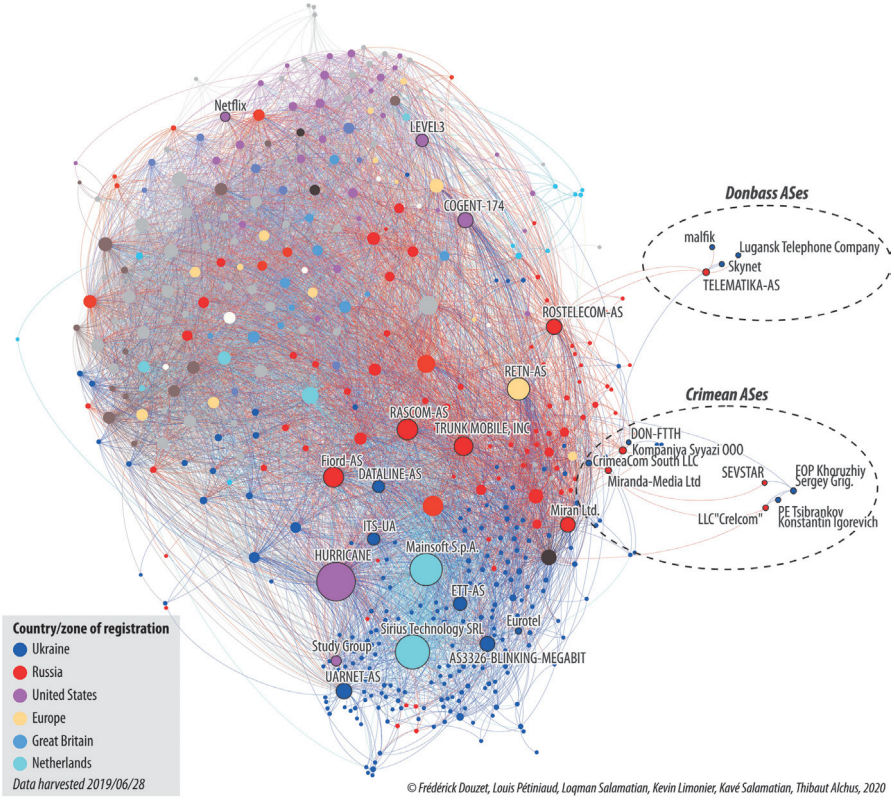

The size of the nodes (ASes) in the graph depends on their betweenness (the state of being between) centrality, i.e., the proportion of the shortest paths between all nodes of the graph that go through this link. The betweenness centrality measures the impact of disconnecting a link for the global connectivity of the network (Ma et al. 2008) and points to the most important nodes in the routing architecture of a country.

Finally, we used Force Atlas 2, a visualization algorithm for a representation of our graph. This algorithm is based on a concept of repulsion – with nodes pushing each other away but links attracting nodes closer – simulates the dynamics of a physical system to spatialize the network. In other words, the closer the nodes, the more connections they share.

The graph shows that Ukraine possesses a very rich network, with nearly 2,200 allocated ASes, among which over 1,800 are active (i.e., announced in our BGP data). This profusion is characteristic of Ukraine and Russia, which counts 5,176 active ASes. Both countries aggregate a high proportion of ASes compared to their population: about 24,300 users per AS on average in Ukraine and 30,485 in Russia, compared to about 110,000 users per AS on average in other European countries. However, most of the Ukrainian ASes are of small size on the graph, which reflects their low centrality in the network, i.e., the fact that they do not have many neighbors and therefore do not attract much of the traffic. Most of them are stub ASes (Tier 3) and serve a limited, sometimes very small, area.

This disproportion has historical roots and can be explained by the relative anarchy in which the Internet was developed in the post-Soviet republics during the 1990s and 2000s while European countries were structuring their network around major historical telecom operators, such as France Telecom in France. This profusion is reinforced by the competition between multiple economic actors with diverging interests in a rather opaque system controlled by oligarchs (Limonier 2018). It makes the network particularly resilient, but also complex and difficult to control, as we will see below.

B. The Polarization of Ukraine's Cyberspace Between Russian and Western Routes

The Ukrainian network is clearly structured around two poles: Russia on the one hand, the United States on the other hand, along with a myriad of other (mostly European) countries. Two Italian ASes (Sirius and Mainsoft) are highly visible due to their aggressive peering policy, but are less relevant when looking closely at the results. The structure of the network offers a great diversity of paths to the global Internet, but they are under the control of either Russian ASes (Rostelecom, Rascom-AS, Transtelecom) or major American or European ASes (GTT, Level 3, Cogent, Hurricane). Therefore, the architecture of the network reflects the geopolitical situation of Ukraine: split between major powers.

FIGURE 2. SIMPLIFIED REPRESENTATION OF UKRAINIAN ASES AND THEIR NEIGHBORS, JUNE 2019

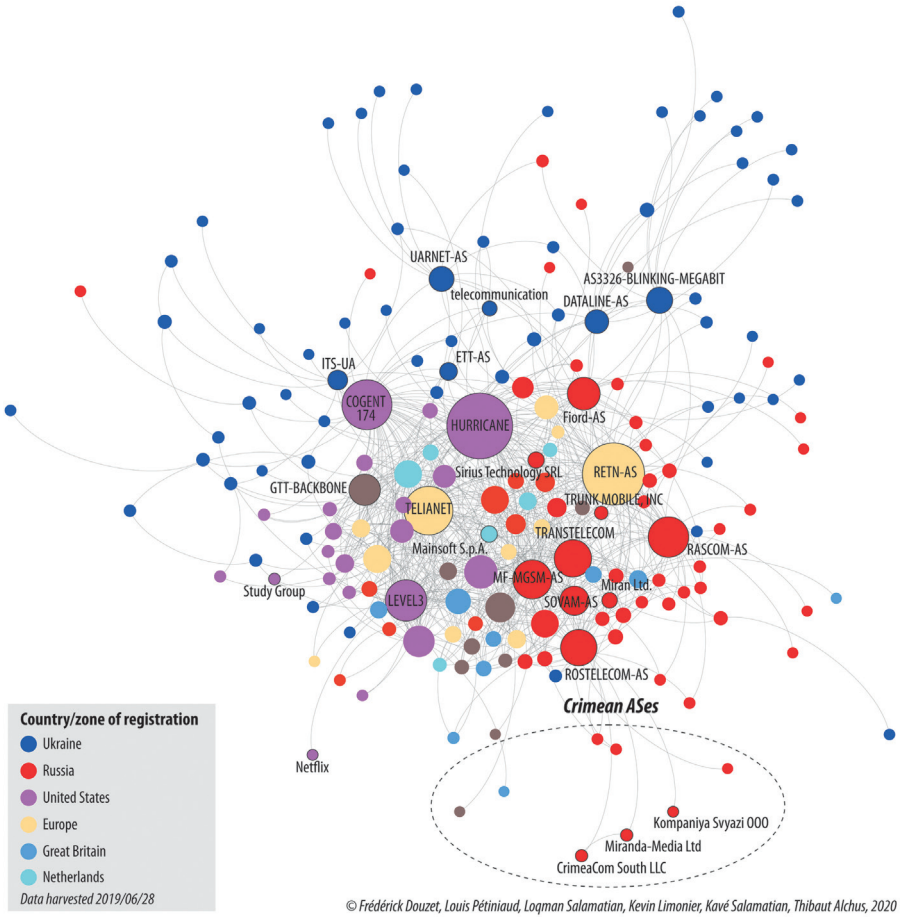

A simplified view of the network (Figure 2) gives a clearer understanding of who the major players are. In Figure 2, we only kept the Ukrainian ASes, their six most important neighbors and the links that appeared the most often in our routing table, thus eliminating 94% of the links, along with the less central ASes.[11] We can see that most Ukrainian ASes have disappeared due to their small size. Only the most important ASes remain in our graph, which are mainly foreign ones. The divide between the two poles is even more distinct.

The two Italian ASes are less central, which means that despite their many connections, they do not capture most of the traffic. They are fully integrated into the galaxy

11 We selected the top 6% of the links that appeared the most often in our routing table (i.e., more than 484 times).

of American and European ASes that connect the main Ukrainian ASes to smaller Ukrainian ASes and to foreign ASes of medium centrality. The UK, Germany and the Netherlands are important, yet usually not essential, points of transit. The place of RETN on the graph seems inconsistent, but is not surprising. Registered in Europe, RETN was once declared Ukrainian, but is currently administered by a major telecom company based in Saint-Petersburg; hence, the proximity to Russian ASes.[12]

FIGURE 3. DISTRIBUTION BY COUNTRY OF REGISTRATION OF UKRAINIAN ASES' NEIGHBORS, JUNE 2019

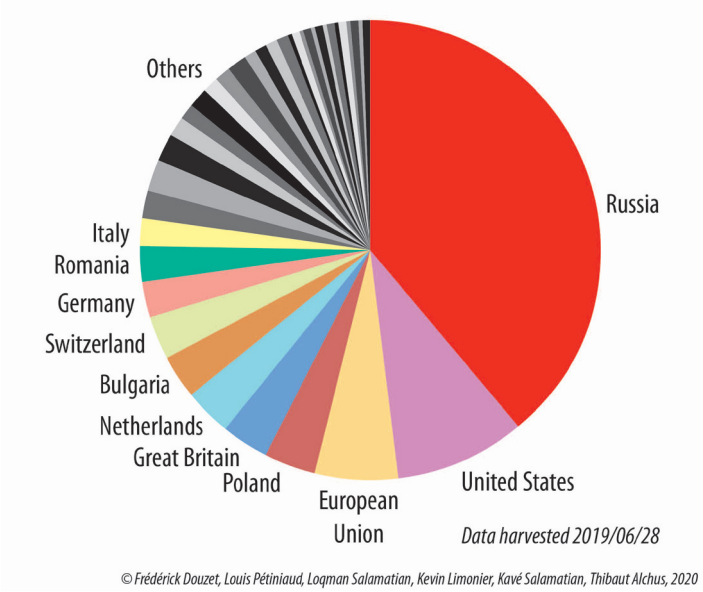

Overall, Russia occupies a major place in the graph, with 95 ASes connected to a Ukrainian AS, representing nearly 40% of all neighbor ASes (Figure 3). In comparison, the United States has only 22 ASes connected, but a number of them are major Tier 1 ASes. The American presence has strongly increased since June 2019 in our graph, with the direct connection of Hurricane Electric – a (near[13]) Tier 1 AS – to Ukrainian ASes. This observation might be explained by the strategic competition between the US and Russia over Ukraine, but could also be part of a wider phenomenon of centralization around major providers which are directly connected to smaller ASes without intermediaries. Some call this the "flattening of the Internet" (Böttger et al. 2019). It requires further investigation to confirm our hypotheses.

12 "RETN network map", RETN, accessed March 9, 2020, https://retn.net/networkmap/.

13 A Tier 1 network can reach every other network on the Internet solely via peering links. Hurricane Electric can reach "only" 85% of the Internet via peering links alone.

Last but not least, we notice on the margins of the graph a couple of clusters of ASes that represent the sub-spaces of Ukraine's cyberspace, namely Crimea and the Donbass regions, which are dealt with in the next section.

FIGURE 4. UKRAINE'S PATHS TO THE GLOBAL INTERNET, JUNE 2019

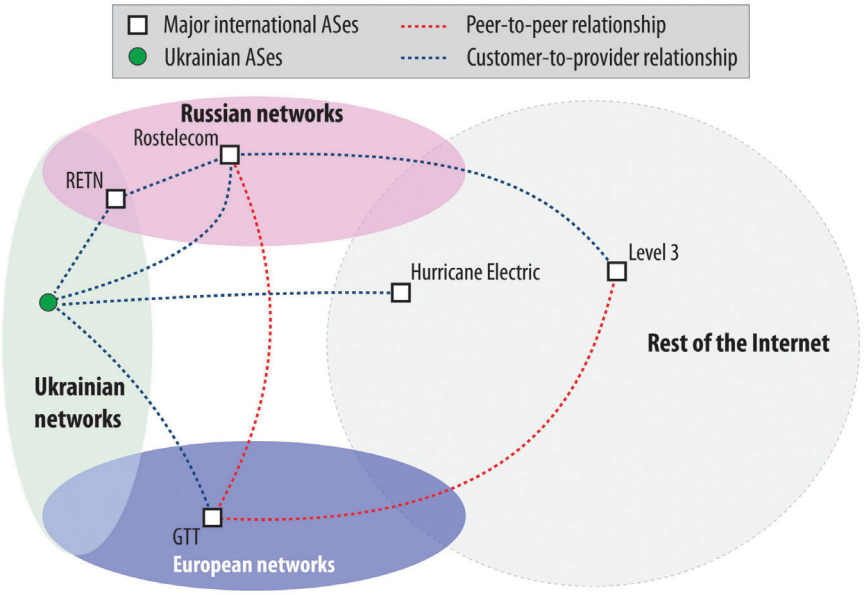

C. A Complex Network, Hard to Control

The Ukrainian network is therefore two-headed, with a few major ASes providing most paths toward the global Internet (Figure 4). It also appears to be rich and distributed from the heart of the country, with some peripheral ASes, on average two jumps away from a major Ukrainian AS. The disputed territories are exceptions, as seen below. Following the Berkman Center of Internet and Society, we measured the complexity score of the network (Roberts et al. 2011) to better understand its architecture. This metric captures the complexity of a network within a country by looking at the diversity in the announcements of IP addresses assigned to the country. A high complexity score means the possibility of a larger set of routing paths, through more providers, to connect ASes to each other or to the global Internet. A low complexity score (below 1) indicates with more certainty a network that is easy to control and to protect by periphery defense (like gatekeepers or firewalls). Also, high complexity means that it will be more difficult to introduce major changes, for example through a cyber-attack, into the structure of the country's network. In other words, changing

the structure of a complex network involves putting in a lot of effort to overcome the native resilience resulting from the complexity of the network.

The results of the complexity score and control value calculation (Table 1) show that both Russia (141) and Ukraine (79) have very high complexity scores compared to other countries of the region. This means that if important changes are observed in the connectivity structure of these two countries, this would likely result from a deliberate effort to implement such a transformation.

TABLE I: COMPLEXITY SCORE AND CONTROL VALUE IN THE BLACK SEA REGION, DECEMBER 2019

Countries	Number of ASes	Complexity Score	Control Value
Ukraine	1821	79.2	0.18
Russia	5049	141.3	0.10
Bulgaria	620	23.6	0.14
Turkey	448	2.7	0.06
Georgia	91	1.8	0.46
Romania	1040	39.8	0.41
Moldova	136	3.6	0.49

We calculated another metric proposed by the Berkman Center: the control value (Roberts et al. 2011). This metric leverages the notion of "points of control," defined as the minimal set of ASes needed to connect 90% of advertised IPs in the country to the external world. The lower the control value, the greater the centralization of the network (Salamatian et al. 2019).

Ukraine requires only 18% (about 328 ASes) of its total number of ASes to announce 90% of its allocated IP addresses. This means that controlling these 328 ASes could be enough to control almost all traffic, considering the small size of ASes. Although the control value is not very high, the profusion of ASes makes the network particularly complex and therefore difficult to control. Russia's control value is lower (10%), but the number of ASes and the complexity score are much higher.

Overall, Ukraine's network is diverse and very complex with a multiplicity of actors involved and a few powerful foreign neighbors who ensure most of the external paths. The strategies of territorial appropriation developed by Russia in Crimea and the development of geopolitical conflicts on the ground therefore constitute a major challenge in cyberspace.

4. THE FRAGMENTATION OF CYBERSPACE IN UKRAINE

Crimea and Donbass have been forcefully fought over since 2014. Crimea was annexed in March 2014 and is now ruled by the Russian Federation, although Ukraine continues to claim sovereignty over the oblast. Russia controls the territory and its main infrastructures, including the Kertch bridge and supply channels for water, energy and Internet access. The two self-proclaimed republics of Donetsk and Luhansk are in a very different situation, since they pit separatists backed by unofficially involved Russian forces against Ukrainian military forces that have been joined by independent volunteers.

Although these conflicts are still active, the territorial limits have stabilized and the Ukrainian government has lost power in both territories; to Russia in Crimea and to independently elected bodies in Donbass. Network control is part of the territorial disputes that redefine power relationships. This process enhances the loss of sovereign control by Ukraine's government and reinforces the dependency of these territories on external actors.

A. The Emergence of Crimea and Donbass as Separate "Territories" in Cyberspace

Our graphs reveal the fragmentation of Ukraine's cyberspace. In 2013, the ASes of Crimea and Donbass were fully integrated into the Ukrainian network (Figure 5), as illustrated by the central position of the Crimean ASes Crelecom and CrimeCom at the heart of Ukraine's network. There was no identifiable region in Ukraine's cyberspace. These ASes are indeed quite distant from each other on the graph, which means that they did not share the same connections. In 2019, we can see sub-regions clearly, characterized by clusters of ASes at the periphery of the networks, which reflect two different geopolitical situations: the annexation of Crimea by Russia, on the one hand, and the marginalization of Donbass in the wake of the separatist uprising, on the other hand.

FIGURE 5. REPRESENTATION OF UKRAINIAN AND RUSSIAN ASES, SEPTEMBER 2013

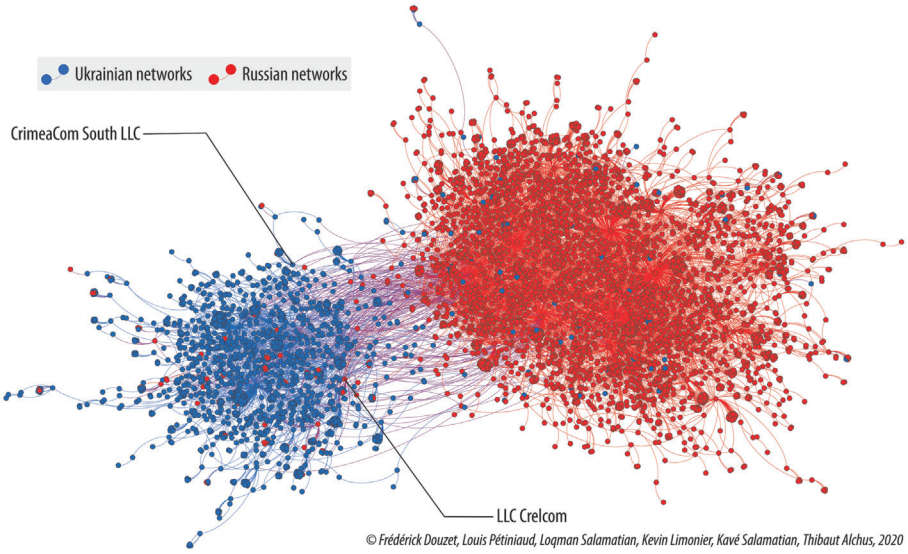

The close-up in our first graph (Figure 6) demonstrates the successful appropriation by Russia of Crimea's connectivity. Most Crimean ASes are now registered in Russia and are connected to Russian ASes. Russia managed to capture nearly all the traffic and there are almost no paths left to Ukraine's main ASes. Despite this amalgamation, Crimea remains at the periphery of the Russian network, as illustrated by the position of Crimean ASes on the graph. This spatial distance means that the number of connections between Crimean ASes and Russian ASes is limited. This observation could be explained by a deliberate strategy to isolate Crimea's network to better control it. This hypothesis requires further research.

FIGURE 6. DONBASS AND CRIMEA, "SCATTERED" TERRITORIES OF CYBERSPACE, JUNE 2019

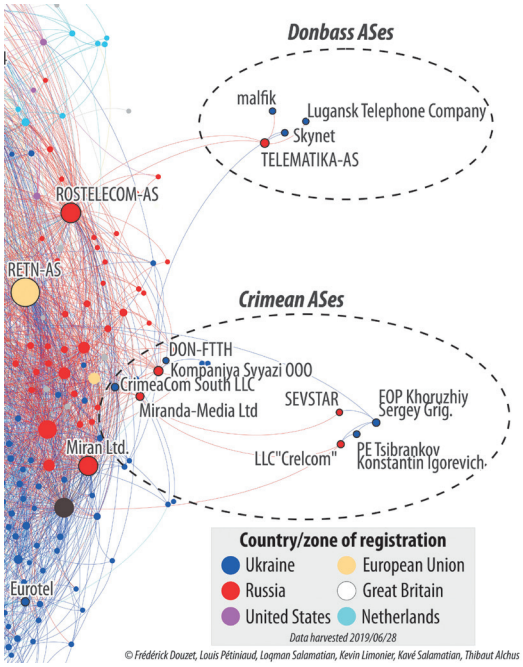

Donbass sits in a different position, apparently at the interface of Crimea and Ukraine. Its ASes have clearly migrated toward Russia but still share many connections with Ukrainian ASes. Donbass has become marginalized in the Ukrainian network, but not fully integrated into the Russian network.

How did this happen? We chose to focus on the case of Crimea to uncover the steps that led to the split between Crimea and Ukraine's cyberspace.

B. Russia's Strategies of Territorial Appropriation of Crimea in Cyberspace

Russia demonstrated its will to control the network as early as February 28, 2014, when a Russian commando force seized the building and equipment of the Ukrainian company Ukrtelekom and cut its cables that linked Crimea to Ukraine, thus disconnecting the largest part of the peninsula from the Internet. But the complexities of Internet connectivity required a more sophisticated strategy to address the concerns of Russian officials, as expressed by the Prime Minister in a tweet on March 24, 2014: data transit between Crimea and Moscow could not be provided by foreign companies.

Russia restored access to the Ukrainian connectivity, but put in place a progressive strategy that led three years later to the digital annexation of Crimea and the

marginalization of Donbass. At that time, it was physically impossible for Russia to ensure that all the Crimean traffic could go directly to the Russian mainland through the Kerch stretch, which explains why access to Ukrainian connectivity had to be restored in 2014. For at least three years, Russia had to rely partially on two fiber-optic cables connecting Crimea to the rest of the world via the isthmus of Perekop, a wide strip of land connecting the peninsula with the Ukrainian mainland. To avoid this situation, Rostelecom, the Russian company in charge of its implementation, bought 1,700 km of cables from the two main ISPs, Datagroup and Atrakom, and unveiled 46 km of new cables through the Kertch Strait, the only option to avoid transit through Ukrainian hubs, on April 25, 2014.

Meanwhile, in mid-April 2014, Rostelecom invested 15 million rubles in one of its branches, Miranda Media, to run operations in Crimea[14] and the first cable was activated on July 17, to secure strategic military communications in priority. Miranda Media (AS201776) popped up in the routing tables. A second cable (905 km long) was deployed on May 15, 2017 to absorb the traffic of Internet users as Miranda Media became more central. In July meanwhile, Ukraine's government decided to stop providing Internet access to Crimea through the two optic cables that linked them together.[15]

In addition, Russian companies pursued an active strategy of buying local ISPs and convincing others to use their services. Many ISPs became Russian as a result of pressure, to avoid potential problems, or out of loyalty to the government (Ermoshina 2018). The following graph (Figure 7) shows the overtime evolution of registration of ASes in Crimea.

FIGURE 7. DISTRIBUTION OVER TIME OF CRIMEA'S ASES BY COUNTRY OF REGISTRATION

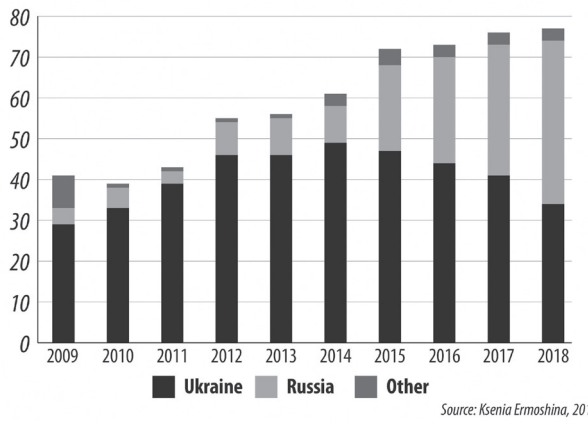

Source: Ksenia Ermoshina, 2018

14 ""Ростелеком" потратил $30 млн на покупку инфраструктуры в Крыму," *Comnews*, May 8, 2014, accessed June 6, 2019, https://bit.ly/2QOnpEw.

15 "Украина прервала связь с Крымом," *Comnews*, July 24, 2017, accessed June 6, 2019, http://www.comnews.ru/content/108850/2017-07-24/ukraina-prervala-svyaz-s-krymom.

This situation is also due to Ukrainian sanctions against companies that continued to provide Internet connectivity to Crimea after the annexation. Major ASes, Russian ones included, were forced to withdraw from Crimea to avoid jeopardizing their activities elsewhere. As a result, smaller Crimean ASes started growing bigger and more central in a network that became structured around three major ASes: Miranda Media (AS201776), Crelcom (AS6789) and CrimeaCom (AS28761), all registered in Russia. A graph of Crimean ASes and their direct neighbors (Figure 8) shows the centrality of these three providers in the network. At the heart of this graph are three major Russian ASes: Rostelecom, SPBNIT and Fiord. A few Tier 1 American ASes are present, but are not central in the graph (Hurricane, Level 3).

FIGURE 8. REPRESENTATION OF CRIMEAN ASES AND THEIR DIRECT NEIGHBORS, JUNE 2019

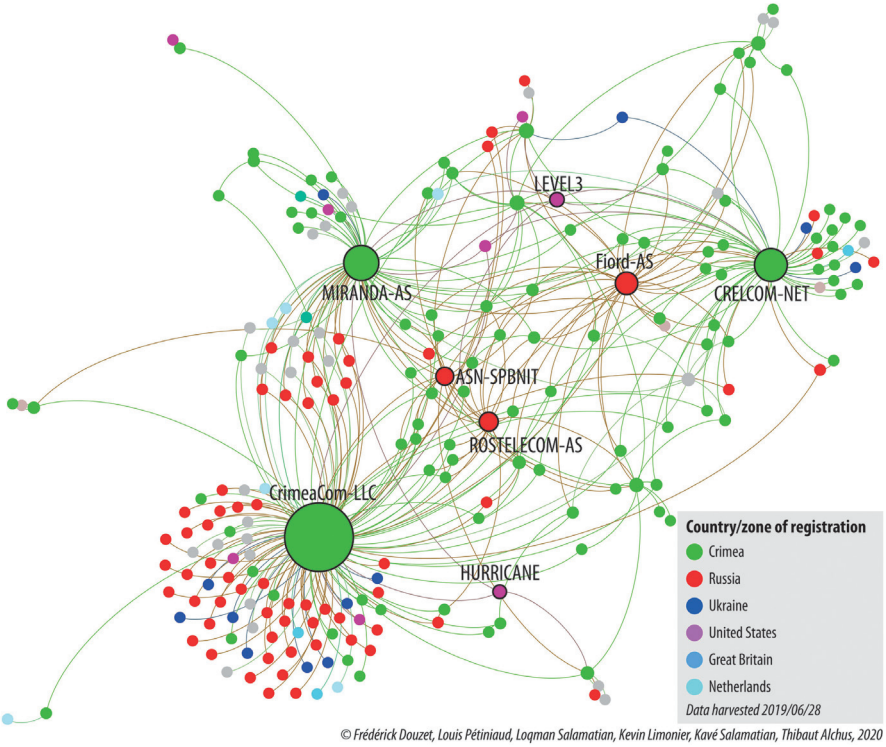

Crimea's network became increasingly centralized around three major actors close to the Russian power. Although we could not measure it, we can hypothesize that the level of complexity of Crimea has decreased as a result of these changes.

The amalgamation of Crimea with the Russian network is confirmed by a measure of latencies. We used the Atlas network to target two IP addresses, one in Simferopol (Crimea) and one in Nova Kakhovka (Kherson), and sent over 900 pings from Ukraine, Russia, Romania, Georgia, Moldova, Bulgaria and Belarus. The results presented in the two following maps (Figure 9 and 10) clearly show the difference of connectivity between these two points in the network: Crimea is in the privileged access zone of Moscow, no longer in Kiev's.

FIGURE 9. CRIMEA'S TOPOLOGICAL PROXIMITY TO MOSCOW, MEASURED BY LATENCIES, 2019

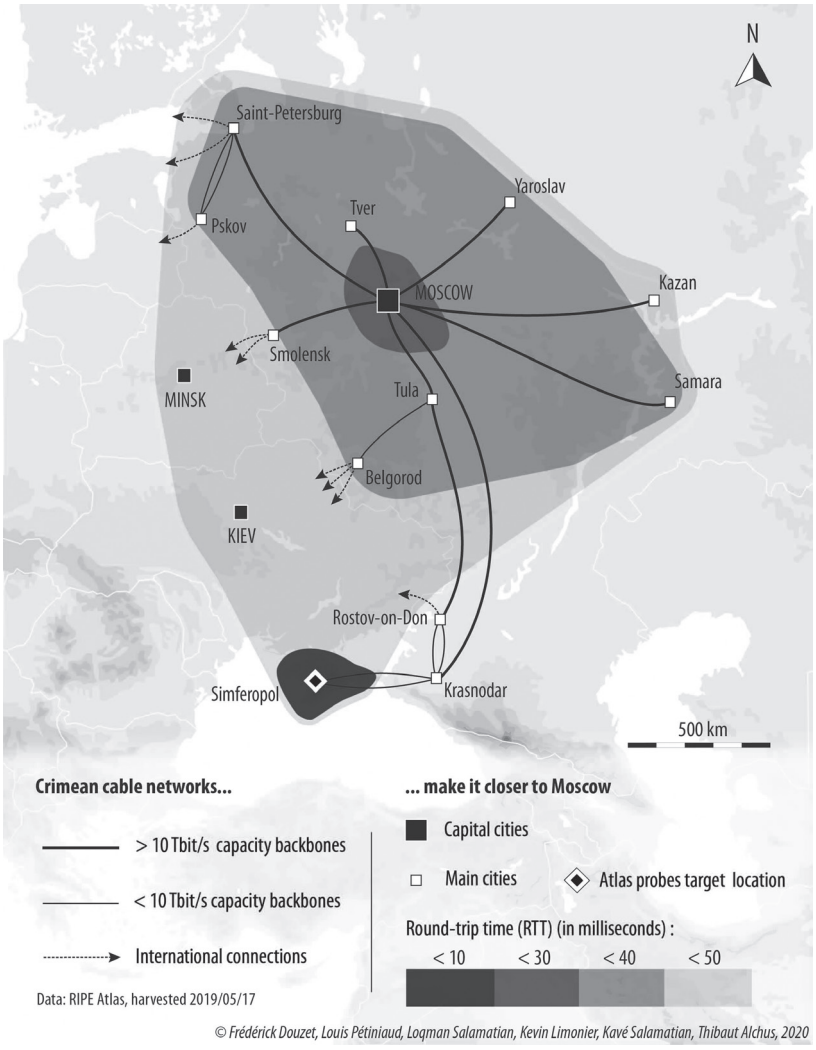

FIGURE 10. CRIMEA'S TOPOLOGICAL MARGINALIZATION FROM UKRAINE, MEASURED BY LATENCIES, 2019

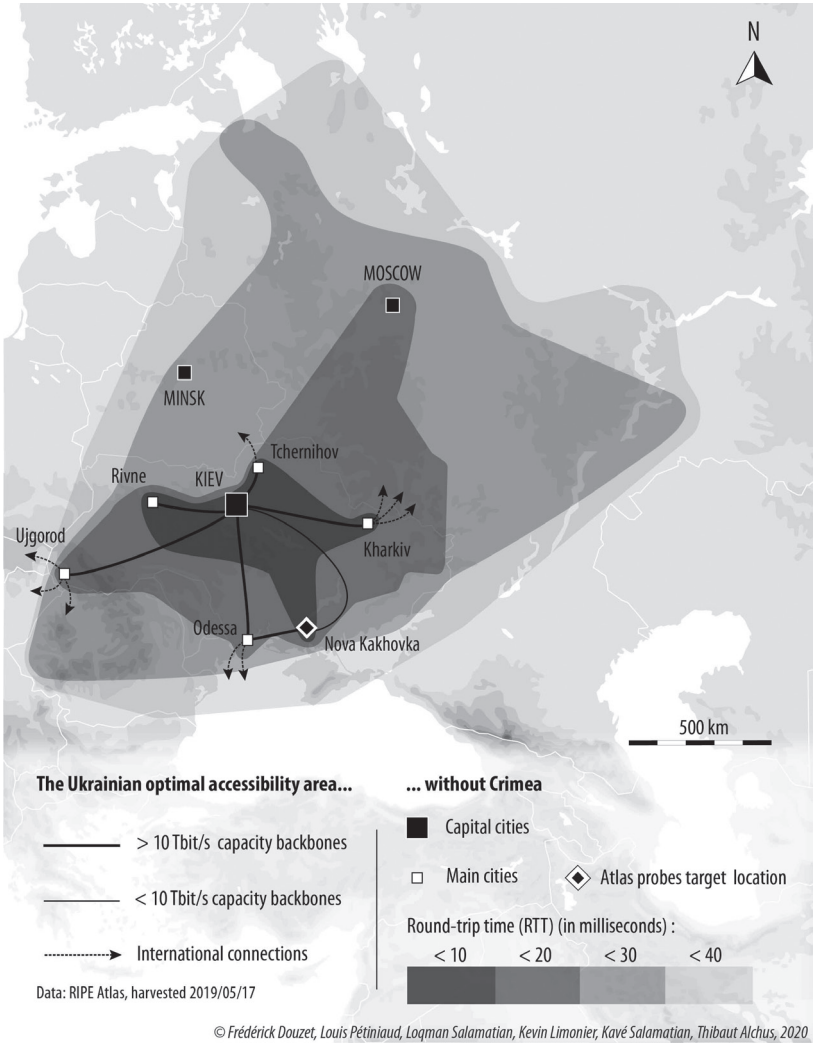

C. Longitudinal Study of Ukraine and Russia ASes

The fragmentation is also well reflected by the antagonism between Ukraine and Russia. A selection of graphs representing ASes of Russia and Ukraine at different times of the crisis show the clear relationship between the evolution of the topography and topology of Crimea and the unfolding geopolitical events (Figure 11). We observe three tendencies: 1. the break-up and progressive integration of Crimea into the

Russian network; 2. the marginalization of Donbass; 3. the gradual increase in the distance between the two countries.

FIGURE 11. THE FRAGMENTATION OF UKRAINE'S CYBERSPACE, 2014–2018

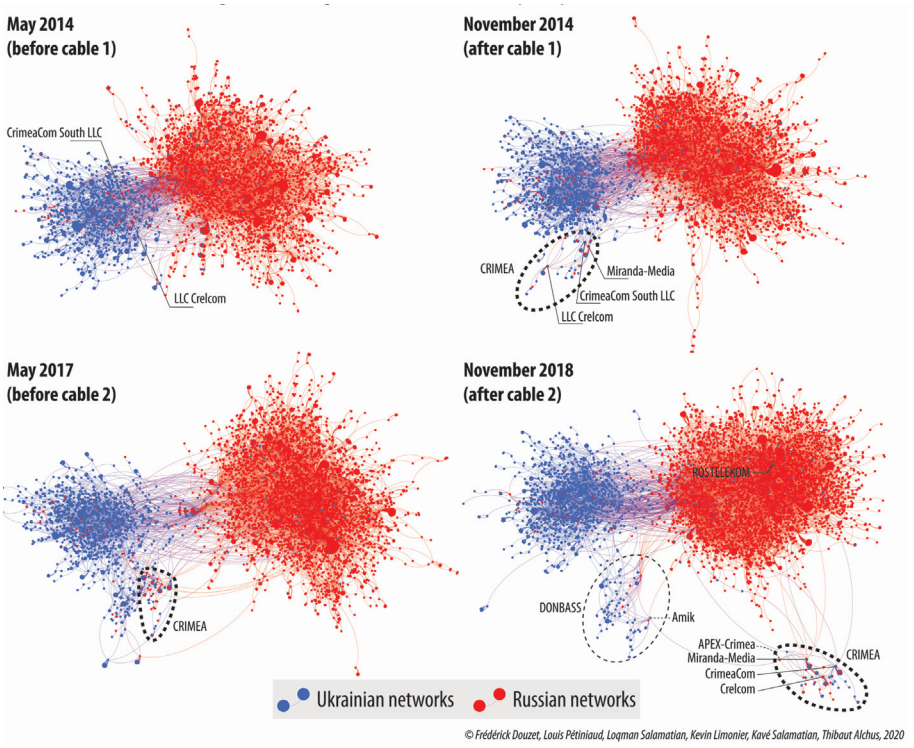

5. CONCLUSION

Our study shows that geopolitical conflicts over territories do have a clear impact on the shape of cyberspace, and that the same dynamics of annexation and fragmentation can be observed. In Crimea and Donbass, Russian authorities and separatist forces were able to attract digital traffic into their respective networks and modify BGP routes in order to divert the local Internet traffic from continental Ukraine, drawing a kind of "digital frontline" consistent with the military one. This resulted in the fragmentation of Ukraine's cyberspace, leading to the emergence of separate sub-spaces. The study of the Crimean Peninsula and of Donbass leads to important methodological findings that can allow us to: (1) define and map digital borders at the routing level; (2) analyze the strategies of actors conducting actions via BGP; (3) categorize these strategies, from

traffic re-routing to cutting off entire regions for intelligence or military purposes; and (4) anticipate future uses for BGP manipulations by identifying strategic bottlenecks within the network.

The ability to demonstrate a government's influence and deliberate strategies of territorial appropriation requires further work. Through the combination of BGP data and fieldwork-based research, we were able to demonstrate that the case of Crimea reveals a clear intent, on the part of Russia, to achieve a control of the connectivity in addition to the physical territory in the peninsula. This case study also reveals the role played by Ukraine in this dynamic of fragmentation through its decision to sanction companies providing connectivity to Crimea.

As a result, the cartography of routing paths should be seen as an additional tool to observe geopolitical conflicts, and their consequences on cyberspace, that should be used in combination with other methodologies to obtain a more complete picture.

ACKNOWLEDGMENTS

The authors thank Maxime Cherveaux for proof-reading this paper.

REFERENCES

Ager, Bernhard, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. 2012."Anatomy of a large European IXP." In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication, Helsinki, Finland, August 2012*, 163–174. New York, NY: Association for Computing Machinery. DOI: 10.1145/2342356.2342393.

Allen, John. 2011. "Topological Twists: Power's Shifting Geographies." *Dialogues in Human Geography* 1, no. 3: 283–298.

Benton, Kevin, and L. Jean Camp. 2016. "Firewalling Scenic Routes: Preventing Data Exfiltration via Political and Geographic Routing Policies." In *SafeConfig '16: Proceedings of the 2016 ACM Workshop on Automated Decision Making for Active Cyber Defense, 2016*, 31–36. DOI: 10.1145/2994475.2994477.

Böttger, Timm, Gianni Antichi, Eder L. Fernandes, Roberto di Lallo, Marc Bruyere, Steve Uhlig, Gareth Tyson, and Ignacio Castro. 2019. "Shaping the Internet: Ten Years of Internet Growth." arXiv:1810.10963v3.

Butler, Kevin, Toni R. Farley, Patrick McDaniel, and Jennifer Rexford. 2010. "A Survey of BGP Security Issues and Solutions." In *Proceedings of the IEEE, Chennai, India, January 2010* 98, no. 1, 100–122. https://ieeexplore.ieee.org/abstract/document/5357585/.

Chiu, Yi-Ching, Schlinker Brandon, Radhakrishnan Abhishek Balaji, Katz-Bassett Ethan, and Govindan Ramesh. 2010. "Are We One Hop Away from a Better Internet?" In *Proceedings of the 2015 Internet Measurement Conference, Tokyo, Japan, October 2015*. 523–529. DOI: 10.1145/2815675.2815719.

Cohen, Rami, and Danny Raz. 2006. “The Internet Dark Matter-on the Missing Links in the AS Connectivity Map.” In *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications, 2006*. DOI: 10.1109/INFOCOM.2006.234.

Dimitropoulos, Xenofontas, Dmitri Krioukov, Marina Fomenkov, Bradley Huffaker, Young Hyun, George Riley, and Kimberly C. Claffy. 2007. “AS Relationships: Inference and Validation.” *ACM SIGCOMM Computer Communication Review* 37, no. 1: 29–40. DOI: 10.1145/1198255.1198259.

Dodge, Martin, and Robert Kitchin. 2003. *Mapping Cyberspace*. London: Routledge.

Edmundson, Anne, Roya Ensafi, Nick Feamster, and Jennifer Rexford. 2018. “Nation-State Hegemony in Internet Routing.” In *COMPASS ‘18: Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, Menlo Park, United States, June 2018*. DOI: 10.1145/3209811.3211887.

Ermoshina, Ksenia. 2018. “A Routing Interregnum: Internet Infrastructure Transition in Crimea after Russian Annexation.” *35th Chaos Communication Congress* (35C3). DOI: 10.5446/39274.

Ermoshina, Ksenia, and Francesca Musiani. 2017 “Migrating Servers, Elusive Users: Reconfigurations of the Russian Internet in the Post-Snowden Era.” *Media and Communication* 5, no. 1. DOI: 10.17645/mac.v5i1.816.

Faravelon, Aurélien, Stéphane Frénot, and Stéphane Grumbach. 2016. “Chasing Data in the Intermediation Era: Economy and Security at stakes.” *IEEE Security and Privacy Magazine*, Part 2, 14, no. 3: 22–31. DOI: 10.1109/MSP.2016.50.

Feamster, Nick and Anirudh Ramachandra. 2006. “Understanding the Network-level Behavior of Spammers.” *ACM SIGCOMM Computer Communication Review* 36, vol. 4: 291. DOI: 10.1145/1151659.1159947.

Gao, Lixin. 2001. “On Inferring Autonomous System Relationships in the Internet.” *IEEE/ACM Transactions on Networking* 9, no. 6: 733–745. DOI: 10.1109/90.974527.

Gregori, Enrico, Alessandro Improta, Luciano Lenzini, Lorenzo Rossi, and Luca Sani. 2012. “On the Incompleteness of the AS-level Graph: A Novel Methodology for BGP Route Collector Placement.” In *IMC ‘12: Proceedings of the 2012 Internet Measurement Conference, Boston, MA, United States, November 2012*, 253–264. DOI: 10.1145/2398776.2398803.

Howard, Philip N., Bharath Ganesh, Dimitra Liotsiou, John Kelly, and Camille François. 2018. “The IRA, Social Media and Political Polarization in the United States, 2012-2018.” *Project on Computational Propaganda*, Oxford Internet Institute.

Latour, Bruno. 1987. *Science in Action*. Washington, DC:Howard University Press.

Latour, Bruno. 2005. *Reassembling the Social*. Oxford: Oxford University Press.

Leguay, Jeremie, Matthieu Latapy, Timur Friedman and Kave Salamatian. 2007. “Describing and Simulating Internet Routes”. *Computer Networks, Elsevier Science*, vol. 51, n°8: 2067-2085.

Limonier, Kevin. 2017. “Guerre hybride russe dans le cyberespace.” *Hérodote* 166–167, no. 3: 145–163. DOI: 10.3917/her.166.0145.

Limonier, Kevin. 2018. *Ru.Net: Géopolitique du cyberespace russophone*, Paris: L’Inventaire.

Ma, Nan, Jiancheng Guana, and Yi Zhao. 2008. “Bringing PageRank to the Citation Analysis.” *Information Processing & Management* 44, no. 2: 800–810. DOI: 10.1016/j.ipm.2007.06.006.

Moriano, Pablo, Soumya Achar, and L. Jean Camp. 2016. “Macroeconomic Analysis of Routing Anomalies.” In *TPRC 44: The 44th Research Conference on Communication, Information and Internet Policy, Arlington, VA, United States, September-October 2016*. https://ssrn.com/abstract=2755699.

Musiani, Francesca, Derrick L. Cogburn, Laura DeNardis, and Nanette S. Levinson. 2016. *The Turn to Infrastructure in Internet Governance*. Houndmills: Palgrave Macmillan.

Painter, Joe. 2010. "Rethinking Territory." *Antipode*, 42, no. 5: 1090–1118.

Piper, Andrew. 2013. "Reading's Refrain: From Bibliography to Topology." *ELH* 80, no. 2: 388.

Poese, Ingmar, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. "IP Geolocation Databases: Unreliable?" *ACM SIGCOMM Computer Communication Review* 41, no. 2: 53–56. DOI: 10.1145/1971162.1971171.

Roberts, Hal, David Larochelle, Rob Faris, and John Palfrey. 2011. *Mapping Local Internet Control*. Harvard: Berkman Klein Center, Harvard Center for Internet & Society. https://cyber.harvard.edu/netmaps/mlic.pdf.

Roughan, Matthew, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 2011. "10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems." *IEEE Journal on Selected Areas in Communications* 29, no. 9. DOI: 10.1109/JSAC.2011.111006.

Salamatian, Loqman, Dali Kaafar, and Kavé Salamatian. 2018. "A geometric approach for Real-Time Monitoring of Dynamic Large Scale Graphs. AS-Level Graphs Illustrated." Paper given at *ACM SIGCOMM Internet Measurement Conference, Boston, MA, United States, October-November 2018*.

Salamatian, Loqman, Frédérick Douzet, Kavé Salamatian, and Kevin Limonier. 2019. "The Geopolitics Behind the Routes Data Travels: A Case Study of Iran." arXiv:1911.07723.

Severo, Marta, and Tommaso Venturini. 2016. "Enjeux topologiques et topographiques de la cartographie du Web." *Réseaux*, 195, La Découverte.

Sowell, Jesse H. 2012. "Empirical Studies of Bottom-Up Internet Governance." In *Proceedings TPRC, Arlington, VA, September 2012*.

Van Beijnum, Iljitsch. 2002. *BGP: Building Reliable Networks with the Border Gateway Protocol*. O'Reilly Media, Sebastopol, CA, United States.

Vervier, Pierre-Antoine, Olivier Thonnard, and Marc Dacier. 2015. "Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks." In *Proceedings 2015 Network and Distributed System Security Symposium, San Diego, CA, United States, February 2015*. DOI: 10.14722/ndss.2015.23035.

Wählisch, Matthias, Thomas C. Schmidt, Markus de Brün, and Thomas Häberlen. 2012. "Exposing a Nation-Centric View on the German Internet - A Change in Perspective on AS-Level." *In: Taft N., Ricciato F. (eds) Passive and Active Measurement. PAM 2012. Lecture Notes in Computer Science, vol 7192. Springer, Berlin, Heidelberg*.

Wong, Joon Ian. 2016. "How Streaming Video Changed the Shape of the Internet." *Quartz*, October 5, 2016. https://qz.com/742474/how-streaming-video-changed-the-shape-of-the-internet/.

Zimmer, Jameson. 2018. "Google Owns 63,605 Miles and 8.5% of Submarine Cables Worldwide." *BroadbandNow*, September 12, 2018. https://broadbandnow.com/report/google-content-providers-submarine-cable-ownership/.