JOURNAL OF CYBERSECURITY

Journal of Cybersecurity, 2021, 1–19
https://doi.org/10.1093/cybsec/tyab018
Research paper

Research paper

The geopolitics behind the routes data travel: a case study of Iran

Loqman Salamatian[1], Frédérick Douzet[2,*], Kavé Salamatian[3] and Kévin Limonier[4]

[1]Department of Computer Science, Columbia University Mudd Building, 500 W 120th St, New York, NY 10027, USA, [2]GEODE, Université Paris 8, Campus Condorcet, Bâtiment Nord, 14, cours des Humanités, 93322 Aubervilliers, France, [3]Polytech Annecy-Chambery, Université de Savoie, 5 chemin de Bellevue, 74940 Annecy, France and [4]GEODE, Université Paris 8, Campus Condorcet, Bâtiment Nord, 14, cours des Humanités, 93322 Aubervilliers, France

*Correspondence address: Campus Condorcet, Batiment Nord, 14, cours des Humanités, 93322 Aubervillers, France. Tel: +33-6-87-14-44-61; E-mail: fdouzet@gmail.com

Received 16 December 2020; revised 17 June 2021; accepted 23 July 2021

Abstract

In November 2019, in the wake of political demonstrations against the regime, Iran managed to selectively cut off most traffic from the global Internet while fully operating its own domestic network. It seemingly confirmed the main hypothesis our research had led us to, based on prior observation of data routing: Iran's architecture of connectivity enables selective censorship of international traffic. This paper examines, through the case of Iran, how states can leverage the Border Gateway Protocol (BGP) as a tool of geopolitical control and what are the trade-offs they face. This question raises a methodological question that we also address: how the analysis of BGP can infer and document these strategies of territorialization of cyberspace. The Internet is a network of networks where each network is an autonomous system. Autonomous systems (ASes) are independent administrative entities controlled by a variety of actors such as governments, companies and universities. Their administrators have to agree and communicate on the path followed by packets travelling across the Internet, which is made possible by BGP. Agreements between ASes are often confidential but BGP requires neighbouring ASes to interact with each other in order to coordinate routing through the constant release of connectivity update messages. These messages announce the availability (or withdrawal) of a sequence of ASes that can be followed to reach an IP address prefix. In our study, we inferred the structure of Iran's connectivity through the capture and analysis of these BGP announcements. We show how the particularities of Iran's BGP and connectivity structure can enable active measures, such as censorship, both internally and externally throughout the network. We argue that Iran has found a way to reconcile a priori conflicting strategic goals: developing a self-sustaining and resilient domestic Internet, but with tight control at its borders. It thus enables the regime to leverage connectivity as a tool of censorship in the face of social instability and as a tool of regional influence in the context of strategic competition.

Key words: geopolitics, cyberspace, Iran, Border Gateway Protocol, data routing

Introduction

In November 2019, in the wake of severely suppressed political demonstrations against the regime, Iran managed to cut off most traffic from the global Internet while still fully operating its own domestic network. Within the course of 24 h, the regime was able to selectively block access to the outside Internet for most users, except for a small portion of traffic vital to its economy such as banking data. This single event tends to reinforce the main hypothesis that

our research had led us to, based on prior observation of data routing: Iran's architecture of connectivity enables selective censorship of international traffic.

This initiative is not isolated. On 23 December 2019, Russia claimed to have successfully ran a test disconnecting its network from the global Internet,[1] on a par with the law no. 608767-7 on the creation of a 'sovereign Internet'—known as the RuNet—which came into force the month before.[2] Russia has therefore clearly set as a strategic goal to redesign its architecture of connectivity in order to be able, just like Iran, to better enforce control over data routing and information.

These operations reflect the growing geopolitical importance of data routing, as cyberspace has now become a strategic domain. In most states, the development of the Internet happened through the initiative of multiple stakeholders acting in dispersed order [1], motivated by the fantastic promises of the technology and bright economic opportunities. While the US government played a major role in fostering the initial development of the network of networks, it handed its global expansion and commercialization to the private sector. Although many states invested in the development of the Internet in their own countries, very few did it with security and control at heart. China is one of the few states that initially perceived the strategic dimension of the Internet and built its architecture around the priority of government control, designing its Internet from the start as an intranet with tight control at its borders [2]. Instead, in most other countries, the process often resulted from a stack of micro-decisions from a large community of actors and the wide adoption of protocols that were initially conceived as a quick fix to keep the flow of data going [3]. The Border Gateway Protocol (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 [4] or by countries to block access to some contents or exclude some users from the Internet, for malicious and strategic purpose [5–75–7]. Numerous studies have revealed some inherent fragilities of this protocol—also known as the 'three napkins protocol'—invented in 1989 as a quick fix for data routing in the context of the exponential growth of the Internet [8, 9]. Early works show that spammers have used traffic diversions on the network [10]. Since then, several countries have opted for a network architecture facilitating the definition of a BGP strategy.

With the massive development of the Internet and the proliferation of cyberattacks since the mid-2000s, states have seen their sovereign powers challenged by multiple actors, be they criminals, hacktivists, private corporations, dissidents, non-state actors or other states [11]. They have also discovered new opportunities to increase and assert their power, making cyberspace the new frontier of state power [12–14]. Robert Kaiser argues that the 2007 attacks against Estonia catalysed the materialization of cyberwar as a new policy object [15], leading to the representation of cyberspace as a warfighting domain [16]. In 2016, at the Warsaw Summit, NATO member nations recognized cyberspace as a new 'operational domain in which NATO must defend itself as effectively as it does in the air, on land and at sea' [17], thus following the United States and a wide range of countries that had already identified cyberspace as a new military domain. In the wake of the Arab spring and the rise of social media influence, cyberspace has become a matter of the high politics of national security and core concerns of governments [18]. This securitization of cyberspace [19, 20] has led states to develop strategies of territorial appropriation of cyberspace in order to assert their power and defend their sovereignty, leading to concerns about the fragmentation of the Internet [21]. Daniel Lambach however argues that a multitude of actors engage in territorializing practices and that there are multiple ways to territorialize cyberspace [22]. The architecture of data routing is one.

The question we ask in this paper is the following: How can states leverage BGP as a tool of geopolitical control and what are the trade-offs they face? This question raises a methodological question we also address: How can the analysis of BGP infer and document these strategies of territorialization of cyberspace?

These strategies of territorial appropriation and control are antagonistic with the need for states to enjoy the benefits of global connectivity. This fundamental tension between these contradictory objectives—control and openness—requires states to arbitrate on key strategic decisions. This paper addresses some of these trade-offs.

This question sits at the intersection of several bodies of literature. First, the international relations literature on the control of cyberspace for strategic purposes has largely focused on the control of contents [23] and issues of global governance, particularly in the context of the domain name system [24, 25]. Nazli Choucri and David D. Clarke however provide a useful conceptual framework for analysing control points, identifying where processes could be influenced depending on which actor controls action at each level [26]. Based on his extensive study of BGP operations, Ashwin J. Mathew argues that the fully decentralized nature of Internet routing is a myth and that the process has always included centres of power [27]. In his account of the evolution of Internet routing, Bradley Fidler discusses the challenges of control over routing through an historical perspective, arguing that control has been difficult to maintain as the Internet became more global [28].

Second, the literature on BGP mostly focuses on its technical aspects [29, 30], its security [9], its evolution over time [27, 28] or its economic dimensions [31]. But its strategic dimensions are largely understudied. Several articles have however explored BGP strategies of nation-states. Edmundson *et al.* [32] elaborated a methodology based on traceroutes quantifying the importance of a specific set of nations for global routing. Karlin *et al.* [33] describe on a coarse level the impact of extra-territorial Internet routing and the associated risk in terms of sovereignty. Wählslich *et al.* [34] develop a taxonomy of the autonomous system connectivity that aims at understanding the role of the main actors of connectivity in Germany.

Third, the literature in geopolitics has recently acknowledged cyberspace as a new space of geopolitical conflicts and strategic competition between great powers [11, 13, 14]. According to the definition of Yves Lacoste, geopolitics studies the rivalries of power and influence over a territory at various levels of analysis [35]. It analyses the dynamics of a conflict over a territory, the contradictory representations and strategies of stakeholders to assert control and ownership of a territory, and defend their interests within this territory.

There are however few empirical studies in the geopolitics literature documenting the strategies state and non-state actors develop to assert territorial control over cyberspace through connectivity. Efforts to understand and map the architecture of cyberspace have focused primarily on the physical layer (cables, servers and other material equipment) grounded in the physical territory [36, 37]. These can be adequately understood and easily mapped with the traditional tools of political and physical geography. Some studies have also attempted to capture the overall data traffic [38]. The information layer has been the object of much attention since the 2010s, due to jihadist

1 Cantalin Cimpanu, 'Russia successfully disconnected from the internet', Zero Day, 23 December 2019, https://www.zdnet.com/article/russia-successfully-disconnected-from-the-internet/

2 Federal Law n 608767–7 'On information, information technologies and information defense', https://sozd.duma.gov.ru/bill/608767–7.

propaganda and manipulations of information during elections, leading to innovative cartographies of social networks and of the modes of content propagation [39, 40]. The cartography of the logical layer and the strategic dimension of the architecture of connectivity and data routing, however, have been given much less attention in the scientific literature [41, 42].

This paper offers a detailed analysis of Iran's architecture of connectivity. The goal of our study was to objectively characterize the strategies adopted by decision-makers in a chosen region with regard to BGP architectures in order to infer their strategic goals. We present the results of an empirical 3-year-long research project conducted by GEODE,[3] a multidisciplinary team made of geographers, computer scientists, mathematicians and area specialists. Our team has developed a new methodology to map cyberspace in its lower layers (infrastructures and routing protocols) in order to measure and represent the structure of connectivity in areas of geopolitical tensions using BGP.

We decided to focus on Iran for several reasons. First, there are clear indicators that over the past decade, Iran has sought to develop a strategy to better control its Internet connectivity. Iran has a long history of communication networks control that initially overlooked the strategic dimension of Internet connectivity until the 2010s. Before the 1979 revolution, the second bureau of the army in charge of military intelligence had implemented a centralized control of communications, seeking full control for surveillance purposes. After the revolution, this process was continued and augmented by the implementation of a kill switch for all communication means for strategic security[4]. The government, however, largely disdained the advent of the Internet in the mid-1990s, enabling private actors to invest into Internet infrastructure and build communication links that were not under the direct control of the government. The Stuxnet attack against the nuclear facility of Natanz in June 2010 came as a wake-up call for the authorities. The year before, the Iranian Green Movement emerged in the wake of the disputed 2009 presidential election and triggered a campaign led by religious conservatives meant to teach young people about the risks of the Internet, perceived as a major catalyst of revolt [43].

The 2010s therefore constituted a turning point in the development of Iran's cyber strategy. In November 2010, Iran set up a 'Cyber Defence Command' under the supervision of the 'Passive Civil Defence Organization', a subdivision of the Iranian Armed Forces Joint Staff. The Passive Civil Defence Organization was tasked with restructuring the Iranian Internet in order to make it controllable [44]. In addition, Iran's supreme leader created the 'Supreme Council of Cyberspace' by decree, in 2012, hosted in the 'National Cyberspace Centre', with a mandate to maintain up-to-date knowledge of internal/external cyberspace and to decide on 'how to deal with the harms of the Internet'.[5] The Supreme Council was first supervised by the former head of Iranian Armed Forces Joint Staff who became head of the 'Passive Civil Defence Organization'.

Second, from a technical standpoint, Iran holds a central position in the connectivity of the Middle East, the region of the world that has seen the largest growth in Internet penetration over the past decade [45]. This central position could therefore be leveraged to create strategic dependencies regarding Internet connectivity. In this regard, the Internet can be both a major source of risk and also a strategic asset for the Iranian regime to spread its influence.

Third, from a geopolitical point of view, Iran is a major actor in the Middle East and it is at the centre of several ongoing geopolitical rifts. The withdrawal of US troops from Iraq in 2011 allowed Iran to become a more important player in the region and the country clearly aspires at becoming an undisputed major regional power [46]. This desire involves consolidating domestically the stability of the regime while asserting its power on the regional scene.

In this paper, we argue that Iran's architecture of connectivity allows the regime to leverage BGP to achieve strategic goals, both domestically—to control content and suppress contestation—and externally—to enhance the autonomy and resilience of its domestic network in case of attack, and exert influence on its partners through connectivity. These different strategic goals create somewhat contradictory expectations in terms of network architecture. A tight control over connectivity is facilitated by a highly-centralized routing and therefore a low complexity in the network, with only a few control points that can be easily manipulated. But these few control points constitute major source of vulnerability in case of attack or incident. A high level of resilience of the domestic network—i.e. the faculty of a network to respond and recover after a failure—suggests instead a more distributed architecture and a higher level of complexity of the network, which is therefore more difficult to control. In addition, a total shutdown of the network through a kill switch would have adverse economic consequences for the country, which requires building an architecture that allows more sophisticated control. Similarly, leveraging the Internet as a tool of influence involves creating the physical and logical infrastructure connecting the country to its neighbours, thus further increasing the complexity of the network in order to be able to support their connectivity. When neighbouring countries have other options for accessing the Internet, the challenge for Iran is to be able to attract their traffic while maintaining tight control over global connectivity.

In order to understand the trade-offs made by Iran to achieve its strategic goals, we elaborated a methodology based on BGP data analysis (see the section 'Methodology'). We developed a cartography of the architecture of connectivity of Iran demonstrating how Iran's domestic network is connected to the global Internet (see the section 'Iran's BGP Tree Architecture: A Strong Control at Its Borders'). We then calculated the complexity of the network and compared it to neighbouring countries, in order to assess the balance between resilience and control reached by the regime for its domestic network (see the section 'A Complex Network: Low Government Control but Strong Resilience and Opacity'). We then studied how Iran, through practices, leveraged BGP as a tool of censorship (see the section 'BGP As a Tool of Censorship: The Dream of a "Halal" Internet') and as a tool of regional influence (see the section 'Connectivity as a Tool of Influence?').

We are well aware that BGP data alone are not enough to prove that all the characterizations of the network result from a deliberate and coordinated strategy implemented by the regime. However, given the political context, the public elements of Iran's cyber strategy and our empirical observation of practices, we find it difficult to consider some of these features are not intentional.

Based on these analyses, we argue that Iran has found a way to build an architecture of connectivity and leverage BGP to reconcile a priori conflicting strategic goals: developing a self-sustaining and resilient domestic Internet—but with tight control at its borders, thus enabling the regime to leverage connectivity as a tool of censorship in the face of social instability—and turning it into a tool of regional influence in a context of strategic competition.

3 Craig Timberg, 'The three-napkins protocol: Quick fix for early Internet problem left web open to attack', *The Washington Post*, 2 June 2015.

4 <Geode.science>

5 History of Iranian telecommunications, Lorestan telecommunication company newsletter, January 2013 ≪*'1Ì.†G E.'(1'* 'Ì1'F≫. ۲۸ دی ۱۳۹۲. نانترنل تارباخم تکرش یکینورتکلا

Methodology

A network of networks made of autonomous systems

The Internet is a network of networks made of around 118 000 nodes[6] (as of June 2021) called autonomous systems (ASes), each identified by a unique number. An autonomous system is itself a network. It owns and manages a set of contiguous IP addresses (or prefixes)[7] allocated by a Regional Internet Registry (RIR),[8] answering to the Internet Corporation for Assigned Names and Numbers (ICANN), a major regulatory body of the Internet.

These autonomous systems vary in size and importance. The largest AS is AS3356 and belongs to Level 3 Parent LLC; it announced 1 328 298 841 IPv4 addresses in May 2019. The second largest is AS721, owned by the US Department of Defense, with 89 384 192 addresses. In contrast, some ASes contain a single computer only.

Each autonomous system (AS) is managed by a single administrative authority—either public or private—that decides its own routing policies and has therefore full control and authority on internal routing within its network and over the access policies for traffic transiting through its network. However, an AS has to interact with its neighbours in order to exchange traffic with them. Indeed, data transiting through the Internet has to cross several independent ASes to travel from one part of the globe to the other [47]. These policies therefore require an administrator to decide with which ASes to establish connections and which routes to choose to forward its data across the Internet. These relationships—called BGP agreements—are contractual relationships based either on peering—agreements between ASes of similar size or importance on a volume of traffic exchange without monetary exchange—or commercial agreements.

The political dimension of BGP

Because of the technical and commercial nature of BGP agreements, their political dimension has largely been overlooked by the scholarly literature. And yet, BGP is political in many ways [41]. These agreements along with the algorithms that determine the priorities for data routing are guided by technical criteria—usually the shortest route in terms of number of crossed AS [48]—and economic choices (the cheapest route) but also by security and geopolitical concerns because of the vulnerability of the system, which is easy to manipulate for malicious or strategic purpose [9, 49]. In addition, the structure of connectivity can also be critical to the resilience of a network [48] and create dependency relationships between territories, providing some countries with a form of influence or even spatial power over other territories [50]. In other words, when there are different possible paths from one point to another, the AS determines which route to use to transport traffic according to its own policies that might depend on economic, technical or political considerations. Similarly, the decision to advertise a route and let traffic cross the AS to follow this route depends on the operator's commercial policy, its strategy, its competitive environment, as well as technical considerations. However, these policies are unknown to an external observer [51]. Thus, the path data packets take to move from one point to another on the Internet may change according to trade agreements and competition between economic and/or political actors. BGP is therefore a field of friction between the different actors of the network. Finally, routing policies define the routes and therefore the shapes of cyberspace [52].

Empirical studies based on BGP data face a number of methodological challenges caused by the highly dynamic nature of the Internet and the high level of technical incidents affecting routers that can fail and be restarted at any given time. Information about an autonomous system also changes frequently: ownership of ASes changes, relationships between ASes evolve fast and routing policy changes are constantly announced through updates. These challenges require an ongoing collection of data in order to infer the main structure of connectivity in a given region but also to track its evolution overtime.

Our paper focuses mostly on BGP data to infer inter-AS contracts that are in most cases not publicly available. BGP data could theoretically be controlled for or completed by many other means, including fieldwork and analysis of available commercial data. For example, we could use, if available, a sample of inter-AS contracts and identify any meaningful variation with our results, knowing that there might be differences between the legal agreement and its implementation. In the case of Iran, however, none of these options were available, which is why we relied exclusively on BGP data for this paper.

Data collected and used

In order to understand the links between ASes that determine the paths available for data transit across the Internet, we need to use inference methodologies. Indeed, ASes relationships are often confidential and not publicly available. We used graphs to represent the connectivity between autonomous systems.

We have developed a BGP observatory that generates every minute a full graph of ASes relationships obtained by processing up to 30 BGP flows coming from different routers across the network. In order to generate an AS graph, the observatory captures and processes the path updates advertised by the routers running BGP to update neighbouring routing tables [53, 54]. This real-time snapshot contains about 89 000 nodes and 200 000 links. We used the largest source of publicly available BGP routing data in 2019, RouteViews[9] and the RIPE Routing Information Service (RIS),[10] which aggregates BGP messages from BGP monitors at cooperating ASes [55]. Over a period of 3 years, we have collected >10 Terabytes of snapshots of AS graphs. In addition, we have augmented BGP announcements by adding relevant information like (i) the name associated with each AS, (ii) the country where the AS was registered, (iii) the number of IP address prefixes announced by the AS and (iv) the number of times a connection has appeared on the routing table. We used the Potaroo blog to get statistics about the number of prefixes and ASes associated with each country year after year.[11] GDP data and Internet accessibility statistics across the globe come from the World Bank website and are from 2017.[12]

The policies of BGP actors result in patterns and topologies of BGP connectivity that are represented by AS-level graphs. It is impor-

6 Introduction to the high council of cyber space, Tebyan newspaper, 16 February 2017 https://article.tebyan.net/411376/آی-روش-اب-یی-انش-یزاجم-یاضف-یاع

7 https://www.potaroo.net/tools/asn32/

8 In this paper, we focus on IPv4 since IPv6 addresses are not widely used yet in the Middle-East. See https://www.akamai.com/us/en/resources/our-thinking/state-of-the-Internet-report/state-of-the-Internet-ipv6-adoption-visualization.jsp

9 The Regional Internet Registry (RIR) for Europe, the Middle East and Central Asia is RIPE, based in Amsterdam in the Netherlands.

10 Routeviews. http://www.routeviews.org/routeviews (9 March 2020, last accessed).

11 Routing Information Service—RIS, RIPE. https://www.ripe.net/analyse/Internet-measurements/routing-information-service-ris (9 March 2020, last accessed).

12 BGP Routing Table Analysis Reports, Houston G. Blog. https://bgp.potaroo.net/(9 March 2020, last accessed).

tant to note that while each of these ASes depends administratively from a given state, many of them extend far beyond their national territories and may even be the aggregation of machines at different geolocations.

These AS graphs are however known to be incomplete, which is why we need to combine the data we collect with other sources such as active measurements (traceroutes with IP to AS mapping techniques such as bdrmapIT [56]) and IXP Looking Glass datasets when feasible and relevant. In particular, peer-to-peer links are harder to observe than customer-to-providers links [57, 58]. BGP path filtering policies do not expose less-preferred paths that would be chosen if the preferred announced paths were not available [59].

We also used a dataset compiled in 2009 by the Berkman Klein Center for Internet & Society to quantify the complexity of Internet connectivity (a notion that we define below) [60]. We have recalculated this metric for the graphs we have derived. Details on the methodology and its limitations are available in a previous publication [41].

Our methodological contribution

Our methodology differs from CAIDA's[13] work in different ways. We collect data and study the overall structure of the entire AS graph at the national, regional and global levels, but also its evolutions over time. Our goal is to infer the features of the architecture of connectivity at different levels of analysis in order to assess potential strategic choices. CAIDA primarily studies local changes in AS relationships in order to monitor Internet disruptions, using a platform called IODA.[14] This platform focuses on generating time-series capturing the local connectivity between actors through an assessment of the responsiveness of targeted ASes. Our work, instead, analyses BGP graphs to understand how the architecture of connectivity has been shaped in order to make these disruptions possible.

For the case study of Iran, we chose not to use traceroute techniques, which allow us to map the path taken by a specific packet of data between two endpoints of the Internet. First, these techniques require active measurements, and we have encountered a lack of responsiveness—and obviously wrong results—when targeting the Iranian territory. We scheduled some measurements from RIPE Atlas Probes and noticed that they were often unresponsive or returning private IP addresses. This pattern is known and the need for new topological measurement tools is currently being discussed within the networking community. The second reason is that traceroute targets a specific IP and therefore offers a very local view of the topology of the network as perceived by the end user, i.e. AS where the specific IP address is hosted. BGP gives us a more global vision of the network.

Iran's BGP Tree Architecture: A Strong Control at Its Borders

We first looked at the architecture of Iran's ASes and the way they are connected to the rest of the world. Our first finding is that Iran is connected to a limited number of its neighbours and to Tier-1 and Tier-2 ASes registered in the United States.

Figure 1 offers a representation through a graph of all the ASes in the Middle East and their direct neighbours, where nodes are an autonomous system and links are a BGP agreement. The Iranian network appears clearly on the margin of the graph because of its tightly limited number of connections with the rest of the graph. Similarly, Lebanon emerges as an independent component whose connectivity solely depends on Europe. We note a dense component constituted of Gulf states on the right of the graph that indicates a rich connectivity between the different networks. Finally, we notice a set of very densely connected nodes at the centre of the graph, composed of international ASes. They are the gateways to the international Internet.

We examined more closely the distribution of countries conterminous to the Iranian network (Fig. 2), i.e. countries in which the ASes directly connected to the Iranian network are registered. We observe that Iran is directly connected to a very limited number of Middle Eastern registered ASes. We also notice that the United States is heavily underrepresented among the registered ASes that Iran is directly connected to, considering that most of the Tier-1 and Tier-2 ASes (the highways of the Internet) are registered in the United States.

We then studied more precisely Iran's ASes and how they were connected to international ASes. With a population of >80 million and an Internet penetration rate above 50%, Iranians make up a significant share of Internet users in the Middle East. In terms of the number of ASes, Iran ranks 29th globally with 0.71% (750 ASes) of the overall ASes allocated in the world. However, only 448 ASes—representing 68% of all ASes allocated to Iran—were advertised in the network and about 90% of them contained <50 prefixes, which means that almost one third of all ASes allocated to Iran were not yet in use. In our inferred graphs, Iran therefore consists of a total of 472 ASes. In addition, >12 700 000 IP addresses are registered in Iran, ranking the country 32nd globally with about 0.35% of the whole IP addresses space in the world. Among these addresses, 98.5% were announced, which represents about 0.34% of the overall announced IP addresses[15] in the world.

Figure 3 represents the major Iranian ASes—i.e. ASes advertising at least five prefixes—along with their international neighbours. In this graph, we discard edges—i.e. ASes that are announced <3 times in different routing tables—that are often not operated because they are used as backup links or for private peering. Domestic service providers are represented in green, American providers in blue, Europeans in red and other suppliers from the Middle East in yellow.

Figure 3 provides interesting insights. First, we observe a relative lack of direct connectivity between Iran and most of the neighbouring countries. For example, there is no direct connectivity between Iran and Saudi Arabia, Bahrain or Kuwait. While communication is possible between these countries, it goes through intermediary network providers, mainly attached to the United States or the United Kingdom. This *de facto* situation clearly results from a geopolitical situation that has led to minimal economic interactions and infrastructural development between Iran and most of its neighbouring countries (except Qatar, Oman and Turkey).

Second, there are only three Iranian ASes that connect most of Iran's traffic to the rest of the world: The Information Technology Company (ITC—AS12880, AS60148), the Telecommunication Infrastructure Company (TIC—AS48159, AS49666) and Institute for Research in Fundamental Sciences (IPM—AS6736). ITC is under the aegis of the Ministry of Information and Communication Technologies (ICT) of Iran and the Telecommunication Infrastructure Company is directly affiliated to the ICT. IPM is attached to the Ministry of Higher Education. IPM is the historical provider that brought the Internet to Iran in 1993 and it used to be a connectivity point for the Iranian academic network; it remains the maintainer of the Iranian domain name registrar managing the '.ir' domain.

13 World Bank Open Data. https://data.worldbank.org/(9 March 2020, last accessed).

14 Center for Applied Internet Data Analysis based at the University of California's San Diego Supercomputer Center

15 https://www.caida.org/projects/ioda/

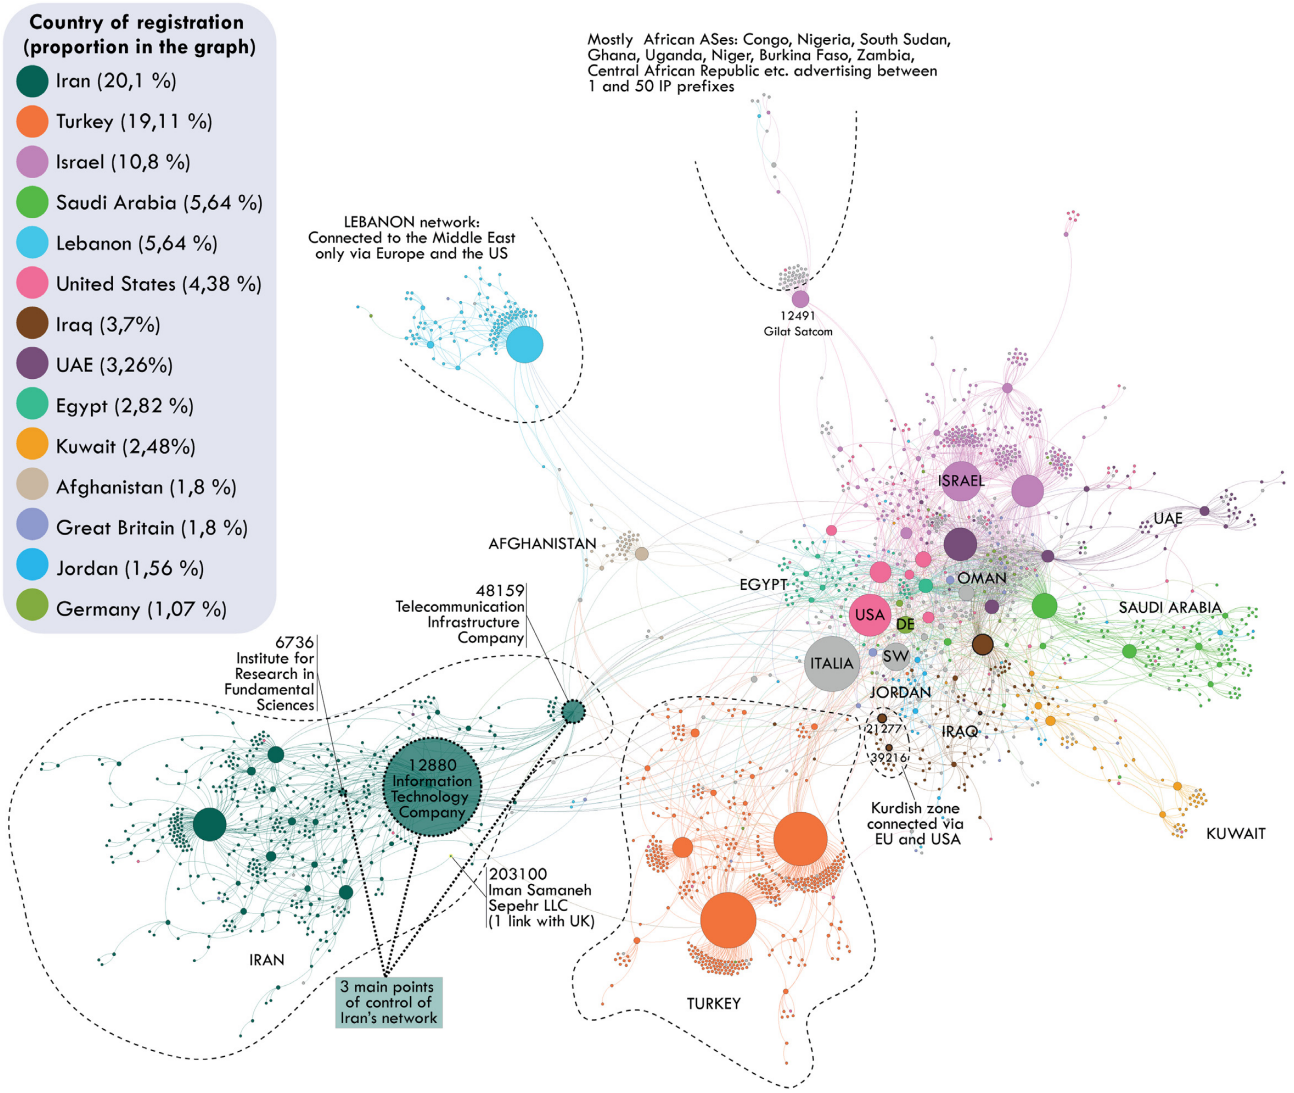

Figure 1: Representation of all the ASes in the Middle East and their direct neighbours using the ForceAtlas 2 algorithm [61].

Iran has built its architecture of connectivity in order to connect its domestic network to the global Internet through only three highly controlled ASes. Indeed, the Iranian network is akin to a tree. We can delineate (i) a trunk consisting of highly interconnected government-owned ASes that open a path to foreign networks, and (ii) branches managed by private Internet service providers (ISPs). While these branches are well connected to each other, they do not have a great diversity of paths linking their traffic to the outside and they have to pass through the trunk. Such a backbone allows Iran to control the information exchanged with the international network since all traffic goes through government ASes that may very well decide to stop it. There is a bottleneck between the Iranian Internet and the rest of the Internet; and government-controlled ASes ITC, TIC and IPM play the role of gatekeepers that control the access to foreign content and decide what traffic passes through. We believe that the control of these three points allowed the regime to selectively disconnect its domestic network from the global Internet in November 2019, allowing only a small and critical share of the traffic to go through.

The paths linking the domestic network to the global Internet are therefore highly centralized around these three points to allow control. This situation is very unlikely to change in the near future given, on the one hand, the embargo against Iran and, on the other hand, the tight control and the regulatory framework exerted by the regime on the private sector and civil society, preventing AS administrators from freely establishing peering relationships with international ASes.[16] We looked at the structure of connectivity 'within' the domestic network to understand whether the domestic routes were also centralized to allow control.

Zooming into the Iranian AS ecosystem (Fig. 4), we can observe that the most central ASes are almost evenly disseminated throughout the network. The domestic network is therefore not centralized. Respina (AS42337) is the most central ISP, connecting >100 ASes. We also notice that TIC is underrepresented within Iran's domestic network connectivity graph despite the fact that it plays a fundamental role as a connection to the rest of the network (Fig. 3). This means that internal ASes are not directly connected to TIC but rather connected through proxies that aggregate the traffic and can potentially implement filtering (see the section 'BGP As a Tool of Censorship: The Dream of a "Halal" Internet').

16 BGP Routing Table Analysis Reports. https://bgp.potaroo.net (9 March 2020, last accessed).

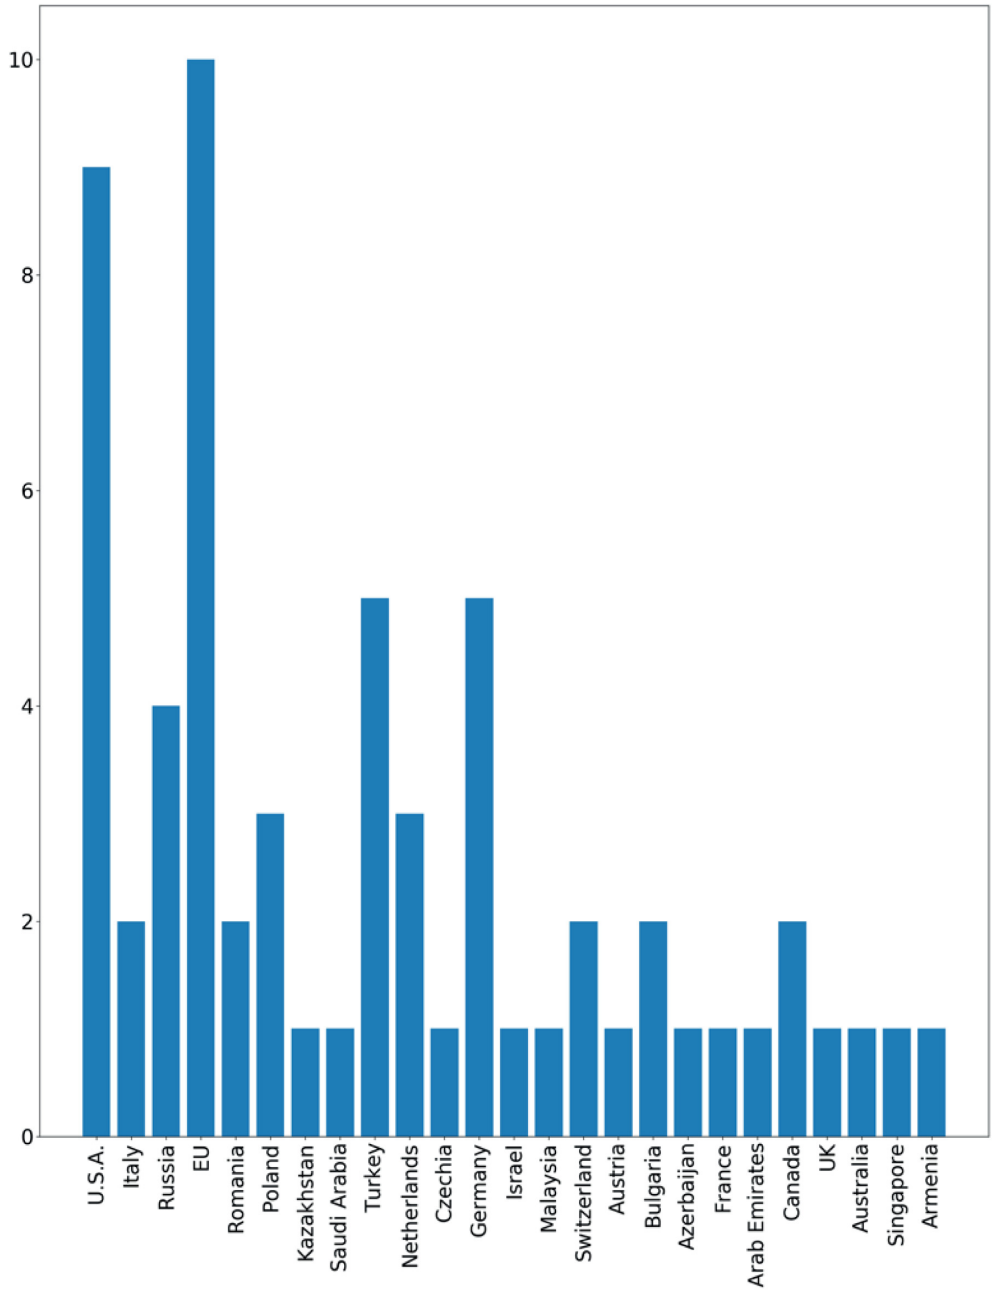

Figure 2: Histogram of the nationalities associated with the direct neighbours of ASes registered in Iran at RIPE.

We chose to compare Iran with other countries in the Middle East. Table 1 presents the total number of external and internal BGP connections of Iran and other comparable countries in the Middle East. We observe that Iran has a relatively larger proportion of internal connections, i.e. connections linking two ASes within the same country, as opposed to Israel for example, which has a larger proportion of external connections. We also note that Iran has a large number of connections overall, contrasting with other countries in the region such as Saudi Arabia where the number of edges is restricted. This translates into a richer connectivity internally.

The Iranian domestic network architecture is therefore different from other countries' networks. Our hypothesis is that Iran, while limiting its number of connections to external ASes, has sought to develop a self-sustaining and resilient domestic network, involving multiple internal connections and therefore multiple routing paths within its network. In order to test this hypothesis, we decided to calculate the complexity score of its network.

A Complex Network: Low Government Control But Strong Resilience and Opacity

In order to have a better understanding of the architecture and the dynamics of the network, we use the complexity score initially developed by the Berkman Klein Center for Internet and Society [60]. This metric captures the complexity of the network within a country by looking at the diversity in the announcements of IP addresses assigned to the country. A country where all or a major part of the

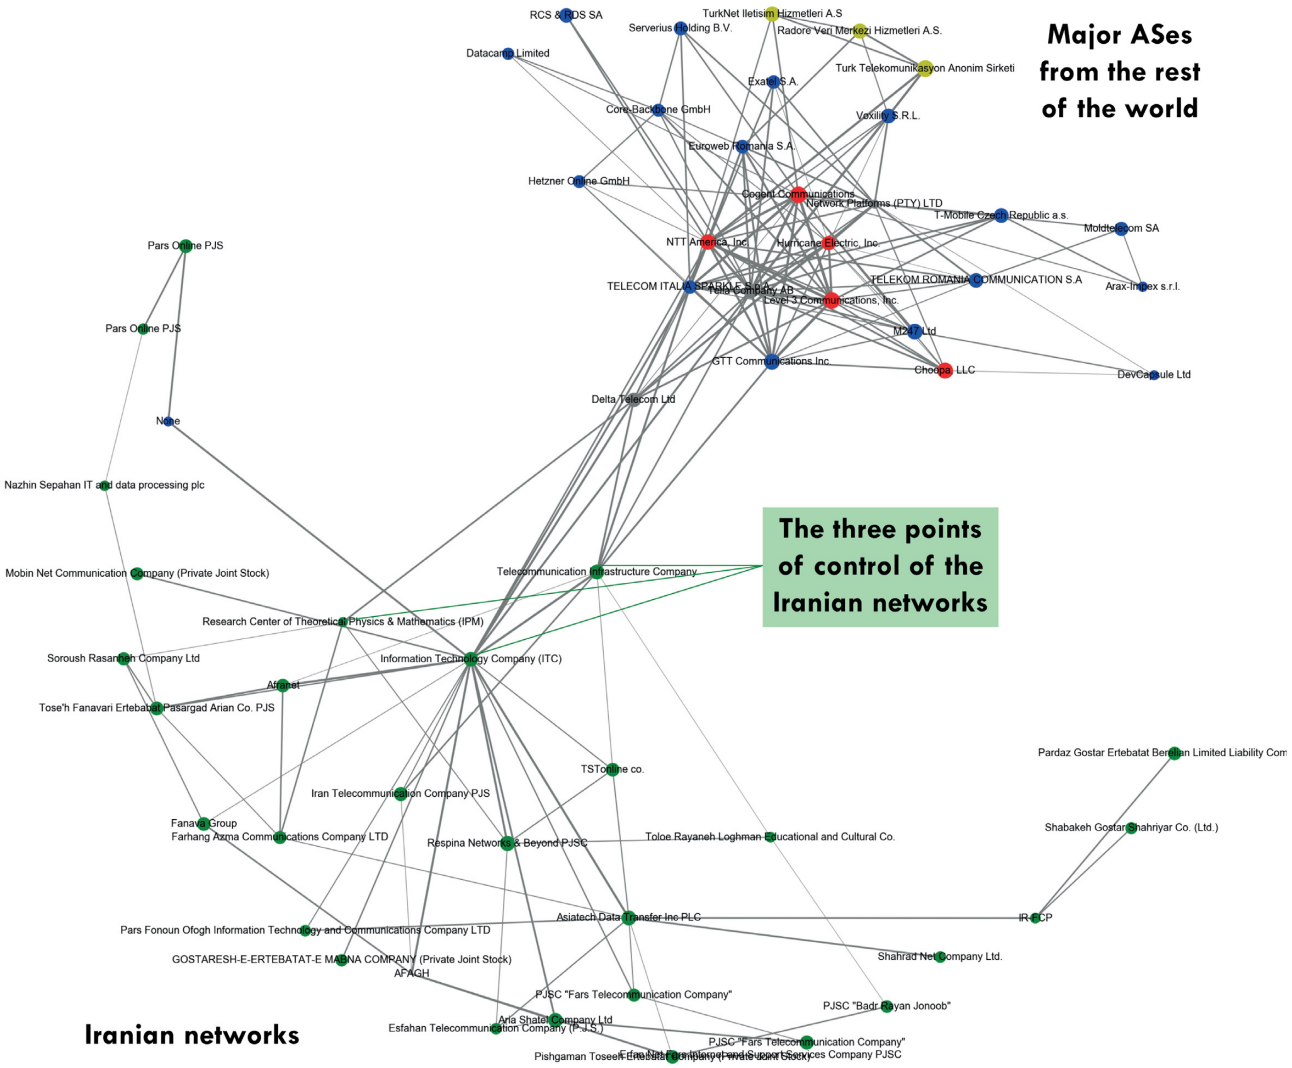

Figure 3: Simplified representation of the Iranian ASes.

available IP addresses visible from the outside[17] belong to a limited number of ASes will have a low complexity score. This means that internal actors have a limited choice of network providers. Conversely, a large diversity in the ownership of IP addresses with a large number of ASes is a sign of a more complex ecosystem within the country and it infers the possibility of a larger set of routing paths through a greater number of providers to connect to each other—or to the global Internet if the AS connectivity structure allows it.

As an illustration of a small network, let us suppose that a country wishes to implement a strong control of the Internet and permits only one single AS to provide all IP addresses visible to the outside. This enables the country to exert a perfect control on the traffic, as it would be impossible to access the network without passing through this AS, but it would also entail a very fragile network both internally and externally as this AS would become a single point of failure. A more complex network, on the contrary, would be more resilient as the diversity of existing paths and visible addresses makes it possible to circumvent the potential points of failure. The higher the complexity score, the more resilient the network is—but it is also harder to control.

To achieve a more complex and therefore more internally resilient domestic network requires to increase the number of domestic ASes and to enable their IP addresses to access the global Internet through a richer set of alternative paths. An analysis conducted by Mahsa Alimardani demonstrated the clear impact of the Stuxnet attack on the sensitivity of Iran concerning the resilience of its network and, by extension, on the creation of new ASes [62]. The ongoing international embargoes have also driven the Iranian government to develop the Iranian National Information Network project,[18] which could work as an Intranet separated from the global Internet and protect its domestic network from foreign interference [63]. Over the past 10 years, Iran has indeed increased its AS fleet at great speed and more rapidly than its competitors in the Gulf. Figure 5 shows the evolution of the number of ASes in Iran between 2009 and 2018 (in pink) and the share of ASes announced in Iran compared with the rest of the world (in green). Managing an AS requires a good level of expertise; this evolution therefore demonstrates that Iran has been able

17 See https://www.menog.org/presentations/menog-15/322-Update-IR-IX.pdf where it is mentioned that any entity willing to peer in Tehran-IX requires an agreement from TIC (p. 6).

18 IP addresses visible from the outside might be accessed from outside the domestic network, through announced BGP paths.

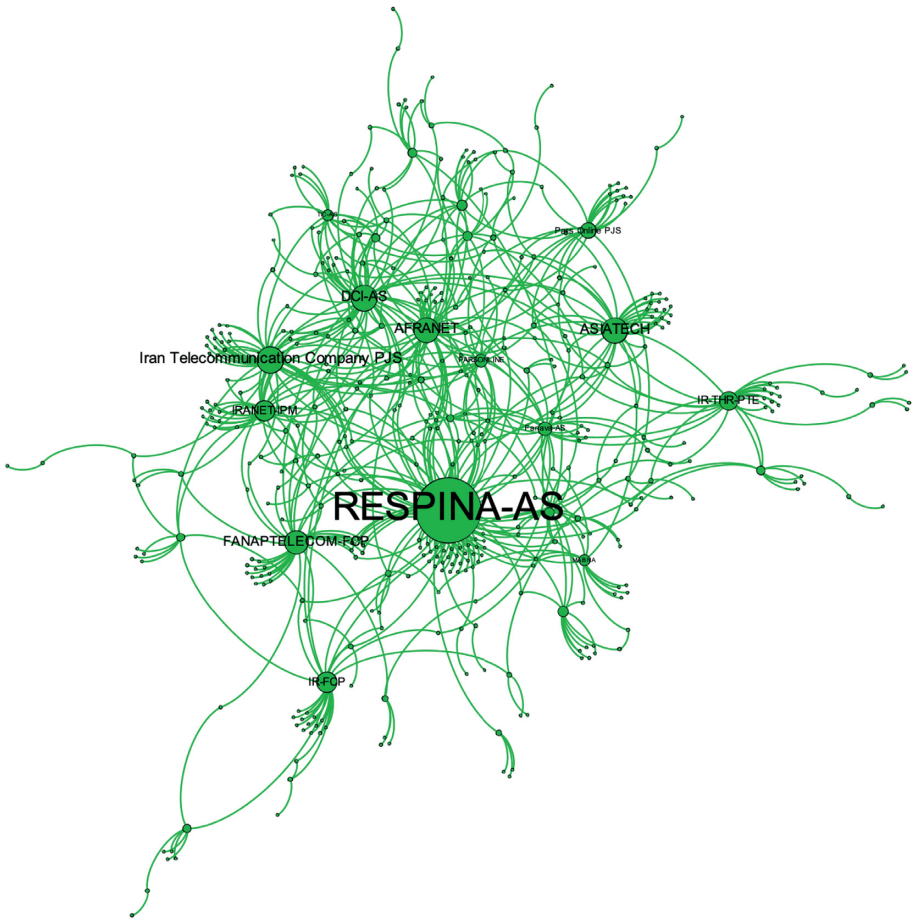

Figure 4: Zoom into the Iranian BGP landscape.

Table 1: External and internal BGP connections in Middle Eastern countries

External connection	Internal connection	Total ASes observed	Country
79	643	472	Iran
113	191	140	Saudi Arabia
113	180	109	Iraq
97	12	12	Oman
52	11	10	Qatar
40	136	132	Lebanon
128	18	22	Bahrain
24	41	37	Jordan
26	74	61	Kuwait
162	93	63	Egypt
15	51	49	Afghanistan
83	67	76	U.A.E.
395	326	261	Israel
130	567	473	Turkey

to gather the economic and technical resources that are necessary to manage this high number of ASes.

However, this rapid increase in the number of ASes comes in contradiction with the wish to control all traffic, since controlling a larger number of ASes and outgoing paths is more difficult. The complexity score developed by the Berkman Klein Center comes as a way of quantifying this trade-off between internal resilience and control.

The complexity metric presents some limitations as it tends to overly-simplify the way routing functions and does not account well for the diversity of situations and behaviours that can occur in practice. This creates a weight distortion that can exaggerate the importance of some ASes. However, we consider that this metric applies relatively well to the Middle East where ASes are often bound to national or regional usage and where the network is less densely irrigated than in Europe or in the United States.

The Berkman Klein Center offers another intuitive metric: the control value [60]. This metric leverages the notion of 'points of control' defined as the minimal set of ASes needed to connect 90% of

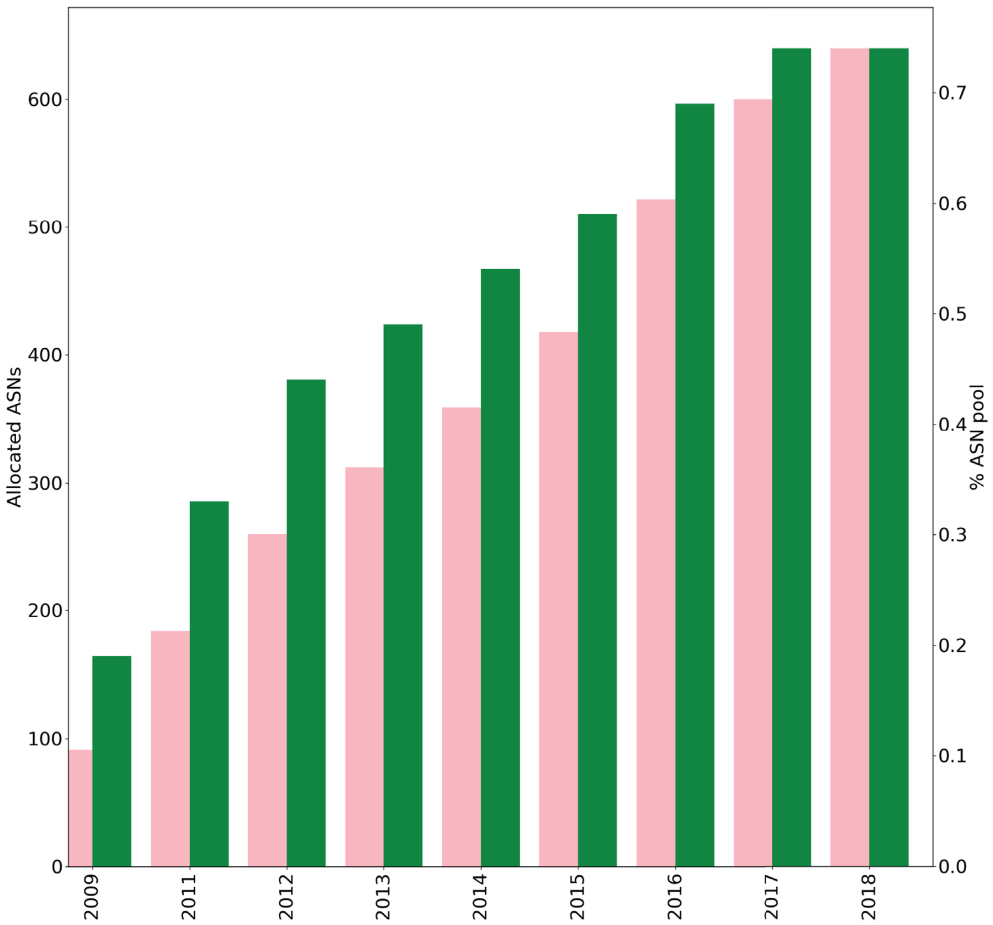

Figure 5: Evolution of the number of ASes registered at RIPE in Iran and share of Iranian ASes among all countries.

advertised IPs in the country to the external world. The proportion of the country's ASes in the set of points of control defines the control value. For example, if a country A possesses 300 ASes and if it takes 30 ASes to connect 90% of the IP addresses announced, then the size of the set of points of control is 30 and the control value is 30/300 = 10%. The lower the control value, the greater the centralization of the network.

These two scores, complexity and control value, are somewhat redundant. Control value focuses on accessibility to the Internet while the complexity score focuses more on the internal complexity. Looking at these two metrics gives us methodological tools to compare the network architecture of different countries by evaluating where they stand on the control vs resilience trade-off. The control value is a measure of the concentration of IP addresses within a small number of ASes and, by extension, of how concentrated the routing architecture is. The complexity value, on the other hand, measures the routes data actually travel through. Complexity is about control and it quantifies how the country's users may have their traffic exposed to observation, manipulation and disruption. Both values are complementary as shown in Table 2. To sum up, a network with high complexity and control value scores is decentralized with a diversity of routing paths and therefore more resilient but harder to control. A network with low scores of complexity and control value, on the contrary, is highly centralized (most IP addresses are handled by few ISPs) and therefore easier to control but less resilient.

We present the complexity metric calculated by the Berkman Klein Center for Middle Eastern countries in 2011 along with our own calculations for 2019. Our sources of data are slightly different from the one used in the original paper: we have used the full BGP connectivity graphs while Roberts *et al.* [60] used CAIDA's AS relationship data alone. Table 2 shows that, compared with other countries in the Middle East, Iran has a highly complex network with a great number of interactions between its ASes. More interestingly, Iran—like Turkey—has maintained a high level of complexity from

Table 2: Complexity in Middle Eastern countries in 2011 and 2019

	Complexity		Control value		Number of ASes	
Country	2011	2019	2011	2018	2011	2019
Iran	**3.82**	**3.75**	**2%**	**34%**	**96**	**437**
Saudi Arabia	3.74	0.43	5%	10%	66	139
Iraq	6.46	4.93	75%	55%	4	107
Oman	1.06	0.05	50%	25%	2	12
Syria	0.85	0.00	33%	50%	3	2
Bahrain	10.20	0.26	22%	37%	18	19
Kuwait	4.70	0.52	20%	17%	30	61
Egypt	1.25	0.04	8%	9%	36	58
Afghanistan	NA	4.14	NA	52%	NA	46
U.A.E.	0.58	0.31	20%	20%	8	65
Turkey	2.72	2.67	1%	7%	226	450
Israel	3.24	2.41	2%	10%	165	251
Qatar	1.55	0.02	40%	29%	5	9
Lebanon	11.99	7.81	22%	42%	32	133

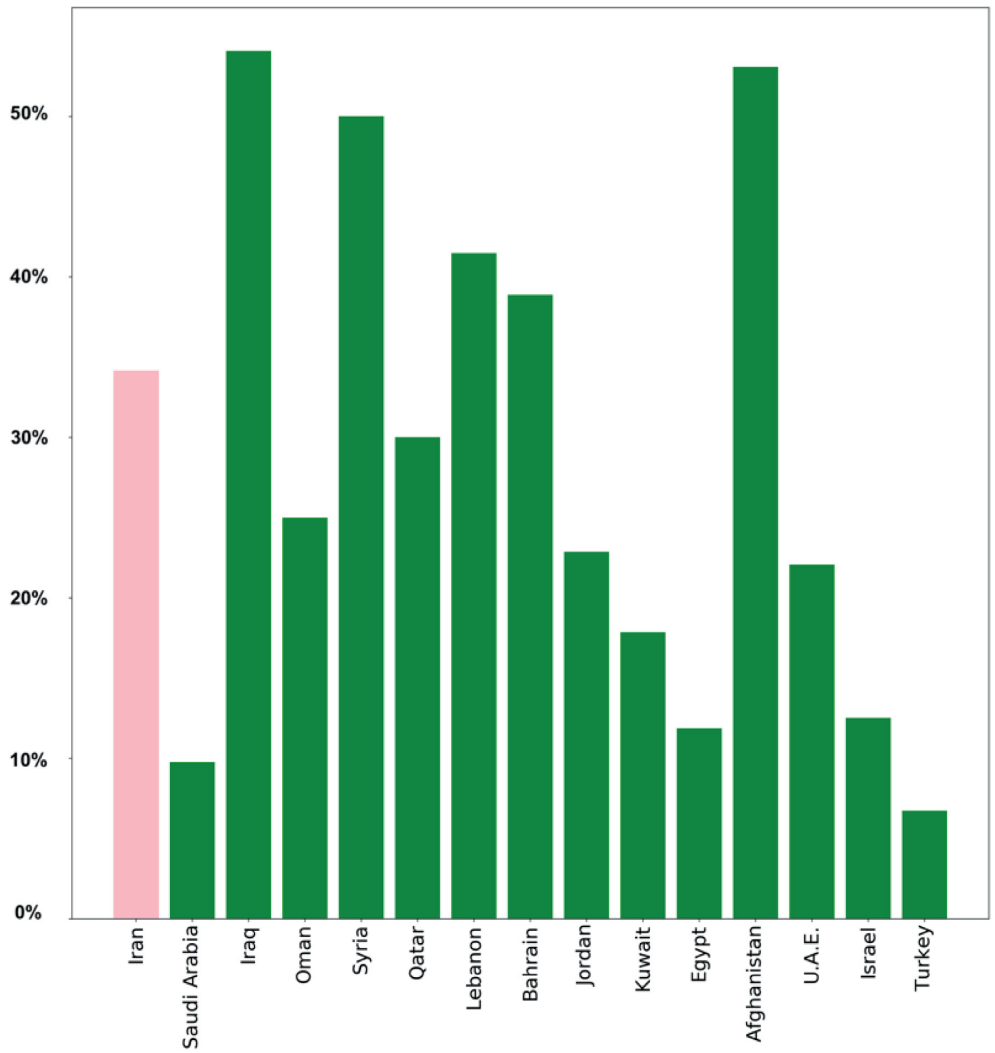

Figure 6: Network Control Values among Middle Eastern countries in May 2019.

2011 (3.83) to 2019 (3.75) despite its effort to exert tighter control over the network, while other Middle Eastern countries have drastically reduced the complexity of their networks over the past 8 years, a sign that the Iranian strategy of development of its internal network has been significantly different from that of the other countries. This tendency is exacerbated in the Gulf countries where we see that most values are hovering around a value of 0.5.

The control value scores (see Table 2 and Fig. 6) confirm the above results. In 9 countries out of 15, the control value is below 30%, meaning that 90% of the traffic circulates through less than one-third of the network. In Saudi Arabia, Turkey, Egypt and Israel, the level of centralization is particularly high, with control values close to 10%. This architecture concentrates traffic into a few ASes, making the network vulnerable to traffic congestion, potential failures or cyberattacks targeting these few ASes. But it does facilitate controls over contents and users within the boundaries of the network. Iran, on the other hand, has a much higher complexity score and therefore a much lower level of centralization, with a much higher control value of 34%. This means that the connectivity inside Iran is better distributed among the existing ASes. More importantly, this value drastically increased from a 2% in 2011. Initially, most IP addresses were concentrated into very few ASes but as Iran built its network over the past decade, IP addresses were distributed among a much greater number of ASes, thus spreading out the IP addresses across the domestic network.

In the case of Iraq and Afghanistan, we note sets of control values hovering above 50%, which reflects the lack of strong central authority capable of controlling and shaping the architecture of the domestic network. There are therefore multiple points of connection to the global Internet. It is noteworthy that even if Israel and Turkey have relatively large complexity scores, the control value there is much lower than in Iran, showing that these two countries have a stronger grip over their national network infrastructure, despite a larger number of connections to foreign countries.

These trends reflect different representations of Internet sovereignty and different strategies of control. Iran emphasizes controls at the border of the domestic network: alternative paths for traffic flows exist within its borders whereas most Middle Eastern countries have obviously increased the centralization of their domestic networks to better control Internet users. As a result, most domestic traffic is forced through a few controlled ASes, with the risk of creating domestic congestion and low resilience. Iran instead has managed to ensure a robust internal connectivity while maintaining the ability to isolate its network from the global

Internet, in a way that does not lead to congestion within the central ASes.

The above discussion shows that Iran has succeeded in building a national network that reconciles two seemingly incompatible properties. On the one hand, the Iranian domestic network is highly resilient, with a relatively large number of ASes and a rich ecosystem of internal paths. On the other hand, this network is highly controlled, with all the outgoing traffic flowing through the three main government-controlled ASes and a relatively low control value. Iran has therefore managed to isolate its network from the global Internet while ensuring a robust internal connectivity in a way that does not imply congestion among the central ASes. The bottleneck therefore occurs externally.

This implies that the Iranian network can completely cut off access to the rest of the Internet for Iranian users and isolate its network without modifying its internal state. Moreover, Iran, through its rich domestic network, can modulate the level of disconnection according to its interests and strategic objectives. A complete disconnection of the Internet comes with high collateral economic costs but partial or temporary disconnection can bear some advantages. From this perspective, the Iranian network architecture is becoming similar to the Chinese one, with the notable fact that this evolution has happened relatively quickly[19] [2].

In contrast, other Middle Eastern countries have strongly decreased the complexity of their domestic network in order to achieve a better control but at the cost of a higher risk of disruption and lower resilience that can result in domestic congestion, accidental failures or, worse, cyberattacks.

In this section, we have demonstrated how Iran's BGP architecture facilitates controls at the border while preserving the resilience of its network. The following section documents how BGP has been actively used as a tool of censorship by the regime.

BGP As a Tool of Censorship: The Dream of a 'Halal' Internet

As the Iranian authorities work towards the creation of a 'halal' Internet[20]—i.e. a domestic Intranet—it implies the establishment of a particularly sophisticated censorship system in Iran. As of 2012, about 27% of Internet sites had been blocked [64]. If we look into the details of the blocked domains, we note that a large majority of them are news media and websites dealing with human rights issues. The most censored category of websites is unsurprisingly pornographic content. It is interesting to note that >50% of the most visited websites across the world are not accessible to Iranian Internet users [65].

It is also very important to consider that content access control and censorship in Iran is not only caused by the actions of the Iranian government but increasingly implemented by foreign companies that fear being infringing on the rules of the US Office of Foreign Assets Control, thus opting for nationwide bans on accessing their own contents and services, even in cases where waivers exist. Among examples of these self-imposed content access controls are IEEE in 2003 [66], reversed in 2004 but reimplemented in 2019, Amazon Web Service in 2019,[21] GitHub in 2019 [67]. While censorship has been implemented through traditional web filtering techniques, like DNS poisoning [68]—i.e. giving wrong answers to some DNS requests making some websites inaccessible—or web proxy filtering [69], more advanced techniques like Deep Packet Inspection [70] have been deployed. However, the existence of embargoes toward Iran have reduced access to these technologies running on high speed links. While Chinese companies are very active in the Iranian telecommunication market—in particular Huawei and ZTE[22]—they are not selling filtering technologies, even though China is a technological leader in this domain. Indeed, China has implemented a very drastic exportation control over its security and filtering techniques. These technologies are classified as 'items related to the maintenance of national security and national interests' in its export restriction regulations.[23] A reason for this classification is the fear of reverse engineering and a loss of control on the usage of the technology.

The ever-increasing Internet bandwidth has made it difficult for Iranian censorship to keep up in terms of filtering capability. The regime has been forced into distributing its filtering process closer to the customer, in order to spread the load over the network and therefore to reduce the volume of traffic to process. The responsibility to implement filtering has been endowed to ISPs.[24] Nonetheless, this approach does enable rapid actions of control while adding new filters requires the cooperation of all ISPs and takes time. Yet it has been a strategic goal for Iran to implement 'kill switches' for global communications that would enable disconnecting all Internet connections, even mobile communications, from a central point in the network, in order to allow immediate censorship of contents. The specific domestic network architecture we described above, along with mindful BGP manipulations, provide the tools to achieve this goal.

In the following paragraphs, we describe how BGP architecture has been concretely used for censorship and also as an active tool for interfering with traffic. There are globally two classes of AS-related incidents: outages and hijacks. An outage happens when a prefix is no longer announced by any AS. It usually corresponds to a technical problem that is symptomatic of the fragility of an ISP. Yet the past few years have seen a rise in country-wide Internet outages caused by national censorship [71].

A hijack is the illegitimate takeover of prefixes by corrupt Internet routing tables across the graph. The traffic then follows paths that it should not be taking and transits through a new AS. This allows the new AS it crosses to analyse the nature of the traffic or to suppress it. In practice, hijacks are not all malicious.[25] They might result from unintentional misconfigurations, inducing a change on the path data follows over the Internet. Hijacks are frequently unwanted and result from misconfiguration, but they are also used with censorship in mind.

In Fig. 7, we show the correlation between the number of ASes in different countries and the number of BGP-related events, outages and hijacks, that have taken place in these countries. As we can see, there is a direct correlation between the number of ASes in a country and the number of observed outages. This results from the fact that most BGP events are configuration errors and the rate of these errors

19 RRK. http://www.rrk.ir/Laws/ShowLaw.aspx?Code=1640. (15 October 2019, last accessed).

20 Daniel Anderson, 'SplInternet Behind the Great Firewall of China', https://dl.acm.org/doi/pdf/10.1145/2390756.2405036

21 https://www.theguardian.com/world/2012/jan/05/iran-clamps-down-Internet-use

22 Center for Human Rights in Iran: More Iranians Forced to Rely on Unsafe Online Hosting after Amazon Ban. https://iranhumanrights.org/2019/08/more-iranians-forced-to-rely-on-unsafe-online-hosting-after-amazon/(15 October 2019, last accessed).

23 Su, J. Analyst: China's ZTE Shuts Down After U.S. Tech Ban Over Iran Sales. *Forbes* (9 May 2019).

24 https://www.china-briefing.com/news/chinas-export-control-law-explainer-china-briefing-news/

25 Ahmed Shaheed, 'Layers of Internet Censorship in Iran', 7 May 2014, http://www.shaheedoniran.org/english/blog/layers-of-internet-censorship-in-iran/

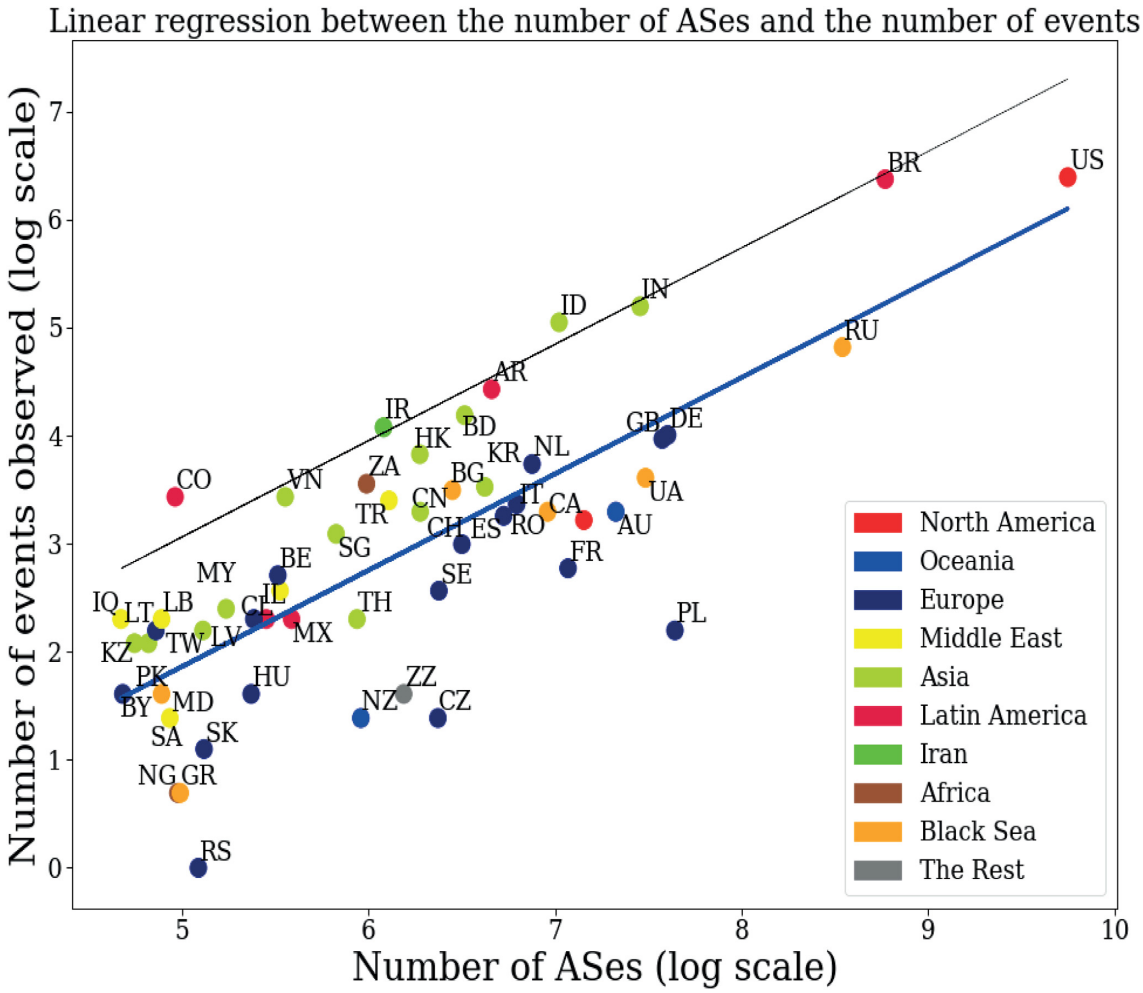

Figure 7: Linear regression between the number of ASes and the number of BGP events. The different colours represent different regions of the world.

is roughly constant over all ASes regardless of the country. However, looking at the position of Iran in this figure, we can observe that it is above the 95% regression confidence line, i.e. the likelihood that BGP events in Iran have the same explanation than the bulk of the other countries with random configuration errors is below 5%. Only two other countries are similar to Iran in this respect: Indonesia and Colombia. But both of these countries have a large number of maritime cables (from the Caribbean Sea in Colombia) and have mainly witnessed outages resulting from issues pertaining to these cables (accidental cable cuts, etc.). Iran is therefore over-represented among BGP events with respect to the size of its AS ecosystem and we interpret this as a sign that Iran is misusing BGP for specific purposes.

A focus on Iranian outages is particularly instructive [72]. In January 2018, Iran tested a holistic approach to censorship through the systematic elimination of any outgoing connection via BGP tampering and a suspension of the ASes' outgoing traffic.[26] This was most likely no coincidence. Iran experienced a wave of protests against the regime between December 2017 and January 2018. On 1 January 2018, BGPStream[27] detected that >44% of the Iranian prefixes were no longer accessible and that a majority of Iranian ASes had disappeared from BGP graphs, all of this while demonstrations were in full swing in the streets of the largest cities. This was no accident. By disabling a population's access to the Internet, Iran showed its aptitude at conducting a sophisticated form of Internet censorship. While the straightforward approach usually consists in physically disconnecting critical infrastructures, Iran's control over its network allowed a more elaborate approach based on BGP to disrupt the routing of packets.

On 26 June 2019, >80% of Iran's ISPs were disrupted and disconnected. A more thorough BGP-level analysis shows that around 8 p.m., ITC AS (AS 48159) stopped announcing externally most of its prefixes (which account for 26% of all the Iranian prefixes).[28] This led to a complete reshaping of the Iranian network with specific ASes reacting by changing their associated paths to certain prefixes. Most users within the country started noticing slow speeds and disruptions to the overall Internet. Even more interesting, some users on Twitter described the network as a 'True National Internet'[29] as connections

26 They are then often called leakages.

27 Reported by BGP Stream website as Event #2110094. https://bgpstream.com/event/210094 (26 July 2019, last accessed); Widespread Internet disruption in Iran amid geopolitical crisis. https://netblocks.org/reports/widespread-Internet-disruption-in-iran-amid-geopolitical-crisis-3AnwGkB2 (10 August 2019, last accessed).

28 BGP stream. https://bgpstream.com/(10 August 2019, last accessed).

29 IODA Signals for AS48159. https://ioda.caida.org/ioda/dashboard#view=inspect&entity=asn/48159&lastView=overview (10 August 2019, last accessed).

to local services were not as severely affected as connections to the global Internet.[30]

In November 2019, a drastic rise in fuel prices (50–200%) triggered nationwide civil protests. In the midst of popular unrest, we observed a large-scale Internet outage that lasted around 6 days and denied Internet users access to social media platforms, preventing them from sharing information about the movement. Interestingly, during this blackout, a large set of government-related services along with strategic economic services such as banking remained operational, while the Internet was completely unavailable for civil society and some targeted geographical areas.[31]

BGP observations also show incidents that can be interpreted as hijacks, following the example of a very well-documented case of BGP hijack of YouTube by Pakistan Telecom[32] in 2008. In January 2017, a similar event happened in Iran when the Iranian ITC took part in hijacking prefixes that contained pornographic websites.[33] While BGP hijack announcements were meant for the Iranian network solely, they got out of Iran because of an ITC configuration error. The accidental announcement outside the Iranian network spread the hijack throughout the network and suspended the activity of pornographic websites on the entire Internet. This generated a reaction from BGP traffic monitoring the services, such as BGPmon[34] and Dyn,[35] which led to a correction of these advertisements. In July 2018, Iran hijacked traffic for the Telegram application from all over the world to make it transit through Iran.[36]

These examples illustrate how Iran misused BGP as a tool to intervene on the Internet and to control its contents, in a way permitted by its architecture. In recent years, Iran has developed a highly opaque Internet that facilitates outages on a national scale such as the one we have just described. Furthermore, the control of the exit points of a network makes it possible to hide most hijacking operations within the network [73]. For Iran, controls at the borders allow not only to regulate the content entering the country but also to manipulate the transit of requests. The high complexity induces an extra layer of 'thickness', which means that traffic within the border appears to be *foggy* to an external observer. In the absence of monitors located within the country's network, it becomes very tricky to get a precise idea of the dynamics behind routing modifications. Most of the major censorship events originating from Iran, such as the ones described above, were visible only due to configuration errors that spread to foreign ASes.

Finally, the last section demonstrates how Iran has been able to leverage its architecture of connectivity in order to attract regional Internet traffic, and could use its connectivity as a tool of influence.

Connectivity As a Tool of Influence?

In the decades following the Iran–Iraq war, Iran progressively acquired a favourable geopolitical situation in the Middle East. The invasion of Kuwait by Iraq, the collapse of the Soviet Union, the First and Second Gulf Wars, and the War in Afghanistan in the aftermath of the 9/11 attacks all resulted in a massive involvement of US-led coalitions in the region that largely contributed to mitigating threats to Iran from its direct neighbours. International embargoes, in addition, encouraged Iran to invest in the development of sovereign strategic assets and infrastructures, and to consolidate its regional position. The Stuxnet attack against Natanz nuclear facilities triggered Iran's interest in improving and better controlling its Internet infrastructure [62]. In this section, we argue that the development of physical and logical Internet infrastructure allows Iran to leverage its connectivity as a tool of regional influence that can be documented by a BGP data analysis.

Iran began to improve its domestic infrastructures and in particular its telecommunication backbone in the early 2010s [74]. The improvement came with the deployment of a large-scale fibre optic network, between different large cities first, and thereafter even in rural regions. Based on this network, Iran developed a fibre optic industry and an expertise in deploying long-distance cables[37]. This investment produced dense fibre connectivity networks extending all over the Iranian territory, putting the country in a strategic position to provide network connectivity to its neighbouring countries and beyond.

Currently, two major international cables cross the Iranian network: the Europe-Persia Express Gateway (EPEG) and the Trans-Asia-Europe (TAE). In addition, Iran also has access to the Fibre-Optic Link Around the Globe cable in the south via a direct connection to the UAE. The TAE cable runs from Azerbaijan to Turkmenistan through the northern part of Iran, with an extension to Georgia. Yet it has been mostly used to provide connectivity between Turkmenistan and Azerbaijan at a maximal capacity of 2.08 Gbps. The EPEG cable, with a total length of ∼10 000 km, goes from Frankfurt to Barka, in the Sultanate of Oman, via Eastern Europe, Russia, Azerbaijan, Iran and the Hormuz Strait, where it is connected to a rich network of maritime cables. This cable is owned by a consortium of four carriers: Cable & Wireless, Rostelecom, Omantel and TIC (Telecommunication Infrastructure Company of Islamic Republic of Iran). The maximal capacity of this cable is 3.2 Tbps but its current operational capacity is 500 Gbps.[38] This cable is very interesting from a geopolitical perspective as it is currently the only viable Internet traffic transit route path from East Asia to Europe that represents an alternative to crossing the Red Sea and the Suez Canal. It is also shorter in length and can decrease by 10 ms. the delay between Tokyo and Frankfurt. It is noteworthy that the cable became operational in 2013, 2 years only after the signature of the memorandum of understanding[39] establishing it, showing the maturity of the Iranian fibre network. This cable is currently operational and its bandwidth has gradually increased. There have been discussions with Qatar to provide it with Internet connectivity that will not depend on the other Gulf countries in the context of the embargo imposed by Saudi Arabia for over 3 years against the emirate. Figure 3 shows the central presence of Delta Telecom Ltd (AS29049), the Azeri entry point to EPEG and TAE, and therefore the importance of the EPEG and TAE cables.

30 Twitter. https://twitter.com/Pouyan_01001010/status/1143989760299585541 (10 August 2019, last accessed).

31 BGPMon. https://bgpmon.net ; IODA CAIDA.

32 Article 19, Iran: Tightening the Net 2020. After Blood and Shutdowns, September 2020. https://www.article19.org/ttn-iran-november-shutdown/

33 Youtube Hijacking: A RIPE NCC RIS case study. https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study (10 August 2019, last accessed).

34 Oracle: Strategic Acquisition. https://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/(10 August 2019, last accessed).

35 BGPMon. https://bgpmon.net (10 August 2019, last accessed).

36 Oracle: Internet Intelligence. https://dyn.com/monitoring-analytics/(10 August 2019, last accessed).

37 Howell O'Neill, P. Telegram traffic from around the world took a detour through Iran. *Cyberscoop* (30 July 2018).

38 Fibre Optic: Iran will deploy a 14 000 km Fibre-optic Network within two months. http://www.fiberopticom.com/news/iran-will-deploy-a-14–000km-fibre-optic-networ-21105894.html (14 August 2019, last accessed).

39 Vodafone: EPEG access points. https://www.vodafone.com/business/carrier-services/connectivity/submarine-terrestrial-cable/EPEG (14 August 2019, last accessed).

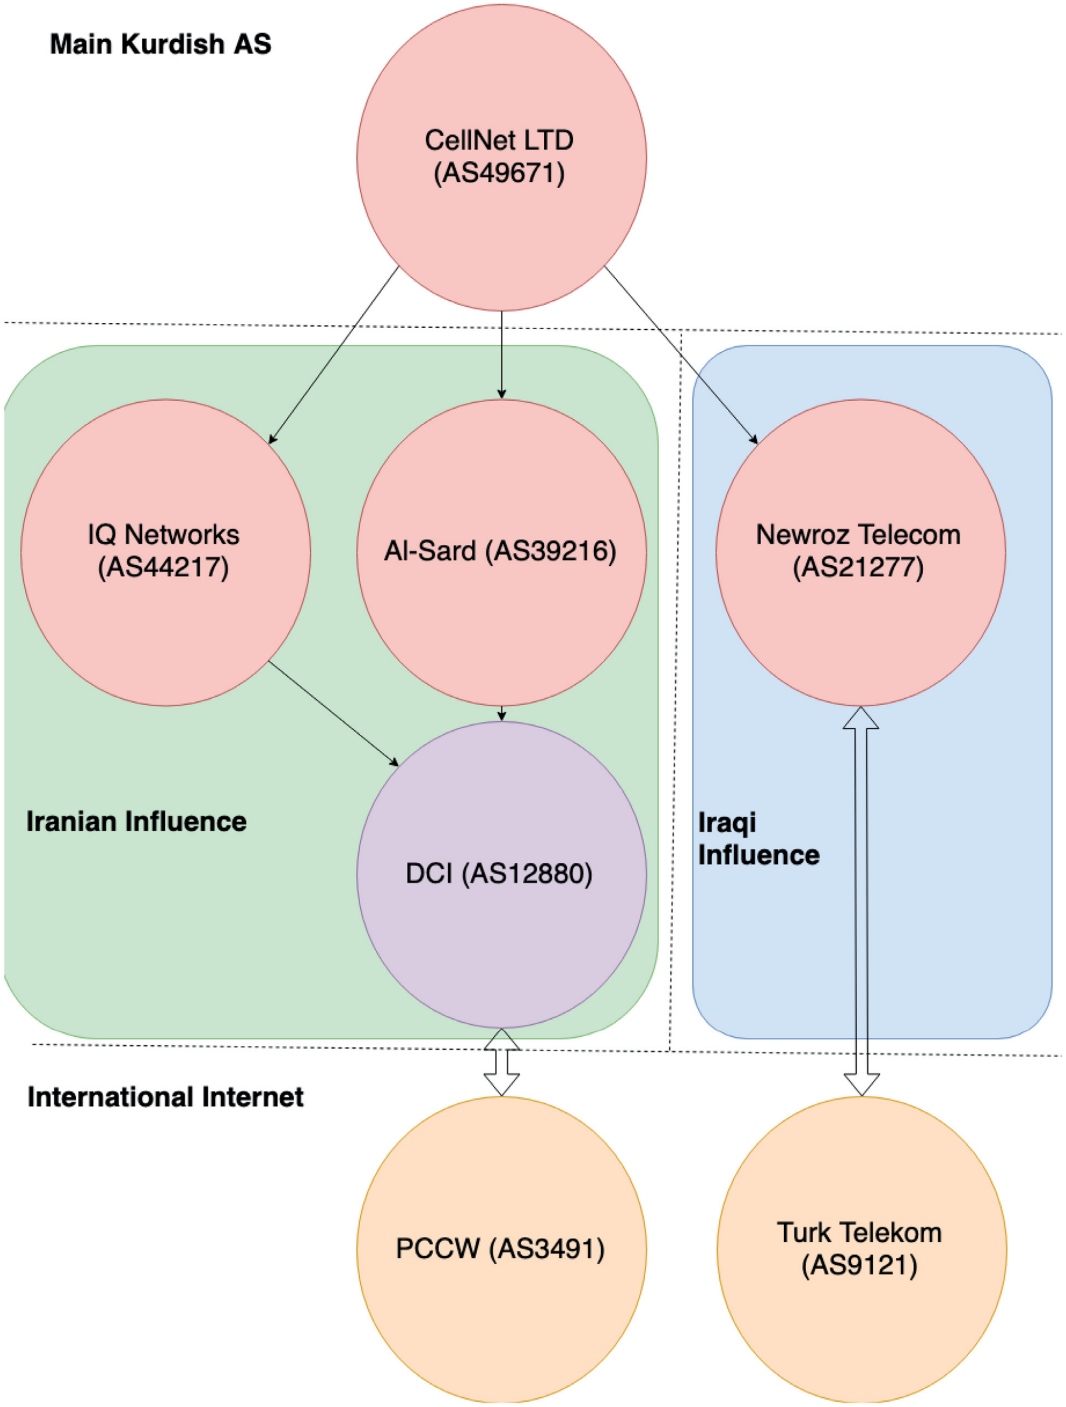

Figure 8: Representation of connectivity among the Iraqi Kurdish ASes Al-Sard and IQ Networks, both using the Iranian AS ITC as an intermediary towards the rest of the Internet [75].

The deployment of international cables in Iran has multiple benefits. First, it provides Iran with the international connectivity it needs. Moreover, this investment brings direct economic benefits to Iran's TIC but also indirect economic growth to the country through the development of a digital ecosystem (e.g. datacentres, access infrastructures crossing Iranian ASes, engineering services, etc.).

The third benefit is strategic. The traffic flowing through EPEG is, for a notable part of its path, under the direct control of Iran, making it possible for the country to observe, monitor and interact with the data. Even if data encryption may prevent access to the full content of the traffic crossing Iranian territory, metadata remains accessible. Moreover, the ability to interact (i.e. block or disturb) with the traffic gives an edge to the actor that upstreams the traffic. This strategic advantage creates an incentive for countries to deploy Internet terrestrial cables through their territory and attract international traffic.

Iran has leveraged the physical and logical structure of its network to attract traffic from its neighbouring countries, which can be

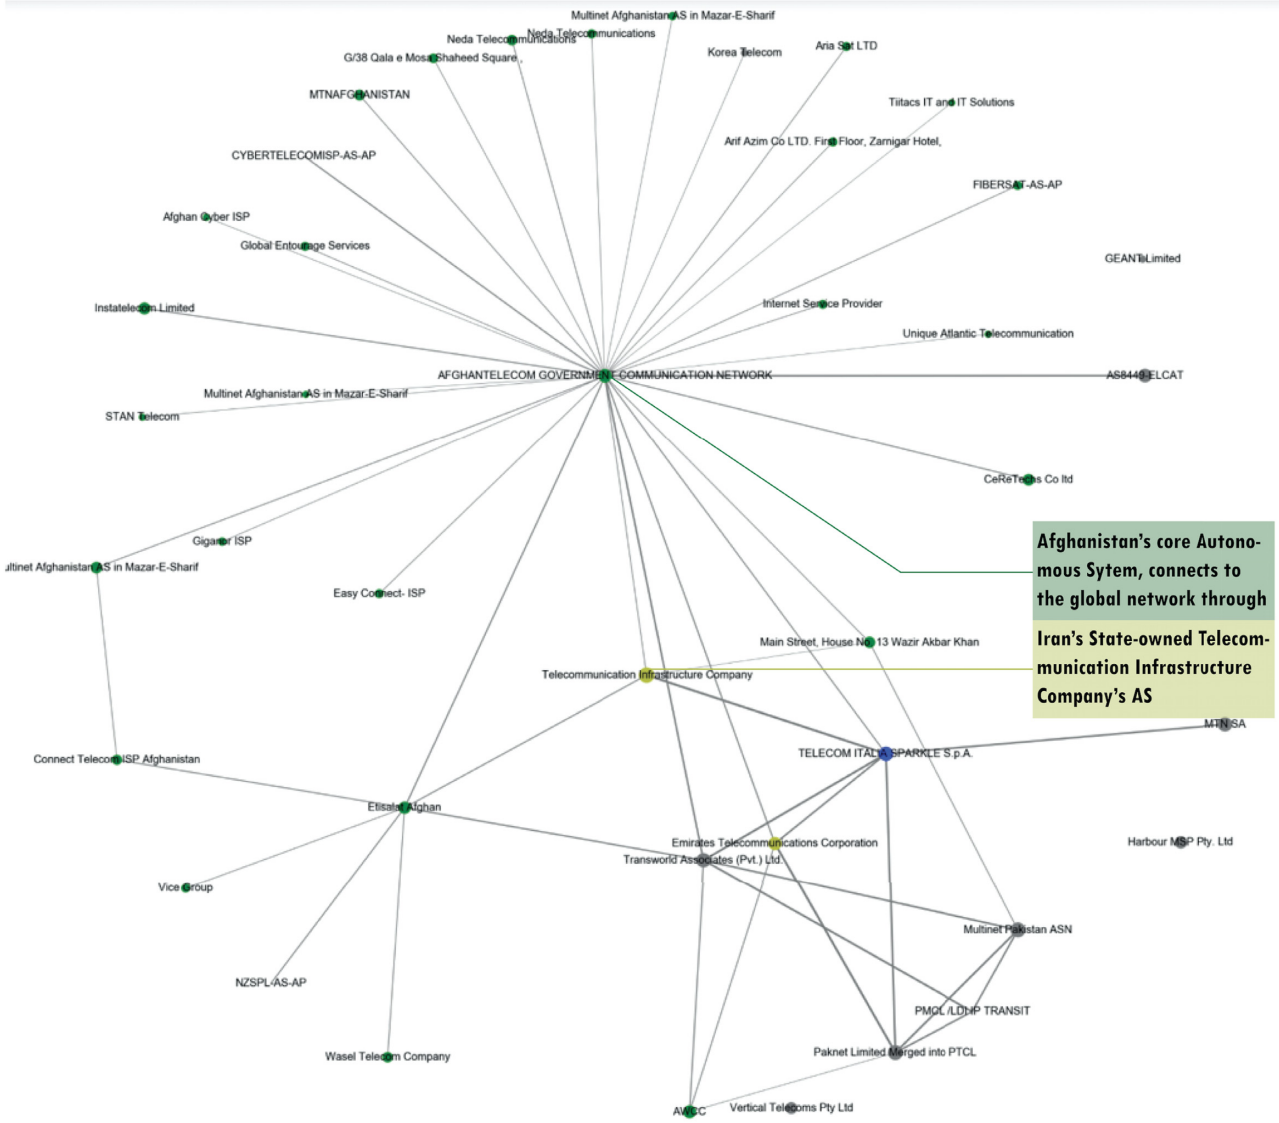

Figure 9: Reduced representation of the Afghanistan graph in February 2018.

documented by a BGP feed analysis. For example, the Iraqi Kurdistan depended on connectivity from Iran for its access to the Internet between 2010 and 2012. This region in the north of Iraq gained a large level of autonomy from Baghdad after the Second Gulf War [76] and decided that its Internet traffic should go through Iran. This decision was partly motivated by the desire to avoid transiting directly through Turkey and to follow a geographically closer path. This resulted in the emergence of new paths originating from Kurdish ASes, going through Iranian ASes to access the global Internet. We show in Fig. 8 some of these paths. The situation has since evolved and the majority of the Kurdish traffic transits now through Nowruz, an Iraqi AS, and Turkey. Yet part of the traffic still transits through Iran to access Azeri ASes.[40] The evolution of routing in the Iraqi Kurdistan sheds light on how the geopolitical context shapes connectivity and reciprocally how the AS-level connectivity can be used as a tool in power relationships within the region.

Until early 2018, Afghanistan, another neighbour of Iran, was in a situation of dependency toward Iranian connectivity. Its network was highly centralized around the ASes Afghan Telecom AS59295 and Etisalat AS131284 and it accessed the global Internet through the main Iranian AS (ITC) and few other international ASes (specifically Emirates Telecom Corporations and Telecom Italia), as shown in Fig. 9. While the architecture of the Afghani domestic network was not entirely dependent on the Iranian network, the failure of an AS or a blockage of some traffic by Iran could have led to a partial blackout of the Afghani Internet. Furthermore, TIC is the only international AS connected to the two main afghan ASes. Thus, Iran plays a fundamental role in granting some of its economic partners access to the global Internet and can use this access as a tool of influence.

Conclusion

This paper offers new perspectives on the use of BGP, routing policies and network architectures to understand the strategies developed by

40 Ministry Of ICT Telecommunication Infrastructure Company: Memorandum of Understanding signed at the sideline of EPEG quadrilateral agreement. https://www.tic.ir/en/news/1600/Memorandum-of-Understanding-signed-at-the-sideline-of-EPEG-quadrilateral-agreement (14 August 2019, last accessed).

states to 'territorialize' cyberspace. We used an in-depth analysis of the Iranian domestic network and its connectivity to the global Internet as a case study to demonstrate how an authoritarian regime could leverage BGP for geopolitical control, both domestically and regionally. The methodology we have developed for this study shows how the capture and analysis of BGP announcements can help infer and document these strategies.

We demonstrated that Iran's BGP structure is organized around three ASes controlled by the government that connect the domestic network to the global Internet, providing the government with a sophisticated 'kill switch' to fully or selectively disconnect Iran from the global Internet. These results were corroborated by the actual disconnection operated by the regime in November 2019. By limiting the number of ISPs and ASes directly connected to the outside network, Iran has therefore created a frontier between its domestic network and the rest of the Internet. These features provide the regime with strategic assets, as well as vulnerabilities.

At the same time, however, Iran has built a lively ecosystem of ASes with a rich set of paths within the country. The Iranian domestic network is therefore very resilient within its borders, due to its low level of centralization and high complexity. We compared the evolution of the Iranian network with other countries in the Middle East and observed that while Iran is not the only country to have increased controls over its domestic network, it has succeeded in increasing its internal complexity while increasing its control. This later point distinguishes Iran from other countries in the Middle East and demonstrates the existence of a strategy of control through BGP architecture.

We thereafter observed and evaluated how Iran has leveraged this domestic network structure to implement an active strategy of censorship based on BGP hijacks and outages. The complexity of its domestic network can also help opacify cyberattacks or local BGP manipulations. Through the mere observation of BGP architecture and incidents, we identified several ways in which Iran could have deliberately used cyberspace to achieve its strategic goals of asserting its own power both domestically and more widely in the Middle East. Although technical observations alone cannot tell whether all our inferred characterizations of the Iranian network resulted from a coordinated strategy from the regime, they nevertheless show interesting features of the Iranian cyberspace.

We also documented how Iran has leveraged its physical and logical network to attract traffic from its neighbours, emerging as a major connectivity provider in the Middle East. This can provide Iran with significant strategic gains in terms of ability to monitor and alter the traffic that goes through its territory.

Our contribution therefore speaks to significant geopolitical effects of decisions by the authoritarian regime in terms of the architecture of connectivity in the face of social instability and strategic competition.

What is notable about our findings is that we solely used BGP and routing information in our analysis, illustrating how these tools are relevant to develop a geopolitical analysis of cyberspace. The present methodology, applied to the Iranian cyber strategy, can be extended to other geographical contexts. Our methodology can also be enriched from other classical sources of data, like online content and fieldwork. Our methodological contribution can lead to a more comprehensive understanding of the geography of cyberspace and of the cyber strategies of Internet actors who try to exert greater control on cyberspace. As the environment is highly dynamic, longitudinal observations could provide an interesting window into the evolution of states' strategies according to their geopolitical contexts and how these routing strategies eventually contribute to shaping cyberspace.

Acknowledgements

The authors acknowledge Louis Pétiniaud for contributing to the methodology and helping with the plots and Maxime Chervaux for proofreading the document.

Conflict of interest statement. None declared.

References

1. Herold D. Escaping the world: a chinese perspective on virtual worlds. *J Virtual Worlds Res* 2012;**5**:1–16.
2. Goldberg S. Why is it taking so long to secure Internet routing? Routing security incidents can still slip past deployed security defenses. *Queue* 2014;**12**:20–33.
3. Clark D. The design philosophy of the DARPA Internet protocols. *ACM SIGCOMM Comput Commun Rev* 1988;**18**:106–14.
4. Feamster N, Ramachandra A. Understanding the network-level behavior of spammers. *ACM SIGCOMM Comput Commun Rev* 2006. doi: 10.1145/1151659.1159947.
5. Aryan S, Aryan H, Halderman JA. Internet censorship in Iran: a first look. In: *3rd Workshop on Free and Open Communications on the Internet*. Washington, DC: USENIX, 2013.
6. Vanbever L, Li O, Rexford J. *et al.* Anonymity on QuickSand: using BGP to compromise Tor. In: *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, October 27–28, 2014. HotNets-XIII, Los Angeles, CA.
7. Anderson D. SplInternet behind the great firewall of china. *ACM Queue* 2012;**10**:40–9.
8. Vervier PA, Thonnard O, Dacier M. Mind your blocks: on the stealthiness of malicious BGP hijacks. In: *Proceedings 2015 Network and Distributed System Security Symposium*, February 8–11, 2015. San Diego, CA: NDSS.
9. Butler K, Farley TR, McDaniel P. *et al.* A survey of BGP security issues and solutions. *Proc IEEE* 2009;**98**:100–22.
10. Ramachandran A, Feamster N. Understanding the network-level behavior of spammers. In: *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. New York, NY: Association for Computing Machinery, 2006, 291–302.
11. Douzet F. Understanding cyberspace with geopolitics. *Hérodote* 2014;**152–3**:3–21.
12. Douzet F. Cyberspace: the new frontier of state power. In: Moisio S, Koch N, Jonas AEG, *et al.*(eds). *Handbook on the Changing Geographies of the State*. Chelteman: Edward Elgar, 2020, 325–39.
13. Rid T. *Cyber War Will Not Take Place*. Oxford: Oxford University Press, 2013.
14. Maurer T. *Cyber Mercenaries*. Cambridge: Cambridge University Press, 2018.
15. Kaiser R. The birth of cyber war. *Polit Geogr* 2015:**46**: 11–20.
16. Libicki M. Cyberspace is not a warfighting domain. *J Law Policy Inf Soc* 2012;**8**:325–40.
17. Brent L. NATO's role in cyberspace. *NATO Rev* 2019.
18. Choucri N, Clark D. Who controls cyberspace? *Bull At Sci* 2013;**69**:21–31.
19. Buzan B, Wæver O, De Wilde J. *Security: A New Framework for Analysis*. Boulder: Lynne Rienner Publishers, 1998.
20. Dunn Cavelty M. *The Militarisation of Cyberspace: Why Less May Be Better*. In: 4th International Conference on Cyber Conflict 2013, Tallinn, Estonia.
21. Mueller M. *Will the Internet Fragment? Sovereignty, Globalization and Cyberspace*. Hoboken: Wiley, 2017.
22. Lambach D. The territorialization of cyberspace. *Int Studies Rev* 2020/09;**22**:482–506.
23. Deibert RJ. The geopolitics of Internet control: censorship, sovereignty, and cyberspace. In: Chadwick, A, Howard, P-N (eds). *Routledge Handbook of Internet Politics*. London: Routledge.

24. Denardis L. *The Global War for Internet Governance*. New Haven: Yale University Press, 2014.
25. Mueller M. *Networks and States: The Global Politics of Internet Governance*. Cambridge: MIT Press, 2010.
26. Choucri N, Clark D. *International Relations in the Cyber Age. The Co-Evolution Dilemma*. Cambridge: MIT Press, 2019.
27. Mathew AJ. The myth of the decentralised Internet. *Int Policy Rev* 2016;**5**. doi: 10.14763/2016.3.425.
28. Fidler B. The evolution of Internet routing: technical roots of the network society. *Internet Hist* 2019;**3**:364–87.
29. Gao L, Rexford J. Stable Internet routing without global coordination. *IEEE/ACM Trans Netw* 2000;**9**: 307–17.
30. Wang N. *et al.* An overview of routing optimization for Internet traffic engineering. *IEEE Commun Surv Tutor* 2008;**10**: 36–56.
31. Gao L. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans Netw* 2001;**9**:733–45.
32. Edmundson A, Ensafi R, Feamster N. *et al.* Nation-state hegemony in Internet routing. In: *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*. New York, NY, USA: Association for Computing Machinery, 2018, 1–11.
33. Karlin J, Forrest S, Rexford J. Nation-state routing: censorship, wiretapping, and BGP. *ArXiv*, 2009. abs/0903.3218.
34. Wählisch M, Schmidt TC, de Brün M. *et al.* Exposing a nation-centric view on the German Internet: a change in perspective on AS-level. In: *Passive and Active Measurement*. Berlin, Heidelberg: Springer, 2012, 200–10.
35. Lacoste Y. *Dictionnaire de géopolitique*. Paris: Larousse, 1993.
36. Dodge M, Kitchin R. *Mapping Cyberspace*. London: Routledge, 2003
37. Musiani F, Cogburn DL, DeNardis L, Levinson NS. *The Turn to Infrastructure in Internet Governance*. Houndmills: Palgrave Macmillan, 2016.
38. Faravelon A, Frénot S, Grumbach S. Chasing data in the intermediation era: economy and security at stakes. *IEEE Secur Priv* 2016;**14**:22–31.
39. Limonier K, Gérard C. Guerre hybride russe dans le cyberespace (fr). *Herodote* 2017;**166-167**:145–63.
40. Howard PN, Ganesh B, Liotsiou D. *et al.* The IRA, social media and political polarization in the United States, 2012–2018. Project on Computational Propaganda, University of Oxford, 2018.
41. Douzet F, Pietiniaud L, Salamatian L. *et al*. Measuring the fragmentation of the Internet: the case of the Border Gateway Protocol (BGP) during the Ukrainian crisis. In: *2020 12th International Conference on Cyber Conflict 20/20 Vision: The Next 2020*. Tallinn, Estonia: NATO CCDCOE.
42. Limonier K, Douzet F, Pétiniaud L. *et al*. Mapping the routes of the Internet for geopolitics: the case of Eastern Ukraine. *First Monday*, 2021. doi: 10.5210/fm.v26i5.11700.
43. Moghanizadeh S. *The role of social media in Iran's Green Movement*. Ph.D. Thesis. University of Goteborg, 2013.
44. Farahan J, Mirrafi A. Presentation of the civil defense strategy against cyber threads. *J Strateg Def Stud* 2019;**75**:259–28.
45. Gelvanovska N, Rogy M, Rossotto CM. Broadband networks in the Middle East and North Africa: accelerating high-speed Internet access. *World Bank* 2014, doi: 10.1596/978-1-4648-0112-9.
46. Giblin B. L'Iran : un acteur majeur au moyen-orient (fr). *Herodote* 2018;**169**:3–13.
47. Leguay J, Latapy M, Friedman T. *et al.* Describing and simulating Internet routes. In: *NETWORKING 2005. Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications Systems*. Berlin, Heidelberg: Springer, 2005, 659–70.
48. Chiu YC, Schlinker B, Balaji Radhakrishnan A. *et al*.In: *ACM Internet Measurement Conference – IMC*. Tokyo, Japan: ACM SIGCOMM, 2015.
49. Benton K, Camp LJ. Firewalling scenic routes: preventing data exfiltration via political and geographic routing policies. In: *Proceedings of the 2016 ACM Workshop on Automated Decision Making for Active Cyber Defense*. New York, NY, USA: Association for Computing Machinery, 2016.
50. Allen J. Topological twists: power's shifting geographies. *Dialogues Hum Geogr* 2011;**1**:283–98.
51. Luckier M, Huffaker B, Dhamdhere A. *et al.* As relationships, customer cones, and validation. In: *Proceedings of the 2013 Internet Measurement Conference – IMC*. Barcelona, Spain, 2013.
52. Robine J, Salamatian K. Peut-on penser une cybergéographie (fr)? *Hérodote* 2014;**123**:152–3
53. Roughan M, Willinger W, Maennel O. *et al*. 10 lessons from 10 years of measuring and modeling the Internet's autonomous systems. *IEEE J Sel Areas Commun* 2011;**29**:1810–21.
54. Salamatian K, Kaafar D, Salamatian L. A geometric approach for real-time monitoring of dynamic large scale graphs: AS-level graphs illustrated. *Cornell University*, 2018, doi: 10.1109/JSAC.2011.111006.
55. Orsini C, King A, Giordano D. *et al.* BGPStream: a software framework for live and historical BGP data analysis. In: *Proceedings of the 2016 Internet Measurement Conference*. New York, NY: Association for Computing Machinery, 2016.
56. Marder A, Luckie M, Dhamdhere A. *et al.* Pushing the boundaries with bdrmapIT: mapping router ownership at Internet scale. In: *Proceedings of the Internet Measurement Conference*. New York, NY: Association for Computing Machinery, 2018.
57. Ager B, Chatzis N, Feldmann A. *et al.* 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: Association for Computing Machinery, 2012.
58. Cohen R, Raz D. 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*. Barcelona, Spain: Institute of Electrical and Electronics Engineers, 2006.
59. Gregori E, Improta A, Lenzini L. *et al.* On the incompleteness of the AS-level graph: a novel methodology for BGP route collector placement. In: *Proceedings of the 2012 Internet Measurement Conference*. New York, NY: Association for Computing Machinery, 2012.
60. Roberts H, Larochelle D, Faris R. *et al. Mapping local Internet control*. Harvard: Berkman Klein Center, 2011.
61. Jacomy M, Venturini T, Heymann S. *et al.* ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLoS One* 2014;**9**:e98679.
62. Alimardani M. Stuxnet, American sanctions, and cyberwar are legitimizing Iranian Internet controls. *VICE*. 2019, https://www.vice.com/en/article/vb9859/stuxnet-american-sanctions-and-cyberwar-are-legitimizing-iranian-internet-controls, (accessed August 10, 2021).
63. Marchant J, Robertson B. Chaos & control: the competing tensions of Internet governance in Iran. *Internet Policy Observatory* 2015. https://repository.upenn.edu/cgi/viewcontent.cgiarticle=1014&context=internetpolicyobservatory, (1 August 2021, date last accessed).
64. Shaheed A. Layers of Internet censorship in Iran. 2014. https://www.shaheedoniran.org/english/blog/layers-of-internet-censorship-in-iran/. (1 August 2021, date last accessed).
65. Tor Blog. *New Blocking Activity from Iran*. https://blog.torproject.org/new-blocking-activity-iran (15 October 2019, date last accessed).
66. Brumfiel G. Publishers split over response to US trade embargo ruling. *Nature* 2004;**427**:663.
67. Raadi M. If you don't know, now you know: GitHub is restricting access for users from Iran and a few other embargoed countries. *DEV*. 2019. https://dev.to/mjraadi/if-you-don-t-know-now-you-know-github-is-restricting-access-for-users-from-iran-and-a-few-other-embargoed-countries-5ga9 (accessed August 10, 2021).
68. Levis P. The collateral damage of Internet censorship by DNS injection. *SIGCOMM Comput Commun Revi* 2012;**42**:21–7.
69. Chen T, Wang V. Web filtering and censoring. *Computer* 2010;**43**:94–7
70. Bendrath R, Mueller M. The end of the net as we know it? Deep packet inspection and Internet governance. *New Media Soc* 2011;**13**:1142–60.
71. Dainotti A, Squarcella C, Aben E. *et al.* Analysis of country-wide Internet outages caused by censorship. In: *Proceedings of the 2011 ACM SIGCOMM Internet Measurement Conference*. New York, NY: Association for Computing Machinery, 2011.
72. Chaabane A, Chen T, Cunche M. *et al.* Censorship in the wild: analyzing Internet filtering in Syria. In: *Proceedings of the 2014 Internet Measurement Conference*. New York, NY: Association for Computing Machinery, 2014.

73. Madory D. Iran leaks censorship via BGP hijacks. *Oracle DYN*. 2017. https://blogs.oracle.com/internetintelligence/iran-leaks-censorship-via-bgp-hijacks-v3 (accessed August 10, 2021).
74. Barzegar K. Iran's foreign policy in post invasion Iraq. *Middle East Policy* 2008;**15**:47–58.
75. Cowie J. Iran: exporting the Internet. *Oracle DYN* 2010. https://blogs.oracle.com/internetintelligence/iran%3a-exporting-the-internet-1 (accessed August 10, 2021).
76. Gunter MM. A de facto Kurdish state in Northern Iraq. *Third World Quarterly* 1993;**14**:295–319.