Military Cyber Affairs
The Journal of the Military Cyber Professionals Association
ISSN: 2378-0789

Volume 3 | Issue 1 Article 7

2018

China's Maxim – Leave No Access Point Unexploited: The Hidden Story of China Telecom's BGP Hijacking

Chris C. Demchak
U.S. Naval War College, chris.demchak@usnwc.edu

Yuval Shavitt
Tel Aviv University, shavitt@eng.tau.ac.il

China's Maxim – Leave No Access Point Unexploited: *The Hidden Story of China Telecom's BGP Hijacking*

Chris C. Demchak[1]
Yuval Shavitt[2]

China's Workaround: Hijacking Internet Traffic not covered by the anti-theft 2015 Xi-Obama Agreement

Surprisingly, the voluntary 2015 Xi-Obama agreement stopping military forces from hacking commercial enterprises for economic gain did appear to reduce Chinese theft from western targets. China's technological development process, however, was still dependent on massive expropriation of foreign R&D. This necessitated new ways to get information while still technically adhering to the agreement. Since the agreement only covered military activities, Chinese corporate state champions could be tasked with taking up the slack. But even Chinese multinationals, such as Huawei or ZTE, were already being viewed with suspicion. Instead, the government opted to leverage a seemingly innocuous player – one that is normally viewed as a passive service provider – to target the foundational infrastructure of the internet to bypass the agreement, avoid detection, and provide the necessary access to information.

Enter China Telecom[3] – a large state champion telecommunications company. While the 2015 agreement prohibited direct attacks on computer networks, it did nothing to prevent the hijacking of the vital internet backbone of western countries. Conveniently, China Telecom has ten strategically placed, Chinese controlled internet 'points of presence'[4] (PoPs) across the internet backbone of North America. Vast rewards can be reaped from the hijacking, diverting, and then copying of information-rich traffic going into or crossing the United States and Canada – often unnoticed and then delivered with only small delays.

This article will show how this hijacking works, and how China employs its conveniently distributed points of presence (PoPs) in western democracies' telecommunications systems to redirect internet traffic through China for malicious use. It will show the actual routing paths, give a summary of how one hijacks parts of the internet by inserting these nodes, and outline the major security implications.

---

[1] Dr. Chris C. Demchak is the RDML Grace M. Hopper Chair of Cyber Security at the U.S. Naval War College and Director, Center of Cyber Conflict Studies, US Naval War College.
[2] Dr. Yuval Shavitt is a Professor of Electrical Engineering at Tel Aviv University and a member of its Blavatnik Interdisciplinary Cyber Research Center. He is also the CTO and original founder of BGProtect LTD.
[3] China Telecom owns ChinaNet in America.
[4] A 'point-of-presence' (PoP) is a major point of connection where a long-distance telecommunications carrier such as Verizon or British Telecom connects to a local network and picks up the local traffic – or transit traffic – to move it onwards towards its various destimations.

These Chinese PoPs are found all over the world including Europe and Asia. The prevalence of and demonstrated the ease with which one can simply redirect and copy data by controlling key transit nodes buried in a nation's infrastructure requires an urgent policy response. To that end, we recommend an 'Access Reciprocity' strategy for vulnerable democracies – one that is then collectively coordinated across allies. The goal is to restrict China's internet hijacking options and fix the imbalance in information access and potential losses. Any single nation can unilaterally pursue this policy, but it will take the sum of democratic civil societies to have the scale to effectively deter this behavior over the longer term.

How to Hijack the Net

Successfully hijacking the net requires understanding how to manipulate key structures in contractual and regulatory agreements about who moves information packets to whom across the internet. The Internet consists of tens of thousands of independently managed networks, interconnected through contractual peer-or-pay arrangements by which the data packets are exchanged. Each of these networks is called an 'Autonomous System' (AS), meaning that network independently controls the access to and from all its internal network nodes. Users inside that AS connect to other users in other networks through that AS' own gateway servers. A good example is a university's own network whose students connect by routers to other students staying wholly inside the university's 'intranet' or to others globally by passing the university's gateway servers to reach the wider internet.

For data traffic to move, addresses of senders and recipients are needed. These ASs are each assigned a unique 'Autonomous System Number' (ASN) to identify itself globally for receipt of information packets. Each AS controls a set of 'internet protocol' (IP) addresses assigned in blocks of consecutive numbers.[5] These blocks are assigned much like telephone number area codes; for example, blocks in the US are now regulated in the US by the Federal Communications Commission (FCC). If the AS is also an Internet Service Provider (ISP), it then further assigns some of the individual IP addresses it manages to home customers and others in chunks of an address block to business customers. Examples of ASNs are AS3356, which belongs to Level3, a tier-1 Internet Service Provider (ISP); AS5400, which belongs to British Telecom, a tier-2 ISP; AS8551, which belongs to Bezeq International, a Israeli ISP; AS25046, which belongs to Check Point Software, a leading cybersecurity company; and AS15169, which belongs to Google.

In the internet, information is sent across intervening ASs as small data 'packets' with their destination IP addresses attached. Each router in the transited networks looks at the destination IP address in the packet and forwards it to the next and closest AS according to a 'forwarding table.' The 'glue' holding the Internet together uses two forms of software 'protocols'- the Internet Protocol (IP) [RFC971] and the Border Gateway Protocol (BGP) [RFC 4271]. The IP defines how information is exchanged between end systems at the network level and requires that every device connected to the Internet (such as a computer or a router) will have a unique global address, its IP address. The source and destination IP addresses are placed in each packet of information which is sent out

[5]Internet Protocol (IP) are assigned to an Autonomous System (AS) by its Regional Internet Registry (RIR) such as ARIN in North America or APNIC in the Asia Pacific. The RIRs, in turn, receive their regional blocks of IP addresses from the Internet Assigned Numbers Authority (IANA) which is a department of the nonprofit Internet Corporation for Assigned Names and Numbers (ICANN).

DOI: https://doi.org/10.5038/2378-0789.3.1.1050

across the internet through the network of interconnected ASs. The process is similar to how letters have to and from addresses and are moved between post offices and hubs before reaching their destination.

Occupying critical nodes at the top of the global internet data exchange system are the 'tier 1' providers whose influence in the paths taken by information flows can be enormous. The global internet's information exchange has never been free; the entire structure is a variable peer-or-pay system. A small number of the very large ASs form the 'tier 1' or 'backbone' set of global 'peers' who contract among each other to share massive volumes of traffic reciprocally without paying transit fees. The tier 1 set of global peers may each have more than one ASN as part of their holdings. For example, Verizon Enterprise Solutions (formerly UUNET (MCI) and XO Communications) has over a dozen ASNs (e.g., AS701, AS702, AS703, AS2828). All other ASs must pay for – or specially negotiate – packet traffic transiting arrangements. The long-distance carriers – i.e., the telecommunications corporations or agencies – own and operate the major PoPs connecting the traffic across all the ASs, and thereby control the major nodes of the entire internet traffic flow.

While the paths built for any set of messages across ASNs are based on multiple economic and engineering criteria, a key requirement is to select the shortest route to its destination IP address. Critical to moving traffic across the sea of tier 1 and other ASNs are the 'forwarding tables' which show the next – and closest – AS router for a given packet to cross. The servers hosting the 'Border Gateway Protocol' (BGP) – the key Internet routing protocol – build these forwarding tables which are shared across each contributing AS. Within the BGP forwarding tables, administrators of each AS announce to their AS neighbors the IP address blocks that their AS owns, whether to be used as a destination or a convenient transit node.

Errors can occur given the complexity of configuring BGP, and these possible errors offer covert actors a number of hijack opportunities. If network AS1 mistakenly announces through its BGP that it owns an IP block that actually is owned by network AS2, traffic from a portion of the Internet destined for AS2 will actually be routed to – and through – AS1. If the erroneous announcement was maliciously arranged, then a ***BGP hijack*** has occurred. The amount of traffic routed from AS1 to AS2 depends on a variety of factors, and it can have almost global effects. A fundamental presumption behind the current internet protocol is that geography and physics still matter. The routing is biased to shorter routes simply because the transfer of electrons across a longer distance takes more time and incurs greater risk of routine and basic distortions in the data. Thus, if a routing table falsely specifies what the shortest distance is, the data will automatically attempt to move that way.

Building a successful BGP hijack attack is complex, but much easier with the support of a complicit and preferably largescale ISP that is more likely to be included as a central transit point among a sea of ASs. As a result, today most BGP hijacks are the work of government agencies or large transnational criminal organizations with access to, leverage over, or control of strategically placed ISPs. For example, in 2008, Pakistan Telecom – the tier 1 AS for Pakistan – accidentally hijacked all Youtube traffic for several hours as administrators make mistakes in using routing to censor a clip considered non-Islamic. Two years later, on April 8th, 2010 China Telecom hijacked 15% of the Internet traffic for 18 minutes in what is believed to be both a large-scale experiment and a demonstration of Chinese capabilities in controlling the flows of the internet.

Over the past few years, researchers at BGProtect LTD based on the DIMES project [DIMES] at the Tel Aviv University built a route tracing system monitoring the BGP announcements and distinguishing patterns suggesting accidental or deliberate hijacking[6] across many routes simultaneously and with a granularity down to the individual city. Using this technique, the two authors of this paper noticed unusual and systematic hijacking patterns associated with China Telecom.

The Security Implications

Hijack attacks expose a network to potentially critical damage because it is not a hack of the endpoint but of the critical exchanges carrying information between endpoints. The rerouted traffic flows sensitive data across the collection points of an intervening adversary without any human clicking on suspicious links or a network administrator seeing any surges in unexplained data transfers. This gives the malicious attacker access to the organization's network, to stealing valuable data, adding malicious implants to seemingly normal traffic, or simply modifying or corrupting valuable data. If diverted and copied for even small amounts of time, even encrypted traffic can be broken, as shown in the well known, recent 'DROWN' and 'Logjam' encryption attacks.

A man-in-the-middle (MITM) attack can neutralize an organization's firewall, for example. In this form of attack, a bad actor inserts its covert collection method between the sender and the real desired destination, between the endpoints. For another example, with the traffic rerouted into an adversary's cache, the attacker can learn enough to impersonate trusted sources in or to the attacked network, especially valuable in obtaining validated certificates. The data can be used for widely successful phishing attempts through email, voice, or texting attacks. [Rexford] Impersonation attacks can allow the malicious attacker to harvest passwords of the company's web users. With those keys to the victim's network in hand, attackers can distort, disconnect, or destroy any part of the company's network accessible from the Internet, increasingly to include critical financial and physical systems and their backups.

Despite all the discussion of how geography has been defeated by the global cyberspace, the closer a network is to the attacker or its complicit ISP, the more likely an attack will succeed because defending administrators are less likely to have enough time to detect, analyze, and mitigate the attack. Thus, if an attacker wants its attack to be more potent, they need to use a network that has a global presence, or in other words, a network that is not too far from any potential victim network. For a government, the wider the geographical spread of its own and controlled networks, the more their global reach can help with orchestrating such attacks.

China Telecom Well Placed in North America

China Telecom (CT) entered North American networks at the beginning of the 2000s and has since grown to have 10 PoPs, eight in the US and two in Canada, spanning both coasts and all the major

[6] No technical details will be provided in this piece. For more technical information, contact Dr. Yuval Shavitt, Tel Aviv University.

DOI: https://doi.org/10.5038/2378-0789.3.1.1050

exchange points in the US. Few other non-American ISPs has such a wide-spread presence on US soil.

Figure 1: China Telecom large presence in North America (image taken from the CT web site)

Using these numerous PoPs, CT has already relatively seamlessly hijacked the domestic US and cross-US traffic and redirected it to China over days, weeks, and months as demonstrated in the examples below. The patterns of traffic revealed in traceroute research[7] suggest repetitive IP hijack attacks committed by China Telecom. While one may argue such attacks can always be explained by 'normal' BGP behavior, these, in particular, suggest malicious intent, precisely because of their unusual transit characteristics – namely the lengthened routes and the abnormal durations. The following are a set of such unusual cases.

Canada to Korea, 2016 – traffic to Government Site

Starting from February 2016 and for about 6 months, routes from Canada to Korean government sites were hijacked by China Telecom and routed through China. Figure 2a shows the shortest and normal route: Canada-US-Korea. As shown in figure 2b, however, the hijacked route started at the

[7] Traceroute research involves tracing the routes along which traffic is sent across the internet and uses a variety of data sources including especially the globally published routing tables. The process involves acquisition and analysis of enormous quantities of traffic data.

China Telecom PoP in Toronto, the traffic was then forwarded inside the Chinese network to their PoP on the US West Coast, from there to China, and finally to delivery in Korea. This is a perfect scenario for long-term espionage, where the victim's local protections won't raise alarms about the long-term traffic detours. Note that the shortest route between the originators and the destination is definitely not through two China Telcom PoPs in North America to China and only then to Korea. That this pattern continued for six months is good evidence that this was no short-term misconfiguration or temporary internet conditions disruption. This attack repeated later for shorter time durations.

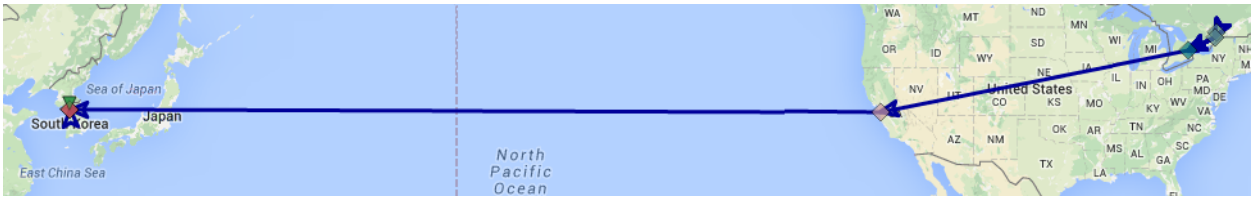

Figure 2a: The normal and shortest route from Canada to Korea before the attack.

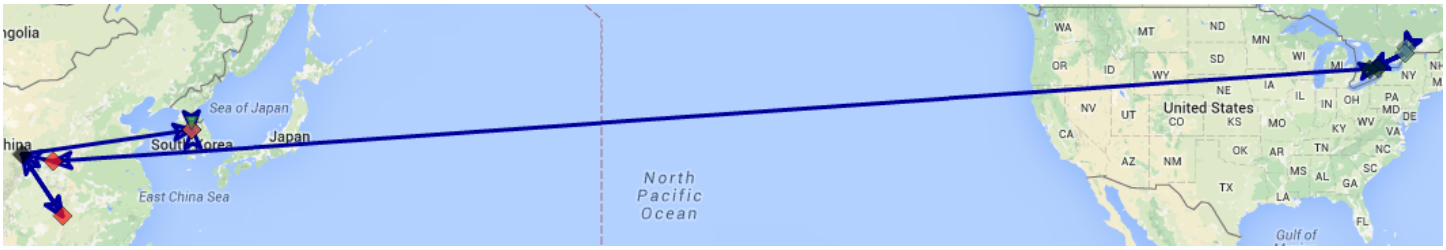

Figure 3b: The hijacked route through the CT POP in Maryland – long way from Canada to Korea .

US to Italy Oct. 2016 – Banking and Money

On October 2016, traffic from several locations in the USA to a large Anglo-American bank headquarters in Milan, Italy was hijacked by China Telecom to China. The normal route is shown in figure 3a and the hijacked route in figure 3b. The attack started at the ChinaNet[8] PoP near Los Angeles and, while it lasted for 9 hours, it did not seem well planned. ChinaNet actors seemed to have difficulties in routing the traffic back to Milan. The route inside the Chinese network changed several times as the attackers worked to try and redirect the traffic back. Ultimately, they seemed to give trying to send it on, and the traffic never arrived.

[8] Wholly owned unit of China Telecom.

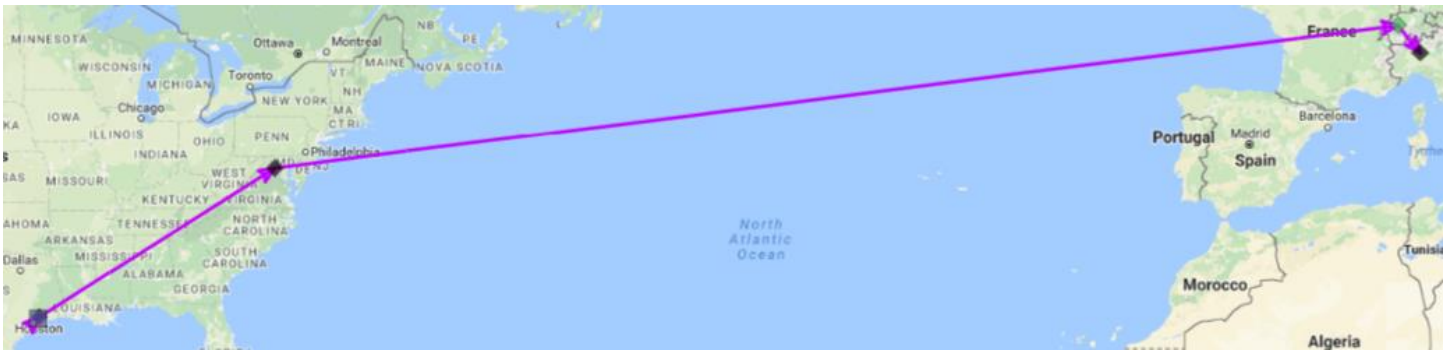

Figure 4a: US large bank to Italy normal route

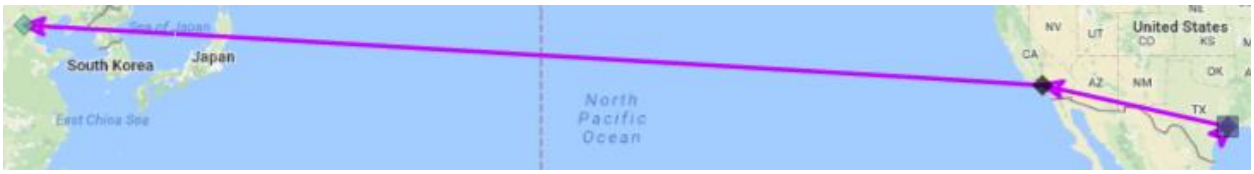

Figure 5b: US large bank to Italy but after hijack, traffic never arrives, seems to terminate in China.

Scandinavia to Japan, April-May 2017 – News

Traffic from Sweden and Norway to the Japanese network of a large American news organization was hijacked to China for about 6 weeks in April/May 2017. As shown in figure 4, the hijack started in China Telecom PoP in Maryland and forwarded to their PoP in California. From there traffic was redirected to China and then through Hong Kong to Japan. By no stretch could this period of disjointed routing have been accidental.

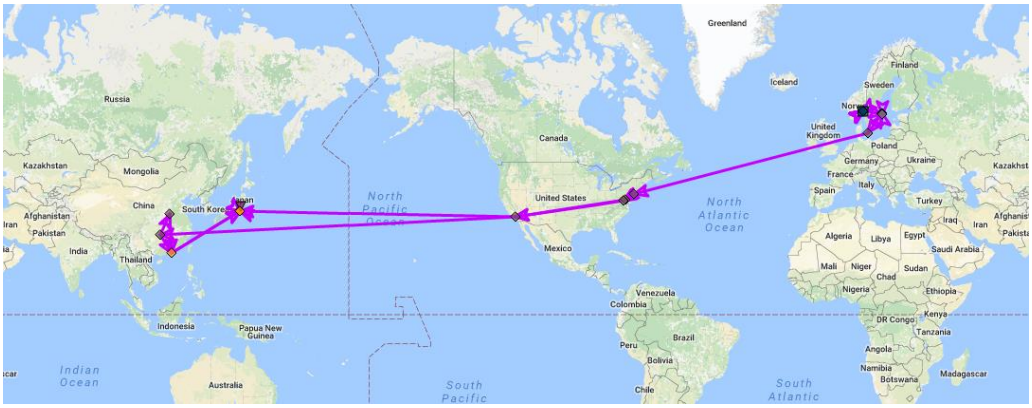

Figure 6: A deflected route from Oslo, Norway to Tokyo, Japan.

Italy to Thailand April-July 2017 – ISPs

Traffic to the mail server (and other IP addresses) of a large financial company in Thailand was hijacked several times during April, May, and July 2017.Some of the hijack attacks started in the USA. As shown in figure 5, traffic sent from Milan, Italy to Bangkok was hijacked by a ChinaNet PoP

in California. This hijack affected at least two large International American based providers: Cogent and Level3. In parallel, there was an attack on providers in South Korea.

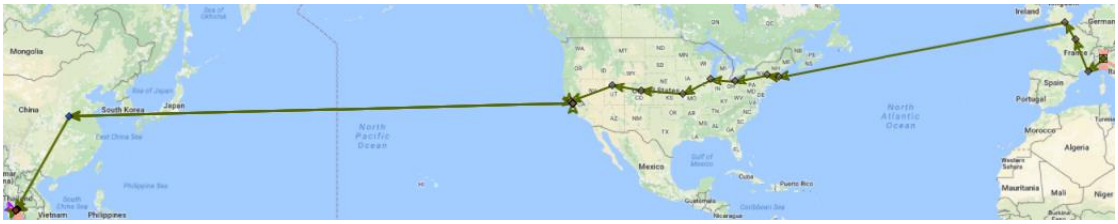

Figure 7: Traffic from Milan, Italy during hijack to China.

US Telecomms Blackballed from China – No Reciprocity

China's own national network is fairly isolated from the world, protecting it from foreign hijacking of its own domestic or transit traffic. There are, in principle, only three major internet gateways into China, located in Beijing, Shanghai, and Hong Kong. Hong Kong serves as a large international exchange, a legacy of the time it was ruled by Great Britain. Many International companies have PoPs in Hong Kong, but this network is isolated from the rest of China. In fact, the Hong Kong major internet hub presents a great opportunity for China to hijack traffic that traverses it, usually with one endpoint of the communication being in the Asia Pacific region. Elsewhere in China, US-based ISPs have no presence. AT&T has publicized that is has a presence in China, but this seems to be only in collaboration with a local player, and not an AT&T directly owned and managed operation.

The Policy of 'Access Reciprocity' to Curb Hijacks

Today China has ten POPs in North America (eight in the US and two in Canada) while the US has none in China. That imbalance in access allows for malicious behavior by China through China Telecom at a time and place of its choosing, while denying the same to the US and its allies. Note that the hijacked routes come from – or are traveling to – allied states, but the traffic stumbles on China Telecom's PoPs due to the shortest route bias in BGP rules and then is hijacked in the US by the Chinese network. If China Telecom had only one PoP – say in Los Angeles at most – then hijacking would be very difficult to achieve and to obscure from oversight. One could even argue that fairness dictates that China Telecom should not extend beyond Hong Kong unless other global peers were given equivalent access to have PoPs in China itself.

A new policy is needed: an "Access Reciprocity" policy on internet PoPs located in North America or, indeed, even with allied democratic nations. One could use many metrics to establish the PoPs allowed, including a population metric for example. That is, the US at 350 million citizens currently hosts eight China Telecom PoPs. With China at three times that population, the US Telecoms should be allowed three times that number of PoPs in China. The advantage of such a metric is that it makes evident the imbalance of one nation having multiple PoPs in another nation or region, while the latter have none and are not allowed any in the first nation. Or, if the demand for access reciprocity is refused, then an appropriate defense policy in response could state that no traffic to or

DOI: https://doi.org/10.5038/2378-0789.3.1.1050

from or across the US or ally is allowed to enter a CT PoP in the US or in the ally's networks. That policy could be inserted in BGPs routing tables as required and automatically implemented.

The advantages of a stated 'Access Reciprocity' policy is that it embodies interstate fairness, enhances the cybersecurity of the US and its allies and can be implemented into existing routing tables. Any single nation can unilaterally pursue this policy, but it will take the sum of democratic civil societies acting in agreement to have the scale to effectively deter this malicious behavior over the longer term.

Furthermore, if such an allied 'Access Reciprocity' agreement emerges in the form of coordinated national policies and institutions, the possibility arises for a regional and possibly international IT norm emerging from practice in other domains. Over time, basic reciprocal fairness in digital transnational exchanges can be viewed as desirable, clarifying, and effective in nurturing cooperation in a hostile, "asocial" global environment. [Axelrod] Imagine if reciprocal fairness included security and privacy scrutiny of a Chinese manufacturer's source code before its product or any updates may be imported into the US or its allies – as is now the law in China. More balance between democratic and authoritarian information technology systems by enforcing reciprocal fairness likely has a significant positive influence on the currently deleterious trends in international cyber insecurity. This would be the first step in making hijacking internet traffic much more difficult and costly for adversaries. If it were tied to a broader cyber operational resilience alliance (CORA) among democracies, then it provides another legal and feasible tool for use by democracies in defending their wellbeing and survival in a contested cyber domain. [CORA]

References

[RFC791] J. Postel. "Internet Protocol. RFC 791". Internet Engineering Task Force. Sep. 1981.

[RFC 4271] Y. Rekhter, T. Li, and S. Hares. "A Border Gateway Protocol 4 (BGP-4)". RFC 4271. Internet Engineering Task Force (IETF). Jan. 2006.

[DIMES] Yuval Shavitt and Eran Shir. DIMES: Let the Internet Measure Itself. ACM Computer Communications Review, 35(5):71--74, Oct. 2005.

[Rexford] Henry Birge-Lee, Yixin Sun, Annie Edmundson, Jennifer Rexford, and Prateek Mittal. "Using BGP to Acquire Bogus TLS Certificates". Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs 2017), Minneapolis, MN, USA, July 2017.

[Axelrod] Axelrod, Robert, and William D. Hamilton. 1981. "The Evolution of Cooperation". Science. 211:4489. March 27. pp.1390-1396.

[CORA] Demchak, Chris C. 2017. "Defending Democracies in a Cybered World". Brown Journal of World Affairs. 24:1. fall/winter. pp.1-19.