



MPTCP Robustness Against Large-Scale Man-in-the-Middle Attacks

Chi-Dung Phung, Benevid Felix Silva, Michele Nogueira, Stefano Secci

► To cite this version:

Chi-Dung Phung, Benevid Felix Silva, Michele Nogueira, Stefano Secci. MPTCP Robustness Against Large-Scale Man-in-the-Middle Attacks. *Computer Networks*, 2019, 164, pp.106896. 10.1016/j.comnet.2019.106896 . hal-02287761

HAL Id: hal-02287761

<https://hal.science/hal-02287761>

Submitted on 13 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MPTCP Robustness Against Large-Scale Man-in-the-Middle Attacks

Chi-Dung Phung^{a,b}, Benevid Felix Silva^c, Michele Nogueira^c, Stefano Secci^a

^a*Cnam, Cedric, Paris, France*

^b*Orange Labs, Chatillon, France*

^c*Federal University of Paraná, Paraná, Brazil*

Abstract

Multipath communications at the Internet scale have been a myth for a long time, with no actual protocol being deployed at large scale. Recently, the Multipath Transmission Control Protocol (MPTCP) extension was standardized and is undergoing rapid adoption in many different use-cases, from mobile to fixed access networks, from data-centers to core networks. Among its major benefits – i.e., reliability thanks to backup path rerouting, throughput increase thanks to link aggregation, and confidentiality being more difficult to intercept a full connection – the latter has attracted lower attention. How effective would be to use MPTCP, or an equivalent multipath transport layer protocol, to exploit multiple Internet-scale paths and decrease the probability of Man-in-the-Middle (MITM) attacks is a question which we try to answer. By analyzing the Autonomous System (AS) level graph, we identify which countries and regions show a higher level of robustness against MITM AS-level attacks, for example due to core cable tapping or route hijacking practices.¹

Keywords: MPTCP, Man-in-the-Middle attacks, communication robustness

1. Introduction

2 The Multipath Transmission Control Protocol (MPTCP) [24] is an ex-
3 tension of TCP to concurrently use multiple network paths for a given con-
4 nection. Among many proposals to support these features at the transport

¹A preliminary version of the content of this paper was presented in [43].

layer, it is considered as the one having attracted the largest interest and deployment [44]. One of the main reasons for this success is the incremental deployability adopted in its design, with the required signaling transparently reusing existing features of the TCP options.

MPTCP employs multiple ‘subflows’ to route traffic from a source to a destination in an IP network via different network interfaces and/or TCP ports at the transmitting and/or receiving endpoints. Subflow IP traffic can then be routed independently in the network segment. However, besides the usage of multiple network interfaces at the source or destination, the presence of flow-level load-balancers sensible to port numbers, or multipath proxies aware of the network topology [8] can differentiate the route followed by the subflow packets.

MPTCP is being adopted by major operating systems; it is already happening for Apple OSX and IOS, where it is used for some applications. Its integration in the mainstream Linux kernel is expected for the upcoming versions [48]. Among the motivations pushed forward in support of MPTCP, there are [40]: (i) bandwidth aggregation, i.e., the increased network bandwidth offered to a connection; (ii) connection reliability, i.e., the possibility to use an alternative path in case of failure along the primary path or at the primary network interface level; (iii) communication confidentiality, i.e., the decreased ability for a Man-in-the-Middle (MITM) attacker to intercept all the traffic of a same connection.

While the first two aspects above have been largely explored in the last decade, the latter was marginally studied to date. In this paper, we report the results of an extensive measurement campaign aimed at assessing the degree of confidentiality one can expect using MPTCP. In particular, we focus on confidentiality from large-scale, i.e., Autonomous System (AS) level, MITM interception, i.e., looking at the empirical probability that a single connection can be intercepted by an organization or an attacker able to capture all the traffic going through an AS on a given direction (most of Internet communications being asymmetric). Such attacks can happen either by remote access to routing devices of an AS or even by Border Gateway Protocol (BGP) route hijacking. In our analysis, we focus on the case of MPTCP-capable source devices using two edge providers, analyzing measurement results on a geographical basis to identify which countries and regions MPTCP may grant higher confidentiality with respect to large-scale MITM threats.

An important assumption of our analysis is that the MPTCP scheduler behavior of endpoints or multipath converters can be tuned so that it does

not only look for throughput maximization, but also for path diversity exploitation for increased confidentiality, as investigated in [17]. Solutions offering programmability of the MPTCP scheduler are making surface, as notably [25, 16].

It is worth noting that, despite we refer to MPTCP as our reference multipath transport-layer protocol, our study can apply as well to other functionally equivalent protocols, such as for instance multipath QUIC (Quick User Datagram Protocol Internet Connections) [14].

The paper is organized as follows. Section 2 gives a background on MPTCP and related security concerns. In Section 3, we describe our measurement methodology. Section 4 presents the results, different application scopes of this work are discussed in Section 5, and in Section 6 we conclude the paper.

2. Background

In this section we provide the necessary background on the MultiPath TCP (MPTCP) protocol and on Internet-scale Man-In-The-Middle (MITM) attacks.

2.1. MultiPath TCP (MPTCP)

MPTCP extends TCP and allows fragmenting a data flow from a single connection into multiple paths (subflows TCP) [24, 46], as illustrated in Figure 1. At the application layer, a connection appears as a normal TCP connection. At the network layer, each subflow looks like a regular TCP flow whose segments carry in their header a new type of TCP option [24]. The protocol improves the performance offered by a single flow and makes the connection more reliable using concurrent and redundant paths.




Figure 1: Multipath TCP Connection: Overview

The initial TCP connection handshake carries an option, the MP_CAPABLE option, to enable MPTCP capability discovery and subflow creation. The handshake can carry additional information, such as a cryptographic key employed to authenticate the end-hosts and set up new subflows [24]. The establishment of additional subflow may employ also a token and random numbers (nonces), to prevent replay attacks on the authentication method. Further, an additional address identifier may be employed to identify the source IP address of a packet. Hence, even if the IP header has been changed by a middlebox (e.g. NATs, firewalls), end-hosts can identify an address without any doubt or ambiguity.

MPTCP can overcome some weaknesses inherent to TCP, achieving (i) a *greater throughput*, (ii) *higher reliability*, and (iii) *higher confidentiality*. Indeed, a multipath connection can improve the throughput aggregating bandwidth over different paths by concurrent data transmission across all available paths. Moreover, a multipath connection can quickly overcome one path failure by sending data to another available path, increasing the data delivery reliability [47]. Finally, fragmenting data flow across different subflows makes complete connection interception difficult because attackers would need to capture the transmitted content through all the subflows to build the content.

Therefore, MPTCP can provide a greater level of confidentiality than a regular TCP transmission if the subflows of a connection are routed along disjoint paths: the higher the level of disjointedness, the higher the confidentiality guarantee, and furthermore the higher the level of robustness against such attacks. The goal of this paper is to precisely quantify the level of robustness in use-cases where MPTCP is adopted not (only) to improve communication performance or reliability, but (also) to improve confidentiality. When addressing this aspect, router-level path disjointedness can be considered as being too weak in particular against AS-level traffic capturing and route hijacking. This is the reason why we focus instead on a larger scale of path disjointness, i.e., AS-level path disjointedness, which do make sense in practical scenarios as elaborated here after. Running an analysis on an even larger scale than AS-level scale (e.g., regional or country level) would likely be either infeasible or not sufficiently realistic.

2.2. Internet MITM Attacks

In Internet-scale communications, MITM attacks can happen when the attacker gains access to all the traffic transiting through an AS, or at least a

105 portion of it that is enough to reconstruct the transmitted data. In practice,
106 it can be possible by optical layer or BGP route hijacking MITM attacks.

107 At the optical layer, an attacker is able to split cables by using fiber op-
108 tical taps, as described in [58], with a low probability of being detected if
109 peculiar strategies are adopted as explained in [27, 52]. Moreover, one can
110 intercept the traffic by exploiting coupling and out-of-the-fiber light propa-
111 gation phenomena [57], despite the fact that this is particularly challenging
112 when performing wavelength-division-multiplexing.

113 At the BGP layer, MITM attacks exploit the natural way BGP works,
114 stealthily hijacking Internet routes to modify or capture the traffic before it
115 reaches the destination. BGP-based MITM attacks have been quite deeply
116 studied for about twenty years; in a recent survey [15] we have a detailed
117 description of such attacks, their effects as well as mitigation and defense
118 strategies.

119 This type of attack gained special attention in 2008, when a major provider
120 in central Asia hijacked Youtube traffic to apply local policies. In the same
121 year, a practical BGP MITM attack was demonstrated during the DefCon
122 hacking conference [3]: authors successfully intercepted traffic bound for the
123 conference network and redirected it to a system they controlled before rout-
124 ing it back to DefCon. A recent notable attack happened in 2014, attackers
125 injected BGP routes to redirect traffic from Bitcoin miner nodes to a com-
126 promised host [30]; it was estimated that at least \$83,000 worth of Bitcoins,
127 Dogecoins, HoboNickels, and Worldcoins were stolen over a period of four
128 months. More recently, in 2017 all traffic heading to Visa, MasterCard and
129 other service providers was hijacked for a short period of few minutes [54].
130 The cost of such BGP incidents could be even more than what have been re-
131 ported. Notable ones are documented in [29, 51]; often they are not reported
132 because they cannot be always detectable, they have limited scope, last for
133 a short time etc.

134 At the transport layer, the advent of MPTCP raised new security specifi-
135 cation questions and challenges [5, 6]. In [36], cryptography based solutions
136 are proposed against eavesdropping. The authors in [6] present an analysis of
137 residual threats in the MPTCP signaling and propose some fixes. Recently,
138 an extension of MPTCP to secure multipath communications was proposed
139 in [33], to offer authentication and encryption mechanisms not only to the
140 connection but also to single TCP options. This prevents different types
141 of MITM attacks where an attacker could force all the traffic to be sent
142 only over the path under his control by hijacking the traffic and erasing the




Figure 2: Adopted measurement methodology

143 MP_CAPABLE option.

144 In general, most of the works at the state of the art aim at either in-
 145 vestigating security threats for MPTCP or proposing solutions for them. It
 146 is worth mentioning the rising interest in using MPTCP to further enhance
 147 confidentiality when using Internet over-the-top Virtual Private Networks
 148 (VPN) services such as ToR and OnionCat [31]: MPTCP is used in the up-
 149 stream direction from the client to many gateways accessible via the VPN,
 150 on the way to the server, thus increasing the confidentiality level of the con-
 151 nection. Nevertheless, such practices can have a gain which can be hard
 152 to assess: how can you ensure the upstream source-destination traffic does
 153 follow disjoint paths, hence decreasing MITM efficiency, if not at the router-
 154 level, at the AS level? In this paper, for the first time at the state of the
 155 art to the best of our knowledge, and going beyond the preliminary study
 156 presented in [43], we attempt to provide a response to such questions.

157 3. Methodology

158 In this section, we first give a description on the datasets used for con-
 159 structing a representative AS-level graph of Internet, the basis for our anal-
 160 ysis. Then, we describe our approach for computing the number of valid
 161 vertex-disjoint paths between two arbitrary nodes over the constructed graph.
 162 Finally, we detail how we evaluate path diversity at different geographical
 163 scopes. The datasets we employed as well as our scripts are given in [41]
 164 for the sake of reproducibility. Figure 2 gives a schematic illustration of the
 165 different blocks of our measurement methodology.

166 It is worth noting that our methodologies imply that there is a way for a
 167 single MPTCP connection to have access to the network path diversity, by
 168 means of ad-hoc signaling or specific APIs. Solutions exist in this direction,
 169 as described in [35, 18].

170 3.1. Graph construction

171 We extract 2015 AS-level BGP-derived routing data from [4], and couple
172 it with the inter-AS relationship data [4] (i.e., indicating which AS is provider,
173 client or peer in an inter-AS link)². The result is a new dataset containing all
174 the AS links along with their frequency of occurrence and relationship type.

175 We choose this approach because, comparing with other resources [12] [32],
176 the topological data from [4] revealed to be more reliable and able to cap-
177 ture a broadened view of the Internet topology. Indeed, it integrates data
178 not only from Routeviews [50], but also from other resources such as RIPE
179 RIS [49]. It is worth noting that the alternative traceroute-based approach
180 employed in [12] has known issues [45] when converting router-level paths
181 into AS-level.

182 Employing measurements over a long period allows us to capture inter-
183 domain connection dynamics as well as inter-AS economic relationships. For
184 instance, in a one month period, only 85% of inter-AS links appear more
185 than 20 days, the remaining links with lower frequency of occurrence being
186 those used for backup operations or during BGP convergence periods. For
187 the sake of consistency, we removed these unstable links.

188 3.2. Path diversity computation

189 In order to have a measure of the path diversity, we need to enumerate all
190 the paths connecting two nodes over a graph that satisfy given routing prop-
191 erties. This problem is often referred to as policy-compliant path diversity
192 computation in the literature [23, 37]. The common approach [23] to this
193 problem is to convert the original graph into a type-of-relationship (ToR)
194 graph [21], i.e., a directed graph in which (i) the relationship between two
195 adjacent vertexes is expressed via the direction of the edge connecting them,
196 then (ii) maximizing the total number of vertex-disjoint paths between nodes
197 in this graph. However, the time-complexity experienced in such methods is
198 relatively high hence intractable for a graph as big as the AS graph.

199 In order to better scale, we introduce a novel path search algorithm lever-
200 aging the scale-free characteristics [2] of the input AS graph (i.e., a graph
201 with relatively few hubs capturing the majority of the paths) to optimize the

²The inter-AS relationship data from [4] is extracted monthly from the Cyclops database [20], which combines BGP data with Internet eXchange Point (IXP) data and adopts inference techniques proposed in [45].

202 execution time. In such a scale-free graph, the diameter (i.e., the length of
 203 the longest path among all the shortest paths) is not too high. Thus, the
 204 average path length (measured in number of AS hops) connecting any pair
 205 of nodes in the AS-level graph of Internet is around 5 as of today [38] (note
 206 that it is a bit lower with IPv6).

207 Searching for paths in a scale-free graph, i.e., a graph with a large minority
 208 of hub nodes connecting the rest of the nodes, is a problem of controllable
 209 complexity when adopting breadth or depth-first search algorithms with a
 210 limited depth; indeed, fixing a limited depth to a graph search, and that for a
 211 scale-free graph that has a limited diameter, strongly decreases the number
 212 of explored branches in the graph exploration³. From the constructed AS
 213 graph G , the breadth-first search algorithm we describe in Alg. 1 can be
 214 applied to discover all the policy-compliant paths between two nodes s and
 215 d , in a reasonable time.

216 Alg. 1 works as follows: (i) starting from the origin s , the algorithm
 217 explores every adjacent node n of s . (ii) A queue P is introduced to keep
 218 track of the explored paths; initially, it includes all the paths from s to n .
 219 (iii) Following these paths, the algorithm continues discovering the adjacent
 220 nodes to look for destination d . (iv) For a path p dequeued from P , the last
 221 node n is extracted, all of its neighbors are checked in sequence to determine
 222 the valid next hops towards d . (v) Once a neighbor is determined as valid,
 223 link to that neighbor will be added into the current path forming a new valid
 224 path toward destination. This new explored path is then enqueued into P for
 225 the next discovering phase. (vi) A node is considered as valid once the path
 226 through it does not violate the valley-free routing property [28]⁴; we express
 227 such policy-compliant path (i.e., a path that complies with the valley-free
 228 routing policy), using the following regular expression $c2p * p2p ? p2c *$ [37] in
 229 which $c2p$, $p2p$ and $p2c$ denote the relationship between interconnected nodes
 230 (where ? means that you can have one or none p2p link).

³breadth-first search explore first all the neighbors of a node, and then explore deeper in the graph; depth-first search, instead, explore first in depth starting from a given neighbor, and proceeds to the next neighbour only when the exploration in depth from the first one has terminated.

⁴A valley-free path is defined as a path that does not cross more than one peering agreements, which are agreements over which two ASes exchange only routes towards respective customers, which is justified by the fact that peering agreement are meant to be free-of-charge for both ASes

231 It is worth noting that, within G links are labeled according to their
 232 inferred relationship. For example, assuming that n_1, n_2, n_3 are neighbors of
 233 node s , in which s is customer (' c ') of n_1 , provider (' p ') of n_2 and peer with
 234 n_3 ; the links (s, n_1) , (s, n_2) , and (s, n_3) are labeled as ' $c2p$ ', ' $p2c$ ' and ' $p2p$ ',
 235 respectively. With these labels, the regular expression for policy-compliant
 236 path then could be leveraged to determine the validity of next hop toward
 237 the destination. For instance, taking the customer-type neighbors among
 238 the neighbors of s (i.e., n_2), and looking at their neighbors x in turn, those
 239 (n_2, x) links are not validated if they are either $c2p$ or $p2p$ because a customer
 240 is not expected to grant transit towards its other provider(s) to one among
 241 its providers, and a customer is not expected to give access to its peer(s) to
 242 its provider(s). By checking the labels of links along the explored path, the
 243 validity of next hops can be determined. Once a valid path is discovered,
 244 it is enqueued into P for the next discovering phase. The same exploration
 245 and validation processes are repeated for all the paths in P until reaching
 246 destination d or the path length goes over a given threshold τ .

247 The path validation executes at run-time to ensure that non-compliant
 248 paths are detected at the early stage, thus avoiding wasting time exploring
 249 invalid paths. By reducing the number of paths needed to be explored in
 250 the following phases, the search space is continuously optimized. Moreover,
 251 a proper choice of τ not only limits the time and space complexity, but can
 252 also avoid selecting long paths to be avoided in practice.

253 As a result of the path search algorithm, policy-compliant paths between
 254 two endpoints may share common nodes. To get the final set of vertex-
 255 disjoint paths, we run a simple off-line filtering linear algorithm to capture
 256 the shortest disjoint paths. Since the original list of valid paths turned out
 257 to be quite small most of the time and already sorted, the complexity of such
 258 a filtering operation is negligible.

259 3.3. Source-destination pairs

260 Within the constructed AS-level graph, multipath connections could be
 261 simulated by simply attaching end-hosts as virtual nodes into AS nodes of
 262 the original graph. For instance, a multi-homed device can be emulated by
 263 adding a new node, then linking it with at least two AS nodes. The connec-
 264 tion from that node to any other virtual nodes forms a multipath transport-
 265 layer communication. Our approach for emulating multipath communication
 266 can therefore be simply referred to as a process of source-destination pair se-
 267 lection. In the following, we define the target set of AS nodes which we

Algorithm 1: Path Search Algorithm

input : source s , destination d , graph g
output: ValidPathSet
 $VisitedNodes \leftarrow \emptyset$
 $queue.append([s])$
while $queue$ not empty **do**
 $path \leftarrow queue.pop()$
 $v \leftarrow path.LastNode()$
 if $v \notin VisitedNodes$ **then**
 for $n \in v.NeighborSet$ **do**
 if $n \notin VisitedNodes$ **and** $(label(v,n)='p2c'$ **or**
 $label(v,n)='p2p')$ **then**
 for $x \in n.NeighborSet$ **do**
 if $label(n,x)='c2p'$ **or** $label(n,x)='p2p'$ **then**
 $g.RemoveEdge(n,x)$
 end
 end
 end
 $NewPath \leftarrow list(path)$
 $NewPath.append(n)$
 if $n = d$ **then**
 $ValidPathSet.append(NewPath)$
 end
 if $length(NewPath) = \tau + 1$ **then** break
 $queue.append(NewPath)$
 end
 $VisitedNode.add(v)$
 end
end

268 consider for attaching the end hosts. A simulation process is then described
269 in details explaining which communication scenarios are covered in our study.

270 The current Internet ecosystem is composed of more than 70 thousand
271 ASes, out of which the large majority are stub ASes, i.e., ASes that are only
272 origin or destination ASes. About 13% are Tier-3 or small Tier-2 ASes, we
273 arbitrary define in this paper as those appearing at most in the third from last
274 position and at least penultimate position in BGP AS paths; we refer to such
275 ASes as ‘edge provider’ ASes, which can be considered as a representative set
276 of national Internet Service Provider (ISPs). Such ASes are often referred to
277 as ‘eyeball’ ASes. In this paper, an edge provider AS is not a stub AS, but is
278 rather expected to be a regional or national ISP, most of the time (rare are
279 the cases where an international/intercontinental ISP gives Internet access
280 to end-users).

281 Rather than taking into account all possible communications, we tar-
282 get the connections among hosts at the edges, performing connections us-
283 ing multiple sub-flows such as done with MPTCP. Considering connections
284 between hosts in different countries, we precisely address the MITM robust-
285 ness of Internet connections crossing multiple ASes. To precisely determine
286 which communications to cover in our study, we define a target set of source-
287 destination pairs that address, in a reasonable yet arbitrary way, the commu-
288 nications that may be more sensitive to communication privacy. Our choice
289 of source-destination pairs is as follows:

- 290 • the source is interconnected to two edge providers in a country.
- 291 • the destination is not multi-homed, i.e., it is reachable via a single ISP,
292 the one given by the best BGP path from each source edge provider,
293 and belongs to an AS at another country than the one of the source.

294 Figure 3 illustrates an example of how we simulate multipath communi-
295 cations accordingly the above policy. For each two arbitrary edge provider
296 ASes in a same country, one source is created (i.e., a dual-homed source).
297 For each edge provider in another country, one destination is paired with
298 the source. Such a pair dual-homed source - single-homed destination de-
299 fines the two endpoints of a multipath communication. Listing all pairs, i.e.,
300 combining a given source with every destination, all possible (international)
301 communications of a dual-homed host can be covered.

302 Besides reducing the number of pairs to a reasonable and treatable num-
303 ber (requiring about one week of computation), it is worth noting that, in




Figure 3: Representation of the source-destination pair selection process.

such a way, we consider communication in a single direction: from source to destination. That is, under such a path election strategy, we cover the case when a multi-homed device *uploads* to a single-homed server, as well as the case when a single-homed device *downloads* contents from multi-homed servers.

The scenarios that are not covered in our study include: (i) multi-homed devices *downloading* from single-homed server; (ii) single-homed devices *uploading* contents to multi-homed servers; (iii) a multi-homed device communicating with another multi-homed device. A dual analysis, quite expensive computationally, covering these additional cases may be performed as well in future works.

3.4. MiTM robustness metric aggregations

The ability to split traffic over different paths allows multipath protocol to realize MiTM attacks more difficult. Thus, the chance for an attacker to capture all the traffic sent by a source is reduced in proportion to the number of disjoint paths between the source and the destination. Path diversity is therefore a proper indicator to evaluate the MiTM robustness of a multipath communication.

Rather than considering the robustness against MiTM attacks of every connection individually, we are more interested in evaluating such a robustness at the end-host level, thus measuring the degree of robustness offered by a multipath-capable source device to secure its data sending over the Internet.

3.4.1. Source-specific MiTM robustness metric

With regard to the aforementioned approach for source-destination pair selection, we define the *source-specific MiTM robustness metric* as the aver-

age number of disjoint paths over all the destination edge providers that are in a different country than the source. Such a metric can be considered as a level of unlikelihood that a MiTM attack takes place for that source configuration; the higher the value of the robustness metric, the more difficult it is for an attacker to capture traffic from that source.

3.4.2. Source country-specific MiTM robustness metric

Aggregating results from all the sources within a given country we can obtain a *source country-specific MiTM robustness metric*. Such a definition allows us to characterize the robustness level offered by different source countries to multipath communications.

3.4.3. Country-level source-destination MiTM robustness metric

As another way to aggregate the MiTM robustness metric computation, we also study a country-level source-destination based aggregation, i.e., leading to a robustness metric for a pair of source and destination countries. Given a source (a pair of edge providers in a country) and a destination country, its MiTM robustness metric is defined as the average number of disjoint paths from the source over all edge providers belonging to the destination country.

3.4.4. Country-pair MiTM robustness metric

By grouping together the results from all the sources within a source country, we can define the *country-pair MiTM robustness metric* for the corresponding pair of countries.

3.4.5. Metric computation

Let us more precisely characterize the aforementioned source-destination pair selection process with respect to the two MiTM robustness metric aggregations we study in the following, i.e., the source country-level one and the country-pair one. We segment the set of edge providers, E , in country-specific subsets, E_c , where c denotes a country in the set of countries C , i.e., $E = \bigcup_{c \in C} E_c$. We employ the AS-to-country mapping given by the CIDR Report [7]. Let us indicate with $E_{\tilde{c}}$ the restriction to a specific country $\tilde{c} \in C$. Overall, for a given country \tilde{c} , the number of source-destination pairs is therefore equal to the number of pairs of edge providers for the given

country multiplied by the number of edge providers of other countries, i.e.

$$\frac{|E_{\tilde{c}}| \times (|E_{\tilde{c}}| - 1)}{2} \times \sum_{c \neq \tilde{c}} |E_c| \quad (1)$$

For a given source and destination countries, s and d respectively, the number of source-destination pairs connecting them is equal to:

$$\frac{|E_s| \times (|E_s| - 1)}{2} \times |E_d| \quad (2)$$

Doing so, we target a lower bound, pessimistic analysis, since we only take into consideration international communications and we suppose the destination is not multi-homed. The filter we set on the destination enumeration allows us to target communications that may need a higher level of confidentiality due to their international connotation. Moreover, in this way we also avoid a huge bias potentially due to the fact that a large majority of the AS paths available at the national level are not visible in backbone BGP routing tables such as the Routeviews ones (typically because of Internet exchange points, as recently shown in [1]). We believe having a lower bound stand is more appropriate than an upper bound one, while allowing us to scientifically qualify the value of the relative trends.

4. Results

We report the results obtained for a set of 147 countries, i.e., those countries from the United Nations statistics [56] that appear to have at least two distinct edge providers officially based in the country; this automatically excludes Greenland territories, very small city-state countries, many African countries and Indonesia. The geographical coverage is given in Figure 5. In the following sections, we present the statistics for two different MiTM robustness metric aggregations, the country source specific one and the country pair one.

4.1. Source country aggregation

Let us recall the measurement approach for source country-specific MiTM robustness analysis:

- For each country, we generate all possible dual-homed sources, i.e., all possible pairs of edge providers.

- 390 • For each such source configuration, we compute the number of disjoint
391 paths to each destination. For each edge provider that is a different
392 country than the source country, one destination is generated.
- 393 • For a given source, we compute its corresponding robustness metric by
394 taking the average of the number of disjoint paths over all the destina-
395 tions.
- 396 • For each country, a series of MITM robustness metrics is hence gener-
397 ated, one for each source.

398 We characterize the resulting series using boxplot distributions (using a
399 0.1% outliers threshold). We overlay over the boxplots the average of the
400 corresponding series with a red square, order them with increasing averages⁵
401 from left to right. We report the results in Figure 4, and with a geographical
402 view in Figure 5. We express three different viewpoints:

- 403 • *device view* (Figure 4a): the MITM robustness is computed with the
404 source node integrated in the AS graph as an ‘artificial’ node, i.e., the
405 path search algorithm finds the number of AS-disjoint paths from this
406 source node toward the destination. It provides therefore a device view;
407 obviously, in this view the upper bound of the robustness is 2, i.e., the
408 number of edge providers used by the source.
- 409 • *edge provider view* (Figure 4b): the MITM robustness is computed
410 counting the number of disjoint paths from the first and the second
411 edge provider, then decreased by those paths that share an AS hop.
412 Taking into account such a view, we assume that additional AS paths
413 can be made available to MPTCP subflows acting at the edge providers
414 level, e.g., by forms of flow path steering and load-balancing.
- 415 • *differential view* (Figure 4c): the differential robustness results, i.e.,
416 the edge provider view robustness minus the device view robustness,
417 computed for each source configuration individually. This view more
418 precisely quantifies the gain achievable for MPTCP communications
419 when inter-AS load-balancing is enabled at the edge providers.

⁵Average values do include outliers.




Figure 4: MITM robustness distribution for 147 countries.

420 The above viewpoints also reflect different levels of trust on the providers.
 421 That is, while the edge provider view assumes MITM attacks do not happen
 422 at the source and destination edge providers (i.e., there is a high level of trust
 423 on those providers), the device view assumes that attacks can happen at the
 424 source edge providers, hence revealing a low level of trust in source direct
 425 providers.




Figure 5: Countries covered with corresponding MiTM robustness distribution

426 As a general assessment, Figure 4 shows a distribution to be interpreted.
 427 For example, one could consider 1.5 as the rough threshold above which the
 428 likelihood of MiTM is to be considered low, and conversely high if lower
 429 than 1.5. Only about 5% of the countries show good chances of being ro-
 430 bust against MITM from a device viewpoint, while looking at the maximum
 431 instead of the average and median values one could speculate that careful
 432 choice of the edge providers could make the MiTM likelihood low for a ma-
 433 jority of the countries. From an edge provider viewpoint, this ratio grows to
 434 roughly 60%, and higher than 90% looking at the maximum, that is if the
 435 edge provider choice can be influenced by confidentiality concerns.

436 Moreover, the average number of paths connecting a dual-homed node to
 437 international destinations has a significant variance depending on the origin

country. The average robustness ranges from 1 (and less) to 1.6 from a device viewpoint, and from 1 (and less) to 2.5 from an edge provider viewpoint. It is worth noting that the reason why some minimum, and even average values, are below 1, is the partial view over the Internet topology and the incompleteness of inter-AS relationship inference; in fact, these factors make some destinations unreachable (counted as 0 path), but we left the 0 values in the series to also give an index of the level of topology incompleteness for different countries. In any case, the boxplot median is a metric robust against such outliers to look at.

In addition, observing the distributions in Figure 4, we can also remark that:

- Within a country, a high inter-quantile range indicates that the path diversity strongly depends on how the two upstream edge providers are selected for the source.
- The gap between the min and max robustness is another interesting fitness metric to observe. Some countries maintain a small gap (below 1) while others have a very big gap (up to 2). In other words, the deployment of multipath transport-layer communications for securing international communications in some countries can statistically yield a much better result than in other countries, where this gap is smaller. Particularly interesting is the case of Angola (AO), Venezuela (VE) and Namibia (NA), with small robustness gaps, which may be correlated to the presence of inter-continental cables landing in or close to the country [11].
- The median is mostly higher than the average in the device view, and lower than the average in the edge provider view. This is essentially due to outliers, counted in the average and not in the median.
- From the edge provider viewpoint, the maximum value is higher than 2 in the most of the countries, suggesting that with a proper choice of trusted source providers, one can adopt multipath communications to statistically expect high confidentiality for its communications. Particularly alerting are the cases of Uzbekistan (UZ), Nepal (NP) and Lebanon (LB), with quite low maximum values.
- From the device viewpoint, in most of the cases the maximum robustness is not higher than 1.6, both averages and medians are quite far

473 from the desirable target of 2. Hence, without the support of inter-AS
474 load-balancing at source providers, path diversity from a dual-homed
475 node is reduced significantly, indicating a non negligible probability of
476 paths joining on the way to the destination.

- 477 • Considering the differential robustness, we can remark that among the
478 countries that have the lowest device view MITM robustness, those that
479 could most benefit from inter-AS load-balancing practices are Mongolia
480 (MN), Pakistan (PK) and Korea (KR). However, the majority of those
481 countries with low robustness do not improve much the situation going
482 from the device view to the edge provider view.

483 Looking at macro geographical regions, many European countries seem to
484 grant better security than countries in other regions. In order to look at con-
485 tinental characteristics, the plots in Figure 6 show the boxplot results (with
486 1% outliers) aggregated on a macro-region basis (a and c, sub-continental
487 level) and on a relative position basis (b and d, in terms of seacoast and
488 inland borders). We can remark that:

- 489 • Western Europe appears to be the best off, followed by Northern Eu-
490 rope and Northern America. In almost 50% of Western Europe coun-
491 tries there can be 2 disjoint paths from the source edge providers to
492 Internet destinations.
- 493 • Central Asia shows the worst robustness, followed by Australia and
494 New Zealand; the reasons are likely network centralization practices
495 and geographical isolation. It is interesting to notice the relevant gap
496 between Central and South-Eastern/Western Asia.
- 497 • within Europe, Western countries do offer a better diversity over North-
498 ern countries, and especially over Eastern and Southern countries. with
499 a small range of variation and a high median value show the best result.
- 500 • A high variance is recorded at Southern Asia, Northern Europe and
501 Sub-Saharan Africa, which indicates high differences among the coun-
502 tries within these areas.
- 503 • We could not find a strong correlation between the relative continental
504 position, and the robustness metric, yet a positive correlation exists,
505 with countries at the boundaries of oceans, with inter-continental cable

506 landing and that are sea-oriented (most of the border on the coast)
 507 that offer higher robustness than fully internal and continental-oriented
 508 ones.

509 4.2. Source-destination country pair aggregation


510 As we may notice, the MiTM robustness level of a multipath communi-
 511 cation could be affected not only by the country where the communication
 512 starts but also by the choice of upstream providers at that country. Besides
 513 that, within a source country, the robustness level for different destination
 514 countries can significantly vary. To evaluate this latter aspect further, we
 515 perform a source-destination country pair aggregation.

516 Over the set of 147 countries, we evaluate the robustness metric for 1547
 517 directional country-to-country communication pairs in which the MiTM ro-
 518 bustness metric for one pair is computed as follows:


- 519 • For a given source country, we generate all possible dual-homed sources,
 520 i.e., all possible pairs of edge providers.
- 521 • For each such source configuration, we compute the number of disjoint
 522 paths to each edge providers located in the destination country.
- 523 • For a given source, we take the average of the number of disjoint paths
 524 over all the destinations to get its source-destination based MiTM ro-
 525 bustness metric.
- 526 • For a given source-destination country pair, a series of MiTM robust-
 527 ness metrics, one for each source, is therefore created.

528 In Figure 7, we report the CDF of the average MiTM robustness, for all
 529 the 1547 pairs. The high range of variation (between 0.4 and 6) shows us
 530 the big robustness gap between pairs. Only 20% of the country pairs show
 531 an average of two or higher. For the remaining pairs, approximately 73%
 532 of them have the average range from 1 to 2. The remaining 7% are country
 533 pairs with very low robustness, below one; besides the specific context related
 534 to a country pair, a factor behind such bad performance can be the already
 535 discussed topology view incompleteness.


536 To better understand the impact caused by different destinations, we
 537 further characterize the top 147 and bottom 147 pair in the CDF distribution,
 538 i.e., roughly the top 10% and the bottom 10% cases. The results are presented




(a) device view: macro-regions grouping



(b) position grouping



(c) edge provider view: macro-regions grouping



(d) position grouping

Figure 6: MITM robustness metric with continental subregion grouping.




Figure 7: CDF of average MiTM robustness for 1547 pairs of source-destination country

539 in Figure 8, where the country pairs in each group are ordered from left to
 540 right with an increasing average (the average do include the outliers). We
 541 report the MiTM robustness distribution of each pair using the boxplot (with
 542 0.1% outliers) overlaid with a red square representing the average.

543 Figure 8a reports the MiTM robustness metric distribution for the top 147
 544 country pairs. The high inter-quartile range (IQR) with a pair highlights the
 545 strong impact caused by edge providers choice at the source to the robustness
 546 metric. Besides that, there are also some source countries, such as Morocco
 547 (MA), Madagascar (MG), Gibraltar (GI), Guam (GU), Jersey (JE), Namibia
 548 (NA), Liechtenstein (LI) and Belize (BZ), that suffer from the presence of
 549 only one edge provider pair; these countries result in pairs with a collapsed
 550 robustness point in the box. In addition, within these top 147 pairs, there
 551 are some destinations, like Namibia (NA), Guam (GU) and Belize (BZ),
 552 that appear to show high sensibility to the destination choice on the MiTM
 553 robustness.

554 In Figure 8b, we report the results for the bottom 147 country pairs. The
 555 majority of them have Montenegro (ME) as the destination. The second
 556 destination is Republic of Congo (CG). That highlights again the impact of
 557 destination choice on the MiTM robustness level. Unlike the top 10% case,
 558 we see a small inter quartile range (IQR) for most of the pairs, showing that
 559 even a careful choice on the edge providers at the source country cannot
 560 improve much the level of robustness for such connections. In other words,
 561 regardless of the origin country as well as the choice of source edge providers,
 562 the possibility of employing MPTCP to secure the communications destined




Figure 8: MITM robustness distribution for the top and bottom 147 pairs of country (with respect to their average MITM robustness)

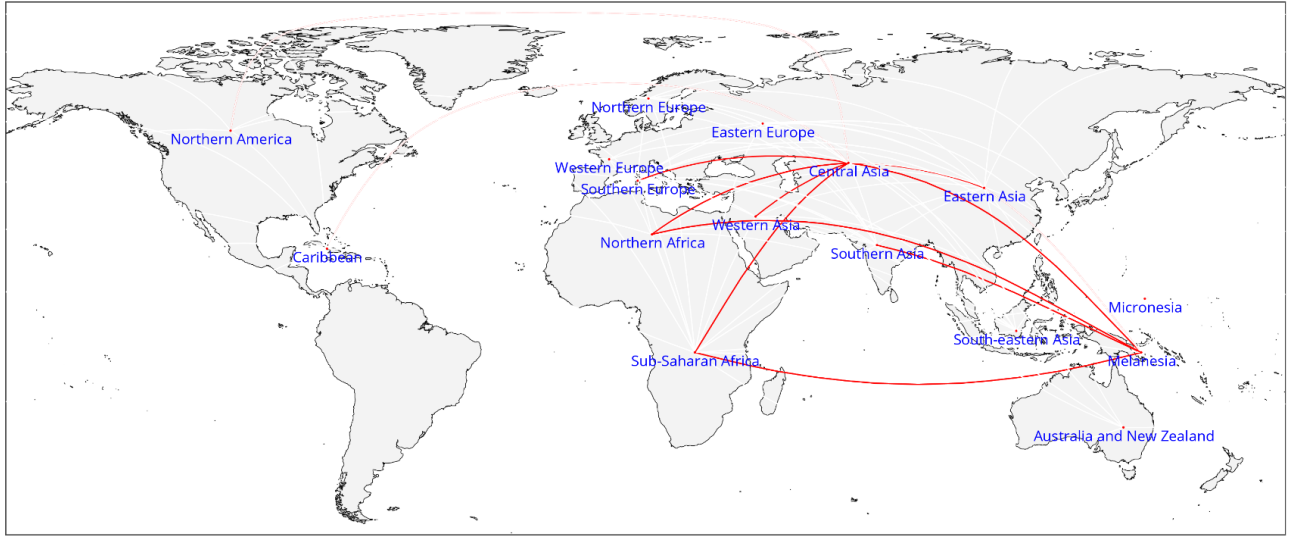
563 to, e.g., Montenegro and Republic of Congo is extremely low.

564 Considering 1 and 2 as the thresholds for very low (zero) and high (suffi-
565 cient) robustness, respectively, a source-destination pair can be classified as:
566 (1) highly robust against MiTM if it has the average robustness level of at
567 least 2, and (2) weak against the MiTM once maintaining the average of 1
568 or lower.

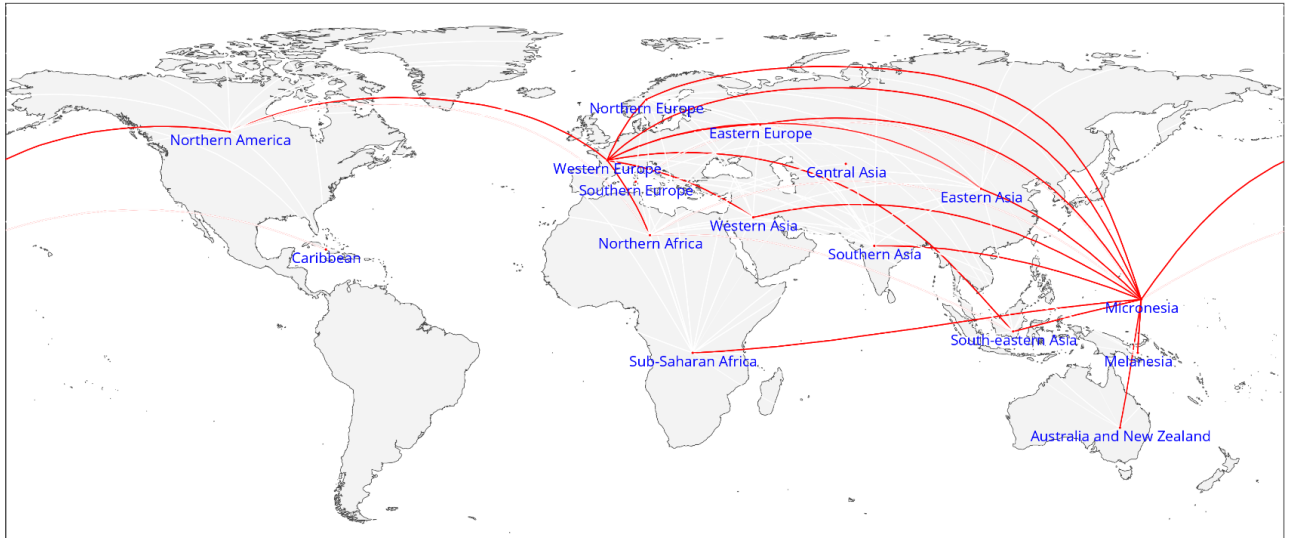
569 We visualize the country-to-country communications in these two classes
570 by mapping them into a geographical map in Figure 9. To avoid too many
571 lines, we first group countries with respect to their subregion, then converting
572 these country-to-country connections into the corresponding subregion-to-
573 subregion connections. Finally, the subregional connections are expressed
574 using lines with different opacity reflecting the portion of country-to-country
575 communications between subregions having the MiTM robustness level less
576 than or equal to 1 as in Figure 9a, and equal to or higher than 2 as in Figure
577 9b.

578 In Figure 9a, we only show the connections between subregions when
579 there are more than 30% of the country-to-country communications with a
580 robustness metric of at most one. For subregion pairs with less than 30% of
581 their country-to-country communications having a robustness metric lower
582 than one, the connection lines are hidden. In other words, the lines point out
583 the subregions where the deployment of MPTCP cannot offer any protection
584 against large-scale MiTM attacks. As presented in the map, the area of Cen-
585 tral Asia and Melanesia are the two subregions having the worst performance,
586 most of their MPTCP communications with other subregions are classified
587 as zero-robust. Thus, most of the subregions could not be benefit from the
588 deployment of MPTCP to secure their communications with Central Asia.

589 In the sub-regional view of the high robustness group presented in Figure
590 9b, we show the connection lines between sub-regions with more than 50%
591 of the country-to-country communications having robustness level of 2 or
592 higher. In such a view, Micronesia and then Western Europe are the two ar-
593 eas that outperform the others in term of MiTM robustness. As depicted in
594 the plot, except for a few low connected regions, like Central Asia, Caribbean
595 and Northern Africa, etc., most of the multipath communications from and
596 to Micronesia can profit from a high level of robustness. It is worth noting
597 that in the region of Micronesia, Guam is the only country covered by our
598 study. The high robustness result captured for communications from and to
599 this region is therefore directly related to the highly connected network in-
600 frastructure of Guam being a crucial node in the Internet cable network [53].



(a) regions with more than 30% of country-to-country communications having at most one path



(b) regions with more than 50% of country-to-country communications having at least two paths

Figure 9: Regional view of the source-destination based MiTM robustness

601 5. Practical aspects

602 We focused our study on MPTCP-based communications. More precisely,
603 it covers the following cases:

- 604 • *MPTCP capable endpoints*: both source and destination, client and
605 server (or vice versa), are MPTCP capable, and the MPTCP commu-
606 nication is not filtered by middle-boxes. As argued in Section 3.3, the
607 multi-homed endpoint can be either the server or the client.
- 608 • *MPTCP proxied endpoints*: at least one endpoint is not MPTCP ca-
609 pable, but the TCP communications are handled by MPTCP proxies,
610 converting TCP packets into MPTCP packets and vice versa, as ex-
611 plained in [8, 10], possibly routed via Internet disjoint paths as pro-
612 posed in [19, 9]. The multipath conversion proxies can sit at endpoint
613 premises (customer premises equipment for the client, hypervisor or
614 middle-box at the server) or at the edge provider level borders.

615 Besides MPTCP-based communications, other protocols offering Internet-
616 scale multipath, connection flow-level load-balancing could also be covered
617 by our study. The following protocols are either not deployed, or they have
618 only undergone a limited deployment at the Internet scale so far; they are:

- 619 • *SCTP*: the Stream-Control-Protocol (SCTP) [55] is another multipath
620 transport protocol absolving the same function as MPTCP, but less
621 deployed than MPTCP due to the limited retrocompatibility.
- 622 • *LISP*: the Locator/Identifier Separation Protocol (LISP) [26] is able to
623 perform inter-AS inbound load-balancing by means of encapsulation,
624 routing locator mapping, and appropriate traffic engineering (TE) pol-
625 icy configuration. LISP primary scope is the edge provider one, hence
626 results with the edge provider view are readily applicable. Further-
627 more, deployment of LISP as an intra-AS TE tool can also allow us
628 to perform inter-AS multipath on the outbound direction as proposed
629 in [42].
- 630 • *MultiPath BGP*: in BGP, the routing decision process only allows us to
631 take one route per network prefix. The selected path can be inefficient
632 in terms of global routing. Recently, forms of *Multipath BGP* were dis-
633 cussed in standardization fora, but finally not standardized; however,

634 some recommendations have been published [39], and implemented by
635 some vendors (see, e.g., [34] and [13]). Such multipath mode can be
636 adopted at the edge provider scope to enable load-balancing at the
637 egress direction. Despite the study [22] on core routing tables reports
638 that in 2010 multipath BGP was practically not used, speculations
639 report that it is used by major cloud providers.

640 The above protocols are a selection of those protocol communication con-
641 texts where load-balancing can affect the AS-path selection. There are also
642 other load-balancing protocols which can potentially influence the egress AS
643 selection as well, as for instance in data-center environments. In the case
644 of MPTCP communications, these protocols, operated at the edge provider
645 view, are able to perform inter-AS load-balancing in such a way that the
646 path diversity exposed in our edge provider view can be made available to
647 MPTCP devices, hence giving them the full potential of MPTCP in terms
648 of communication confidentiality and robustness against MITM attacks.

649 Finally, additional multipath transport-layer protocols are making sur-
650 face, as for example the already mentioned multipath extension to the QUIC
651 protocol [14], nicknamed MPQUIC. As MPTCP that authenticates its op-
652 tions to avoid interference from one path to the others as already discussed,
653 MPQUIC also has a similar protection by default, because every control in-
654 formation in QUIC is authenticated.

655 6. Conclusion

656 We explored in this paper how Internet path diversity could be exploited
657 by means of multipath transport-layer protocols such as MPTCP, when look-
658 ing at increased security against man-in-the-middle attacks. We focused on
659 such attacks acting at the autonomous system level, and at the robustness
660 of multipath communications in what appear as a reasonable configuration
661 where at least one endpoint is multi-homed with two edge providers.

662 We reported extensive, specific and aggregated results for most of the
663 world countries and regions, looking at macro trends that could inspire fur-
664 ther research in the area. Results show that, statistically speaking, a mul-
665 tipath protocols such as MPTCP does not help in guaranteeing robustness
666 against MiTM attacks hence high confidentiality, unless (i) the choice of
667 the edge provider is carefully taken, or (ii) one can rely on inter-AS load-
668 balancing features offered implicitly or explicitly by edge providers. Some

continental regions are strongly more robust than others, and there seems to be a positive correlation with inter-continental cable landing proximity. Moreover, the results show that there are countries surprisingly less well connected than one could think of, such as Northern America countries, and countries that are more obviously less robust against such attacks due to network centralization practices.

It is worth mentioning that the methodology we propose to measure MiTM robustness could be instrumental also for other types of analysis. For instance, having a high MiTM robustness may also represent an increased sensibility to distributed denial of service attacks (DDOS), as the set of possible sources not sharing a network bottleneck can be expected to increase with the AS-level path disjointness. This aspect may be object of further work.

Acknowledgement

We acknowledge the support and relevant effort of Ho Dac Duy Nguyen in the preparation of the data processing and methodology. This work was partially funded by the French Investissement d’Avenir FED4PMR project. We thank Xenofontas Dimitropoulos for his valuable feedback.

References

- [1] Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., Willinger, W., 2012. Anatomy of a large european IXP. In: ACM SIGCOMM 2012 Conference, SIGCOMM ’12, Helsinki, Finland - August 13 - 17, 2012. pp. 163–174.
URL <http://doi.acm.org/10.1145/2342356.2342393>
- [2] Albert, R., Barabási, A., 2001. Statistical mechanics of complex networks. CoRR cond-mat/0106096.
URL <http://arxiv.org/abs/cond-mat/0106096>
- [3] A.Pilosov, T.Kapela, 2017. Stealing the internet - an internet-scale man in the middle attack.
URL <https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf>
- [4] AS-level Topology Archive, 2015. As-level topology archive.
URL <http://irl.cs.ucla.edu/topology>

- 702 [5] Bagnulo, M., March 2011. Threat analysis for tcp extensions for multi-
703 path operation with multiple addresses. RFC 6181, RFC Editor.
704 URL <https://tools.ietf.org/rfc/rfc6181.txt>
- 705 [6] Bagnulo, M., Paasch, C., Gont, F., Bonaventure, O., Raiciu, C., July
706 2015. Analysis of residual threats and possible fixes for multipath tcp
707 (mptcp). RFC 7430, RFC Editor.
708 URL <https://tools.ietf.org/rfc/rfc7430.txt>
- 709 [7] Bates, T., Smith, P., Huston, G., 2017. Cidr report.
710 URL <http://www.cidr-report.org>
- 711 [8] Behaghel, D., Secci, S., Vinapamula, S., Seo, S., Cloetens, W., Meyer,
712 U., Contreras, L., Peirens, B., March 2017. Extensions for network-
713 assisted mptcp deployment models. Internet-draft.
714 URL [https://tools.ietf.org/id/
715 draft-boucadair-mptcp-plain-mode-10.txt](https://tools.ietf.org/id/draft-boucadair-mptcp-plain-mode-10.txt)
- 716 [9] Benchaïb, Y., Secci, S., Phung, C., 2015. Transparent cloud access per-
717 formance augmentation via an MPTCP-LISP connection proxy. In: Pro-
718 ceedings of the Eleventh ACM/IEEE Symposium on Architectures for
719 networking and communications systems, ANCS 2015, Oakland, CA,
720 USA, May 7-8, 2015. pp. 201–202.
721 URL <https://doi.org/10.1109/ANCS.2015.7110140>
- 722 [10] Boucadair, M., Jacquenet, C., Bonaventure, O., Henderickx, W., Skog,
723 R., December 2016. Network-assisted mptcp: Use cases, deployment
724 scenarios and operational considerations. Internet-draft.
725 URL [https://tools.ietf.org/id/
726 draft-nam-mptcp-deployment-considerations-01.txt](https://tools.ietf.org/id/draft-nam-mptcp-deployment-considerations-01.txt)
- 727 [11] Cable Map, 2017. Greg’s cable map.
728 URL <http://cablemap.info>
- 729 [12] CAIDA, 2017. Archipelago (ark) measurement infrastructure.
730 URL <http://www.caida.org/projects/ark>
- 731 [13] Cisco, Sept 2016. Bgp best path selection algorithm. Cisco guide, Trou-
732 bleshooting tech notes, Document ID:13753.
733 URL [https://www.cisco.com/c/en/us/support/docs/ip/
734 border-gateway-protocol-bgp/13753-25.html](https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html)

- 735 [14] Coninck, Q. D., Bonaventure, O., 2017. Multipath QUIC: design and
736 evaluation. In: Proceedings of the 13th International Conference on
737 emerging Networking EXperiments and Technologies, CoNEXT 2017,
738 Incheon, Republic of Korea, December 12 - 15, 2017. pp. 160–166.
739 URL <http://doi.acm.org/10.1145/3143361.3143370>
- 740 [15] Conti, M., Dragoni, N., Lesyk, V., 2016. A survey of man in the mid-
741 dle attacks. IEEE Communications Surveys and Tutorials 18 (3), 2027–
742 2051.
743 URL <https://doi.org/10.1109/COMST.2016.2548426>
- 744 [16] Corbillon, X., Aparicio-Pardo, R., Kuhn, N., Texier, G., Simon, G.,
745 2016. Cross-layer scheduler for video streaming over MPTCP. In: Pro-
746 ceedings of the 7th International Conference on Multimedia Systems,
747 MMSys 2016, Klagenfurt, Austria, May 10-13, 2016. pp. 7:1–7:12.
748 URL <http://doi.acm.org/10.1145/2910017.2910594>
- 749 [17] Coudron, M., Nguyen, H. D. D., Secci, S., 2016. Enhancing buffer di-
750 mensioning for multipath TCP. In: 7th International Conference on
751 the Network of the Future, NOF 2016, Búzios, Brazil, November 16-18,
752 2016. pp. 1–7.
753 URL <https://doi.org/10.1109/NOF.2016.7810142>
- 754 [18] Coudron, M., Secci, S., Maier, G., Pujolle, G., Pattavina, A., Nov 2013.
755 Boosting cloud communications through a crosslayer multipath protocol
756 architecture. In: 2013 IEEE SDN for Future Networks and Services
757 (SDN4FNS). pp. 1–8.
- 758 [19] Coudron, M., Secci, S., Pujolle, G., Raad, P., Gallard, P., 2013. Cross-
759 layer cooperation to boost multipath TCP performance in cloud net-
760 works. In: IEEE 2nd International Conference on Cloud Networking,
761 CloudNet 2013, San Francisco, CA, USA, November 11-13, 2013. pp.
762 58–66.
763 URL <https://doi.org/10.1109/CloudNet.2013.6710558>
- 764 [20] Cyclops, 2017. Cyclops.
765 URL <https://cyclops.cs.ucla.edu>
- 766 [21] Di Battista, G., Erlebach, T., Hall, A., Patrignani, M., Pizzonia, M.,
767 Schank, T., 2007. Computing the types of the relationships between

- 768 autonomous systems. *IEEE/ACM Trans. Netw.* 15 (2), 267–280.
 769 URL <http://doi.acm.org/10.1145/1279660.1279662>
- 770 [22] Elena, E., Rougier, J., Secci, S., 2010. Characterisation of as-level path
 771 deviations and multipath in internet routing. In: *Next Generation In-*
 772 *ternet (NGI)*, 2010 6th EURO-NF Conference on, Paris, France, June
 773 2-4, 2010. pp. 1–7.
 774 URL <https://doi.org/10.1109/NGI.2010.5534468>
- 775 [23] Erlebach, T., Moonen, L. S., Spieksma, F. C. R., Vukadinovic, D., 2009.
 776 Connectivity measures for internet topologies on the level of autonomous
 777 systems. *Operations Research* 57 (4), 1006–1025.
 778 URL <https://doi.org/10.1287/opre.1080.0677>
- 779 [24] Ford, A., Raiciu, C., Handley, M., Bonaventure, O., Jan 2013. Tcp
 780 extensions for multipath operation with multiple addresses. RFC 6824,
 781 RFC Editor.
 782 URL <https://tools.ietf.org/rfc/rfc6824.txt>
- 783 [25] Frömmgen, A., Rizk, A., Erbschäuer, T., Weller, M., Koldehofe, B.,
 784 Buchmann, A. J., Steinmetz, R., 2017. A programming model for
 785 application-defined multipath TCP scheduling. In: *Proceedings of the*
 786 *18th ACM/IFIP/USENIX Middleware Conference*, Las Vegas, NV,
 787 USA, December 11 - 15, 2017. pp. 134–146.
 788 URL <http://doi.acm.org/10.1145/3135974.3135979>
- 789 [26] Fuller, V., Farinacci, D., Jan 2013. Locator/id separation protocol (lisp)
 790 map-server interface. RFC 6833, RFC Editor.
 791 URL <https://tools.ietf.org/rfc/rfc6833.txt>
- 792 [27] Furdek, M., Skorin-Kapov, N., 2012. Physical-layer attacks
 793 in transparent optical networks, optical communications sys-
 794 tems. [Http://www.intechopen.com/books/optical-communications-](http://www.intechopen.com/books/optical-communications-systems/physical-layer-attacks-in-transparent-optical-networks)
 795 [systems/physical-layer-attacks-in-transparent-optical-networks.](http://www.intechopen.com/books/optical-communications-systems/physical-layer-attacks-in-transparent-optical-networks)
- 796 [28] Gao, L., 2001. On inferring autonomous system relationships in the
 797 internet. *IEEE/ACM Trans. Netw.* 9 (6), 733–745.
 798 URL <https://doi.org/10.1109/90.974527>

- 799 [29] Goldberg, S., 2014. Why is it taking so long to secure internet routing?
800 Commun. ACM 57 (10), 56–63.
801 URL <http://doi.acm.org/10.1145/2659899>
- 802 [30] Greenberg, A., 2014. Hacker redirects traffic from 19 internet providers
803 to steal bitcoins.
804 URL <https://www.wired.com/2014/08/isp-bitcoin-theft>
- 805 [31] Hackdopi, 2017. Hackdopi.
806 URL <http://hackdopi.wikidot.com>
- 807 [32] IRR, 2017. Internet routing registries.
808 URL <http://www.irr.net>
- 809 [33] Jadin, M., Tihon, G., Pereira, O., Bonaventure, O., 2017. Securing mul-
810 tipath TCP: design & implementation. In: 2017 IEEE Conference on
811 Computer Communications, INFOCOM 2017, Atlanta, GA, USA, May
812 1-4, 2017. pp. 1–9.
813 URL <https://doi.org/10.1109/INFOCOM.2017.8057011>
- 814 [34] Juniper, 2019. Configuring bgp to select multiple bgp paths. JUNOS 8.2
815 Routing Protocols Configuration Guide in Chapter 33 BGP Configura-
816 tion Guidelines pp:571-572.
817 URL [https://www.net.t-labs.tu-berlin.de/teaching/ss08/RL_](https://www.net.t-labs.tu-berlin.de/teaching/ss08/RL_labcourse/docs/04-juniper-bgp.pdf)
818 [labcourse/docs/04-juniper-bgp.pdf](https://www.net.t-labs.tu-berlin.de/teaching/ss08/RL_labcourse/docs/04-juniper-bgp.pdf)
- 819 [35] Kabbani, A., Vamanan, B., Hasan, J., Duchene, F., 2014. Flowben-
820 der: Flow-level adaptive routing for improved latency and throughput
821 in datacenter networks. In: Proceedings of the 10th ACM International
822 on Conference on emerging Networking Experiments and Technologies
823 (CoNext). ACM.
- 824 [36] Kim, D.-Y., Choi, H.-K., 2016. Efficient design for secure multipath tcp
825 against eavesdropper in initial handshake. In: Proceedings of the 2016
826 International Conference on Information and Communication Technol-
827 ogy Convergence (ICTC 2016).
828 URL <https://doi.org/10.1109/ICTC.2016.7763559>
- 829 [37] Klöti, R., Kotronis, V., Ager, B., Dimitropoulos, X. A., 2015. Policy-
830 compliant path diversity and bisection bandwidth. In: 2015 IEEE

831 Conference on Computer Communications, INFOCOM 2015, Kowloon,
832 Hong Kong, April 26 - May 1, 2015. pp. 675–683.
833 URL <https://doi.org/10.1109/INFOCOM.2015.7218436>

834 [38] Kühne, M., 2012. Update on as path lengths over time.
835 URL [https://labs.ripe.net/Members/mirjam/](https://labs.ripe.net/Members/mirjam/update-on-as-path-lengths-over-time)
836 [update-on-as-path-lengths-over-time](https://labs.ripe.net/Members/mirjam/update-on-as-path-lengths-over-time)

837 [39] Lange, A., March 2012. Issues in revising bgp-4. Internet-draft.
838 URL [https://tools.ietf.org/id/](https://tools.ietf.org/id/draft-ietf-idr-bgp-issues-06.txt)
839 [draft-ietf-idr-bgp-issues-06.txt](https://tools.ietf.org/id/draft-ietf-idr-bgp-issues-06.txt)

840 [40] Li, M., Lukyanenko, A., Ou, Z., Ylä-Jääski, A., Tarkoma, S., Coudron,
841 M., Secci, S., Fourthquarter 2016. Multipath transmission for the inter-
842 net: A survey. IEEE Communications Surveys Tutorials 18 (4), 2887–
843 2925.

844 [41] LIP6-MPTCP, 2017. open source project repository.
845 URL <https://github.com/lip6-mptcp>

846 [42] Misseri, X., Rougier, J., Saucez, D., 2012. Internet routing diversity for
847 stub networks with a map-and-encap scheme. In: Proceedings of IEEE
848 International Conference on Communications, ICC 2012, Ottawa, ON,
849 Canada, June 10-15, 2012. pp. 2861–2866.
850 URL <https://doi.org/10.1109/ICC.2012.6363982>

851 [43] Nguyen, H. D. D., Phung, C., Secci, S., Felix, B., Nogueira, M., 2017.
852 Can MPTCP secure internet communications from man-in-the-middle
853 attacks? In: Conference on Network and Service Management, CNSM
854 2017, Tokyo, Japan, November 26-30, 2017. pp. 1–7.
855 URL <https://doi.org/10.23919/CNSM.2017.8255970>

856 [44] O.Bonaventure, C.Paasch, G.Detal, Jan 2017. Use cases and operational
857 experience with multipath tcp. RFC 8041, RFC Editor.
858 URL <https://tools.ietf.org/rfc/rfc8041.txt>

859 [45] Oliveira, R. V., Pei, D., Willinger, W., Zhang, B., Zhang, L., 2010. The
860 (in)completeness of the observed internet as-level structure. IEEE/ACM
861 Trans. Netw. 18 (1), 109–122.
862 URL <http://doi.acm.org/10.1145/1816288.1816297>

- [46] Peng, Q., Walid, A., Low, S. H., 2013. Multipath TCP algorithms: theory and design. In: ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS '13, Pittsburgh, PA, USA, June 17-21, 2013. pp. 305–316.
URL <http://doi.acm.org/10.1145/2465529.2466585>
- [47] Raiciu, C., Barré, S., Pluntke, C., Greenhalgh, A., Wischik, D., Handley, M., 2011. Improving datacenter performance and robustness with multipath TCP. In: Proceedings of the ACM SIGCOMM 2011 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Toronto, ON, Canada, August 15-19, 2011. pp. 266–277.
URL <http://doi.acm.org/10.1145/2018436.2018467>
- [48] Raiciu, C., Paasch, C., Barre, S., Ford, A., Honda, M., Duchene, F., Bonaventure, O., Handley, M., 2012. How hard can it be? designing and implementing a deployable multipath tcp. In: Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation. NSDI'12. USENIX Association, Berkeley, CA, USA, pp. 29–29.
URL <http://dl.acm.org/citation.cfm?id=2228298.2228338>
- [49] RIPE RIS, 2017. Ripe routing information service.
URL <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>
- [50] Routeviews, 2017. University of oregon route view projects.
URL <http://www.routeviews.org>
- [51] Sankar, A. U. P., Poornachandran, P., Ashok, A., Krishnan, M. R., Hrudya, P., 2016. B-secure: A dynamic reputation system for identifying anomalous BGP paths. In: Proceedings of the 5th International Conference on Frontiers in Intelligent Computing: Theory and Applications - FICTA 2016, Volume 1. pp. 767–775.
URL https://doi.org/10.1007/978-981-10-3153-3_76
- [52] Shaneman, K., Gray, S., 2004. Optical network security: Technical analysis of fiber tapping mechanisms and methods for detection and prevention. In: Proceedings of the IEEE MILCOM 2004. pp. 711 – 716.

- 895 [53] Starosielski, N., Yoo, D., Leong, R., Camacho, K., 2011. Critical nodes,
896 cultural networks: Re-mapping guam’s cable infrastructure. *Amerasia*
897 *Journal* 37 (3), 18–27.
898 URL <https://doi.org/10.17953/amer.37.3.m700712731t62754>
- 899 [54] Toonk, A., 2017. Bgpstream and the curious case of as12389.
900 URL <https://bgpmon.net/bgpstream-and-the-curious-case-of-as12389/>
- 901 [55] Tuexen, M., Stewart, R., Jan 2011. Stream control transmission protocol
902 (sctp) chunk flags registration. RFC 6096, RFC Editor.
903 URL <https://tools.ietf.org/html/rfc6096>
- 904 [56] UNSD, 2017. Standard country or area codes for statistical use (m49).
905 URL <http://unstats.un.org/unsd/methods/m49/m49regin.htm>
- 906 [57] White, J., Pilbeam, A., 2011. An analysis of coupling attacks in high-
907 speed fiber optic networks. In: *Proceedings of the Enabling Photonics*
908 *Technologies for Defense, Security, and Aerospace Applications VII*.
- 909 [58] Witcher, K., 2005. Extensions for Network-Assisted MPTCP Deploy-
910 ment Models. White Paper, SANS Institute.
911 URL [https://www.sans.org/reading-room/whitepapers/](https://www.sans.org/reading-room/whitepapers/physical/fiber-optics-security-vulnerabilities-1648)
912 [physical/fiber-optics-security-vulnerabilities-1648](https://www.sans.org/reading-room/whitepapers/physical/fiber-optics-security-vulnerabilities-1648)