

Measuring Node Decentralisation in Blockchain Peer to Peer Networks

Andrew Howell^{a,1}, Takfarinas Saber^{b,d}, Malika Bendecheche^{b,c,d,*}

^a*Blockdaemon, Los Angeles, California, USA*

^b*Lero – The Irish Software Research Centre*

^c*ADAPT – Science Foundation Ireland Research Centre*

^d*School of Computer Science, National University of Ireland, Galway, Ireland*

Abstract

New blockchain platforms are launching at a high cadence, each fighting for attention, adoption, and infrastructure resources. Several studies have measured the peer-to-peer network decentralisation of Bitcoin and Ethereum (i.e., two of the largest used platforms). However, with the increasing demand for blockchain infrastructure, it is important to study node decentralisation across multiple blockchain networks—especially those containing a small number of nodes. In this paper, we propose NodeMaps, a data processing framework to capture, analyse, and visualise data from several popular P2P blockchain platforms such as Cosmos, Stellar, Bitcoin, and Lightning Network. We compare and contrast the geographic distribution, the hosting provider diversity, and the software client variance in each of these platforms. Through our comparative analysis of node data, we found that Bitcoin and its Lightning Network layer 2 protocol are widely decentralised P2P blockchain platforms, with the largest geographical reach and a high proportion of nodes operating on The Onion Router (TOR) privacy-focused network. Cosmos and Stellar blockchain have reduced node participation with nodes predominantly operating in large cloud providers or well-known data centres.

*Corresponding author

Email addresses: andrew.howell12@mail.dcu.ie (Andrew Howell),
takfarinas.saber@nuigalway.ie (Takfarinas Saber), malika.bendecheche@dcu.ie
(Malika Bendecheche)

Keywords: Blockchain, Peer-to-peer, Decentralisation, Networks, Bitcoin, Lightning Network, Cosmos, Stellar, TOR

1. Introduction

Distributed ledger technology (DLT) has seen tremendous growth over the last decade, starting with Bitcoin, the first widely deployed blockchain technology, claims to be “A Peer-to-Peer (P2P) Electronic Cash System” [1]. In a recent study, Shrivastava et al. [8] discuss the dramatic rise of DLT technology, the variety of new platforms, tools, programming languages. Alongside this rise, we have witnessed the application of blockchain technology in various domains [44, 45].

Decentralisation is a cornerstone of P2P blockchain platforms. Srinivasan [31] discusses how it is important to quantify the term as it can be used to refer to the subsystems that comprise a blockchain platform. It is sufficiently difficult to bootstrap a new decentralised blockchain network with independent node operators without some centralised infrastructure [13]. Incentives must be balanced to promote engagement while taking into consideration network health by reducing centralisation of important node infrastructure [7].

Several studies have previously investigated the decentralisation in the blockchain. Some of them are theoretical and applicable to a wide range of blockchain platforms (e.g., Kwon et al. [2]), the rest of these studies/frameworks are: (i) restricted to identifying the widest number of active nodes [36], (ii) specific to limited blockchain platforms [14, 10, 9, 34, 36, 48], (iii) proposing intrusive tools and adaptations to the open-source blockchain client software to mimic a functioning blockchain node in peer discovery [9, 10, 11, 48], or (iv) focused on geographical distribution [14, 15], network topology [48] or mining power [49, 50].

Interacting with a blockchain requires a node, it’s not fair to assume that each developer or organisation will operate the latest node version and at their own infrastructure. This leads to an interesting scenario where Cloud and service provider companies fill the vacuum offering specialised blockchain infrastructure services. Therefore, it is crucial not to only consider the distribution of nodes

28 geographically, but also based on their software version and on the number of
29 entities operating their infrastructure.

30 In this work, we define and study decentralisation in: (i) node geographical
31 distribution, (ii) diversity of node hosting vendor, and (iii) variations in the
32 software version running on the node. We propose *NodeMaps*; a simple, user-
33 friendly and extensible framework for collecting, processing, and analysing
34 snapshots of various blockchain platforms at once. In this work, we start with four
35 well-know blockchain platforms (i.e., Bitcoin, Lightning Network, Cosmos, Stellar)
36 for which we have access to node data either from public sources (e.g., [14, 15]) or
37 using state-of-the-art efficient and non-intrusive scrappers [34, 36]. Our ultimate
38 goal is to grow the capability of our framework overtime by increasing the number
39 of networks as we identify suitable data sources/scrappers to collect their data.

40 Furthermore, NodeMaps offers several features that are interesting both
41 technically and scientifically:

- 42 • It enables the investigation of the geographical (and country-wise) distri-
43 bution of various blockchain platforms.
- 44 • It allows the identification of server infrastructure providers hosting the
45 nodes associated with the peer IP address (identifying potential availabil-
46 ity/reliability issues and threats from server providers working cahoot).
- 47 • It allows the identification and analysis of client software versions and
48 the investigation of the user agent to assess the variances in deployed
49 software—thus identifying potential legacy and fragmentation issues (sig-
50 nalling increased management, maintenance, and security hidden debts).

51 The remainder of this paper is organised as follows. Section 2 describes
52 the proposed NodeMaps framework. Section 3 details the data collection and
53 processing in our NodeMaps framework. Section 4 presents the results of the
54 data analysis. Section 5 presents our background, describes our used blockchain
55 networks, and reviews our related work. Section 6 concludes the work.

56 **2. NodeMaps: the Proposed Data Processing Framework**

57 In this section, we describe the composition of the proposed NodeMaps
58 framework¹, design choices (particularly in terms of data abstraction), and
59 implementation/deployment details.

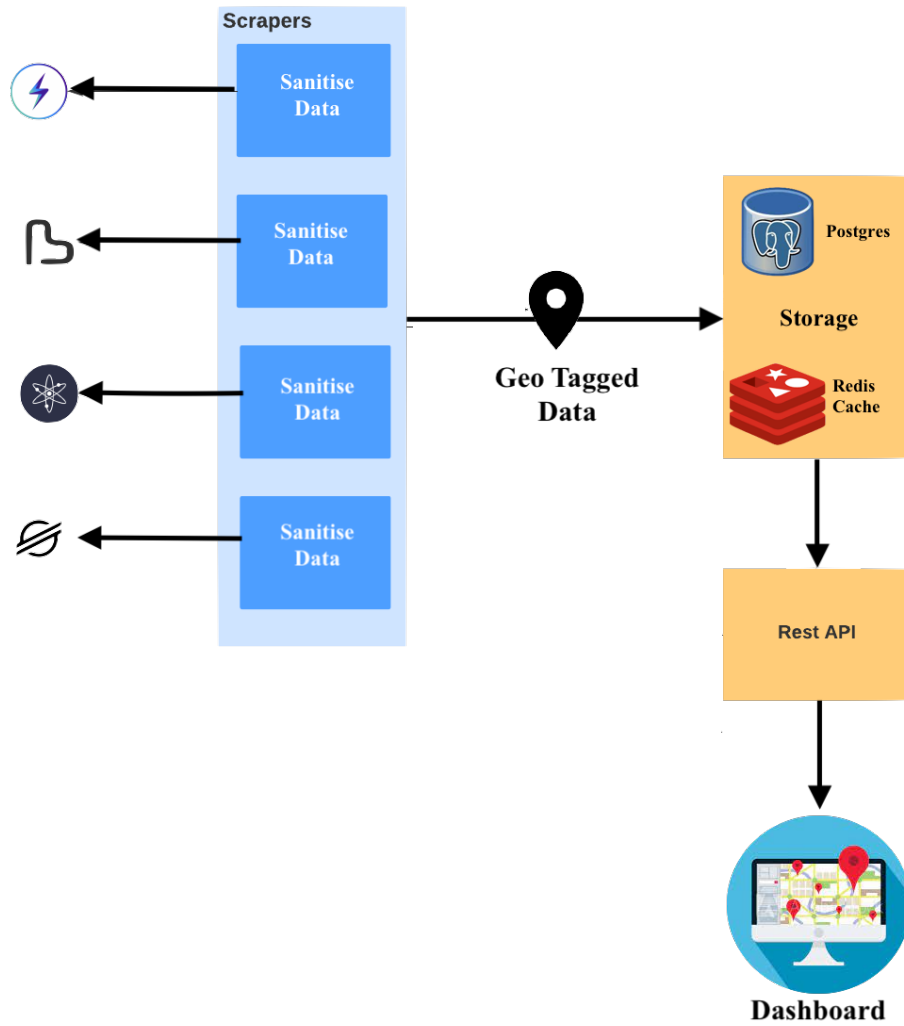


Figure 1: Overview of the Proposed Data Processing Framework Architecture

¹The source code of our NodeMaps framework will be made publicly available upon acceptance of this manuscript.

60 *2.1. NodeMaps Components*

61 Figure 1 shows an overview of the proposed framework (i.e., NodeMaps) to
62 capture, sanitise, store and present node data gathered from blockchain P2P
63 networks.

64 *2.1.1. Scrapers:*

65 A set of blockchain specific scrapping services periodically connect to a
66 remote data source to gather data. Each scrapping service performs a pipeline
67 of data processing actions in order to sanitise the gathered data in preparation
68 for storage. The designed scrappers are non-intrusive and keep the captured
69 data to a minimum. In Section 3, we discuss the scraping process for each of
70 our chosen protocols.

71 *2.1.2. Storage:*

72 Sanitised node data is written to a Postgres SQL database [42]. A Redis
73 Cache [43] is used to store daily snapshots of processed data for each protocol.
74 Redis is a caching system that works by temporarily storing information in a
75 key-value data structure, thus allowing for optimised retrieval of data during
76 analysis. Redis cache is popular as it is available in almost all major programming
77 languages. Therefore, it will also facilitate a future extension of the NodeMaps
78 platform to enable the analysis of different blockchain networks over time.

79 *2.1.3. REST API:*

80 REST API retrieves data from the storage systems, the API is utilised in
81 our analysis process.

82 *2.1.4. Dashboard:*

83 A basic dashboard details the project's goals and features charts demon-
84 strating the analysis (see Figure 2). The dashboard is built with ReactJS and
85 Highmaps modules. It also provides a visual representation of nodes that have
86 been Geo-tagged. Filters allow different datasets to be toggled on and off so

89 *2.2. Common Data Points Schema*

90 The objective of our work is to develop an extensible framework where
91 multiple blockchain platforms can be plugged with ease, thus enabling us to
92 handle/analyse data across several/different blockchains. Therefore, we built
93 NodeMaps in the form of a generic framework that abstracts data scraping and
94 sensitisation processes from the data storage, retrieval and presentation process.

95 We define a common schema that is capable of satisfying the data points
96 required for analysis activities. The scheme is inspired by the Bitnodes API [14],
97 which exposes a data object capable of storing information relating to a Bitcoin
98 node, its geographical and Autonomous System Number (ASN) properties. We
99 extend this data model with additional fields, identified after performing a gap
100 analysis of available node data from our chosen blockchain platforms.

101 *2.3. Implementation/Deployment*

102 The NodeMaps framework is implemented as a set of GoLang services that are
103 based on data models and objects defined with the Swagger Interface Description
104 Language. The node data model defines the common data model and it is
105 codified in a YAML syntax.

106 Open-source libraries are used to auto-generate GoLang packages and API
107 stubs based on the node data model. The choice of tooling allows for easy
108 extension of the node data model to new fields by simply updating the Swagger
109 definition and running a code-generator.

110 We deploy NodeMaps on a single server. All NodeMaps services (including
111 scrapers, data processors, and visualisation) execute as micro services on docker
112 containers within the same server.

113 **3. Data Collection/Processing**

114 In this section, we will discuss the varying approaches we took to capture
115 data snapshots of node data from Bitcoin, Lightning Network, Cosmos, and
116 Stellar blockchain platforms.

117 *3.1. Data Scraping From Bitcoin*

118 Past research has proposed novel ways of extracting P2P node information
119 from the Bitcoin network [10], we choose not to rehash these techniques, in favour
120 of utilising a network snapshot available via the Bitnodes API [14]. Bitnodes [33]
121 operate a network crawler that recursively sends `getaddr` messages to network
122 peers. Starting with a set of seed nodes the system recursively crawls all peers.
123 While there is no node coverage study on the scrapper, Bitnodes claims to capture
124 all Bitcoin nodes that are considered reachable (i.e., if they accept incoming
125 connections from their peers).

126 Bitnodes maintains frequent snapshots and makes them available via the
127 REST API. The Bitnodes scraper connects to the Bitnodes REST API and calls
128 the snapshots endpoint. This endpoint returns a list of available Bitcoin node
129 data snapshots denoted by a Unix Timestamp. The scraper identifies the latest
130 snapshot and issues another REST call to retrieve the node snapshot JSON
131 data.

132 The next stage of the data pipeline involves processing each node record
133 returned from Bitnodes. The application iterates over each node recorded in
134 the retrieved snapshot and performs data sanitisation (i.e., mapping Bitnodes
135 fields to those defined in the framework data model and setting default values
136 for empty fields).

137 Once each record has been sanitised, we store the data in the Postgres
138 database and update the Redis cache with a timestamped snapshot containing
139 all processed records which can be used for further analysis.

140 *3.2. Data Scraping from Lightning Network (LN)*

141 We use a similar technique to that described by Romiti et al. [34], thus taking
142 a full advantage of the various coverage assessments and validations performed
143 by the authors on their scraper. We deploy a Lightning Network Daemon (LND)
144 node with an infrastructure provider. Over time the node builds up a graph
145 of all other nodes that it has learned about during its P2P operations. LND

146 exposes an API endpoint that returns a JSON file which describes the LN node
147 graph.

148 The scraper starts by connecting to the endpoint exposed on our LND node
149 retrieve the built graph data and commence the sanitisation process. Firstly we
150 parse the graph data and we detect if the node record being processed uses Tor
151 (i.e., has a `.onion` address). Some nodes have no address or a private IP. Since
152 the aim of our work is to only process nodes that are exposed publicly to the
153 internet, we filter out nodes that do not have an `IPV4`, `IPV6` or a `.onion` address.
154 However, the framework could easily be adapted to keep all nodes for processing.
155 In our analysis below, a total of 11340 LN nodes were identified as having either
156 `IPV4`, `IPV6` or `.onion` addresses in the network snapshot (out of 12359 nodes).

157 Following the pre-checks, we perform data sanitisation to build our final data
158 model. First, we search for the ASN associated with node IP address using
159 MaxMind's `GeoLite2-ASN` database [35]. We also search for the geographical loc-
160 ation (i.e., `city`, `country_code`, `timezone`, `latitude`, and `longitude`) associated with
161 the node IP address by means of a lookup against the MaxMind `GeoLite2-city`
162 database [35]. Finally, the sanitised data set is saved to a Postgres database and
163 a snapshot is stored in Redis.

164 *3.3. Data Scraping from Cosmos*

165 We run a node with an infrastructure provider for our Cosmos P2P data
166 gathering. We set the configuration of the `gaiad` daemon to function as a seed
167 node (i.e., a special configuration of a Cosmos node to share its peer info and
168 discover as many peers as possible) and populated its initial seeds with the known
169 nodes from the infrastructure provider. We have also adjusted the number of
170 peers to which the node can connect to 1000 at a time. Furthermore, we have
171 increased the maximum number of connections and open file handlers on the
172 host operating system to 90000 in anticipation of a high connection count.

173 The Cosmos scraper is similar in design to those described in the previous
174 sections. However, the captured data requires recursive processing to reveal
175 additional data about the detected P2P nodes. The process we use is similar

176 to the open-source project developed by Chainlayer (i.e., cosmos-crawler [36]).
177 We query each known node that exposes its network information to collect its
178 details and the list of peers to which it is connected. We continue this process
179 recursively until no more peers are newly discovered. We have confirmed that the
180 number of identified nodes was similar to what was reported by CosmoScan1 on
181 the date the snapshot was taken. Currently (as of July 7th, 2022), CosmoScan
182 reports 398 nodes. Therefore, we are adding clarifications about the validation of
183 the number of nodes in Cosmo. While we were unable to find a coverage study
184 for the Cosmos crawler, we have confirmed that the number of identified nodes
185 was similar to what was reported by CosmoScan ² on the date the snapshot was
186 taken.

187 The first stage of data processing involves removing duplicate peer records
188 that might exist as peers can be the source of multiple connections with different
189 destinations. The next stage involves data sanitisation: peers with no public
190 IP addresses or those that could not be reached on their IP/port details are
191 discarded. Analysis of the address book data suggests that many peers are not
192 contactable after thousands of attempts. The final stage performs ASN and
193 Geo-Tagging analysis on the IP addresses of contactable peers as described in
194 the LN node data pipeline and persists the resulting data into the database.

195 *3.4. Data Scraping from Stellar*

196 The Stellar data scraper performs a similar technique as used for bitcoin
197 data capture. The process requests a network snapshot from a REST API
198 provided by Stellar Beats [15]. The API returns a comprehensive set of Stellar
199 network data points, information about the Quorum sets, high level statistics
200 and an array of all discovered nodes. Stellar Beat states that their data is
201 gathered every 3 minutes from the network using a network scraper that gathers
202 information about all nodes [15], thus offering us a greater confidence on the
203 achieved node coverage. There is little reason to re-implement the Stellar Beat

²<https://cosmoscan.net/cosmos/validators-stats>

204 system for our research, instead, we process a snapshot and perform our data
205 sanitisation pipeline. To achieve consistency in our results we ignore the Stellar
206 Beat Geo-Location and ASN data and leverage the common scraper utilities as
207 discussed in the previous sections. Once all data mapping has been completed
208 the snapshot is saved.

209 **4. Analysis of Findings**

210 In this section, we will review a snapshot of captured data for Bitcoin, LN,
211 Cosmos, and Stellar blockchain networks collected on 07/08/2021. We perform
212 an analysis of the data to ascertain the geographical distribution of nodes in
213 each network and hosting vendor diversity.

214 It has to be noted that it is common for nodes in a P2P network to join
215 and exit at any moment. Network latency, maintenance, service disruption and
216 human intervention are all reasons a node may exit a network or appear offline.
217 The view of peers in a network constantly changes and to our knowledge, there
218 is no 100% guaranteed way to discover all nodes at a given time. Additionally,
219 the quality of the collected data are as good as the quality of the data collectors
220 (i.e., crawlers and publicly available data through APIs). Therefore, the (lack of)
221 node coverage represents an important threat to the validity of our data analysis
222 and obtained insight.

223 *4.1. Top Autonomous System Numbers*

224 Our investigation yielded 1172 unique ASNs across the four blockchain
225 platforms we investigated. Table 1 shows the top ASNs in terms of aggregated
226 number of nodes.

227 Table 1 highlights that of all the detected nodes, a large portion, 54.28%
228 operate on the TOR network. TOR has the capability of executing programs
229 as hidden services, shielding the source IP address of the server running the
230 application. This provides a form of anonymity to the service operator and the
231 server hosting the applications, making it difficult to perform IP address analysis

ASNs	Node Count	Percentage
Tor network	14085	54.28
Hetzner Online GmbH	1337	5.15
OVH SAS	648	2.50
AMAZON-02	591	2.28
DIGITALOCEAN-ASN	535	2.06
AMAZON-AES	479	1.85
COMCAST-7922	418	1.61
Online S.a.s.	327	1.26
SHRD SARL	322	1.24
GOOGLE	309	1.19
Contabo GmbH	281	1.08
COGENT-174	247	0.95
ATT-INTERNET4	172	0.66
UUNET	169	0.65
Vodafone GmbH	143	0.55
Deutsche Telekom AG	136	0.52
AS-CHOOPA	106	0.41
Vodafone Libertel B.V.	93	0.36
Alibaba US Technology Co., Ltd.	88	0.34
Hangzhou Alibaba Advertising Co.,Ltd.	86	0.33
Others	5379	20.73

Table 1: Top ASNs for All Blockchain Platforms

232 of the target. Attempting to de-anonymise TOR services is beyond the scope of
 233 this research. For the remainder of this section, we will treat TOR as a provider.
 234 The private nature of TOR means it is possible that some unknown percentage
 235 of the nodes could operate on any of the other ASNs identified.

236 When we look at the aggregated ASN data across all the investigated block-
 237 chains in Table 1, we see that Hetzner Online GmbH ASN is the second largest
 238 provider of servers for blockchain platforms, with 5.15% of all nodes. OVH SAS
 239 comes in third place of all detected ASNs and have just under half the amount
 240 of nodes in Hetzner.

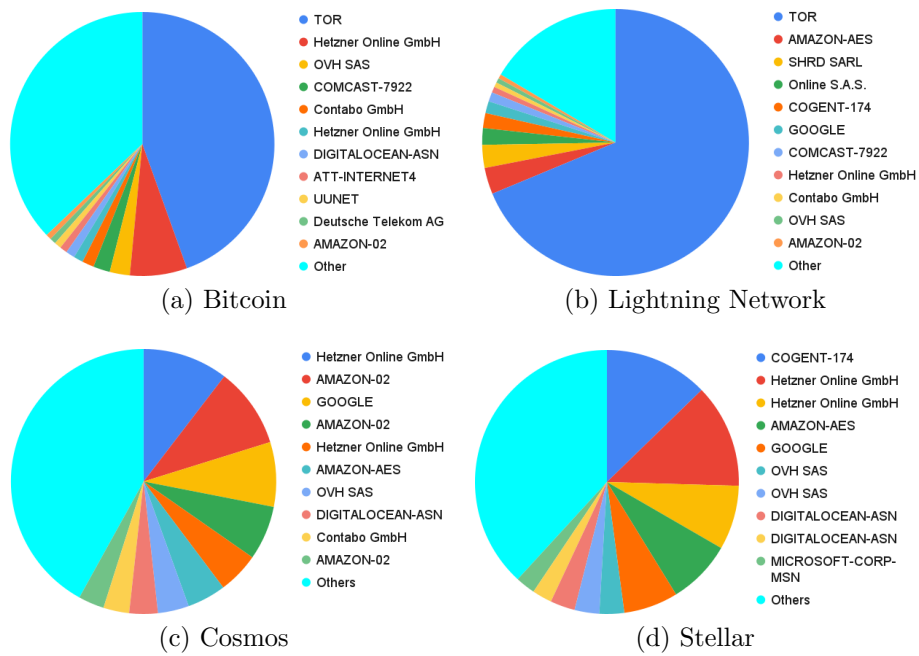


Figure 4: Distribution of Nodes Per ASN in Each Blockchain Platform

241 Figure 4 shows that Hetzner Online Gmb is also the top detected provider
 242 (other than TOR) in both Stellar and Bitcoin, also coming in second for Cosmos
 243 nodes and eleventh in LN. We will discuss the breakdown of each blockchain
 244 platform in the following sections.

245 *4.2. ASNs Per Country*

246 Excluding TOR nodes, Table 2 places the United States (US) as the top
247 node location globally, followed by Germany (DE) and France (FR).

Country	n/a	US	DE	FR	CA	NL	GB	FI	RU	CH	SG	Other
#Nodes	14085	3272	2281	1237	649	607	412	255	236	219	214	2484
%Nodes	54.28	12.61	8.79	4.77	2.50	2.34	1.59	0.98	0.91	0.84	0.82	9.57

Table 2: Top Countries for Node Location

248 Table 3 provides a summary of the top 20 ASNs identified per country
249 aggregated from each blockchain network snapshot. From this data, we can see
250 that Hetzner Online GmbH in Germany has the largest concentration of nodes
251 per provider and per country at 4.32%. When we compare the top node countries
252 and the top ASNs per country, we can correlate the number of different ASNs
253 per region, Table 3 shows that there are 7 ASNs in the US, 4 in Germany and 3
254 in France where the majority of nodes are deployed, indicating that blockchain
255 node infrastructure tends to centralise around a few key providers in specific
256 regions.

257 *4.3. Findings in Bitcoin Platform*

258 We analyse the finding in the Bitcoin platform in terms of ASNs, geographical
259 locations, and software versions.

260 *4.3.1. ASNs:*

261 A total of 14129 Bitcoin nodes were identified in the network snapshot. Table
262 4 shows a breakdown of the top 10 ASNs detected when analysing the data.
263 TOR commands the highest overall slice of the network with 6290 nodes equating
264 to 44.52% which indicates that a major part of Bitcoin users is privacy-focused.

265 The remaining ASNs are a mix of cloud, bare metal server data centers, and
266 Internet Service Providers (ISP's). The ability to easily run a node on basic
267 hardware and the culture surrounding the project are many factors that explain
268 why the node proliferation is so high.

ASN	Country Code	Node Count	Percentage
TOR	n/a	14085	54.28
Hetzner Online GmbH	DE	1121	4.32
AMAZON-AES	US	479	1.85
OVH SAS	FR	422	1.63
COMCAST-7922	US	418	1.61
SHRD SARL	FR	322	1.24
Online S.a.s.	FR	310	1.19
Contabo GmbH	DE	281	1.08
GOOGLE	US	242	0.93
COGENT-174	CA	238	0.92
DIGITALOCEAN-ASN	US	225	0.87
Hetzner Online GmbH	FI	215	0.83
AMAZON-02	US	178	0.69
ATT-INTERNET4	US	172	0.66
UUNET	US	169	0.65
Vodafone GmbH	DE	143	0.55
Deutsche Telekom AG	DE	136	0.52
OVH SAS	CA	123	0.47
DIGITALOCEAN-ASN	DE	109	0.42
AMAZON-02	IE	93	0.36
Others		6470	24.93

Table 3: Top ASNs Per Country

269 Hetzner Online GmbH is the largest detected ASN with 6.99% (i.e., 987) of
 270 publicly identifiable Bitcoin nodes. The remaining network nodes result in 910
 271 different detected ASNs, consisting of a varied mix of cloud providers and global
 272 ISP's with 13.84% of ASNs hosting less than the 10 nodes and 464 unique ASNs
 273 hosting only a single node.

ASN	Country	#Nodes	%Nodes
TOR	n/a	6290	44.52
Hetzner Online GmbH	DE	987	6.99
OVH SAS	FR	347	2.46
COMCAST-7922	US	289	2.05
Contabo GmbH	DE	205	1.45
Hetzner Online GmbH	FI	164	1.16
DIGITALOCEAN-ASN	US	155	1.10
ATT-INTERNET4	US	129	0.91
UUNET	US	119	0.84
Deutsche Telekom AG	DE	105	0.74
AMAZON-02	US	94	0.67
Other		5245	37.12

Table 4: Bitcoin ASNs Per Country

274 *4.3.2. Geographical Locations:*

275 Table 5 provides a geographic breakdown of node locations by country. 6290
 276 nodes operate on TOR meaning the location is unknown. 13.79% of nodes are
 277 located in the United States followed closely by Germany at 12.82%. China
 278 where Bitcoin miners were recently pressured to shut down [37] still operates
 279 1.06% of nodes. The remaining 13.61% of publicly identifiable nodes are spread
 280 across 81 countries where 42 of those countries feature less than 10 nodes.

Country	TOR	US	DE	FR	NL	CA	GB	RU	FI	CN	Others
#Nodes	6290	1948	1811	575	429	319	271	219	194	150	1923
%Nodes	44.52	13.79	12.82	4.07	3.04	2.26	1.92	1.55	1.37	1.06	13.61

Table 5: Bitcoin Nodes Per Country

281 *4.3.3. Software Versions:*

282 Table 6 presents the top ten software versions in the network snapshot. The
 283 Bitcoin `user agent` property indicates the software version of the node (i.e.,
 284 the field exchanged during the node peering process). 43% of nodes (i.e., 6184)
 285 reported the latest Bitcoin Core client software version `Satoshi:0.21.1`, while
 286 21.23% of nodes featured the previous release. 77% of nodes detected reported
 287 releases from within the last year and a half [38] (i.e., versions 0.20.* and 0.21.*),
 288 indicating high engagement from node operators to keep nodes up to date.

289 It is worth noting that 7 nodes reported `Satoshi:0.8.1` which was released
 290 in 2013 [38] and it is unclear if these nodes operate as expected. 10.49% (1482)
 291 of reported node user agents contain a varying mix of reported software clients
 292 and versions. 116 nodes appear to have manually modified user agent strings
 293 and utilise the field as a form of P2P digital graffiti, where the field contains a
 294 personalised message.

Version	#Nodes	%Nodes
Satoshi:0.21.1	6184	43.77
Satoshi:0.21.0	3000	21.23
Satoshi:0.20.1	961	6.80
Satoshi:0.20.0	757	5.36
Satoshi:0.18.0	397	2.81
Satoshi:0.18.1	389	2.75
Satoshi:0.15.0.1	285	2.02
Satoshi:0.19.1	265	1.88
Satoshi:0.19.0.1	229	1.62
Satoshi:0.17.1	180	1.27
Others	1482	10.49

Table 6: Bitcoin User Agent Version

295 *4.4. Findings in Lightning Network Platform*

296 We analyse the finding in the LN platform in terms of ASNs and geographical
 297 locations. Note that we do not analyse software versions in the LN platform as
 298 we were unable to gather such information from LN nodes.

299 *4.4.1. ASNs:*

300 A total of 11340 LN nodes were identified in the network snapshot (out of
301 12359 nodes, or 91.75%). The remaining 8.25% of the nodes (i.e., 1019 nodes,)
302 had neither IPV4, IPV6 nor .onion addresses—thus we filtered them out. Table 7
303 shows a breakdown of the top 10 ASNs detected when analysing LN graph data.
304 Similar to Bitcoin, TOR commands the highest overall slice of the network, with
305 7795 nodes equating to a substantial 68.74%. The remaining ASNs combined
306 add up to 16.13% of the network. As with Bitcoin, the providers are a mix of
307 cloud, bare metal server data centers and Internet Service Providers (ISP's). The
308 remaining 15.13% of the network nodes result in 500 different detected ASNs,
309 consisting of a varied mix of cloud providers and global ISP's. 462 (7.81%) of
310 ASNs host less than 10 nodes, 278 unique ASNs host only a single node.

311 AMAZON-AES has the second-highest node count after TOR with 3.20% of
312 LN nodes, on further investigation of the captured data we identify 318 nodes
313 that share 33 IP addresses. In some cases, there are over ten unique LN node
314 public keys reporting the same IP address and port 9735, LN's P2P port. To
315 our knowledge, it is not possible to run multiple LND instances on the same
316 port on the same machine with the same IP address at any given time. Each
317 peer record appears to share a similar signature for the alias field suggesting
318 they are created by an automated process. Attempting to connect an LND node
319 to a set of peers fails. We assume these anomaly peers are stale data that has
320 not been pruned from the LN graph.

321 SHRD SARL makes up 2.78% of the LN nodes, analysing the alias field for
322 these nodes, it appears that the majority are operated by an LN node as a
323 service provider, Nodl. On further inspection of the alias field of all nodes in
324 the LN `graph.json` snapshot, Nodl also appears to operate a large portion of
325 nodes in Online S.a.s. (2.07%), the third major ASN detected in the snapshot
326 data, suggesting that SHRD SARL and Online S.a.s. are their primary hosting
327 partners.

ASN	Country	#Nodes	%Nodes
TOR	n/a	7795	68.74
AMAZON-AES	US	363	3.20
SHRD SARL	FR	315	2.78
Online S.a.s.	FR	229	2.02
COGENT-174	CA	205	1.81
GOOGLE	US	164	1.45
COMCAST-7922	US	128	1.13
Hetzner Online GmbH	DE	80	0.71
Contabo GmbH	DE	66	0.58
OVH SAS	FR	65	0.57
AMAZON-02	US	59	0.52
Other		1871	16.50

Table 7: Lightning Network ASNs Per Country

328 *4.4.2. Geographical Locations:*

329 Table 8 provides a geographic breakdown of LN node locations by country.
330 7795 or 68.74% of nodes operate on TOR meaning the location is unknown.
331 10.53% of nodes are located in the United States, although many of the node
332 records appear to be stale and uncontactable as discussed previously. 5.71%
333 of LN nodes operate from France primarily in SHRD SARL and Online S.a.s,
334 followed by Germany at 3.32%. According to the LN graph snapshot, 31.26% of
335 nodes can be detected in 76 different countries.

Country	TOR	US	FR	DE	CA	NL	GB	CN	IT	JP	Others
#Nodes	7795	1194	648	376	282	157	134	67	43	43	601
%Nodes	68.74	10.53	5.71	3.32	2.49	1.38	1.18	0.59	0.38	0.37	5.30

Table 8: Lightning Nodes Per Country

336 *4.5. Findings in Cosmos Platform*

337 We analyse the finding in the Cosmos platform in terms of ASNs, geographical
338 locations, and software versions.

339 *4.5.1. ASNs:*

340 A total of 317 Cosmos nodes were identified in the network snapshot. Ident-
341 ified nodes were hosted in a total of 39 unique ASNs, with the top ten ASNs

342 hosting 58.04% of them.

343 Table 9 shows a breakdown of the top 10 identified ASNs with the number
344 (percentage) of nodes running in their infrastructure. In contrast to Bitcoin
345 and LN, there were no nodes detected to be operating on the TOR network.
346 AMAZON-02 has top 10 ASNs in 3 different countries, hosting 62 Cosmos nodes
347 in total (or 19.55%), followed by Hetzner Online GmbH with 15.46% and Google
348 with 7.89%. Cosmos continues to diverge from Bitcoin and LN data as there are
349 no ISP’s in the top ten, only server infrastructure providers.

350 Cosmos proof of stake blockchain has some complexities in its design as a
351 central hub of an “Internet of Blockchains”. Node operators must stake ATOM
352 tokens to a validator node in order to become a block producer. The value of
353 stake requires operating a node in the active set [47] which could have the effect
354 of limiting node operations to entities that are well capitalised. Additionally,
355 the usability and the steep learning curve for developers creating blockchain
356 applications for the first time has been a major issue that keeps developers away
357 from using Cosmos. However, it is true that recently, Cosmos has been fixing this
358 issue as the Cosmos SDK (a modular framework for creating secure blockchain
359 applications on top of Tendermint) is changing rapidly (e.g., with scaffolding
360 being upgraded)—although, there is still the issue that some changes break other
361 changes, thus affecting the overall system stability [46].

ASN	Country	#Nodes	%Nodes
Hetzner Online GmbH	DE	33	10.41
AMAZON-02	IE	31	9.78
GOOGLE	US	25	7.89
AMAZON-02	US	21	6.62
Hetzner Online GmbH	FI	16	5.05
AMAZON-AES	US	15	4.73
OVH SAS	CA	12	3.79
DIGITALOCEAN-ASN	US	11	3.47
Contabo GmbH	DE	10	3.15
AMAZON-02	SG	10	3.15
Others		133	41.96

Table 9: Cosmos ASNs Per Country

362 *4.5.2. Geographical Locations:*

363 Table 10 provides a geographic breakdown of node locations by country. 90
 364 or 28.39% of nodes operate in the United States, 68 (21.45%) in Germany, and
 365 31 (9.78%) from Ireland. In total Cosmos nodes were detected in 23 countries,
 366 20.19% of those countries operate less than 10 nodes. These numbers are not
 367 indicative of the overall Cosmos network as a secure validator network would
 368 consist of public sentry nodes and private validator nodes not publicly addressable
 369 as described in Section 5.2.3.

Country	US	DE	IE	SG	CA	FI	NL	FR	CN	KR	Others
#Nodes	90	68	31	17	17	16	14	9	9	7	39
%Nodes	28.39	21.45	9.78	5.36	5.36	5.05	4.42	2.84	2.84	2.21	12.30

Table 10: Cosmos Nodes Per Country

370 *4.5.3. Software Versions:*

371 Table 11 shows the reported software versions in the Cosmos network snapshot.
 372 66.88% (212) of nodes reported version `v0.34.11`, while `v0.34.9` was detected
 373 2.21% (7) of the time. Due to the node scraping technique used we were unable
 374 to find data associated with 30.91% (98) of the known network. Data related to
 375 peers contained in the `gaiad` address book file does not have much-identifying
 376 data. Similar to LN software versions, further development would be required to
 377 probe the remote peer during the P2P process. Alternatively running a larger
 378 set of seed nodes may provide a deeper view of the network.

379 Analysis of the captured versions suggests high engagement by node operators
 380 to keep nodes up to date. Many newer POS protocols like Cosmos often require
 381 nodes to all run similar software versions as there are often network-wide upgrades
 382 that may break some nodes (e.g., the recent Stargate update [39]).

383 *4.6. Findings in Stellar Platform*

384 We analyse the finding in the Stellar platform in terms of ASNs, geographical
 385 locations, and software versions.

Version	#Nodes	%Nodes
v0.34.11	212	66.88
n/a	98	30.91
v0.34.9	7	2.21

Table 11: Cosmos Software Versions

386 *4.6.1. ASNs:*

387 A total of 165 Stellar nodes were identified in network snapshot. All detected
388 nodes are hosted in a total of 26 unique ASNs.

389 Table 12 shows a breakdown of the top 10 ASNs detected. Similar to Cosmos
390 there were no nodes identified as operating on the TOR network. Hetzner Online
391 GmbH is the top ASN with 34 nodes clocking in at 20.61% of the public
392 addressable network. COGENT-174 features 21 nodes or 12.73% and Google
393 US with 11 nodes or 6.67% of the network.

394 It is worth noting that all ASNs hosting more than one node are all tier 1
395 cloud platforms or data centres (hosting a total of 90.30% of all detected Stellar
396 nodes). The remaining 9.70% of nodes are hosted on 15 providers (one node per
397 provider).

ASN	Country	#Nodes	%Nodes
COGENT-174	CA	21	12.73
Hetzner Online GmbH	DE	21	12.73
Hetzner Online GmbH	FI	13	7.88
AMAZON-AES	US	13	7.88
GOOGLE	US	11	6.67
OVH SAS	CA	5	3.03
OVH SAS	FR	5	3.03
DIGITALOCEAN-ASN	NL	5	3.03
DIGITALOCEAN-ASN	US	4	2.42
MICROSOFT-CORP-MSN	US	4	2.42
Others		63	38.18

Table 12: Stellar ASNs Per Country

398 *4.6.2. Geographical Locations:*

399 Table 13 provides a geographic breakdown of node locations by country. 40
400 (24.24%) of nodes operate in the United States, 31 (18.79%) in Canada, and
401 26 (15.76%) operate in Germany. In total Stellar nodes were detected in 19
402 countries, 5 of which operate 73.33% of the public-facing nodes.

403 Overall, the data shows that the network underpinning Stellar does not have
404 the same geographical node distribution and hosting provider decentralisation as
405 Bitcoin. This is probably due to the consensus design choices which centralise
406 validation activities to a low number of node operators.

Country	US	CA	DE	FI	SG	NL	JP	FR	BE	GB	Others
#Nodes	40	31	26	13	11	7	5	5	4	4	19
%Nodes	24.24	18.79	15.76	7.88	6.67	4.24	3.03	3.03	2.42	2.42	11.52

Table 13: Stellar Nodes Per Country

407 *4.6.3. Software Versions:*

408 Table 14 shows the reported software versions in the Stellar network snapshot.
409 A strong 39.39% of nodes reported the latest release at the time of writing this
410 paper (i.e., `stellar-core 17.3.0`). In total 66.67% of nodes reported releases
411 from within the last six months, indicating a high engagement from Stellar node
412 operators.

413 It is worth noting that 19 nodes reported a software version that appears to
414 be custom-built (denoted by the `-dirty` postfix in the captured version string)
415 which we assume to be based on Stellar Core to some degree. 12 nodes reported
416 operating on unique software versions.

417 **5. Background and Related Work**

418 *5.1. Blockchain Generations*

419 The successful application of Nakamoto’s Bitcoin paved the way for a myriad
420 of alternative blockchain platforms, each with a twist on the original design.

Version	#Nodes	%Nodes
stellar-core 17.3.0 (0b4c12a...)	65	39.39
stellar-core 17.1.0 (fbc0325...)	25	15.15
stellar-core 17.2.0 (e47d483...)	20	12.12
v17.1.0	9	5.45
stellar-core 17.0.0 (096f6a7...)	8	4.85
v16.0.0-129-gb0671b82-dirty	4	2.42
stellar-core 17.1.0 (6c86d89...)	4	2.42
e0ae42ee-dirty	3	1.82
ee87cdcb-dirty	3	1.82
c848b944-dirty	3	1.82
Others	21	12.73

Table 14: Stellar Software Versions

421 **Generation 1** Bitcoin is considered first-generation blockchain technology,
422 it is effectively a P2P transaction settlement system with its own native currency
423 that requires no central entity to operate. Many of Bitcoin’s core principles
424 remain in newer systems, yet developers have introduced a key alterations.

425 **Generation 2** blockchains like Ethereum, innovated on the original P2P
426 ledger system. Buterin discusses in the Ethereum whitepaper that potentially
427 the most important part of Bitcoin was the underlying blockchain technology
428 and its mechanisms for achieving distributed consensus [4] while highlighting
429 the limitations of Bitcoins on-chain scripting language. Buterin proposed a new
430 blockchain platform with a Turing-complete programming language built in.
431 Ethereum’s enhancement would lead the way to decentralised smart contracts
432 that reside and execute on the blockchain and can be triggered by state transitions
433 associated with Ethereum accounts [4].

434 **Generation 3** blockchains like Polkadot [40], are designed to meet the scaling
435 challenges faced with prior iterations of the technology. In order to handle high
436 throughput global scale use cases demand, the core blockchain technologies have
437 been refined or rethought to address future demand [16].

438 The current blockchain landscape is a complex collection of competing net-
439 works mostly operating in their own sandbox with some purported unique
440 differentiating feature. There is also an increasing range of Layer 2 protocols like

441 the Bitcoin Lightning Network (LN), that complement the underlying blockchain
442 and generally add some form of scaling solution. Bitcoin and Ethereum despite
443 their reported scaling issues [18] can still be considered the base layer of the
444 whole blockchain industry due to their huge market share and garnered attention.
445 With new blockchain networks emerging every day and generation of blockchains
446 being developed for scale, we are witnessing the transition from multiple distinct
447 blockchains into an internet of blockchains where we will see cross protocol
448 bridges, asset transfer and decentralised application (DApp) portability [17].

449 The resulting future could lead to a complex interconnected blockchain fabric,
450 where specialised skills are required to participate in network operations.

451 *5.2. Investigated Blockchains*

452 *5.2.1. Bitcoin*

453 Bitcoin is a P2P blockchain that leverages Proof of Work (POW) to mint
454 new blocks by using the SHA256 algorithm repeatedly; new blocks are generated
455 if the result conforms to a specific signature. Historically any node possessed the
456 capability to mine new blocks using the central processing unit (CPU) of the host
457 device. Driven by the Bitcoin block rewards paid to miners for creating a new
458 block, the mining industry is now dominated by Application Specific Integrated
459 Circuits (ASIC) computers, block production is centralised in a number of mining
460 pools [19, 20]. CPU mining is no longer a viable option for Bitcoin mining.

461 If a mining node discovers a new block it broadcasts it to its network of peers,
462 each peer then propagates the message to its peers using its gossip protocol, the
463 decentralised consensus process then determines the block validity. If a block is
464 accepted by the network, the miner is rewarded with freshly minted Bitcoin. In
465 a POW system network security is the conversion of large quantities of energy
466 into cryptographic hashes, the incentivisation of the mining function has evolved
467 into a pay-to-play arms race where more hashing power and fervent desire to
468 reduce operational expense equates to greater rewards. Research suggests that
469 over the past decade the complexity of mining has led to the centralisation
470 of mining activity [2, 19, 20]. Mining pools are responsible for multiplexing

471 mining resources into a shared pool, cryptographic computation is split between
472 pool participants and rewards are shared proportionally to contributed compute
473 power.

474 Roughly every two weeks (2,016 blocks) the Bitcoin network automatically
475 re-tunes the difficulty associated with block minting. The network sets a target
476 of difficulty that equates to roughly a ten-minute block time. Block space is
477 also limited to 1MB (although, since SegWit compresses transactions, upgraded
478 nodes do see blocks larger than 1 MB [41]), transactions have to be at least 250
479 bytes. Block size coupled with the target block time of 10 minutes means that
480 the network can handle roughly 7 transactions per second [18]. This limitation
481 has led to scaling issues and much debate [21]. Increasing the blocksize and
482 shortening block times might seem like an immediate solution to increasing
483 performance yet there are trade-offs in every decision. Keeping the parameters
484 at the current values means there is some form of predictability in ledger storage
485 and node operating requirements, a Bitcoin full node can operate comfortably
486 on an ARM mini computer like a Raspberry PI [22, 23]. Increasing parameters
487 could have a knock-on effect of increasing hardware requirements to run a node,
488 smaller node operators could be priced out resulting in node centralisation.

489 *5.2.2. Lightning Network*

490 The Lightning Network (LN) is a Bitcoin Layer 2 transaction scaling solution
491 that leverages off chain payment channels between two parties. Poon et al. [24]
492 describe the Bitcoin network as a gossip protocol, where each ledger state modi-
493 fication is propagated to each node via a gossip mechanism. This node chatter
494 ensures that the node has the required information to form a consensus [24].
495 This type of network communication is expensive as all nodes must validate the
496 transactions, the solution is vital for consensus but limits transaction scalability.
497 Poon et al. [24] proposed a solution where transactions between two parties
498 could move to an off-chain payment channel where only Alice and Bob know
499 about transactions between each other [24]. In order for a payment channel
500 to be created, each actor must participate in a series of transactions to create

501 and fund the channel using on-chain Bitcoin. The Bitcoin is locked in a 2-of-2
502 multi signature address with conditions that allow for each party to unlock their
503 respective balance if specific conditions are met. Once funds are locked on the
504 base chain, the LN allows each party to transfer the value of the payment channel
505 between each other without having to broadcast any data on-chain. The channel
506 can be settled if either party wishes to exit, or specific states are detected.

507 The LN consists of multiple nodes each can have numerous channels, the
508 system is capable of routing multi-hop payments to other system participants
509 by leveraging the network of interconnected payment channels. All transactions
510 are backed by on-chain Bitcoin secured at the base layer, no transaction data
511 is broadcast to the Bitcoin blockchain meaning transactions can happen at
512 lightning speed, fees are kept low as mining is not required, participants only
513 pay the network routing fee defined by the intermediary nodes.

514 LN is a Layer 2 solution and thus requires access to a Bitcoin node for specific
515 activities that require communication with the base chain. Lightning Nodes
516 typically run a hot wallet, meaning that the nodes indirectly have access to
517 Bitcoin funds, this also adds a custodial aspect to operating an LN node, which
518 could lead to many participants opting to operate their own node [22, 23]. LN
519 nodes must communicate with other peers, operating a node on the LN with a
520 publicly facing address could lead to an increased security risk as it is trivial to
521 correlate an IP address to the LN nodes balance.

522 *5.2.3. Cosmos*

523 Cosmos is a blockchain platform built on the Tendermint consensus al-
524 gorithm [26]. The Cosmos ecosystem consists of many independent blockchain
525 zones, the first of which is called the Cosmos Hub. Each zone is capable of
526 communicating with each other via a novel Inter-Blockchain Communication
527 Protocol (IBC), parallel blockchains can all interact, transferring assets from
528 one zone to another [6].

529 The Cosmos blockchain utilises Proof-of-Stake (POS) in favour of POW
530 mining, the block minting process is a similar exercise although, in place of

531 physical ASIC miners, there exist validators. Cosmos network participants can
532 bond or delegate ATOM tokens (the native currency of the Cosmos blockchain)
533 to a validator. The validator is a special type of node that has the power to
534 vote on block proposals, its voting rights are proportionally weighted based on
535 the validator cumulative stake. The validator broadcasts signed cryptographic
536 signatures to the network when voting on the next block, in exchange for
537 confirmed validating activities nodes are paid a block reward.

538 The Tendermint consensus protocol utilised by Cosmos requires a fixed known
539 set of validators [6]. Currently, the network has 125 validator nodes in the active
540 set [27]. In order to become an active validator, a node must reach at least
541 position 125, currently 03/08/2021, this would require 33,000 ATOM [27].

542 POS networks often adopt a slashing mechanism to keep nodes honest. This
543 is effectively achieved using a slash which involves burning a portion of the
544 node's stake and preventing the node from voting on blocks for some time. If
545 a node double signs a block, the protocol interprets this as an attack on the
546 consensus system and the validate node will be identified and a portion of the
547 node's stake will be slash [28]. Sustained network downtime is also considered a
548 slash-able event, based on this, validator node operators must conform to best
549 practices when deploying node infrastructure. Typically Cosmos validator node
550 infrastructure consists of a public layer of sentry nodes connected by a private
551 link to a protected validator node [29], operating a highly available Cosmos
552 validator deployment has a high cost associated with the entry requirements.

553 *5.2.4. Stellar*

554 Stellar is an open network for money [5] aiming to introduce competition into
555 the international payment markets by leveraging DLT to send money around
556 the world quickly and cheaply. Its protocol nativity supports various trading
557 features (e.g., order books, cross-asset payments [30]).

558 The Stellar Consensus Protocol (SCP), a Byzantine agreement protocol [5]
559 introduced a new consensus mechanism that the Stellar blockchain uses to
560 facilitate secure transactions across a network of un-trusted intermediaries.

561 Organisations in the Stellar network choose other specific organisations to
562 interact with. The system mirrors the interconnected nature of the traditional
563 financial system where inter-bank relations are commonplace. Stellar network
564 nodes operate in Quorum Sets. Nodes only see others that are part of the
565 quorum, a view of the quorum can be ascertained as each node will learn of all
566 others [30]. The key innovation in SCP is the open-membership approach taken
567 to quorum sets constantly evolving with new participants joining the system [30].

568 *5.3. Investigation of Blockchain Decentralisation*

569 Several works have investigated the decentralisation in different blockchain
570 platforms or have put forwards means/tools to help us study their decentralisation.

571 Some of the works are theoretical enabling them to be applicable to a wide
572 range of blockchain platforms. For instance, Kwon et al. [2] proves the lack of
573 decentralisation in permissionless platforms.

574 There is also a large number of applied works. However, they are mostly
575 specific to very limited blockchain platforms—mostly Bitcoin and Ethereum. For
576 instance, Bitnodes [14] investigate node geographical localisation in the Bitcoin
577 blockchain, Kim et al. [9] measure Ethereum network peers, whereas, Gencer
578 et al. [10] measure decentralisation in Bitcoin and Ethereum Networks. This
579 said, some works are also studying other platforms such as Cao et al. [48] who
580 investigate the Monero network.

581 In terms of conducted work, the starting point of all these works is to draw
582 the best picture of various networks by developing tools to crawl, identify, and
583 scrap as much node data as possible with or without much validation. For
584 instance, Romiti et al. [34] put forward a method to identify nodes on the
585 lightning network and performed various validations for their crawler, whereas
586 Chainlayer put forward the Cosmos-Crawler [36] but without any guaranty of
587 coverage.

588 Some works propose intrusive tools and adaptations to the open-source block-
589 chain client software to mimic a functioning blockchain node in peer discovery.
590 For instance, Kim et al. [9] developed *NodeFinder*; an open-source scanning and

591 monitoring of Ethereum’s P2P network based on Geth (a command-line interface
592 for running Ethereum node implemented in Go Language) to identify active
593 nodes and periodically retrieves their client information. Gencer et al. [10] built
594 *Falcon Relay Network* to serve as a backbone for ferrying blocks and to measure
595 decentralisation in Bitcoin and Ethereum Networks. Venati [11] adapted the
596 *NodeFinder* proposed by Kim et al. [9] for the purposes of node counting in the
597 Ethereum private network and Ethereum public network, performed connections
598 at a higher rate, and measured their impact on the network.

599 In terms of decentralisation study, most of the works are focused on geo-
600 graphical distribution (e.g., [14, 15, 34]). Cao et al. [48] push the boundary
601 as they attempt to infer Monero’s topology, size, node distribution, and node
602 connectivity and found that it is highly centralised.

603 There are also studies which assess other aspects of blockchain platforms.
604 For instance, Alzayat et al. [49] analyse inefficiency and inequality in the mining
605 process in POW Blockchains, whereas Cao et al. [50] attempt to define metrics
606 to characterise the impact of network delay on Bitcoin mining.

607 **6. Conclusion**

608 In this paper, we proposed NodeMaps, an extensible framework to capture,
609 analyse, and visualise decentralisation data from several popular blockchain
610 platforms such as Bitcoin, Lightning Network, Cosmos, and Stellar. Leveraging
611 NodeMaps, we also performed an IP address analysis a snapshot of each of these
612 blockchain platforms to compare and contrast the geographic, ASN and version
613 distributions of their nodes.

614 Our analysis showed the decentralisation in Bitcoin and Lightning Network
615 with identifiable nodes hosted by several ASNs in numerous countries and using
616 a wide range of software versions. However, it also highlighted the user focus on
617 privacy as a large percentage of nodes run on TOR. Our analysis also showed
618 that Stellar (a competing network which claims similar open money principles)
619 does not have the same geographical and ASN decentralisation (the majority of

620 nodes run out of just ten providers). It also highlighted the significant number of
621 custom-built nodes that are active in Stellar. Furthermore, our analysis showed
622 that Cosmos has a limited number of nodes operated by only 39 small ASNs.

623 In the future, we would like to extend the NodeMaps data scraping pipelines
624 to handle more modern blockchain platforms (e.g., Substrate and Cosmos SDK)
625 and to enable the collection of other types of data (topological/connectivity
626 structures) as it will offer more insight in the decentralisation. Moreover, we
627 would like to introduce a time perspective with P2P node tracking to assess the
628 evolution of blockchain platforms over time and how they respond to various
629 real-life events. Furthermore, we strive to offer in NodeMaps a wide range
630 of off-the-shelf and ready to use metrics for the assessment of non-functional
631 properties for the different blockchain networks.

632 **Acknowledgement**

633 This work was supported in part by the Science Foundation Ireland grants
634 13/RC/2094_P2 (Lero) and 13/RC/2106_P2 (Adapt).

635 **References**

- 636 [1] Nakamoto, S. (n.d.). Bitcoin: A Peer-to-Peer Electronic Cash System.
637 Retrieved from www.bitcoin.org
- 638 [2] Kwon, Y., Liu, J., Kim, M., Song, D., & Kim, Y. (2019). Impossibility of
639 full decentralization in permissionless blockchains. AFT 2019 - Proceedings
640 of the 1st ACM Conference on Advances in Financial Technologies, 110–123.
- 641 [3] Mazières, D. (2015). The Stellar Consensus Protocol: A Federated Model
642 for Internet-level Consensus.
- 643 [4] Ethereum Whitepaper — [ethereum.org](https://ethereum.org/en/whitepaper/). (n.d.). Retrieved August 2, 2021,
644 from <https://ethereum.org/en/whitepaper/>
- 645 [5] Stellar - an open network for money. (n.d.). Retrieved August 1, 2021, from
646 <https://www.stellar.org/?locale=en>

- 647 [6] Kwon, J., & Buchman, E. (n.d.). Cosmos Whitepaper A Network of Distributed
648 Ledgers. Retrieved from <https://cosmos.network/resources/whitepaper>
- 649 [7] No Incentive? Algorand Blockchain Sparks Debate at Cryptography
650 Event - CoinDesk. (n.d.). Retrieved August 2, 2021,
651 from [https://www.coindesk.com/no-incentive-algorand-blockchain-sparks-](https://www.coindesk.com/no-incentive-algorand-blockchain-sparks-debate-cryptography-event)
652 [debate-cryptography-event](https://www.coindesk.com/no-incentive-algorand-blockchain-sparks-debate-cryptography-event)
- 653 [8] Shrivas, M. K., & Yeboah, T. (2018). The Disruptive Blockchain:
654 Types, Platforms and Applications Blockchain powered platform for
655 social enterprises: A case of Entreprise Sociale de Marché Commun
656 (ESMC) View project Banker’s Draft Management and Verification
657 System using Distributed Ledger Technology View project The Dis-
658 ruptive Blockchain: Types, Platforms and Applications. ©TWCS. <https://doi.org/10.21522/TIJAR.2014.SE.19.01.Art003>
- 660 [9] Kyun Kim, S., Ma, Z., Murali, S., Mason, J., Miller, A., &
661 Bailey, M. (2018). Measuring Ethereum Network Peers, 14. <https://doi.org/10.1145/3278532.3278542>
- 663 [10] Gencer, A. E., Basu, S., Eyal, I., Van Renesse, R., & Sirer, E. G. (n.d.).
664 Decentralization in Bitcoin and Ethereum Networks.
- 665 [11] Reddy Venati, Y. (2019). COUNTING NODES IN ETHEREUM BLOCK-
666 CHAIN.
- 667 [12] Kim, S. (2017). Measuring Ethereum’s peer-to-peer network. Retrieved from
668 <https://www.ideals.illinois.edu/handle/2142/99409>
- 669 [13] Dinger, J., & Waldhorst, O. P. (2009). Decentralized Bootstrapping of P2P
670 Systems: A Practical View. Lecture Notes in Computer Science (Including
671 Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in
672 Bioinformatics), 5550 LNCS, 703–715. [https://doi.org/10.1007/978-3-642-](https://doi.org/10.1007/978-3-642-01399-7_55)
673 [01399-7_55](https://doi.org/10.1007/978-3-642-01399-7_55)

- 674 [14] Global Bitcoin Nodes Distribution - Bitnodes. (n.d.). Retrieved August 1,
675 2021, from <https://bitnodes.io/>
- 676 [15] API - Stellarbeat.io. (n.d.). Retrieved August 2, 2021, from [ht-](https://stellarbeat.io/api)
677 [tps://stellarbeat.io/api](https://stellarbeat.io/api)
- 678 [16] The Blockchain Generations — Ledger. (n.d.). Retrieved August 2,
679 2021, from <https://www.ledger.com/academy/blockchain/web-3-the-three->
680 [blockchain-generations](https://www.ledger.com/academy/blockchain/web-3-the-three-)
- 681 [17] From Ethereum to Binance Smart Chain - BSC Explained. (n.d.). Retrieved
682 August 2, 2021, from <https://defiprime.com/binance-smart-chain>
- 683 [18] The problem of scalability in the Bitcoin network — Bit-
684 panda Academy. (n.d.). Retrieved August 2, 2021, from
685 <https://www.bitpanda.com/academy/en/lessons/the-problem-of->
686 [scalability-in-the-bitcoin-network/](https://www.bitpanda.com/academy/en/lessons/the-problem-of-)
- 687 [19] Cong, L. W., He, Z., Li, J., Thank, W., Baldimtsi, F., Bonneau, J., ...
688 Zhang, X. (2021). Decentralized Mining in Centralized Pools. *The Review*
689 *of Financial Studies*, 34, 1191–1235. <https://doi.org/10.1093/rfs/hhaa040>
- 690 [20] Tovanich, N., Soulié, N., Heulot, N., & Isenberg, P. (2021). In-
691 teractive demo: visualization for Bitcoin mining pools analysis. [ht-](https://doi.org/10.1109/ICBC51069.2021.9461124i)
692 [tps://doi.org/10.1109/ICBC51069.2021.9461124i](https://doi.org/10.1109/ICBC51069.2021.9461124i)
- 693 [21] Atluri, R. (2018). A Survey into the Bitcoin Scaling Dilemma.
- 694 [22] rootzoll/raspiblitz: Get your own Bitcoin & Lightning Node running - on
695 a RaspberryPi with a nice LCD. (n.d.). Retrieved August 2, 2021, from
696 <https://github.com/rootzoll/raspiblitz>
- 697 [23] How to Set Up a Bitcoin Node, With Lightning - CoinDesk. (n.d.). Retrieved
698 August 2, 2021, from <https://www.coindesk.com/becoming-self-sovereign->
699 [how-to-set-up-a-bitcoin-node-lightning](https://www.coindesk.com/becoming-self-sovereign-)

- 700 [24] Poon, J., & Dryja, T. (2016). The Bitcoin Lightning Network: Scalable
701 Off-Chain Instant Payments.
- 702 [25] Let's Get Ready to Wumbo! - by Lightning Labs - The Lightning Lab. (n.d.).
703 Retrieved August 3, 2021, from [https://lightninglabs.substack.com/p/lets-](https://lightninglabs.substack.com/p/lets-get-ready-to-wumbo)
704 [get-ready-to-wumbo](https://lightninglabs.substack.com/p/lets-get-ready-to-wumbo)
- 705 [26] Kwon, J. (n.d.). Tendermint: Consensus without Mining.
- 706 [27] Mintscan - COSMOS Validators. (n.d.). Retrieved August 3, 2021, from
707 <https://www.mintscan.io/cosmos/validators>
- 708 [28] Slashing Overview — Cosmos SDK. (n.d.). Retrieved August 3, 2021, from
709 <https://docs.cosmos.network/master/modules/slashing/>
- 710 [29] Validators — Tendermint Core. (n.d.). Retrieved August 3, 2021, from
711 <https://docs.tendermint.com/master/nodes/validators.html>
- 712 [30] Lokhava, M., Losa, G., Mazières, D., Hoare, G., Barry, N., Gafni, E., Jove,
713 J., Malinowsky, R., McCaleb, J., & Mc-Caleb, J. (n.d.). Fast and secure
714 global payments with Stellar. Proceedings of the 27th ACM Symposium on
715 Operating Systems Principles, 17. <https://doi.org/10.1145/>
- 716 [31] Quantifying Decentralization. We must be able to measure blockchain... —
717 by Balaji S. Srinivasan — news.earn.com. (n.d.). Retrieved August 4, 2021,
718 from <https://news.earn.com/quantifying-decentralization-e39db233c28e>
- 719 [32] Crosby, S. A., & Wallach, D. S. (n.d.). An Analysis of BitTor-
720 rent's Two Kademia-Based DHTs. Retrieved August 4, 2021, from
721 <http://azureus.sourceforge.net>
- 722 [33] ayeowch/bitnodes: Bitnodes is currently being developed to estimate the
723 size of the Bitcoin network by finding all the reachable nodes in the network.
724 (n.d.). Retrieved August 4, 2021, from <https://github.com/ayeowch/bitnodes>
- 725 [34] Romiti, M., Victor, F., Moreno-Sanchez, P., Nordholt, P. S., Haslhofer,
726 B., Maffei, M. (2021). Cross-layer deanonymization methods in the

- 727 lightning protocol. In International Conference on Financial Cryptography and Data Security (pp. 187-204). Springer, Berlin, Heidelberg.
728 https://github.com/MatteoRomiti/lightning_study
729
- 730 [35] GeoLite2 Free Geolocation Data — MaxMind Developer Portal. (n.d.).
731 Retrieved August 4, 2021, from [https://dev.maxmind.com/geoip/geolite2-](https://dev.maxmind.com/geoip/geolite2-free-geolocation-data)
732 [free-geolocation-data](https://dev.maxmind.com/geoip/geolite2-free-geolocation-data)
- 733 [36] chainlayer/cosmos-crawler. (n.d.). Retrieved August 4, 2021, from
734 <https://github.com/chainlayer/cosmos-crawler>
- 735 [37] China to shut down over 90% of its Bitcoin mining capacity after
736 local bans - Global Times. (n.d.). Retrieved August 9, 2021, from <https://www.globaltimes.cn/page/202106/1226598.shtml>
737
- 738 [38] Releases bitcoin/bitcoin. (n.d.). Retrieved August 10, 2021, from <https://github.com/bitcoin/bitcoin/releases?after=v0.21.0rc4>
739
- 740 [39] Cosmos Upgrades to Stargate: Another 2017 ICO Very Nearly Com-
741 pletes Its Vision - CoinDesk. (n.d.). Retrieved August 10, 2021,
742 from [https://www.coindesk.com/cosmos-upgrades-to-stargate-another-2017-](https://www.coindesk.com/cosmos-upgrades-to-stargate-another-2017-ico-very-nearly-completes-its-vision)
743 [ico-very-nearly-completes-its-vision](https://www.coindesk.com/cosmos-upgrades-to-stargate-another-2017-ico-very-nearly-completes-its-vision)
- 744 [40] Wood, G. (2016). Polkadot: Vision for a heterogeneous multi-chain frame-
745 work. White Paper, 21, 2327-4662.
- 746 [41] The Blocksize War – Chapter 5 – SegWit. (n.d.). Retrieved July 6, 2022,
747 from <https://blog.bitmex.com/the-blocksize-war-chapter-5-scaling-ii-segwit>
- 748 [42] Smith, G. (2010). PostgreSQL 9.0: High Performance. Packt Publishing
749 Ltd.
- 750 [43] Reagan, R. (2018). Redis Cache. In Web Applications on Azure (pp. 257-300).
751 Apress, Berkeley, CA.
- 752 [44] Abou Jaoude, J., and Saade, R. G. (2019). Blockchain applications–usage
753 in different domains. IEEE Access, 7, 45360-45381.

- 754 [45] Bendeache, M., Saber, T., Muntean, G. M., and Tal, I. (2020). Application
755 of blockchain technology to 5g-enabled vehicular networks: survey and
756 future directions. International Symposium on High Performance Mobile
757 Computing & Wireless Networks for HPC (MCWN 2020). IEEE
- 758 [46] What is the Cosmos SDK? (n.d.). Retrieved June 7, 2022, from
759 <https://crypto.writer.io/p/what-is-the-cosmos-sdk>
- 760 [47] Guide to Cosmos. (n.d.). Retrieved June 7, 2022, from
761 [https://www.coinbase.com/cloud/discover/protocol-guides/guide-to-](https://www.coinbase.com/cloud/discover/protocol-guides/guide-to-cosmos)
762 [cosmos](https://www.coinbase.com/cloud/discover/protocol-guides/guide-to-cosmos)
- 763 [48] Cao, T., Yu, J., Decouchant, J., Luo, X., Verissimo, P. (2020). Exploring
764 the monero peer-to-peer network. In International Conference on Financial
765 Cryptography and Data Security (pp. 578-594). Springer, Cham.
- 766 [49] Alzayat, M., Messias, J., Chandrasekaran, B., Gummadi, K. P., Loiseau,
767 P. (2021). Modeling Coordinated vs. P2P Mining: An Analysis of In-
768 efficiency and Inequality in Proof-of-Work Blockchains. arXiv preprint
769 arXiv:2106.02970.
- 770 [50] Cao, T., Decouchant, J., Yu, J., Esteves-Verissimo, P. (2021). Characterizing
771 the Impact of Network Delay on Bitcoin Mining. In 2021 40th International
772 Symposium on Reliable Distributed Systems (SRDS) (pp. 109-119). IEEE.