GMI-AIMS-3

Challenges in parsing BGP data

Johann Schlamp

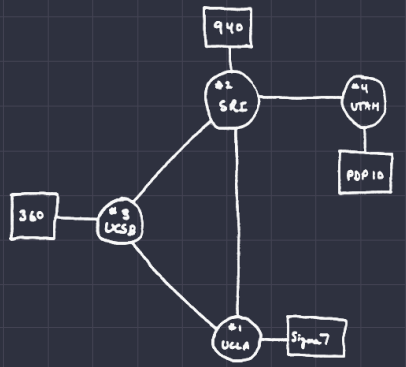

THE ARPA NETWORK, DEC 1969

`0:0:742e:2401:4900::/79` | `AS2936451170`

MOTIVATION

Challenges in parsing BGP data

Selected problems

RFC1997 | RFC2042 | RFC2545 | RFC2858 | RFC2918 | RFC4271 | RFC4360 | RFC4456 | RFC4486
RFC4493 | RFC4724 | RFC4760 | RFC5065 | RFC5291 | RFC5492 | RFC5512 | RFC5543 | RFC5701
RFC6037 | RFC6368 | RFC6396 | RFC6397 | RFC6514 | RFC6608 | RFC6793 | RFC6938 | RFC7311
RFC7313 | RFC7447 | RFC7752 | RFC7911 | RFC8050 | RFC8092 | RFC8093 | RFC8205 | RFC8277
RFC8538 | RFC8654 | RFC8669 | RFC8810 | RFC8950 | RFC9003 | RFC9015 | RFC9026 | RFC9072
RFC9234 | RFC9384

More problems

- **Implementation** pointer arithmetic, IPv6 ll-next-hop, AS_PATH > 255, AS4, ADD-PATH
- **Problem chaining** BGP standard ⇔ BGP speaker ⇔ BGP exporter ⇔ BGP parser
- **Conflicting goals** be conformant with standard ⇔ extract the most information
- **Differing use cases** interactive/bulk, standalone/ecosystem, research/operations

The case for yet another BGP parser (I)

Critical feature requests

- Support for **all** MRT entries/BGP messages and attributes
- Customizable in terms of **selecting record/attribute types**
- **Raw values** and human-readable output (integers vs. strings)
- Native processing of BGP records (+JSON/CSV serialization)

Nice-to-have features

- Transparent support for looking glass text formats (**show bgp** output)
- Rapid prototyping and high-performance modes (namedtuple vs. tuple)
- Built-in statistics and **flexible error handling** (no unexpected aborts)

ftlbgp

- Implemented in Python3 / PyPy3 (fast)
- Zero-Copy operations on all data items (really fast)
- Work in progress: **open-source release** and **paper submission** soon

The case for yet another BGP parser (II)

```
from ftlbgp import BgpParser

with BgpParser(named_records=True, human_readable=True, serialize=False) as parse:

    for record in parse("rib.20240101.0000.bz2"):
        print(record)
```

BgpRouteRecord(*type=, source=, sequence=, timestamp=, peer_protocol=, peer_bgp_id=, peer_as=, peer_ip=, nexthop_protocol=, nexthop_ip=, prefix_protocol=, prefix=, path_id=, aspath=, origin=, communities=, large_communities=, extended_communities=, multi_exit_disc=, atomic_aggregate=, aggregator_protocol=, aggregator_as=, aggregator_ip=, only_to_customer=, originator_id=, cluster_list=, local_pref=, attr_set=, as_pathlimit=, aigp=, attrs_unknown=*)

BgpPeerTableRecord(...)
BgpStateChangeRecord(...)
BgpKeepAliveRecord(...)
BgpRouteRefreshRecord(...)
BgpNotificationRecord(...)
BgpOpenRecord(...)
BgpStatsRecord(...)
BgpErrorRecord(...)

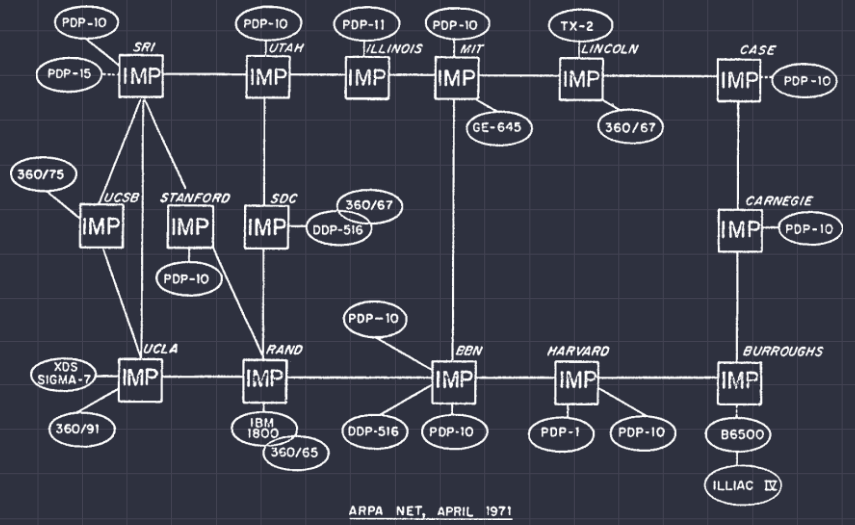

Testing the parser

GENERAL RESULTS

Testing the parser (I)

MRT statistics

- **283,259** input files for **Jan 01, 2024** [Packet Clearing House / RIPE RIS / RouteViews / Leitwert]
- **3,074,186,495** BGP routes extracted from **1,001,812,640** MRT entries

MRT entry	MRT type		MRT entry	MRT type	
BGP4MP	MESSAGE_AS4	610,466,399	TABLE_DUMP_V2	RIB_IPV4_UNICAST	58,273,937
BGP4MP	MESSAGE_AS4_ADD-PATH	9,986,631	TABLE_DUMP_V2	RIB_IPV6_UNICAST	11,361,757
BGP4MP	STATE_CHANGE_AS4	9,814,995	TABLE_DUMP_V2	RIB_IPV4_UNICAST_ADDPATH	167,191
BGP4MP	MESSAGE	3,213,340	TABLE_DUMP_V2	RIB_IPV6_UNICAST_ADDPATH	59,555
BGP4MP	STATE_CHANGE	521,953	TABLE_DUMP_V2	PEER_INDEX_TABLE	76
BGP4MP_ET	MESSAGE_AS4	297,631,200			
BGP4MP_ET	MESSAGE	315,606			

Testing the parser (II)

BGP statistics

- **56,052,068** (5.6%) MRT entries not RFC-compliant
- **16** (!) MRT entries not recoverable

BGP attribute		BGP attribute	
ORIGIN	1,490,360,675	ORIGINATOR_ID	1,123,820
AS_PATH	1,490,360,675	CLUSTER_LIST	1,123,820
COMMUNITIES	1,063,804,129	AS4_PATH	871,562
NEXT_HOP	946,850,166	RESERVED_FOR_DEV	276,517
MP_REACH_NLRI	543,506,876	LOCAL_PREF	89,503
MULTI_EXIT_DISC	383,193,827	AS4_AGGREGATOR	22,190
LARGE_COMMUNITIES	274,134,820	ATTR_SET	3,629
AGGREGATOR	134,804,472	AS_PATHLIMIT	2,809
EXTENDED_COMMUNITIES	107,298,479	CONNECTOR	471
ATOMIC_AGGREGATE	87,433,377	DOMAIN_PATH	357
MP_UNREACH_NLRI	75,182,734	VENDOR_243	2
ONLY_TO_CUSTOMER	5,629,805		

Testing the parser (III)

Collector statistics

- Routes received from **3,835** / **2,755** IPv4/IPv6 peers [1,354 / 940 ASNs]
- Routes containing **16,206** / **3,347** IPv4/IPv6 next-hops [4,184 / 2,324 ASNs]

BGP capabilities

ROUTE_REFRESH	109,199
BGP4MP	82,172
PRESTD_ROUTE_REFRESH	81,555
AS4	77,937
ENHANCED_ROUTE_REFRESH	33,074
GRACEFUL_RESTART	30,932
LLGR	16,764
ADDPATH	11,694
EXTENDED_NEXT_HOP	8,143
FQDN	1,472
BGP4MP_ET	1,469
BGP_ROLE	435
PRESTD_MULTISESSION	1

BGP messages

UPDATE	880,733,106
KEEPALIVE	39,357,481
NOTIFICATION	1,335,554
OPEN	180,046
ROUTE_REFRESH	6,989

NOTE: Random selection of peers

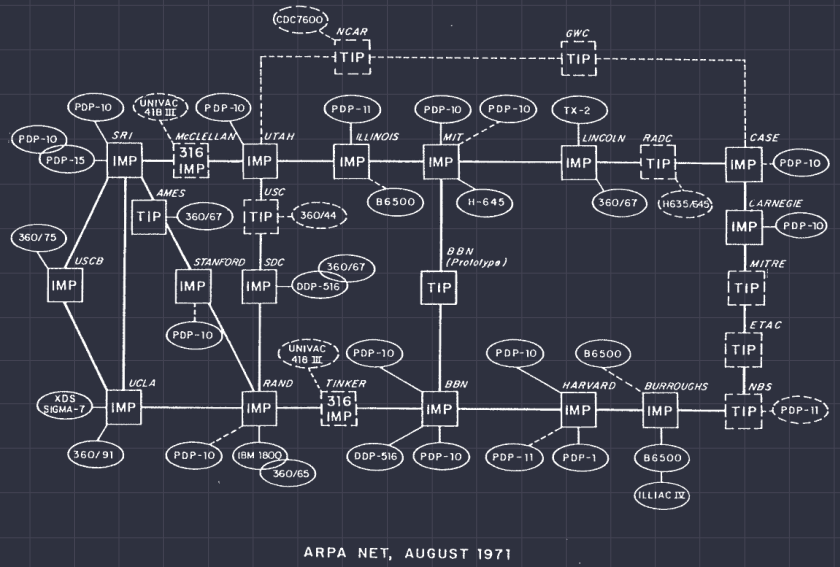

ARPA NET, AUGUST 1971

Comparison with other tools

WAS IT WORTH THE EFFORT?

Comparison with other tools (I)

Our selection of BGP parsers

	Organization	Language	Output
bgpdump	RIPE	C	csv
microbgp	RIPE	C	csv
bgpscanner	Isolario	C	csv
bgpreader	CAIDA	C	csv
FGBGP	bgp.tools	Go	native
BGPKit	bgpkit.com	Rust	csv/json
mrtparse	VMware (?)	Python	native/csv
ftlbgp	Leitwert	Python	native/csv/json

Comparison with other tools (II)

BGP information included in the data set

	RIB entries	Announcements	Withdrawls
bgpdump	693,043,683 [+1]	2,094,588,421 [+10,637,190]	303,962,238 [+6,770,656]
microbgp	693,043,750 [+68]	2,084,950,087 [+998,856]	297,545,819 [+354,237]
bgpscanner	676,562,154 [-16,481,528]	2,094,605,158 [+10,653,927]	305,515,450 [+8,323,868]
bgpreader	692,473,966 [-569,716]	2,175,818,683 [+91,867,452]	427,650,781 [+130,459,199]
FGBGP	693,023,522 [-20,160]	1,353,657,721 [-730,293,510]	183,396,323 [-113,795,259]
BGPKit	693,043,389 [-293]	2,041,801,757 [-42,149,474]	277,792,991 [-19,398,591]
mrtparse	692,892,052 [-151,630]	1,763,545,229 [-320,406,002]	262,661,675 [-34,529,907]
ftlbgp	693,043,682	2,083,951,231	297,191,582

Comparison with other tools (III)

Support for ADD-PATH – extended NLRIs with multiple paths [RFC7911]

	Routes	ADD-PATH	Path Identifiers
bgpdump	**3,091,594,342** [+17,407,847]	**40,158,861** [-10,677,316]	**1,856,496** [-976,935]
microbgp	**3,075,539,656** [+1,353,161]	**37,468,074** [-13,368,103]	**1,788,883** [-1,044,548]
bgpscanner	**3,076,682,762** [+2,496,267]	**40,158,861** [-10,677,316]	**1,856,496** [-976,935]
bgpreader	**3,295,943,430** [+221,756,935]	**0** [-50,836,177]	**0** [-2,833,431]
FGBGP	**2,230,077,566** [-844,108,929]	**0** [-50,836,177]	**0** [-2,833,431]
BGPKit	**3,012,638,137** [-61,548,358]	**40,702,158** [-10,134,019]	**2,703,077** [-130,354]
mrtparse	**2,719,098,956** [-355,087,539]	**46,975,518** [-3,860,659]	**2,728,606** [-104,825]
ftlbgp	**3,074,186,495**	**50,836,177**	**2,833,431**

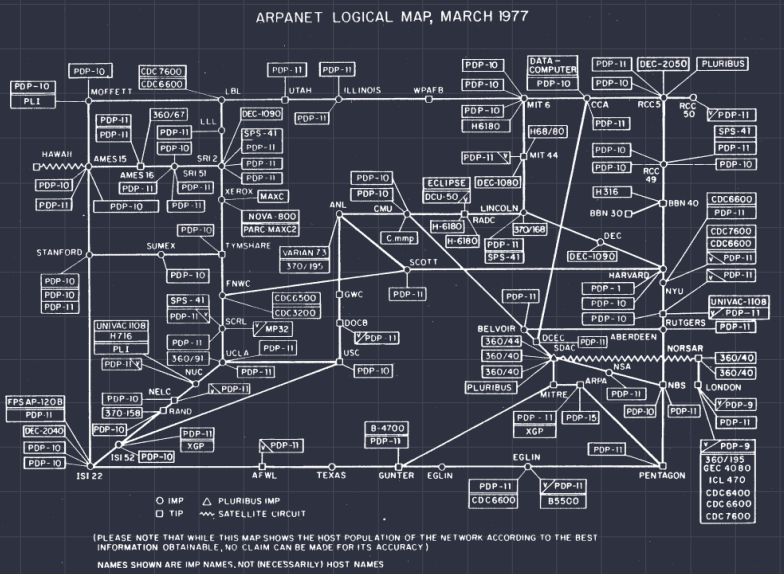

How are ASes interconnected?

TOPOLOGICAL CHARACTERISTICS

How are ASes interconnected?

Analysis of topological characteristics

	AS numbers	AS links	AS triplets	AS paths
bgpdump	83,988 [-2]	674,362 [+2|-241]	10,283,033 [+69|-528]	69,210,069 [+686|-630]
microbgp	83,989 [+1|-2]	676,390 [+1,838|-49]	10,288,006 [+5,603|-1,089]	69,191,395 [+30,776|-49,394]
bgpscanner	83,934 [-56]	668,074 [+2|-6,529]	10,018,224 [+69|-265,337]	66,743,530 [+687|-2,467,170]
bgpreader	83,986 [-4]	674,212 [+2|-391]	10,282,167 [+69|-1,394]	69,200,154 [+686|-10,545]
FGBGP	83,880 [-110]	669,907 [-4,694]	10,204,877 [-78,615]	66,103,452 [+1|-3,106,562]
BGPKit	83,988 [-2]	674,284 [+2|-319]	10,282,397 [+69|-1,164]	69,199,739 [+687|-10,961]
mrtparse	83,987 [-3]	673,113 [+2|-1,490]	10,250,101 [+53|-33,444]	67,616,507 [+445|-1,593,951]
ftlbgp	83,990	674,601	10,283,492	69,210,013

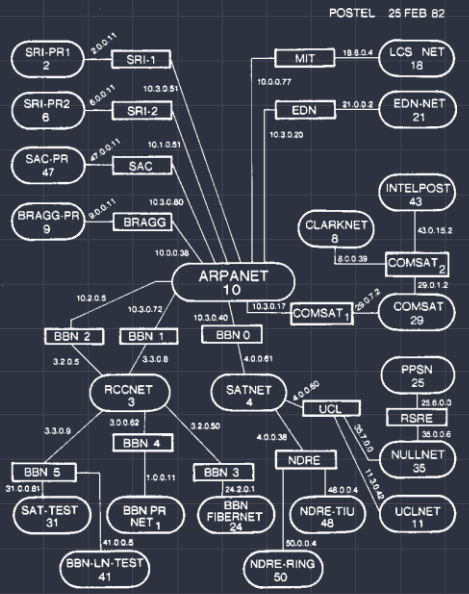

How big is the Internet?

ADDRESS SPACE CHARACTERISTICS

How big is the Internet? (I)

Analysis of address space characteristics

	IPv4 prefixes	IPv6 prefixes	Multi-CC[1]	Large[2]	Invalid[3]
bgpdump	**1,201,133** [+15,381]	**311,101** [+71,752]	**1,357** [+1,270]	**8,092** [+8,085]	**90,063**
microbgp	**1,186,961** [+1,210|-1]	**251,095** [+11,795|-49]	**365** [+278]	**1,376** [+1,369]	**0**
bgpscanner	**1,187,210** [+5,597|-4,139]	**301,199** [+62,924|-1,074]	**1,218** [+1,131]	**3,605** [+3,598]	**148,116**
bgpreader	**1,188,585** [+2,906|-73]	**283,460** [+44,117|-6]	**593** [+506]	**1,076** [+1,069]	**0**
FGBGP	**1,172,953** [+28|-12,827]	**245,723** [+6,459|-85]	**105** [+20|-2]	689 [+682]	**0**
BGPKit	**1,185,695** [-57]	**239,346** [-3]	**87**	**7**	**0**
mrtparse	**1,185,452** [-300]	**239,320** [-29]	**87**	**7**	**0**
ftlbgp	**1,185,752**	**239,349**	**87**	**7**	**0**

[1] *Aggregated prefixes containing RIR netblocks of multiple countries*

[2] *Large prefixes with net mask* ***<8*** *(IPv4) and* ***<16*** *(IPv6)*

[3] *Invalid prefixes with net mask* ***>32*** *(IPv4) and* ***>128*** *(IPv6)*

How big is the Internet? (II)

Analysis of address space characteristics

	IPv4 equivalent (/32)		IPv6 equivalent (/64)	
bgpdump	3,747,340,158	[+684,686,403]	96,487,259,934,925,906	[+95,537,952,283,138,727]
microbgp	3,423,032,903	[+360,379,148]	2,568,829,368,500,655	[+1,619,521,716,713,476]
bgpscanner	3,661,328,689	[+598,674,934]	88,573,694,404,733,341	[+87,624,386,752,946,162]
bgpreader	3,581,595,648	[+518,941,893]	75,739,805,433,388,915	[+74,790,497,781,601,736]
FGBGP	3,175,978,297	[+113,324,542]	22,202,803,563,846,678	[+21,253,495,912,059,499]
BGPKit	3,062,653,742	[-13]	949,307,651,787,179	
mrtparse	3,062,653,499	[-256]	949,303,356,754,341	[-4,295,032,838]
ftlbgp	3,062,653,755		949,307,651,787,179	

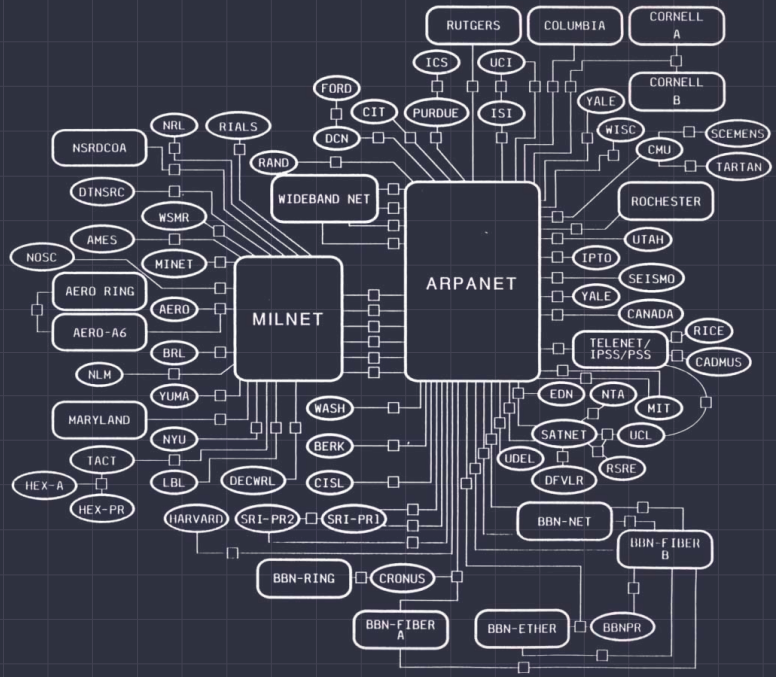

Lessons learned

WORK IN PROGRESS

Summary

Lessons learned

- Raw BGP data requires interpretation and interpolation – we have **dialects** and **artifacts**
- Knowledge of **peer capabilities** can be helpful – but there is no way for direct access
- Adding new features to the BGP/MRT standard can lead to **data loss** (c.f. ADD-PATH)
- Crafting BGP messages with certain attributes may **conceal routes** or even **crash parsers**

Work in progress

- **Global peer database** with explicitly observed and implicitly derived capabilities
- Support **RPKI origin/path validation** results in MRT format in the future?
- Historical data analysis over large timeframe (currently running)
- **Open-source release** and **paper submission** soon

Dr. Johann **SCHLAMP**

schlamp@leitwert.net EMAIL

F958 5A39 FCDC 383E E007
A911 E6CC 7F59 8B24 15A9 PGP

+49 841 93768493 PHONE
+49 174 4944947 MOBILE

Leitwert GmbH ADDRESS
Donaustrasse 17
85049 Ingolstadt
GERMANY

Thank you
QUESTIONS