

Analysis of RouteViews BGP data: policy atoms

Andre Broido and kc claffy

Abstract—

In this paper we introduce a framework for analyzing BGP connectivity, and evaluate a number of new complexity measures for a union of core backbone BGP tables. Sensitive to resource limitations of router memory and CPU cycles, we focus on techniques to estimate redundancy of the merged tables, in particular how many entries are essential for complete and correct routing.

Keywords— BGP Analysis

I. MOTIVATION

Global routing in today’s Internet is negotiated among individually operated sets of networks known as Autonomous Systems (AS). An AS is an entity that

- a) connects one or more networks to the Internet
- b) applies its own policies to exchange of traffic
- c) has a globally recognized and unique identifier

AS policy is used to control routing of traffic from and to certain networks via specific connections. These policies are articulated in router configuration language(s) and implemented by the Border Gateway Protocol (BGP) [Rekhter, Li 1995].

A basic BGP exchange consists of a message regarding reachability of a single network via certain router. The reachability information includes an AS path, i.e. a sequence of ASes through which the sending AS can reach the specified network. In all likelihood the reachability message itself also propagated via this AS path. It is assumed that all ASes in the path forwarded the message deliberately, in accordance with their policies, and *ipso facto* agree to accept traffic destined to the advertized network. As messages about a network traverse the Internet, policies applied for advertisement selection and further announcement accumulate in the system of AS paths.

A BGP table shows the AS path(s) through which each network is considered reachable. This table is an important ingredient of the packet forwarding process in a BGP-enabled router. In addition to the AS path, the table contains metrics associated with the paths, which are used for the best path selection in accordance with AS policies. Some of these metrics are specified locally; others are received from neighbours.

All authors are with CAIDA, San Diego Supercomputer Center, University of California, San Diego. E-mail: {broido, kc}@caida.org.

Support for this work is provided by the Defense Advanced Research Project Agency (DARPA), through its Next Generation Internet program, and by the National Science Foundation (NSF). CAIDA is a collaborative organization supporting cooperative efforts among the commercial, government and research communities aimed at promoting a scalable, robust Internet infrastructure. CAIDA is based at the University of California’s San Diego Supercomputer Center (SDSC). More information is available at the CAIDA website.

Many thanks to David Meyer of U. Oregon, Sean MsCreary of Packets Clearing House, and Brad Huffaker of CAIDA for their help and encouragement.

In addition to packet forwarding, information stored in a BGP table can be used to monitor aspects of the Internet’s architectural evolution, since the data reflects consumption of vital and finite Internet resources [Houston 2001b]:

1. IP addresses (current limit is about 4.3 billion)
2. AS numbers (current limit is 64,412)
3. network identifiers, limited by router memory
4. BGP table size, limited by router memory
5. CPU cycles for processing BGP updates (which can impinge on routers’ forwarding capabilities, if updates are generated too often)
6. Network bandwidth consumed by BGP messages.

The number of networks in the table has bearing on both router memory and CPU cycles. Routing flaps, i.e., advertisements and withdrawals of prefixes, tends to increase with the number of prefixes in the table [Houston 2001b] [Doran 2001]. Reduction of prefix count is typically seen as beneficial to infrastructural integrity, and finding mechanisms to do so is thus important architectural research.

BGP inter-domain routing table data enables two different aggregations of IP address space:

- 1) from IP address (32-bit integers) to longest matching IPv4 network prefix, and
- 2) from network prefix to AS originating that prefix into the global routing mesh.

An origin AS is an AS that appears at the end of an AS path for a prefix. It is thought to be the AS that originally advertized that prefix. In today’s interdomain routing tables, almost 99% of prefixes consistently have a single origin AS.

The mappings above are useful for converting IP addresses to ‘home’ prefixes and AS numbers. This usage of BGP data is helpful for tasks such as visualization of Internet AS topology [HBCFKLM 2000] or for analysis of trends in routability of IP space given by registries to Internet service providers (ISPs) and customers [BC 2001a].

Analyzing BGP tables can support studies of many questions regarding Internet properties discussed in the Internet community (e.g., NANOG, IETF):

1. How much IP address space is routable?
2. Are allocated IP addresses actually being routed?
3. Which ASes are the most important?
4. How many ASes are single-/multihomed?
5. How many prefixes do ASes advertize?
6. Does the core table grow due to multihoming?
7. Are more-specific prefixes driving up the table size?
8. Are IP addresses within an AS topologically connected?
9. What is the IP hop diameter of an AS?
10. Can Internet topology, especially particularly highly connected, central, or vulnerable points, be inferred from BGP tables?

Generally speaking, a BGP table is useful for answering

many types of questions about the Internet, when other data sets are prohibitive in size, diversity, or aggregatability. BGP data also has an advantage that its two most basic concepts, *prefix* and *AS* are close approximations to real-life notions of *network* and *administrative domain*.

AS interconnections given by BGP AS paths represent contractual relations between autonomous systems, from which one can infer business models and relations between ISPs [L.Gao 2000], as well as assess statistics of inter-AS connectivity (peering sessions) [Faloutsos 1999]. These uses are theoretical in nature since the presence of an AS in a path does not guarantee that the traffic will be actually carried by this AS; BGP connectivity is declared rather than observed. Nonetheless, such analyses still provide considerable, and in some cases otherwise unavailable, insight into the global routing mesh, especially when taken in conjunction with other types of data.

Original specification of BGP [Rekhter, Li 1995] states that

Information reduction may imply a reduction in granularity of policy control - after information is collapsed, the same policies will apply to all destinations and paths in the equivalence class.

In this paper we develop this idea and introduce a framework for analyzing complexity of global routing policies. We define and evaluate an important new complexity measure – the number of *policy atoms* – for a union of core backbone BGP tables. With resource limitations of router memory and CPU cycles in mind, we focus on techniques to estimate redundancy of the merged tables, mostly how many entries are essential for complete and correct routing. These complexity measures can answer questions such as:

1. How many prefixes does it take to cover the whole routable address space?
2. What is the complexity of the system of AS paths associated with individual prefix?
3. How many different types of networks are globally distinguishable with respect to routing policies?
4. How many routing policies are applied by the Internet to addresses originated by one AS?

In addition to BGP table size reduction, another application for finding aggregate units of routing is the design of Internet active measurement systems. A system that covers many BGP prefixes is more likely to provide relevant information, e.g. about network topology, than a system that covers just a few. However, measurement resource limitations (bandwidth and CPU) require coarser probing frequency. A system that covers provably distinct routes could help eliminate redundancy in path probing, thus maximizing coverage and sampling frequency at the same time.

The most significant contribution to the analysis of BGP tables was done by Geoff Houston [Houston 2001a,b], with whom we share many motivations. Our results differ in two ways: we analyze a larger set of tables (up to 33), not just one; and we take only prefixes common to backbone ASes, which avoids the bias of locally maintained routes. We thus capture a more complete inventory of distinct units of

routing policy. The extra resolution is moderate in numerical terms (a factor of about $\sqrt{2}$) but significant in a policy context.

A. Roadmap of paper

In section 2 we describe the BGP data we use in this paper. Section 3 deals with important basics of BGP table syntax and with analysis of more specifics. Section 4 introduces the notion of *policy atoms* and describes features of atoms we are currently investigating. Section 5 studies variations of path selection and announcement policy across an AS with the notion of *ramified atoms*. Section 6 summarizes insights we have gained and discusses current and future directions.

II. BGP DATA AVAILABILITY.

There are several publicly available sources of raw and processed BGP data. Summary plots of several interesting metrics of Australian ISP Telstra’s BGP routing table are updated daily [Houston 2001a].

Samples of BGP data from diverse sources are available via a variety of *looking glasses* [CAIDA 2001], a globally distributed, independent set of servers that support examination of BGP AS paths for individual IP addresses. Looking glasses are primarily intended as debugging tool and cannot provide a full BGP table upon request.

RIPE NCC (Reseaux IP Europeens) recently started collecting 75 or so BGP route views, mostly from European ISPs, including those accredited at LINX (London Internet Exchange). We intend to parse RIPE’s BGP tables in the near future.

Another source of data on inter-AS relations, though not on AS paths, is the *whois* service provided by RIPE, APNIC and some other registries [IRR 2001]. These registries store AS traffic policies in succinct database entries: e.g.

```
from AS701 100 accept ANY AND NOT 0.0.0.0/0
from AS8399 105 accept AS8399 AND NOT 0.0.0.0/0
to AS8399 announce AS-IAXIS
```

This data allow analysis of policies as relations between triples of AS: the third AS is announced by the second to the first with an intent to invite the first to use the second AS for transit to the third AS.¹

Yet another relevant public collection of data is the records of reservations (to a country), allocations (transfers to ISPs) and assignments (transfers to customers) of address blocks by three authoritative Internet registries, ARIN, RIPE and APNIC [ARIN 2001]. We have analyzed this data in another study [BC 2001a].

A. University of Oregon RouteViews data

The data used for this paper comes from a public data source based at the University of Oregon’s Advanced Network Technology Center [Meyer 2001a]. RouteViews data

¹We analyze this data in a companion paper where we introduce and study the *constrained dual AS graph*, a policy-oriented representation of inter-AS connectivity in the form of a graph where nodes are ordered pairs of ASes and nodes AB and BC are linked when B was observed announcing C’s prefixes to A.

is a union of several dozen unpruned backbone BGP tables [Meyer 2001b]. Note that there are other mechanisms to obtain this and similar data, including the RIPE registry database project described above.

Participating peer ASes often share more than one routing table view. It is an unsolved question which selection of views yields the best picture of global routing. The question is important since tables are 10M gzipped (250M uncompressed) as of May 2001, rapidly consuming disk space if one needs more than one snapshot a day. One of our goals is to find metrics that can help with peer selection when one cannot handle processing data from all peers, such as the marginal utility of adding another peer after N peers have already been merged.

RouteViews daily sampled tables stored at Hans-Werner Braun's server `moat.nlanr.net` [NLANR 2001] start at 1997-11-08 and end in March 2001. PCH [PCH 2001] began storing daily snapshots on 2001-02-19, and RouteViews itself started storing samples every 2 hours from 2001-04-20 [Meyer 2001b].

For this paper, we processed RouteViews BGP tables 1999-04-30, 2000-05-02, 2001-05-01 and 2001-05-03. The dates differ slightly so as to use the largest table within a few days of May 01 of respective year.

RouteViews peer participation has been on the rise since inception. All except a few peers contribute a full default-free backbone routing table. Remaining peers contribute a table less than half of full size, typically just a few thousand prefixes.

The following table shows the number of participating RouteViews peers and 'large' peers (with cutoff at 90% of maximum prefix count)

Year	1999	2000	2001
All peers	21	23	35
Large peers	15	17	27

The major difficulty in analyzing trends in RouteViews data is that growth of individual tables is accompanied by growth of the number of contributing peers. Occasional drops and changes in peer addresses further contribute to making it almost impossible to find a representative collection of peers with large enough table present for long enough to do trend analysis. For example, there are only six large peers in common in the three tables of 1999-2001 presented here. Furthermore, diversity measures that depend upon the number of peers (e.g. 'atoms' discussed below) are influenced by the number and choice of peer tables, masking the trends one intends to analyze. We thus devote section III.B to the selection of peer tables for analysis.

III. BGP TABLES: PROPERTIES AND CAVEATS

An entry in a BGP table looks like

Network	Next Hop	Mc	Path
12.0.0.0	204.29.239.1		6066 3549 7018
12.0.48.0/20	204.29.239.1		6066 3549 209 1742
	213.200.87.254	20	3257 13646 1742

First entry is the network to be reached. Second is the peer who contributed the view. Third is the metric (which is usually Multi-Exit Discriminator, see [HM 2000]) and the

fourth is AS path. (We did not show four other parameters which have missing, 0 or almost constant values, and we will not analyze metric either.)

The table is ordered by network prefix. Repeated prefixes are not shown. Networks which align on classful boundaries (/8 in Class A, /16 in Class B and /24 in Class C space) are shown without prefix mask length.

In the example above, the address block 12.0.48.0/20 is originated by AS 1742. Two peers advertize reachability to it via two AS paths. One of them contains 4 ASes (3 hops) and another 3 ASes (2 hops).

Note that this block is a subset (or *more specific*) of address block 12.0.0.0/8; nonetheless, their origin ASes are different and AS paths diverge a few hops away from a common peer. We will discuss more specifics later in this section.

A. AS path length

The AS path length is the fundamental metric of a BGP routing table. Internet providers generally over-provision capacity for bandwidth and packet processing in their networks [Atkinson 2001], and generally indicate that packet loss is most likely to occur on AS boundaries, either when passing packets to the customer or to another provider. If packet loss events on the boundaries are independent, the probability of receiving a packet through an AS path is the product of probabilities of crossing each AS boundary without packet loss. In the absence of reasonable estimates, one can assume probability of such loss-free crossing to be a constant p . The probability of getting a packet through an AS path with n AS crossings would then be equal to p^n . Since $p < 1$, this expression decreases as n grows. It therefore makes sense to minimize AS path length n in order to get more packets delivered.

And in fact, BGP does. Although it is possible to override this metric on the basis of vendor-specific weights [Cisco 2001] or local preference [Stewart1999], when several paths with equal local preference exist for the same prefix, BGP selects the shortest AS path. This algorithm ostensibly minimizes packet loss properties and also theoretically minimizes the hassles involved when something goes wrong between two arbitrary ASes.

The fundamental nature of the AS path length metric in the BGP decision algorithm gives rise to a common practice of prepending extra copies of an AS number to the beginning of the path, to reduce the likelihood of selection of that (now longer) AS path for forwarding traffic. Over 10% of the AS paths in the RouteView table currently exhibit prepending. Prepending can even be explicitly articulated as policy, e.g., in a line from RIPE's whois database:

```
Community Definition
```

```
NNNN:3062 To LINX PEERS prepend NNNN NNNN
```

where NNNN is an AS number.² The practice of prepending results in mistyped lines in most RouteViews tables. Two apparent AS loops have been present in Route-

²A *community* is an BGP attribute associated with a set of network prefixes.

Views for about 2 months. This type of noise adds fictitious connectivity to AS graphs constructed from BGP tables.

After path selection occurs, repeated AS names are no longer relevant; we remove redundant prepended instances of an AS in our table before analysis. As far as we can tell, existing analyses of BGP data do not show any sign of awareness of these and some other idiosyncrasies in the RouteViews data.

Selection of shortest AS path is not a policy; it is part of the BGP standard. But there are several mechanisms that can override this standard; we will assume these *are* policies:

1. prepending (own) AS number more than once;
2. use of BGP attributes that take precedence over AS path length
3. traffic engineering tools that update the output of BGP selection process [Meyer 2001d];

We can study the use of these policies and their ramifications by analyzing BGP AS paths.

B. Peer selection among RouteViews

We next describe RouteViews BGP data. Our aim is to find the set of prefixes generally agreed as routable. This set will serve as an input to path identification algorithms, which analyze routing diversity of the globally visible systems of prefixes. To evaluate this diversity, we need the maximum number of complete peer views. However, as the number of views increases, the number of prefixes shared among them decreases, complicating analysis.

To make peer contributions comparable and to avoid influence by different number of entries in the tables, we will in most cases drop all peers contributing less than 90% of the maximum count of prefixes with lengths 0-24 bits. Prefixes longer than /24 are typically not globally routable and should be filtered at AS boundaries. A large number of local (customer) prefixes can upset the balance of arithmetic relations among prefixes in the table. As an example, a single AS in the Oregon data was observed as carrying 12K to 13K prefixes longer than /24 through March-April 2001. These prefixes were not seen by outside peers.

We need to be careful with the exact value of the table size cutoff since an odd peer or two can have many prefixes that no one else carries which will artificially raise the threshold. Even within 90% threshold, the sequence of peer table sizes often has large (a few thousand prefixes) gaps, and it is not always clear which peers to omit. It hinges on the estimates of global object counts, which are sensitive to the number of prefixes shared by all peers.

The counts of prefixes of length /24 or shorter in the May 03, 2001 data are shown below:

#pref,K	105	103	101	100	99	98	92	85	1-8
#peers	3	2	1	6	13	2	2	4	4

There are three types of backbone tables (and policies of prefix selection), resulting in 98-105K, 92K and 85K entry tables. It turns out that the 85K table is obtained by filtering on prefix length as described in [Houston 2001a]. The most common size table is 99K. An 98K threshold leads to

the choice of 27 peer tables. An 85K prefix cutoff leaves 33 peers.

The number of prefixes shared by different number of peers in May 01, 2001 data is highly variable:

Peers	1	2	3	4	5	6	7
prefixes	6675	1897	4154	2703	855	421	1498

(the number of prefixes shared by 8-26, or 36 peers is 1592; for 37 peers (maximum possible) it is 0.)

Peers	27	28	29	30-32	33	34	35
prefixes	6166	305	6368	1926	72507	7659	3653

6166 prefixes are shared by exactly 27 peers (and not shared by peers with table sizes of 92K and less) correspond to the table size drop from 98K to 92K. These prefixes consist mostly of long ranges of prefixes in /22-/24 range that closely follow each other. (In fact, 3028 of them are immediately adjacent to the previous prefix from the same set of 6166.) Many of these ranges are also announced by the same AS. After cutting the list of peers at 98K prefixes, the counts of shared prefixes look as follows:

Peers	1	2	3	4	5	6	7
prefixes	6454	1159	4821	2332	196	576	1496
Peers	21	22	23	24	25	26	27
prefixes	37	22	7	386	536	760	97250

In other words, almost all prefixes (97.2K out of possible 98.8K, or 98.4%) are shared by all peers, which feels right: there should be, after all, a set of generally agreed routable prefixes, even if some of them are redundant.³

C. The challenge of more specifics

Between 1998 and 2000 the size of the routing table expanded so fast that it threatened to outgrow router memory and require a major overhaul of Internet infrastructure. By 2001 the growth abated somewhat, but traffic-engineering-based sources of routing table growth, particularly multihoming and load balancing, remain strong [Houston 2001b].

In a typical scenario, a multihomed customer connects to the Internet through more than one provider. To maximize the value of this connectivity, the customer advertizes its addresses on both connections. These addresses are taken from the first provider's block and cannot be aggregated (merged) with the blocks of the second provider. The second provider must announce a smaller block. However, the first provider is compelled to advertize the smaller block as well, since otherwise packets will be attracted to the more specific announcement via the second provider, defeating the traffic engineering objectives. This condition results in an extra globally visible block where the first provider's aggregate (larger block) would normally suffice.

In one type of load balancing scenario, a customer advertizes half of addresses on one link and half on the other. The range advertized on first provider's link can be aggregated. However, if the first provider's link goes down, half of the customer's space becomes unreachable. To increase reliability, the customer announces both the whole and half block on the second connection. To avoid starving the first

³Note, however, that the maximum carried by a peer is much higher: 105.5K on 03 May 2001.

connection, the customer convinces the first provider to announce at least the whole block. Hence, two to three extra globally visible blocks are likely to appear in the table.

This can even leave a trace in documented policy of the first provider, like the following excerpt from RIPE whois database:

```
NNNN:1000 - leak this route
(apply to routes from NNNNN CIDR blocks)
```

These *multi-homing* conditions imply the presence of simultaneous multiple connections from a network to the Internet, and predict an abundance of more specifics, i.e. address blocks that are parts of larger aggregates (provider blocks.) In the two years covered by the 3 tables we analyze, the trend is clear: singly homed ASes moved from first to second place, and triply and more highly homed AS expanded their share from 25% to 30% of all multihomed ASes.

Year	1999	2000	2001
AS count	4893	7482	10911
Indeg.1	48.2%	41.4%	39.9%
Indeg.2	38.2%	39.8%	42.7%
Indeg.>2	13.5%	18.8%	18.4%

Let us examine how many more specifics exist in the table. The table of May 03, 2001 contains 118379 prefixes in the union of 37 peer views. Of those, 102095 are common to six or more peers. (Each AS contributes to RoveView tables at most two views, except one who contributes 3; this guarantees at least 3 ASes carrying a prefix.)

Among those prefixes, more specifics make up 55210, or 54% of all common prefixes. Another 41731 prefixes are non-specifics, in that their address space is covered by no other prefix. 5154 prefixes are *least specific*; they are not more specifics of any other prefix, but themselves have more specific prefixes present in the table. We will call these least specific prefixes *roots*. The set of more specific prefixes of a root has the structure of a tree, which is induced by the similar structure of CIDR subdivisions of the whole IPv4 address space. The distribution of tree height (height one corresponds to a tree with only one level of hierarchy, e.g. a prefix and a single more specific) is as follows:

Tree height	1	2	3	4	5
No.roots	4295	721	125	12	1

The number of nodes in the trees (excluding the root node) varies from 1 (1611 trees) and 2 (954 trees) to large unique trees of 1132 and 1243 nodes (with tree height 2) and 1324 nodes (tree height 4). We conclude that there are prefixes with up to 5 less specifics, a fact rarely mentioned in the literature, even though it is easy to observe (we saw it first in May 2000; it occurs in the table of 2001-04-03, but not of 2001-04-20 and 2001-04-29, where the maximum height is 4.)

The distribution of prefixes by specificity level for 03 May 2001 data is given in the table below. The count for level 0 includes roots and non-specifics; the count for level 1 reflects immediate more-specifics of the roots; the count for level 2 includes more-specifics of the latter, and so on.

Level	0	1	2	3	4	5
prefixes	46885	42758	11366	1013	64	9

The address space consumption by prefixes of different specificity level is given in the following table:

Level	0	1	2	3	4	5
IPs	1122M	106M	9.44M	0.67M	24200	2304

Each level of specificity takes at least one decimal order of magnitude less address space than the previous level. In particular, most IP addresses are covered by only one prefix.⁴ The extended version of this report describes further implications of specificity [BC 2001b]

According to the BGP specification [Rekhter, Li 1995]

The set of destinations described by the overlap (of two prefixes) represents a portion of the less specific route that is feasible, but is not currently in use. If a more specific route is later withdrawn, the set of destinations described by the overlap will still be reachable using the less specific route.

...If a BGP speaker accepts the less specific route while rejecting the more specific route from the same peer, then the destinations represented by the overlap may not forward along the ASs listed in the AS_PATH attribute of that route. Therefore, a BGP speaker has the following choices:

- Install both the less and the more specific routes*
- Install the more specific route only*
- Install the non-overlapping part of the less specific route only (that implies de-aggregation)*
- Aggregate the two routes and install the aggregated route*
- Install the less specific route only*
- Install neither route*

If a BGP speaker chooses e), then it should add ATOMIC_AGGREGATE attribute to the route. A route that carries ATOMIC_AGGREGATE attribute can not be de-aggregated. That is, the NLRI of this route can not be made more specific. Forwarding along such a route does not guarantee that IP packets will actually traverse only ASs listed in the AS_PATH attribute of the route. If a BGP speaker chooses a), it must not advertise the more general route without the more specific route.

The upshot of this reasoning is that ⁵

- less specific prefix is responsible for forwarding all traffic to more specific;
- to avoid forwarding responsibility, less specific can be de-aggregated, i.e. split into complementary blocks. ⁶
- BGP requirement guarantees propagation for more-specifics as they "ride on the backs" of less-specifics;

This means that if routing policies were to be ignored, half of the table (47K roots and non-specifics) would suffice for reaching all of the Internet. Since top prefixes are non-

⁴Some addresses have 6-fold coverage, which is probably as safe as buying five insurance policies for the same type of accident.

⁵Note again that BGP allows AS path not to be the path followed by packets.

⁶The smallest number of blocks is then is $b_m - b_l$, where b_l and b_m are bit lengths of prefixes. It may be possible that stretches of closely spaced splinter blocks are caused not by multi-homing per se, but deaggregation into $2^{b_m - b_l} - 1$ blocks, e.g. complementing a /24 in a /16 by 255 other /24s. This can indeed cause exponential table growth.

overlapping, match between IP and top prefix is always unique, which could also simplify router design.

D. Other caveats

There are several other caveats in dealing with analysis of BGP tables, the most prominent being the wide fluctuations of more specific prefixes from peer to peer and from day to day. In the extended version of this report we cover other caveats as well: AS loops, tangles, private ASes, AS sets, and multiple origin ASes [BC 2001b].

IV. BGP ATOMS AND CLASSIFICATION OF ROUTING POLICIES.

In the previous sections, we have shown how to reduce the number of routing table entries using root prefixes and non-specifics. We next explore a new way of grouping addresses on the basis of their global routing properties, from a single AS perspective.

We group prefixes (address blocks) according to the AS path to which they map in the BGP table. The number of such groups is then equal to the number of different AS paths. All path counts refer to reduced (non-prepended) paths with no repeated adjacent ASes. Geoff Houston's report [Houston2001a] also includes this grouping, finding 14081 AS paths to 108K prefixes on 2 May 2001.

We seek a more effective way to group prefixes, which would reflect properties of the global routing system rather than just a single AS. Otherwise we risk policy biases inherent in any one specific view, e.g., a strong preference for some AS paths or an excessive number of prefixes. Moreover, a single AS may not see a policy difference in the routing of two prefixes, simply because they are routed through the same path from it – the projection of data to the ‘observation plane’ of this AS loses some fine detail.

We offer a generalization of prefix grouping by AS path:

DEFINITION. Two prefixes are said to be *path equivalent* if we cannot find a BGP peer who sees them with different AS paths. An equivalence class of this relation is called a *BGP atom*.

Convenience associated with description of global routing in terms of atoms derives from the fact that an atom captures the part of routing policy relevant to AS paths and applicable to many prefixes at once. Intuitively, consider an atom as a double-sided coin: with a system of AS paths on its head and set of prefixes on its tail.

Algorithmically, we construct BGP atoms as follows:

1. Find all prefixes common to a chosen system of peers.
2. Associate a system of peer AS paths with each prefix.
3. For each system of AS paths, find all prefixes that share this system of paths.

The following table shows relevant statistics for the three May tables from 1999-2001 (using May 01 table for 2001) choosing the 6 peers present throughout this period: AS 1, 7018, 3561, 2828 (US) 1755, 3333 (Europe).

Year	1999	2000	2001
AS count	4893	7482	10832
common prefixes	57720	75174	99009
AS paths, max	6603	9859	14207
atoms	8615	12327	17474
atoms/paths	1.30	1.25	1.23
atoms with 1 prefix	3912	5814	8582
largest atom, pf.	1152	1799	2290
crown atoms	4697	6684	9465

The AS and prefix counts in the table differ slightly from those given elsewhere in the paper since we used only 6 peers. Prefix counts refer to the number of prefixes common to all peers, which were used to compute atoms.

The table shows that the growth in the number of atoms seen by 6 peers closely follows the growth in the number of ASes, prefixes and AS paths. All four numbers approximately doubled from 1999 to 2001.

The table suggests strong potential for atomic reduction of routing tables. Atoms generalize the grouping of prefixes by AS paths; yet as time goes by, their counts approach the AS path count, at least the maximum values (which for this peer selection are consistently from AS 3333 RIPE.) One may even hope that the diversity of AS paths seen by one peer will eventually become close to what can be found with global analysis including dozens of them. For the time being, however, this is not yet so, as we will see when analyzing tables with 27 and 33 peers.

About 50% of atoms contain just one prefix, though there are also atoms with many prefixes, which is where reduction of the BGP table occurs. The number of one-prefix atoms and the maximum prefix count for an atom have doubled in 1999-2001, matching the evolution of the table size. The cumulative distribution of atoms by size (figure 1) is close to a Weibull curve [Extreme 2000]

$$P\{n > x\} = \exp(-ax^b)$$

where n is the number of prefixes in an atom, with $b \approx 0.15$ for $x \leq 100$ prefixes. Frequencies of individual counts satisfy the relation

$$P[n = x] \leq cx^d, \quad x \leq 100$$

with $d \approx 1.8$. It is also possible to approximate $P\{n \geq x\}$ with power function $cx^{-1.3}$ [Faloutsos 1999].

With regard to address space content, sizes under /16 accumulate about 5%. Atoms of total address range of a /16 (65.5K) contribute about 10% of the address space, and a comparable number of atoms. Atoms larger than /16, but smaller than /8, accumulate address space in a logarithmically uniform way, i.e. each binary order of magnitude size contributes approximately an equal number of addresses, for a total of 45% of routable address space. Yet another 40% of the routable address space is contributed by atoms containing large blocks (/8s, i.e. 16.8M, and more.)

Fig.2 also shows that most atoms are small in terms of address size: more than 60 and about 3/4 have sizes under 65K. The only size in which both many atoms exist and significant portion of address space is covered is /16.

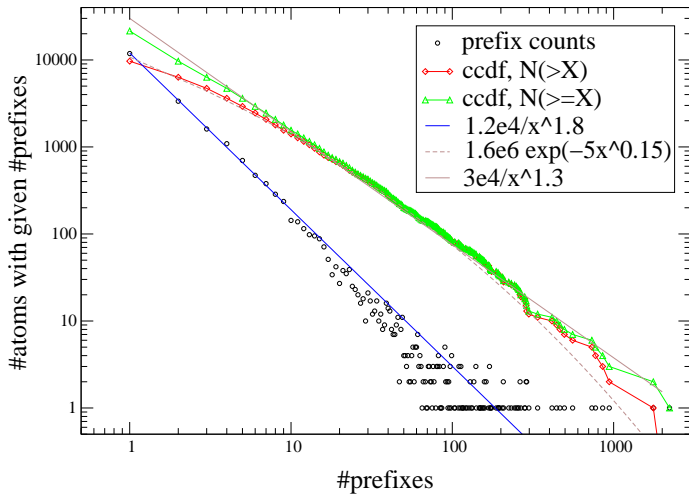


Fig. 1.

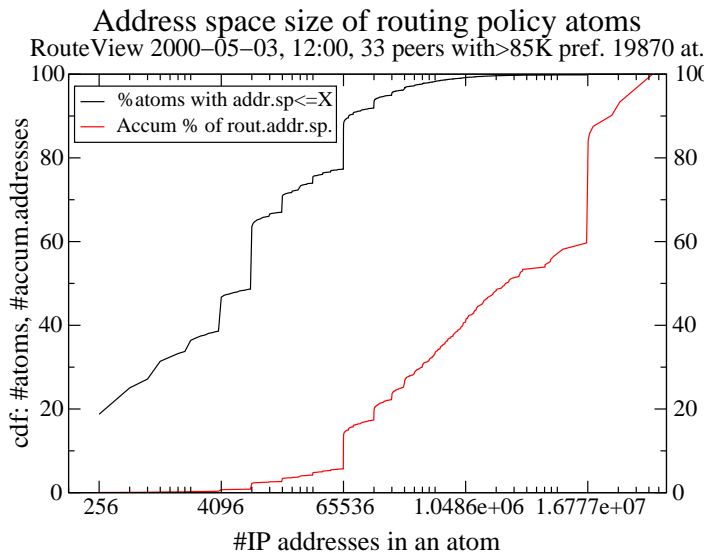


Fig. 2.

Note also that both distributions show bends after spikes at powers of 2, where atoms with one large prefix and one or more smaller blocks are accumulated.

A. IP renumbering and BGP table size reduction

BGP standard encourages carrying multiple prefixes in an UPDATE message [Rekhter, Li 1995]:

The BGP protocol allows for multiple address prefixes with the same AS path and next-hop gateway to be specified in one message. Making use of this capability is highly recommended.

BGP is therefore able to carry atoms in its messages at the connection setup when the whole table is exchanged. The standard size of TCP data segment, 1460 bytes, will have space for at least 300 prefixes, that is, for every atom out of 21.6K except 13 (99% of atoms have 50 or less prefixes.)

This can be a substantial savings on processing overhead

and on network bandwidth. If the atoms are agreed upon in advance, – e.g. through the use of globally defined communities or a new transitive attribute, the savings could be leveraged throughout the whole Internet without the need to collect prefixes with equal AS paths each time updates are sent, as suggested in BGP specification. (Assembling updates could even be done in parallel with an ASIC or content-addressable memory.)

However, as prefixes change status, individual advertisements are required and the savings in BGP transmission may be harder to realize at that stage.

Recall though that BGP is designed for advertizing aggregates which may be only partly reachable, so as to reduce route flapping. Atoms can be aggregated in several possible ways, the simplest (although not the easiest) being the renumbering of IP addresses.

The table of 03 May 2001 contains 27 peers who carry over 90% of the maximum prefix count. The number of atoms for this system of peers is 21570, 22.2% of the common prefix count.

Each atom contains AS paths to one or more prefixes and in most cases the atom has a unique origin AS. In other words, almost every atom, no matter how many prefixes it contains, has some AS responsible for it. It appears that providers could renumber and to reduce the set of address blocks to the set of existing global routing policies. Assigning equi-routable devices IP addresses from the same CIDR block, could then reduce the number of globally visible prefixes to 1994 levels [Houston 2001b].

A possible side effect of such renumbering could be an increase in the required address space. The granularity of CIDR blocks is coarse; supporting 1025 devices in one block requires a block of size 2048, with half the space wasted. Since one half of IP space is already either allocated to ISPs or assigned to customers [BC 2001a], this policy potentially could exhaust the IPv4 address supply.

Taking into account the distribution of top-level prefixes and their sizes among atoms, actual increase is smaller than that worst case scenario. Both for the 27 peer selection of May 03, 2001, and for 33 peer selection, the increase totals 39.6% of address space which is currently routable, or 10% of all IP address space (2^{32} addresses.)

If assignment is done in a more careful way, so that the largest prefix in an atom is not renumbered and remaining space is placed into one CIDR block, the increase is about 17% of current space and 4.4% of available IP space.

The situation is complicated however by the fact that some atoms contain more specifics, and some atoms contain nothing else but more specifics. For 33 peers (i.e. for filtered prefix set) the number of former is 9579 out of 19870, and the number of latter is 5879. The prefixes in an atom with more specifics only can have their less specifics scattered over up to 7 other atoms, although in 90% of cases they are concentrated in just one atom.

#at.w.less sp.	1	2	3	4	5	6	7	Total
more sp.only	5269	528	66	10	3	2	1	5879

For 27 peers and corresponding unfiltered prefix set the numbers are of course much higher. There are 7489 atoms

containing more specifics only, out of 21573. The fraction of those which have all their less specifics in one atom is 86.4%.

#at.w.less sp.	1	2	3	4	5	6	14	Total
more sp.only	6557	808	96	20	4	3	1	7489

Prefixes in the atoms whose less specifics are concentrated in one atom can be renumbered in agreement with their less specifics, whereas prefixes in other atoms may run into problems. This matter will be further investigated in the full version of the report.

To conserve address space, we can use arbitrary intervals in IP space as address blocks. In order to support arbitrary block sizes, we must impose stringent control of allocations to guarantee that the routing algorithms find forwarding entries in constant (or logarithmic in the size of the routing table) time.

To preserve the current general setup, we could use a quasi-CIDR allocation approach, selecting block sizes from the sequence of numbers that make up an addition chain [Knuth 1997], and subdividing IP space from top to bottom into a binary tree using these blocks. The sequence can be refined at certain lengths to provide efficient and convenient use of address space, while also enabling fast address-to-block mapping in routers. In particular, one can use the sequence of lagged Fibonacci numbers, thus reducing the worst case scenario for possible address space use to $q - 1$, where q is the unique positive root of $X^k - X^{k-1} - 1 = 0$. For example, $k = 3$ requires in the worst case 47% extra space, instead of CIDR's 100% extra. It may be reduced further, depending on lag k . The imbalance in the tree can be exploited to make some addresses accessible faster than others. The extended version of this report has further analysis of this approach [BC 2001b].

B. Crown atoms

We next define crown atoms in pursuit of further potential reduction in the number of routable objects.

Despite recent growth in the prevalence of multihoming, a significant number of ASes are still singly homed, i.e., with only one inbound link, at least as observed in RouteViews data. Section III provided statistics of multihomed ASes.

On the other hand, from the global routing standpoint, if all paths to AS B pass through AS A, it would make no difference whether B's networks were advertised by A. This is one reason for the existence of multi-origin prefixes. One necessary condition for truncation of paths at B is that it occur in paths consistently always before A and other ASes closer to A than B (this is not always the case, but exceptions are rare, not more than a dozen atoms.)

DEFINITION. An AS B is called a *focal point* for an atom A if B is present in all paths and every other AS consistently either follows B or precedes B in all paths in which this AS appears.

A *crown point* is an AS that has the largest number of following ASes among all focal points.

A *crown atom* is an atom whose system of AS paths has been truncated at the crown point.

Note that if an atom has a unique origin AS, this origin is a focal point (with 0 following), and therefore most atoms have an associated crown atom.

If there is no focal point, we will let the crown atom coincide with the original atom by definition.

When AS paths are truncated, systems of paths that differ only in the truncated part may become equal, which will decrease the number of atoms. In fact, for the 01 May 2001 data, the number of crown atoms equals 15174, which is 30% less than the number of atoms, and about 7 times less than the number of prefixes.

The reduction is much larger when the number of peers is small. For example, for the 8 peers with the largest number of AS paths, the number of atoms is 20109 and the number of crown atoms is 12452. For the 6 peers present in all data sets, it is 9465. These values are not surprising since the growth in the number of peers results in growth in the number of observed AS links and in a decrease in single-homing. Crown atoms are likely to partly lose their advantage as the number of peers grows large. On the other hand, single-homed networks do exist in reality, and for these, truncation to crown atoms makes sense. The problem is, however, that truncated atoms might not have existed in the first place, and thus truncation of crowns may not reduce the number of existing atoms. (More explanation in [BC 2001b].)

C. Dependence of atoms on peer choice

An inherent limitation of our definition of atoms is that it relies on a specific collection of available peer views, and the results will differ using different views. We need to show that for a sufficiently complete set of peers, this dependence diminishes. Ideally, the addition of new peers should not change the system of atoms after we have 'complete' coverage; the equivalence relation is maximally refined, and the set of classes converged to its *projective limit*.⁷

For the data of May 01, 2001, we analyzed four other choices of peers:

- A. Top 8 peers by AS path count
- B. Top 8 peers by prefix count (without AS repetitions)
- C. 8 non-US peers

An implicit question for (C) is how much coverage it would cost if all US peers were removed from our data set.

The next table shows the dependence of atom counts on the choice of peers.

IPs	AS	Selection	Pref.	AS pth	Atoms
27	24	many prefixes	97940	14566	21512
8	8	many paths	99140	14566	20109
8	8	many prefixes	99256	14566	19368
9	8	non-US peers	98764	14207	18376

The number of atoms using by 27 big RouteViews peers is 90% reached with the 8 peers with the largest prefix and/or path count. International providers have somewhat smaller resolution, 85.4% of the maximum attained by 27 peers.

⁷A projective limit is the limit of subdivision of existing objects; its dual is the *injective limit*, which is the limit of adding completely new objects [Lang 1992].

We now seek to answer the question: how many different routing policies can exist in the Internet with regard to address blocks in the same AS?

The following table shows the data for May 03, 2001, computed with 27 peers seeing a total of 21573 atoms.

Atoms	1	2	3	4	5	6	7	8-65
#AS	6851	2188	856	430	217	150	89	262

In this data, there are 36 different atom counts that can occur for an AS. The maximum number of routing policies applied in the Internet to prefixes from the same AS is therefore 65; however, it is reasonable to assume the actual upper bound is 30, which is the last atom count assumed by two different AS (counts of 32,33,41,43,46,52,54 and 65 are assumed by one AS each. The last count with over 10 AS sharing it is 14 atoms.)

These results change slightly with 85K prefixes cutoff which enables taking 33 peers:

Atoms	1	2	3	4	5	6	7	8-60
#AS	6491	1905	738	391	185	146	85	214

E. More specifics in atoms

The prevalence of more specific prefixes is generally agreed to contribute to routing table growth. It is therefore important to know, whether more specifics are routed any differently than their less specific counterparts. The results should help explain why the use of more specifics is so popular in the first place.

We examine the distribution of more specifics through atoms. There are 97250 prefixes common to 27 peers in the table of 03 May 2001, out of a total 118K. The total number of more specifics in the union of peer tables is 71.3K, of which 51K (52.3% of 97K) are in the atoms. Among those, 15.5% are in the same atom as one of their less specifics, and 84.5% are in an atom that contains no other less specifics of the same prefix.

For the 33 peers with over 85K entries in RouteViews⁸, the result does not significantly change despite a large drop in the count of more specifics, to 38K, or 45.3% out of 83.7K common prefixes. The number of more specifics present in their atom together with a less specific prefix is 13.64% of 38K, and the number of those that have a less specific in the same atom is 86.45%.

It seems clear that more specific prefixes are introduced into the routing table for a purpose, namely for to express different routing policies from those carried by larger aggregates, as previously suggested in [Houston 2001b]. Our result shows also strong quantitative agreement with the data given by Houston for January 2001, where he states that among 37.5K more specific announcements, 30K, or 80% of them use different AS paths than their corresponding aggregate and thus are introduced to express different routing policies than the larger aggregates.

⁸This table size results from applying a specific common filtering policy [Houston 2001a].


A fundamental premise of BGP is that an AS shares only its best paths with its neighbours. If an AS implements uniform routing policies throughout its networks, then packets to a given network from anywhere in that AS will always follow the same AS path to the destination. There would be at most one outgoing link at any node of the atom's AS graph, and the graph would be an (inbound) tree.

In reality, an AS can be spread geographically and logically, and different networks within an AS may have different policies due to differences in address assignment, peer interconnection, or loose intra-AS coupling (e.g. companies acquired by a large ISP may still follow their old policies [McCreary 2001].) This policy dependence on local properties of the AS that propagated the connectivity message, though hidden in addressing by AS numbers, can still be observed through *ramification* of atoms.

Ramification is a phenomenon where several paths going to the same prefix branch or loop at an AS, resulting in the system of paths mapping to a graph with positive outdegrees and/or positive cyclomatic number, which we define below.

For the BGP table of 03 May 2001 (with selection of 27 peers with over 98K prefixes), 7532 atoms, or 35% of all atoms are ramified. This ramification, not surprisingly, is mostly observed at a handful of ASes, which either participate in RouteViews, or are known to be large providers, or both. ASes which are ramified in the largest number of atoms have counts of ramified atoms equal to 3295, 1990, 1533, 1047, 538, 481, 135 and 124. All other ASes are ramified in 35 atoms or less.

For example, one ramified atom in the table for May 01⁹ consisted of 7 prefixes with origin AS 8472, with two distinct AS paths:



Path 1 is seen by AS 1740 and 3 other peers. Path 2 is seen by AS 7018 and 5 other peers.

AS 7018 is a global provider, present almost everywhere though more dense on the US East Coast, whereas AS 1740 is present only in California. 'Hot potato' routing dictates that an AS should send traffic destined for a non-customer via the point closest to where that traffic entered its network. Consistent with this policy, AS 1740 reaches AS 8472 via a different intermediate provider, and most likely using different parts of AS 7018's infrastructure than those peering with 7018 at the East Coast.

The rest of the path system for this atom is comprised of 40 other ASes that form one or more inbound (fan-in) trees, i.e. graphs with all nodes of outdegree 1 except one (root) node with outdegree 0. We provide a quantitative analysis of atoms' cyclic complexity below.

⁹This atom is not found in the table on May 03; prefixes from AS 8472 are no longer common to all 27 peers.

The next example represents an atom with two cycles in undirected graph. Here, ramified AS paths have different lengths:

1-1239-852-838 6453-1239-852-838
 1-701-1691-852-838 6453-701-1691-852-838

We call an atom a *tangle* if it has a directed loop. The table of May 03, 2001 contains 9 tangles, 5 of them being a result of a typo in prepended sequence. Here is an example of a tangle which most likely results from traffic engineering in two cross-continental backbones. In an atom with cyclomatic number (see below) 7, paths 8 and 9 contain a directed loop between AS 1239 and AS 6461.

8) 6453-1239-2914-11908-6461-4926-4270-4387
 9) 6461-1239-5511-4000-7303-4270-4387

There are several causes of ramification. An AS may have a valid engineering reason, e.g. load balancing, to announce different AS paths to different peers, even if peerings occur within geographical and topological proximity of each other. An AS may not announce to neighbours the best available path, of announce the best path to only some of them. Another possibility is that an AS announces a path that has nothing to do with where the traffic is being sent, e.g. to keep knowledge of business relations private. All these scenarios run contrary to the basic assumptions of BGP, but BGP does not prevent them from happening. On the contrary, such flexibility is often consciously leveraged by those who know exactly what they are doing.

This reasoning suggests that we should treat ramified atoms as a rule rather than exception. The more observation points we establish, the higher the chance that some paths will enter and exit infrastructure of some global provider at different points, which will manifest as a ramification in the atom's path system. Ramification is then associated not so much with individual atoms, but with multihomed transit ASes with highly diverse policies. It is also an indication that an atom's view provided by contributing peers is reasonably complete.

A. Cyclomatic numbers

A diamond-shaped configuration which remained after removing unramified nodes is by far the most common case in the current RouteViews inventory of ramified atoms. This remainder has just one cycle when viewed as a non-directed graph.

To measure complexity of ramification, we use the notion of *cyclomatic number* [Harary 1975]:

DEFINITION. The cyclomatic number of a graph equals its count of links minus the count of nodes plus the count of connected components.

Cyclomatic number is a convenient parameter of the graph since it does not change when attached trees are removed. The cyclomatic number of a tree is 0, so it can be computed with or without preliminary stripping of trees off the graph. It measures the dimension of the first homology group of the graph viewed as a simplicial complex [Mac Lane 1995], which means that a graph with cyclomatic number c has 2^c different cycles and combinations of cycles.

On 03 May 2001, the same set of 27 peers as that used on May 01 (with 90% or more of the maximum prefix count, i.e. 99K or more) have the set of prefixes split into 21573 atoms (61 more than on May 01). The number of ramified atoms is 7532 (27 more). The whole set of atoms has the following distribution of cyclomatic number:

Cyc.num.	0	1	2	3	4	5	6-10
#atoms	14152	5678	1568	145	16	9	5

(The count for 0 includes 111 ramified atoms, containing multi-origin prefixes.)¹⁰ We will now find out how this number changes when more peers are added.

B. Dependence on peer selection. Routing policy detail.

To check how ramification in general and cyclomatic number in particular depends on the choice of peers, we analyzed the data for 03 May 2001. One ISP¹¹, who contributes RouteViews tables from two geographically distant routers, routers at MAE-East and MAE-West, has a policy of filtering prefixes larger than certain length, mostly /19 in Class A and /16 in Class B, which results in a BGP table with 85K prefixes. [Houston 2001a].

For that reason, the system of atoms obtained when we include this ISP and three other providers (6 more peers with over 85K prefixes, for a total of 33 tables), the number of atoms drops to 19870. Out of those, 19206 are the atoms of 27 peer set with prefixes missing from 33 peer set cancelled. (All 83784 prefixes of 33 peers are of course present in 27 peer set.) In other words, the addition of 6 peers refines the system of atoms by 664 units, or 3.3%. This means that the system of atoms has by a large converged to the limit when 27 peers are taken.

We compare this to the drop in the number of common prefixes (for 03 May) from 97250 to 83784, when 6 peers are added: the loss of 13466 (13.85%) prefixes leads to the loss of 2267 (11%) atoms. The reduction in atoms due to filtering is of the same order as the reduction in prefixes, and more than three times larger than the refinement of atoms due to addition of peers. This means that filtering resulted in a loss of as much of fine-grained detail about global routing policies and traffic engineering designs as the reduction in the number of prefixes that has been attained.

We further corroborate this result using the counts of atoms with a single prefix: 11808 (54.7% of total) for 27 peers; 10650 (55.4% of total) for 27 peers with filtered prefix set; and 11160 (56.2% of total) for 33 peers. Once again, the reduction in the number of one-prefix atoms is completely in line with the reduction of prefixes and all atoms: filtered set has 9.8% single-prefix atoms less than unfiltered set. The number of atoms with 2 prefixes shows a comparable drop, about 13.6%, and the number of atoms with 3 prefixes drops by 11.7%.

C. Dependence on peer selection. Ramified atoms.

The distribution of cyclomatic number for 33 peers with 83784K common prefixes and 19870 atoms, of which 11113

¹⁰This number changes as more peers are added.

¹¹The ISP read and approved of our publishing this data.

(about 50% more than for 27 peers) are ramified and 8757 unramified, is shifted towards atoms with 2 and more cycles, compared with corresponding distribution for 27 peers. Both changes are mostly contributed by two AS, with more coming from relatively less known East Coast provider whose website states that it is 7-th ISP in the US.

Cyc.num.	0	1	2	3	4	5	6-10
#atoms	8854	6333	3138	1185	290	53	17

(The count for 8854 atoms with cyclomatic number 0 includes 97 ramified atoms.)

@@ We conclude that prefix length filtering sweeps away amount of routing policy detail (in terms of atoms counts) which is proportional (equal as a ratio) to the reduction in the number of prefixes. In other words, filtering does not significantly impair the Internet's ability to formulate and express different routing policies. (Should it be otherwise, the reduction in the number of atoms would be much higher than the reduction in the number of prefixes.) Many routing policies are formulated in terms of prefixes that will not be filtered, and those that will be in many cases treated by Internet as extras attached to existing policies. Furthermore, the estimates for the number of distinct policies are reasonably independent of the selection of particular peers.

On the other hand, details of internal structure of global routing policies depend heavily upon the observation system. Same set of data can appear to have just a handful (about 2%, for the selection of 6 peers) or over 50% of ramified atoms, when peer AS are chosen from providers who maintain rich connectivity and diversity in their networks and run them using highly optimized traffic engineering solutions.

VI. CONCLUSION

We have described some the basic background needed for analyzing typical BGP tables, but we have only been able to touch on several topics that are expanded more completely in the extended version of this report [BC 2001b]. We have covered several idiosyncracies and architecturally revelant trends in current core BGP tables, and introduced the notion of *policy atoms* as the grounding for a new framework in routing table analysis. We found that the number of atoms and individual counts of atoms with given number of prefixes properly scale (change in proportion accordingly) both with the Internet's growth and with filtering of prefixes by length. The set of atoms thus represents Internet properties in an accurate way, yet with much smaller complexity. We also found that atoms' AS path systems can have rich internal structure, and that complex routing policies used by major backbone networks result in large number of ramified atoms (those with non-tree AS graphs), some of them even having directed cycles. We are continuing to investigate the use of atoms as a framework for evolving the Internet routing in the next decade.

We recognize that any attempt to capture an 'Internet route map' in its entirety will inevitably produce noise-like phenomena that render parts of the data irregular, incomprehensible or chaotic. One can never take too many precautions when analyzing Internet data. Neither common-

sense assumptions nor specifications and standards should be taken at their face value. The more we study the routing, the more astounded we are that a system with such diversity tends to work reasonably well so much of the time.

References

- [ARIN 2001] American Registry for Internet Numbers. <ftp://ftp.arin.net/pub/stats/>
- [Atkinson 2001] R.J. Atkinson, end2end-interest posting, Apr. 24, 2001
- [BC 2001a] Broido Andre and kc claffy, 'Analysis of available IPv4 address space allocation, assignment, routing data', <http://www.caida.org/broido/addr/addr.html>
- [BC 2001b] Broido Andre and kc claffy, 'Preliminary analysis of BGP routing tables: extended report' <http://www.caida.org/broido/nrdm/nrdmp.ps>
- [BC 2001c] Broido Andre and kc claffy. Internet topology, in preparation.
- [CAIDA 2001] <http://www.caida.org/tools/measurement/reversetrace/>
- [Cisco 2001] BGP Best Path Selection Algorithm. <http://www.cisco.com/warp/public/459/25.shtml>
- [Doran 2001] Sean Doran. Routing System Scaling – Disaster Looming, but Medium-Term Fixes Known. Posted to nanog@merit.edu, 2 Apr 2001.
- [Extreme 2000] Extreme value distributions. In: Engineering statistics handbook, Ch.8. National Institute of Standards, 2000. <http://www.itl.nist.gov/div898/handbook/apr/section1/apr163.htm>
- [Faloutsos 1999] Michalis Faloutsos, Petros Faloutsos, Christos Faloutsos, "On Power-Law Relationships of the Internet Topology", ACM SIGCOMM'99.
- [Lang1992] Serge Lang, Algebra, 3rd edition, Addison-Wesley, 1992.
- [L.Gao 2000] On Inferring Autonomous System Relationships in the Internet. IEEE Global Internet, Nov 2000. <http://www-unix.ecs.umass.edu/lgao/globalinternet.ps>
- [HM 2000] Sam Halabi, Danny McPherson. Internet Routing Architecture, 2nd ed, Cisco Press, 2000, 498 p.
- [Harary 1975] F. Harary. Graph Theory. Addison Wesley, 1975.
- [Houston 2001a] Geoff Houston, BGP routing table statistics, updated daily. <http://www.telstra.net/ops/bgp/>
- [Houston 2001b] Geoff Houston, 'Analyzing the Internet's BGP Routing Table', *The Internet Protocol Journal*, Volume 4, Number 1, March 2001. <http://www.telstra.net/gih/papers/ipj/4-1-bgp.pdf>
- [HBCFKLM 2000] B. Huffaker, A. Broido, kc claffy, M. Fomenkov, K. Keys, E. Lagache, D. Moore, Skitter AS Internet Graph. Published by CAIDA, 2000.
- [IRR 2001] Internet Routing Registry. List of routing registries. <http://www.irr.net/docs/list.html>
- [Knuth 1997] Donald Knuth, 'The Art of Computer Programming: Seminumerical Algorithms (Vol 2, 3rd Ed)', Addison Wesley, 1997.
- [Mac Lane 1995] Saunders Mac Lane. Homology. Classics in Mathematics, Springer-Verlag, 1995
- [McCreary 2000] Sean McCreary, private communication, 2000.

[Meyer 2001a] U. Oregon's Advanced Network Technology Center <http://www.antc.uoregon.edu/>

[Meyer 2001b] RouteViews, U. Oregon's RouteViews project, <http://www.antc.uoregon.edu/route-views/>

[Meyer 2001c] RouteViews, daily updates

<http://archive.routeviews.org/bgp>

[Meyer 2001d] David Meyer, private communication, 2001

[NLANR 2001] <http://moat.nlanr.net/Routing/rawdata/>

[PCH 2001] Sean McCreary, Bill Woodcock. PCH RouteViews archive. <http://www.pch.net/documents/data/routing-tables/>

[Rekter, Li 1995] Y.Rekhter, T.Li. A Border Gateway Protocol 4 (BGP-4) RFC 1771, March 1995.

<ftp://ftp.isi.edu/in-notes/rfc1771.txt>

[RIPE 2001] BGP data

[Stewart 1999] J.W.Stewart III. BGP4: Inter-Domain routing in the Internet. Addison-Wesley, 1999, 137 p.