An Information-Theoretic Approach to Routing Scalability

Gábor Rétvári, Dávid Szabó, András Gulyás, Attila Kőrösi, János Tapolcai Department of Telecommunications and Media Informatics Budapest University of Technology and Economics Email: {retvari,szabod,gulyas,korosi,tapolcai}@tmit.bme.hu

ABSTRACT

Many of our computer networks, not the least of which the Internet, are built upon hop-by-hop routing. At the moment, it is not clear whether we will be able to scale these networks into the future economically. In this paper, we propose a new information-theoretic model to study routing scalability, we present preliminary analysis suggesting that hop-by-hop routing tolerates network growth surprisingly efficiently, and we sketch the scalability map of the Internet which we then use to make some bold predictions.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Store and forward networks*; E.4 [Coding and Information Theory]: Data compaction and compression

Keywords

Internet scalability; hop-by-hop routing; forwarding table; data compression; entropy

1. INTRODUCTION

Throughout its several decades of history, the Internet has evolved from an experimental academic network to a full-blown global communications infrastructure. Most of the architectural transitions that have taken place in the background, from the ARPANET protocols to IP, from classful addressing to classless, from IP version 4 to version 6, were (and are) largely fueled by concerns regarding the ability of the network to accommodate future growth. And growth has been anything but light lately; according to our Internet data plane measurements, as of Jun, 2014, the size of the routed

Hotnets XIII, October 27-28 2014, Los Angeles, CA, USA

Copyright 2014 ACM 978-1-4503-3256-9/14/10...\$15.00

http://dx.doi.org/10.1145/2670518.2673863

IPv4 address space has grown by more than 7 percent throughout the last 8 months to over half a million prefixes and counting¹.

The Internet control plane, for one, presents its fair share of scalability matters [1]. Yet, it seems very likely that we can cope in the long term, by leveraging Moore's Law or, say, moving BGP entirely to the cloud [2]. In reality it is the data plane, with its custom line card ASICs and SRAM modules and stringent real-time performance requirements, where Internet scalability issues manifest themselves most markedly [3, 4]. And even though subject to concentrated research efforts [5] and heated debates [6,7], there is still no clear answer to the question as to whether the Internet routing architecture will scale to the foreseeable future.

In this paper, we argue that we need a fundamentally new model for routing state scalability to settle this debate. This model must be sufficiently and necessarily rich, omitting just the uninteresting subtleties while still delivering meaningful answers, and it must also be amenable to quantitative and qualitative argu*mentations*, admitting large-scale simulations as well as mathematical analysis providing, say, universal lower and upper bounds or asymptotic scalability characterizations. We argue furthermore that the model should not be of worst-case nature. It has been shown earlier that on certain instances of graphs some routers must store prohibitively large amounts of routing information and there is no way out of this without forwarding paths or signaling load going out of control [5,8]. Unfortunately the results do not bring over to the Internet as these worst-case graphs do not have anything in common with reality. Instead, the right model should deliver correct and verifiable scalability predictions, even when subjected to concrete topologies and routing policies. Last but not least the model should also be *imple*mentable, lending itself readily to be deployed on reallife routers.

After so much ado it should not come as particularly

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

¹We have created a website we dubbed the *Internet Routing Entropy Monitor* to publish browsable daily statistics and downloadable data sets for the community, see http://lendulet.tmit.bme.hu/fib_comp.

surprising that what we offer in this paper is such a model. Our model lifts network scalability research into the realm we believe it naturally belongs to: *informationtheory*. In particular, the model allows to reason about the scalability of general graph families or particular networks through analyzing the entropy of certain strings arising as forwarding tables at network nodes.

Below, we report on some initial results we obtained with the model. We find that in a broad class of graphs as well as in the Internet AS-level topology routing table entropy is unexpectedly low, allowing to compress routing state to a fraction of its original size. We further argue that the key to routing scalability is *heterogeneity*, either topology-wise or policy-wise. We support this claim by evaluating real routing table traces taken from operational Internet routers.

Unfortunately, without sufficient data it is impossible to make overall predictions. Thus, we feed the model into Internet-scale simulations and, surprisingly, our rudimentary analysis already produces meaningful results that nicely align with ground truth. We then use the simulations to sketch the scalability map of the Internet, highlighting the areas worst affected by network growth and those unlikely to experience any scalability stress at all, and we attempt to use our insights to make some actual predictions on the Internet's future.

The model is described in Section 2. We present the mathematical analysis in Section 3 and the simulation studies in Section 4, and finally in Section 5 we discuss our findings some further.

2. THE MODEL

The model is based on the idea that, after omitting some complicating details, forwarding tables (the data structures that maintain forwarding state in hop-by-hop routing and hence shape scalability in the large) eventually boil down to simple strings, which can then be subjected to information-theoretic analysis.

One of the most complicated aspects of Internet routing is the inherent structure in the IP address space, whereas hosts are collected on an administrative or geographical basis into a hierarchy of subnets and supernets. For one, it is notoriously difficult to reason about the information content of such structured data [9]. Second, it is not immediately clear how much information is encoded intrinsically within the address space itself and what remains to be encoded in the routing state. But more importantly, even the optimal encoding of IP forwarding tables ends up separating structure from flat labels [10]; containing a separate component encoding the prefix hierarchy whose size depends only on the number of IP prefixes and hence, by and large, constant across the Internet; and another component, unique to each router, encoding the association of next-hops to flat IP prefixes. Correspondingly, we chose to go with

a flat address space, to which structure can be added back in a separate step.

So we are given an undirected graph G(V, E) on nnodes, with each node $v \in V$ labeled with an id id(v), a unique integer in [1, n], and each outgoing port $(v, u) \in$ E of node v labeled with a locally unique id $port(v, u) \in$ $[1, \delta_v]$, where δ_v denotes the node degree of v. Port ids are local and hence we are free to permute these in the range $[1, \delta_v]$; contrariwise, vertex ids are global and thereby essentially random from our perspective. We further assume that routing policies (like shortestpath or valley-free) for each node have been fixed in advance and packet forwarding must strictly obey the paths emerging from these policies, but as far as our model is concerned we do not assume any particularity about the routing policies themselves whatsoever.

Packets contain a header with the id of the destination node, which is then fed to the local routing function s_v of each node v along the forwarding path to determine the output port to pass the packet on. The routing function $s_v : [1, n] \mapsto [1, \delta_v]$ itself maps a destination node id (as found in the packet header) to the corresponding outgoing port. We suppose that each node $v \in V$ is aware of its own id (this immediately imposes lg(n) bits² lower space bound for storing s_v) and hence can identify packets destined to itself, so we shall usually set the symbol of s_v at position id(v) arbitrary. A simple way to look at s_v is to think of it as a string of length n on the alphabet $\Sigma_v = [1, \delta_v]$, so that the symbol $s_v(i)$ at position $i \in [1, n]$ gives the output port to be used to forward packets towards the node with id *i*. Then, our task is to find lower and upper bounds on the number of bits needed to encode this string in general graphs, like the Internet.

A naïve encoding of the routing function s_v at any node v would require $n \lg \delta_v$ bits (n times the number of bits needed to distinguish the ports of v). This is also a lower bound on the size of s_v when we do not have any prior knowledge on s_v that we could exploit to compress it (i.e., the distribution of port ids is essentially uniform). Hence, we shall view this quantity (the so called *information-theoretic limit*) as the uncompressed size of the routing state. However, if we possess some knowledge on the input, say, we know a zero-order model (symbols of the input are independent, identically distributed random variables from a known distribution), then we can leverage that knowledge to compress beyond this limit. We can then give firm lower and upper bounds on the expected size of an optimal encoding of the input and, as shown below, this argumentation goes over to forwarding tables seamlessly.

DEFINITION 1. The zero-order entropy bound for storing the routing function s_v at node v is $nH_0(v)$ bits,

²The notation $\lg x$ is shorthand for $\log_2(x)$ (or $\lceil \log_2(x) \rceil$ wherever it makes more sense).

where

$$H_0(v) = \sum_{j \in \Sigma_v} p_j \log_2 \frac{1}{p_j}$$

is the Shannon-entropy of the distribution $p_j : j \in \Sigma_v$ with each p_j describing the probability of port j to appear as a next-hop port in s_v .

The significance of the entropy bound is multi-faceted. First, note that $H_0(v) \leq \log \delta_v$ with equality if and only if the distribution of port ids in the routing function is completely uniform. Second, nH_0 is an upper bound on the amount of space needed to encode s_v ; one could use Huffman-coding to realize the entropy bound but this might make fast lookups on the compressed form impossible; to both encode to zero-order entropy and facilitate fast lookup one can use for instance wavelet trees [11], providing $O(\log n)$ random access and $nH_0 + o(n)$ bits storage space, or the simple scheme in [11] to get (roughly) O(1) lookups. But most importantly, the entropy bound is also a firm lower bound for realizing the routing function, as shown next.

THEOREM 1. Given a graph G(V, E) on n nodes with each node assigned a unique id in [1, n] randomly, any encoding of the routing function $s_v : v \in V$ needs at least $nH_0(v)$ bits.

PROOF (SKETCH). The proof is based on the idea that any encoding B that could store s_v on fewer than nH_0 bits would violate the (noiseless) source coding theorem due to Shannon. This is because for any string swe can construct a graph that reproduces s at some node as a routing function, whose encoding with B below nH_0 bits would now lead to a contradiction. \Box

Observe the generality of this result: for any graph and *any* routing policy there is simply no way to encode the routing function below the entropy bound without actual information loss (i.e., a violation of the routing policy). Furthermore, the lower bound remains firm as long as we do not have control over the allocation of node ids and thereby the order of symbols in s_v is essentially random³. The entropy bound is *computable* for any graph at least empirically: we just need to obtain the forwarding functions and then evaluate the entropy over the port id distribution, which can certainly be done with enough computing power at hand; the results are *verifiable* as we can easily evaluate the entropy on real forwarding table traces from production routers (see later); and the model is also *implementable* in actual router hardware, even though current wavelet tree implementations' performance might lag behind a bit (but see [10] on how to overcome this issue).

3. ANALYSIS

Are there networks that admit hop-by-hop routing without overflowing routers' memory, even as the underlying network grows larger? And, on the contrary, are there networks and/or routing policies for which forwarding state suddenly explodes after a certain point? Below, we present some rudimentary analysis to answer (some of) these questions. For simplicity, in the below analysis we shall fix the routing policy as shortest-path routing over positive link costs (but we note that the analysis would go similarly for any routing policy [8]).

Our main scalability measures will be the entropy bound of $nH_0(v)$ bits that, recall, gives the minimum and maximum amount of information a node v must store to encode its routing function. We also look for the compression efficiency $\eta(v) = H_0(v)/\lg \delta_v$, describing the ratio of the size of the compressed routing function to the uncompressed one (this is the reciprocal of the usual compression ratio metric).

We first present the bad news: it is very simple to construct graphs where s_v cannot be encoded efficiently.

LEMMA 1. Let K_n be the complete graph on n nodes with unit cost on each edge. Then, for any node v we have $H_0(v) = \lg n$.

The proof is trivial by observing that each port id appears exactly once in the routing function (omitting a node's own entry). Thusly, the port id distribution is uniform and the entropy bound matches the informationtheoretic limit. This on the one hand means that hopby-hop routing is unscalable on complete graphs (at least in a zero-order setting, that is, when node ids are random) as state grows superlinearly with size and, on the other hand, this is also an example for the case when s_v cannot be compressed at all (i.e., $\eta = 1$).

In some sense, the above negative result stems from the fact that a complete graph with unit link costs is totally *homogeneous*, so each port appears in s_v with equal probability. We observe, nevertheless, that if the network exhibits some form of *heterogeneity* then the port id distribution in s_v becomes skewed, which tends to reduce entropy in most cases.

We first break the symmetry by introducing a slight random link cost variation into our complete graphs. That is, instead of a uniform cost model we turn to a model where link costs are chosen randomly from a known distribution.

LEMMA 2. Let K_n be the complete graph on n nodes and let the cost on each edge be an i.i.d. random variable chosen from Exp(1). Then, for a randomly picked node v we have $\mathbb{E}(H_0(v)) = \lg e$ as n grows to infinity.

The proof is based on the observation that under this random graph model the size of the branches of the shortest path tree is described by the Chinese restaurant

³If, on the other hand, we are free to assign node addresses then we can design some structure into the address space, which opens the door to higher-order encoders to go beyond the zero-order entropy bound; such "name-dependent" routing schemes are, however, beyond the scope of this paper [5].

model [12]. Note also that the result remains true for a wide variety of i.i.d. link cost distributions.

What is remarkable in this finding is that the slight diversity introduced by random link costs already reduces the entropy of the routing function to $\lg e \approx 1.44$ bits. This is not even dependent on the network size, even though the size of the alphabet that underlies the routing function (i.e., the number of ports) grows without limit. This is despite that the expected value of the link costs is still 1 like above and also despite that locally the graph is completely homogeneous. Also note that the information-theoretic limit remains $\lg n$, so s_v in this case is *infinitely compressible* $(\eta \to 0)$.

What if we go the other way around and we alter the topology instead of the link costs? The natural procession of the above line of thought would be to fix link costs at unit and this time let the randomness shape the existence of edges, i.e., to go with an Erdős-Rényi random graph model. Unfortunately, this model has this far resisted our analytical efforts. Instead, we go with another uniform cost random topology that is, similarly to an Erdős-Rényi graph, locally completely homogeneous: the hypercube.

LEMMA 3. Let G_d be the hypercube graph in d dimensions with unit cost on each edge and suppose there is a strict order on the ports for each node to break ties between equal cost shortest paths. Then, for any node v we have $H_0(v) \leq 2$.

The proof is based on the observation that, under the above conditions, the next-hop distribution in the routing function forms a geometric sequence with common ratio 1/2, and so $H_0 = \sum_i 2^{-i} \lg(2^i) = \sum_i i 2^{-i} < 2$.

So it seems that there is something very interesting going on here. Intuitively, one would not expect hopby-hop routing to scale particularly well: why would certain next-hops concentrate more traffic than others, after all? This is indeed the case for completely homogeneous networks, like the complete graph, that exhibit uniform next-hop distribution and therefore maximum entropy. However, if we introduce heterogeneity either topology-wise or policy-wise then dominant next-hops begin to emerge, producing highly skewed next-hop distributions and very small forwarding state entropy, all in all, excellent scalability and infinite compressibility.

4. SIMULATIONS

A compelling question now arises: Does the mathematical analysis apply to the Internet in some way?

So, we set out to simulate the Internet! Rather an ambitious undertaking, given that we do not know the router-level Internet topology, let alone the routing policies deployed by service providers. Thus, we went with a gravely simplified simulation methodology: we used the valley-free policy routing model over the Autonomous System (AS) level Internet graph topology with inferred AS policies. Our goals were likewise modest, in that we merely asked how much information would *any* routing protocol need (not necessarily BGP) to store at Internet routers to encode the inter-domain path system. Consequently, our conclusions only remotely apply to fullfledged Internet routing, complete with more-specifics, unannounced addresses, etc. Or perhaps they do, but we shall certainly need a reality check to see this.

Our simulation methodology was as follows: take the CAIDA AS-graph with inferred AS-AS business relationships [13]; calculate all the valid inter-domain paths for each AS-pair according to the valley-free routing model; select the best path according to the Gao-Rexford conditions (prefer customer paths over peer paths over provider paths) with ties broken by shortest AS-path length; create a "per-AS" routing function for each AS (with one entry per destination AS) by writing down the port id corresponding to the best paths one after the other⁴; convert this to a per-prefix routing function (one entry per IP address prefix) with repeating the next-hop entry for each AS as many times as the number of prefixes that AS announces into the inter-domain routing system [14]; and finally take the resultant strings and calculate the entropy as of Definition 1.

Essentially every single aspect of this methodology is arguable, ranging from using an AS-level graph to estimate a router-level measure (the routing function), to the inherent imprecision of the AS-graph [15], the inferred routing policies, and the Gao-Rexford conditions (but see also [16]), all the way to the fact that our model is for a continuous unstructured address space while the IP address space is fundamentally hierarchical and contains "holes".

Even so, it turns out that the results still match ground truth to some extent, in the sense that the routing function entropy produced by the simulations is in the same order of magnitude as that of real IP forwarding tables (FIBs) measured in terms of the FIB-entropy metric of [10]. As of June 1, 2014, on a set of 10 IP FIBs downloaded from the Internet2 (AS11537) entropy is varying between 1.3 and 1.7 bits (with a mean of 1.45) bits) and our simulations yield 1.72 bits; for a set of 10 FIBs obtained from the HBONE (AS1955) we get 1.28-2.56 bits of entropy (mean: 2.16) in contrast to our simulation results of 1.55 bits. Most of these FIBs, however, contain default gateways and hence are not quite representative for inter-domain routing; for a restricted set of 3 full-BGP FIB instances from the HBONE (containing on average 502,060 entries) the mean entropy is 1.28 bits. Strikingly, our simulations quite precisely reproduce this result; we get 1.26 bits of entropy for

⁴Note that in a zero-order model like ours the order of nexthop ports within the routing function does not matter, only their relative frequency counts.



Figure 1: Entropy and compression efficiency by tier.

the HBONE after manually correcting the CAIDA ASgraph using the routing policies published at the Routing Arbiter Database. The simulations are robust across a wide range of input parameters and local preference settings, and even compensating for blackhole routes ("holes" in the IP address space) do not alter the results substantially.

Easily, this close match is more probably an accidental simulation artifact than tangible proof: after all, we didn't expect the simulations to deliver *quantitatively* correct results. Until we manage to collect enough IP FIBs and we validate the results comprehensively, all further conclusions drawn from the simulations should be taken with a certain (large) grain of salt.

Fig. 1 gives a visualization of the simulation results, depicting the routing function entropy and compression efficiency per AS tier. Here, Tier 1 means the topmost clique of 16 ASes (as of the CAIDA AS graph) connected into a peering mesh without upstream provider, and subsequent tiers mark ASes located in increasing distance from this top tier. Note that the Internet2 and the HBONE are Tier 2 in this tier-based classification. In addition, we collected all the results into a giant *Internet scalability map*, rendering the Internet inter-domain routing system within an illustrative hyperbolic layout (taken from [17]), with smaller blue vertices marking the ASes of low entropy (0–2 bits) and bigger red vertices marking the ASes that suffer high routing table entropy. See Fig. 2.

The conclusions stand out. ASes in the Internet Tier 1 experience consistently higher entropy than those at lower tiers (most of the red nodes on the map are Tier 1), with values as high as 7–8 bits. What is worse, the routing state at the Tier 1 does not seem really compressible; the uncompressed size is only about twice as large as the compressed one, indicating a close to uniform next-hop distribution in the forwarding tables. The reason appears to be that from the top tier we see a vastly homogeneous Internet, with no second tier ASes attracting significantly more prefixes than others.

In contrast, below the Tier 1, that is, at essentially 99.9% of all ASes the entropy of routing state is below 1 bit and compression efficiency is much better. This means that forwarding tables at lower tiers may be more regular and more compressible (to the point that a single-homed AS's routing state is compressible to zero bits in theory), suggesting that these ASes suffer way lower scalability stress than those at top tiers. In other words, from the lower tiers we see a much more heterogeneous Internet, with most of the prefixes available via only a few primary upstream providers (which therefore concentrate lot of prefixes, introducing a huge bias into the next-hop distribution in the routing function) and customer and peer ASes contributing very little to the entropy.

5. DISCUSSION

So, does the Internet routing system scale after all? We would say that the answer is generally affirmative, at least as long as our analysis goes.

It seems that we do not need that much information to encode the current inter-domain routing state of the Internet. For the roughly half million prefixes announced into the routing system today, the simulations provide entropy bounds (which, recall, indicate the absolute lower and upper bound on the information content of ASes' forwarding state) ranging from 50–70 Kbytes at lower tiers to about half a megabyte at the Tier 1. For the Internet2 forwarding tables (note that these contain only about 14,000 entries) the simulations give an entropy bound of 20 Kbytes (15 Kbytes empirically), and for the HBONE the bound is 80 Kbytes (70 Kbyte on the real data).

But is 80 Kbytes, an equivalent of answering about 600,000 yes-or-no questions, a lot to encode the route to every single host on the Internet? After all, this is barely as large as any average PDF document, low resolution photo image, or small novel, and way smaller than any computer encoded music file, photo of reasonable resolution, or most of the executable files on your PC for that matter. And this is well within the reach of the first ARPANET IMPs (which later evolved to Internet routers) or popular home computers of the 80s. And still, this much information is enough to know the whereabouts of about half a million IP prefixes we see in the routed Internet from the HBONE today. Note also that there are actual working FIB encoders capable to attain this (or a small multiple of this) size near reasonable lookup and update performance [10].

If we extrapolate from the results to the near-future, we do not see too much reason to ring the alarm bells either. It seems likely that we will be able scale our routers to the foreseeable future, provided that current



Figure 2: The Internet scalability map: routing function entropy visualized in a hyperbolic layout of the AS graph, with vertex size and color (from blue to red) varying proportionally with entropy.

growth trends continue, i.e., the appearance of new prefixes follows a linear or slow quadratic trend [18].

But what if things go really wrong all of a sudden? What if the secondary IPv4 address market goes berserk, what if /26-s and /28-s begin to appear *en masse* in the Default Free Zone, what if IPv6 really catches on and every home and small enterprise network obtains its provider independent multi-homed /48 or /64, all in all, what if, say, 10 million IP prefixes show up in the interdomain routing system overnight? Then again, this is only 10 Mbytes worth of memory (at the Tier 1), which would fit into the line card memory of basically every off-the-shelf WAN router marketed today. And even if we were to announce every single IPv4 address into the routing system, this still would impose only about 4 Gbytes of routing state at the Tier1, which is again within the reach of modern data plane technologies [19].

There are, however, certain worrying trends worth being aware of. For instance, it seems that ASes at top tiers are the ones who really pay the price for Internet growth, with Tier 1 ASes paying one order of magnitude more (roughly 1 byte per prefix) than those at low tiers (less than 1 bit per prefix) for the addition of every single new prefix to the inter-domain routing system. This can concentrate scalability stress to a few large ASes, which may cause new tussles between Internet stakeholders [20]. Easily, all this is true only as long as we can extrapolate from the simulation results. We believe that some simple refinements to the methodology would could go a long way to increase our confidence in these results; for instance address space structure could be re-introduced easily by transitioning to the corresponding FIB-entropy notion of [10]; and we could also get more realistic routing table size estimates by observing that, as far as the information-theoretic model is concerned, it is enough to take *router-level topology* into account only *locally*, while the remote portion of the Internet can still be represented by its aggregate AS-level map.

Even so, the single most intriguing conclusion from the above seems to be that regularity somehow tends to miraculously emerge in the forwarding state associated with large-scale hop-by-hop routing. The reasons, emergence, and the characterization of this vast regularity, and the role network heterogeneity plays in this, are for interesting further study.

6. ACKNOWLEDGEMENTS

J. T. is with the MTA-BME Future Internet Research Group and A. K. is with the MTA-BME Information Systems Research Group. This work was supported by the HSNLab and the TAMOP 421B-09/1/KMR-2010-0002, OTKA/PD-104939, and OTKA/108947 grants.

7. REFERENCES

- Ahmed Elmokashfi and Amogh Dhamdhere. Revisiting BGP churn growth. SIGCOMM Comput. Commun. Rev., 44(1):5–12, December 2013.
- [2] Debayan Gupta, Aaron Segal, Aurojit Panda, Gil Segev, Michael Schapira, Joan Feigenbaum, Jennifer Rexford, and Scott Shenker. A new approach to interdomain routing based on secure multi-party computation. In *HotNets*, pages 37–42. ACM, 2012.
- [3] Xiaoliang Zhao, Dante J. Pacella, and Jason Schiller. Routing scalability: an operator's view. *IEEE JSAC*, 28(8):1262–1270, 2010.
- [4] Varun Khare, Dan Jen, Xin Zhao, Yaoqing Liu, Dan Massey, Lan Wang, Beichuan Zhang, and Lixia Zhang. Evolution towards global routing scalability. *IEEE JSAC*, 28(8):1363–1375, 2010.
- [5] Dmitri Krioukov, k c claffy, Kevin Fall, and Arthur Brady. On compact routing for the Internet. SIGCOMM Comput. Commun. Rev., 37(3):41–52, July 2007.
- [6] D. Meyer, L. Zhang, and K. Fall. Report from the IAB Workshop on Routing and Addressing. RFC 4984, 2007.
- [7] Kevin Fall, Gianluca Iannaccone, Sylvia Ratnasamy, and P. Brighten Godfrey. Routing tables: Is smaller really much better? In ACM HotNets-VIII, 2009.
- [8] Gábor Rétvári, András Gulyás, Zalán Heszberger, Márton Csernai, and József J. Bíró. Compact policy routing. In Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC '11, pages 149–158, 2011.
- [9] Rajeev Raman, Venkatesh Raman, and S. Srinivasa Rao. Succinct indexable dictionaries with applications to encoding k-ary trees and multisets. In ACM-SIAM SODA, pages 233–242, 2002.

- [10] Gábor Rétvári, János Tapolcai, Attila Kőrösi, András Majdán, and Zalán Heszberger. Compressing IP forwarding tables: towards entropy bounds and beyond. In ACM SIGCOMM 2013, pages 111–122, 2013. Revised version available online: http://arxiv.org/abs/1402.1194.
- [11] Paolo Ferragina, Raffaele Giancarlo, and Giovanni Manzini. The myriad virtues of wavelet trees. *Inf. Comput.*, 207(8):849–866, 2009.
- [12] Svante Janson. One, two and three times log n/n for paths in a complete graph with random weights. *Comb. Probab. Comput.*, 8(4):347–361, July 1999.
- [13] The CAIDA AS relationships dataset (2013-08-01). http://www.caida.org/data/ active/as-relationships.
- [14] Routeviews prefix to as mappings dataset for IPv4 and IPv6 (2014-05-01). http://data.caida.org/ datasets/routing/routeviews-prefix2as.
- [15] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. Anatomy of a large european IXP. In *SIGCOMM*, pages 163–174. ACM, 2012.
- [16] Phillipa Gill, Michael Schapira, and Sharon Goldberg. A survey of interdomain routing policies. *Computer Communication Review*, 44(1):28–34, 2014.
- [17] Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. Sustaining the Internet with hyperbolic mapping. *Nature Communications*, 1:62, 2010.
- [18] G. Huston. BGP in 2013. http://www.potaroo.net/ispcol/2014-01/bgp2013.html.
- [19] Intel Corporation. Intel data plane development kit: Programmers guide, November 2012.
- [20] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden. Tussle in cyberspace: Defining tomorrow's Internet. In ACM SIGCOMM'02, pages 347–356, 2002.