A Survey of Interdomain Routing Policies

Phillipa Gill
Stony Brook University
phillipa@cs.stonybrook.edu

Michael Schapira
Hebrew University of Jerusalem
schapiram@huji.ac.il

Sharon Goldberg
Boston University
goldbe@cs.bu.edu

Abstract

Researchers studying the inter-domain routing system typically rely on models to fill in the gaps created by the lack of information about the business relationships and routing policies used by individual autonomous systems. To shed light on this unknown information, we asked ≈ 100 network operators about their routing policies, billing models, and thoughts on routing security. This short paper reports the survey's results and discusses their implications.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols
General Terms: Measurement, Experimentation.

1. INTRODUCTION

BGP enables autonomous systems (ASes) to customize their *routing policies* to select AS-level routes through the Internet. Because these routing policies are used to realize an AS's economic, performance, security, and traffic engineering goals, the details of these routing policies are often kept secret. However, information about interdomain routing policies is crucial for answering important research questions (*e.g.,* network reliability [22, 41], routing convergence [11, 12], incentive compatibility [8, 15, 27], geopolitical control [25], secure routing [2–4, 13, 29], and the design of tools for debugging routing problems [21]). Many studies [2, 3, 13, 16, 29] resort to using routing policy *models* developed over a decade ago [11, 12, 17–19]. How accurate are these models? Researchers have highlighted exceptions to modeled routing policies [10, 14, 32, 37], or aspects of interdomain routing that these models omit [35].

To shed more light on the routing policies and on other operational issues pertaining to interdomain routing with BGP, we circulated a survey to multiple network operator communities, collecting almost 100 responses. As with any survey, the results should be treated with caution; they are highly dependent on who was asked and who answered, and are subject to "noise" that arises when questions are misunderstood. However, survey responses do provide anecdotal evidence of the presence of certain operational behaviors in routing, give rise to new research questions, and motivate more thorough surveys and analysis. Responses to the survey (Section 3, Tables 1-2) highlight the following:

Revisiting models of routing policies. The Gao-Rexford model [11, 12, 17–19] is perhaps the most popular model of interdomain routing. The model supposes the routing policies obey constraints imposed by business relationships between neighboring autonomous systems. While a number of measurement studies have used public data to *infer* that these constraints are sometimes violated [10, 14, 32, 37], our survey presented a unique opportunity to ask network operators *directly* about when and how they violate these constraints; responses suggest that violations are the exception, rather than the rule (Section 4.1). These exceptions may arise from incorrect initial assumptions or changes in routing policies over the past decade.

Existing models of BGP dynamics suppose that routes selected by ASes are memoryless functions of the network topology and ASes' routing policies. Responses to our survey call this into question; indeed, a majority of respondents use routing policies that also depend on the history of interactions between neighboring routers (Section 4.3).

We also asked about other restrictions on routing policies (namely, next-hop preferences and consistent-export policies (Section 4.2)), and about discrepancies in routing policies between routers in the same AS (Section 4.4).

MRAI timers and security. We asked about two changes to BGP currently being debated by the standards community. First, we asked about the Min Route Advertisement Interval (MRAI) timer, which limits the rate of update messages between neighboring BGP-speaking routers. Early BGP RFCs [33] suggest a default MRAI interval of 30 seconds, but both vendors and the standards community are moving towards shorter intervals, and even disabling MRAI timers completely. We found that a vast majority (90%) of respondents have disabled MRAI timers (Section 5). Second, we asked operators how they would incorporate information provided by BGPSEC [26] — a secure variant of BGP that is currently being standardized — into their routing policies; responses highlight a lack of consensus on how this should be done. In light of recent results [29], responses also suggest that the policies favored by operators can greatly limit the effectiveness of BGPSEC (Section 6).

Revisiting billing models. Finally, we asked operators about their billing models, and found that billing based on traffic volume remains popular (Section 7)).

2. BACKGROUND: ROUTING WITH BGP

BGP is the Internet's interdomain routing protocol. A BGP-speaking router uses BGP to learn AS-level paths to destination IP prefixes via routing announcements from its neighboring routers. For each destination IP prefix, the router selects a single path from the set of routes learned from its neighbors, according to its configured *routing policies.* An *export policies* is then applied to announce the route to a subset of its neighbors.

Table 1: Questions asked in the survey and (aggregated) responses. The survey also accompanied Questions 3-12 with a free-form text boxes labeled "Why? (optional)" where operators could elaborate on responses.

Question	Response
1. What kind of network do you operate?	Table 3
2. On what continent is your network?	Section 3
3. Do you always assign a higher LocalPref (see Step 1 in the table) to a path through your customer than to a path through your peer or transit provider? (Note: exclude cases where routes through customers are tagged as backup.)	Y(77) N(16) NA(4)
4. Does your LocalPref configuration depend only on the next-hop AS (and not on other ASes on the path)?	Y(54) N(34) NA(9)
5. Do you use the same LocalPref configuration across all BGP-speaking routers in your network?	Y(72) N(22) NA(3)
6. Is the "prefer oldest path" step (see Step 7 in the table) enabled on your BGP-speaking routers? (Note: this step is enabled by default on Cisco routers in the last few years.)	Y(71) N(22) NA(4)
7. If path validation (eg BGPSec) was deployed in your network, **before** what step (1-8) in the table would you place the following step: "Prefer secure paths (validated paths) over insecure paths"? Select a number from 1-8.	Section 6
8. Do you do any neighbor-specific best path selection e.g., select a different best path for different customers for policy reasons (and **not** due to hot-potato routing, etc.)?	Y(38) N(54) NA(5)
9. Do you announce paths from peers and providers to other peers and providers?	Y(21) N(73) NA(3)
10. If you are willing to announce a certain path to a neighboring AS, are you also willing to announce any path with higher LocalPref (should it become the best path) to the same AS?	Y(47) N(30) NA(20)
11. What MRAI timer value is used in your network (in seconds)? (0 specifies that MRAI timers are not used).	Section 5
12. If you are at an ISP, do you use 95/5 percentile pricing with your customers?	Y(46) N(38) NA(13)

The BGP decision process. The BGP decision process is a set of steps that the router uses to select a route from a set of routes it learns from its neighbors. Unlike protocols based on shortest-path routing (*e.g.,* OSPF), the BGP decision process presents ASes with flexibility to realize arbitrary routing policies that maximize their local objectives (*e.g.,* profit, performance *etc.*). A simplified version of the BGP decision process, based on [6, 24], is shown in Table 2.

Perhaps the most important "knob" provided by the BGP decision process is the LocalPref step. Each route a BGP router learns about is tagged with a LocalPref attribute based on the network's routing policies, and the route with the highest LocalPref is selected. LocalPref can, for example, be used to ensure that revenue-generating routes are preferred over expensive (revenue-depleting) routes, or that routes containing a particular 'undesirable' AS are avoided. The specific details of the LocalPref policies used by individual ASes are largely kept private, often because they are related to business agreements between neighboring ASes.

The "AS-path" step follows the LocalPref step (Table 2). In this step, a BGP router that has a choice between multiple routes with the same LocalPref will select routes that are shortest in terms of the number of ASes on the path. If there are multiple such routes, the router applies the remaining steps in the BGP decision process to select a single route. These steps are based on intradomain and tiebreaking criteria; for example, the MED attribute (step 4) allows an AS with multiple entry points to influence the entry point chosen by a neighboring AS [7]. The many steps in the BGP decision process can create significant complexity, both for researchers seeking to understand the process, and for operators seeking to use it; indeed, operators sometimes respond to this complexity by disabling some of the steps, *e.g.,* by ignoring MEDs or disabling the "prefer oldest path" step.

Export policy. Once a router selects a route, its configured *export policy* determines the subset of its neighbors to which it announces the route. Export policies are often determined by business relationships, and therefore often kept private as well.

3. THE SURVEY

The survey asked the questions in Table 1 and included the information on the BGP decision process in Table 2. Information about the survey was sent in Fall 2011 to a number of network operator mailing lists. (Specifically, `nanog@nanog.org` on Sept 8 and Sept 13, `juniper-nsp@puck.nether.net` on Oct 5, `ripe-list@ripe.net` on Oct 17, and `apnic-talk@lists.apnic.net` on Oct 27, 2011.) We collected a total of 100 responses from 98 unique IP addresses. After removing an anomalous response (two consecutive responses from the same IP),we ended up with a total of 97 responses from unique IP addresses. The survey allowed respondents to omit questions that they did not understand or did not want to answer. We report results normalized by the number of respondents to a given question.

Table 2: Simplified BGP decision process [6, 24]. This table was also provided with the survey.

#	Criteria
1	Highest LocalPref
2	Lowest AS Path Length
3	Lowest origin type
4	Lowest MED
5	eBGP-learned over iBGP-learned
6	Lowest IGP cost to border router (hot-potato routing)
7	If both paths are external, prefer the path that was received first (i.e., the oldest path) [6]
8	Lowest router ID (to break ties)

Breakdown of respondents. We had 44 responses from the ARIN (North America) region, 34 from RIPE (Europe), and 19 from APNIC (Asia Pacific). We did not obtain responses from LACNIC (Latin America) or AfriNIC (Africa), likely because we did not target the right mailing lists for these regions. Table 3 shows the breakdown of responses by network type. Larger transit providers are over-represented in our responses, likely because operators of larger networks tend to be most active on mailing lists. As such, some of our results might be biased towards highlighting exceptions to standard modeling assumptions (Section 4.1-4.2), since these highly-skilled operators are more likely to use "exotic" routing policies to manage their complex networks. Specifically, 79% of responses come from small/medium/large transit providers and tier 1 networks, and only 7% come from content providers. Stub ASes (*i.e.,* AS without customers, that do not transit traffic for other networks) are under-represented in the responses we collected (*i.e.,* only 12% of respondents operate stubs, but 85% of networks in published AS topologies [5] are stub ASes). To account for some of this bias, we will often break down responses by network type.

Table 3: Responses to Question 1 (network type).

Responses	Network type
30	medium transit (11-100 customer ASes, have providers)
27	small transit (< 10 customer ASes, have providers)
16	large transit (< 100 customer ASes, have providers)
12	stubs (no customer ASes)
5	large content provider
4	tier 1 (no provider ASes)
2	small content provider
1	not-specified

4. ROUTING & EXPORT POLICIES

We now present background on routing policies models and discuss the results of the Q3-Q6 and Q9-Q10.

4.1 The Huston / Gao-Rexford Model

The now standard model of routing policies was developed by Gao and Rexford [11, 12] based on seminal work by Griffin, Sheppard, and Wilfong [17] and Huston [18, 19]. The simplest version of the model supposes that a pair of neighboring ASes have one of two business relationships: (1) *customer-to-provider*, where the customer AS purchases connectivity from its provider AS, or (2) *peer-to-peer*, where the two ASes transit each other's traffic for free. The Gao-Rexford model supposes that:

GR preference. All ASes use a LocalPref where routes through a neighboring customer (*i.e.,* "customer routes") are preferred over routes through a neighboring peer or provider (*i.e.,* "peer routes" or "provider routes").

GR export. A customer route may be exported to all neighboring ASes. A peer or provider route may only be exported to customers.

The preference condition captures ASes' incentives to send traffic along revenue-generating customer routes, as opposed to peer routes (which do not increase revenue), or provider routes (which come at a monetary cost). The export condition captures ASes' willingness to transit traffic from one neighbor to another only when paid to do so by a customer.

The appeal of the Gao-Rexford model lies not only in its simplicity, and the belief that it (mostly) aligns with ASes routing and export policies, but also in its important implications BGP stability: when all ASes adhere to the above two Gao-Rexford constraints, and under the assumption that no AS is an indirect provider of itself (there are no directed cycles of customer-to-provider relationships in the AS-level hierarchy), BGP is guaranteed to converge to a stable routing outcome [11, 12, 17]. Consequently, the GR model is used in a large number of studies concerned with route selection in BGP (*e.g.,* [3, 13, 16, 25, 41]) and also spawned a cottage industry of research (see [28] and references therein) that uses publicly-available data (from *e.g.,* route collectors [1, 34]) to infer AS business relationships.

Most surveyed networks follow GR Preference. Figure 1 shows the breakdown of 76 responses from transit providers (*i.e.,* small/medium/large transit and Tier 1 per Table 3) that answered Q3 and Q9 of the survey. (We omit responses by stub ASes and content providers ASes because they generally do not have customers that pay for transit.) The majority of surveyed transit providers (68%) indicate that they employ both GR Preference and GR Export in their network and a total of 87% of the transit providers use GR Preference (but not necessarily GR Export). A look at the free-form comments from Q3 confirms the intuition behind GR preference: one operator commented that *"Traffic on customer ports = increased revenue/profit. Traffic on supplier ports = increased cost"* and another said customer routes are preferred *"because they pay us"*.

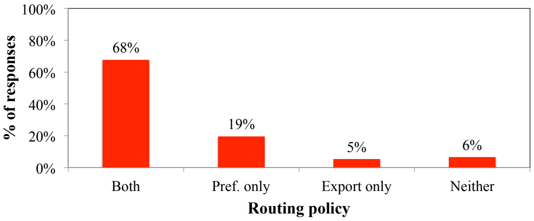

Figure 1: Responses to survey Q3 (GR Preference) and Q9 (GR Export) by transit providers (small/medium/large transit and Tier 1s.)

Exceptions to GR export. We also found that the most common exception to the Gao-Rexford model was networks that implement GR preference but not GR export. While all four Tier 1 respondents indicated that they follow GR Export, we found that 6-7 respondents from small/medium/large transit providers (Table 3) indicate that they violate GR export, which ranges from 23% to 37% of such respondents who answered Q9. (We note, however, that this might be an overestimate of the number of violations of GR Export, since we intended the term "peer" in Q9 to imply "settlement-free peer", and the comments suggest that some operators interpreted it as "neighboring AS".) The reasons cited for violating GR Export included: *"To be a good neighbor ... providing good paths to Akamai / Google caches"*, *"Mutual backup transit"* and *"Secret sauce, secret agreements."* which suggest that sometimes, an AS that violates GR export might do so only for a small number of destination IP prefixes, or only when the network is under duress and backup is needed.

However, the following comment indicates systematic violations of GR export: *"[we don't it] but it does go on. Peer/transit swaps in exchange for someone else's peer/transit routes. ... Some networks try to increase the size of themselves this way without being asked to do so. ... This generally breaks routing and it's horrible that people do it."*

4.2 Next-Hop Policy & Consistent Export.

We now turn our attention to the following orthogonal restrictions on routing policies:

Next-hop LocalPref (NH). An AS a uses a next-hop LocalPref if AS a assigns the same LocalPref to *every* path announced by a given neighbor AS b for *a given destination IP prefix* d. (*e.g.,* paths $abPd$, $abRd$ have equal LocalPref, even if path segments P and R differ; also, b is the *next hop.*)

Consistent export (CE) [8]. Consistent export intuitively means that ASes' export rules align with their routing policy. Specifically, if AS a exports path R to some neighbor, then it must also export every other path Q that has equal or higher LocalPref than path R to that neighbor.

The NH assumption has been leveraged to design BGP debugging tools [21], and many simulation studies [3, 13, 16, 29, 41] use models that are a special case of NH, CE and GR preference and export (Section 4.1).[1]

[1] The GR preference condition does not imply Next-Hop LocalPref. This is because while GR preference requires an AS to base its LocalPref on its business relationship with the

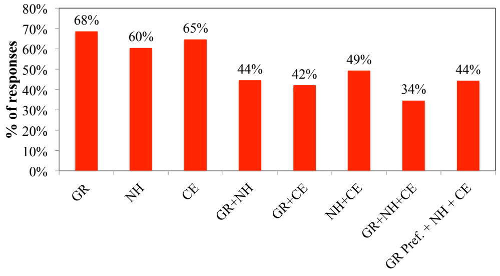

Figure 2: Percent of respondents that follow combinations of Gao-Rexford (GR) (Q3,Q9), Next-hop (NH) (Q4), and Consistent Export (CE) (Q10). Normalized by number of respondents.

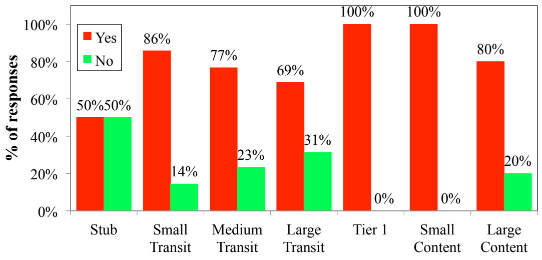

Figure 3: Q5: Do you use the same LocalPref configuration across all BGP-speaking routers in your network? Normalized by the number of respondents to Q5. (2 stubs did not respond).

(In fact, the literature on incentive compatibility of BGP [8, 9, 15] considers an even more extreme version of next-hop policies that also requires ASes to disable step 2 of the BGP decision process (Table 2), preventing an AS from using AS-path length to rank routes. [36] even makes the case that the routing system should exclusively use consistent export and (the more extreme version of) next hop policies!)

Figure 2 shows the percent of survey respondents who indicate that they adhere to the Gao-Rexford (GR), Next-hop (NH, Q4 (Table 1)), or Consistent Export (CE, Q10) criteria. Here we only include networks that transit traffic (small/medium/large transit providers and tier 1 networks), and we normalize by the number of respondents that answered a given question (or set of questions). While this reduces our sample size slightly, we still have a considerable number (61) transit providers that answered all four relevant questions (Questions 3,4,9,10). Two thirds of these respondents use at least one of GR, NH, or CE with half using *both* NH and CE. Many operators indicated that they use next-hop routing to avoid *"policies [becoming] tedious and complicated to manage"*. By contrast, one operator commented that he/she does not use NH because if *"a downstream path is bad (congestion, etc), then you modify it accordingly."*

We found that 34% of transit providers indicate that they use GR, NH, and CE. This grows to 44% if we consider networks that have at least implemented GR preference (but not necessarily GR export). The wording of Q4, however, could lead to *underestimates* of the number of respondents that follow NH. Specifically, one operator indicated that he/she violates NH because *"[i]n most cases, LocalPref is based on prefix-lists ... rather than AS paths"* (*i.e.,* LocalPref is based only on destination IP prefix), which does not actually violate the definition of NH.

4.3 Prefer oldest path

The seventh step in the BGP decision process of Table 2 is often ignored in models of BGP dynamics. According to this step, after applying the first six steps in the BGP decision process, the AS should prefer the route it learned first, *i.e.,* the "oldest path". This seemly innocuous step means that routing outcomes are no longer memoryless/deterministic, an assumption that shaped Griffin, Sheppard, and Wilfong's seminal model of BGP dynamics [17] and many subsequent works. Specifically, with "prefer oldest path", the route selected is not just a function of the AS's routing policy and the set of available paths, but is also affected by history.

76% of respondents answering Q6 indicated that they have enabled this step. While one respondent indicated that he/she disabled this step and was *"using router-id tie breaker instead for deterministic outcome"*, another respondent wrote that the step was enabled for *"Perceived stability improvements, no driver to change from default."* Another claimed that *"Oldest path = Stable path. Who cares if the router-id is better? I'll gladly take the route that isn't flapping."*

While some network operators associate the "prefer oldest path" step with improved stability, the implications of this step for BGP routing stability is poorly understood (note: this is different from route-flap damping [30]). Understanding the implications on BGP dynamics (and stability) calls for rethinking established assumptions on determinism.

4.4 Consistency within an AS

While studies of interdomain routing often treat an AS as an atomic entity [3, 4, 13, 16, 31], in practice, an AS consists of many (often geographically distributed) routers that must be configured individually. Q5 was designed to understand discrepancies in LocalPref configuration across routers within a single AS. Figure 3 shows the responses to Q5. The majority of respondents (77% of those that answered Q5) indicated that they use the same LocalPref configuration across all routers; many indicated that this is done for *"[s]implicity, consistency"*.

Stub networks are most likely to violate this assumption. Of the 10 respondents (to Q5) that operate stub ASes, 5 of them indicated that they do not use use the same LocalPref configuration across all routers; however, one respondent indicated that this is because *"we don't use Local Pref"*, which could explain this observation.

While all 4 Tier 1 respondents indicated that they use a consistent LocalPref, there is evidence of exceptions to consistency at other networks, with 5 large transit, 7 medium transit, 4 small transit, and 1 large content provider indicating that they do *not* use consistent LocalPref. One operator explained that this is the result of *"[r]egionalisation"* and another said *"As far as weights for customers, yes [to consistency]. For peers/providers, no. We have various provider link bandwidths and costs to deal with, so some routing policy is dictated by those constraints."* Another explanation was *"[s]ome route paths require manual intervention as other factors dictate WHEN we want to route differently, other than simple path availability."*

next-hop AS on the AS path, nothing prevents an AS from basing its routing decisions on distant ASes along the AS path as well (*e.g.,* by prioritizing customer paths that do not traverse a distant, undesirable AS over customer paths that do traverse that AS).

5. MRAI

The Min Route Advertisement Interval (MRAI) timer limits the rate of update messages between BGP neighbors. Configuring the MRAI timer is a trade off between (a) minimizing the number of BGP update messages with (b) the need to react quickly to routing changes. Early BGP RFCs [33] recommend a default minimum MRAI time interval of 30 seconds. However, both router vendors and the standards community are moving towards much lower time intervals, and even disabling MRAI entirely [20, 23]. Q11 asked network operators about MRAI timer values in their networks. Interestingly, most respondents, 75 out of 83, do not use MRAI timers at all. Only 6 responses stated that they use (the Cisco default) MRAI timer value of 30 seconds.

One operator explained: *“There is no point in sitting on an update. Push it out, even if it creates additional churn in terms of updates.”* Others explained that MRAI *“just makes a mess”* and another commented *“One of our upstreams implements “default Cisco” MRAI timer values, I wouldn’t wish this on my competitors, let alone my customers!”*

6. INCORPORATING SECURITY.

Q7 considered the (hypothetical) incorporation of information from a secure routing protocol, BGPSEC [26], into routing policies. BGPSEC is a protocol that the IETF is standardizing to validate BGP paths. BGPSEC allows an AS to validate if an AS-level path announced by its neighbor was actually announced by all the ASes on that path.

For BGPSEC to limit attacks on routing, path validity information must influence path selection; thus, information provided by BGPSEC must be incorporated into the BGP decision process. The BGPSEC standard provides flexibility in how this is done. Instead of specifying how to incorporate this information in the BGP decision process, the RFCs indicate that this “is a matter of local policy” [26].

To that end, Q7 asked operators at what step in the BGP decision process they would incorporate information about route validity. Results are shown in Figure 4. The results indicate that only 9% of respondents would prioritize security first, above all other considerations, and the most popular response (21%) was to place security between the LocalPref and AS path step. One operator explained this as follows: *“LP [LocalPref] is still king for policy decisions. AS-Path length isn’t all that relevant these days. I’m going to kick up the LP of as-paths that I care about.”*

However, a large fraction of operators place security considerations much lower in the BGP decision process, with 40% indicating that it would be placed below AS-path length (step 2 in Table 2). One operator explained: *“... definitely after LocalPref. Placing this before AS Path would require a lot more analysis before I’d be comfortable, and we’re currently honoring MEDs on select peers, so initially at least it would go after all of those.”* Another explained his ambivalence for prioritizing route validity by stating that *“until everyone uses a validation mechanism, there’s no point in disrupting policies with that”* and another said *“not willing to affect routing policy or adversely affect performance.”’*

There was also a large number of operators that declined to state where they would rank BGPSEC in the decision process. Some felt that their existing security solutions were adequate stating *“we use static prefix lists on customers instead. no operational interest or vendor support for bgpsec.”*

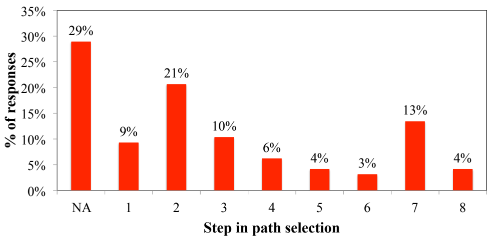

Figure 4: Responses to Q7: If path validation (*e.g.,* BGPSec) was deployed in your network, before what step (1-8) in the table would you place the following step: “Prefer secure paths (validated paths) over insecure paths”?

While others felt that BGPSEC state should be included in existing steps: *“needs to be configurable, BGPsec state should influence metrics like localpref”*.

These observations have implications of the efficacy of BGPSEC; recently [29] showed that if operators are unwilling to place security *first*, above all other steps in BGP decisions, then BGPSEC provides only limited protection.

7. BILLING MODELS

Our final question concerns transit pricing models. We were interested in whether ASes can increase revenues by increasing the volume of customer traffic they attract, because this assumption played a key role in our prior work on the economics of deploying BGPSEC [13]. According to Stanojevic *et al.* [38], 95/5-percentile pricing is a volume-based billing method that is “the most prevalent method that transit ISPs use for charging their customers.” With 95/5-percentile billing, a billing cycle is split in constant-size intervals and number of bytes transferred by the customer in each interval is recorded; the customer is then billed for the 95th-percentile of the distribution of recorded intervals. Answers to Q12 (Table 1) indicate that 95/5 percentile billing is common, with 55% of survey respondents at non-stub ISPs reporting that they use it because it is the *“industry standard”* and *“purely historical. It has always been done so no one wants to change it.”* The responsents support Stanojevic *et al.*’s [38] claim, but also indicate that there are exceptions where alternate billing models are used. (See *e.g.,* Valancius *et al.* [39].) One alternate volume-based billing scheme was to *“... price based on destination. If we don’t have to carry a customers traffic across the US or Atalantic or Pacific they can get cheaper pricing”*.

We find that 95/5 billing is used by 3 out 4 of respondents from Tier 1 networks, 56-60% of medium to large transit providers, and 46% of small transit providers. Reasons smaller transit providers did not use 95/5 percentile billing included serving limited/specialized populations, *e.g.,* being a *“[n]on profit [with] no bills”* and *“[not having] customers that are burstable. If we did, we would.”*

8. CONCLUSION AND FUTURE WORK

Models of BGP routing are intended to help us reason about today’s vast and complex interdomain routing system. Hence, good models of routing policies should be both simple and yet expressive enough to distill crucial aspects of real-life policies. Our survey sheds light on routing policies used in practice and on the extent to which common modeling assumptions actually hold on the Internet. While we find that standard modeling assumptions (e.g., the Gao-Rexford path preference policies) are well-grounded in reality, some of the responses beg for closer scrutiny (*e.g.,* using empirical measurements) and motivate new routing models.

While our survey provides a useful starting point, it also calls for more rigorous study of routing policies. Indeed, like any survey, our results suffer from biases (*e.g.,* larger networks are over represented) and "noise" resulting from misunderstandings of terminology. (Indeed, some comments we collected indicated that operators interpreted the term "peer" to mean "neighbor", rather than "settlement-free peer" as we had intended. Similarly, we had to discard the results of Q8 because our use of the term "neighbor-specific" (see [40]) was not understood.) Further exploration of the operational issues highlighted by this paper through targetted and larger-scale surveys is of value. More extensive empirical analyses and the development of new routing models are also valuable directions for future research.

Acknowledgments

We thank the network operators who completed our survey. This work was supported by a gift from Cisco, NSF Grant S-1017907, ISF grant 420/12, a Marie Curie CIG and I-CORE.

9. REFERENCES

[1] University of Oregon Route Views Project. `http://www.routeviews.org/`.

[2] I. Avramopoulos, M. Suchara, and J. Rexford. How small groups can secure interdomain routing. Technical report, Princeton University Comp. Sci., 2007.

[3] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *SIGCOMM*, 2007.

[4] H. Chang, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocol. In *SIGCOMM*, 2006.

[5] Y.-J. Chi, R. Oliveira, and L. Zhang. Cyclops: The Internet AS-level observatory. *SIGCOMM CCR*, 2008.

[6] Cisco. BGP best path selection algorithm: How the best path algorithm works. Document ID: 13753, May 2012. `http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml#bestpath`.

[7] Cisco. How BGP routers use the multi-exit discriminator for best path selection. Document ID: 13759, March 2012. `http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094934.shtml`.

[8] J. Feigenbaum, V. Ramachandran, and M. Schapira. Incentive-compatible interdomain routing. In *Conference on Electronic Commerce*, pages 130 – 139, 2006.

[9] J. Feigenbaum, M. Schapira, and S. Shenker. *Algorithmic Game Theory*, chapter Distributed Algorithmic Mechanism Design. Cambridge University Press, 2007.

[10] T. Flach, E. Katz-Bassett, and R. Govindan. Quantifying violations of destination-based forwarding on the internet. IMC '12, 2012.

[11] L. Gao, T. Griffin, and J. Rexford. Inherently safe backup routing with BGP. *IEEE INFOCOM*, 2001.

[12] L. Gao and J. Rexford. Stable Internet routing without global coordination. *Trans. Netw.*, 2001.

[13] P. Gill, M. Schapira, and S. Goldberg. Let the market drive deployment: A strategy for transistioning to BGP security. *SIGCOMM'11*, 2011.

[14] V. Giotsas and S. Zhou. Valley-free violation in Internet routing - analysis based on BGP community data. IEEE ICC, 2012.

[15] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright. Rationality and traffic attraction: Incentives for honest path announcements in BGP. In *SIGCOMM'08*, 2008.

[16] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols? In *SIGCOMM'10*, 2010.

[17] T. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *Trans. Netw.*, 2002.

[18] G. Huston. Peering and settlements - Part I. *The Internet Protocol Journal (Cisco)*, 2(1), March 1999.

[19] G. Huston. Peering and settlements - Part II. *The Internet Protocol Journal (Cisco)*, 2(2), June 1999.

[20] P. Jakma. Revisions to the BGP 'Minimum Route Advertisement Interval'. Internet-Draft: `http://tools.ietf.org/html/draft-ietf-idr-mrai-dep-04`, 2012.

[21] U. Javed, I. Cunha, D. R. Choffnes, E. Katz-Bassett, T. Anderson, and A. Krishnamurthy. Poiroot: Investigating the root cause of interdomain path changes. In *SIGCOMM*, 2013.

[22] J. John, E. Katz-Bassett, A. Krishnamurthy, T. Anderson, and A. Venkataramani. Consensus routing: The Internet as a distributed system. In *NSDI*, 2008.

[23] Juniper. `https://www.juniper.net/techpubs/software/junos/junos57/swconfig57-routing/html/bgp-summary32.htm`.

[24] Juniper. Selecting the best path, 2005. `http://www.juniper.net/techpubs/software/erx/erx51x/swconfig-routing-vol2/html/bgp-config10.html`.

[25] J. Karlin, S. Forrest, and J. Rexford. Nation-state routing: Censorship, wiretapping, and BGP. *CoRR*, 2009.

[26] M. Lepinski. BGPSEC protocol specification: draft-ietf-sidr-bgpsec-protocol-06. Internet-Draft, 2012.

[27] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *ACM STOC*, May 2008.

[28] M. Luckie, B. Huffaker, A. Dhamdhere, and V. Giotsas. AS relationships, customer cones, and validation. 2013.

[29] R. Lychev, S. Goldberg, and M. Schapira. Is the juice worth the squeeze? BGP security in partial deployment. In *SIGCOMM'13*, 2013.

[30] Z. Mao, R. Govindan, G. Varghese, and R. Katz. Route flap damping exacerbates Internet routing convergence. In *SIGCOMM CCR.*, 2002.

[31] D. Pei, M. Azuma, D. Massey, and L. Zhang. BGP-RCN: Improving BGP convergence through root cause notification. *Computer Networks*, 48(2):175–194, 2005.

[32] S. Y. Qiu, P. Mcdaniel, and F. Monrose. Toward valley-free interdomain routing. IEEE ICC, 2007.

[33] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 BGP-4. RFC 4271, January 2006.

[34] RIPE Network Coordination Center. RIPE Routing Information Service. http://www.ripe.net/data-tools/stats/ris/routing-information-service.

[35] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 lessons from 10 years of measuring and modeling the Internet's autonomous systems. *JSAC*, 2011.

[36] M. Schapira, Y. Zhu, and J. Rexford. Putting BGP on the right path: a case for next-hop routing. In *HotNets*, 2010.

[37] N. Spring, R. Mahajan, and T. Anderson. The causes of path inflation. SIGCOMM, 2003.

[38] R. Stanojevic, N. Laoutaris, and P. Rodriguez. On economic heavy hitters: Shapley value analysis of 95th-percentile pricing. IMC '10, pages 75–80, 2010.

[39] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani. How many tiers?: Pricing in the Internet transit market. In *ACM SIGCOMM*, 2011.

[40] Y. Wang, M. Schapira, and J. Rexford. Neighbor-specific BGP: more flexible routing policies while improving global stability. In *SIGMETRICS'09*, pages 217–228. ACM, 2009.

[41] J. Wu, Y. Zhang, Z. M. Mao, and K. Shin. Internet routing resilience to failures: Analysis and implications. In *CoNEXT*, 2007.