"I Can't Get No Satisfaction": Helping Autonomous Systems Identify Their Unsatisfied Inter-domain Interests

Juan Camilo Cardona, Stefano Vissicchio, Paolo Lucente, and Pierre Francois

***Abstract*—Given the distributed and business-driven nature of the Internet, economic interests of Autonomous Systems (ASes) may be incompatible. Previous works studied specific effects of incompatible interests, especially BGP policy conflicts leading to routing and forwarding anomalies. In this paper, we rather focus on the effects of incompatible interests that do not trigger such anomalies. We take the perspective of a single AS: We show that incompatible interests can have a tangible impact on its business, and provide a classification of its unsatisfied interests. Since incompatible interests cannot be solved automatically, our effort is directed to support network managers in their business decisions. Hence, we describe algorithms to identify and assess their impact, as well as a prototype of a warning system aimed at signaling the most relevant unsatisfied interests. We evaluate our prototype on real data from two operational networks. In addition to illustrate the potential of our system, our evaluation shows that unsatisfied interest are relatively frequent and likely affect a significant amount of traffic in practice.**

***Index Terms*—Unsatisfied Interests, Inter-domain Routing, Network Management, BGP.**

I. Introduction

Internet routing is business driven. Indeed, each domain (called Autonomous System or AS) is independently managed by an administrative entity, with its own economic *interests*. To pursue their interests, network operators configure routing policies. In BGP, the de-facto standard protocol for inter-domain routing, policies express route preferences (e.g., prefer a route from a given neighboring AS) and filters (e.g., do not propagate routes from one specific neighbor to another), hence controlling ingress and egress points for Internet traffic.

Due to the nature of inter-domain routing and the lack of global coordination, interests of different ASes may be incompatible. Consider, for instance, Fig. 1. In this example, $AS4$ has an interest to receive incoming traffic destined to its prefix $1/8$ from $AS2$. However, $AS3$ prefers to send traffic directly to $AS4$, and $AS1$ favors the link to $AS3$ to forward this traffic. Unfortunately, those interests are actually incompatible, that is, no valid distribution of traffic realizes the interests of all the involved ASes. In this scenario, depending on the specific policies configured by $AS4$, $AS3$ and $AS1$, we have three possible cases, detailed in Figure 2. In the first case (Figure 2a), $AS1$ forwards traffic to $AS2$. This configuration realizes the interests of $AS4$ but neither those of $AS1$ nor of $AS3$. In the second case (Figure 2b), $AS1$ sends traffic to $AS3$, that is forced to forward to $AS2$, thus sacrificing its own interests. In the last case (Figure 2c), both $AS1$ and $AS3$ realize their respective interests, at the expenses of $AS4$.

Juan Camilo Cardona is with Cisco Systems / UC3M / IMDEA Networks.
Stefano Vissicchio is with the Universite Catholique de Louvain.
Paolo Lucente and Pierre Francois are with Cisco Systems.

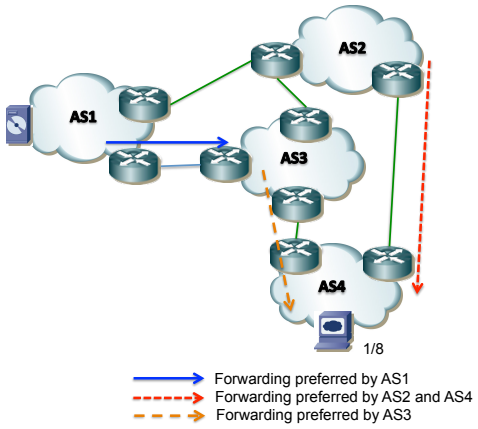

Fig. 1. Example of incompatible interests.

Incompatible interests can and do happen in the Internet. For example, the interests shown in Figure 1 are realistic if (i) AS4 is a customer of AS3 and AS2, (ii) AS1, AS3, and AS4 are all customers of AS2, and (iii) AS1 and AS3 are settlement-free peers. In this case, Figure 2 can reflect policies and router configurations that are common in operational networks [1][2].

In BGP, incompatible interests can trigger well-studied anomalies. Indeed, they may result in so-called policy disputes that can translate into routing (i.e., control-plane instabilities) and forwarding (i.e., inter-domain loops) anomalies [3]. Policy disputes have been the target of numerous research efforts, covering the full range between theoretical (e.g., [4], [5], [6]) and practical (e.g., [7], [8]) contributions.

Nevertheless, much less work has been done on incompatible interests that do not trigger routing and forwarding anomalies (as any configuration in Fig. 2). Nonetheless, even those incompatible interests are relevant from an operational perspective, since it directly affects ASes business. For example, in Fig. 2a, $AS1$ is forced to pay for the traffic forwarded to $AS2$, while its (unsatisfied) interest to send traffic to $AS3$ would have led to no expenses. Similar economic losses occur for $AS3$ and $AS4$ in Fig. 2b and 2c, respectively.

In this paper, we complement the state of the art by focusing on incompatible interests that do not lead to routing or forwarding anomalies. We take the perspective of a single AS, and we study local effects of globally-incompatible interests. We abstract those effects in the concept of *unsatisfied interests*,

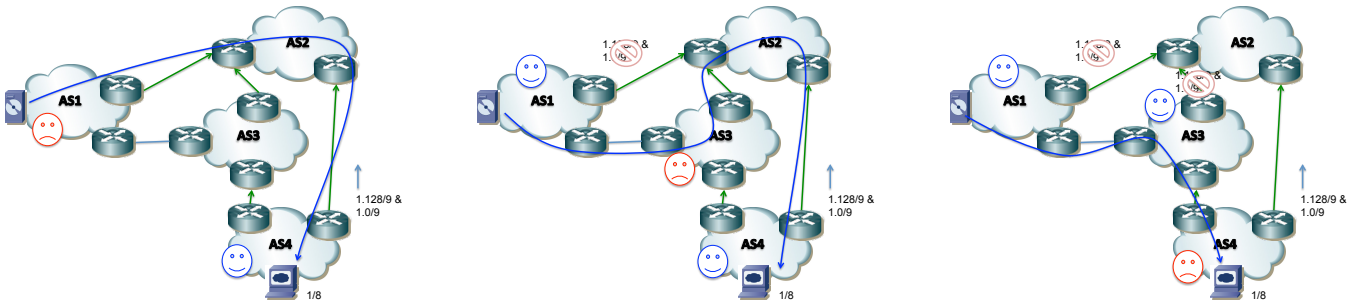

(a) Traffic distribution that dissatisfies the interests of $AS1$ (e.g., resulting from $AS4$ announcing prefixes $1.0/9$ and $1.128/9$ to $AS2$)

(b) Traffic distribution that dissatisfies the interests of $AS3$ (e.g., resulting from $AS1$ filtering prefixes $1.0/9$ and $1.128/9$)

(c) Traffic distribution that dissatisfies the interests of $AS4$ (e.g., resulting from both $AS1$ and $AS3$ filtering prefixes $1.0/9$ and $1.128/9$)

Fig. 2. Enumeration of possible traffic distributions for the network in Figure 1: Every of those configurations forces an unsatisfied interest at some AS.

i.e., stable routing states where traffic is delivered to the destination but without respecting the given AS interests. We show several cases where unsatisfied interest can theoretically and practically lead to *significant economic losses* for individual ASes. Since globally-incompatible interests are (by definition) impossible to resolve automatically, we address the operational need for individual ASes to timely detect, understand, and assess unsatisfied interests. This is not an easy task, since it requires (i) a deep understanding of the BGP routing system, (ii) consideration of the interplay between different ASes (and their respective interests), (iii) integration of routing and forwarding data from different sources, and (iv) effective implementation, to automate data analyses and present the most relevant data to human operators. As a result, manual inspection is not a feasible approach, and even current commercial and research tools do not support unsatisfied interest analyses, to the best of our knowledge.

In order to fill this gap, we make several contributions.

First, we provide the first exhaustive classification of unsatisfied interest affecting desirable routing and forwarding states. Namely, we distinguish between unsatisfied interests affecting outbound and inbound traffic of any given AS X. The former category is related to the BGP routes that neighboring ASes send to X. For example, if a neighbor advertises a de-aggregated prefix p on a specific inter-domain link with X, this will attract all traffic from X for p to that link, independently of the policies (and the interests) of X. Conversely, dissatisfactions at inbound traffic are reflected in policies applied by X's neighbors on routes announced by X. For instance, if a neighbor of X filters some X's routes, this can change the ingress points of traffic traversing X.

Second, we develop new algorithms to (i) automatically detect outbound and inbound unsatisfied interests; and (ii) measure their impact. Our algorithms leverage our domain-specific knowledge of the Internet routing system and the variety of data sources on the Internet traffic typically available to network operators. In particular, they use different data to detect distinct types of unsatisfied interests, and combine multiple BGP views (internal and external to the given AS) and traffic data. Our algorithms quantify the impact of unsatisfied interests in terms of affected traffic volumes. However, they can be easily modified for custom analyses (e.g., security ones based on traffic exchanged with security-sensible prefixes).

Third, we design and prototype a warning system, that builds upon our algorithms and raises alarms for the most critical unsatisfied interests. Our system is meant to support network managers in their strategic business decisions, like changes of commercial agreements with other ASes (e.g., cost-model adjustments) or selection of service providers with more aligned interests. Since unsatisfied interests may coincide with unfulfilled peering contracts, our system also offers technical support for verification of commercial agreements.

Fourth, we use our system to conduct a deep measurement study on unsatisfied interests in real-world ASes. Our measurement campaign and its validation demonstrate (i) the feasibility of our system; (ii) the effectiveness of our algorithms to detect business-affecting unsatisfied interests; and (iii) the high frequency and impact (e.g., in terms of traffic volume) of unsatisfied interests in the real Internet.

The rest of the paper is structured as follows. Sec. II introduces BGP and routing policies. Sec. III classifies unsatisfied interests. Sec. IV presents our algorithms to detect them. Sec. V describes our warning system. Sec. VI details the results of our measurements. Sec. VII discusses related works, and Sec. VIII concludes the paper.

II. Background

BGP [9] is the standard protocol used for inter-domain routing. In a nutshell, it allows ASes to exchange routing information about prefixes reachable through their interconnection. Through BGP, an AS can inform the others on which prefixes are reachable through it, and some characteristics of the corresponding paths. Path characteristics are denominated path attributes, and include information like the sequence of ASes on the way to the origin, the preference of the local AS for each path (Local-preference), or the preference of the external AS to specific links (MED). BGP defines the best path algorithm (Table I), where routers select the best routes from the ones available, based on their path attributes.

Information propagation depends on the BGP peerings. ASes establish peerings between them on the basis of commercial agreements. The basic types of agreements are *transit-customer* and *settlement-free peerings*. In the first, an AS (transit) transports the traffic of the other AS (customer) to and from any other network, in exchange of a fee. In the second case, both ASes agree to share the costs of the interconnection

Step	Criterion
1	Prefer path with highest Local preference.
2	Prefer path originated by local router.
3	Prefer path with shorter AS-path length.
4	Prefer path with lowest origin code.
5	Prefer path with lower MED (For paths from the same neighboring AS)
6	Prefer EBGP to IBGP.
7	Prefer path with closest next-hop.
8	Prefer oldest path, if EBGP.
9	Prefer path in which the Router ID of NH is lowest.

TABLE I
BGP BEST PATH ALGORITHM [9].

while promising connectivity to their own networks, and the ones of their customers. Other business agreements, such as paid-peering or partial-transit, are also established in practice [10]. Except for a Tier-1 networks, ASes usually connect to more than one transit provider, and establish settlement-free peerings with multiple neighbors [11].

One of the tasks of network operators is to manage the inter-domain traffic traversing their AS, that is, to implement their interests via *policies*. Networks are indeed connected through several links, each with different characteristics in terms of cost or quality. Policies are targeted not only to avoid of link congestion, but also to reduce expenses and deliver required performance. Because of the distributed nature of the Internet, setting policies is hard for operators. Indeed, operators have limited control and knowledge over policies and forwarding behavior of other ASes. Moreover, to fulfill their interests over time, they may be required to continuously adjust implemented policies upon the various events affecting inter-domain traffic, like failures, traffic fluctuations or policy changes.

Different techniques are typically used to set up effective outbound and inbound policies. For outbound traffic, operators have control on the preference among routes that they receive from external networks, but cannot force neighbors to propagate specific routes. Even worse, for inbound traffic, operators can only try to influence the decisions of external ASes, by carefully setting attributes in their BGP announcements. We now briefly expand on common techniques used for policies affecting outbound and inbound traffic, respectively.

Outbound policies drive the selection of paths that internal routers should use to forward traffic to external destinations. This is normally achieved by tweaking the attributes of the incoming routes to give priority to the ones more aligned with the interests of the local AS. To this end, the local preference value is commonly changed [12][13]. Some operators also use MED tweaking for this purpose, although it was initially designed for inbound TE. Other strategies rely on special communities to achieve more granular control [14]. In Figure 2a, for example, AS4 can try to balance its outbound traffic over its two links (AS4-AS3 and AS4-AS2), by assigning a higher local preference for certain prefixes on one link.

With **Inbound policies**, operators try to *influence* the routing decisions of external ASes, in order to obtain their desired inbound traffic distribution. However, the fact that each AS selects preferred paths based on its own policies often leads to a trial-and-error process with no guarantees of success [15]. Operators typically use AS-prepending or prefix deaggregation to influence path selection of others ASes [16][17][15]. They can also try to tune MED or append pre-arranged communities[18]. In Figure 2a, AS4 inbound policy consists in announcing more-specific prefixes (1.128/9 and 1.0/9) through only one of its two transit ASes. Note that this policy is effective (fulfilling AS4's interests) in Figure 2a, but it is not in Figures 2b and 2c.

III. CLASSIFICATION OF UNSATISFIED INTERESTS

In this section, we study unsatisfied interests of single ASes and their impact. We define AS *interests* in terms of profit and costs as established by commercial agreements. This definition reflects long-term policies set by the given AS and their effectiveness in the stable routing state. The analysis in this section can however be extended to shorter-term interests by including the consideration of transient network conditions in the definition of interests. For example, interests may encompass specific neighbor preferences (prefer A over B) – or no preferences at all – for given flows, in the presence of given failures (C is not available) or congestion.

We classify unsatisfied interest into outbound and inbound, depending on the traffic that they affect, and describe realistic examples for each class. Our examples also show that (i) our classification covers all types (inbound, outbound and transit) of traffic traversing the considered AS; and (ii) unsatisfied interests may be due to various economic reasons, and be realized via different technical means.

A. Outbound unsatisfied interests

We define outbound unsatisfied interests, or *outbound dissatisfactions*, as follows. An AS X suffers from an outbound dissatisfaction if X is prevented from *sending* some traffic flows through an intended inter-domain link. That is, BGP forces the traffic for given destinations to exit X via an inter-domain link l_1 while X has interest to use another link $l_2 \neq l_1$.

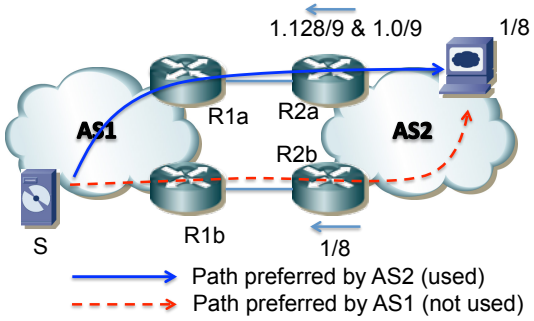

Fig. 3. Example of an outbound unsatisfied interest for $AS1$.

Figure 3 shows a simple scenario in which $AS1$ is affected by an outbound dissatisfaction, as a result of incompatible interests of $AS1$ and $AS2$. The two ASes are connected with two distinct physical links (e.g., in different locations). However, they disagree on which inter-domain link should be used for the traffic from the source S to the destination prefix $1/8$. The dashed (red) and the solid (blue) arrows respectively indicate that $AS1$ would like to forward such traffic through $R1b$, while $AS2$ prefers to receive this traffic

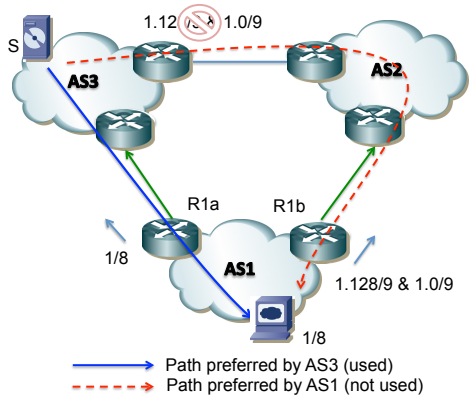

Fig. 4. Example of an inbound unsatisfied interest for $AS1$.

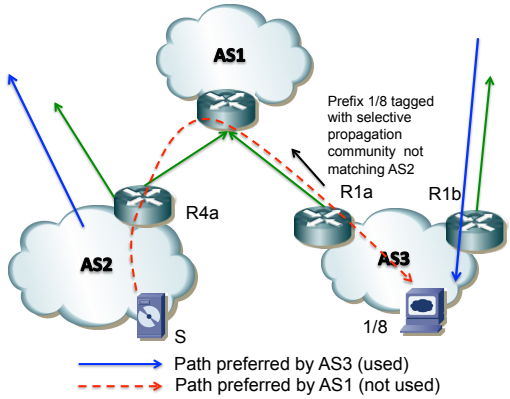

Fig. 5. Example of an inbound dissatisfaction affecting transit traffic of $AS1$.

at $R2a$. Disagreements like the one in Figure 3 realistically happen in the Internet [19]. For instance, they can be the result of the adoption of the *hot-potato* policy by both $AS1$ and $AS2$. With this policy, ASes try to reduce the internal path followed by Internet traffic (e.g., for minimizing resource utilization). Hence, Figure 3 can easily occur if S is geographically closer to $R1b$ and the machines hosting prefix $1/8$ are closer to $R2a$.

In the example, the inbound policy of $AS2$ actually prevails, and $AS1$ is forced to forward outbound traffic against its economic interests. This unsatisfied interest is due to $AS2$ selectively announcing paths to more specific prefixes, on the $(R1a, R2a)$ link. Note that $AS2$ has other ways to enforce its interest, e.g., it can also set different BGP attributes (e.g., AS-path or MED) in the announcements that it propagates on the two inter-domain links (see Section II). All those cases can be categorized as inconsistent advertisements, and are traditionally considered a bad practice in private peerings [20][19], since they typically violate peering contracts. However, inconsistent advertisements do not always translate into contractual violations, e.g., in Internet eXchange Points (IXPs) where the peering ecosystem has become more informal with the proliferation of route servers [21], [22]. Moreover, outbound dissatisfactions are not always originated by inconsistent advertisements. For instance, $AS1$ in Figure 2a suffers from an outbound dissatisfaction due to $AS4$ sending consistent advertisements to both $AS2$ and $AS3$, and selectively announcing more specific prefixes to $AS2$ only.

Finally, observe that outbound dissatisfactions do not necessarily impact outbound traffic only. Indeed, they can also affect transit traffic, that the considered AS did not originate but has to transfer from one neighbor to another. For example, in Figure 2b, $AS3$ suffers from an outbound dissatisfaction that impacts the transit traffic from $AS1$ to $AS4$.

B. Inbound unsatisfied interests

We define inbound unsatisfied interests, or *inbound dissatisfactions*, as complementary to the outbound ones. In particular, we say that an AS X suffers from an inbound dissatisfaction if X is prevented from *receiving* certain traffic over the preferred inter-domain link. That is, BGP forces the traffic to a given destination to enter X from an inter-domain link l_1 while X has interest to receive it over $l_2 \neq l_1$.

An example of inbound unsatisfied interests is displayed in Figure 4. As in the previous example, we take the perspective of $AS1$. For traffic destined to its prefix $1/8$, $AS1$ has an economic interest in receiving it at $R1b$ (dashed red path). However, this clashes with $AS3$ interests to send such traffic directly to $AS1$ (solid blue path in the figure). As the previous one, this example is also realistic. On one hand, $AS1$ interest can depend on the need to balance incoming traffic between its two border routers, i.e., depending on the destination prefix. On the other hand, $AS3$ interests may be due to its commercial agreements with $AS1$ and $AS2$, especially if $AS3$ is a service provider of $AS1$ (getting money for the traffic exchanged on their direct inter-domain link) and is a settlement-free peer of $AS2$ (with free of charge traffic exchange agreement).

In Figure 4, the direct link between $AS3$ and $AS1$ is eventually selected, triggering an inbound dissatisfaction at $AS1$. Indeed, $AS3$ filters more-specific prefix advertisements from $AS2$, hence uses its direct route for $1/8$ received by $AS1$. While implicitly ignoring $AS1$ interests, $AS3$ policy may be compliant with its contractual obligations, since it may not be forced to consider all announcements of $AS1$'s prefixes.

Finally, note that transit traffic can also be impacted by inbound dissatisfactions. Consider the example depicted in Figure 5. $AS3$ has interest to receive the traffic from source S in $AS2$ to a given destination $1/8$ using router $R1b$. Nevertheless, $AS1$ may have economic benefits in forwarding the traffic from $AS2$ to $AS3$, e.g., if both $AS2$ and $AS3$ are customers paying for transit through $AS1$. In the example, $AS3$ indicates into the BGP announcement to $AS1$ (e.g., with pre-agreed communities [23]) that $AS1$ should not propagate the announcement to $AS2$. Since $AS1$ may be forced by contractual agreements to respect such an indication, it ends up not propagating to $AS2$ the announcement from $AS3$, which results into an incoming unsatisfied interest at $AS1$.

IV. Detection of unsatisfied interests

In this section, we propose algorithms to detect unsatisfied interest and assess their economical and technical impact. We designed those algorithms to have the following features.

Our algorithms exploit peculiarities of inbound and outbound dissatisfactions. Outbound and inbound dissatisfactions depend on different aspects of the BGP configuration

(e.g., received routes vs. neighbor policies, see §III). We provide distinct algorithms to detect each of the two classes. For the outbound case, we only rely on control-plane information: we indeed compare the received routes with the expected ones to detect outbound unsatisfied interest, infer their cause, and estimate their impact by analyzing what-if scenarios. For the inbound case, since we cannot assume that policies of external ASes are known, we use data-plane information to detect unexpected ingress points for given flows.

Our algorithms estimate the impact of detected unsatisfied interests and pinpoint most practically-relevant ones. Conversations with network operators confirmed that the importance of unsatisfied interests depend on the affected traffic flows. In the following, we focus on the most generally-accepted impact metric, which is the affected traffic volume. Hence, we correlate the prefixes involved in the unsatisfied interests with the traffic destined to them. Further, we classify unsatisfied interests based on their qualitative impact (reduction of route diversity, occurrence upon failures, etc.).

Our algorithms can be customized according to specific needs of operators, with respect to both new unsatisfied interest types and impact estimation. For example, the algorithms can be easily extended to include more or different impact metrics for unsatisfied interest, or restrict to destinations (like the most popular or security-sensitive ones). In addition, while we focus on the standard BGP implementation [9], our algorithms can be slightly modified to take into account the interaction of different routing and specific implementation or configuration features. For example, they can be easily adapted for configurations implementing policies in internal routing protocols (e.g., iBGP [24], [25]).

Our algorithms are correct provided that their input is truthful. We discuss in Section V how operators can collect the needed information to run the algorithms in real networks. Moreover, we provide a validation of the correctness of our algorithms under realistic assumptions in Section VI.

We separately detail algorithms for outbound and inbound unsatisfied interests in Section IV-A and IV-B, respectively.

A. Detection of Outbound unsatisfied interests

For a given AS, the exit points of inter-domain traffic depend on (i) the routes announced by neighboring ASes, (ii) the locally-preferred ones, and (iii) the intra-domain routing (i.e., IGP or iBGP) configuration. Operators have control of the latter two parameters, but not on the first one.

To detect outbound unsatisfied interests, we designed an algorithm that compares received routes with the set of missing ones, i.e., the routes expected to be announced by neighboring ASes but not announced by them. In this comparison, we automatically assess *whether* and *how much* traffic would differ if the missing paths were announced to the network. For instance, if a network detects that is not receiving a route from a settlement-free peer, the algorithm checks how the traffic of the prefix is currently being routed. If the traffic is currently being routed through a transit provider and its volume is significant, the algorithm would detect this case and rank it high. In the case of inconsistent advertisement (see for example Figure 3), the algorithm checks whether the inconsistency of the neighbor is reducing the next-hop diversity of the network, and rank it based on the outbound volume of these prefixes.

We now provide more details on the algorithm, which is summarized in Algorithm 1.

1) Input: The input consists of (1) the set of all current best routes; (2) the list of missing routes (expected but not received from neighboring ASes); and (3) statistics of outbound traffic per prefix. We discuss in Section V several methods that operators can use in real networks to collect these inputs.

2) Detection of unsatisfied interests: First, the algorithm computes how the traffic distribution of the network would improve if the missing paths where actually received. To this end, we classify every missing path into different categories, depending on its qualitative effect. A missing path could be in none or even in more than one category. Based on private conversations with operators, we identified four main effects of unsatisfied interest on outbound traffic, which we detail in the following.

- **Neighbor preference dissatisfactions**. A missing path is added to this category if its announcement would lead to the selection of a more preferred neighbor. In Algorithm 1, we use function $ASPreference$ to compare the preference of the operators among neighboring ASes. In simple configurations, this function simply checks for the Local Preference values of the routes, since operators tend to configure them consistently for each neighbor [13]. The function, however, can also cover per-session policies, depending on operator's input or router configurations.
- **Next-hop diversity dissatisfactions**. A missing path is added to this category if the exit point of the missing path is equally preferred with respect to the one of the current best paths, hence increasing the number of next-hops (NHs) for some outbound traffic. This may be operationally important, for instance, to balance traffic internally. Inconsistencies advertisements always fall in this category.
- **Back-up path dissatisfactions**. A missing path is added to this category if it comes from a neighbor more preferred than the one of any second best path, for destinations with a single active path. This category therefore covers cases in which a single link failure would let traffic be sent to less preferred neighbors.
- **Unexpected transit dissatisfactions**. Missing paths are added to this category if they are generating transit flows between two non-customer ASes (unexpected transit flows). This is a special case of *Neighbor preference dissatisfaction* paths, and is the problem experienced by AS3 in Figure 2b. Indeed, operators often try to avoid transporting traffic between non-customer neighboring ASes, as this does not provide any economical benefit [26]. To do this, they configure their routers not to advertise routes coming from non-customer neighbors to other non-customer neighbors. However, an operator might receive a route to a prefix *p* from a customer, which it propagates to neighboring ASes, while, it receives from non-customer neighbors a route for a prefix *p'*, more

input : **1.** Missing paths ($MissingPaths$)
2. Current BGP paths per prefix ($CurrentBestPath$ function),
3. Outbound Traffic per Prefix ($OutboundTrafficDemand$).
4. Preference of AS ($ASPreference$ function)
output: For each missing path returns the Impact type(s) (contained in $CurrentLabels$) and Impact level ($ImpactMetric$).

```
/* Go over each missing path and analyze its impact.
   Each Missing path is composed by an NLRI (NLRI) and
   a set of path attributes (MP).                        */
for (MP, NLRI) ∈ MissingPaths do
    /* Store the current best path (CBP) and backup
       path (CBaP) for NLRI.                             */
    CBP = CurrentBestPath(NLRI);
    CBaP = CurrentBackupPaths(NLRI);
    /* Calculate the best paths if the missing path
       (MP) were received for NLRI. Store this best
       path in NBP.                                      */
    NBP = BGPBestPathAlgorithm(CBP ∪ MP);
    /* 1) Detection of unsatisfied interests: Perform
       the classification tests and apply labels. The
       next part can be modified to fit the requirements
       of each operator.                                 */
    CurrentLabels = ∅;
    /* If the preference of any AS in NBP
       (NewPreference) is higher than the preference of
       the current path (CurrentPreference), apply label
       NeighbourPreferenceImprovement.                   */
    NewPreference = Max({ASPreference(AS) | AS ∈
    GetNeighboringASes(NBP)});
    CurrentPreference = Max({ASPreference(AS) | AS ∈
    GetNeighboringASes(CBP)});
    if NewPreference > CurrentPreference then
        CurrentLabels =
        CurrentLabels ∪ {NeighbourPreferenceImprovement};
    end
    /* If the current NHs (CurrentNHs) is a strict
       subset of the new NHs (NewNHs) under the
       missing paths, apply label IncreaseNHDiversity.
       */
    CurrentNHs = GetNHs(CBP);
    NewNHs = GetNHs(NBP);
    if NewNHs ⊋ CurrentNHs then
        CurrentLabels =
        CurrentLabels ∪ {IncreaseNHDiversity};
    end
    /* If there is currently a single active path for
       the prefix, and the missing path improves the
       preference of the back-up AS, apply label
       IncPrefofBKforSingleActivePath.                   */
    if |CBP| == 1 then
        NewBackupPaths =
        BGPBestPathAlgorithm(CBaP ∪ MP);
        CurrentBKPreference = Max({ASPreference(AS) |
        AS ∈ GetNeighboringASes(CBaP)});
        NewPreference = Max({ASPreference(AS) | AS ∈
        GetNeighboringASes(NewBackupPaths)});
        if NewPreference > CurrentBKPreference then
            CurrentLabels = CurrentLabels ∪
            {IncPrefofBKforSingleActivePath};
        end
    end
    /* If we find a path (CoveringP, CoveringNLRI), in
       which CoveringNLRI covers the NLRI, and if
       (CoveringP, CoveringNLRI) is propagated to
       other non-customer ASes, apply label
       UnexpectedTransit.                                */
    if (∃ (CoveringP, CoveringNLRI) which:
     CoveringNLRI Covers NLRI and
     IsPropagatedToNonCustomerneighbors(CoveringP) then
        CurrentLabels = CurrentLabels∪{UnexpectedTransit};
    end
    /*                                                   */
    /* 2) Impact assessment: Assess impact of interest
       conflict for the paths that match at least one
       classification.                                   */
    if CurrentLabels is not ∅ then
        ImpactMetric = OutboundTrafficDemand(NLRIP) ;
        Register (P, NLRI) with labels CurrentLabels
        and Impact ImpactMetric
    end
end
```

Algorithm 1: Detection of outbound unsatisfied interests.

specific than p. Since routers forward packets based on the more specific prefix (p'), the network might start transiting traffic between non-customer neighbors [1]. The missing routes from the customer towards the more-specific prefixes (p') are the ones added to this category.

3) Impact assessment: The second step of the algorithm is to measure the impact of each missing path. This value depends on the amount of outbound traffic and on its classification. We indeed map every dissatisfaction to the corresponding traffic volume as measured in the peak hour of the network. Operators could also employ more complex metrics such as bit-mile calculations, or metrics that estimate the potential revenue reduction due to the missing path. We decided to not implement such metric as setting its parameters is difficult to achieve from a researcher point of view.

Eventually, missing paths, their categories, and their impact value become descriptive features of the detected unsatisfied interests. These features can be used by an alarm system (see for example Section V) to highlight cases that should be analyzed individually by operators.

Note that, since the algorithm is based on re-simulating the BGP decision process, it always correctly provides outbound policy dissatisfactions, provided that the input is correct.

B. Detection of Inbound unsatisfied interests

The distribution of inbound inter-domain traffic into a network depends on the paths announced by the local AS and the policies implemented by the other ASes. The algorithm in this section aims at detecting neighboring ASes whose policies work against local inbound policies. Assume for example that the algorithm runs at $AS4$ in Figure 2c. It detects dissatisfactions for the traffic for prefixes 1.128/9 and 1.0/9 coming from $AS3$, since the resulting routing is in contrast with $AS4$'s interests.

Since operators rarely know the policies of external AS, it is almost unfeasible to detect the inbound dissatisfactions by using control plane data available from external looking glasses. Therefore, we rely only on local data plane information to detect inbound dissatisfactions. This type of test is simpler than the one in Section IV-A, but it also provides less information.

input : **1.** Inbound flow ($InboundFlow$) , with attributes $InboundFlowAttributes$ (containing attributes such as $SourceIP$, $DestinationIP$, Bw over the peak hour, etc.) arriving over link L.
2. Inbound policy contained in a function $IsFlowUndesired$.
output: Returns the inbound flows that are conflicting with the policy of the operator, together with their impact level ($ImpactMetric$).

```
/* For each inbound flows InF on each link L.           */
foreach Link L do
    foreach InboundFlow, with attributes InboundFlowAttributes
    (SourceIP, DestinationIP, BW, etc.) do
        if IsFlowUndesired(InboundFlowAttributes, L, Bw)
        returns True then
            Register InboundFlow, L, Bw;
        end
    end
end
```

Algorithm 2: Detection of inbound unsatisfied interests.

Algorithm 2 describes our algorithm. In a nutshell, it searches for inbound traffic that should be received on another

inter-domain link according to the policy of the local-AS operator. In the following, we provide more details on it.

1) Input: To identify the undesired traffic, we need to know its characteristics in terms of origin AS, origin prefix, or destination prefix. We therefore group traffic in *traffic flows*. On each flow, two input data are provided:

- **Inbound Inter-domain policy**, defined as a function $IsFlowUndesired$. Based on BGP attributes of inbound flows, this function defines the link (or links) through which the operator would like to enter her network. Operators can encode this function using automatic or manual methods. In simple cases, the function looks for source prefixes of ASes for which inbound traffic should not be detected. For example, operators usually do not expect traffic from customers or peering ASes in transit providers links, or traffic from ASes to which propagation is being remotely filtered through special communities [23]. More refined inbound policies may lead to more complex calculation, such as the bit-mile of the flow [27], in order to define whether the flow is undesired or not.
- **Inbound traffic statistics**, disaggregated on a per-flow basis. Specifically, operators should provide statistics of inbound traffic flows per prefix for individual ingress links of the network.

2) Undesired flow detection: The algorithm cycles over each inbound traffic flow. The function $IsFlowUndesired$ is then used to check whether the flow should be present at the link or not, reflecting operator's policy. For simple policies, the function only checks whether the origin AS of the flow is not connected through a more preferred peer. An example of this case is when traffic from a peering neighbor is entering through the link with a transit provider. Beyond making $IsFlowUndesired$ more complex, sophisticated policies also require more checks. For instance, for disaggregated prefixes, our algorithm also checks for traffic towards more specific prefixes entering at other links.

3) Impact assessment: After an inbound unsatisfied interest is detected, the algorithm assesses the impact of the corresponding flow. We follow a similar approach to the outbound case: we account for the actual traffic of the inbound flow in the peak hour of the network. We stress again that other metrics can be easily added or used as replacement, by modifying the ranking step of this algorithm.

Each dissatisfaction is stored as a tuple ($external\_link$, $flow\_attributes$, $traffic\_volume$). The algorithm finally returns an ordered set of such tuples, so that operators can perform detailed analyses on the detected dissatisfactions.

Similar to the outbound case, the inbound unsatisfied interests detection algorithm is correct, provided that its input is also truthful. Indeed, it is centered around the comparison between data-plane measurements of each traffic flow and the intended policies for that flow.

V. System architecture

In this section, we describe the warning system that we designed to detect, rank, and create alerts for inter-domain unsatisfied interests. We first describe the architecture of the system, and then the implementation of all its components.

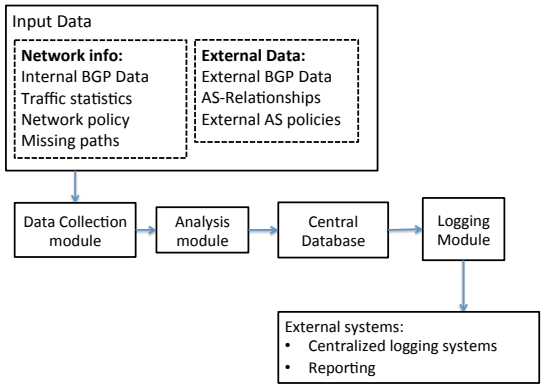

Fig. 6. Architecture of our warning system.

A. Architecture

Our warning system relies on four modules, as shown in Figure 6. A *data collection module* gathers the required *input*, interfacing with network devices or data collectors. The *logging module* communicates with external management systems, triggering warnings and outputting results in a convenient format. Finally, the *analysis and central database modules* implement the logic to detect, rank, and store warnings. By relying on those four modules, this architecture decouples the implementation of dissatisfaction detection algorithms (analysis module) from external interfaces. This facilitates the implementation of detection algorithms, by isolating the complexity of external interfaces in separate modules. Further, it simplifies the adaptation to different networks, i.e., by modifying the external modules (data collection and logging).

We now provide more details on each module.

1) Data collection module: This module provides a standard interface between our system and input data sources. Typically, it interacts with multiple data sources at the same time. Indeed, a key ability for detecting unsatisfied interests consists in correlating different types of data that can only be fetched from multiple monitoring systems (e.g. BGP collectors, network controllers, traffic monitoring, routers), or protocols (e.g. JSON, XML, CSV, etc.). We describe possible data extraction methods in Sec. V-B.

2) Analysis module: The analysis module is the heart of the system. It implements the algorithms described in Section IV using the input provided by the data collection module. Operators can tune the analysis module's parameters to fit the behavior of the algorithms to their needs (e.g. incident classification, frequency of operation, etc.). Observe that our algorithms can also be easily parallelized. For example, different missing routes in the outbound dissatisfaction detection algorithm can be processed in parallel, since their analysis does not require shared information. Similar considerations apply, for instance, to different traffic flows in the inbound dissatisfaction detection algorithm.

3) Central Database: This module stores the output of the dissatisfactions detection algorithms. Such an output contains fine-grained attributes to generate detailed reports on unsatisfied interests (see Section IV). In particular, for every detected unsatisfied interest, it includes its class (inbound or outbound),

its impact according to the implemented metrics, the category to which it belongs (neighbor preference dissatisfaction, next-hop diversity dissatisfaction, etc.), and additional information (for example, attributes of the missing route in the case of outbound dissatisfactions).

4) Logging Module: This module logs warnings obtained from the analysis one, and implements the external interface of our warning system. In concrete, it isolates the other modules to external systems, and translates unsatisfied interest information into specific formats. The resulting warnings could be used directly by operators, for example, to generate alarms upon dissatisfactions on certain traffic flows. Moreover, they can be sent to other management systems (e.g., a general warning system or an SDN controller) deployed in the network. For example, this can be useful to correlate dissatisfaction alarms with contingent network state, and block or adapt warnings under specific network conditions (failures, congestion, etc.).

B. Implementation

We developed a Proof of Concept (PoC) of our warning system and run it on a server with 16 cores and 32GB of RAM. Python was used to implement the algorithms and the logic of the *data collection* and *analysis* modules. We employ MySQL to implement the *central database module*. The logging module generates summary files on CSV format that are later used to generate plots (with MatplotLib [28]). Further, we implemented support for the detection of most dissatisfaction cases (i.e. unexpected traffic dissatisfactions and inbound cases) in Pmacct, as documented at http://wiki.pmacct.net/ImplementationNotes. We plan to release a full implementation of our prototype as part of the future work.

A major challenge from the implementation viewpoint is the collection of **input data** for our detection algorithms. We now discuss methods that operators can use to gather such data.

1) Traffic data: Typically, ISPs collect traffic statistics to perform various business-critical activities, from accounting and billing to traffic engineering. *Netflow* and *sflow* are the two most popular technologies for collecting this data [29].

2) Received BGP routes: All routes received from external neighbors need to be collected to run the outbound dissatisfaction detection algorithm (see Section IV-A). Several methods can be used for this purpose, including (i) the usage of custom scripts, e.g., based on router CLI commands and screen scraping; (ii) the configuration of iBGP sessions with add-path [30] (or similar features to propagate all BGP routes) between edge routers and a route collector (such as [31], [32]); (iii) the usage of monitoring protocols like BMP [33] and (iv) the configuration of selective port mirroring on edge routers, as proposed in [34].

3) Intended Policies: The outbound dissatisfaction detection algorithm needs the relative preference that operators have for different routes on which to send outbound traffic. In many cases, this can be calculated automatically by checking the default local-preference given to neighboring ASes. For more complex configurations, the preference could be provided manually by operators. In contrast, the inbound dissatisfaction detection algorithm needs as input the attributes of the traffic that should not enter the network over specific inter-domain links. Different sources of information can be used to obtain this data automatically. The peering relationships of the network can be used to build a starting policy for unexpected or unwanted inbound traffic. In a typical set-up, for instance, an operator does not want traffic from settlement-free peers, or its customers, on transit links. The peering relationships to neighboring AS can be obtained using router configurations, BGP data, or by fetching information from Internet Routing Registries (IRR), when available. In cases in which ASes are allowed to steer inbound traffic over links with the same AS, using BGP communities or MED, operators would like to check if their neighbor is respecting their commands. This information is typically reflected in router configurations.

4) Missing paths: The algorithm to detect outbound dissatisfactions takes missing paths as input. We recall that missing paths are those that are supposed to be received but are actually not received due to policies of external ASes (see Section IV-A). Our system currently focuses on two general and practically relevant classes of missing paths, that is, inconsistent advertisements and incomplete sets of paths.

Inconsistent advertisements identify BGP messages received from the same neighboring AS on different inter-domain links but not equally preferred by the local AS, e.g., because of different attributes[19], [20], [35]. We gather inconsistent advertisements comparing the routes announced by each peer on different physical location, as in [19].

Incomplete sets of routes represent cases in which a neighboring AS does not announce routes to some prefixes while it was supposed to. Of course, determining incomplete sets of routes depends on operators' expectations. While those expectations can be case-specific, our system currently focuses on two policies that are commonly shared by the large majority of operators [35]. Namely, we check that (i) transit providers announce routes to all destination prefixes, and (ii) peers propagate all routes originated by the peer itself or its customers[1]. To perform those checks, for each neighboring AS X, we compare the routes received from X with those that X announces to other ASes, as exposed by public BGP collectors, like Routeviews [37] or RIPE RIS [38], and AS relationship datasets, like the one provided by CAIDA [39]. In a real deployment, operators can also rely on their own data sources, for example, AS relationships provided by commercial companies [40] or special policies agreed with direct neighbors, to complement public data sources.

The algorithm used to find incomplete sets of routes is detailed in Algorithm 3. For every neighboring AS X, we consider the list of prefixes in which X appears in the AS-PATH of some BGP route. We then compare the list of prefixes obtained from BGP routes received by the local ASes with the one extracted from external BGP sources (e.g., RIS and Routeviews). If X is an eBGP peer, we only need to analyze the routes where the successive AS in the AS-PATH is X or one of its customers.

Note that even if the input is not 100% accurate (e.g.,

[1]Note that our tool supports partial peering [36], in the sense that we can define the subset of customer routes that the ISP is expecting to receive.

input : External BGP paths.
output: Incomplete paths.

```
[1]/* Only analyze those paths where a peering AS is seen:
      */
[2]foreach Path P, with ASPATH containing a neighboring AS Neigh do
[3]   if Neigh is a transit provider; or Neigh is a peer and the path arrives
      from one of the customers of Neigh then
[4]      PathAttributes ← BGP attributes from P;
[5]      BestPaths ← Current Best paths of the network towards
         GetNLRI(P);
[6]      BackupPaths ← Best paths of the current network towards
         GetNLRI(P) when BestPaths are removed;
[7]      if P is better than any path in BestPaths or BackupPaths then
[8]         Return P;
[9]      end
[10]  end
[11]end
```

Algorithm 3: Algorithm used to obtain the incomplete set of routes.

because of exceptions to the assumed AS policies [41], [25]), Algorithm 3 tends to reduce the likelihood of incorrect output, since it considers a limited set of routes (related to the local neighborhood). We validate the impact of inaccurate input on the output of the algorithm in Section VI-D. Moreover, we can avoid artifacts of transient phenomena (e.g., misconfigurations or outages) to influence the output of Algorithm 3 by providing long-lived BGP paths in input to it.

VI. Evaluation

In this section, we present the results obtained after deploying a prototype of the warning system in the network of two service providers. We used our system in an offline mode, for a-posteriori analyses of unsatisfied interests. In the following, we first describe our datasets in Section VI-A. We then discuss unsatisfied interests detected by our system for outbound and inbound traffic in Sections VI-B and VI-C, respectively. We finally validate our results with controlled experiments in Section VI-D.

A. Data-sets

Our evaluation is based on the following two datasets.

1) **Tier-2**. This dataset consists of the BGP routing tables and traffic data from an European Tier-2 network for the month of June 2014. Its network spans several countries, and has BGP connections with around 900 neighboring ASes. The routing tables are taken directly from the border routers of the network using Command Line Interface (CLI) commands. Hence, we avoid hidden BGP paths, which can trigger false alarms in some of our test [42]. Concerning traffic data, we have measurements aggregated per destination and per source prefix, as collected by network core routers.
2) **Academic Network**. This dataset includes the BGP routing tables and traffic data from the Spanish academic network (RedIris) for the month of March 2013. RedIris has around 20 nodes, and is connected to two inter-domain traffic providers and several peers over private links and exchange points. The BGP tables are obtained directly from each of the border routers, as for the Tier-2 dataset. The traffic data consists of Netflow dumps from the routers of the network. Netflows dumps allow us to obtain the total traffic aggregated over different characteristics, such as ingress/egress interface; source/destination prefix or transport protocol.

In the following, we use the Tier-2 dataset of the outbound part, due to large path diversity that this network possesses, which can provide a rich set of results for this type of traffic. We move to the academic network dataset for the inbound dissatisfactions results, as the granularity of the traffic of the Tier-2 dataset is too coarse (cannot be divided in ingress traffic, per origin prefix, per physical link) to run the inbound dissatisfaction detection algorithm.

B. Outbound traffic measurements

Starting from the Tier-2 dataset, our system first searched for missing paths using the procedures described in Section V-B4. We found 645 peers with 654,779 missing paths affecting 232,193 prefixes. 78,876 of these prefixes were affected by inconsistent advertisement, and 192,545 were affected by incomplete paths. Not all of these missing paths are meaningful for our analysis. Our system indeed runs the algorithm described in Section IV-A to identify which of these paths lead to unsatisfied interests, assess their impact, and classify them correctly. After feeding the algorithm with those missing paths, we found that 439,273 of them (about 67%) have some kind of impact to the network. Those paths globally affected 144,622 prefixes. The collection and processing of network data dominates the overall execution time of the system, as it takes almost 10 hours to be completed. After network data is indexed, the algorithm ran in around 1 hour.

Finally, we queried the Central Database storing the output of the detection algorithm, to find the cases with larger operational impact. In the following, we analyze the results of our queries, providing an overview across all unsatisfied interests and a case-by-case analysis of the most impacting dissatisfactions.

1) Results overview: We aggregate results according to different dimensions.

First, we group unsatisfied interests by neighboring AS and impact type. Results are shown in Figure 7. The X axis ranks the different grouped unsatisfied interests based on their impact, while the Y axis shows the amount of traffic impacted by each unsatisfied interest. For confidentiality reasons, we use a non-disclosed value, in the order of Mbps, to scale the graph. Figure 7 displays many interesting things. First, it shows that unsatisfied interests do have significant impact in practice: Globally, they affect an absolute, non-scaled amount of traffic of about 2Gb per second! Moreover, it highlights that the distribution of unsatisfied interests with respect to their impact is quite skewed. In total, 740 grouped unsatisfied interests are found, but only 84 (approximately 11%) have an actual impact. The impact of dissatisfactions on traffic is also skewed: The top 10 account for 85% of affected traffic. Finally, both frequency and impact of different types of dissatisfactions are uneven. Generally speaking, the unexpected transit dissatisfaction is the less common, with only 3 cases in the top 50 of most impacting unsatisfied interests. Nevertheless, this type of dissatisfaction has the

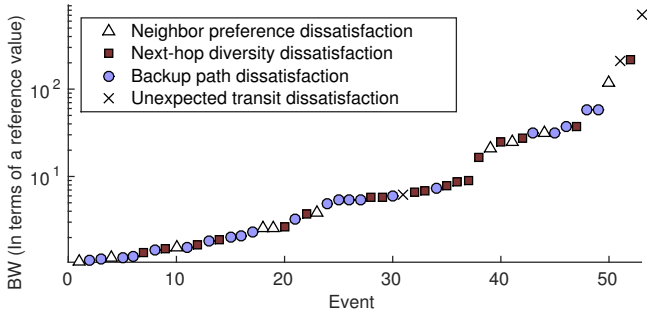

Fig. 7. Outbound unsatisfied interests for the Tier-2 network, grouped by neighboring AS and type of impact.

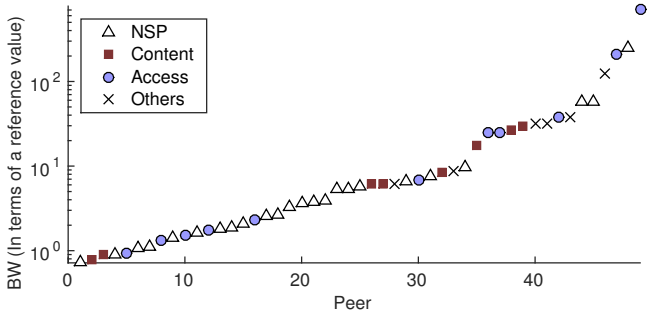

Fig. 8. Outbound traffic affected by unsatisfied interests for every neighboring AS.

overall greatest impact (the first and third most impacting dissatisfactions are of this type). Neighbor preference, next-hop diversity and backup path dissatisfactions complete the top-50 list with 9, 18, and 20 cases, respectively.

Observe that different unsatisfied interests in Figure 7 can pertain to the same neighboring AS. We found that from the top 50 cases with more aggregated traffic from the figure are linked to 41 different neighboring ASes.

To show the impact of specific neighbors, we plot the detected dissatisfactions aggregated only on a per neighboring AS basis in Figure 8. The X axis ranks the neighboring ASes based on the impact of their unsatisfied interests respectively induced by them. The Y axis shows the impact of dissatisfactions in terms of affected traffic, using the same scaling process as in Figure 7. We found outbound unsatisfied interests for about 471 neighboring ASes, for which 66 had an impact larger than zero. Even the distribution of ASes responsible for traffic impacted by dissatisfactions is highly skewed. Indeed, the top 10 ASes account for 87% of the total impacted traffic.

2) Case-by-case analysis: Figures 7 and 8 show that few dissatisfactions and few ASes are responsible for most of the traffic affected by unsatisfied interests. We now take the perspective of an operator, and delve into the dissatisfactions with the highest impact, trying to understand their causes and possible solutions to them. In particular, we focus on the top six AS (peers 45 to 50) in Figure 8.

The first and third AS of Figure 8 have a similar type of impact to the network, mostly due to the first and third unsatisfied interests in Figure 7, which are *Unexpected transit* ones. Both these ASes are multi-homed customers of the Tier-2 AS. The detected dissatisfactions are generated by the selective advertisement (to another neighbor) of more specific prefixes with respect to those sent to the Tier-2 AS. More concretely, each dissatisfaction-inducing AS advertised at least one prefix p to the other transit provider but not to the Tier-2 AS, which is left with a less specific prefix p' covering p. We assume that this is likely done for redundancy purposes. The Tier-2 AS, however, receives the more specific prefixes from non-customer neighbors, and it is forced by them to forward traffic to non-customers ASes rather than directly to its customers. Unfortunately, there is no easy way to solve this situation [1]. Filtering the more specific prefixes could be, in some cases, considered as a contradiction to their policy from the point of view of the customers[2]. This information, however, could be very useful for network operators and peering managers, e.g., to re-negotiate commercial agreements with those customers.

The effect of the second AS of Figure 8 corresponds in most part to the second most-impacting dissatisfactions in Figure 7, which is a *Next-hop diversity* one. This particular AS is a service provider with multiple points of presence in different countries in Europe. By looking at the missing paths creating this dissatisfaction, we find that all of them are due to inconsistent advertisement. This peer is indeed connected to the Tier-2 AS via three links, a private one and two through two different European IXPs. The two sessions through the two IXPs show missing paths. Figure 9 depicts the distribution of those missing paths on each link. The top figure depicts the number of missing paths per ingress link in the network, while the lower figure weights each missing path with the amount of outbound traffic carried by the corresponding prefix. Almost half of all prefixes are consistently announced through IXP1 and about one fourth through all the three links (top figure); however, the corresponding prefixes do not attract much traffic (bottom figure). The other prefixes are inconsistently announced, hence the neighboring AS attracts most of the traffic to the private link. The missing paths are also not covered by a less specific prefix. The neighboring AS could be incurring in inconsistency advertisement because it does not want to transport traffic from the two IXPs to these specific destinations. Depending on the contractual situation between the Tier-2 and this neighboring AS, the Tier-2 may contact the neighboring AS to enforce pre-agreed policies. If the missing paths are due to route server path hiding [21], the two ASes may decide to establish a direct BGP section between each other.

The fourth AS of Figure 8 is the one causing the fourth most-impacting interest dissatisfaction of 7, which is a *Neighbor preference* one. In this case, the dissatisfaction-inducing AS is a settlement-free peer of the Tier-2 AS, hosting a top-30 Alexa site, and being source and destination of a considerable amount of traffic. The missing paths causing the detected dissatisfaction are caused by routes that the neighboring AS receives but does not announce directly to the Tier-2, which

[2]The customers could use the inbound detection algorithm to detect if a provider is ignoring the more specifics.

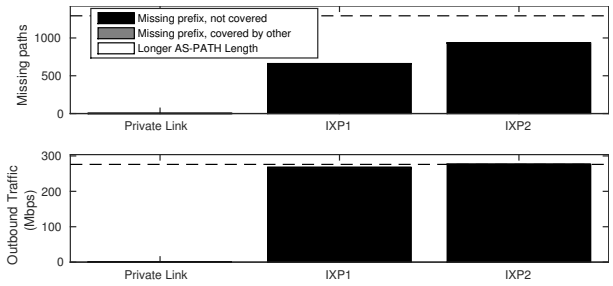

Fig. 9. Inconsistencies from an individual peer. The dashed lines identify the total number of prefixes or total outbound traffic for this neighboring AS.

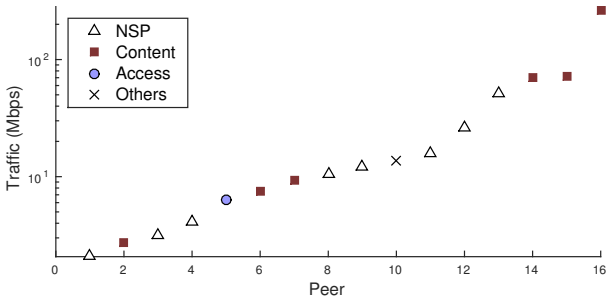

Fig. 10. Peers traffic received over transit links.

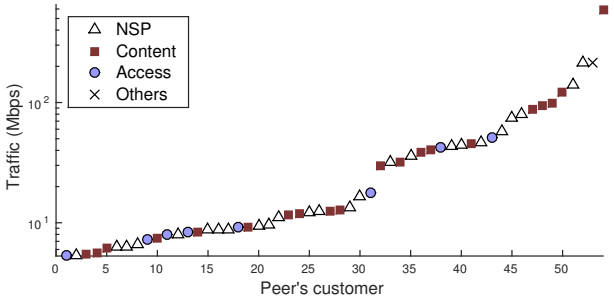

Fig. 11. Peers customers' traffic received over transit links.

receives routes to the corresponding only through its own transit providers. After a quick examination, we found that the origin AS of the missing paths indistinctly belongs to the same organization. The dissatisfaction-inducing AS may then be filtering the paths because it wants to avoid transporting traffic to these prefixes from the Tier-2 to its sibling origin AS. By detecting this case, the Tier-2 could analyze whether it could demand the neighbor to announce this routes directly, for example if the peering agreement explicitly includes this on its terms [35][20].

Finally, the fifth and sixth ASes in Figure 8 are inducing *Backup path* dissatisfactions. These two ASes face a similar type of dissatisfaction, due to a common customer. The problem is that both ASes are not advertising paths that they receive from this common customer. The two neighboring ASes and their customer are all settlement-free peers of the Tier-2 AS, connecting through only one physical link (an IXP). In the case of a failure of this IXP, the traffic towards the customer risks to be forced over transit providers. The reason why these two peers do not forward these prefix is not yet clear.

C. Inbound traffic measurements

The inbound detection algorithm basically loops over the ingress flows per source prefix, and decides whether those flows fit into the interests of the operator. An important part of this algorithm is therefore to obtain, in a (semi-)automated way, the interests of the local network. For our evaluation, we reverse-engineered the applied BGP policy directly from router configurations in our academic network dataset. In particular, we correlated default Local-preference values to the type of neighboring AS. The result was a very basic policy targeting to enforce two key sub-interests, that is, 1. receive peer's traffic only over peer's links, and 2. receive peer's customer's traffic only over peer's links.

1) Results overview: Figure 10 summarizes the cases in which peer's traffic is found over transit provider's links, hence unsatisfying sub-interest 1. We found cases for 33 peers, 16 with more than 1Mbps. The X axis ranks each settlement-free peering ASes based on the amount of traffic originating in those AS found in transit links. The Y axis measures directly that traffic in Mbps. The top 3 ASes account for about 75% of the total affected traffic.

Figure 11 summarizes the cases in which peer's customer's traffic is found over transit provider's links, that is, those unsatisfying sub-interest 2. The X axis ranks the customers of the settlement-free peers based on the amount of traffic originating in those AS found in transit links. The Y axis measures directly that traffic in Mbps. We found traffic of more than 500 peer's customers over the transit links of the network (in the figure, we just show the cases with more than 5Mbps).

2) Case-by-case analysis: Although our system does not aim to provide root causes for the undesired traffic, we inferred reasons for the detected dissatisfactions. Regarding dissatisfactions of sub-interest 1 (see Figure 10), the top three settlement-free peers with more traffic on the transit links are content providers. This type of companies performs Traffic Engineering practices that differ from those of access or transit networks. Indeed, content providers use different sources of information, such as the DNS system [43] or cache location [44], to select the source host and the path towards the end user. Those configurations may work around the BGP routing system, hence enabling those ASes to send traffic ignoring direct connections [44]. Given these circumstances, the Academic-network could decide to expand its infrastructure with these content providers [44], or explore other collaboration techniques, as described for instance in [43], to reduce the amount of traffic of these companies over its transit links. The next three companies with most traffic are service providers (NSPs). They may prefer to send some traffic destined to the academic network to third-party ASes, for instance to limit back-haul transport costs.

In contrast to sub-policy 1, dissatisfactions to sub-policy 2 (Fig. 11) are harder to analyze because of the (routing) distance of ASes causing them to the academic network. Still, there might be cases of undesired policies from the direct neighbors: for instance, they can perform intermediate filtering [45]. The undesired traffic can also be due to the outbound

traffic policies of the origin networks. Note that many of the top contributors are content providers, for which the same analysis as for dissatisfactions of sub-policy 1 applies. Concretely, these companies might select the paths from the caches connected to the transit providers of the Academic network, instead to the ones available through the settlement-free peers. Operators and peering managers can use this data to potentially look for new peering agreements and reduce the inbound traffic through transit providers.

D. Validation

We finally validate our warning system on synthetic data and on the evaluated real-world measurements.

1) Systematic validation on controlled experiments: First, we run controlled experiments on synthetic datasets for which we assume to know the ground truth.

Our algorithms always correctly detect and classify unsatisfied interests if their input is correct. To experimentally confirm the correctness of our prototype system, we take the Internet topology and the AS relationships as reported by Caida [39] and PeeringDB [46]. Note that some ASes are multi-connected. While we have no information on physical AS links, some ASes are indeed connected through multiple links whenever they are both members of the same IXPs. Thanks to those links, we are able to reproduce all the unsatisfied interests presented in Sec. III (including inconsistencies advertisements among peering ASes). Then, we simulate BGP route propagation when compatible policies (i.e., Gao-Rexford ones and consistent advertisement) are set Internet-wide, and generate the corresponding data-plane traffic. In this case, our warning system correctly reports no unsatisfied interests, independently on the AS where it is run. Moreover, in separate experiments, we select a victim AS and modify BGP policies and traffic of a certain set of close ASes, so to cause given unsatisfied interests on the victim. We re-run our warning system on the routes and traffic as received by the victim AS: It always correctly detects and classifies all types of inbound and outbound dissatisfactions.

Our system is robust to incorrect AS policies and relationships. Most input taken by our system can be accurately collected in real networks. This is especially true for all data (traffic measurements, BGP routes, local policies, etc.) locally available at the AS running our system (see Sec. V).

We therefore evaluate the sensitivity of our system to the commercial relationships between external ASes and remote AS policies, that we use to infer missing BGP paths and the corresponding unsatisfied interests. We take the Internet topology and the AS relationships as reported by Caida and PeeringDB as *original topology*. Those relationships have an accuracy of about 99% according to [47]. Hence, we randomly change 1% of the AS relationships in our original topology and obtain a *modified topology*. We then run our system on every AS, using the modified topology as an input. This generates discrepancies between the missing paths on the original topology and those computed by our algorithms. Those discrepancies can reflect both inaccurate commercial relationships (if the original topology represents the ground truth) and exceptions to the Gao-Rexford policies assumed by our algorithms (if the modified topology is correct about commercial relationships, but the paths extracted from the original topology are correct considering actual AS policies). In turn, those discrepancies affect the correctness of computed interest dissatisfactions. For every AS, we therefore compare the result of our algorithms with dissatisfactions that would have happened on the original topology, whose BGP paths are taken as ground truth. We repeat such an experiment 30 times for statistical significance.

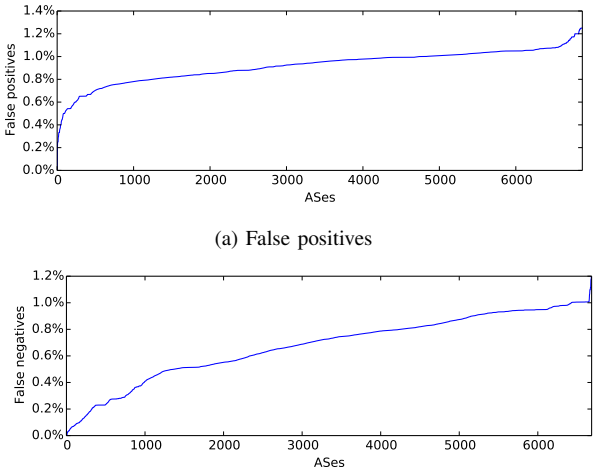

(a) False positives

(b) False negatives

Fig. 12. Sensibility analysis to inaccuracies of relationships of external ASes.

The results of those experiments are shown in Figure 12. In particular, Figures 12b and 12a respectively plot the average number of false positives (dissatisfactions detected by our system but not present in the original topology) and false negatives (dissatisfactions in the original topology not captured by the system). Our results show that both false positives and negatives are under 1% for almost all the ASes, and not more than 1.2% on all of them. This means that the error rate of the system is less than proportional with respect to the inaccuracy of AS relationships, in the vast majority of the cases.

Our system is so robust to inaccurate input mainly because of the high connectivity of the Internet. When running on a given AS X, our system generates false positives or negatives only for routes that (1) traverse a pair of ASes with an erroneous relationship in the input (or assumed policy), and (2) do affect the interests of X. Intuitively, the likelihood that both (1) and (2) occur at the same time is low. Actually, the likelihood that routes traversing the inaccurately-classified AS relationships are received by X is already quite low, given the number of physical paths in the Internet.

2) Informal validation with operators: Further, we conduct an informal validation with the operators of the Tier-2 AS in our evaluation dataset. After showing them the results of our system, they confirmed that the detected dissatisfactions were indeed not expected and problematic. As far as we know, they started to investigate the detected dissatisfactions, by talking with the peers responsible for them.

VII. Related Work

Routing divergence. Several authors have studied the effects of routing conflicts in the Internet. These works have examined the conditions in which uncoordinated policies can cause the BGP algorithm to diverge [3][4][5]; and to propose systems that can detect these cases [7][8], or prevent them [4]. The unsatisfied interests that we analyze in this work do not fit in this category, as they focus only in the policy interest of the operator under a stable state. In other words, we analyze cases in which BGP converges, but where one or more ASes do not obtain the state that they originally intended.

BGP security. BGP is fundamentally insecure and vulnerable to different attacks. Multiple proposals have emerged to secure the Internet [48]. In recent years, the IETF standardized the Resource Public Key Infrastructure (RPKI) [49]. RPKI uses a public key system to validate some of the information included in BGP updates. Our warning system is not specifically designed for security purposes. Nevertheless, it could detect the effect of specific attacks (like forged BGP routes), if they are root causes of unsatisfied interests. Two observations hold in this case. First, the system can be customized for security objectives, e.g., to always report unsatisfied interests for security-sensitive traffic flows, with simple modifications to dissatisfaction impact assessment within the detection algorithms (see Section IV). Second, its output can be correlated with dedicated systems [50].

External peering auditing. The distributed nature of the Internet makes it prone to situations in which the behavior of a single system (either allowed or not) can affect many others. Many authors have analyzed how to detect the cause of specific situations affecting ASes, even under environments in which only partial information can be obtained. [19] describes the problem of inconsistency advertisement from neighboring peers and how they could be detected using local data. [51] uses data-plane data (traceroute) to find disruptions between what is announced in the control-plane and the actual path of each packet. [52] proposes a cryptography system that can test several properties of the received routes from neighboring ASes. [53][54][55] use network information to pinpoint any AS behaving in an unexpected manner, or causing specific route changes. [45] checks for prefix filtering that can limit the visibility of prefixes for other ASes. These systems complement the information provided by the dissatisfaction warning application. Some of these systems (e.g. [19]) can feed our application with missing paths that can be used to run the outbound traffic algorithm. Other systems (e.g. [45]) could help operators find root causes for specific unsatisfied interests that with greatest impact on the network.

Internal configuration checking. Some unsatisfied interests can be explained, not by policies of external ASes, but by mistaken configurations of internal devices. Route leaks, for instance, can arise due to this problem, and can trigger the dissatisfaction warning system (e.g. by detecting traffic to transit prefixes arriving in a settlement-free peer link) [56]. Operators can implement configuration checking systems that can avoid misalignment between internal policies and configurations [57]. If the operators actively use Internet Routing Registry (IRR) to publish their inter-domain policy, systems like [58] can be employed to check the policy against the BGP updates.

Different inter-domain routing protocols. Due to the limitations and problems experienced by the current Internet, different authors have proposed improved inter-domain network architectures. These can either provide flexibility to the systems [59]; improve the behavior of BGP by extending some aspects of the protocol [60]; or working with overlay protocols that provide more features and control [61]. Our system operates abstractly in policy decisions from ASes, and not directly in information generated exclusively by BGP. Therefore, the system can be adapted to any of these inter-domain architectures.

Content providers and network neutrality. Content providers have established themselves as the origin of a large proportion of Internet traffic. These companies have particular operational practices and policies concerning inter-domain routing, which can clash with the interest of eyeballs or transit providers, triggering warnings in our system. For instance, some content providers have no backbone, and decide forwarding routes disregarding any policy reflected in BGP announcements [44][62]. Access and transit networks should be aware of these policies, and use the warning systems as a notification system that can point operators to cases with a large impact on the network. The operators can then tackle each case in conjunction with content provider. Providers can use systems like the ones proposed in [43][63][64], which can be used to establish a collaboration between content and access providers, in order to reduce incompatible interests.

A large part of the discussions around *network neutrality* have been generated by conflicting relations between access and content provider networks [65]. Although specific details of agreements fitting network neutrality clauses is beyond the scope of this paper, we envision that our system could be use by *both sides* of this relationship. The warning system cold provide useful information to specific situations that do not fit points of any peering agreement, such as missing prefix announcement or disrespect for inbound policy.

VIII. Conclusion

In this paper, we studied inter-domain routing configurations in which the (economic) interests of one or more ASes are unsatisfied. Taking an AS-centric perspective, we (i) classified possible unsatisfied interests; (ii) proposed algorithms to detect them and assess their impact; and (iii) described a warning system providing operators with critical input on business-impacting unsatisfied interests. We used our system to perform real-world measurements. Our results show that unsatisfied interests do occur in practice and can affect a non-negligible amount of traffic (a couple of Gbps in one of our cases).

We stress that unsatisfied interests are not a bug of BGP. Rather they are the normal product of the Internet business model, where egoistic players autonomously pursue their own selfish interests. Our results show that unsatisfied interests should receive much more attention from future research works, including proposed evolutions of inter-domain routing

(e.g., [66]). In particular, we do not advocate for tentatives to avoid or prevent conflicts of interests, but for primitives, mechanisms or tools (like our warning system) enabling network operators and peering managers to better manage the unsatisfied interests affecting their networks.

REFERENCES

[1] J. C. Cardona, P. Francois, and P. Lucente, "Impact of BGP filtering on Inter-Domain Routing Policies," *draft-ietf-grow-filtering-threats-06. Work in Progress. IETF Draft.*, 2015.

[2] F. Knzler. (2011) How More Specifics increase your transit bill (and ways to avoid it). https://ripe63.ripe.net/presentations/48-How-more-specifics-increase-your-transit-bill-v0.2.pdf. INIT7.

[3] T. G. Griffin, F. B. Shepherd, and G. Wilfong, "The stable paths problem and interdomain routing," *IEEE/ACM ToN*, 2002.

[4] R. Sami, M. Schapira, and A. Zohar, "Searching for stability in interdomain routing," in *Proc. INFOCOM*, 2009.

[5] L. Cittadini *et al.*, "Wheel+ ring= reel: The impact of route filtering on the stability of policy routing," *IEEE/ACM ToN*, 2011.

[6] M. Suchara, A. Fabrikant, and J. Rexford, "BGP safety with spurious updates," in *Proc. INFOCOM*, 2011.

[7] A. Flavel *et al.*, "BGP route prediction within ISPs," *Computer Communications*, vol. 33, no. 10, pp. 1180–1190, 2010.

[8] L. Cittadini *et al.*, "From theory to practice: Efficiently checking BGP configurations for guaranteed convergence," *IEEE Transactions on Network and Service Management*, vol. 8, no. 4, pp. 387–400, 2011.

[9] S. Hares, Y. Rekhter, and T. Li, "A Border Gateway Protocol 4 (BGP-4)," *IETF RFC 4271*, 2006.

[10] P. Faratin *et al.*, "Complexity of Internet interconnections: Technology, incentives and implications for policy," *TRC*, 2007.

[11] R. V. Oliveira *et al.*, "In search of the elusive ground truth: the internet's as-level connectivity structure," in *ACM SIGMETRICS PER*, vol. 36, no. 1. ACM, 2008.

[12] S. Uhlig and O. Bonaventure, "Designing BGP-based outbound traffic engineering techniques for stub ASes," *ACM SIGCOMM CCR*, vol. 34, no. 5, pp. 89–106, 2004.

[13] P. Gill, M. Schapira, and S. Goldberg, "A survey of interdomain routing policies," *ACM SIGCOMM CCR*, 2013.

[14] P. Verkaik *et al.*, "Wresting Control from BGP: Scalable Fine-Grained Route Control." in *USENIX Annual Technical Conference*, 2007, pp. 295–308.

[15] B. Quoitin *et al.*, "Interdomain traffic engineering with BGP," *Communications Magazine, IEEE*, vol. 41, no. 5, pp. 122–128, 2003.

[16] R. Gao, C. Dovrolis, and E. W. Zegura, "Interdomain ingress traffic engineering through optimized AS-path prepending," in *NETWORKING 2005*. Springer, 2005.

[17] N. Feamster, J. Borkenhagen, and J. Rexford, "Guidelines for interdomain traffic engineering," *ACM SIGCOMM CCR*, 2003.

[18] B. Quoitin *et al.*, "Interdomain traffic engineering with redistribution communities," *Computer Communications*, vol. 27, no. 4, pp. 355–363, 2004.

[19] N. Feamster, Z. M. Mao, and J. Rexford, "BorderGuard: Detecting cold potatoes from peers," in *Proceedings of IMC*, 2004.

[20] W. Norton. (2012) A Brief Study of 28 Peering Policies. http://drpeering.net/AskDrPeering/blog/articles/Peering_Rules_of_the_Road_-_A_Brief_Study_of_28_Peering_Policies.htmls.

[21] E. Jasinska *et al.*, "Internet Exchange Route Server," *draft-ietf-idr-ix-bgp-route-server-04. Work in Progress. IETF Draft.*, 2014.

[22] J. C. Cardona Restrepo and R. Stanojevic, "A history of an Internet eXchange Point," *ACM SIGCOMM CCR*, 2012.

[23] B. Donnet and O. Bonaventure, "On BGP communities," *ACM SIGCOMM CCR*, 2008.

[24] L. Cittadini, S. Vissicchio, and G. Di Battista, "Doing Don'ts: Modifying BGP Attributes within an Autonomous System," *Proc. NOMS*, 2010.

[25] S. Vissicchio, L. Cittadini, and G. Di Battista, "On iBGP routing policies," *IEEE/ACM ToN*, 2015.

[26] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *IEEE/ACM ToN*, 2001.

[27] R. Powell. (2013) XO Joins Level 3 For Bit-Mile Peering. http://www.telecomramblings.com/2013/01/xo-joins-level-3-for-bit-mile-peering/.

[28] J. D. Hunter, "Matplotlib: A 2D graphics environment," *Computing in science and engineering*, 2007.

[29] B. Claise, "Cisco systems NetFlow services export version 9," *IETF RFC 3954*, 2004.

[30] V. Van den Schrieck, P. Francois, and O. Bonaventure, "BGP add-paths: the scaling/performance tradeoffs," *Selected Areas in Communications, IEEE Journal on*, vol. 28, no. 8, pp. 1299–1307, 2010.

[31] P. Lucente. (2015) PMACCT. http://wiki.pmacct.net/.

[32] J. Choi *et al.*, "Understanding BGP next-hop diversity," in *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, 2011.

[33] J. Scudder, R. Fernando, and S. Stuart, "BGP monitoring protocol," *draft-ietf-grow-bmp-07. Work in Progress. IETF Draft.*, 2012.

[34] S. Vissicchio *et al.*, "Beyond the best: Real-time non-invasive collection of BGP messages," *Proc. INM/WREN*, 2010.

[35] CenturyLink. (2011) CenturyLink's North America IP Network Peering Policy. http://www.centurylink.com/legal/peering_na.html. CenturyLink.

[36] I. Castro and S. Gorinsky, "T4P: Hybrid interconnection for cost reduction," in *INFOCOM WKSHPS*, 2012.

[37] D. Meyer *et al.*, "University of oregon route views project," 2005.

[38] R. RIPE, "Raw Data," 2014.

[39] CAIDA. (2014) The CAIDA AS Relationships Dataset. http://www.caida.org/data/active/as-relationships/. CAIDA.

[40] "Dyn," http://dyn.com, 2014.

[41] W. Mühlbauer *et al.*, "In search for an appropriate granularity to model routing policies," *CCR*, 2007.

[42] S. Uhlig and S. Tandel, "Quantifying the BGP routes diversity inside a tier-1 network," in *NETWORKING*. Springer, 2006.

[43] I. Poese *et al.*, "Improving content delivery using provider-aided distance information," in *Proceedings of IMC*. ACM, 2010, pp. 22–34.

[44] C. Kaufmann. (2014) BGP and Traffic Engineering with Akamai. http://www.menog.org/presentations/menog-14/282-20140331_MENOG_BGP_and_Traffic_Engineering_with_Akamai.pdf. Akamai.

[45] A. Lutu *et al.*, "Separating wheat from chaff: Winnowing unintended prefixes using machine learning," *Proceedings of INFOCOM.*, 2014.

[46] PeeringDB, "PeeringDB," http://www.peeringdb.com/.

[47] M. Luckie *et al.*, "AS relationships, customer cones, and validation," in *Proceedings of IMC*, 2013.

[48] G. Huston, M. Rossi, and G. Armitage, "Securing BGPA literature survey," *Communications Surveys & Tutorials, IEEE*, 2011.

[49] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol." *IETF RFC 6810*, 2013.

[50] R. NCC. (2015) RIPE NCC RPKI Validator, Tools, and Resources. https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources.

[51] T. Flach, E. Katz-Bassett, and R. Govindan, "Quantifying violations of destination-based forwarding on the internet," in *IMC*, 2012.

[52] M. Zhao *et al.*, "Private and verifiable interdomain routing decisions," in *Proceedings of SIGCOMM*. ACM, 2012, pp. 383–394.

[53] A. Haeberlen *et al.*, "NetReview: Detecting When Interdomain Routing Goes Wrong." in *NSDI*, 2009.

[54] R. Teixeira and J. Rexford, "A measurement framework for pin-pointing routing changes," in *ACM SIGCOMM workshop on NetT*, 2004.

[55] Y. Wu *et al.*, "Diagnosing missing events in distributed systems with negative provenance," in *SIGCOMM*, 2014.

[56] S. K. *et al.*, "Problem Definition and Classification of BGP Route Leaks," *draft-ietf-grow-route-leak-problem-definition-01. Work in Progress. IETF Draft.*, 2015.

[57] N. Feamster and H. Balakrishnan, "Detecting BGP configuration faults with static analysis," in *NSDI*, 2005.

[58] G. Siganos and M. Faloutsos, "Analyzing BGP policies: Methodology and tool," in *INFOCOM 2004*, vol. 3. IEEE, 2004, pp. 1640–1651.

[59] B. Raghavan *et al.*, "Software-defined Internet architecture: Decoupling architecture from infrastructure," in *HotNets*. ACM, 2012, pp. 43–48.

[60] Y. Wang, M. Schapira, and J. Rexford, "Neighbor-specific BGP: more flexible routing policies while improving global stability," in *Int. joint conference on Measurement and modeling of comp. systems*, 2009.

[61] D. Farinacci *et al.*, "The locator/ID separation protocol (LISP)," *IETF RFC 6830*, 2013.

[62] P. Fiadino *et al.*, "On the detection of network traffic anomalies in content delivery network services," in *Teletraffic Congress (ITC)*, 2014.

[63] M. Wichtlhuber, R. Reinecke, and D. Hausheer, "An SDN-Based CDN/ISP Collaboration Architecture for Managing High-Volume Flows," *Network and Service Management, IEEE Trans. on*, 2015.

[64] A. Flavel *et al.*, "FastRoute: A Scalable Load-Aware Anycast Routing Architecture for Modern CDNs," *NSDI*, 2015.

[65] N. Economides and J. Tag, "Network neutrality on the Internet: A two-sided market analysis," *Information Economics and Policy*, 2012.

[66] V. Kotronis, X. Dimitropoulos, and B. Ager, "Outsourcing the routing control logic: better internet routing based on SDN principles," in *Proc. HotNets*, 2012.