

Integrating Inter-domain Routing Analysis in novel management strategies for large scale IP networks

Pedro A. Aranda Gutiérrez¹, Ilka Miloucheva²

¹Telefónica I+D

C:/Emilio Vargas 6, 28043 Madrid, Spain

²Salzburg Research

Jakob Haringer Str. 5, 5020 Salzburg, Austria

Abstract

This paper presents InterRoute, a technology for inter-domain routing monitoring, topology discovery and analysis that has been developed in the scope of the INTERMON IST project.

InterRoute is designed to track BGP-4 routing activity at intermediate points on an active end-to-end path in the Internet and to obtain inter-domain topology and metrics characterising the quality of inter-domain connectivity. Based on the use of publicly available inter-domain routing information repositories at intermediate points of active end-to-end connection, metrics for inter-domain route quality are evaluated by InterRoute and supplied for management purposes. Concentrating on the active end-to-end path at a given point in time, this approach yields a manageable amount of information which can be used as the basis for Quality of Service (QoS) validation supporting the network operator in connection management tasks.

The integration of the inter-domain stability analysis of InterRoute in the INTERMON architecture allows flexible interaction with QoS and thus Service Level Agreement (SLA) monitoring and analysis tools using common data base to study inter-domain QoS behaviour.

1 Introduction

A lot of research is aimed to study the impact of inter-domain routing on QoS/SLA. Service Providers, corporate users, customers and application service providers need tools to analyse and verify the stability of inter-domain routing and how the end-to-end QoS parameters service level agreements (SLAs) are impacted by inter-domain routing and topology changes [9], [10], [11],[12].

The Border Gateway Protocol (BGP) [26] is central for inter-domain stability and robustness of Internet, therefore one big area of research is the understanding of inter-domain routing stability based on passive

analysis and monitoring of BGP-4 messages [15],[16],[17],[18], [19], [20],[27],[28].

The inter-domain route stability could be defined from different view points. [21] discusses two notions of routing stability based on prevalence, e.g. probability of observing of a given route, and persistence, e.g. probability of the duration of a given route.

[18] defines inter-domain route stability based on metrics showing the reachability to a prefix for a given time period. They consider stability described by prefix availability, e.g. the fraction of time that a prefix is reachable, and prefix steadiness, e.g. the average time that a prefix is continuously reachable.

The route stability discussion in [18],[19] consider primary the route changes.

Our paper extends the notion of the route stability and defines it from the view point of the impact of inter-domain routing on QoS/SLA of end-to-end connection. The proposed router stability metric has been derived from the definition of the BGP-4 routing protocol and shows the potential impact of routing events on end-to-end QoS and SLA.

Adequately mapped, the proposed end-to-end connection inter-domain stability metrics allow a network administrator to locate problems in a network promptly. We describe the integration of the proposed metrics in the InterRoute toolkit to study the inter-domain stability of the end-to-end connection.

The InterRoute technology contributes to the known concepts for study of inter-domain routing and track BGP-4 routing activity for management purposes focussing at intermediate points on an active end-to-end path in the Internet. The design of the InterRoute technology is described and usage is explained in real inter-domain end-to-end QoS/SLA management scenario. The integration of InterRoute in the framework of INTERMON project serves as the starting point for visual data mining processes which help operators understand the relation between topology changes and deviations of QoS parameters in an end-to-end communication using the Internet. The interaction of the InterRoute technology with other tools of INETRMON toolkit is shown.

This paper is organised as follows. Section 2 presents the general approach for end-to-end QoS analysis referencing to the state-of-the-art. Section 3 addresses design of the InterRoute technology and particular functions of the InterRoute tool in QoS/SLA end-to-end management scenarios such as inter-domain routing metrics evaluation. Section 4 discusses integration in the INTERMON project and interaction with other tools of the INTERMON architecture. Section 5 presents future work and conclusions.

2 Analysis of inter-domain route quality for an end-to-end connection

2.1 State-of-the art and general approach

This paper proposes an approach for the study of inter-domain route stability of end-to-end connections based on analysis of *BGP-4 Update messages* at intermediate points of the end-to-end connection, e.g. a border router or exchange point included in the topology of the end-to-end connection.

The *BGP-4 Update message* contains routing information carrying withdrawals as well as announcements of routes. It consists of following parts:

- Withdrawn Routes (lists the destination prefixes for which the sending router is no longer ready to forward packets)
- Path Attributes (contains the route together with additional attributes applying to all destinations from NLRI field. AS Path is a mandatory attribute that describes the route. Local Pref is describing the preference for the advertised route. Next hop is optional and contains the IP address of the router. Multi Exit Discriminator (MED) discriminates among entry points to same neighbouring AS.)
- Network Layer Reachability information, e.g. NLRI (contains destinations given as IP prefixes, that the sender of the *Update message* can reach using some route).

There are approaches with different purpose to analyse BGP-4 behaviour in Internet based on selective monitoring and statistics calculation considering structure and parameters of the *Update message* for exchanges/border router usually contained in common routing repositories [6],[15],[16],[17]. The general approach to evaluate BGP-4 behaviour for route stability analysis is based on counting the announcements and withdrawals [16], [20].

In addition to the fraction of withdrawals and announcements, [20] proposes to select and evaluate more information in the BGP-4 Update message in order to detect and analyse duplicate announcements, route flapping, erroneous announcement, non-locality

of effects of instability and migration, reachability of AS and prefix, BGP-4 convergence.

In order to analyse BGP-4 behaviour in case of DoS attacks, [17] classifies the *BGP-4 updates* into classes indicative of actual route changes. Because there are differences in BGP-4 implementations, the specific BGP-4 behaviour from different senders is considered.

[6] describes different kinds of inter-domain routing anomalies which result from BGP-4 operation in Internet affecting the performance and efficiency of QoS provision. Analysing the *BGP-4 Update message* values and parameters, anomalies like slow BGP-4 convergence [22], withdrawals of dead prefixes and/or duplicate announcements of existing prefixes with no attribute change, link down, inconsistent flap damping policy and other could be detected.

[16] designs a BGP assistant to monitor and analyse BGP traffic based on BGP Updates and Route convergence times using Oregon Route Views [2] data. The Route Explorer based on BGP-4 [15] builds a visual graph of the network, highlighting link failures and route "flapping" problems as they are discovered. The customer can compare that with historical data on the network behaviour, or perform simulations on how the network would behave if certain critical routes were changed. The analysis of BGP-4 messages is recently also focus of IETF, for instance [23] address the single-router convergence time issues.

The novelty of the proposed approach in this paper is to define appropriate inter-domain route stability metrics of end-to-end connections aimed at analysing the effect of inter-domain routing on QoS/SLA provision. We define the inter-domain route stability of the end-to-end connection to be described by:

- *General* inter-domain route stability metrics taking into account the end-to-end route changes caused by BGP-4 messages and the potential impact on end-to-end QoS (SLAs) of BGP-4 messages of the routers included in the end-to-end connection. This definition leads to general differentiation of hot, warm and cold evaluation.
- *Extended* inter-domain route stability analysis metrics for discovery of anomalies of BGP-4 behaviour at the end-to-end connection including study of message sequences and temporal relationships of routing changes, implicit withdrawals, route flapping and duplicate announcements.

Analysis of sequences of BGP-4 messages could detect different kinds of BGP-4 anomalies due to software bugs, false configuration, policy changes. For instance, conflicting policies could be the cause of repeated advertising and withdrawing of routes [9].

2.2 General inter-domain route stability metrics of end-to-end connection

The general metrics take into consideration the inter-domain route changes directly impacting the connection and route changes possible impact on QoS/SLA such as:

- Loss of inter-domain connectivity and route change of end-to-end connection
- Changes in inter-domain topology of the end-to-end connection, i.e. alternative route changes
- probably traffic variations in the end-to-end connection impacting QoS/SLA.

Withdrawals indicate three potentially critical situations with regard to an end-to-end QoS / SLA:

- Connection end provider withdrew the route (either accidentally or intentionally) and the end-to-end connectivity was lost
- A transit provider withdrew the route. This was caused by changes in routing policies or operational errors and might have impacted the end-to-end SLA by establishment of alternative route or loss of end-to-end connectivity.
- A end/transit provider for the connection accidentally injected the route violating the routing policies agreed with his peers and had to withdraw it when the situation was detected. This might have created unwanted alternate paths, which will have impacted the end-to-end SLA.
- An advertisement for a prefix, which is more specific than the prefix installed in the routing table at the beginning of the observation period can redirect traffic to a new interface and thus affect the end-to-end QoS/SLA.

Based on analysis of BGP-4 protocol behaviour, this metric is graded in three different levels: cold, warm and hot, which depends on the parameters of the *BGP-4 Update* Message related to the intermediate border/exchange routers of the end-to-end connection:

- *Cold*: the router has had a stable behaviour during the complete observation period: routes to either endpoints of the end-to-end path were present in the routing table at the beginning of the observation time window and no advertisements related to the end-to-end path were received during it.
- *Warm*: The router has had stable routing entries to either endpoints of the end-to-end connection during the whole observation period but has received announcements from other autonomous systems for either endpoints. These announcements hint potential instabilities in preceding autonomous systems along the path, which are outside of the control of the operator.

- *Hot*: the router has received at least *one BGP-4 Update* message which changes the connectivity of the end-to-end connection. This may occur if the router gets advertisement for a prefix, which is more specific than the prefix used in the initial routing table, as discussed later or if it receives a withdrawal for the route during the observation period.

Considering the *BGP-4 Update* message structure, the table 1 shows the definition of the discussed metrics:

<i>General inter-domain stability metric</i>	<i>Description of BGP-4 Update message parameters found at intermediate border router (exchange points) of connection</i>
Hot	A prefix for the endpoint was present at the beginning of the reservation period <u>and</u> one or more withdrawals with NRLI=prefix appeared <u>or</u> announcements with NRLI for a more specific prefix containing the endpoint address appeared during the observation period <u>OR</u> an announcement for an NRLI leading to an endpoint appeared on a router which initially had no prefix for that endpoint installed in the routing table.
Warm	A prefix for the endpoint was present at the beginning of the observation period <u>AND</u> announcements with NRLI=prefix appeared <u>AND</u> no withdrawals with NRLI=prefix appeared <u>AND</u> no announcements with NRLI for a more specific prefix containing the endpoint address appeared during the observation period
Cold	A prefix for the endpoint was present at the beginning of the observation period <u>AND</u> no announcements with NRLI=prefix appeared <u>AND</u> no withdrawals with NRLI=prefix appeared <u>AND</u> no announcements with NRLI for a more specific prefix containing the endpoint address appeared during the observation period.

Table 1: Definition of general inter-domain stability metrics for end-to-end connection

2.3 Extended inter-domain route stability of end-to-end connection

The frequency of the following kinds of *BGP-4 Update* messages or sequences of them obtained by the intermediate routers of the end-to-end connection in relation to the time window could be considered for connection stability analysis:

- *Route change* is defined by sequence of BGP-4 Updates consisting of a withdrawal for the route of the connection followed by announcements which carry a route (e.g. AS path attribute) for the end-to-end connection which differs from the previous route.
- *Implicit withdrawal* is when the destination for the end-to-end connection is reachable and the announcement carries a route which is different from the current end-to-end connection route..
- *Route Flapping* is defined by number of sequenced “flaps”. “Flap” is a sequence of BGP-4 Updates consisting of withdrawals for the route of the end-to-end connection followed by rapid announcements which re-announces the withdrawn route. Detection of route flapping could be a sign for connection quality degradation, therefore carefully integration of strategies for flap damping is required [12].
- *Route duplicate announcement* is when the destination for the end-to-end connection is reachable and the announcement carries a route which is identical to the current route of the connection. [20] reports that more than 40% of all BGP-4 announcements are duplicates. The study of the BGP-4 duplicates is important in order to reduce excessive BGP-4 traffic of the connection.

3 InterRoute technology

3.1 Design

InterRoute is developed with the goal to:

- (1) provide an interface between public and/or private BGP-4 route and topology repositories
- (2) calculate inter-domain topologies based on analysis algorithms
- (3) retrieve significant inter-domain routing events
- (4) compute inter-domain route stability metrics.

InterRoute implements a processing chain for inter-domain topology information along the actual path measured by the topology discovery tools such as traceroute.

The design of InterRoute and particularly its information exchange is shown in Figure 1.

The topology discovery tools (e.g. traceroute) render the connection as a set of routers, which are used to infer the Internet exchanges (e.g. border router) along the active path.

The use of this strategy is confirmed by other projects which combine different sources of topology data to provide accurate topology maps of the Internet [5].

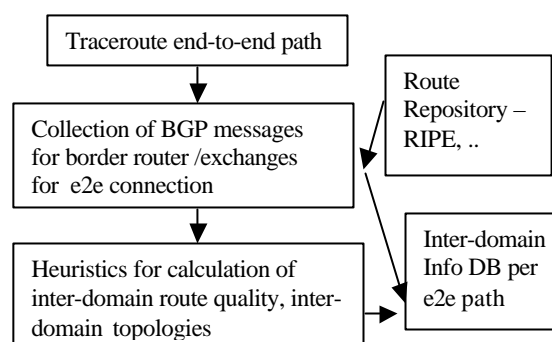


Figure 1: InterRoute design

InterRoute selects the BGP-4 Update messages for border router/ exchanges of the inter-domain path from public route repositories. Public route repositories have been established by the Oregon Route Views [2] project, the RIPE [24] to store the BGP-4 routing activities of selected border routers/Internet exchanges. Currently, InterRoute is based on the RIPE Routing Repositories [24]. RIPE, an European IP network management collaboration, collects fine-grained routing data by archiving BGP-4 announcements from multiple locations in a MySQL data base [24].

The RIPE Repositories are accessed through a web based interface which links the IP address of the border routers/exchanges included in the end-to-end connection with the collected routing information. All route repositories both public and private are based on the standard format defined in Merit's Multi-threaded Routing Toolkit MRT [4]. While more peers would add incrementally to the data coverage, current routing data archives (Route Views and RIPE) provide sufficient data to perform macroscopic analysis [25].

InterRoute collects all inter-domain routing information available for the time window of the observation and stores them into appropriate topology data base. A snapshot of the routing table immediately before the observation time window is obtained and forms the basis for the rest of the computations.

The evolution of the routing table is obtained from the routing update collections which are available for the significant route repositories. This results in a potentially large database. Routing table snapshots - approximately 400k routes long and 600k-1M updates per day are common. InterRoute filters the updates which can be significant for an end-to-end scenario involving two end systems, implementing heuristics considering that border routers work on a longest prefix match basis when selecting the output interface for a given IP address. Upon a route withdrawal affecting the longest prefix match, the router will use the next longest prefix match locally, until the routing

protocol establishes a new best path. This new best path is likely to follow a route which is close to the original one. InterRoute takes this situation into account and stores all possible prefix matches in the database. Despite this relaxed filtering strategy, this reduces the amount of storage required for the topology database by several orders of magnitude, yielding around or less than 100 significant records for a 24 hour observation period.

3.2 User Interface

InterRoute is based on Graphical User Interface (GUI) based on which the user can supply parameters to obtain the inter-domain topology and inter-domain route stability metrics for a given time period.

The inter-domain topology for the end-to-end connection detected by InterRoute shows the autonomous systems included in the end-to-end connection together with their links, yielding a clear picture of the relations among them. The endpoints are also shown as elements of the autonomous systems which connect them to the Internet. Alternative inter-domain topologies are discovered.

The border routers detected by InterRoute are assigned to their autonomous system and included in the 3D visualisation. The 3D visualisation allows border routers to be assigned routing messages and a inter-domain stability metrics. InterRoute tool has been demonstrated in an inter-domain QoS monitoring scenario between Salzburg and Madrid.

Figure 2 shows an Autonomous system which was obtained as part of the topology between Madrid and Salzburg together with the inter-domain routing stability metrics for route change analysis.

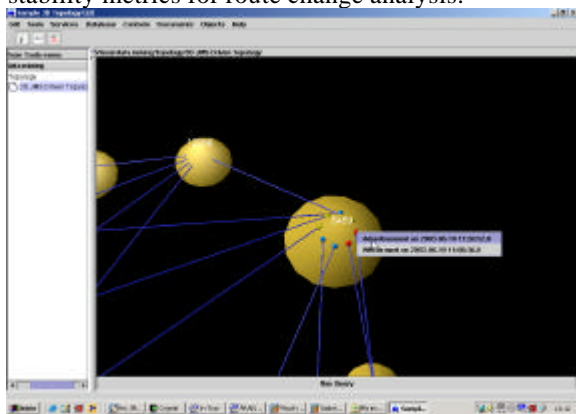


Figure 2: Inter-domain topology and basic inter-domain stability metrics derived by InterRoute

Cold routers are painted blue, warm routers yellow and hot routers red. Inter-domain route quality analysis is based on this metric. Inter-domain routing events are assigned to their corresponding routers. The routing

messages are revealed to the network manager when the object is right-clicked on the GUI. As shown, a list of routing events appears, stating the nature and timestamp of the routing events attached to a router.

4 InterRoute integration in the INTERMON toolkit

The INTERMON project [1] is developing an integrated platform for monitoring, visualisation and modelling of inter-domain environments around an integrated policy based database of different tools.

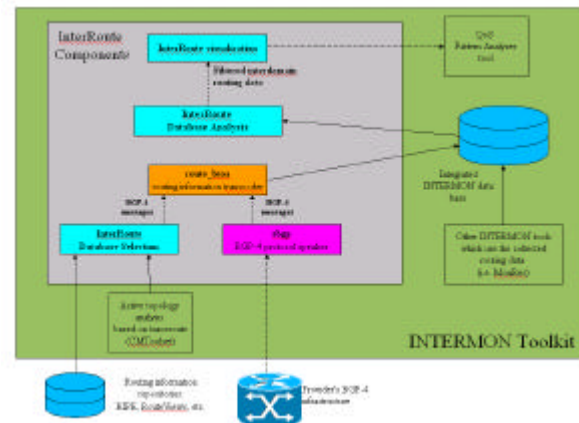


Figure 3: InterRoute integration in INTERMON

The INTERMON toolkit is being designed with several application scenarios in mind based on interaction of different tools for inter-domain QoS analysis. InterRoute is applied in an inter-domain scenario, where the end-to-end QoS between two end systems in the Internet is being monitored, modelled and visualised. As shown in Figure 3, InterRoute is used in close relationship with the QoS Pattern Analyser and CMToolset Topology Discovery, which implements one of the QoS measurement components of the INTERMON toolkit.

InterRoute could use the topology traceroute data collected by the CMToolset in the traceroute data base for a flow. InterRoute yields a collection of routing events, whose timestamps are used in the QoS visualisation process as marks for potentially interesting moments in the measurements.

Apart from the interaction with the CMToolset, the different alternative inter-domain topology generated by InterRoute for an end-to-end connection in inter-domain environment are also used by the simulators integrated in the INTERMON toolkit.

InterRoute could be used with the Spatio-temporal Pattern analyser of INTERMON toolkit in order to discover QoS patterns based on inter-domain topology changes and inter-domain route quality [14].

The interaction of InterRoute with Spatio-temporal analyser is used to obtain the QoS parameter pattern, e.g. sequence of time series data describing measured QoS parameters, in the time window, in which inter-domain routing event or topology change happen. This allows to detect possible dependencies of QoS parameter and inter-domain routing stability.

5 Conclusions and Future work

Routing stability is key to providing a predictable end-to-end-service through the Internet. Metrics to describe inter-domain route stability of an end-to-end connection for management purposes based on analysis of *BGP Update* Message has been proposed in this work and implemented in the InterRoute tool.

Inter-domain stability analysis based on the general stability metrics reflecting route changes and InterRoute tool has been tested in an realistic end-to-end scenarios (Madrid – Salzburg) using data from RIPE Repository for an intermediate Internet Exchange (London PBX) along the path. First trials confirm the basic concept behind this metric.

Current work is based on integration of the other metrics to study sources for inter-domain instability of end-to-end connection.

Further focus is to include metrics for BGP-4 convergence and to study relationships of inter-domain route stability metrics and QoS parameter patterns obtained in inter-domain environment.

6 References

[1]INTERMON IST-2001-34123 <http://www.ist-intermon.org>
[2]Oregon Routeview <http://antc.uoregon.edu/route-views>
[3]RIPE - <http://www.ripe.net>
[4]MRT - <http://www.mrtd.org>
[5] N. Spring, R. Mahajan, D. Wetherall , Measuring ISP Topologies with Rocketfuel, ACM SIGCOMM, 2002
[6] J.Davis, S. Ilnicki, V. Kanevski, L. Tatman, A. Tudor, Inter-domain routing anomalies, Agilent Labs Technical Report, Oct. 15, 2001
[7]C.-N. Chuah, S. Bhattacharyya, C. Diot, Measuring I-BGP Updates and Their Impact on Traffic, SPRINT ATL Technical Report TR02-ATL-051099, 2002
[8]D. Hetzer, I. Miloucheva, U. Hofmann, J. Quittek, F. Saluta, Integrated Information System for Inter-Domain QoS Monitoring, Modelling and Verification, EURESCOM, Heidelberg, 2002
[9]Griffin, T.G., Wilfong, G., An Analysis of BGP Convergence Properties, SIGCOMM 1999

[10]R. Mahajan, D. Wetherall, T.Anderson, Understanding BGP Misconfiguration, ACM SIGCOMM, Aug 2002
[11]C.Boutremans, G.Iannaccone, C.Diot, "Impact of Link Failures on VoIP performance",NOSSDAV, 2002
[12]Z. M. Mao, R. Govindan, G. Varghese, R. Katz, Route Flap Damping Exacerbates Internet Routing Convergence, SIGCOMM 2002
[13]P.A.Aranda Gutiérrez, I. Miloucheva, Analysis of end-to-end QoS behaviour in inter-domain environment, IPS Workshop, Salzburg, March, 2003
[14]I. Miloucheva, U. Hofmann, P.A. Aranda Gutiérrez, Spatio-temporal Analysis in large scale Internet environment, MIPS Workshop, Napoli, Italy, November 18-21, 2003
[15] Loring Wirbel, Packet Design develops route-monitoring tool, May 20, 2002
[16] Yihua Liao and Ke Zhang, "BGP Behavior Monitoring and Analysis", ECS 289M (Advanced Topics in Computer Security) project, Spring 2002
[17] L. Wang, X. Zhao, D. Pei, R. Bush. D. Massey, A.Mankin, S. F. Wu, L. Zhang, Observation and analysis of BGP Behaviour under Stress, ACM SIGCOMM Internet Measurement Workshop, 2002
[18] A. Govindan, A. Reddy, An Analysis of Inter-domain Topology and Route Stability, Proceedings of IEEE INFOCOM 97, 1997
[19] J. Vähäpassi, Internet Routing Stability, <http://keskus.hut.fi/opetus/s38130/k00/Papers/Topic17-Stability.doc>
[20] Y. Jiang, A. Doria, D. Olsson., F. Pettersson, Inter-domain Routing Stability Measurement, HPSR2002 Conference Proceedings, Kobe, Japan,2002
[21] V.Paxson, End-to-end routing behaviour in the Internet. ACM SIGCOMM 96, August 1996
[22] D. Pei, L. Wang, D. Massey, S. F. Wu, L. Zhang, A Study of Packet Delivery Performance during Routing Convergence", International Conference on Dependable Systems & Networks (DSN), June 2003
[23] H. Berkowitz et al, Terminology for Benchmarking BGP Device Convergence in the Control Plane, draft-ietf-bmwg-conterm-05.txt, 2003
[24] RIPE-NCC, "Routing Information Service Project (RIS) Project". <http://www.ripe.net/ripenncc/pub-services/np/ris-index.html>
[25] Routing Analysis and Peering Policy for Enhancing Internet Performance and Security, CAIDA, <http://www.caida.org/projects/ncs/nsfSpec02.html>
[26] Y.Rekhter, T.Li, "A Border Gateway Protocol", RFC 1771, BGP Version 4, March 1995
[27] C.Labovitz, R. Malan, F.Jahanian. "Origins of Internet Routing Stability", INFOCOM, 1997
[28] C. Labovitz, A. Ahuja, F. Jahanian, "Experimentatl Study of Internet Stability and Wide Area Network Failures", FTCS, 1999