

# BGP-4 Protocol Patterns and Their Impact on QoS Behavior

Pedro A. Aranda Gutiérrez

*Telefónica I+D, Spain*

## Abstract

*This paper is aimed to describe the INTERMON approach for inter-domain BGP-4 analysis based on pattern matching on public BGP-4 route repository data.*

*An overview of different BGP-4 protocol analysis methodologies is presented and compared with the approach followed in the INTERMON project.*

*The paper presents the usage of the pattern analysis technology in realistic inter-domain performance management scenarios, where BGP-4 patterns help delimiting the time windows where QoS pattern abstraction is used to evaluate the QoS behaviour from point of view of applications.*

*Finally, based on INTERMON experience, complex causal data mining approach for abstraction of inter-domain routing impact on end-to-end QoS behaviour using BGP-4 protocol patterns is discussed.*

## 1 Introduction

The Internet is known as the Network of Networks. This reflects the fact that different Internet Service Providers interconnect their networks to provide a worldwide infrastructure. This interconnection takes place at a level which preserves the autonomy of the different providers when managing their domains. BGP-4 is the Interdomain Routing Protocol and governs the peering relationships between providers. Contrary to routing protocols used with providers' networks, BGP-4 was designed with the following goals in mind:

- To provide a minimum level of peer validation which reflects the peering relationship between providers
- To be able to scale well beyond the amount of routing information handled by interior routing protocols
- To minimise the exchange of information between peers
- To hide the internals of the autonomous systems.

This yielded a very flexible protocol definition which supports nearly all kinds of networking environments. As mere examples of the applications of BGP-4, just two should be mentioned:

- The transition from IPv4 to IPv6 will be extremely simplified by the fact that both Internet Protocols use BGP-4 as their Interdomain Routing Protocol

- One of today' most popular value added services on top of MPLS backbones is MPLS-VPN's, which can be easily deployed in program because of the use of BGP-4.

But as with any simple protocol, BGP-4 also has some drawbacks which might lead to problems in day to day network management. As pointed out in [2], the behaviour of the protocol itself depends on the (unpredictable) order of the updates which arrives at each and every router for a given section or subsection of the Internet.

Identifying patterns for the update sequence which can help the operator to trace back the history of a given connection in case of SLA violations. The nature of the update sequences will hint the nature of the SLA violation and reciprocally, having information on the kind of SLA violation, the operator should be able to trace back the sequence of BGP-4 events which led to it.

## 2 State of the Art in BGP-4 Protocol Pattern Analysis

BGP-4 protocol analysis started based on the experience of establishing BGP-4 Route Repositories by projects like Oregon Routeviews [3] and the RIPE-NCC RIS [4]. These projects collect BGP-4 updates from selected points in the Internet and store them in repositories which can later be accessed by researchers. Using a common format for the data stored in the different repositories has helped this research [5],[6].

Most recent insight on BGP-4 protocol response in normal situations and under stress and Internet network evolution has only been possible because these data are being made available.

Mostly major incident analysis based on bulk parameters.[7], [8], [9] analyse well known distributed Denial of Service incidents in the Internet. They analyse BGP-4 activity during DDoS incidents, examining the number and distribution of BGP-4 updates exchanged and their impact on the routing tables and router performance. This bulk parameter analysis techniques have also shown positive results in detecting mayor BGP-4 configuration conflicts in the Internet[10].

Other research projects study the evolution of the connectivity of autonomous systems and its evolution in time[11], [12]. These tools concentrate more on the visual representation of the evolution of the connectivity between autonomous systems analysing the BGP-4 route repositories. Heuristics are applied on the update

attributes (mainly the AS\_PATH attribute) to build the Autonomous System Connectivity Graph.

### 3 InterRoute: BGP-4 analysis in INTERMON

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

#### 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[12].

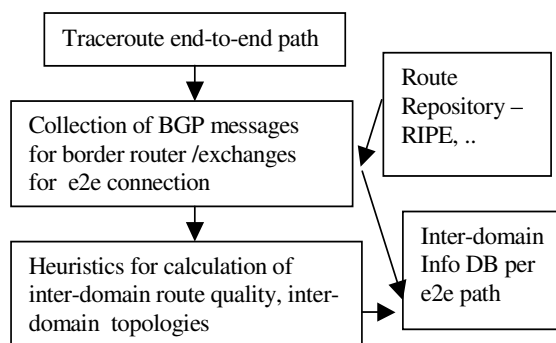


Figure 1: InterRoute design

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[4]. RIPE, an European IP network management collaboration, collects

fine-grained routing data by archiving BGP-4 announcements from multiple locations in the Internet.

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 [5]. While more peers would add incrementally to the data coverage, current routing data archives (Route Views[3] and RIPE[4]) provide sufficient data to perform macroscopic analysis.

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

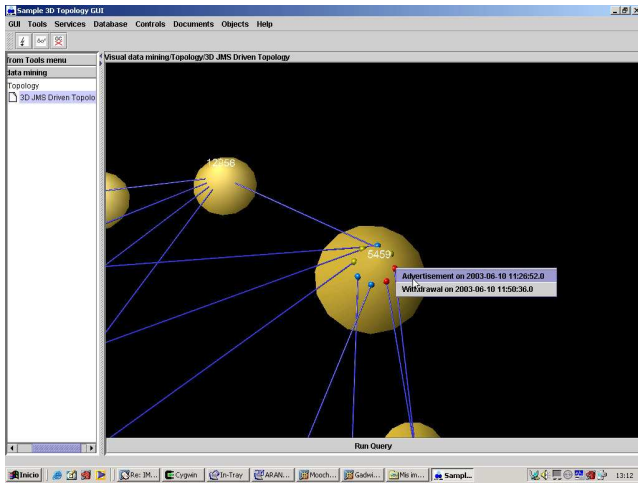
InterRoute is based on Graphical User Interface (GUI) based on which the user can supply parameters to obtain the inter-domain topology and inter-domain route stability metrics for a given time period.

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 a prototype implementation: an Internet Exchange is represented as an Autonomous System. The data where obtained as part of a topology/QoS measurement scenario between Madrid and Salzburg . This topology GUI shows the relationships between the different autonomous systems 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 simple metric. The GUI assigns inter-domain routing events to the routers where they were detected.

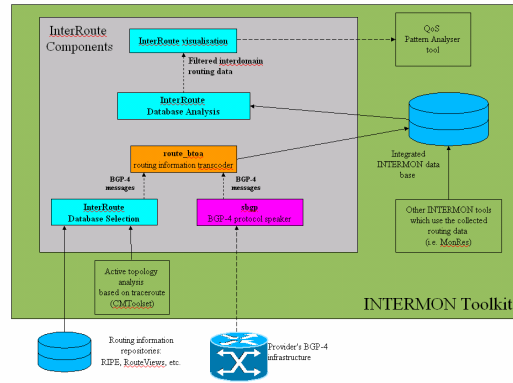
Since cold routers are those which do not have any related BGP-4 events, the user is directed towards the interesting routers with this simple colour code.

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 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.



**Figure 3: InterRoute integration in INTERMON**

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.

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 BGP-4 protocol patterns detection

InterRoute takes the view of the end to end path under study as seen from an intermediate AS. A heavy distillation process needed to obtain the routes involved in the end to end scenario. This results in a manageable set of routes and updates which explain the evolution of the

end to end scenario as seen by the border routers of the intermediate AS.

Following information is available for each BR involved in the scenario:

1. Route which was available at the beginning of the observation period, with prefix and BGP-4 attributes, such as AS\_PATH, etc.
2. BGP-4 Updates, including type (Advertisement or withdrawal), prefix and BGP-4 attributes such as AS\_PATH, etc. when applicable

These series of routing data derive a complete view of the topology of autonomous systems in the Internet that implement the end to end connection, including alternate paths which are programmed as backup routes and its evolution during the period under study.

Applying pattern analysis techniques to this dynamic graph will help fix in time more complex events with potential impact on QoS. These events can be used in several ways:

1. as the entry points for detailed analysis of measurements
2. as a hint to distribute QoS probes in the network

## 5.1 The BGP-4 heuristics

Analysing the routing event series at a given border router, different levels of patterns appear. The first level is the distribution in time. Clusters of events can be identified and since any event implies extra processing power in the router or eventual loss of connectivity, the impact on end to end QoS can be studied. Moreover, in a complex scenario for visual data mining, these events are good triggers for QoS analysis.

The second level of analysis includes the use of heuristics applied to the event series, which allow the discovery of complex network (mis)management manoeuvres to be unveiled. At the moment of writing this paper, several heuristics are being explored.

The most significant events which can be identified are:

1. *Route flapping*
2. *Route dampening*
3. *Manual trial-and-error approach to traffic engineering*

We will discuss the heuristics in the following subsections.

### 5.1.1 Route flapping

Route flapping in the absence of route dampening will manifest itself as an alternate series of advertisements and

withdrawals, which ends when the flapping event is detected. Depending on the origin for such a route flapping event, the time interval between routing updates varies. Short time intervals have been observed in route flapping incidents involving a hardware defect in an interface board. Longer time intervals have been observed when routing loops are generated in certain network configurations

### 5.1.2 Route dampening

In order to reduce route flapping and its negative impact on route stability in the Internet and core router performance, route dampening is introduced as a means to penalize unstable routes. The following figure is taken from a BGP-4 tutorial [13] and illustrates the route dampening mechanism:

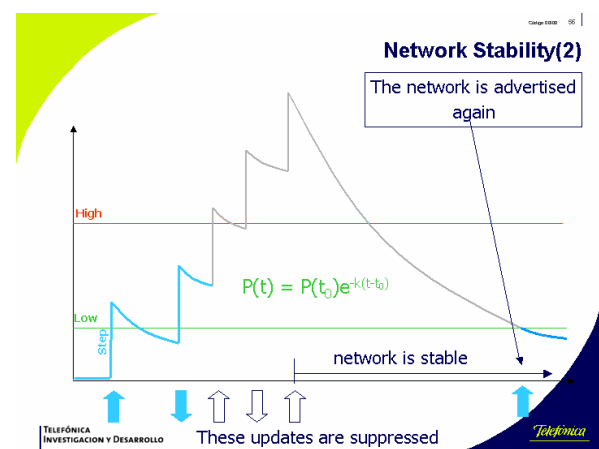


Figure 4: Route dampening illustrated

Route dampening is applied to updates referring to a specific NRI and originated by a specific router, regardless of their nature (i.e. whether they are advertisements or withdrawals).

Thus, the heuristic to detect routing problems suppressed by route dampening can be expressed as follows:

Find short term bursts of updates referred to a specific NRI and originated by one specific router, followed by longer periods without routing activity. All bursts start with an advertisement and end with a withdrawal.

### 5.1.3 Manual trial and error traffic engineering

Manual trial and error traffic engineering occurs when the operator has poor information on the network topology from the BGP-4. When setting the attributes of a BGP-4 advertisement for a given prefix through a backup link, traffic to the prefix is delivered through this link despite being a backup link.

The operator will notice and try to correct the attributes in order to stop the traffic flow. The candidate BGP-4 attribute in this case is the AS\_PATH, which is altered using the *'as-path prepend'* technique.

Figure 5 shows such a sequence of trial and error traffic engineering operations which has to be finally aborted (last three updates between 15:18 and 15:32).



Figure 5: Trial and error Traffic Engineering

Examining this sequence under the optics of the route dampening heuristic proposed in 5.1.2, it is easy to see that it conforms to both heuristics. This is a well known drawback of manual trial and error traffic engineering, which should be highlighted in any training program for network operators handling Internet core routers [13].

The heuristic to detect instances of trial and error network engineering can be described as follows:

Find a sequence of advertisements referred to a specific NRI and originated at one specific router, in which the attributes are changed in order to obtain.

## 6 Further work

The three heuristics described here are not too complex to implement in network management systems, provided BGP-4 routing data repositories are available. Further work in this direction is directed towards automatic detection of human configuration errors with different level of severity. This heuristic is most likely going to need the joint use of QoS monitoring data and BGP-4 routing data.

Configuration errors in general are extremely damaging to the Internet. Often, configuration errors go by undetected for enough time to cause damage to businesses relying on the Internet. A tool which is able to detect configuration errors is a sensible complement for network operators.

## 7 Conclusions

The use of real routing information in network management opens up new possibilities of network supervision. The INTERMON project is using public BGP-4 routing data repositories to analyse the stability of

end to end connections in the Internet, when such data are available. The promising results should be a stimulus for Network Operators to establish their own private BGP-4 routing update storage infrastructure to analyse their own routing domain or domains. This information should be complemented with information available in public routing information repositories.

## 8 Acknowledgements

The work described in this paper has been carried out under the scope of INTERMON, IST project 2001-34123, co-funded by the 5<sup>th</sup> Framework of the EC:

## 9 References

- [1] INTERMON-IST-2001-34123; <http://www.ist-intermon.org>
- [2] Tim Griffin: "An open problem in Internet Routing - Policy Language Design for BGP", [http://www.cambridge.intel-research.net/~tgriffin/talks\\_tutorials](http://www.cambridge.intel-research.net/~tgriffin/talks_tutorials)
- [3] The Oregon Routeviews project
- [4] RIPE-NCC RIS Routing Information Service ; <http://www.ripe.net/ris/>
- [5] The Multithreaded Routing Toolkit; <http://www.mrtd.net>
- [6] Ayesha Gandhi, Analysis of BGP Routing Tables, <http://www.ecse.rpi.edu/Homepages/shivkuma/teaching/fall2002/exp2002-bgp.pdf>
- [7] James Aldridge, Andrei Robashevsky, Henk Uijterwaal, Arife Vural and René Wilhelm. Effects of the DDoS Attack on the RIPE NCC 27 February 2003; <http://www.ripe.net/ttm/worm/ddos.html>
- [8] James Aldridge; Routing Information Service Sapphire/Slammer Worm - Impact on Internet Performance; 7-February-2003; [http://www.ripe.net/ttm/worm/ris\\_index.html](http://www.ripe.net/ttm/worm/ris_index.html)
- [9] TONI FARLEY, PATRICK MCDANIEL and KEVIN BUTLER; A Survey of BGP Security Issues and Solutions; <http://www.patrickmcdaniel.org/pubs/td-5ugj33.pdf>
- [10] Cengiz Alaettinoglu: RIPE/RIS Project: BGP Analysis-CIDR at work, NANOG October 2001
- [11] Giuseppe di Battista, Federico Mariani, Maurizio Patrignani, Marzio Pizzonia, Archives of BGP Updates, Integration and Visualization; IPS'2003 Proceedings
- [12] The Hermes Project, <http://www.dia.uniroma3.it>
- [13] Pedro A. Aranda Gutiérrez: BGP-4 routing tutorial, IPS'2003, Salzburg