1. Artur KRUPA[1], 2. Dariusz MAKOWSKI[1,2]

Warsaw University of Life Sciences (1), MEVSPACE SP. z o.o. (2)
ORCID: 1. 0000-0001-8583-9873

doi:10.15199/48.2025.03.46

An optimisation of Border Gateway Protocol on inter-op links

Abstract. *The article describes an original and dedicated approach to analysing inter-operator traffic data using customised analytical solutions, including a data collection and optimisation tools to reduce the problem of non-intuitive path planning in BGP protocol. The system is based on information about flows from current routes, traffic management data within the local network and at the border of dedicated operators and creating a network of connections verifying the current consumption of output gates on infrastructure links.*

Streszczenie. *Artykuł opisuje dedykowane podejście do analizy danych o ruchu międzyoperatorskim z wykorzystaniem autorskich rozwiązań analitycznych, w tym systemu gromadzenia danych oraz narzędzi optymalizacyjnych redukujących problem nieintuicyjnego planowania ścieżek w protokole BGP. System bazuje na informacjach o przepływach z bieżących tras, danych o zarządzaniu ruchem w sieci lokalnej i na granicy dedykowanych operatorów oraz tworzeniu sieci połączeń weryfikujących bieżące zużycie bram wyjściowych na łączach infrastrukturalnych. (**Optymalizacja protokołu Border Gateway Protocol na łączach międzyoperacyjnych**)*

Keywords: network optimisation, BGP protocol, routing optimisation, inter-operator links
Słowa kluczowe: optymalizacja sieci, protokół BGP, optymalizacja routingu, łącza międzyoperatorskie

Introduction

The Border Gateway Protocol (BGP) is a fundamental part of the modern Internet infrastructure [1-3], having a key role in managing traffic between different autonomous networks (ASNs). It was introduced in 1989 as a successor to the Exterior Gateway Protocol (EGP) in response to the growing need for global communication. BGP enables the exchange of information about available routing paths between network operators, which allows for the efficient distribution of data across the vast and distributed infrastructure of the Internet.

The basic functionality of BGP is that it allows network operators (especially those managing large, independent networks) to communicate and exchange information about the best routes to send data. BGP operates at the application layer level, using TCP as the transport protocol, which ensures the stability and reliability of data transmission between different autonomous systems (ASs). In practice, BGP decides the routes that data packets take across the global network, choosing paths based on numerous attributes, such as path length [4], operator routing policies, and economic priorities [5].

Despite its fundamental importance for the operation of the Internet, BGP has significant limitations [6-9] that are becoming increasingly apparent in the face of the dynamically growing needs of modern networks. This protocol was designed nearly 35 years ago and was not intended to manage the complexity and scale of today's network connections. Its mechanisms must consider modern requirements for low latency, high throughput, and dynamic traffic management, crucial for real-time applications such as financial services, multimedia, and entertainment [10,11]. Moreover, this "trust-based" protocol is vulnerable to various types of attacks and manipulations, which seriously threatens the stability of the global Internet. In response to these challenges, this paper proposes a dedicated approach to analysing and optimising of inter-operator traffic that considers the contemporary realities and requirements of 21st-century networks. The described analytical system is based on advanced data collection, analysis tools, and possible optimisation points, allowing more intelligent and adaptive traffic management decisions. This allows for minimising the problems related to non-intuitive route planning and optimising data transmission in inter-operator networks while considering the imperfections of the BGP protocol.

Limitations

One of the major limitations of BGP is its reliance on manual configuration and management, which often leads to suboptimal routing decisions [12]. In practice, this means that the routes selected by BGP may be far from the shortest possible, resulting in increased latency and inefficient use of network resources. Furthermore, BGP lacks adaptive mechanisms that would allow dynamic route adjustments in response to changes in network load or real-time failures. These deficiencies become particularly acute in modern networks, where stable and efficient connections are crucial for many mission-critical applications [13,14].

An example of these limitations is simply illustrated in the diagram shown in Fig. 1. The diagram represents an example connection diagram with a source server (SSRV) and a destination server (DSRV). Example operator nodes are marked with ASN numbers. The connections between them are described with bidirectional arrows indicating communication in both directions for network traffic. The values above them indicates example Latencies times given in milliseconds. Network has 3-layer communication channel based on the BGP protocol (**L1**: ASN 15001 and 3003, **L2**: 323, 6500 and 123, **L3**: 121 and 211). Based on BGP methodology, the Red Path indicates the shortest possible path from SSRV to DSRV. The Green Path provides the shortest response time and latency despite having more nodes than the default selection. The diagram is based on first attempt to network optimisation process made in our laboratory – ASNs are randomly generated to avoid making the operator's nodes public and compromising the company's confidential information.

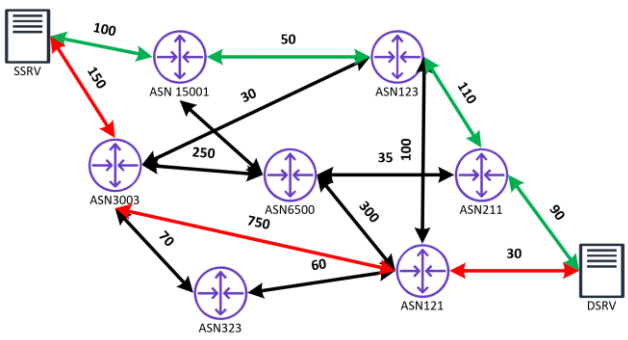

Fig. 1. Simple and example communication path between Source (SSRV) and Destination (DSRV).

However, the problem arises when analysing the network from the point of view of services, which currently require exceptional performance, reliability, and optimisation from the point of view of maintaining communication. The BGP protocol, in its simplicity, which is undoubtedly an advantage, will select the shortest path for the example route in the diagram mentioned above. Table 1 presents selected connections for routes from many possible ones:

Table 1. List of selected paths for BGP and their latencies

Path (AS Numbers)	Hops	Latency [ms]
15001 123 211	4	350
15001 6500 121	4	680
3003 6500 121	4	730
3003 323 121 123 211	6	580
3003 121	3	930

In the case of the BGP protocol, the default route from the available ones will be the one with the fewest hops. Therefore, the last one on the list will be selected. As you can see from its length, it is the shortest possible connection, but it is burdened with a link transmission time of no less than 930 milliseconds, which is almost 1 second. Considering even the longest example route (the penultimate on the list), it is almost half the shorter, although it contains three more hops. Interestingly, the shortest route (almost one-third of the "most optimal" time according to BGP) does not have the largest number of hops at all. Therefore, it is clear that the protocol does not guarantee optimisation but the selection of the shortest route.

Of course, there are methods for planning routes and algorithms that remember the last paths, but they do not consider any of the parameters ensuring effectiveness. For such needs, BGP has been equipped with a list of attributes that determine the priority of importance and significance from the point of view of route selection. As can be seen in Table 2, these include the weighted factor, local exit gateway selection preference, local edge router path, AS preference, origin code, neighbourhood code, preferred mode, the shortest route to the next hop, oldest known path, device classification, and IP addressing.

Table 2. List of attributes and their priorities in BGP

Attribute	Priority
Weight	1
Local Preference	2
Originate	3
AS Length	4
Origin Code	5
Neighbor MED.	6
eBGP over iBGP	7
Shortest IGP to BGP NH	8
Oldest path	9
R-ID	10
Neighbor IP	11

However, these parameters are still only optional parameters and do not take into account the routes in the global network, which means that while we can impose traffic from the source (SSRV) or on the side of the operator with whom we are directly connected (it is our superior operator – partially in the examples given here ASN 15001 and ASN 3003), further traffic is already independent. Any enforcement will be irrelevant if the route selection protocol directs to its "shortest" path.

To find the best path from the optimisation point of view, it seems good to use a greedy algorithm (such as Dijkstra's algorithm), which is also present and used in local traffic routing within the so-called IGP (Internal Gateway Protocol), i.e. in the OSPF (Open Shortest Path First) protocol present in local networks. Considering the optimisation project in the global internet network, this algorithm does not meet its assumptions, especially when problems such as:

- **Scalability**: The problem of a much larger number of network devices
- **Traffic rules**: Inter-operator links include decision-making traffic rules that are not strictly based on the shortest paths
- **Network structure dynamics**: Variability in time and conditions does not allow for multiple recalculations of traffic and routes
- **Loops**: BGP takes into account the risk of loops and allows for their prevention, which the algorithm does not have in the standard method of operation
- **Analysis time**: BGP requires to consider longer path processing, offering improved stability of operator devices, for which the algorithm is again not prepared

Optimisation

To meet these challenges, it is necessary to introduce advanced analytical tools that can be placed locally in the operators' infrastructure, enabling ongoing analysis and optimisation of data flow. During the analysis and research, three possible technologically acceptable methods were developed to optimise the solution, which can contribute to improving the efficiency of network traffic management [15].

A) Local Analysis and Dynamic Redirection (LADR)

The first optimisation stage is the local placement of analytical tools responsible for the current analysis of network traffic in specific segments of the operator's network [16]. These tools can analyse data flows in real-time, identifying potential problems such as congestion or inefficient routes. Based on the collected data, the system can dynamically redirect traffic, choosing optimal paths that minimise delays and maximise the use of available network resources.

During the research, dedicated devices with physical access to nodes were introduced to the operator's network to ensure uninterrupted communication with the physical network collecting data. Data was collected from devices in the laboratory network of the locally run server room and from partner operators with a direct connection to the server room. An example diagram is shown in Fig. 2.

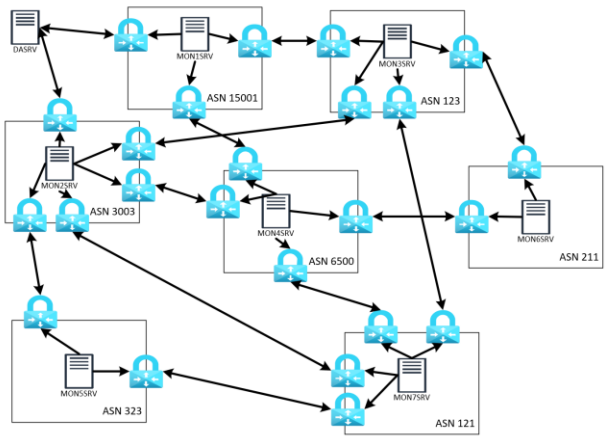

Fig. 2. Reporting data with LADR from monitoring systems (MON1SRV to MON7SRV) to Data Analysis Server (DASRV).

Based on a Fig. 2 (while maintaining the previously provided on a Fig. 1 ASN numbering) monitoring systems (marked MON1SRV to MON7SRV) were installed within the inter-operator networks, whose task was to report Latency times between individual operator nodes, the so-called path "1-to-1", and then send data to the data acquisition server (DASRV). Padlocks symbolize direct output from an operator to another operator (inter-operator link). Therefore, arrows from monitoring systems verify the network status in each direction. The operations verifying the current link performance times were repeated cyclically at set time intervals. These systems had a direct output on the operator's links, thanks to which they could examine the network for many parameters independently and then return the results to the mother server. In the case of a server at a given operator on a given link, the monitoring system does not verify this path but only all others. This made it possible to avoid connecting existing nodes each other and verify already existing paths.

Using such a solution is very expensive from the point of view of systems and tools. Inter-operator compatibility and maintaining continuity of communication are necessary. The need for constant monitoring and updating of optimisation algorithms to adapt to changing traffic patterns is the basis for the existence of the solution. However, such a method certainly provides increased traffic management flexibility, thanks to the possibility of dynamically adjusting routes in response to changing network conditions. In addition, reducing delays and improving service quality through optimal use of available bandwidth ensure optimal routing throughout the network from source to destination. This solution allows for effectively constructing a network of connections and determining, based on expected parameters, which ASN points are optimal for a given case.

B) Local Traffic Decision Points (LTDP)

The second optimisation step is to implement local decision points that integrate network traffic data and make decisions on managing data flows based on them. These points can be located at the interface of operators' networks, where they analyse incoming and outgoing traffic, identifying optimal routes and potential threats [17]. Unlike central management systems, LTDPs can react faster to network traffic changes, increasing operational efficiency.

An undoubted advantage of the solution may be quick reactions to local changes in network traffic, which allows for immediate optimisation decisions. Reducing the load on central traffic management systems allows for a more balanced distribution of tasks between different infrastructure elements. However, it is necessary to synchronise local decisions with the central traffic management policy must also be considered, which can lead to conflicts in the event of data inconsistency. There are also potential difficulties in integrating existing traffic management systems, especially in the case of heterogeneous network infrastructures.

C) Automated Local Traffic Optimisation (ALTO)

The third step of optimisation is the introduction of automated traffic optimisation mechanisms that operate locally without the need for operator intervention [18]. These systems use advanced machine learning algorithms that analyse historical and current network traffic data and automatically optimise routes and resource allocation. This allows current traffic management, forecasting future traffic patterns, and appropriately preparing the network to handle them.

A key aspect of proper functioning is making decisions based on continuous traffic analysis and updating the situation in the network. The efficiency of devices and precise determination of analysis times allow for avoiding delays. At the same time a centralised monitoring system makes it possible to determine traffic directed to services requiring the most efficient network at a given time. Here, algorithms described by IETF (Internet Engineering Task Force) as LSGO (Large Scale Global Optimisation) Project or [19] can be used.

Implementing such a solution involves standardising network devices' protocol and communication methods, often independent operators. On the one hand, such a solution allows for the automation of processes and dispersion of decision-making, increasing the efficiency of the entire system. Additionally, using the state of current changes and historical data will help to increase the efficiency of transmission, especially in environments requiring efficiency and responsiveness about the basic parameter, which is the length of the route. Unfortunately, the lack of easy standardisation of the solution will not allow for effective training of algorithms. As a result, it can lead to suboptimal decisions of the system and disruptions in the network.

Optimisation results

In the tested laboratory network based on a 3-layer infrastructure, achieving a significant increase in performance in key (from the optimisation point of view) parameters was possible. Analysis of Table 3 shows that the parameters of the non-optimized network (column **Before**) are significantly weaker in all conditions than after optimization using the methods proposed earlier.

Table 3. List of achieved parameters in BGP optimisation

Attribute	Before	Lab	Public Network
Peak Balancing [%]	37	15	28
Target Path Latency [ms]	20	3	10
Packet Loss [%]	0.1	0.0	0.03
Network Utilisation [%]	79	93	87

The parameters were measured based on the current traffic of the server room over the last year, averaging the results depending on the type of transmission and data while maintaining the BGP protocol standard rules. In the case of the internal tests (column **Lab**), which, although they may not have reflected the real state of the network, were conducted under full load, generating network traffic at a level exceeding 100 Gb/s. After implementing the model of monitoring systems in the public network in partner operator links with at least 8-layer infrastructure and at least 20 ASNs, the saturation level of the tested network's links reached a throughput level of almost 70 Gb/s. From the point of view of parameters, the level of delays of optimized paths decreased almost twice (from 20 ms to 10 ms), and the average resource load (unusability) increased by almost 8%. Packet losses caused by device unavailability were reduced by almost three times, which was achieved by ongoing monitoring of device status and advance route planning.

The results achieved in public networks are weaker than in the laboratory, but this is due to the greater "inertia" of the network and a smaller impact on the stability of communication by the installed nodes. The network system was designed to provide an optimal route and summary values depending on the selected data, optimizing traffic. However, it was impossible to modify routes outside the network, which resulted in a decrease in the value

compared to the test data and a summary increase compared to the non-optimized network.

Conclusions

Summarising the conducted analyses and initial tests of local optimisation solutions, introducing advanced mechanisms for analysing and managing network traffic at the level of inter-operator operators is a significant step forward in improving the operation of modern networks. In the face of growing requirements for throughput and minimising delays, traditional approaches based solely on the BGP protocol show their limitations, which requires searching for new methods to improve the efficiency and stability of network connections.

The solutions presented in the article focus on local traffic analysis and dynamic data redirection, which enable more precise and responsive network management. A key element of this strategy is the placement of analytical tools directly in the operator's infrastructure, which allows for ongoing monitoring and optimisation of traffic routing. These systems, through flow analysis and optimisation at the local level, can effectively reduce network load and improve the quality of services.

The currently proposed approaches are the results of local analysis, but when effectively carried out, they ensured that it is possible to implement the solution on a much larger scale. Communication and route building at the level of 15-20 "hops" is still possible to analyse and select the optimal path depending on the traffic at a given time. In addition, it is planned to implement software in network devices installed at operator partners with software dedicated to network monitoring. This will ensure effective data collection and verification of local redirects to develop mechanisms that allow for precise control of network traffic through appropriate enforcement at the BGP protocol level, especially in the context of data retransmission.

The results of these tests will be decisive in the selection of the final solution, which will be implemented on a large scale. This will allow for flexible adaptation to changing conditions and network requirements, crucial for the stable and effective management of modern telecommunications infrastructures. The key element of the proposed solution is a system based on information about flows from current routes and data on traffic management within local networks and at the border of dedicated operators. This system enables the creation of a network of connections that continuously verify the current use of exit gateways on infrastructure links, allowing for more effective data flow management.

The contemporary development of the Internet requires operators to use new tools and technologies that will allow for more effective management of the increasingly complex and dynamically growing network infrastructure. The proposed solutions are a step towards more intelligent and flexible network traffic management, which is necessary to meet the challenges the modern Internet faces. This article aims not only to present new methods of analysis and optimisation but also to inspire further research in this area, which will allow for even more effective use of the potential of existing and future networks.

***Authors**: Ph.D. Artur Krupa, Warsaw University of Life Sciences, Institute of Information Technology, Department of Artificial Intelligence, Nowoursynowska 159, 02-776 Warsaw, Poland, E-mail: artur_krupa@sggw.edu.pl; M.Sc. Dariusz Makowski, MEVSPACE Sp. z o.o., Augustyna Locciego 33, 02-928 Warsaw, Poland*

REFERENCES

[1] Van Beijnum, I. (2002). BGP: Building reliable networks with the Border Gateway Protocol" O'Reilly Media, Inc.".

[2] Stewart III, J. W. (1998). BGP4: inter-domain routing in the Internet. Addison-Wesley Longman Publishing Co., Inc.

[3] Rekhter, Y., Li, T., and Hares, S. (Eds.). (2006). RFC 4271: A border gateway protocol 4 (BGP-4).

[4] Sunita, M., and Mallapur, S. V. (2024). Optimal detection of border gateway protocol anomalies with extensive feature set. Multimedia Tools and Applications, 83(17), 50893-50919.

[5] Flach, T., Papageorge, P., Terzis, A., Pedrosa, L., Cheng, Y., Karim, T., and Govindan, R. (2016, August). An internet-wide analysis of traffic policing. In Proceedings of the 2016 ACM SIGCOMM Conference (pp. 468-482).

[6] Lynn, C. (2000). Secure border gateway protocol. IEEE J. Selected Areas in Comm, 18(4), 582-592.

[7] Smith, B. R., and Garcia-Luna-Aceves, J. J. (1996, November). Securing the border gateway routing protocol. In Proceedings of GLOBECOM'96. 1996 IEEE Global Telecommunications Conference (pp. 81-85). IEEE.

[8] Badawy, M. (2017). Improving Performance of Border Gateway protocol for Large-Scale Autonomous Systems.

[9] Navai, N. (2000). Performance evaluation of the border gateway protocol (Doctoral dissertation, University of British Columbia).

[10] Manzoor, A., Hussain, M., and Mehrban, S. (2020). Performance analysis and route optimization: redistribution between EIGRP, OSPF and BGP routing protocols. Computer Standards and Interfaces, 68, 103391.

[11] Griffin, T. G. (2010, January). The stratified shortest-paths problem. In 2010 Second International Conference on COMmunication Systems and NETworks (COMSNETS 2010) (pp. 1-10). IEEE.

[12] Sunita, M., and Mallapur, S. V. (2024). Optimal detection of border gateway protocol anomalies with extensive feature set. Multimedia Tools and Applications, 83(17), 50893-50919.

[13] Alotaibi, H. S., Gregory, M. A., and Li, S. (2022). Multidomain SDN-Based Gateways and Border Gateway Protocol. Journal of Computer Networks and Communications, 2022(1), 3955800.

[14] Douzet, F., Pétiniaud, L., Salamatian, L., Limonier, K., Salamatian, K., and Alchus, T. (2020, May). Measuring the fragmentation of the Internet: the case of the Border Gateway Protocol (BGP) during the Ukrainian crisis. In 2020 12th international conference on cyber conflict (CyCon) (Vol. 1300, pp. 157-182). IEEE.

[15] Chen, K., and Hu, C. (2011). Border gateway protocol monitoring system can be cost effective. IET communications, 5(15), 2231-2240.

[16] Gottlieb, J., Greenberg, A., Rexford, J., and Wang, J. (2003). Automated provisioning of BGP customers. IEEE network, 17(6), 44-55.

[17] Zhao, X., Band, S. S., Elnaffar, S., Sookhak, M., Mosavi, A., and Salwana, E. (2021). The implementation of border gateway protocol using software-defined networks: A systematic literature review. IEEE Access, 9, 112596-112606.

[18] Fabian, B., Baumann, A., and Lackner, J. (2015). Topological analysis of cloud service connectivity. Computers and Industrial Engineering, 88, 151-165.

[19] Liu, J., Sarker, R., Elsayed, S., Essam, D., and Siswanto, N. (2024). Large-scale evolutionary optimization: A review and comparative study. Swarm and Evolutionary Computation, 101466.