BGP Anomalies Classification using Features based on AS Relationship Graphs

Thales Paiva, Yaissa Siqueira, Daniel Batista, Roberto Hirata Jr. and Routo Terada

November 17, 2021

{tpaiva,yaissa.siqueira,batista,hirata,rt}@ime.usp.br

IEEE Latin-American Conference on Communications - LATINCOM 2021

Computer Science Department
IME USP

The Border Gateway Protocol - BGP

The internet consists of several interlinked Autonomous Systems (ASes)

- **Each AS is responsible for a set of IP prefixes**
- **BGP is used by ASes to exchange routing information**
  - AS routers exchange messages with their neighbors
- **The route used by an AS depends on its relationship with its neighbors**
  - The true relationship is usually confidential but may be inferred [Gao01]

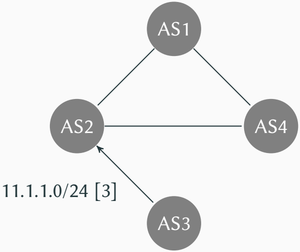

The Border Gateway Protocol - BGP

The internet consists of several interlinked Autonomous Systems (ASes)

- **Each AS is responsible for a set of IP prefixes**
- **BGP is used by ASes to exchange routing information**
  - AS routers exchange messages with their neighbors
- **The route used by an AS depends on its relationship with its neighbors**
  - The true relationship is usually confidential but may be inferred [Gao01]

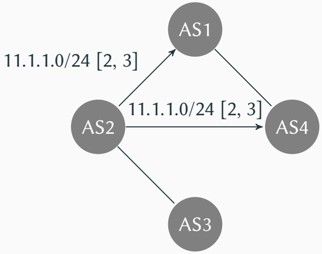

BGP Anomalies

BGP update messages can be drastically affected by anomalous events [AMBA16]

- **Direct anomalies (such as IP hijacking or typos in prefixes)**
- **Indirect anomalies (such as worms spreading)**
- **Link failure (caused by earthquakes or blackouts)**

After identifying anomalous behavior, it is important to classify it

- **Classification allows for BGP operators to respond accordingly**

Previous work

Until recently, work on BGP anomaly was concerned only with detection

Since 2020, LSTM models have been proposed for classification of anomalies [CLL+21, Fon20]

However, we identified the following limitations:

- **Classification is limited to events, not their type**
- **The models are tested against events seen during training**

Our work

We propose

- **A set of features based on the inferred AS relationship graph**
- **Together with an LSTM-based classifier**

The model is trained and tested with *different events*

It achieves reasonably good performance

Our code and data is publicly available

`https://github.com/thalespaiva/bgp-anomaly-classification/`

Anomalous Events

Events for training

AS9121 RTL	Direct
AWS Route Leak	Direct
Code Red	Indirect
Nimda	Indirect
Moscow Blackout	Link Failure

Events for testing

Malaysian Telecom	Direct
Slammer	Indirect
Japan Earthquake	Link Failure

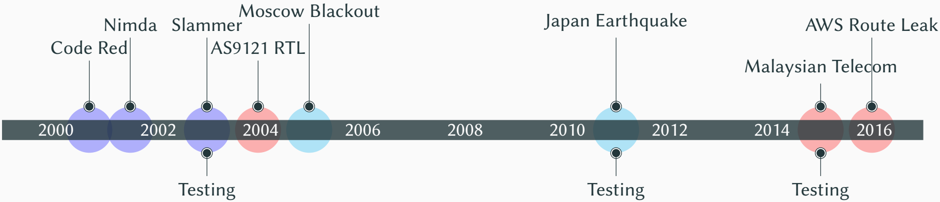

AS relationship graph

How to deal with the dynamic nature of the AS graph?

- **Extract paths seen in updates 2 days before each event**
- **Apply AS relationship inference [Gao01] over the collected paths**

AS paths 2 days before event

```
[3, 2, 4]
[4, 1]
[2, 1]
[3, 2, 1, 4]
[2, 3]
```

Inferred AS graph

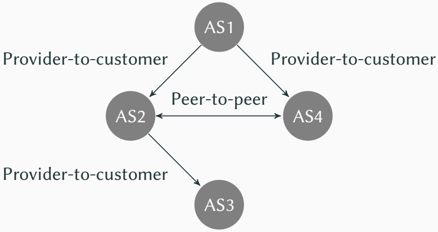

Feature Extraction

- **We downloaded sets of BGP updates exchanged during the events**
- **Data was downloaded from selected collectors from the RIS project**
- **Computed 17 features over sets of 1 minute duration of updates**

Examples of features:

- **Features commonly used for BGP anomaly features**
  - Number of announcements
  - Average length of AS paths
- **Features based on the AS graph**
  - Average degree of the ASes within AS paths
  - Number of edges of each type (e.g. provider-to-customer)

Our LSTM-based Model

Input for the network is a sequence of 10 sets of features (10 minutes)

Layer type	Output dimension
Convolutional 1D	(10, 32)
Max Pooling 1D	(5, 32)
LSTM	(100)
Dropout	(100)
Dense	(3)

- **Training events were split into sets of non-overlapping sequences**
  - 70% for training
  - 20% for validation
  - 10% for preliminary test
- **Model was trained for 10 epochs with batch size 1**
  - 100% accuracy for the preliminary test

Classification Results

Testing with events **not seen during training**

		Predicted Label		
		Direct	Indirect	Link Failure
	Direct	11	0	0
True Label	Indirect	0	88	0
	Link Failure	0	13	26

Discussion

The LSTM model together with these features appears to be promising

The main limitation of this work is the dataset

- **Events are not evenly distributed**
- **We need a larger number of events**

Therefore we encourage researchers to validate our approach using larger and different datasets

Conclusion and future work

We plan to post an extended version of this paper:

- **An analysis of the robustness of the model**

Future work:

- **Use better relationship inference algorithms [JSD+19]**
- **Consider larger sets of BGP anomalies**

Contact:

- **Email: tpaiva@ime.usp.br**
- **Repository:** `https://github.com/thalespaiva/bgp-anomaly-classification`

References i

- Bahaa Al-Musawi, Philip Branch, and Grenville Armitage, *BGP Anomaly Detection Techniques: A Survey*, IEEE Communications Surveys & Tutorials **19** (2016), no. 1, 377–396.
- Min Cheng, Qing Li, Jianming Lv, Wenyin Liu, and Jianping Wang, *Multi-Scale LSTM Model for BGP Anomaly Classification*, IEEE Transactions on Services Computing **14** (2021), no. 3, 765–778.
- Paulo César da Rocha Fonseca, *A Deep Learning Framework for BGP Anomaly Detection and Classification*, Ph.D. thesis, Universidade Federal do Amazonas, 2020.

- Lixin Gao, *On Inferring Autonomous System Relationships in the Internet*, IEEE/ACM Transactions on Networking **9** (2001), no. 6, 733–745.
- Yuchen Jin, Colin Scott, Amogh Dhamdhere, Vasileios Giotsas, Arvind Krishnamurthy, and Scott Shenker, *Stable and Practical AS Relationship Inference with ProbLink*, Proceedings of the 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19), 2019, pp. 581–598.