Worldwide Internet Outage Detection

4th CAIDA Internet Measurement And Political Science (IMAPS) Workshop 2018

Andréas Guillot

Supervisors:

Romain Fontugne Pascal Mérindol Cristel Pelsser

September 07, 2018

University of Strasbourg

Outage Detection

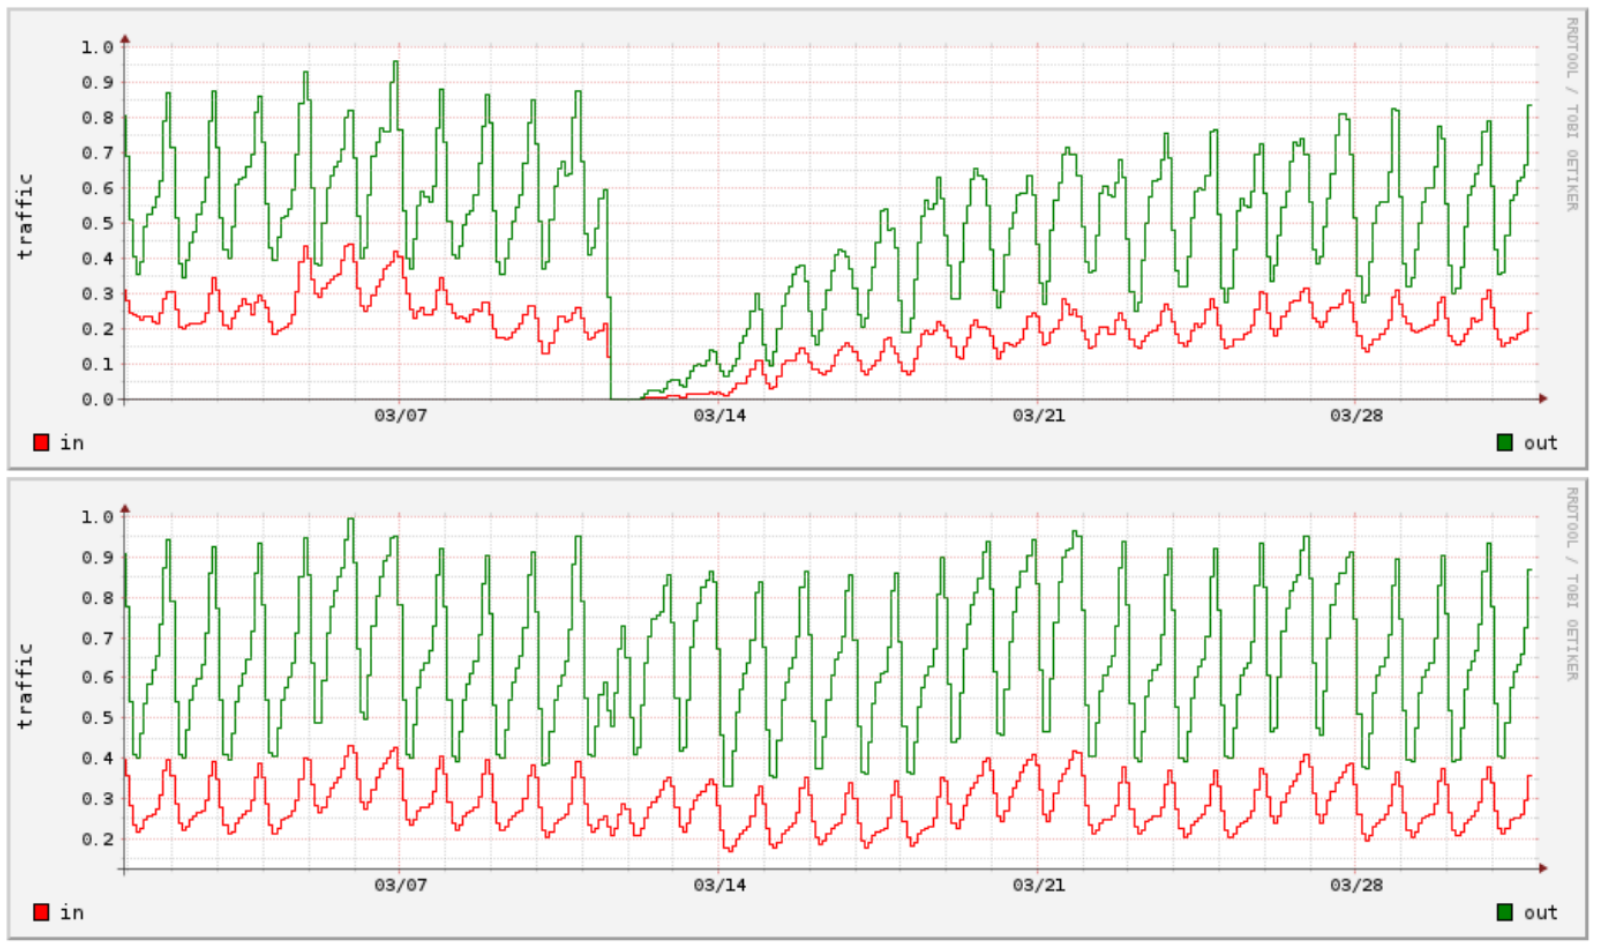

Figure 1: Japanese traffic for the March 2011 earthquake, Miyagi prefecture (top) and nationwide (bottom)[1]

[1] Kenjiro Cho et al. "The Japan earthquake: the impact on traffic and routing observed by a local ISP". . In: *Proceedings of the Special Workshop on Internet and Disasters*. ACM. 2011, p. 2.

Related works

Active monitoring

Technique

- Sending *pings*[2]

Advantage

- Worldwide coverage

Drawback

- Injects traffic

Passive monitoring

Technique

- Existing active data[3]

Advantage

- Generates no traffic

Drawback

- Partial coverage

[2] Lin Quan, John Heidemann, and Yuri Pradkin. "Trinocular: Understanding internet reliability through adaptive probing". In: *ACM SIGCOMM Computer Communication Review*. Vol. 43. 4. ACM. 2013, pp. 255–266

[3] Anant Shah et al. "Disco: Fast, good, and cheap outage detection". In: *Network Traffic Measurement and Analysis Conference (TMA), 2017*. IEEE. 2017, pp. 1–9

Related works

Active monitoring

Technique

- Sending *pings*[2]

Advantage

- Worldwide coverage

Drawback

- ~~**Injects traffic**~~

Passive monitoring

Technique

- Existing active data[3]

Advantage

- Generates no traffic

Drawback

- Partial coverage

[2] Lin Quan, John Heidemann, and Yuri Pradkin. "Trinocular: Understanding internet reliability through adaptive probing". In: *ACM SIGCOMM Computer Communication Review*. Vol. 43. 4. ACM. 2013, pp. 255–266

[3] Anant Shah et al. "Disco: Fast, good, and cheap outage detection". In: *Network Traffic Measurement and Analysis Conference (TMA), 2017*. IEEE. 2017, pp. 1–9

Related works

Active monitoring

Technique

- Sending *pings*[2]

Advantage

- Worldwide coverage

Drawback

- ~~**Injects traffic**~~

Passive monitoring

Technique

- Existing active data[3]

Advantage

- Generates no traffic

Drawback

- **(Partial coverage)**

[2] Lin Quan, John Heidemann, and Yuri Pradkin. "Trinocular: Understanding internet reliability through adaptive probing". In: *ACM SIGCOMM Computer Communication Review*. Vol. 43. 4. ACM. 2013, pp. 255–266

[3] Anant Shah et al. "Disco: Fast, good, and cheap outage detection". In: *Network Traffic Measurement and Analysis Conference (TMA), 2017*. IEEE. 2017, pp. 1–9

Outline

Dataset

Methodology

Validation

Dataset

Internet Background Radiation

CAIDA Network telescope

- *Allocated*, *routeable*, but *unused* IP addresses

Traffic composition

- Misconfigurations
- Spoofed traffic

Use cases

- Attacks [4]
- Censorship [3]
- Local outages [1]

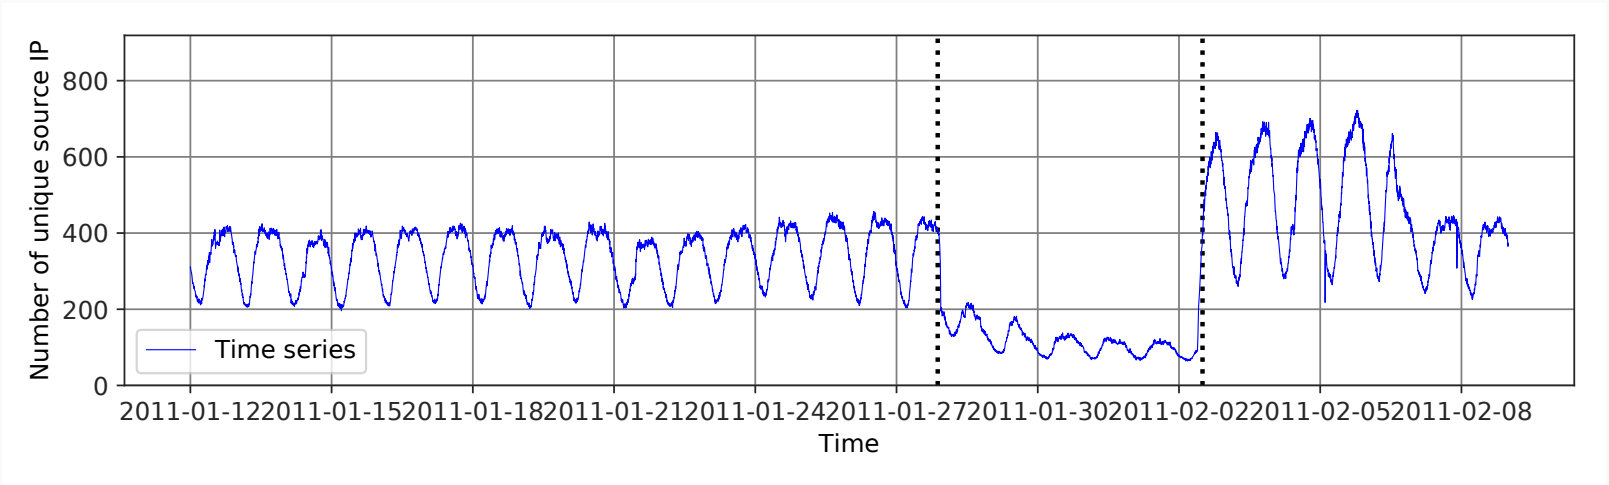

Figure 2: Number of unique source IP addresses from *Egypt* during the Egyptian revolution

Internet Background Radiation

CAIDA Network telescope

- *Allocated*, *routeable*, but *unused* IP addresses

Traffic composition

- Misconfigurations
- Spoofed traffic

Use cases

- Attacks [4]
- Censorship [3]
- Local outages [1]

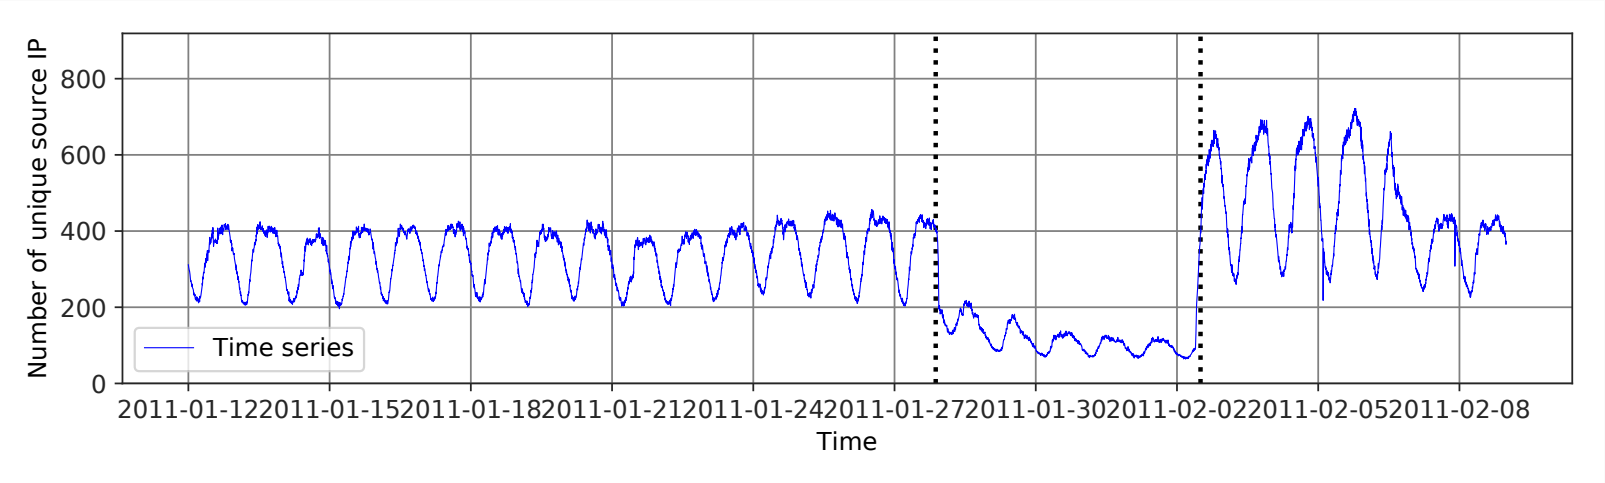

Figure 2: Number of unique source IP addresses from *Egypt* during the Egyptian revolution

→ More than 60 000 time series

Methodology

Current methodology

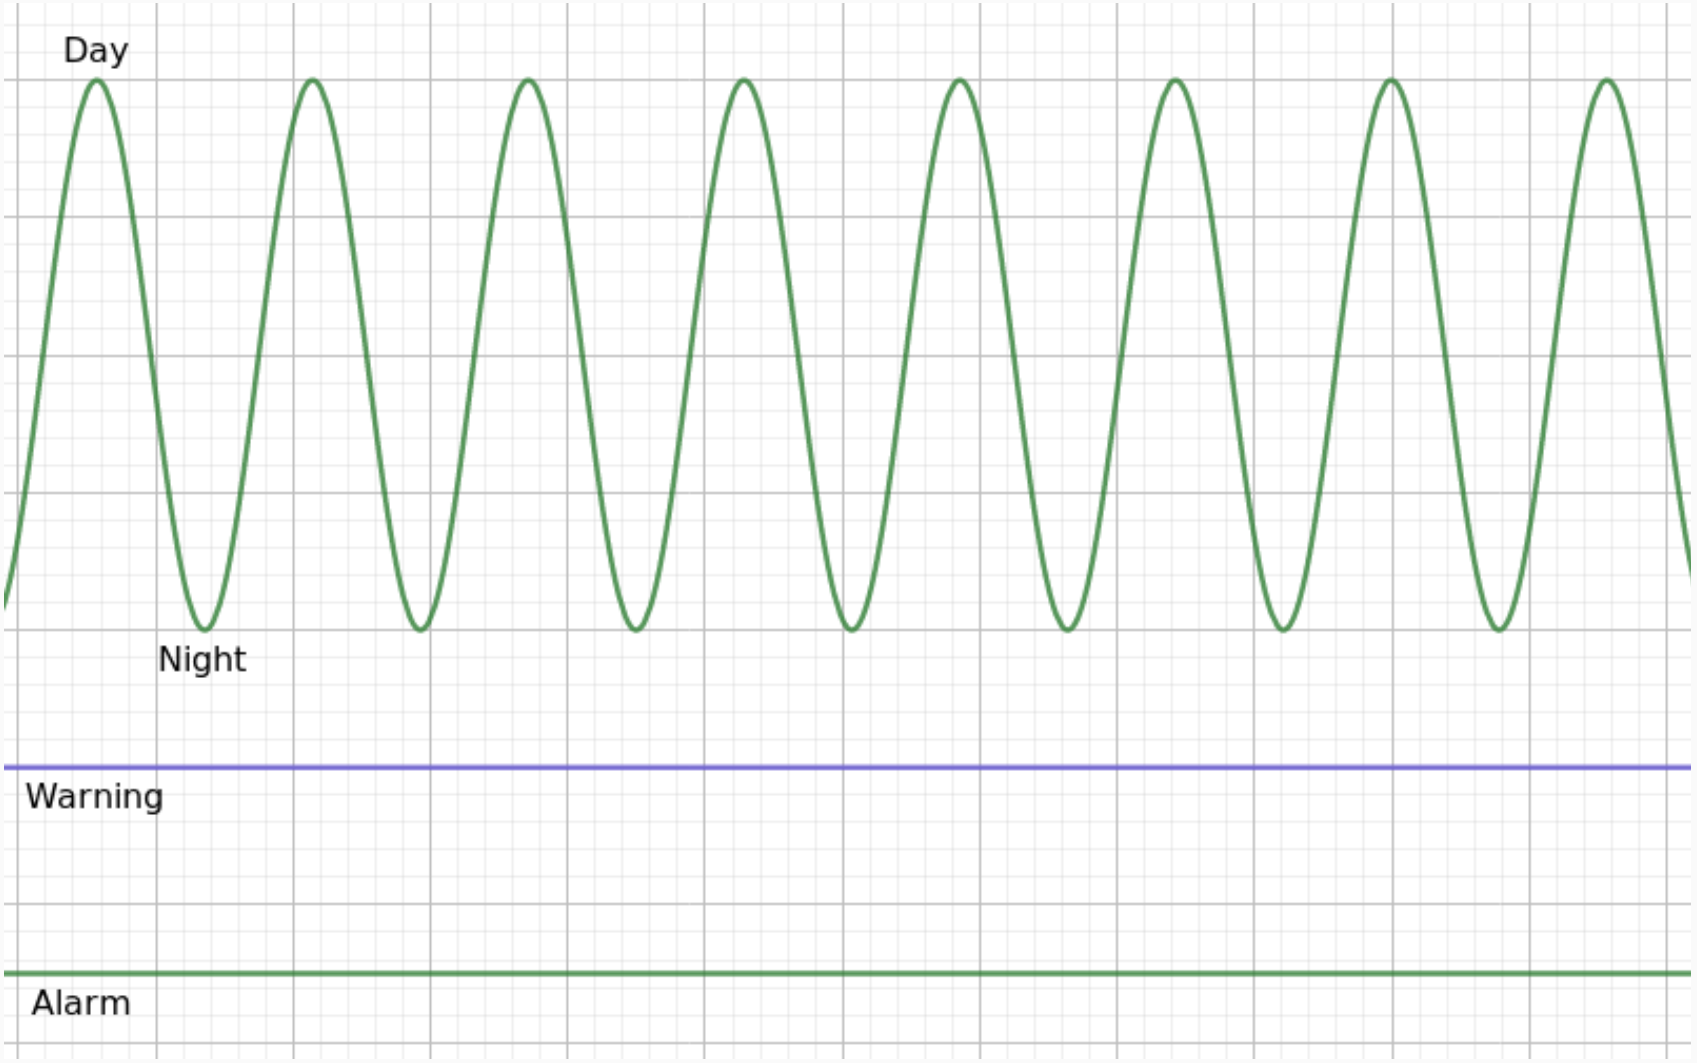

Figure 3: Detecting outages using fixed thresholds

Current methodology

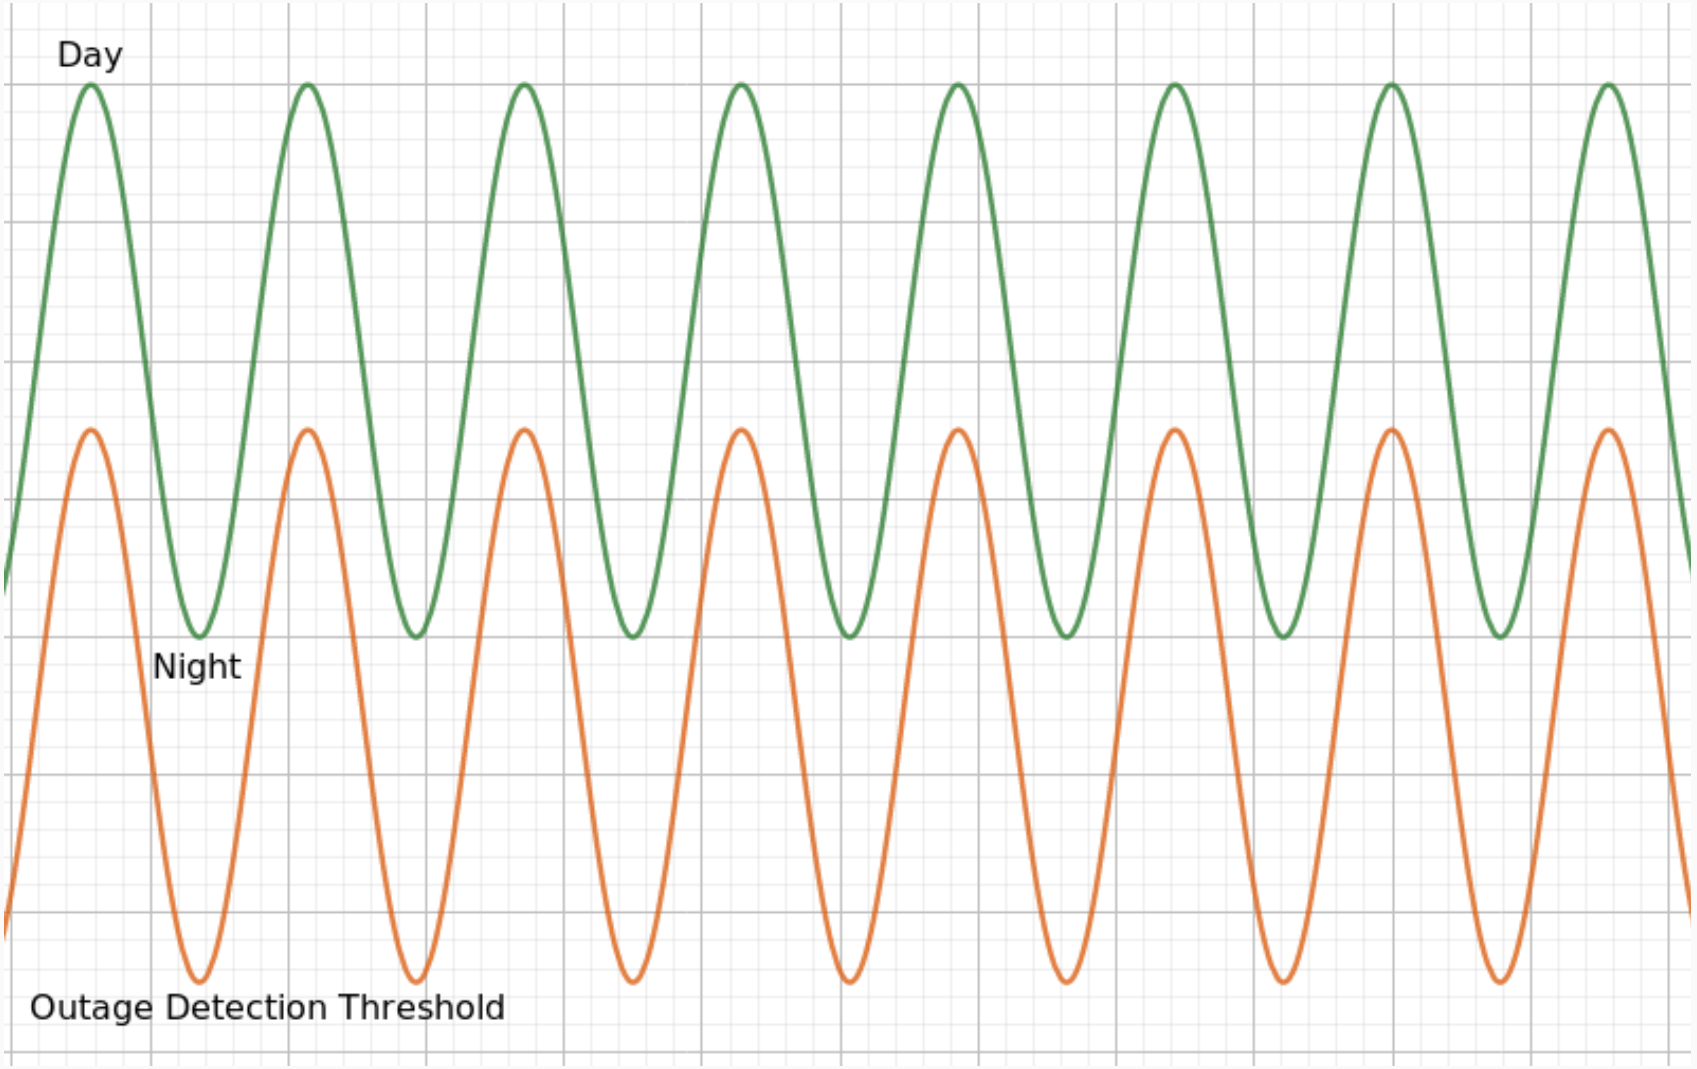

Figure 3: Detecting outages using dynamic thresholds

Current methodology

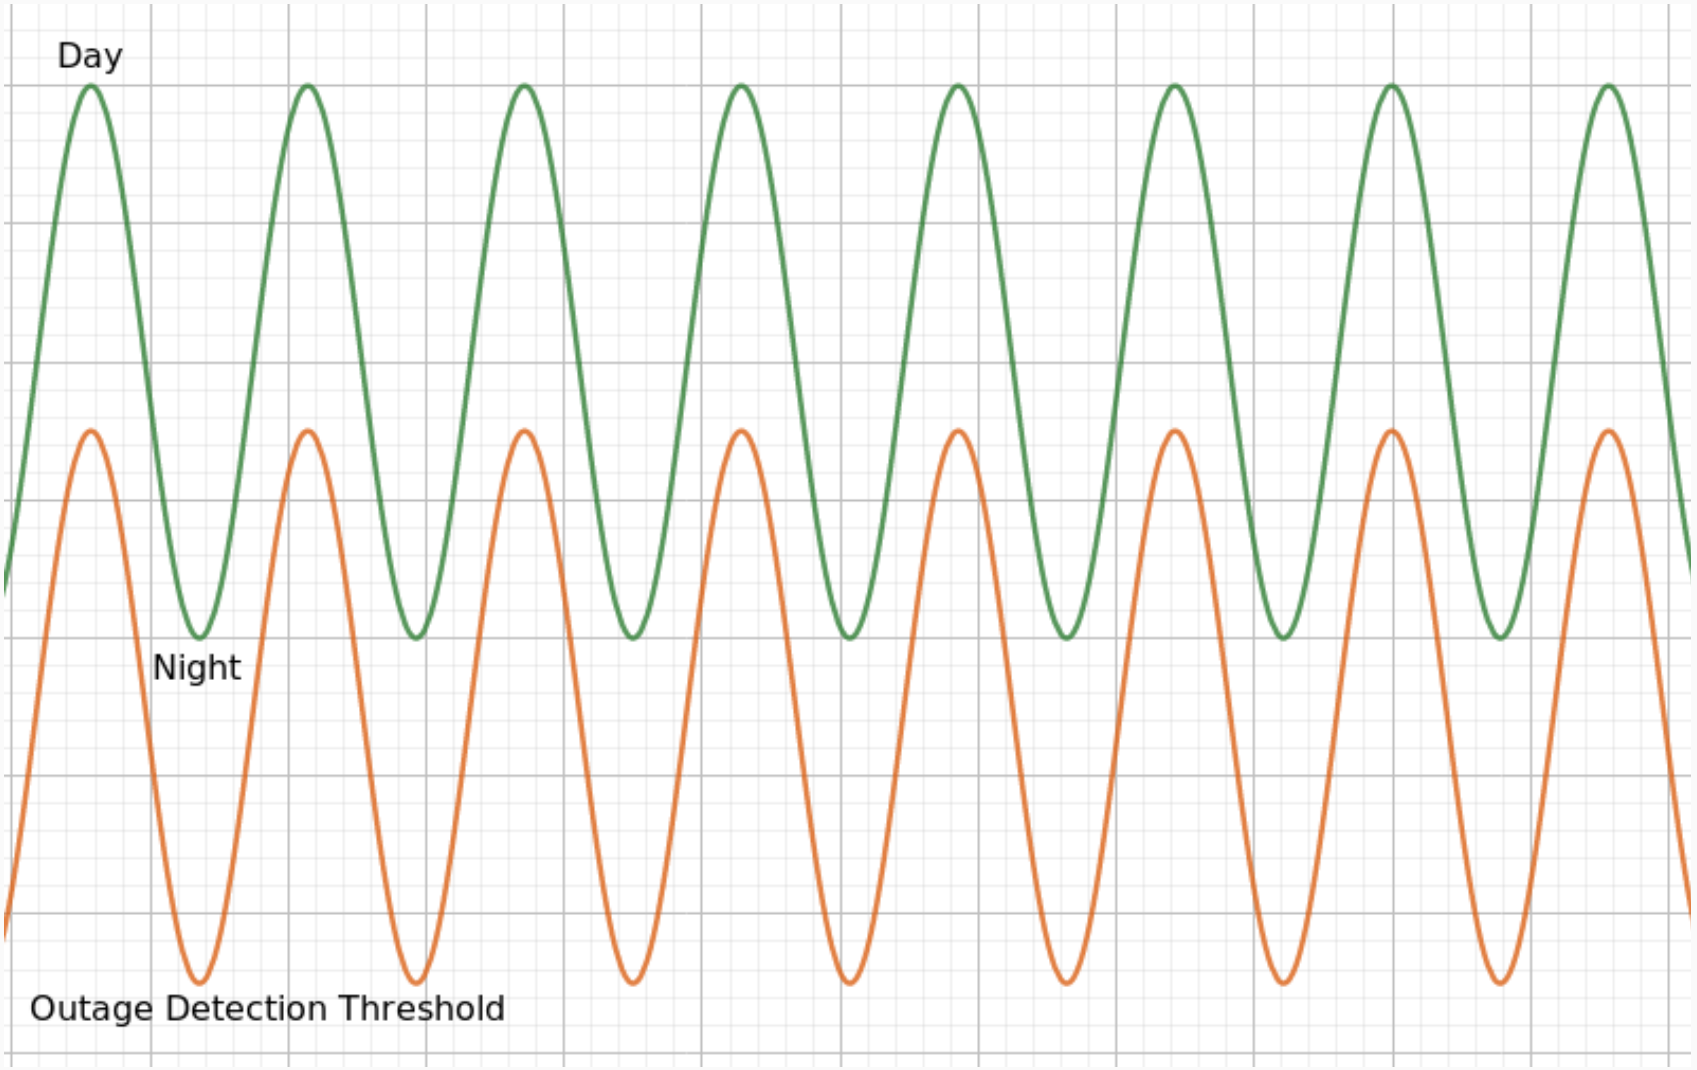

Figure 3: Detecting outages using dynamic thresholds

$\Rightarrow$ This is our goal

Outage detection process

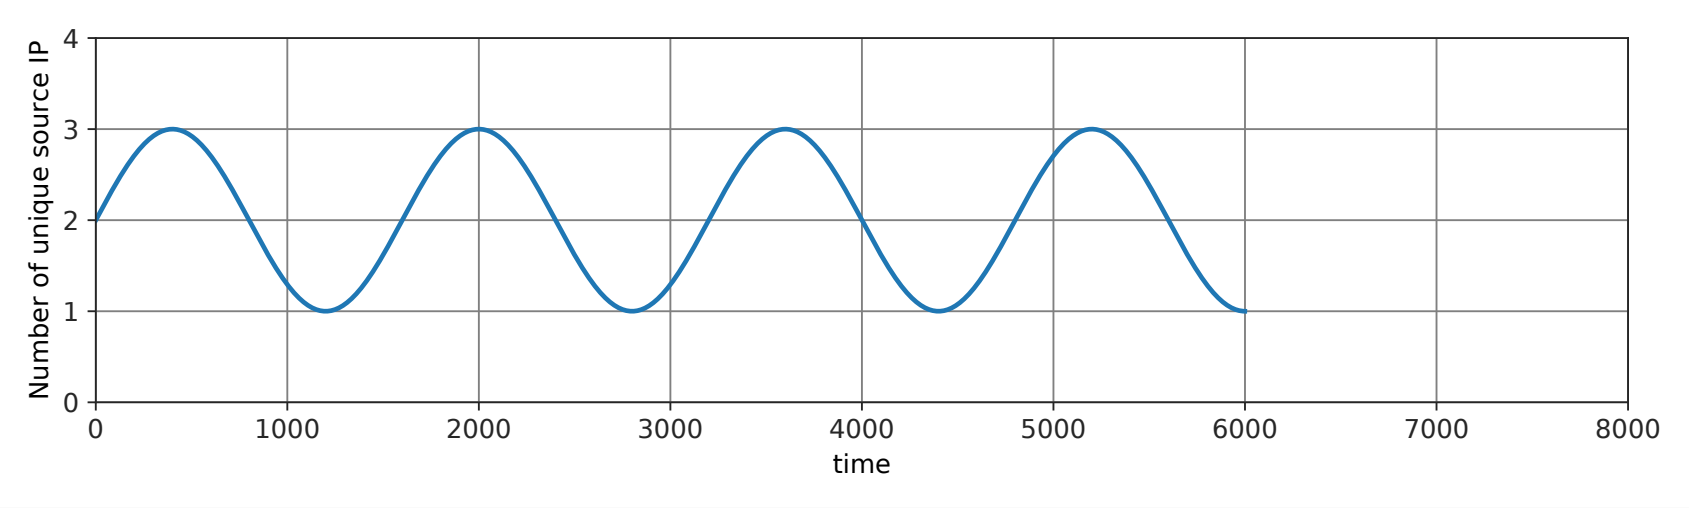

Figure 4: Outage detection process

Outage detection process

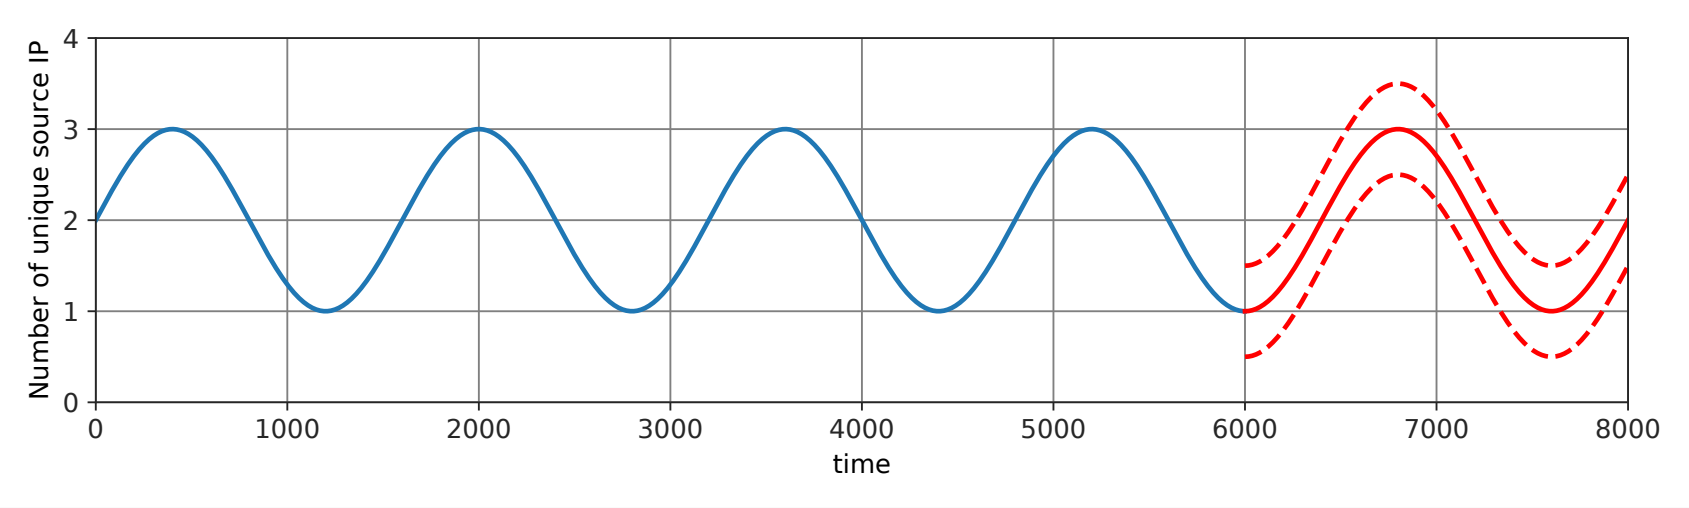

Figure 4: Outage detection process

Outage detection process

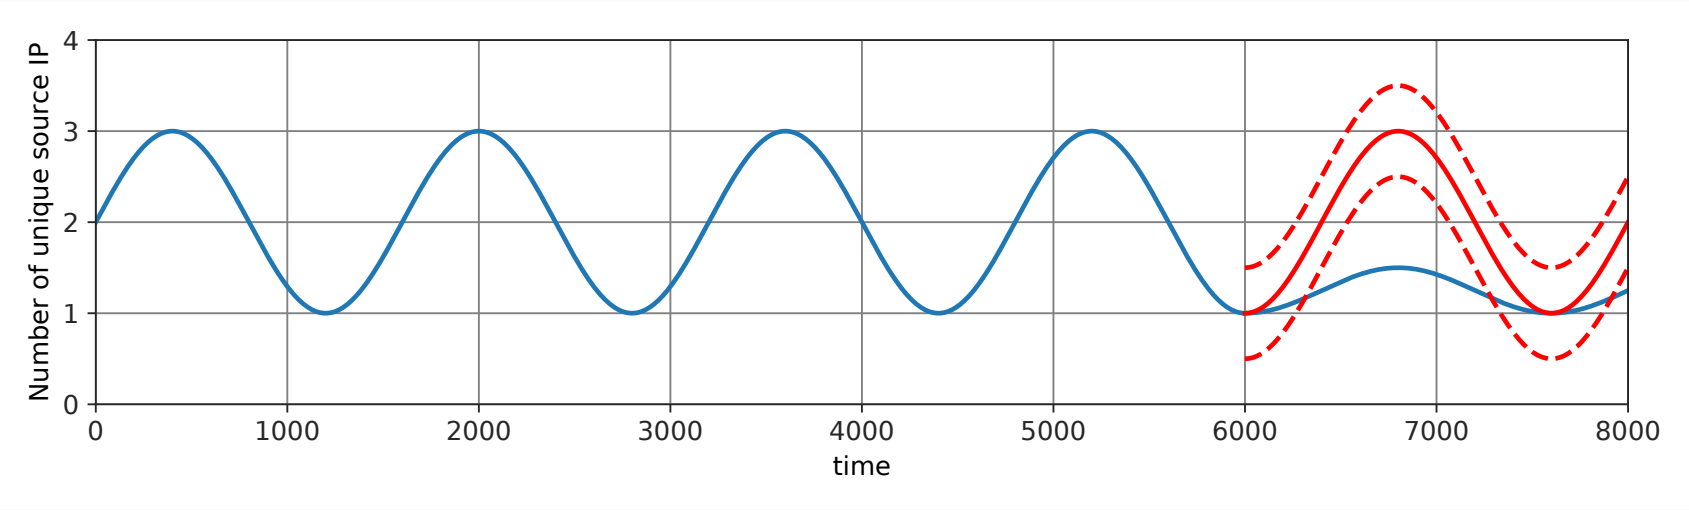

Figure 4: Outage detection process

The SARIMA model

S : Seasonal

AR : AutoRegressive

I : Integrated

MA : Moving Average

The SARIMA model

S : Seasonal $\rightarrow$ Remove *trends*

AR : AutoRegressive

I : Integrated $\rightarrow$ Normalize *mean and variance*

MA : Moving Average

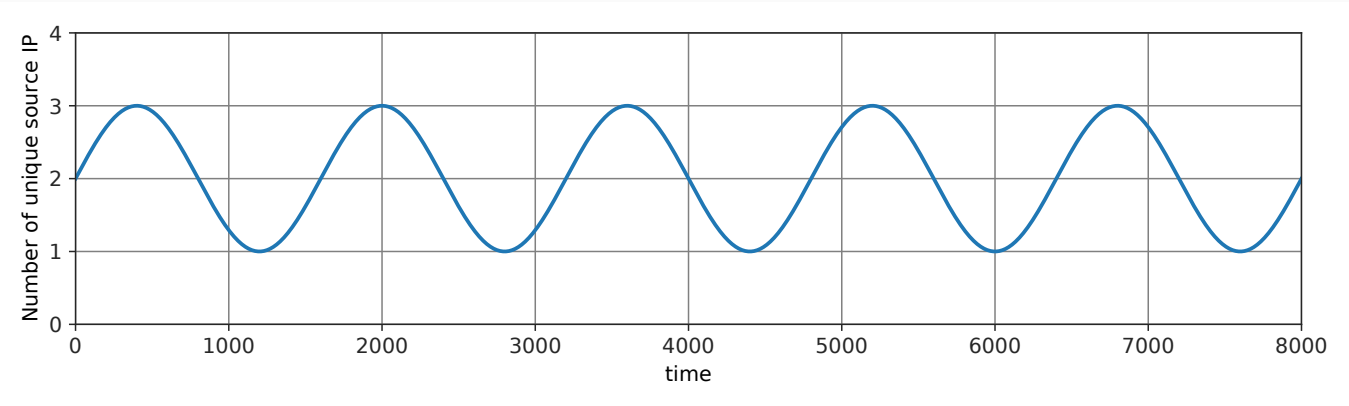

Figure 5: Original time series

The SARIMA model

S : Seasonal $\rightarrow$ Remove *trends*

AR : AutoRegressive

I : Integrated $\rightarrow$ Normalize *mean and variance*

MA : Moving Average

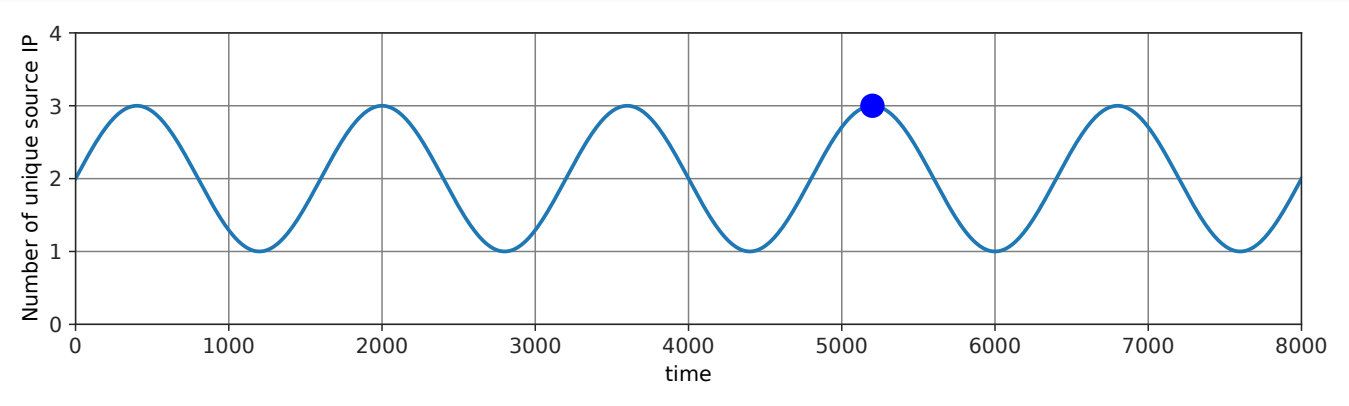

Figure 5: Original time series

The SARIMA model

S : Seasonal $\rightarrow$ Remove *trends*

AR : AutoRegressive

I : Integrated $\rightarrow$ Normalize *mean and variance*

MA : Moving Average

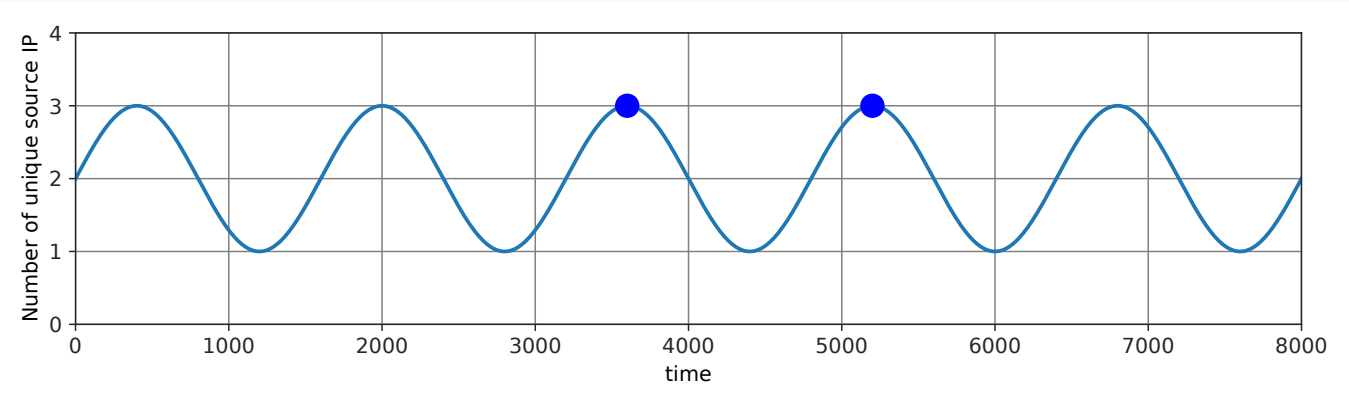

Figure 5: Original time series

The SARIMA model

S : Seasonal $\rightarrow$ Remove *trends*

AR : AutoRegressive

I : Integrated $\rightarrow$ Normalize *mean and variance*

MA : Moving Average

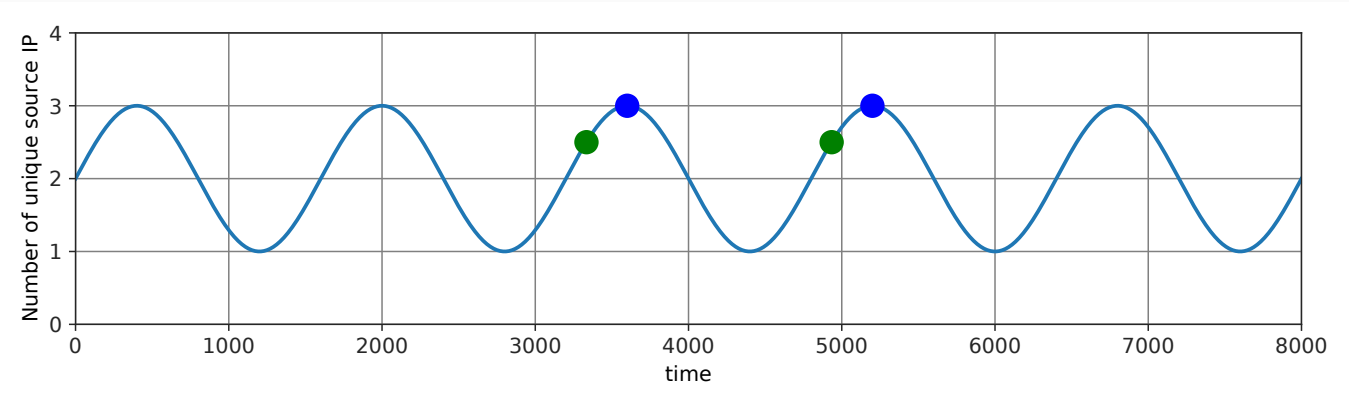

Figure 5: Original time series

The SARIMA model

S : Seasonal $\rightarrow$ Remove *trends*

AR : AutoRegressive

I : Integrated $\rightarrow$ Normalize *mean and variance*

MA : Moving Average

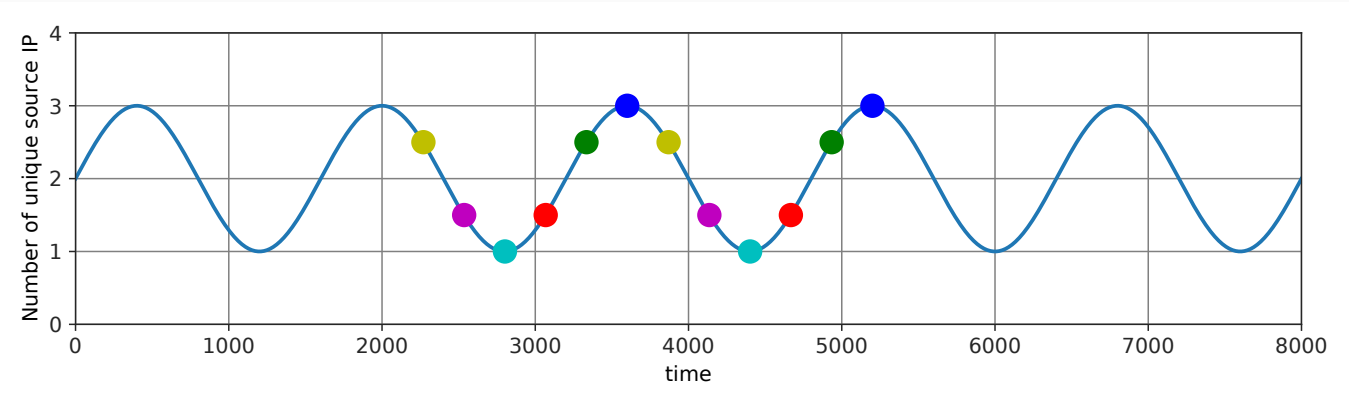

Figure 5: Original time series

The SARIMA model

S : Seasonal $\rightarrow$ Remove *trends*

AR : AutoRegressive

I : Integrated $\rightarrow$ Normalize *mean and variance*

MA : Moving Average

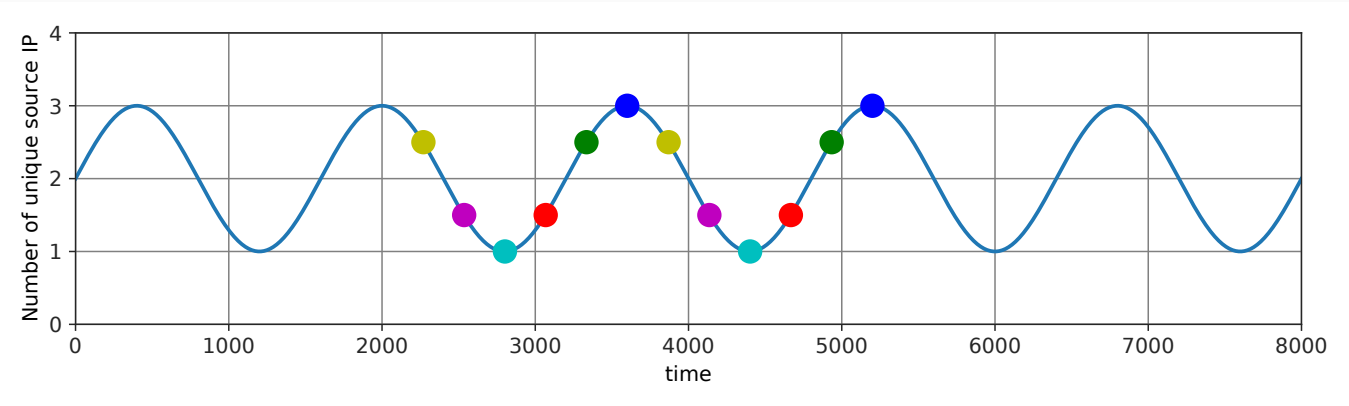

Figure 5: Original time series

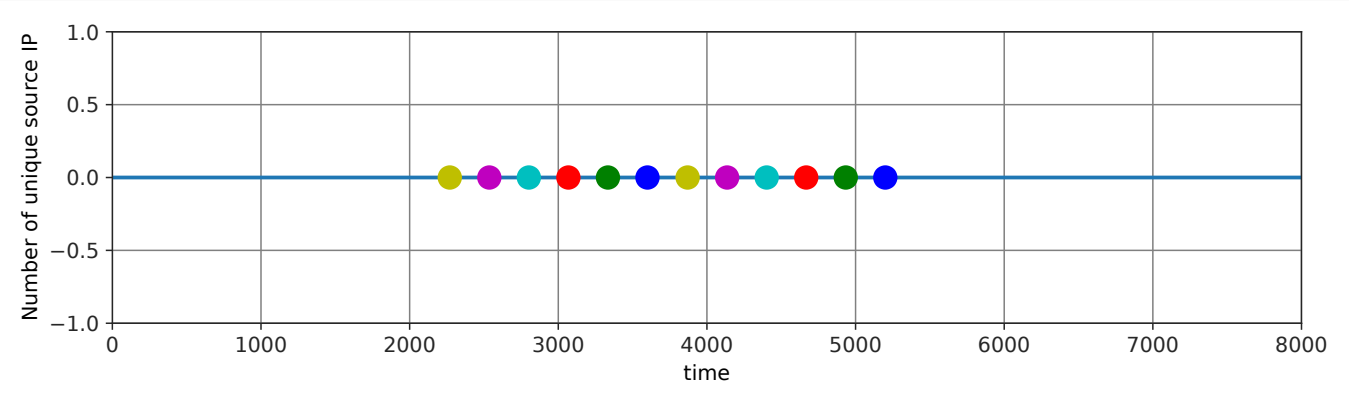

Figure 6: Differenced time series

The SARIMA model

S : Seasonal

AR : AutoRegressive $\rightarrow$ Predict based on *past values*

I : Integrated

MA : Moving Average $\rightarrow$ Predict based on *past errors*

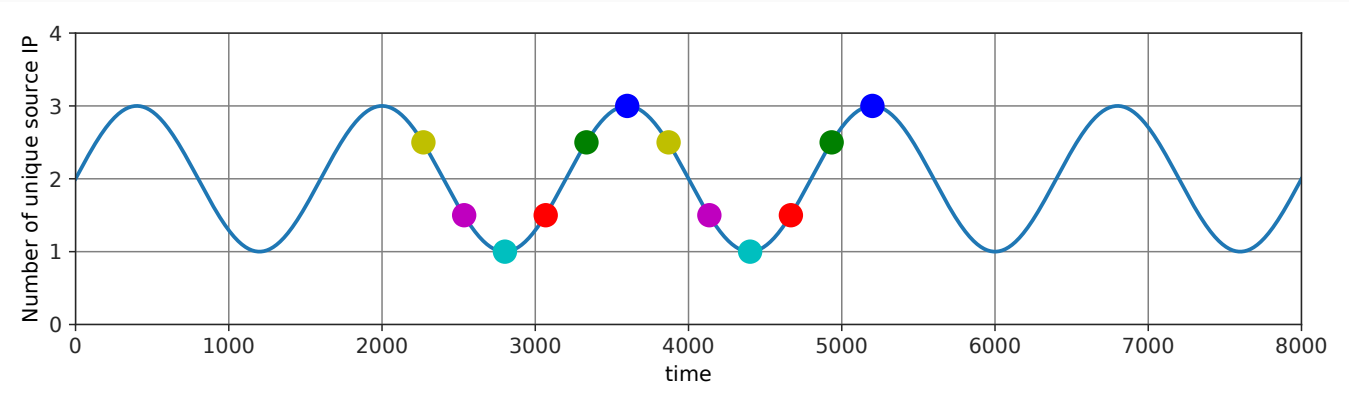

Figure 5: Original time series

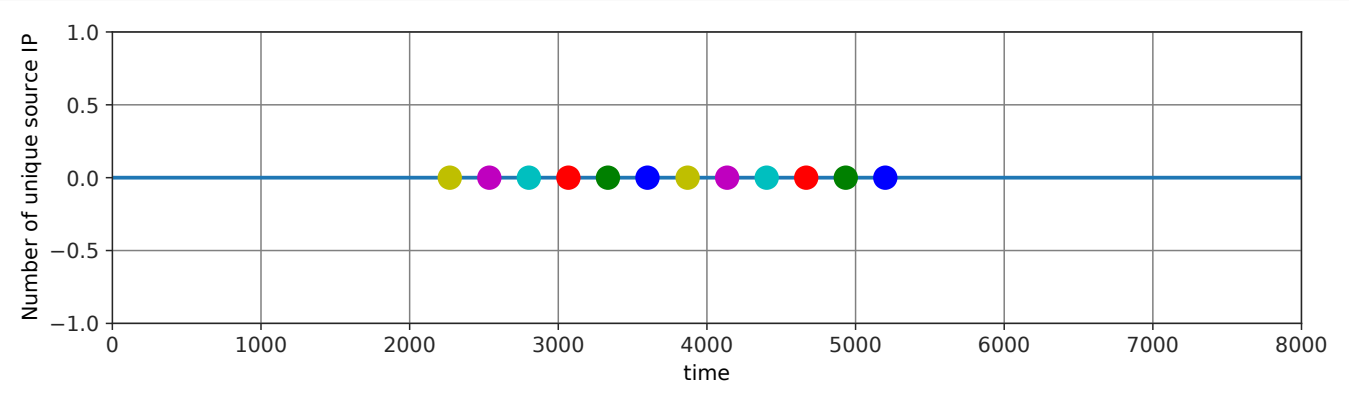

Figure 6: Differenced time series

Signal modelling

1. **Splitting the data set**
   - Training
   - Validation
   - Test
2. Finding the best parameters
3. Predicting the test set

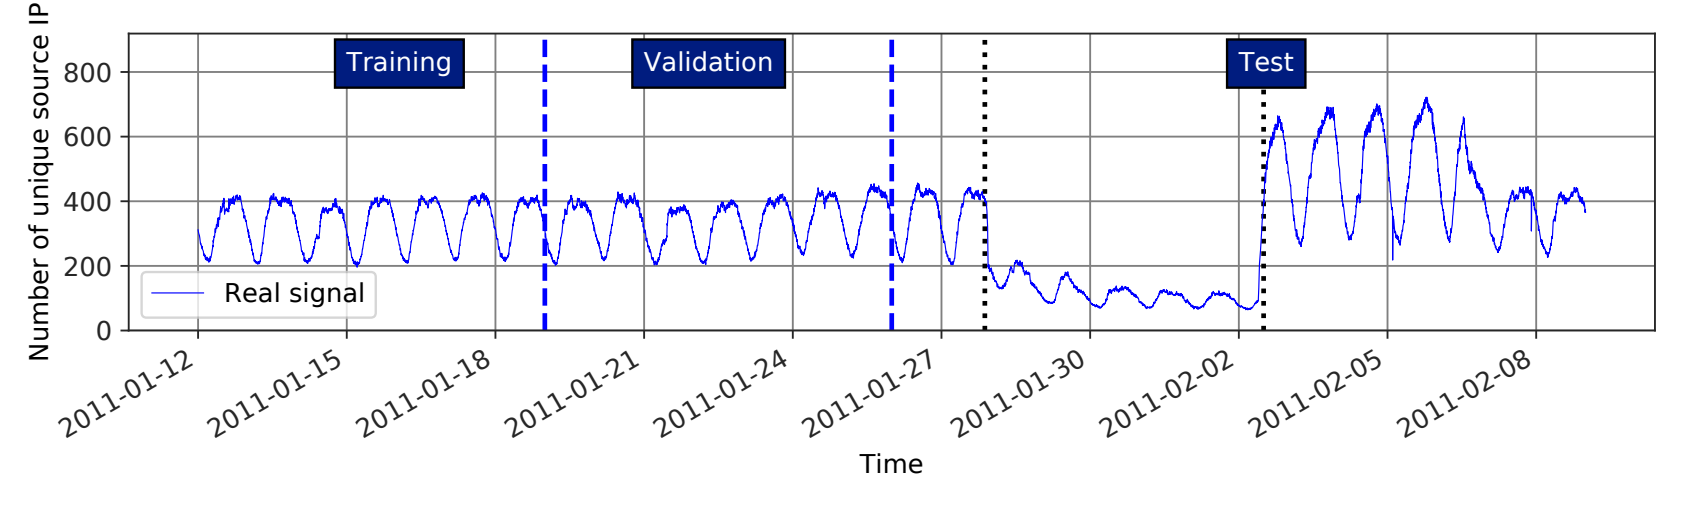

Figure 7: Different data sets

Signal modelling

1. Splitting the data set
2. **Finding the best parameters**
   - Which has the best *regression error*
3. Predicting the test set

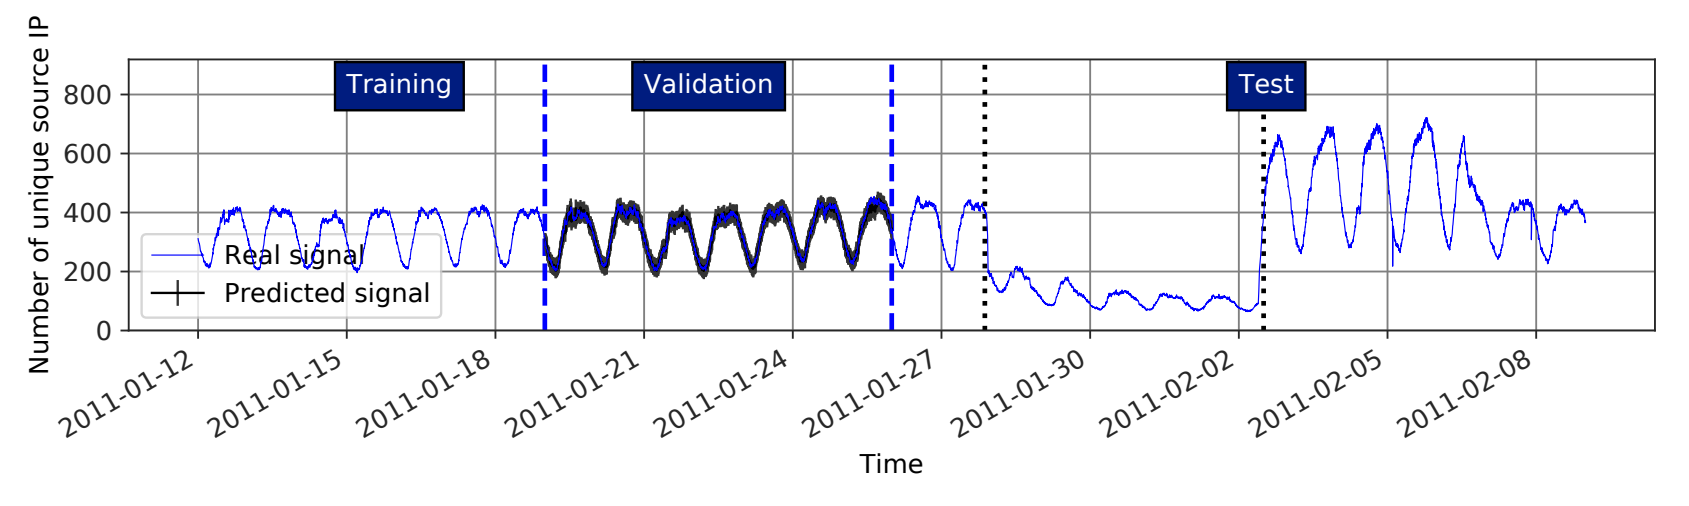

Figure 7: Making predictions on the validation set ($AR = 4, MA = 1$)

$\Rightarrow$ Gives us the $AR(p)$ and $MA(q)$ parameters

Signal modelling

1. Splitting the data set
2. Finding the best parameters
3. **Predicting the test set**
   - Integrating predictions to correct outages

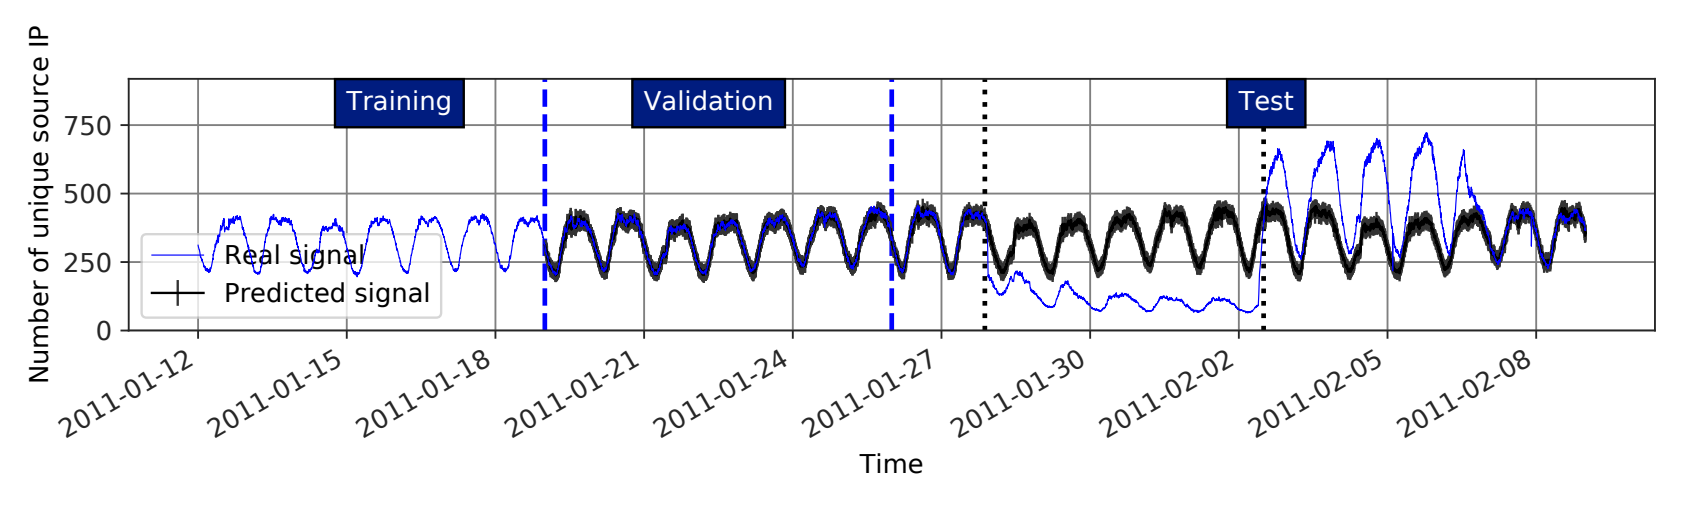

Figure 7: Predicting and inpainting the test set to preserve the integrity of the model

Outage detection

Definition of an outage

1. Outside of confidence interval
2. *actual* $<$ *predicted*

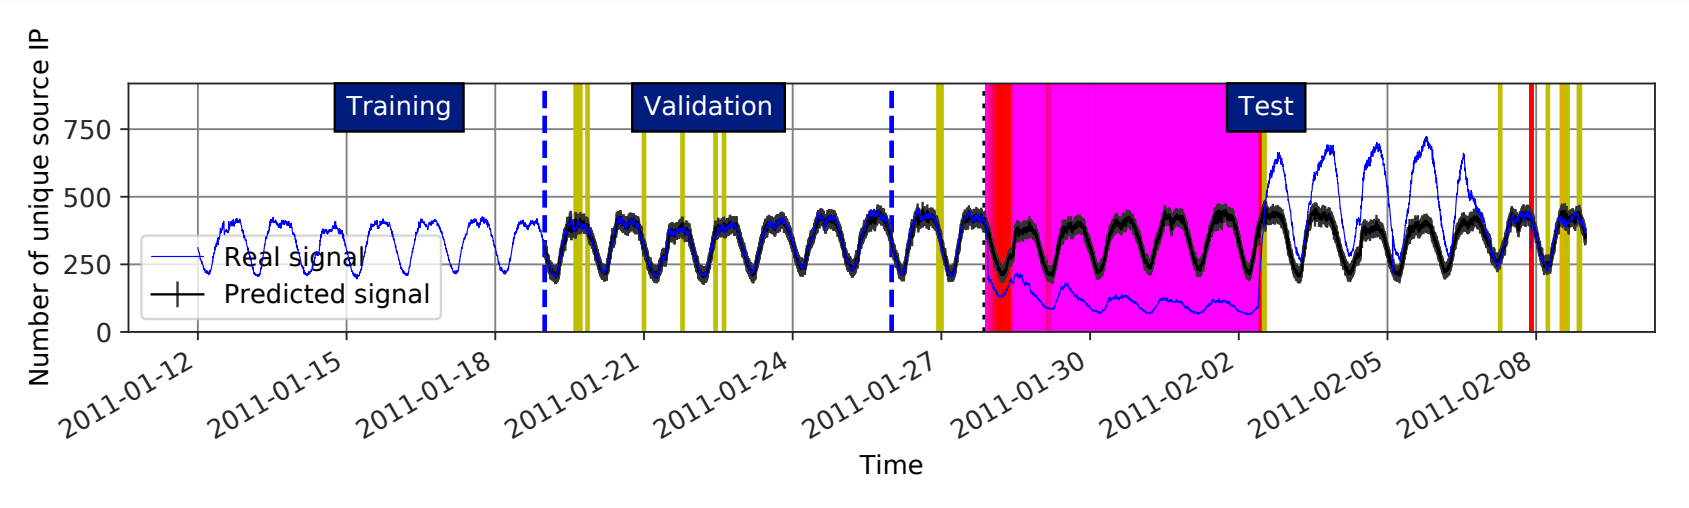

Figure 8: Analyzing the test set with the best model ($AR = 4, MA = 1$)

The 2017 Syrian exams

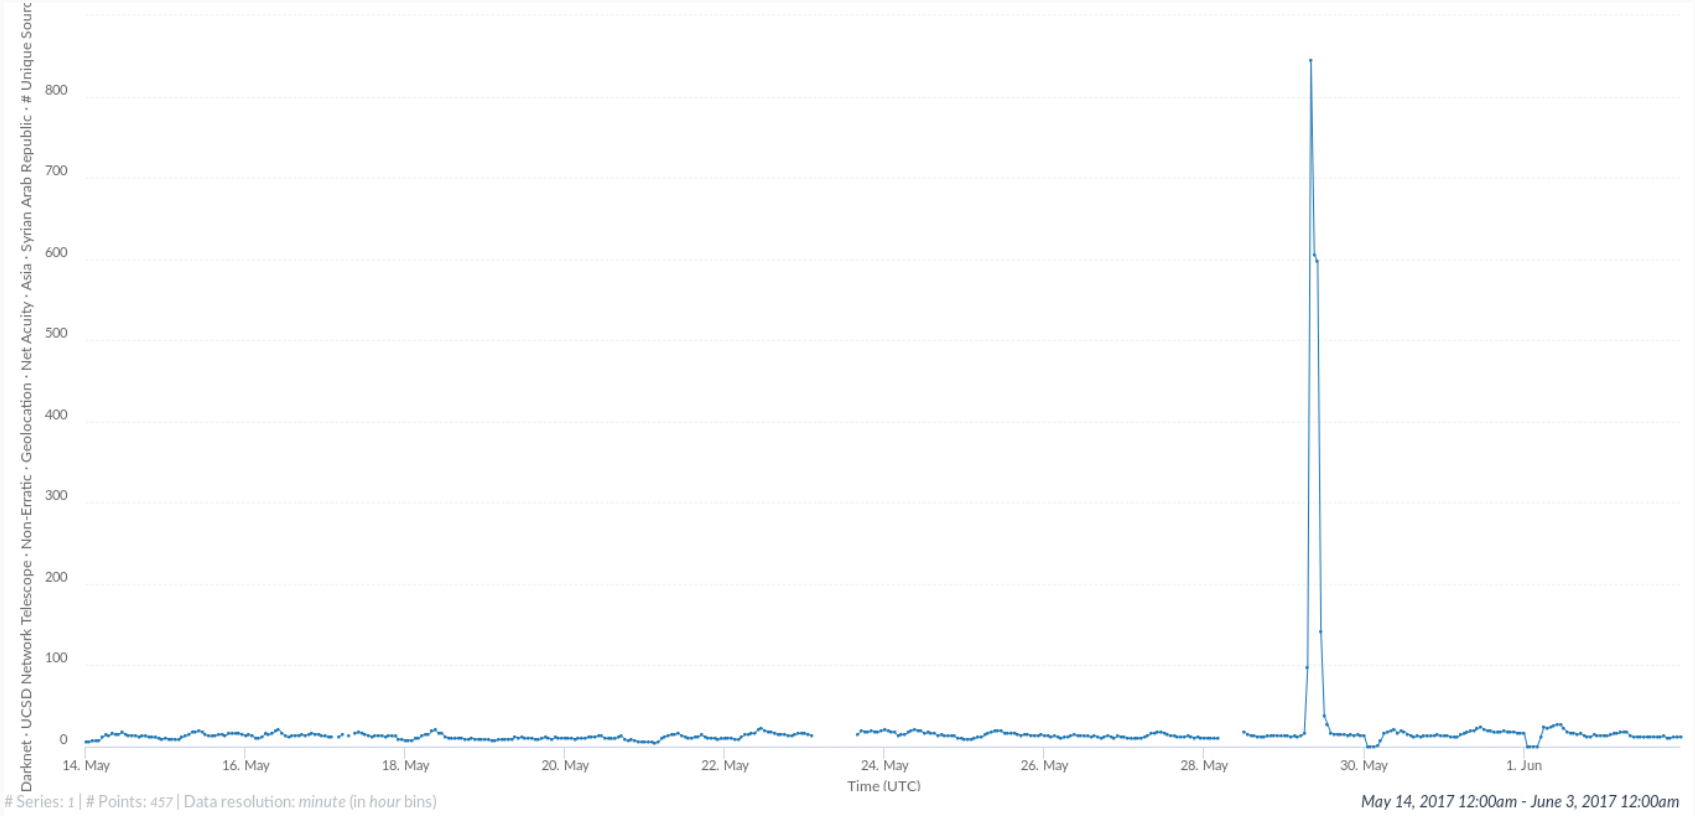

Figure 9: Syrian data — Missing data, extreme value, and outages (source: `ioda.caida.org`)

The 2017 Syrian exams

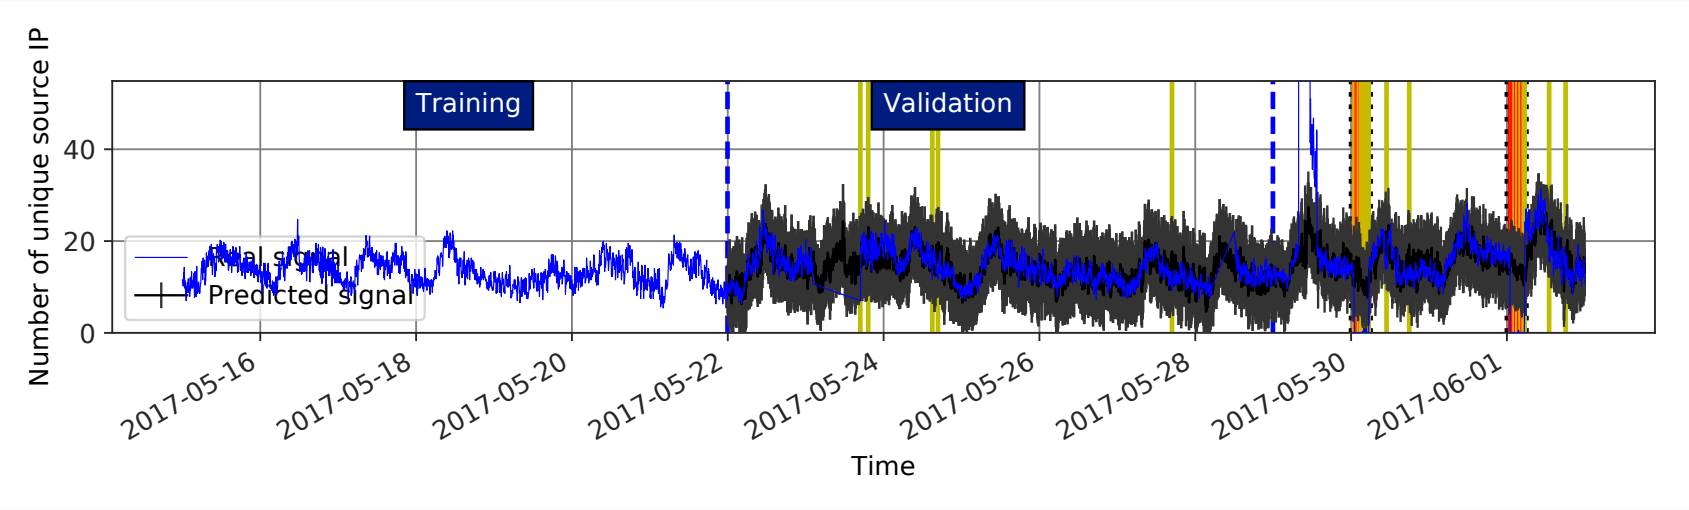

Figure 9: Syrian data — Missing data, extreme value, and outages

The Brazilian power cut

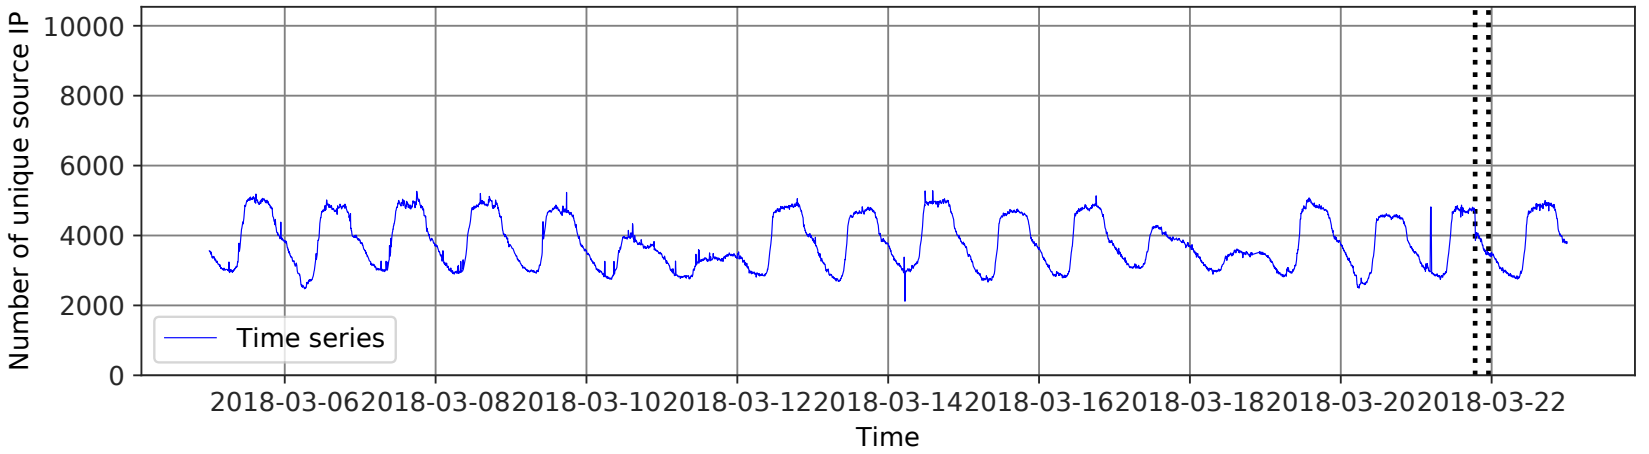

Figure 10: Brazilian data — Entire country

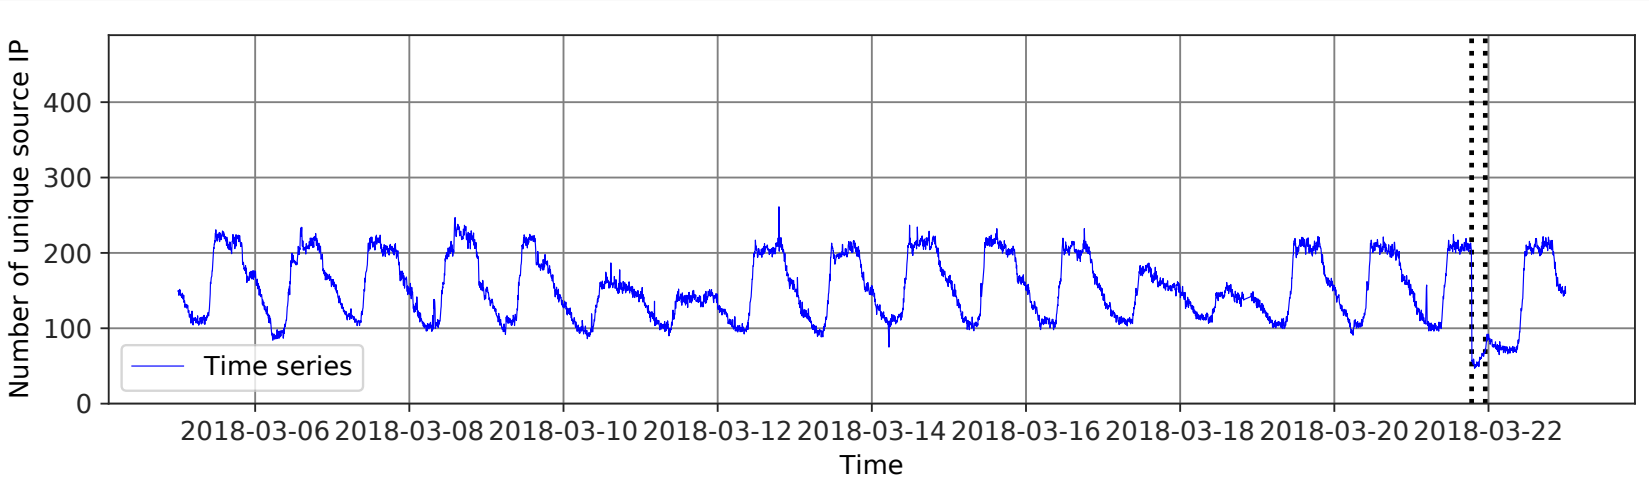

Figure 11: Brazilian data — Northern part of the country

The Brazilian power cut

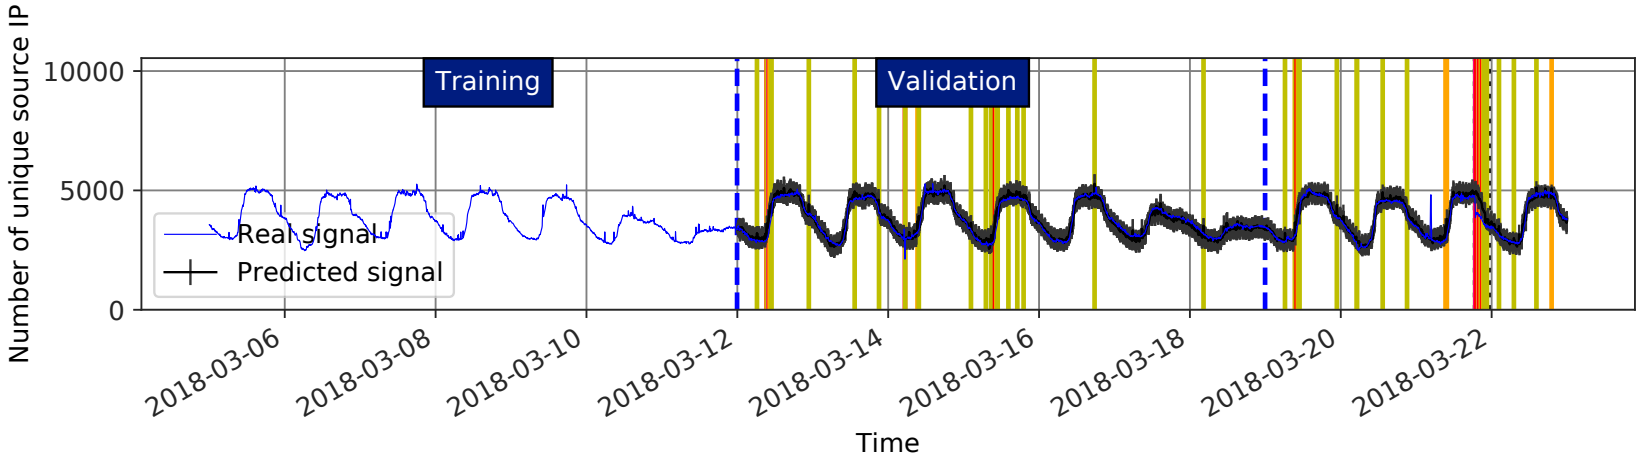

Figure 10: Brazilian data — Entire country

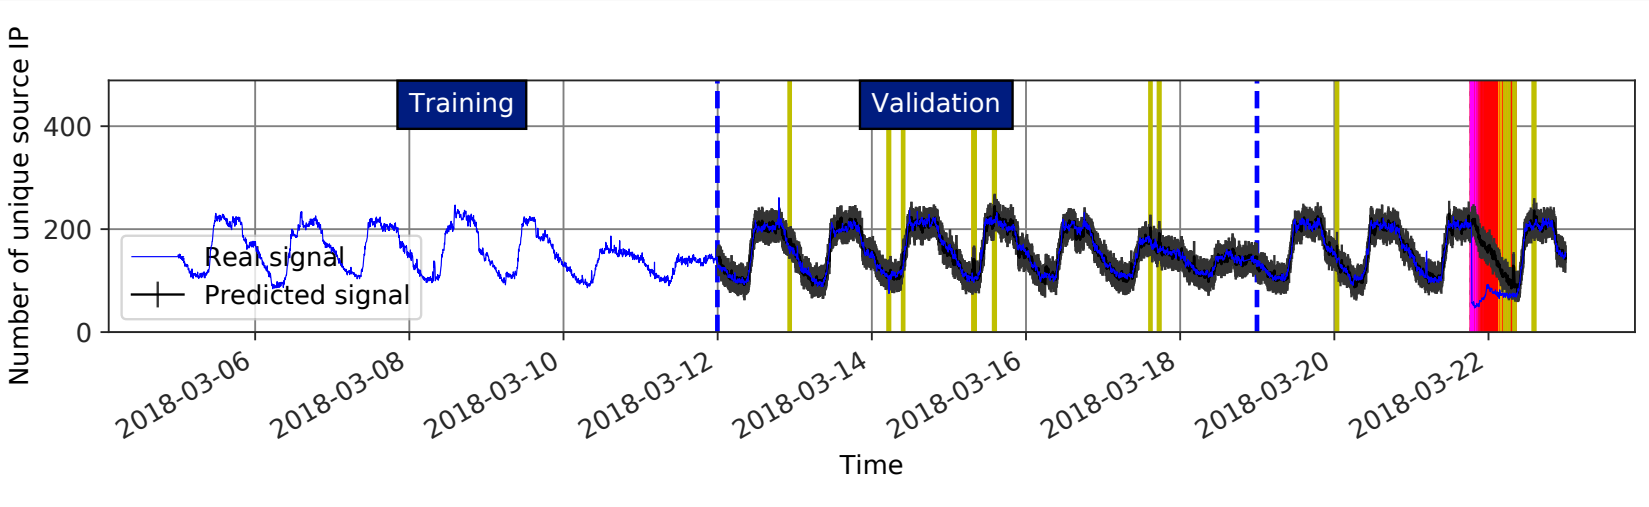

Figure 11: Brazilian data — Northern part of the country

Validation

Validation: ground truth

Objective

- *Evaluate* our solution using *known cases*

Characteristics

- 20 case studies
- Multiple spatial scales
  - Countries
  - Regions
  - Autonomous Systems

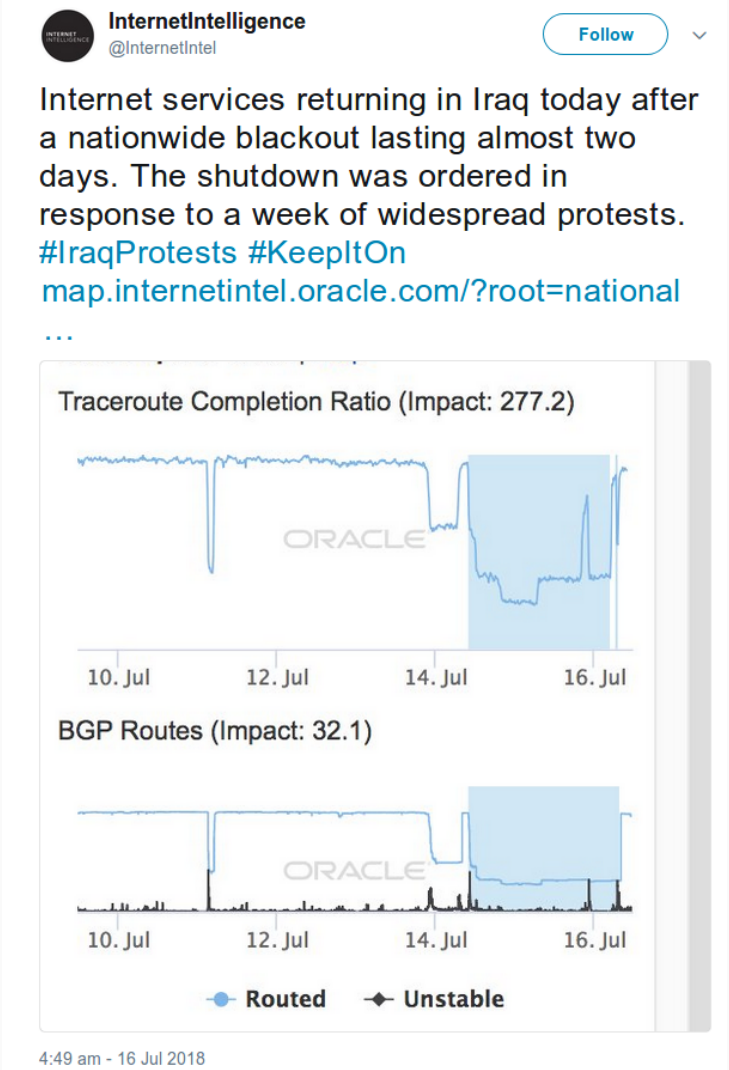

Figure 12: Example of a tweet announcing an outage (source)

Evaluating a time series

	True Positive	False Positive	True Negative	False Negative
Inside an outage	✓	✗	✗	✓
Raises an alarm	✓	✓	✗	✗

Figure 13: Classification of the different events

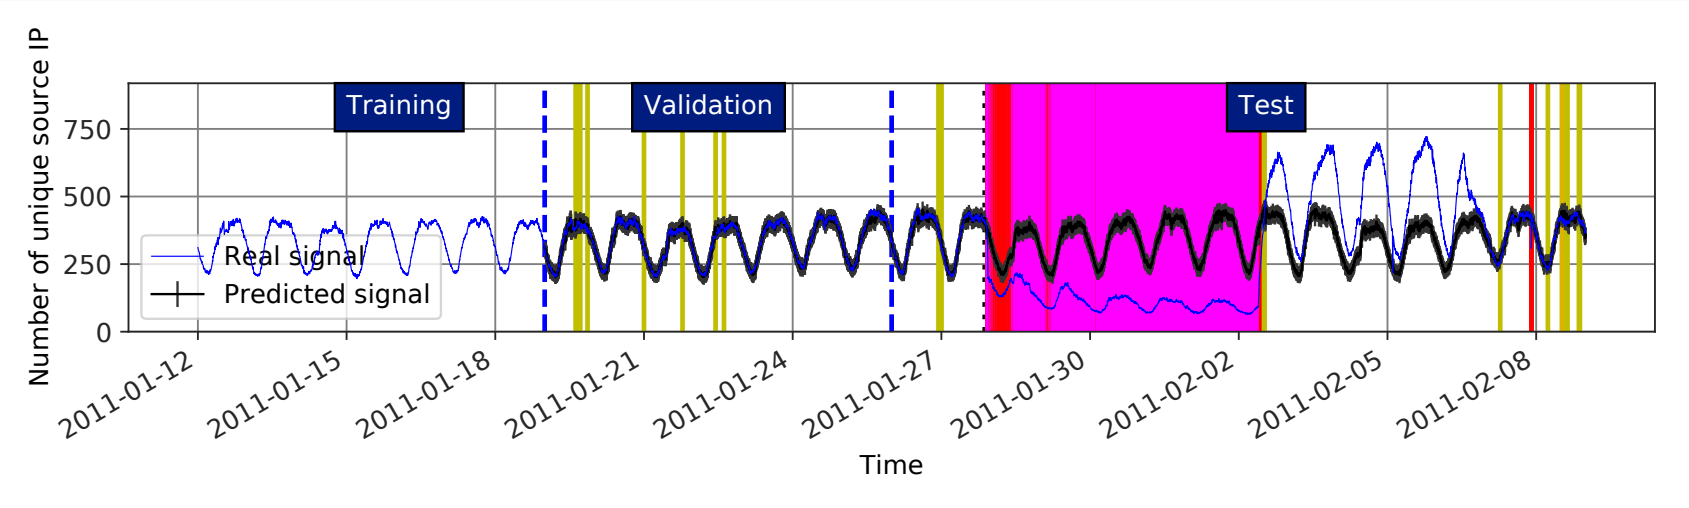

Figure 14: Outage detection on the Egyptian data set

Evaluating our solution

Objective

- Identifying *thresholds*

Thresholds

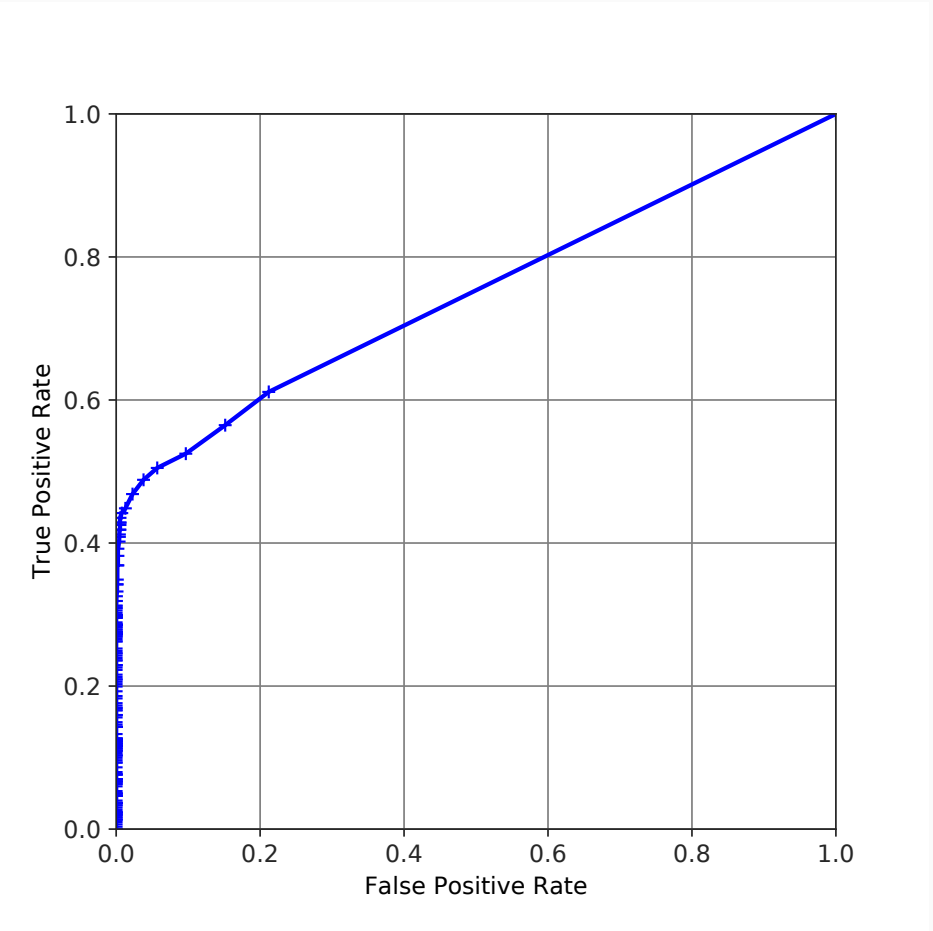

Figure 15: ROC curve

Evaluating our solution

Objective

- Identifying *thresholds*

Thresholds

- 60% — 20%

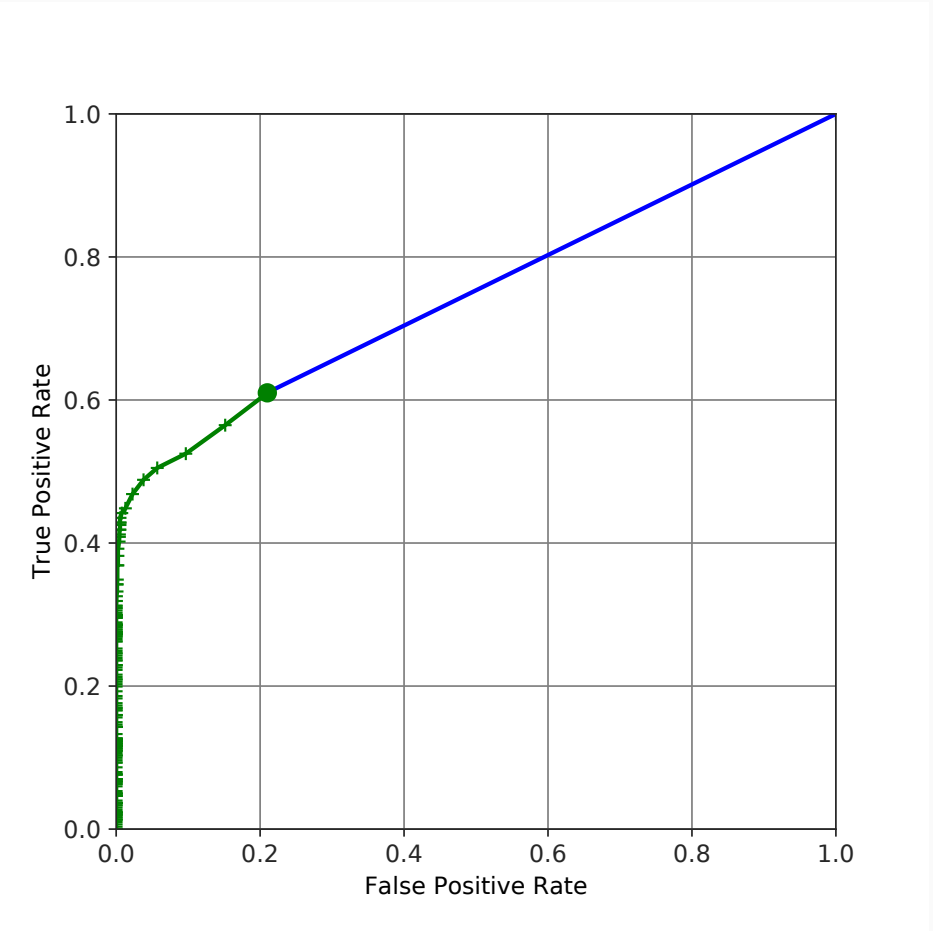

Figure 15: ROC curve

Evaluating our solution

Objective

- Identifying *thresholds*

Thresholds

- 60% — 20%
- 45% — 1%

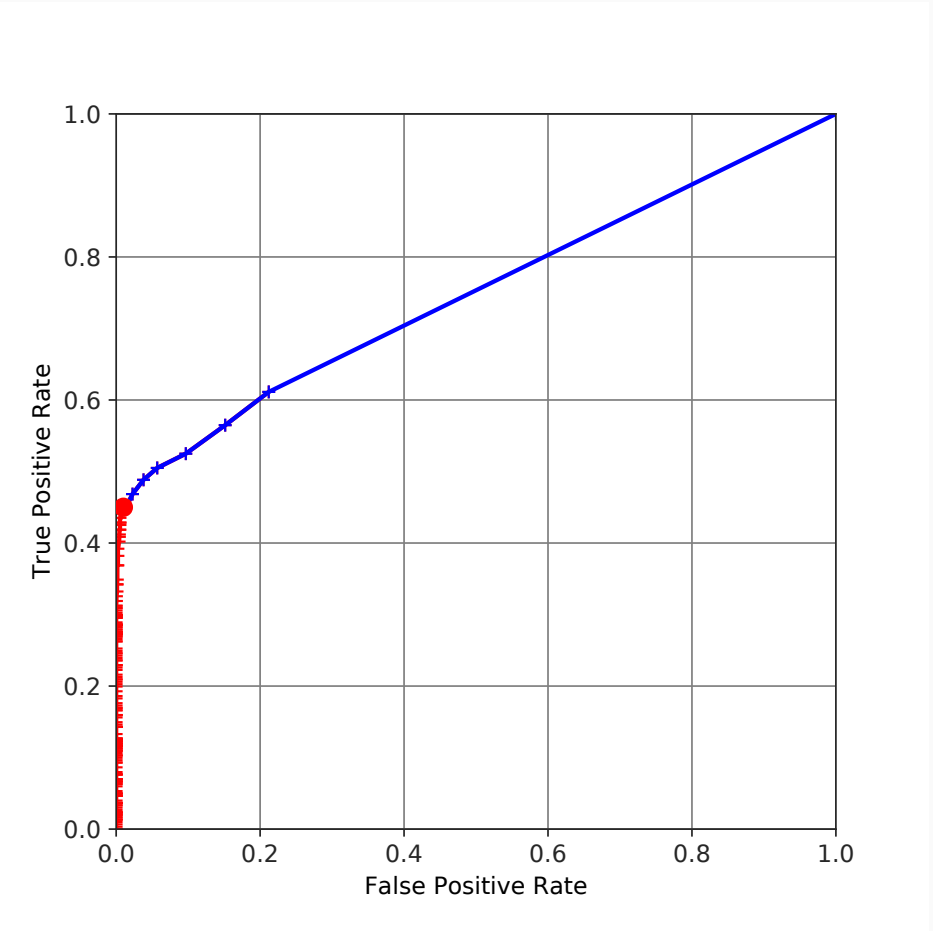

Figure 15: ROC curve

Application of the different thresholds

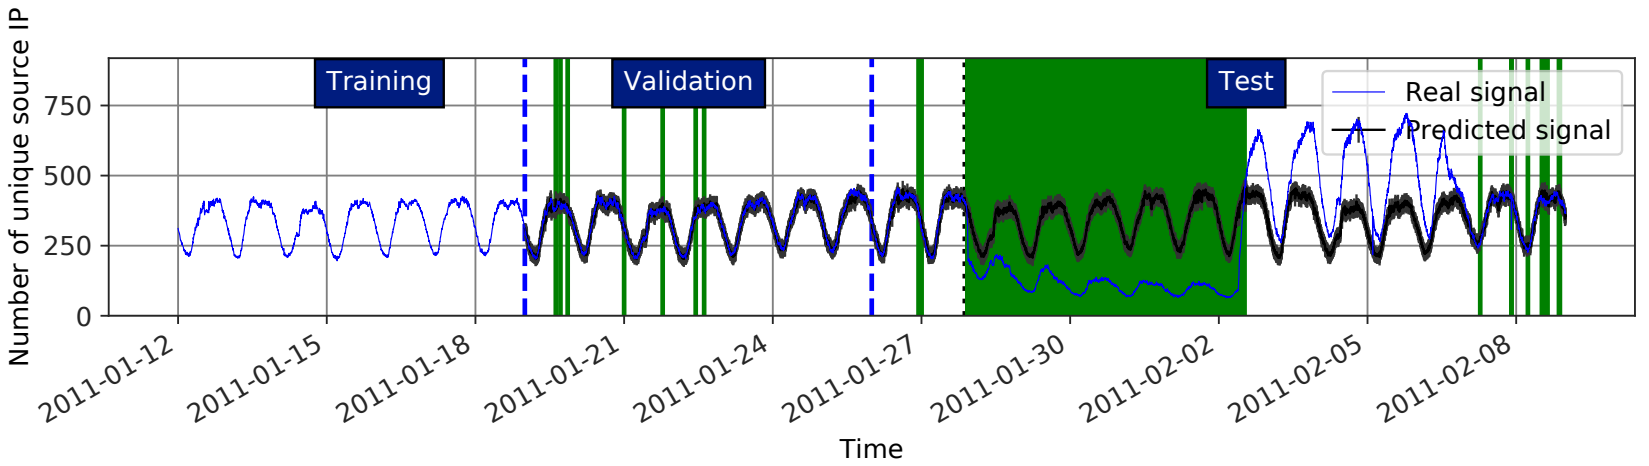

Figure 16: 60% — 20% threshold on the Egyptian data set

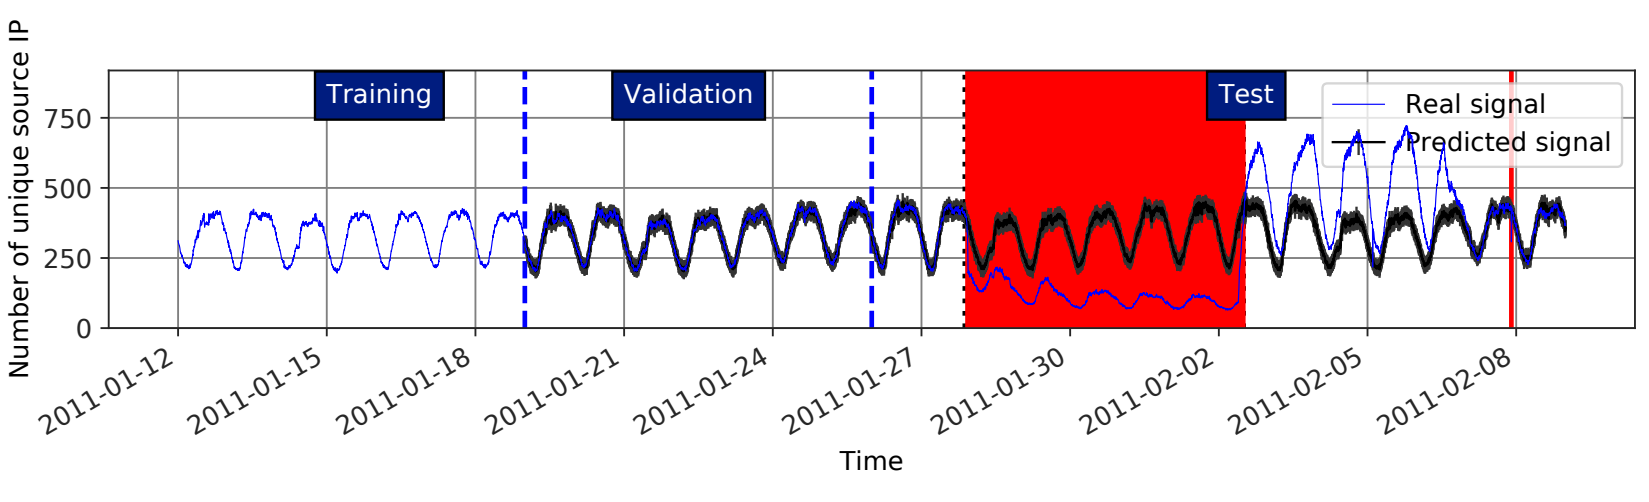

Figure 17: 45% — 1% threshold on the Egyptian data set

Conclusion

Perspectives

- Cross-validation with a different outage detection technique

Perspectives

- Cross-validation with a different outage detection technique
- Using our technique on similar data sets

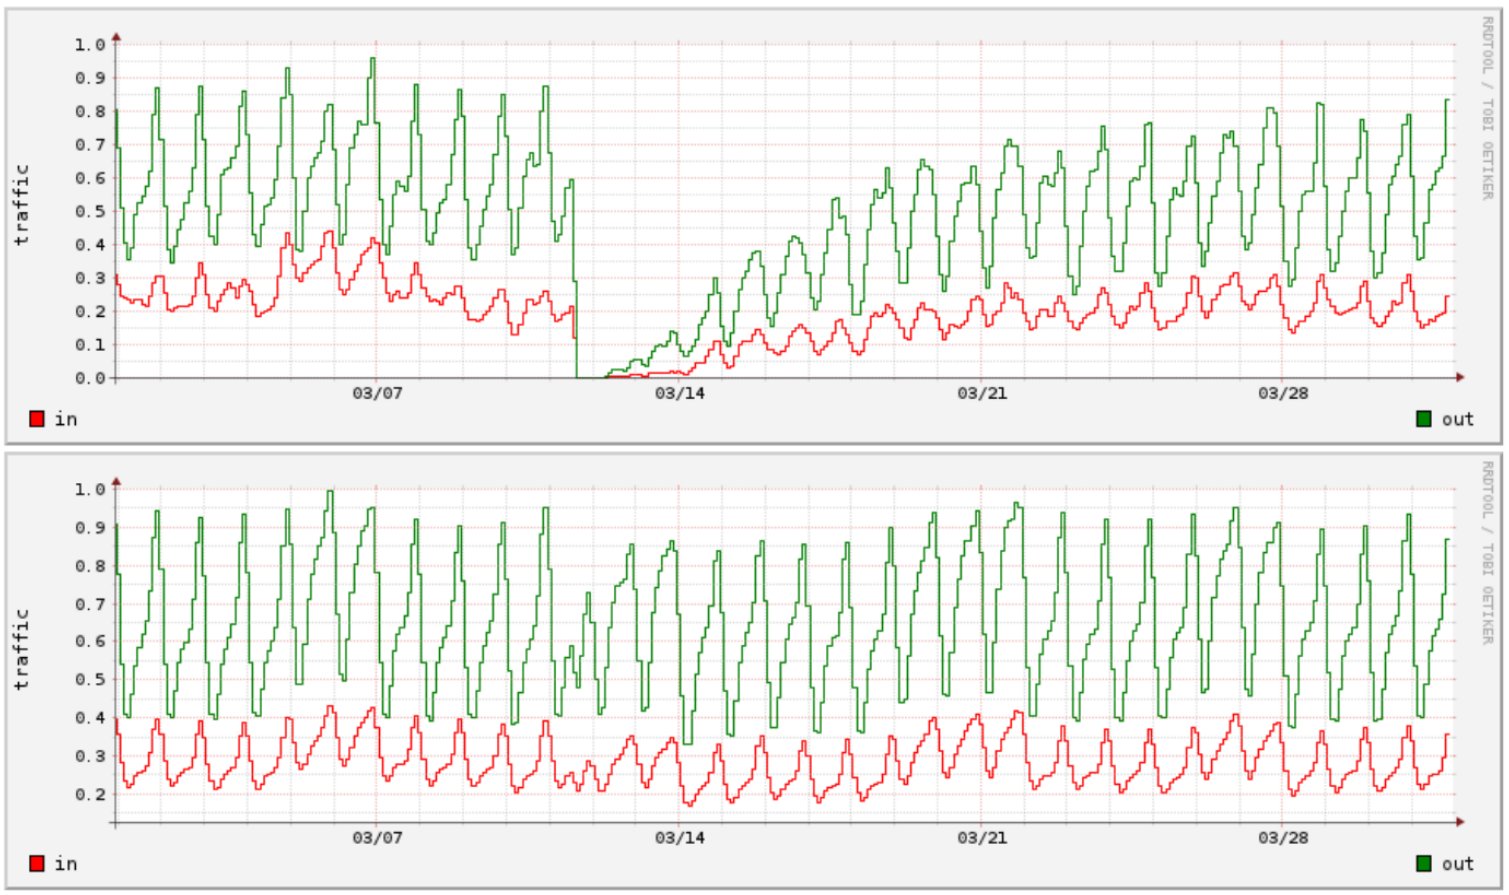

Figure 18: Japanese traffic for the March 2011 earthquake, Miyagi prefecture (top) and nationwide (bottom) [2]

Zachary's examples

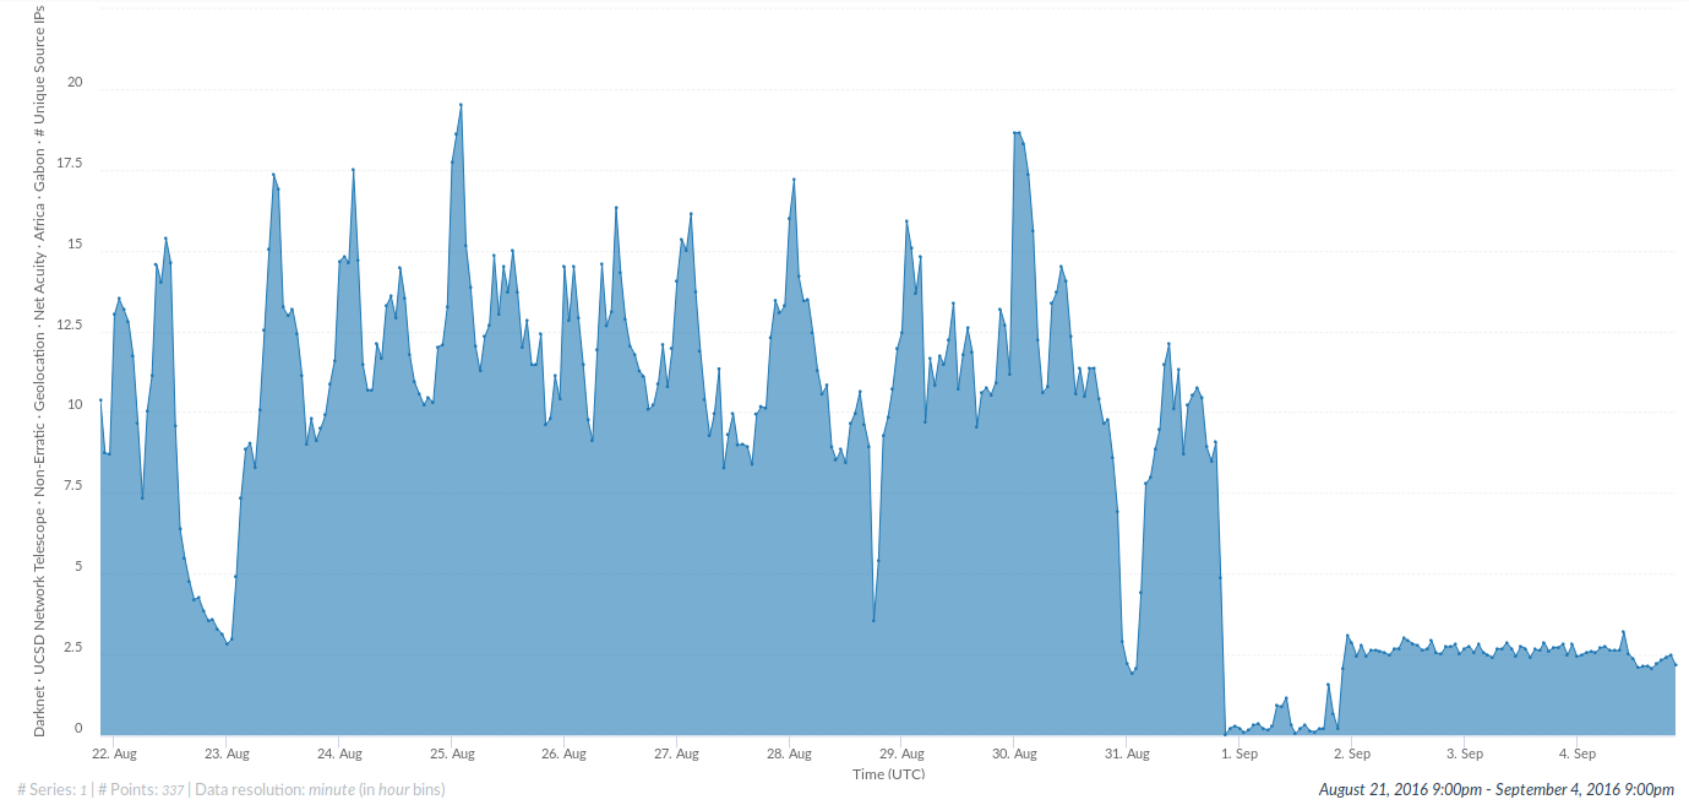

Figure 19: Gabon – August 27, 2016

Zachary's examples

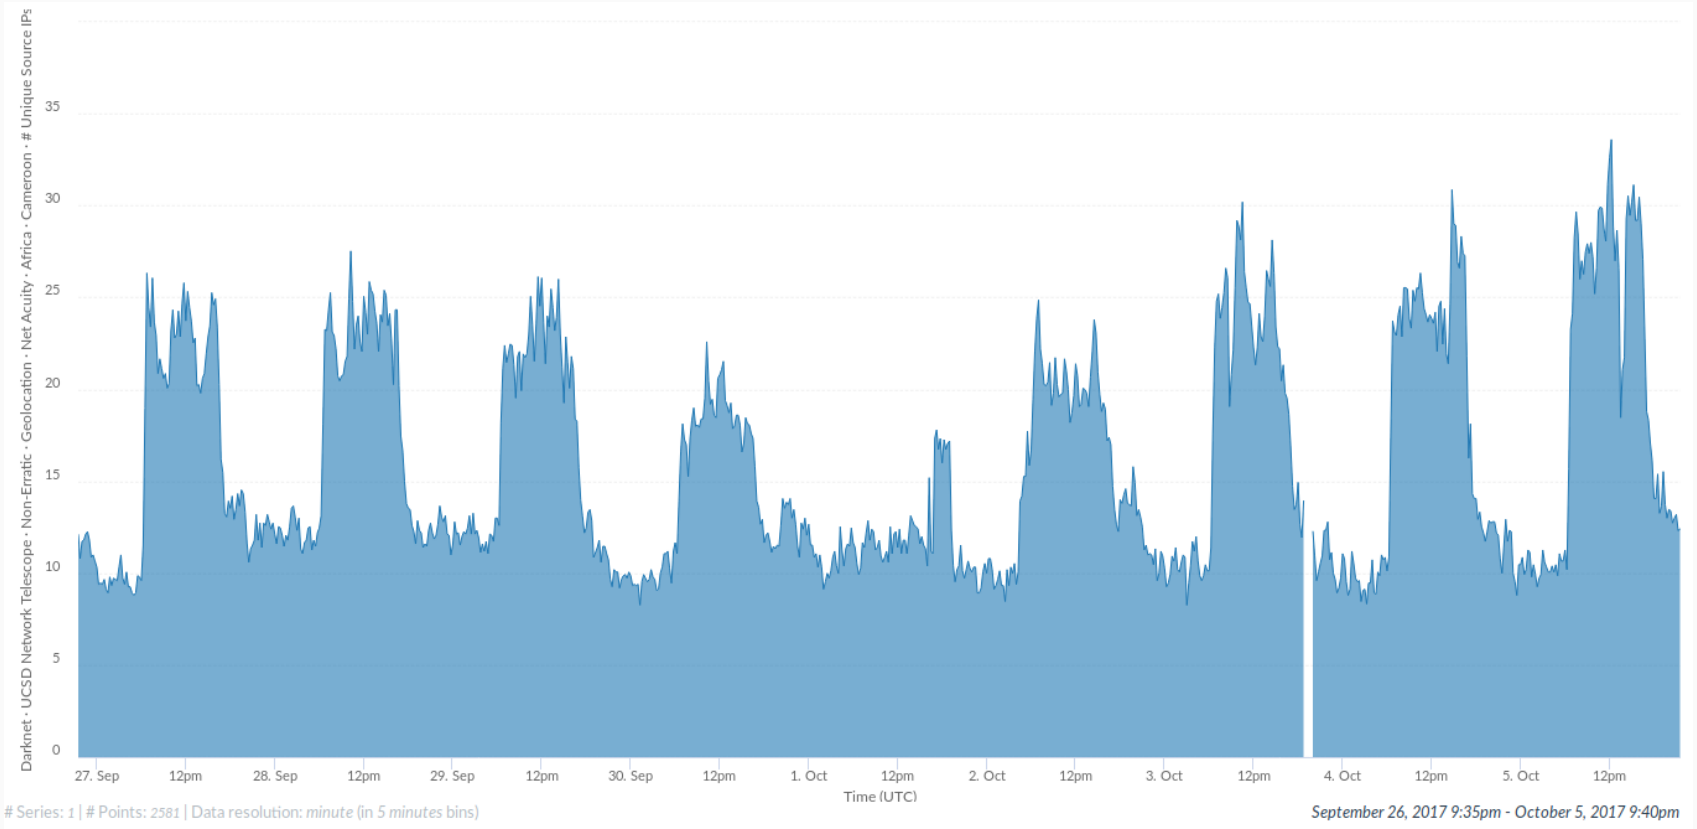

Figure 19: Cameroon – October 1, 2017

Summary

Goal

- Detecting worldwide Internet outages

Data Source

- Internet background radiation, a passive source with global coverage

Solution used

- SARIMA, a time series forecasting technique

Bibliography i

[1] Karyn Benson et al. “Gaining insight into as-level outages through analysis of internet background radiation”. In: *Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on*. IEEE. 2013, pp. 447–452.

[2] Kenjiro Cho et al. “The Japan earthquake: the impact on traffic and routing observed by a local ISP”. In: *Proceedings of the Special Workshop on Internet and Disasters*. ACM. 2011, p. 2.

[3] Alberto Dainotti et al. “Analysis of country-wide internet outages caused by censorship”. In: *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM. 2011, pp. 1–18.

Bibliography ii

[4] Uli Harder et al. “Observing internet worm and virus attacks with a small network telescope”. In: *Electronic Notes in Theoretical Computer Science* 151.3 (2006), pp. 47–59.

[5] Lin Quan, John Heidemann, and Yuri Pradkin. “Trinocular: Understanding internet reliability through adaptive probing”. In: *ACM SIGCOMM Computer Communication Review*. Vol. 43. 4. ACM. 2013, pp. 255–266.

[6] Anant Shah et al. “Disco: Fast, good, and cheap outage detection”. In: *Network Traffic Measurement and Analysis Conference (TMA), 2017*. IEEE. 2017, pp. 1–9.

Backup slides

p, the number of *autoregressive* terms

Definition

An autoregressive model specifies that the output variable depends linearly on its own *previous values* and on an imperfectly predictable *stochastic term*.

Identifying *p*

p can be found using an *autocorrelation plot*, which is obtained by computing the correlation of a value with prior values.

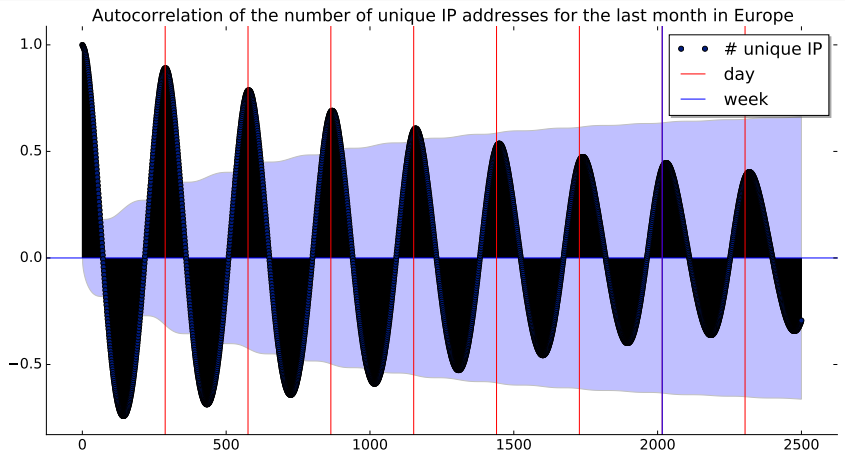

Figure 20: Autocorrelaion plot of the dataset presented on figure 1

q, the number of *moving-average* terms

Definition

An autoregressive model specifies that the output variable depends linearly on the *residual errors* of the previous forecasts and on an imperfectly predictable *stochastic term*.

Identifying *q*

p can be found using a *partial autocorrelation plot*, which is obtained by computing the autocorrelation of a value while accounting for previous dependencies.

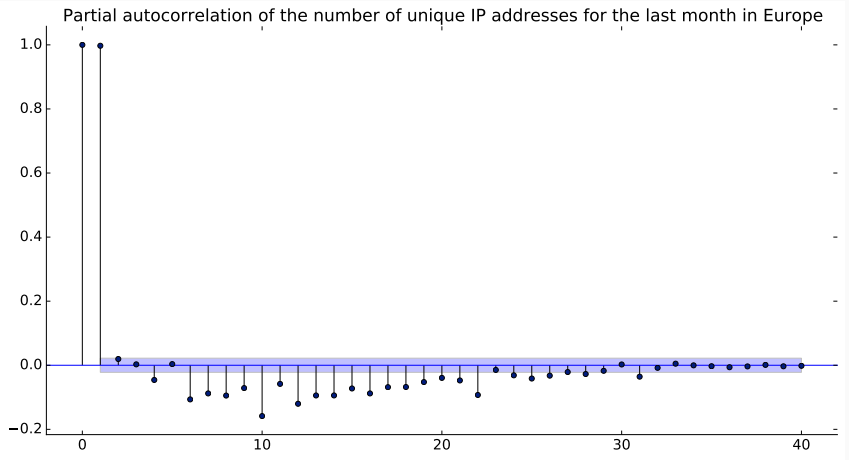

Figure 21: Partial autocorrelaion plot of the dataset presented on figure 1

d, the degree of integration

Definition

d is used to make the time series *stationary*. A stationary time series has a consistent mean, variance, and covariance over time. It also shows *no significant trend*.

Defining trends

A trend occurs when variations occur at specific time intervals

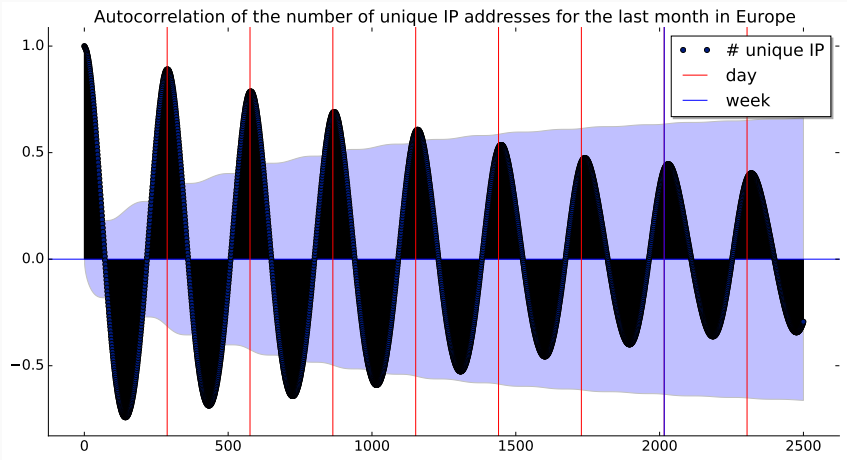

Figure 22: Illustration of a daily trend in the data

Non-seasonality or trends can reduce the accuracy of a predictive model, which is why they need to be removed

Removing Seasonalities

Differencing the data

A seasonality can be removed using a *difference*. A difference of degree d is equivalent to $Y'_n = Y_n - Y_{n-d}$

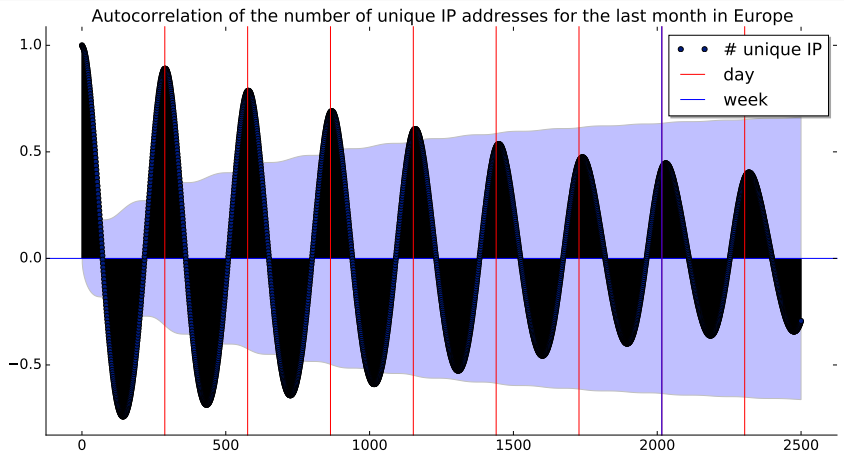

Figure 23: Illustration of a daily trend in the data

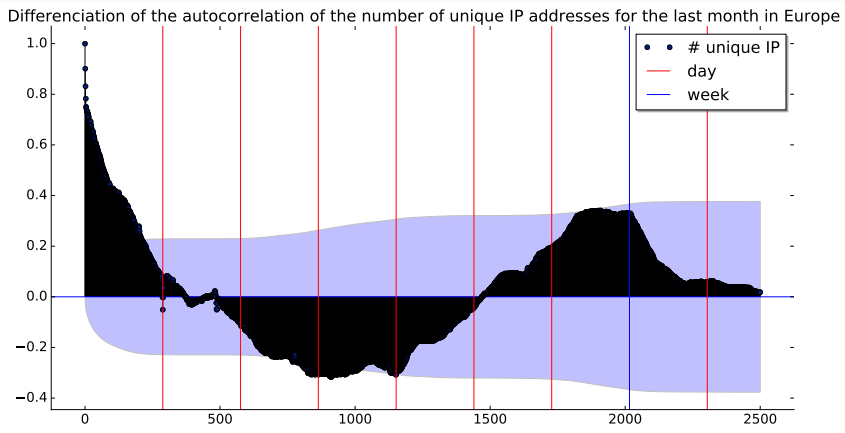

Figure 24: Data from figure 1 that has been differenced by a day

Time series analysis and forecasting techniques

End goal

- *Real-time predictions* on *thousands of* time series

	Temporal Component	Computationally cheap	Fit for real-time predictions
Fourier Transform	✗	✓	✗
Wavelet Transform	✓	✓	✗
Neural Networks	✓	✗	✓
SARIMA	✓	✓	✓

Figure 25: Table comparing the diferent forecasting techniques