Learning network states from RTT measurements

M. Mouchet T. Chonavel S. Vaton

IMT Atlantique, France

RIPE 77, October 2018

IMT Atlantique
Bretagne-Pays de la Loire
École Mines-Télécom

What constitutes the delay on the Internet ?

Orders of magnitude

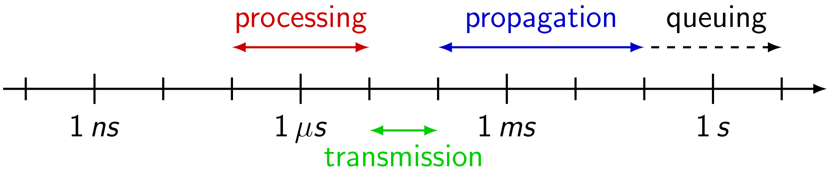

processing 100 $ns \rightarrow 10\,\mu s$

transmission 10 $\mu s \rightarrow 100\,\mu s$

propagation 100 $\mu s \rightarrow 100\,ms$

queuing up to seconds (bufferbloat)

What constitutes the delay on the Internet ?

Impact of traffic level and routing changes

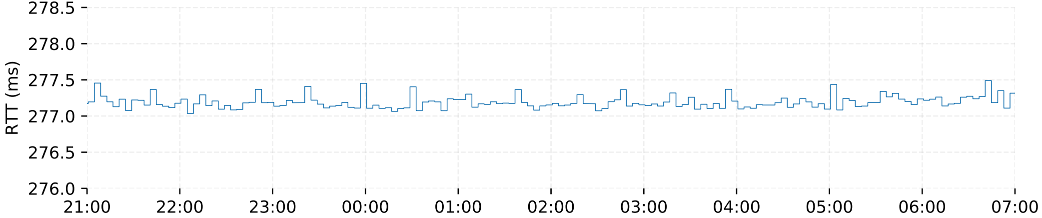

RTT measurements between at-vie-as1120 and vn-sgn-as24176.

What constitutes the delay on the Internet ?

Impact of traffic level and routing changes

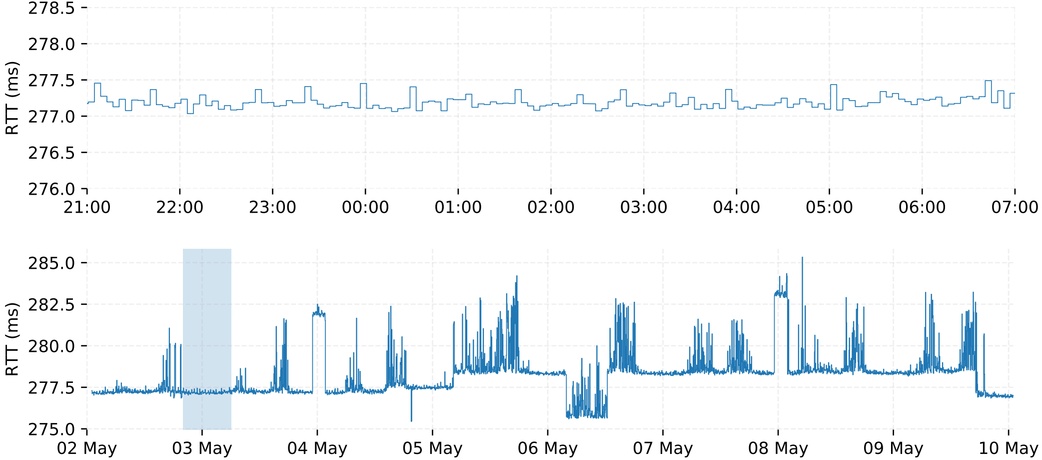

RTT measurements between at-vie-as1120 and vn-sgn-as24176.

Can we find back the hidden network states ?

Objective

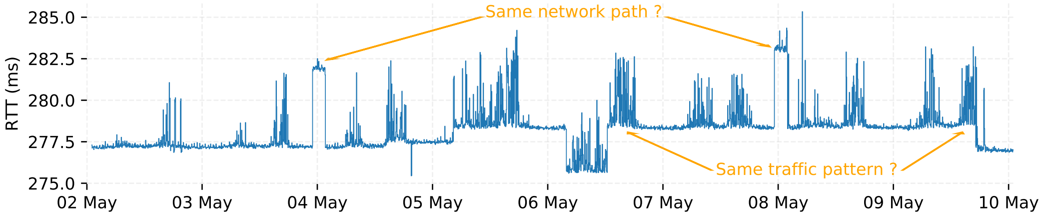

Can we find back the hidden network states ?

Objective

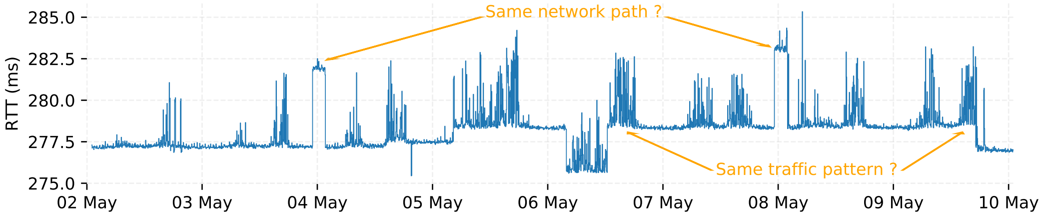

We want to associate each delay observation to a particular network state (network path, traffic level)

Can we find back the hidden network states ?

Use cases

Can we find back the hidden network states ?

Use cases

Detection of new network states

Anomaly detection
Traffic engineering

Can we find back the hidden network states ?

Use cases

Detection of new network states

Anomaly detection
Traffic engineering

A-posteriori analysis of network events

Correlation with incidents reports

Can we find back the hidden network states ?

Use cases

Detection of new network states

Anomaly detection
Traffic engineering

A-posteriori analysis of network events

Correlation with incidents reports

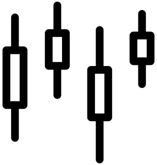

Statistical analysis

Studying patterns
Summarizing measurements

Can we find back the hidden network states ?

Why not using traceroutes ?

Can we find back the hidden network states ?

Why not using traceroutes ?

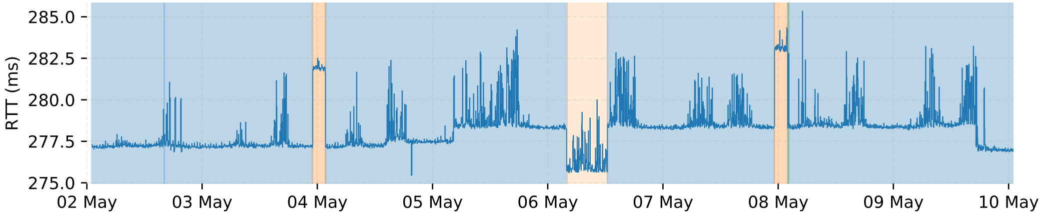

IP paths detected in the forward traceroute (one color per path).

Can we find back the hidden network states ?

Why not using traceroutes ?

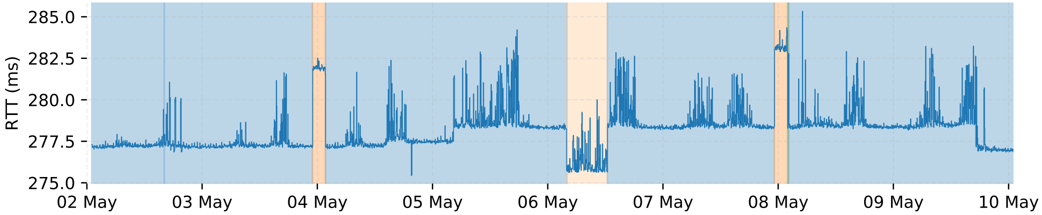

IP paths detected in the forward traceroute (one color per path).

- Traceroutes are more *expensive* and historical data is not always available;

Can we find back the hidden network states ?

Why not using traceroutes ?

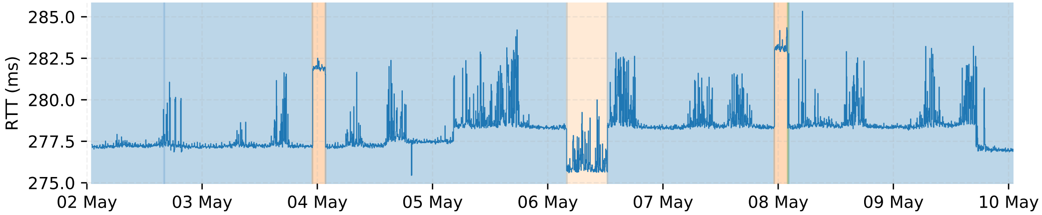

IP paths detected in the forward traceroute (one color per path).

- Traceroutes are more *expensive* and historical data is not always available;
- Forward and reverse traceroutes are needed for a complete view;

Can we find back the hidden network states ?

Why not using traceroutes ?

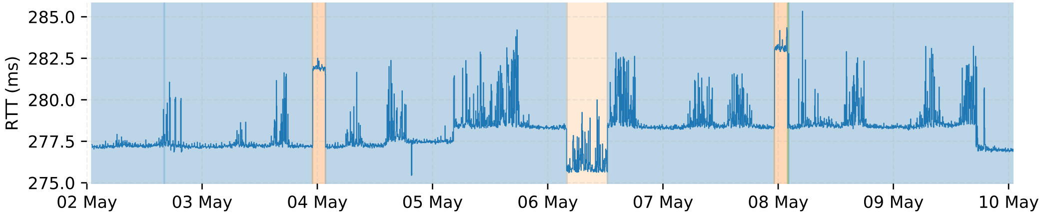

IP paths detected in the forward traceroute (one color per path).

- Traceroutes are more *expensive* and historical data is not always available;
- Forward and reverse traceroutes are needed for a complete view;
- They are blind to congestion and changes under the IP layer;

Can we find back the hidden network states ?

Unsupervised machine learning

Can we find back the hidden network states ?

Unsupervised machine learning

"**Clustering** is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups."[1]

[1] en.wikipedia.org/wiki/Cluster_analysis

Can we find back the hidden network states ?

Unsupervised machine learning

"**Clustering** is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups."[1]

$\Rightarrow$ **unsupervised learning**, in contrast to classification.

[1] en.wikipedia.org/wiki/Cluster_analysis

Can we find back the hidden network states ?

A Bayesian approach

Can we find back the hidden network states ?

A Bayesian approach

Build a generative model

Can we find back the hidden network states ?

A Bayesian approach

Build a generative model → Infer model parameters from observations

Can we find back the hidden network states ?

A Bayesian approach

Build a generative model → Infer model parameters from observations → Assign observations to hidden states

Can we find back the hidden network states ?

A Bayesian approach

Build a generative model → Infer model parameters from observations → Assign observations to hidden states

If we (loosely) know the model that generated the data, there are powerful statistical methods to infer the model parameters from the observed data.

Can we find back the hidden network states ?

Which generative model ?

Can we find back the hidden network states ?

Which generative model ?

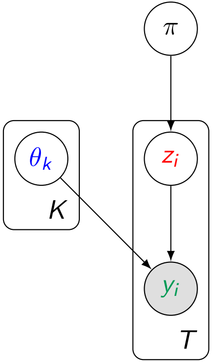

Independent observations

(Mixture model)

z: network state, y: observed delay, θ: delay distribution params., π: proportions

Can we find back the hidden network states ?

Which generative model ?

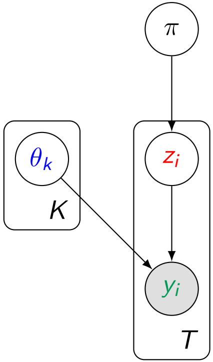

Independent observations
(Mixture model)

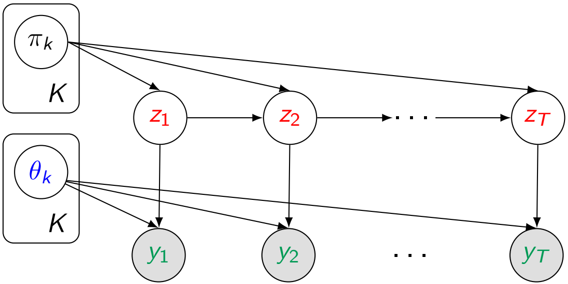

Non-independent observations
(Hidden Markov model)

z: network state, y: observed delay, θ: delay distribution params., π: proportions

Can we find back the hidden network states ?

Hidden Markov models

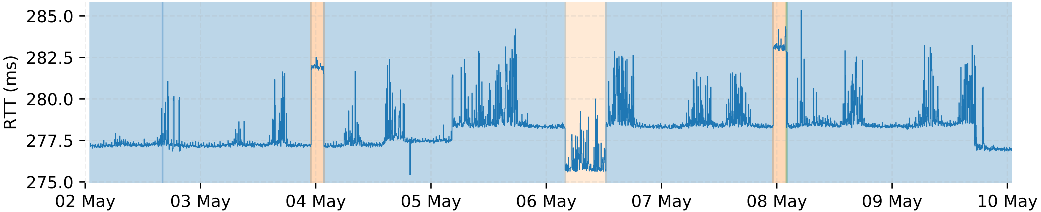

IP paths detected in the forward traceroute (one color per path).

Can we find back the hidden network states ?

Hidden Markov models

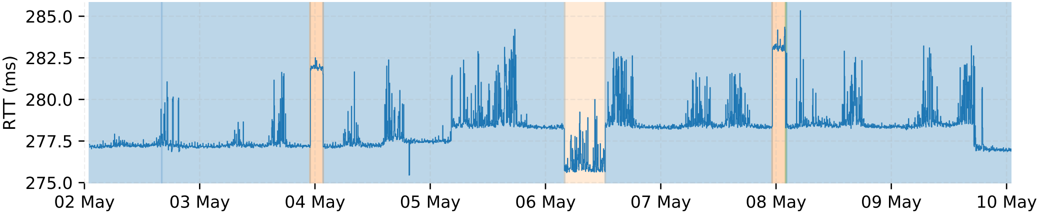

IP paths detected in the forward traceroute (one color per path).

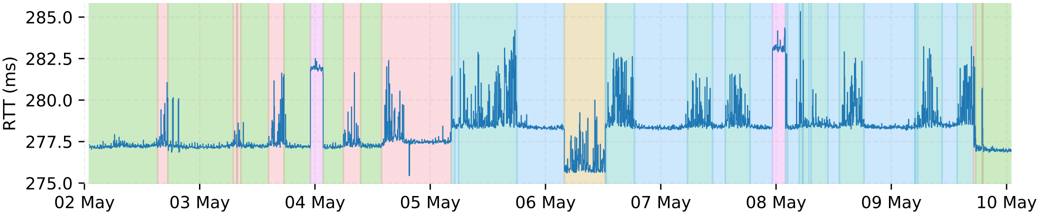

Network states learned using an hidden Markov model (one color per state).

Can we find back the hidden network states ?

Hidden Markov models

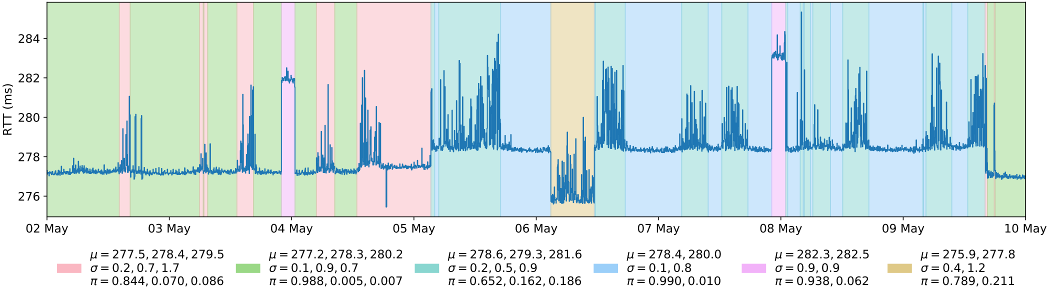

Network states learned using an hidden Markov model (one color per state).

Can we find back the hidden network states ?

Hidden Markov models

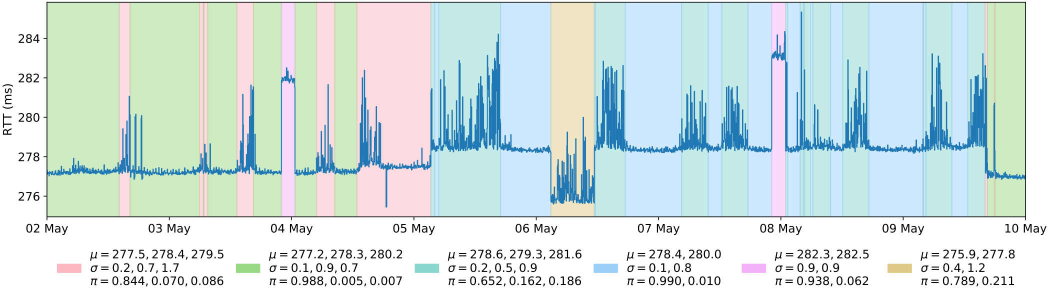

Network states learned using an hidden Markov model (one color per state).

- Accounting for temporal dependencies gives a (visually) better clustering;

Can we find back the hidden network states ?

Hidden Markov models

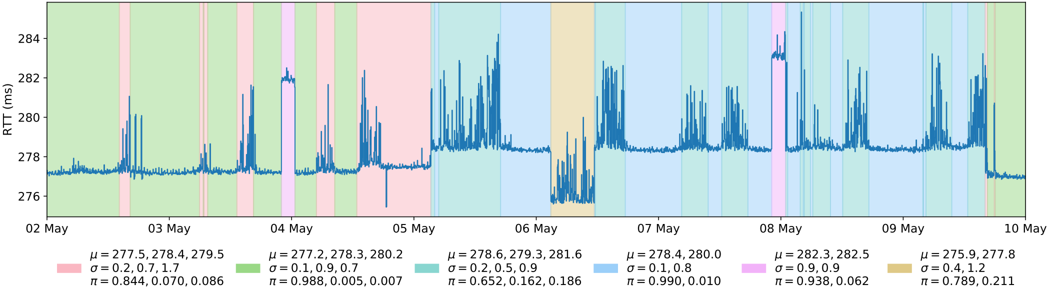

Network states learned using an hidden Markov model (one color per state).

- Accounting for temporal dependencies gives a (visually) better clustering;
- We can observe that any given learned state maps (generally) to only one IP path;

Can we find back the hidden network states ?

Hidden Markov models

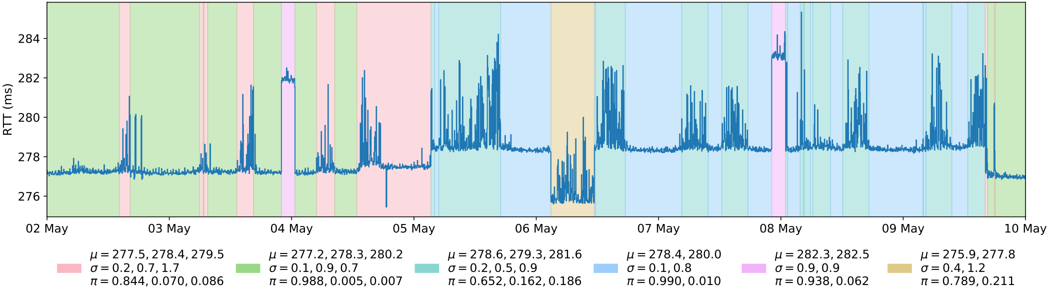

Network states learned using an hidden Markov model (one color per state).

- Accounting for temporal dependencies gives a (visually) better clustering;
- We can observe that any given learned state maps (generally) to only one IP path;
- We now have an information on the average duration of each state, and the relationship between them;

Applications

What is it good for ?

A single model for...

Applications

What is it good for ?

A single model for...

- **Operations**
  - Detect congestion in upstream networks
  - Detect (and react to) significant network changes
  - Study the correlation of some learned states and NOC tickets frequency
  - ...

Applications

What is it good for ?

A single model for...

- **Operations**
  - Detect congestion in upstream networks
  - Detect (and react to) significant network changes
  - Study the correlation of some learned states and NOC tickets frequency
  - ...
- **Analysis & Experiments**
  - A-posteriori study of network events
  - Statistical analysis
  - Parsimonious monitoring
  - ...

Applications

Detecting congestion in upstream networks

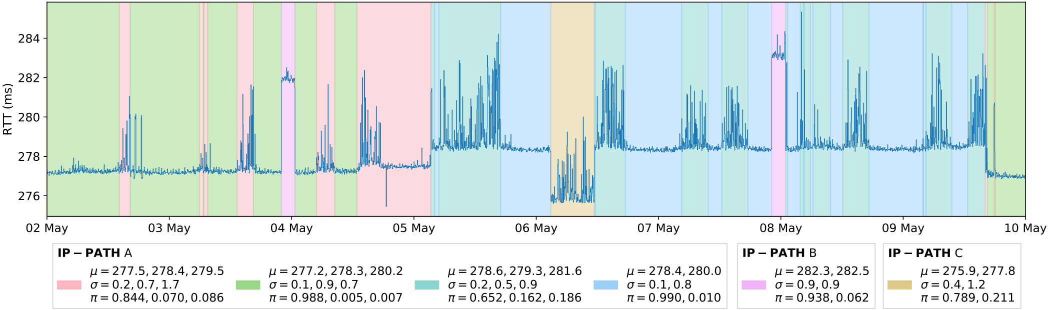

Applications

Detecting congestion in upstream networks

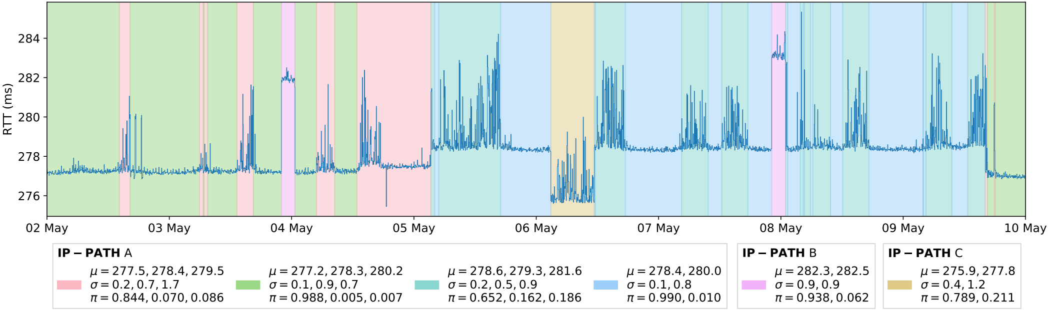

- We group learned states by IP path:
  - 4 states are learned for IP path A.
  - All IP path changes occur in a single AS (Cogent).

Applications

Detecting congestion in upstream networks

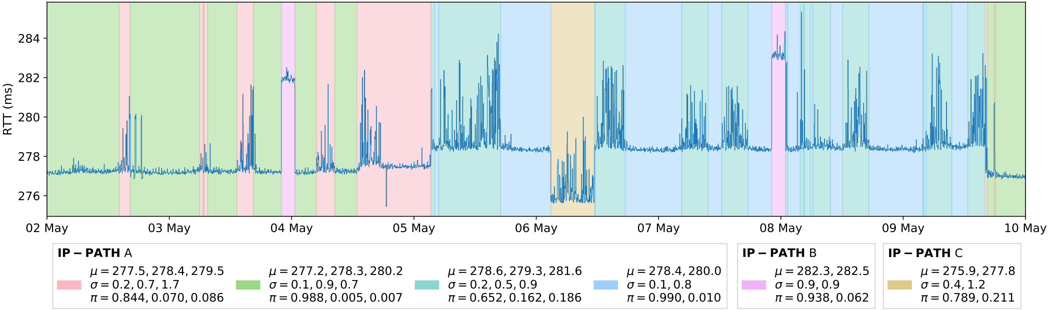

- We group learned states by IP path:
  - 4 states are learned for IP path *A*.
  - All IP path changes occur in a single AS (Cogent).
- Path A seems to experience periodic degradations in the transit AS.

Applications

A-posteriori study of network events

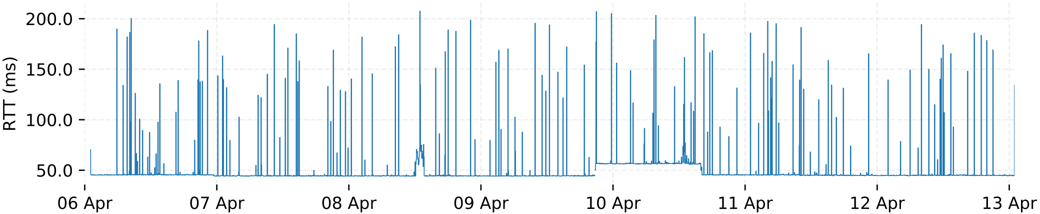

RTT measurements between at-vie-as1120 and ua-iev-as24725 during DE-CIX April 2018 outage.

Applications

A-posteriori study of network events

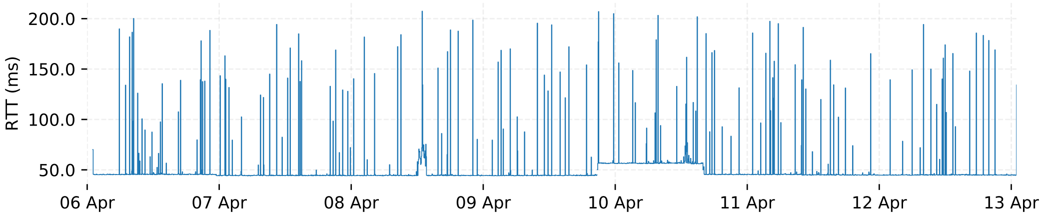

RTT measurements between at-vie-as1120 and ua-iev-as24725 during DE-CIX April 2018 outage.

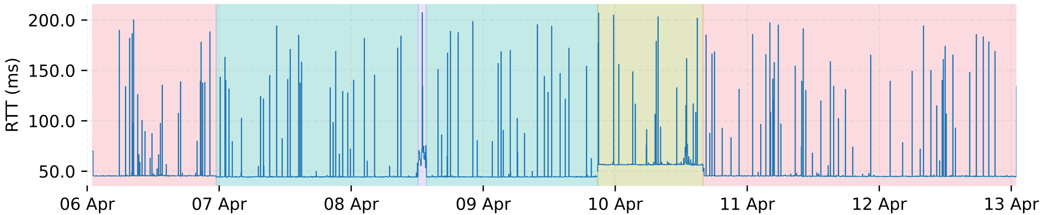

Network states learned.

Applications

A-posteriori study of network events

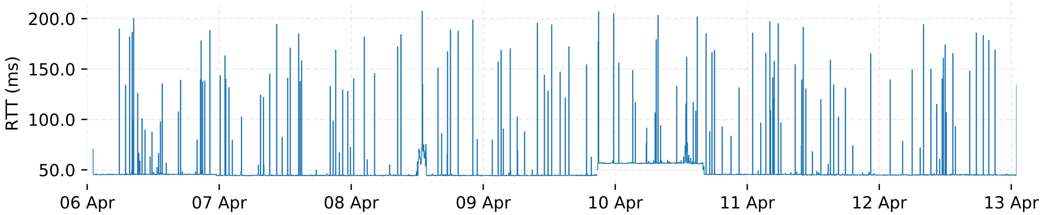

RTT measurements between at-vie-as1120 and ua-iev-as24725 during DE-CIX April 2018 outage.

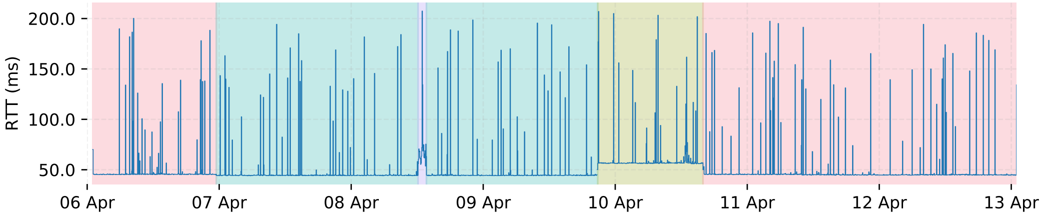

Network states learned.

$\Rightarrow$ paths with a new state during the outage timeframe were potentially affected.

Applications

Summarizing measurements & statistical analysis

Raw measurement

Applications

Summarizing measurements & statistical analysis

Raw measurement

- Delay observations

```
[277.308594, 277.117119, 277.202751, 277.185931, 277.194325,
916, 277.090857, 277.142608, 277.253547, 277.242663, 277.1068
289096, 277.124249, 277.082867, 277.088851, 277.10704, 277.17
77.056224, 277.186694, 277.185636, 277.198256, 277.516988, 27
, 277.239843, 277.243679, 277.199281, 277.133354, 277.142234,
8826, 277.242883, 277.10146, 277.242262, 277.391059, 277.0958
183232, 277.290016, 277.923457, 277.314035, 277.149393, 277.1
277.379276, 277.347547, 277.213029, 277.769285, 277.42185, 27
, 277.452252, 277.552072, 277.236618, 277.426263, 277.28154,
33, 277.432613, 277.220628, 277.164712, 277.365041, 277.40111
192, 277.145654, 277.169332, 277.198846, 277.165107, 277.3337
.225253, 277.154568, 277.161582, 277.334336, 277.134814, 277.
77.219654, 277.249793, 277.215609, 277.174268, 277.281447, 27
277.194162, 277.110135, 277.591023, 277.216204, 277.266436, 2
6, 277.337074, 277.227543, 277.262591, 277.141063, 277.318128
518, 277.311759, 277.210704, 277.27715, 277.239259, 277.38406
234368, 277.317254, 277.303589, 277.347363, 277.279396, 277.6
277.369207, 277.344948, 277.806792, 277.282587, 277.155303, 2
03, 277.219614, 277.350831, 279.124067, 277.922656, 277.55001
5963, 279.926155, 280.12368, 277.283426, 276.961512, 277.2687
.332774, 276.856348, 276.896102, 277.007425, 280.035673, 279.
277.222915, 277.314299, 277.501443, 277.156558, 277.292475, 2
25, 276.905996, 277.278617, 277.227999, 277.171932, 277.19581
```

Applications

Summarizing measurements & statistical analysis

Raw measurement

- Delay observations


```
[277.308594, 277.117119, 277.202751, 277.185931, 277.194325,
916, 277.090857, 277.142608, 277.253547, 277.242663, 277.1068
289096, 277.124249, 277.082867, 277.088851, 277.10704, 277.17
77.056224, 277.186694, 277.185636, 277.198256, 277.516988, 27
, 277.239843, 277.243679, 277.199281, 277.133354, 277.142234,
8826, 277.242883, 277.10146, 277.242262, 277.391059, 277.0958
183232, 277.290016, 277.923457, 277.314035, 277.149393, 277.1
277.379276, 277.347547, 277.213029, 277.769285, 277.42185, 27
, 277.452252, 277.552072, 277.236618, 277.426263, 277.28154,
33, 277.432613, 277.220628, 277.164712, 277.365041, 277.40111
192, 277.145654, 277.169332, 277.198846, 277.165107, 277.3337
.225253, 277.154568, 277.161582, 277.334336, 277.134814, 277.
77.219654, 277.249793, 277.215609, 277.174268, 277.281447, 27
277.194162, 277.110135, 277.591023, 277.216204, 277.266436, 2
6, 277.337074, 277.227543, 277.262591, 277.141063, 277.318128
518, 277.311759, 277.210704, 277.27715, 277.239259, 277.38406
234368, 277.317254, 277.303589, 277.347363, 277.279396, 277.6
277.369207, 277.344948, 277.806792, 277.282587, 277.155303, 2
03, 277.219614, 277.350831, 279.124067, 277.922656, 277.55001
5963, 279.926155, 280.12368, 277.283426, 276.961512, 277.2687
.332774, 276.856348, 276.896102, 277.007425, 280.035673, 279.
277.222915, 277.314299, 277.501443, 277.156558, 277.292475, 2
25, 276.905996, 277.278617, 277.227999, 277.171932, 277.1958
```


HMM

Applications

Summarizing measurements & statistical analysis

Raw measurement

- Delay observations


```
[277.308594, 277.117119, 277.202751, 277.185931, 277.194325,
916, 277.090857, 277.142608, 277.253547, 277.242663, 277.106
289096, 277.124249, 277.082867, 277.088851, 277.10704, 277.1
77.056224, 277.186694, 277.185636, 277.198256, 277.516988, 2
, 277.239843, 277.243679, 277.199281, 277.133354, 277.142234,
8826, 277.242883, 277.10146, 277.242262, 277.391059, 277.095
183232, 277.290016, 277.923457, 277.314035, 277.149393, 277.
277.379276, 277.347547, 277.213029, 277.769285, 277.42185, 2
, 277.452252, 277.552072, 277.236618, 277.426263, 277.28154,
33, 277.432613, 277.220628, 277.164712, 277.365041, 277.4011
192, 277.145654, 277.169332, 277.198846, 277.165107, 277.333
.225253, 277.154568, 277.161582, 277.334336, 277.134814, 277.
77.219654, 277.249793, 277.215609, 277.174268, 277.281447, 2
277.194162, 277.110135, 277.591023, 277.216204, 277.266436, 
6, 277.337074, 277.227543, 277.262591, 277.141063, 277.31812
518, 277.311759, 277.210704, 277.27715, 277.239259, 277.3840
234368, 277.317254, 277.303589, 277.347363, 277.279396, 277.
277.369207, 277.344948, 277.806792, 277.282587, 277.155303, 
03, 277.219614, 277.350831, 279.124067, 277.922656, 277.5500
5963, 279.926155, 280.12368, 277.283426, 276.961512, 277.268
.332774, 276.856348, 276.896102, 277.007425, 280.035673, 279.
277.222915, 277.314299, 277.501443, 277.156558, 277.292475, 
25, 276.905996, 277.278617, 277.227999, 277.171932, 277.1958
```


HMM

- Transition matrix
  - #states × #states matrix
  - Gives information about the average duration of each state and how states are connected

Applications

Summarizing measurements & statistical analysis

Raw measurement

- Delay observations


```
[277.308594, 277.117119, 277.202751, 277.185931, 277.194325,
916, 277.090857, 277.142608, 277.253547, 277.242663, 277.1068
289096, 277.124249, 277.082867, 277.088851, 277.10704, 277.17
77.056224, 277.186694, 277.185636, 277.198256, 277.516988, 27
, 277.239843, 277.243679, 277.199281, 277.133354, 277.142234,
8826, 277.242883, 277.10146, 277.242262, 277.391059, 277.0958
183232, 277.290016, 277.923457, 277.314035, 277.149393, 277.1
277.379276, 277.347547, 277.213029, 277.769285, 277.42185, 27
, 277.452252, 277.552072, 277.236618, 277.426263, 277.28154,
33, 277.432613, 277.220628, 277.164712, 277.365041, 277.4011
192, 277.145654, 277.169332, 277.198846, 277.165107, 277.3337
.225253, 277.154568, 277.161582, 277.334336, 277.134814, 277.
77.219654, 277.249793, 277.215609, 277.174268, 277.281447, 27
277.194162, 277.110135, 277.591023, 277.216204, 277.266436, 2
6, 277.337074, 277.227543, 277.262591, 277.141063, 277.31812
518, 277.311759, 277.210704, 277.27715, 277.239259, 277.3840
234368, 277.317254, 277.303589, 277.347363, 277.279396, 277.
277.369207, 277.344948, 277.806792, 277.282587, 277.155303, 2
03, 277.219614, 277.350831, 279.124067, 277.922656, 277.5500
5963, 279.926155, 280.12368, 277.283426, 276.961512, 277.268
.332774, 276.856348, 276.896102, 277.007425, 280.035673, 279.
277.222915, 277.314299, 277.501443, 277.156558, 277.292475, 2
25, 276.905996, 277.278617, 277.227999, 277.171932, 277.1958
```


HMM

- Transition matrix
  - #states $\times$ #states matrix
  - Gives information about the average duration of each state and how states are connected
- Observations distributions parameters
  - $\propto$ #states
  - Gives information about the statistical distribution of the delay in each states

Applications

Summarizing measurements & statistical analysis

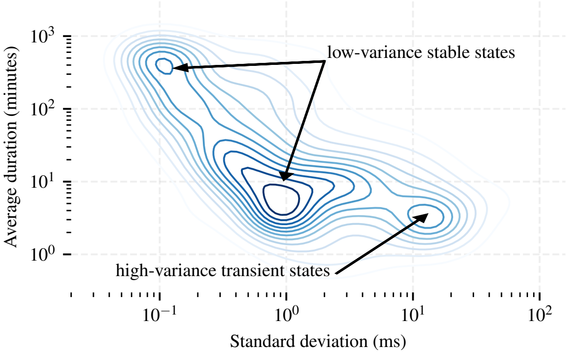

Relationship between state duration and observations standard deviation on RIPE Atlas anchoring mesh measurements.

⇒ We observe an inverse relationship between the average duration of a state and its standard deviation (i.e. stable states last longer).

Conclusion

- RTT observations depend on the underlying network state.
- Those states can be recovered using Hidden Markov models.
- Experiments shows that HMMs are a reasonable model for the RTT on the Internet.
- In comparison to other models (such as neural networks), HMMs parameters are easy to interpret (for a human being).

Conclusion

- **Current works:**
  - $\boxed{\text{RTT timeseries analysis using Bayesian HMMs}}$
  - Parsimonious monitoring[2] (reduction of up to 85% of the monitoring cost in routing overlays)

[2] S. Vaton, O. Brun, M. Mouchet, P. Belzarena, I. Amigo, B. J. Prabhu, and T. Chonavel. Joint minimization of monitoring cost and delay in overlay networks: optimal policies with a Markovian approach. *Journal of Network and Systems Management* (Aug. 2018). `https://doi.org/10.1007/s10922-018-9464-1`

Conclusion

- **Current works:**
  - RTT timeseries analysis using Bayesian HMMs
  - Parsimonious monitoring[2] (reduction of up to 85% of the monitoring cost in routing overlays)
- **Future works:**
  - Online (*real-time*) model learning
  - Learning of monitoring policies from scratch

[2] S. Vaton, O. Brun, M. Mouchet, P. Belzarena, I. Amigo, B. J. Prabhu, and T. Chonavel. Joint minimization of monitoring cost and delay in overlay networks: optimal policies with a Markovian approach. *Journal of Network and Systems Management* (Aug. 2018). https://doi.org/10.1007/s10922-018-9464-1

Conclusion

- **Current works:**
  - $\boxed{\text{RTT timeseries analysis using Bayesian HMMs}}$
  - Parsimonious monitoring[2] (reduction of up to 85% of the monitoring cost in routing overlays)
- **Future works:**
  - Online (*real-time*) model learning
  - Learning of monitoring policies from scratch

Thanks for your attention!

[2] S. Vaton, O. Brun, M. Mouchet, P. Belzarena, I. Amigo, B. J. Prabhu, and T. Chonavel. Joint minimization of monitoring cost and delay in overlay networks: optimal policies with a Markovian approach. *Journal of Network and Systems Management* (Aug. 2018). https://doi.org/10.1007/s10922-018-9464-1

Appendix

Mixture models

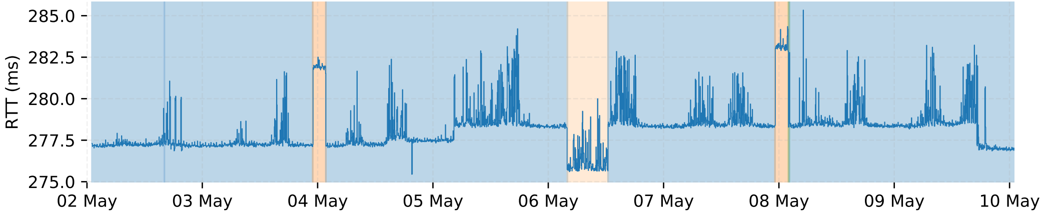

IP paths detected in the forward traceroute (one color per path).

Appendix

Mixture models

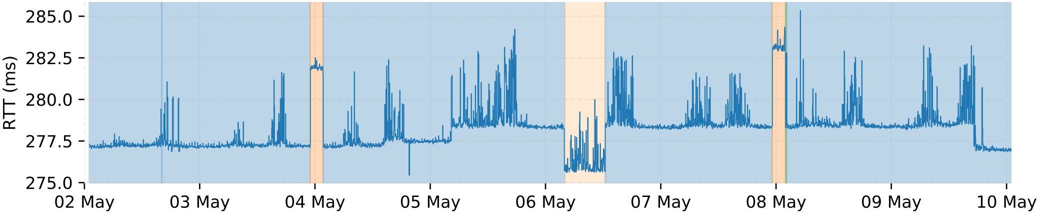

IP paths detected in the forward traceroute (one color per path).

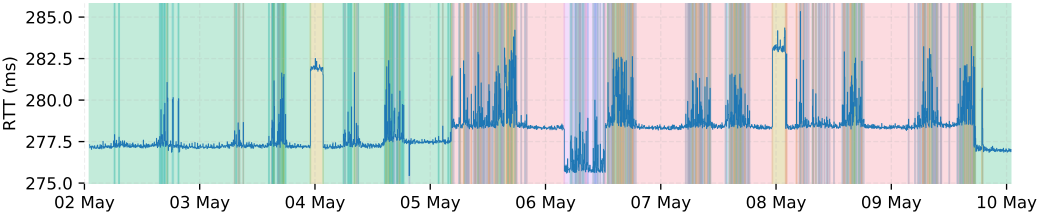

Network states learned using a mixture model (one color per state).

Appendix
Mixture models

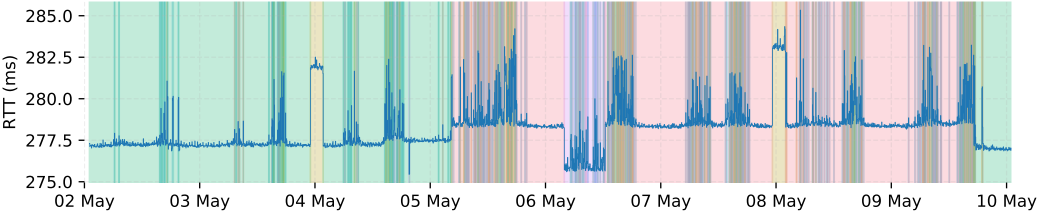

Network states learned using a mixture model (one color per state).

Mixture models does not account for temporal dependencies, thus:

Appendix

Mixture models

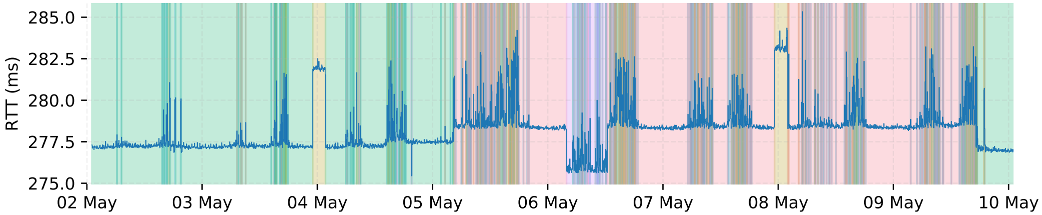

Network states learned using a mixture model (one color per state).

Mixture models does not account for temporal dependencies, thus:

- They fail to correctly cluster delay observations with an high variance;

Appendix

Mixture models

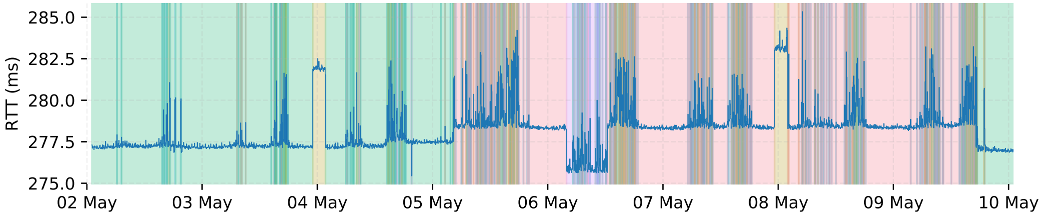

Network states learned using a mixture model (one color per state).

Mixture models does not account for temporal dependencies, thus:

- They fail to correctly cluster delay observations with an high variance;
- If we use state transitions to detect network anomalies, they would generate a lot of false alarms;

Appendix

IP/AS path correlation

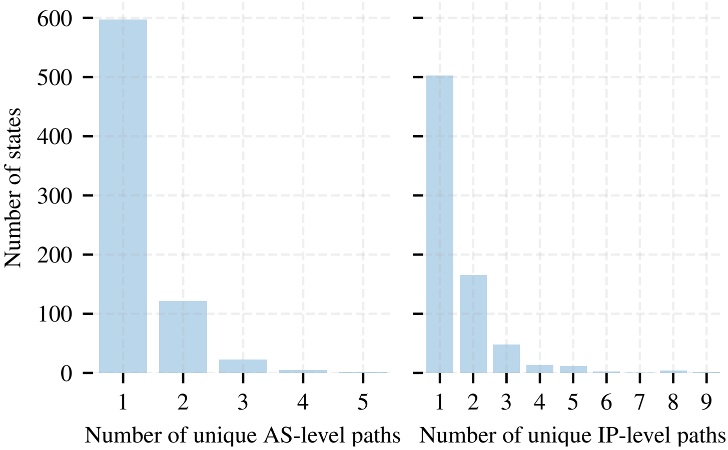

Appendix

Credits

- Creative Commons CC-BY:
  - Radar by IYIKON from the Noun Project
  - analysis by mynamepong from the Noun Project
  - statistics by Adnen Kadri from the Noun Project