ROMA TRE UNIVERSITY
Department of Engineering

Towards an Automated Investigation of the Impact of BGP Routing Changes on Network Delay Variations

Massimo Rimondini

Claudio Squarcella

Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

ROMA TRE UNIVERSITY
Department of Engineering

BGP Routing Changes Network Delay Variations

Massimo Rimondini
Claudio Squarcella
Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

BGP Routing Changes

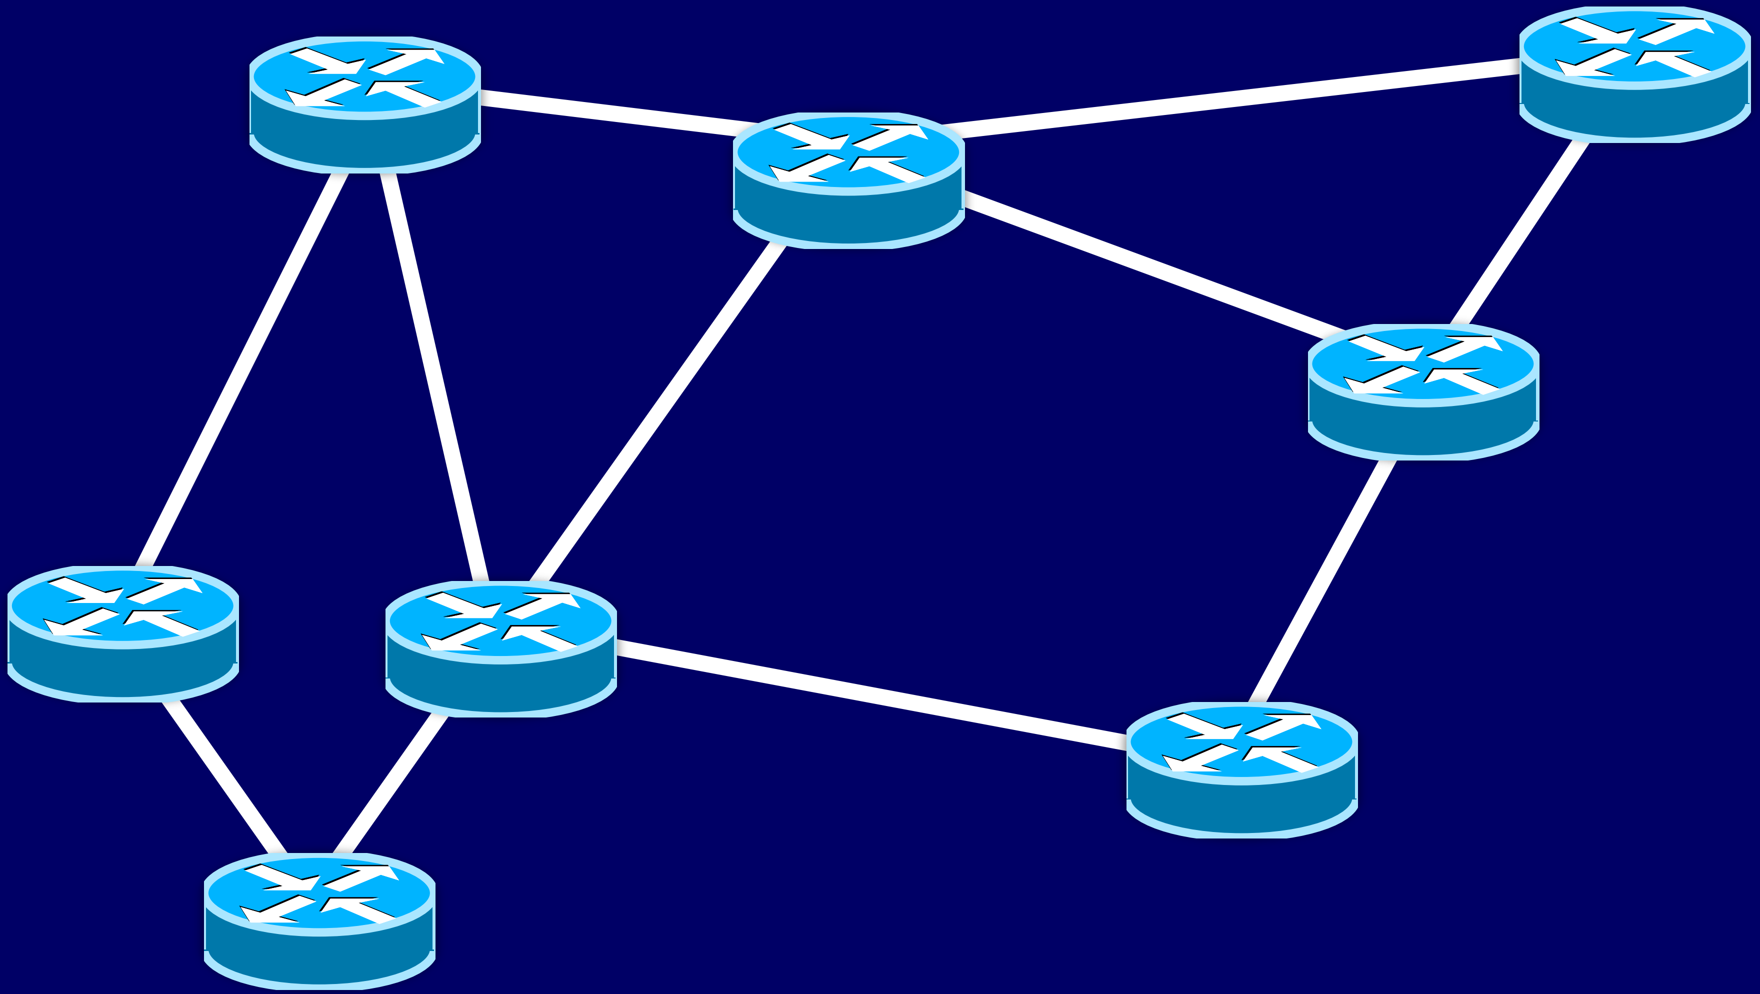

BGP Routing Changes

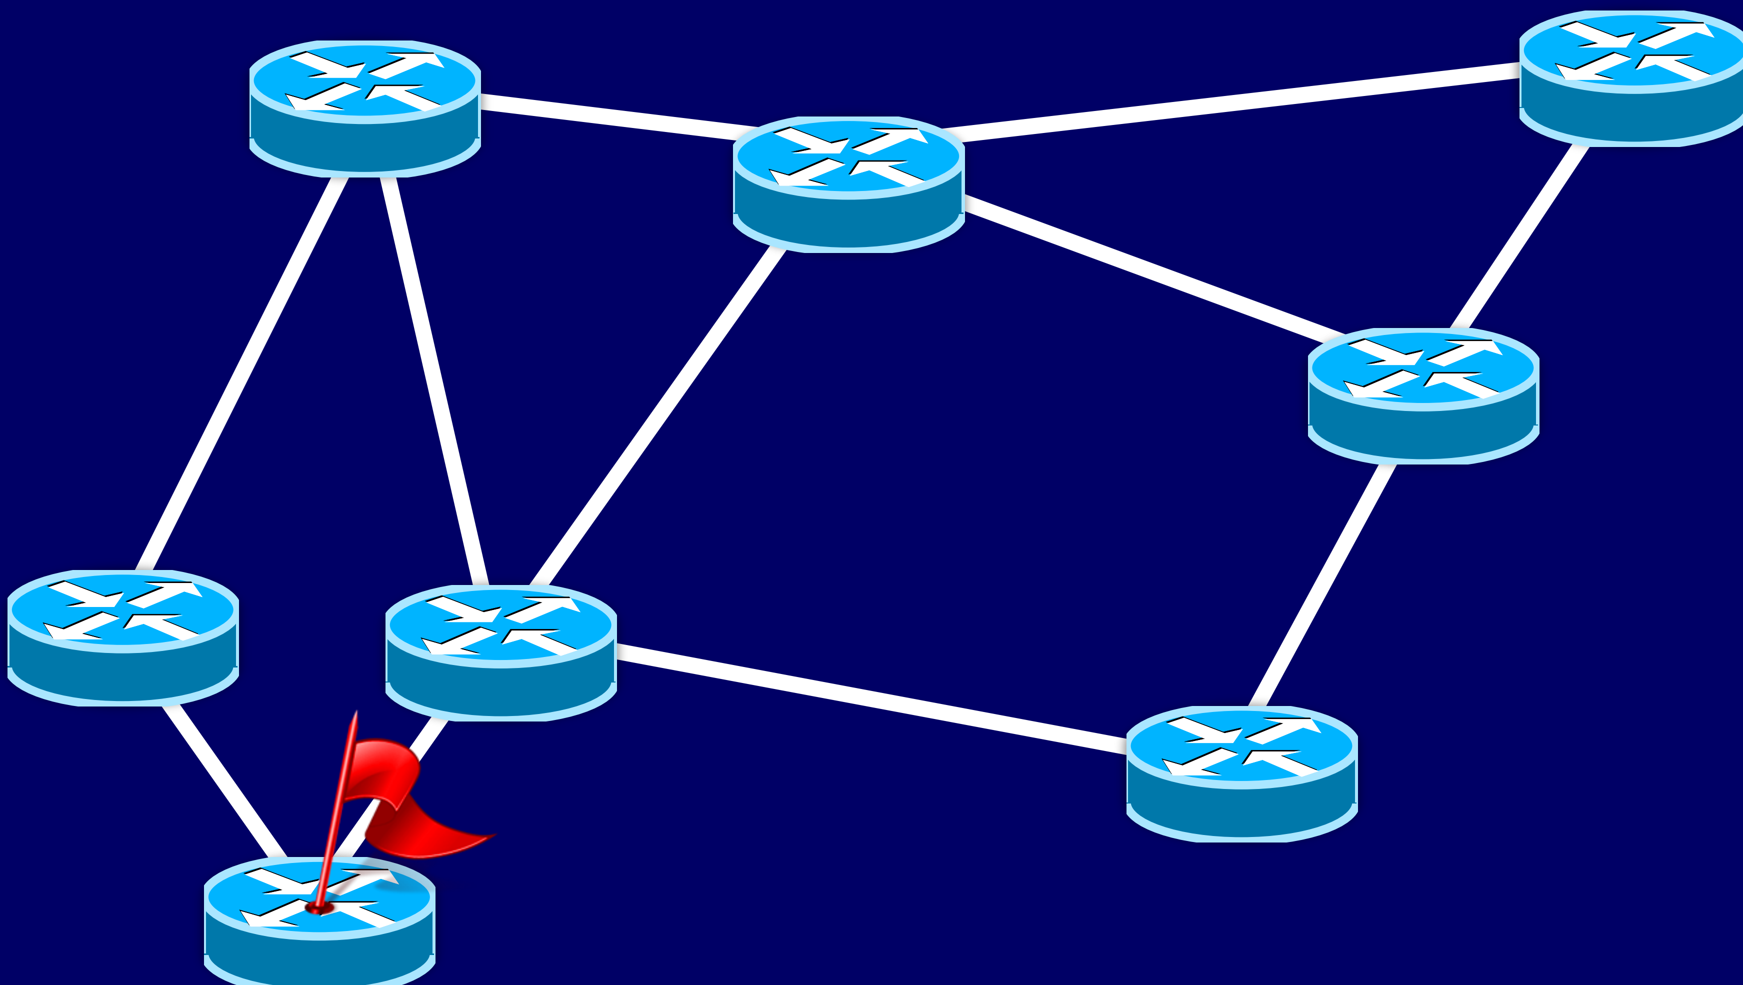

BGP Routing Changes

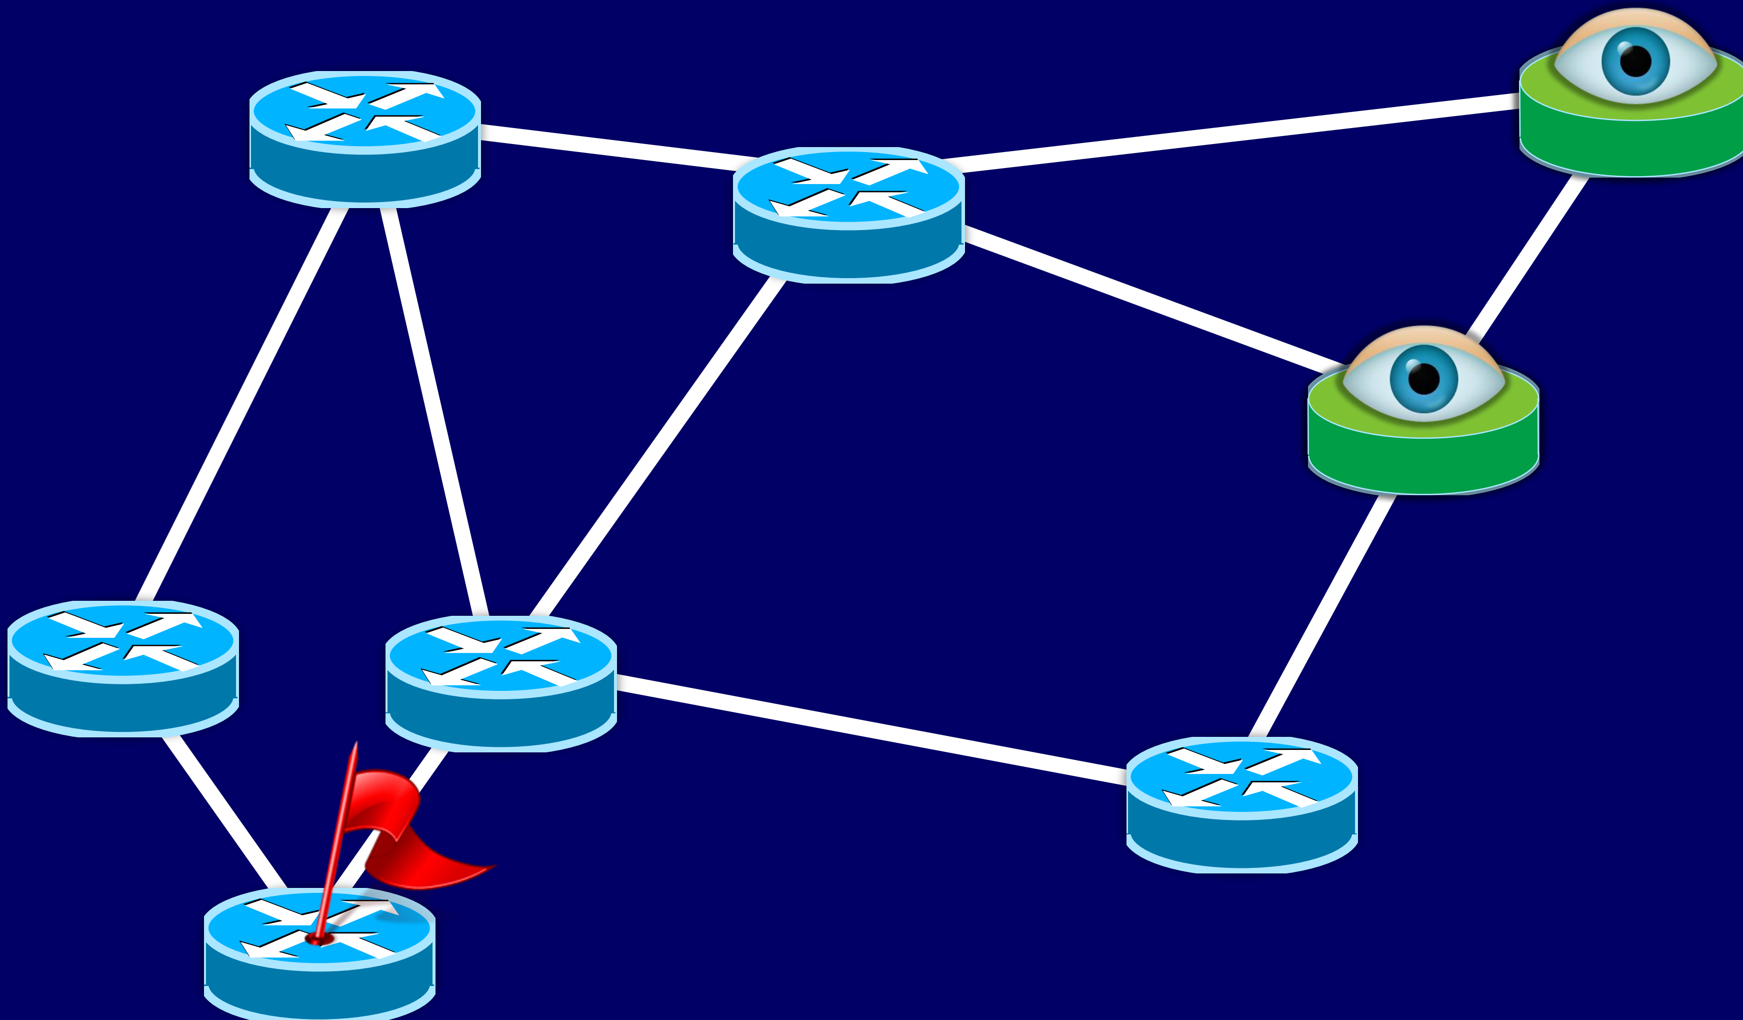

BGP Routing Changes

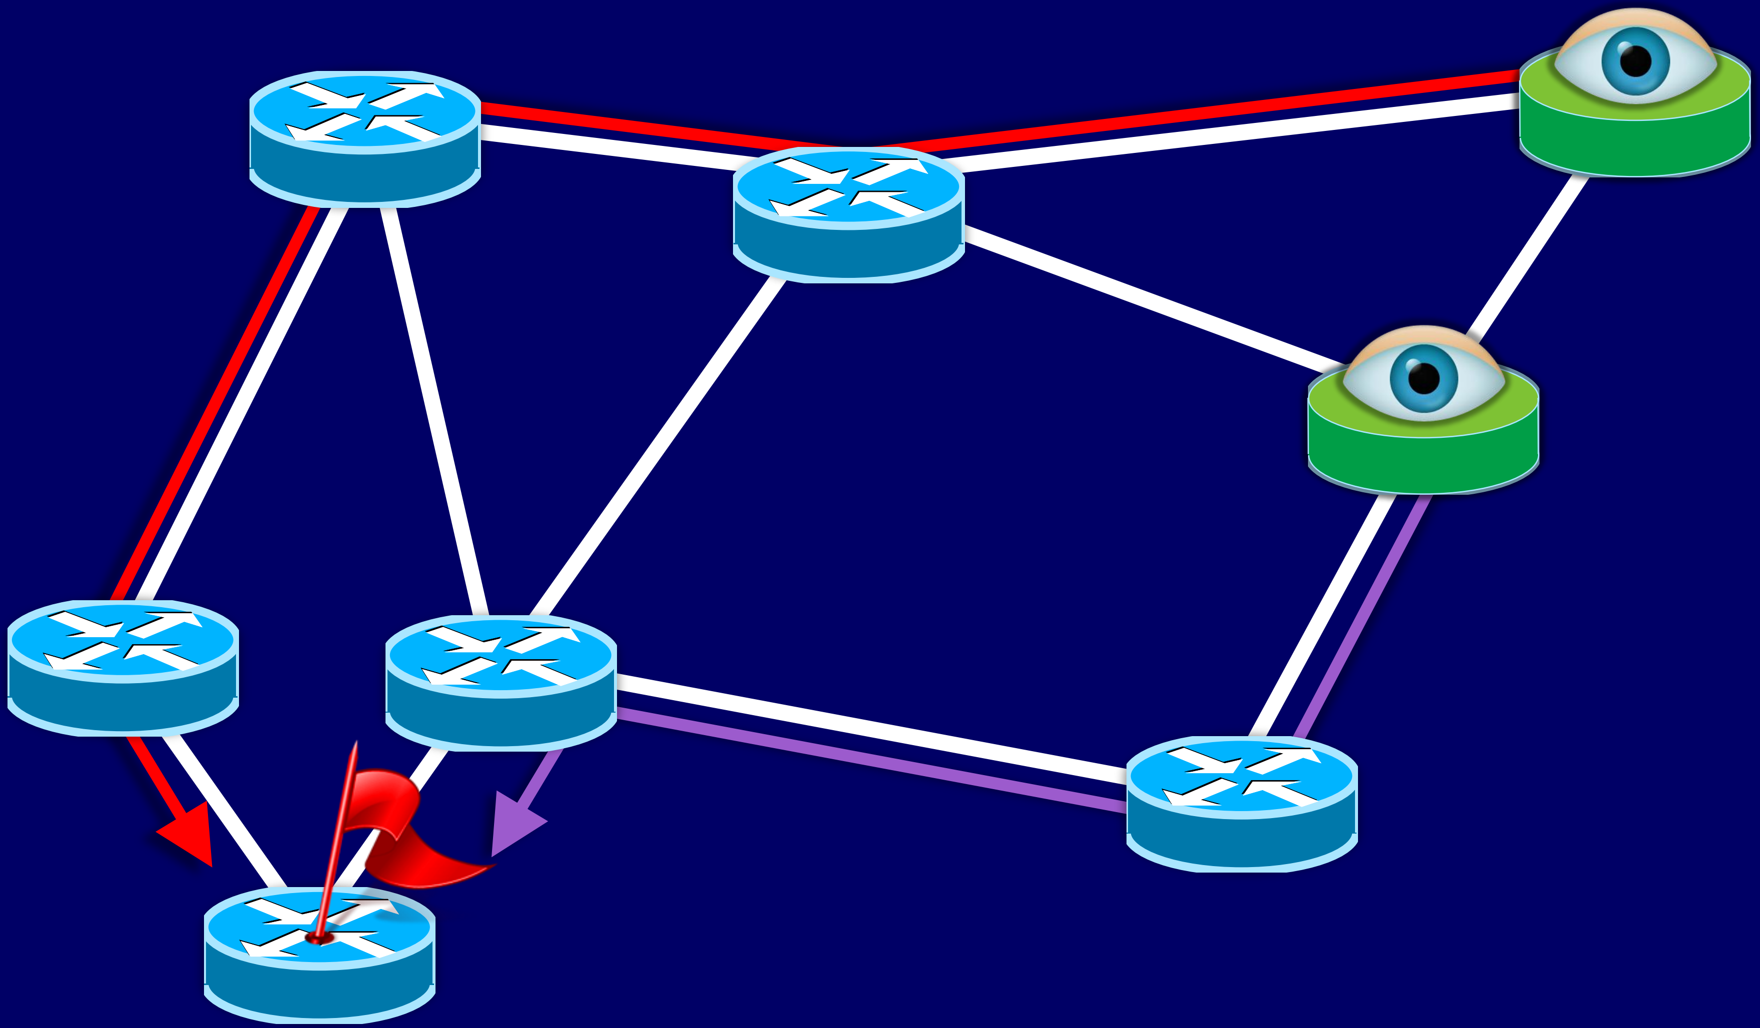

BGP Routing Changes

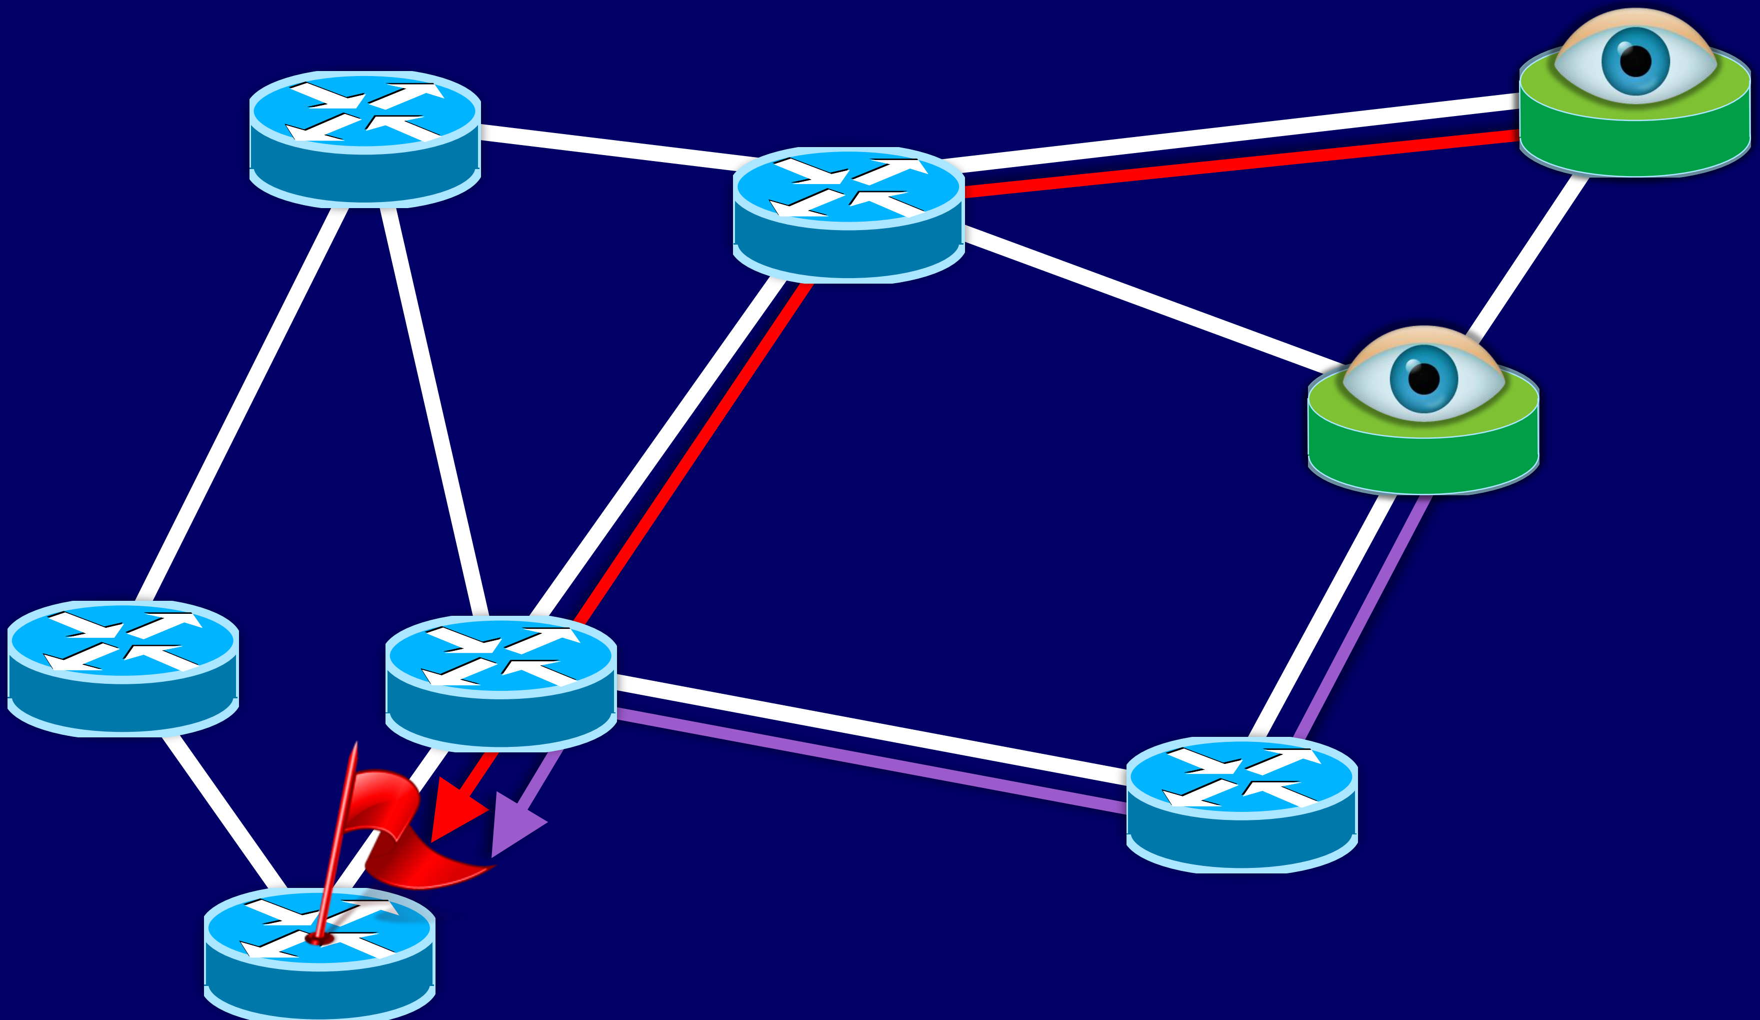

BGP Routing Changes

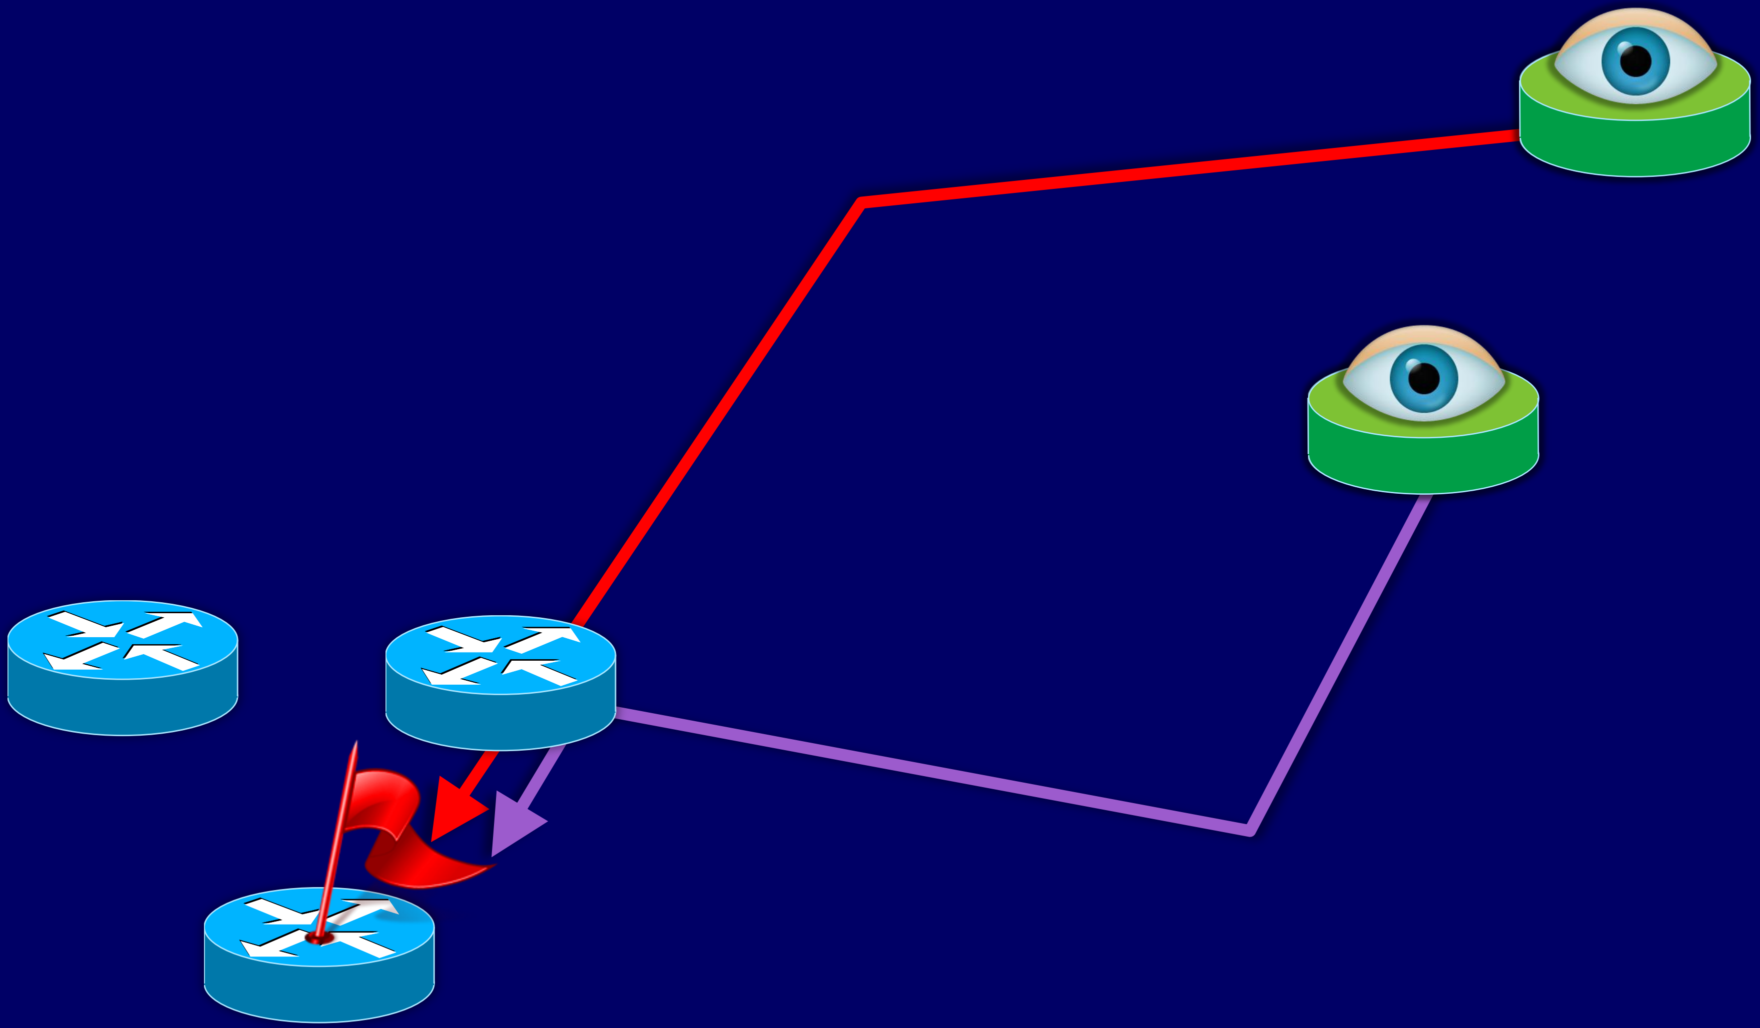

BGP Routing Changes

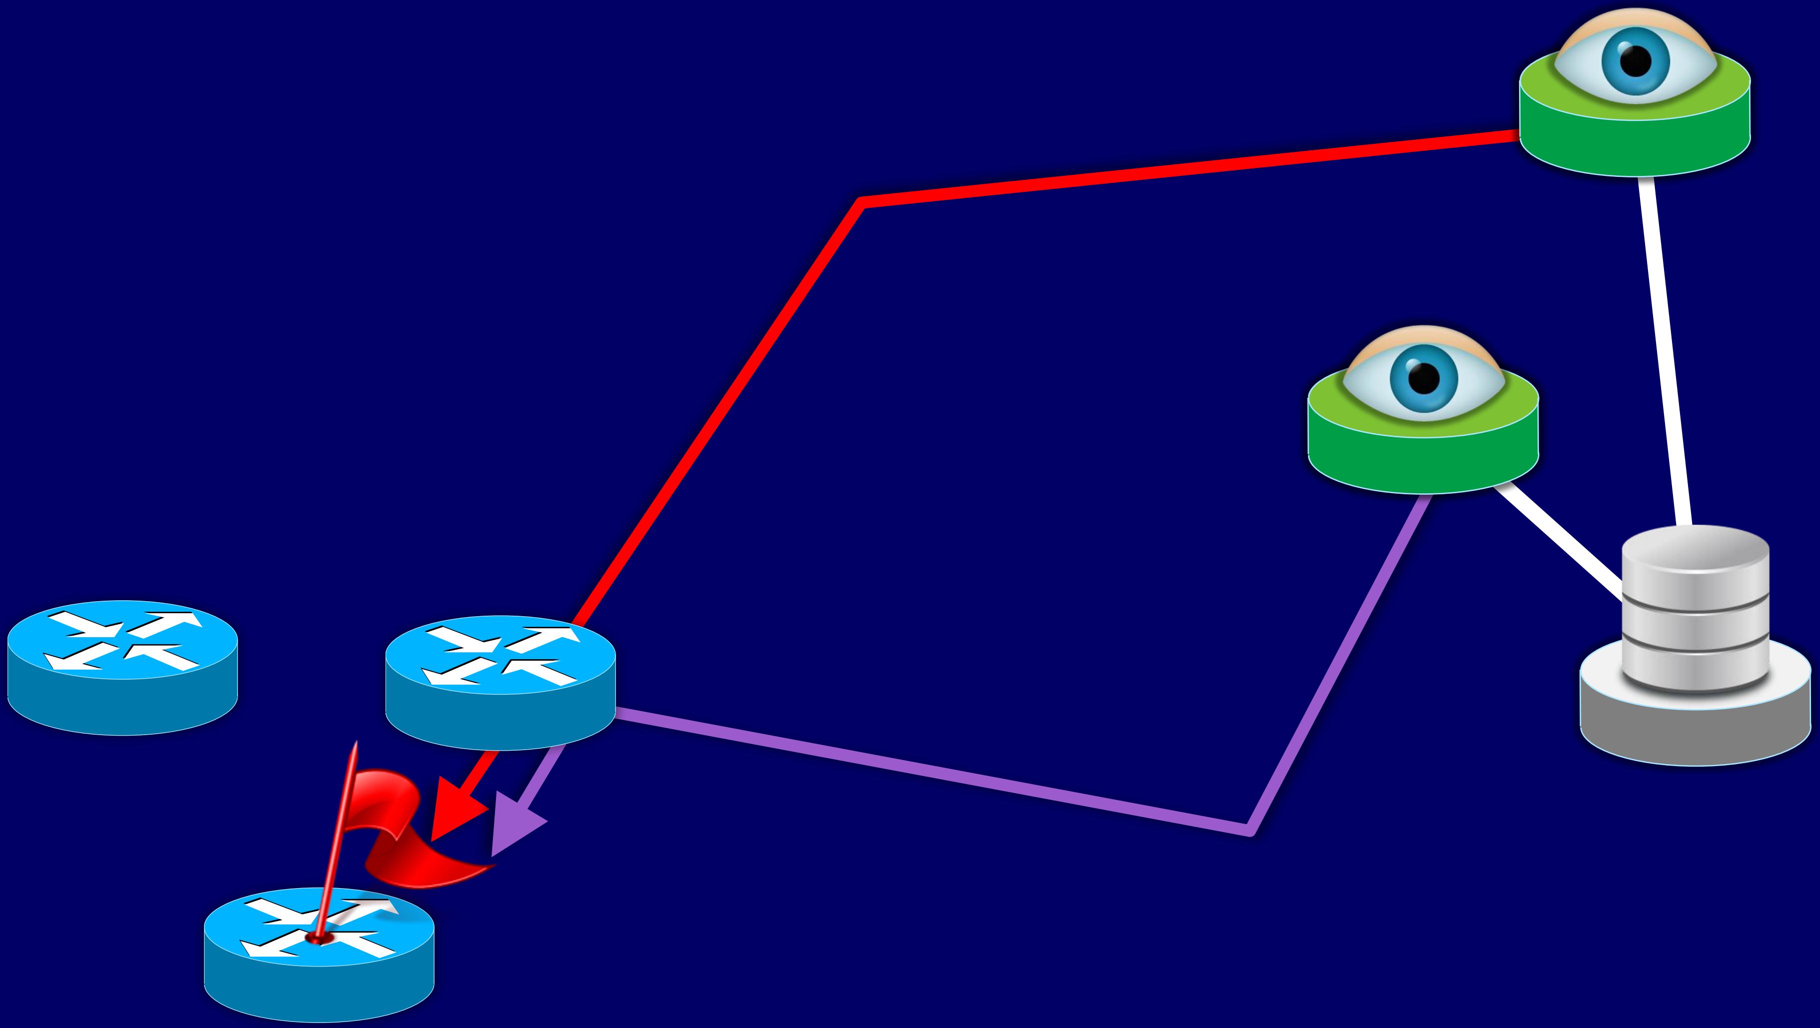

Network Delay Variations

Network Delay Variations

Network Delay Variations

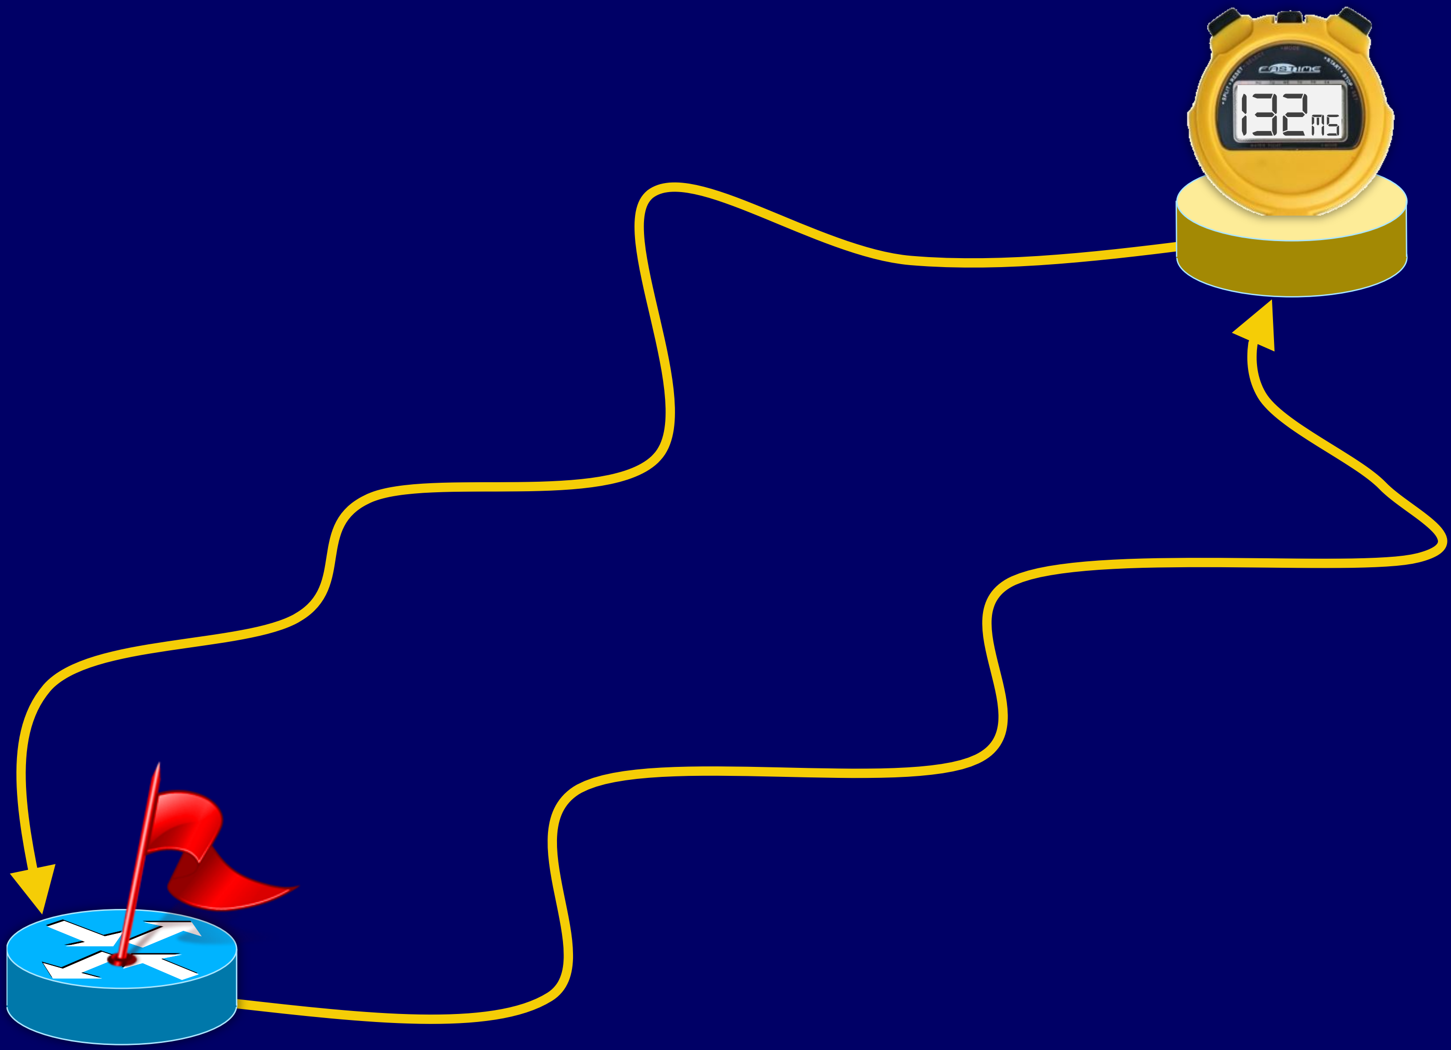

Network Delay Variations

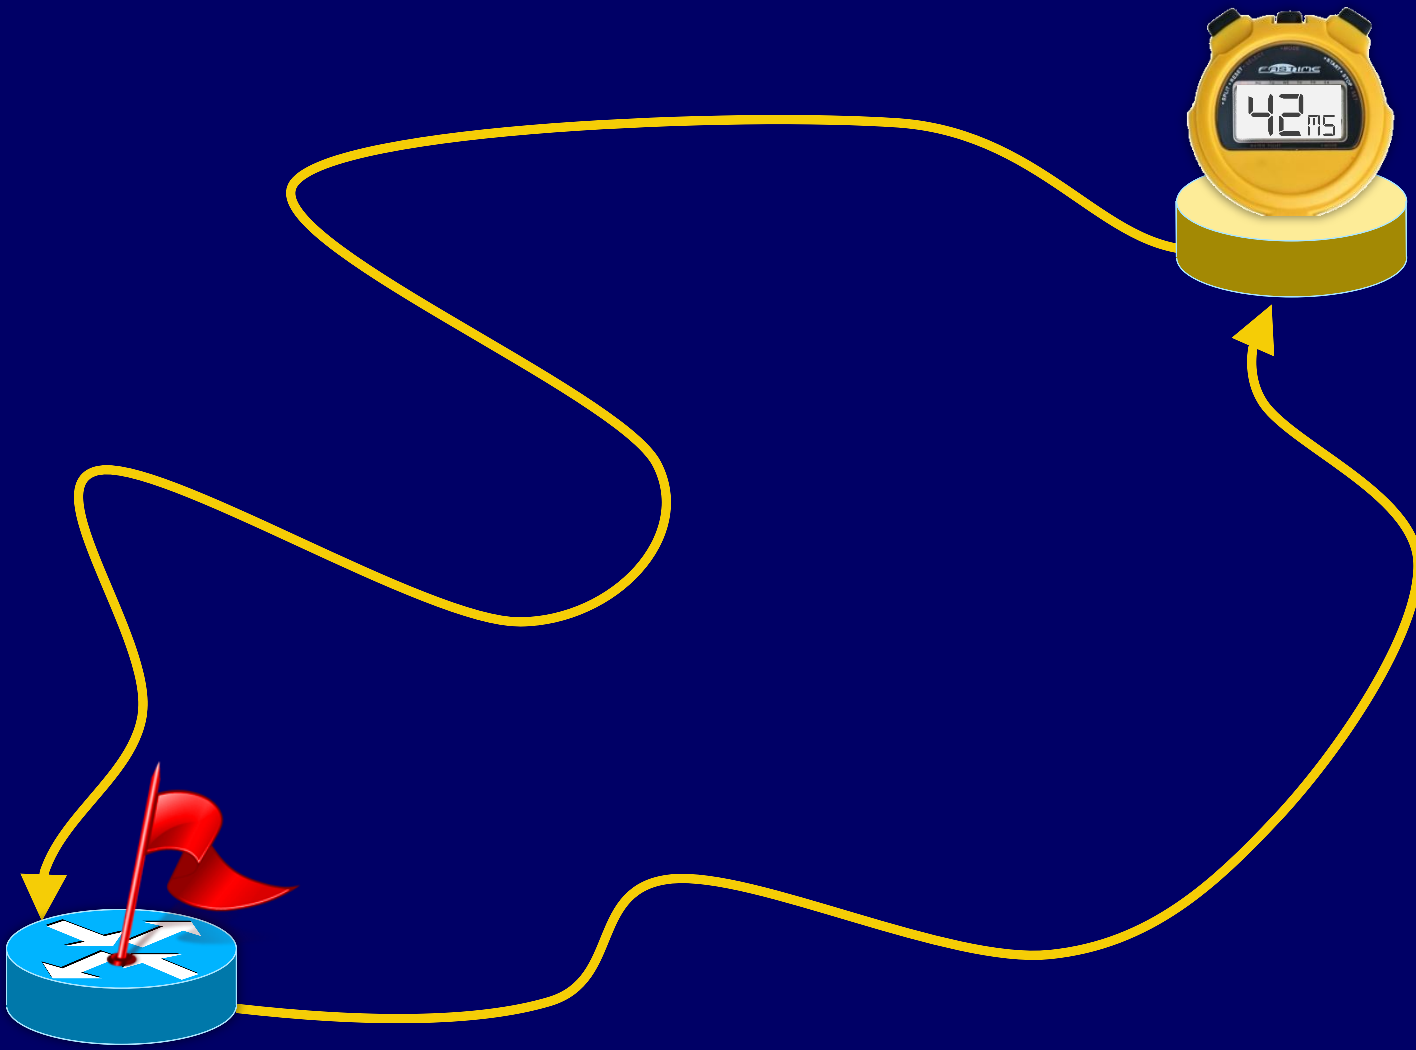

Network Delay Variations

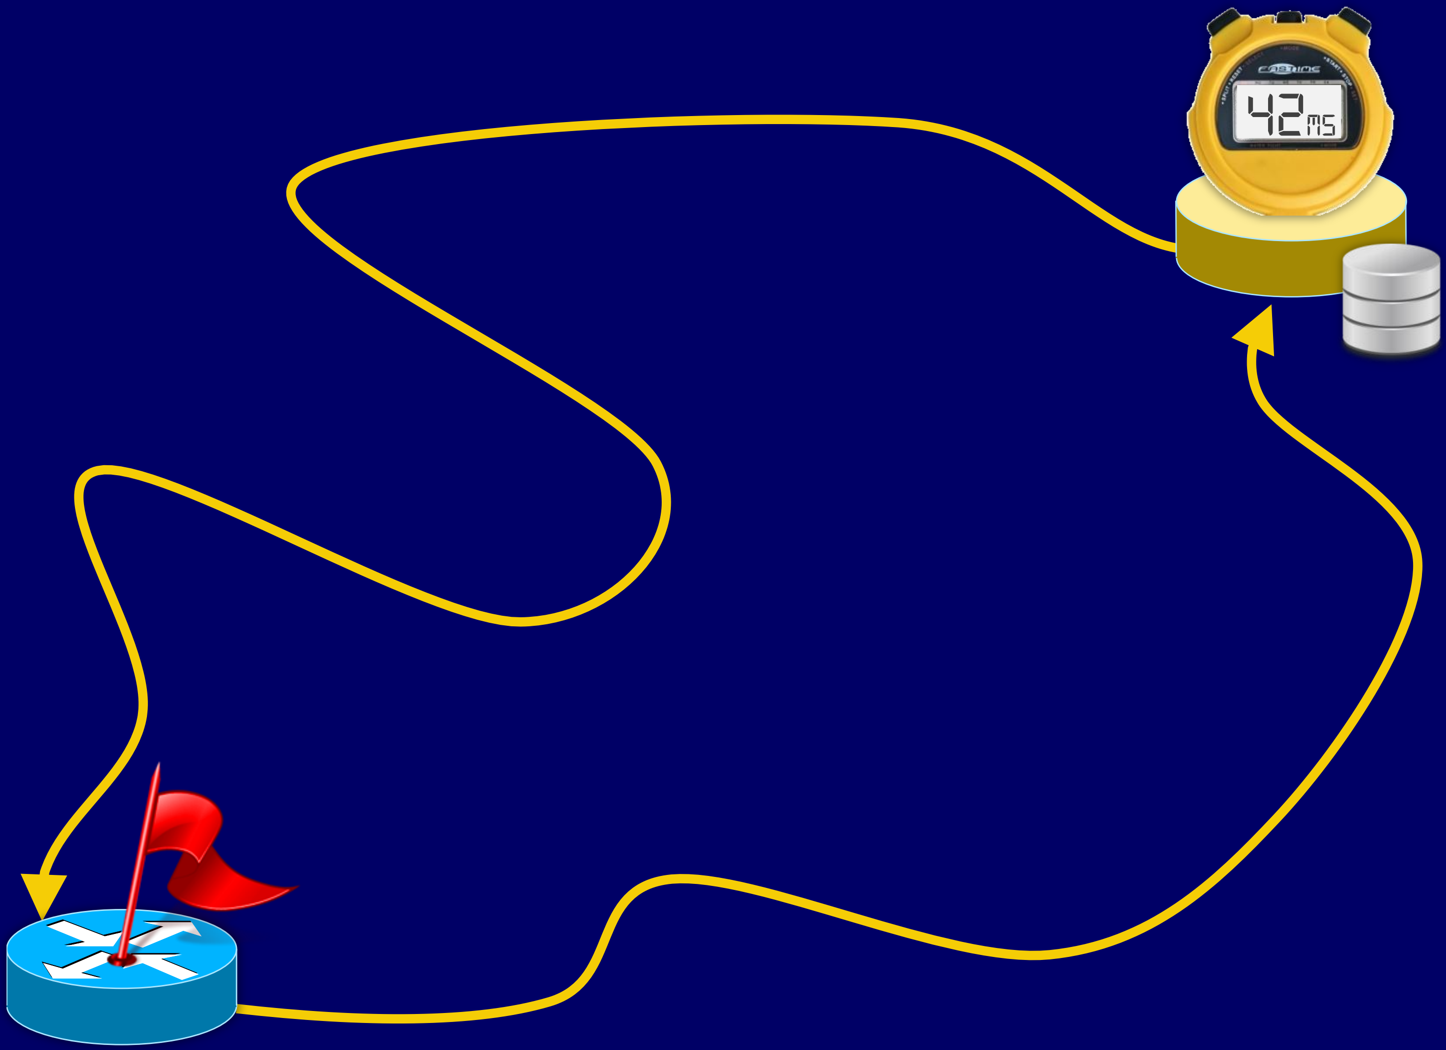

ROMA TRE UNIVERSITY
Department of Engineering

BGP Routing Changes
Network Delay Variations

Massimo Rimondini
Claudio Squarcella
Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

ROMA TRE UNIVERSITY
Department of Engineering

Impact
BGP Routing Changes
Network Delay Variations

Massimo Rimondini
Claudio Squarcella
Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

“Impact”

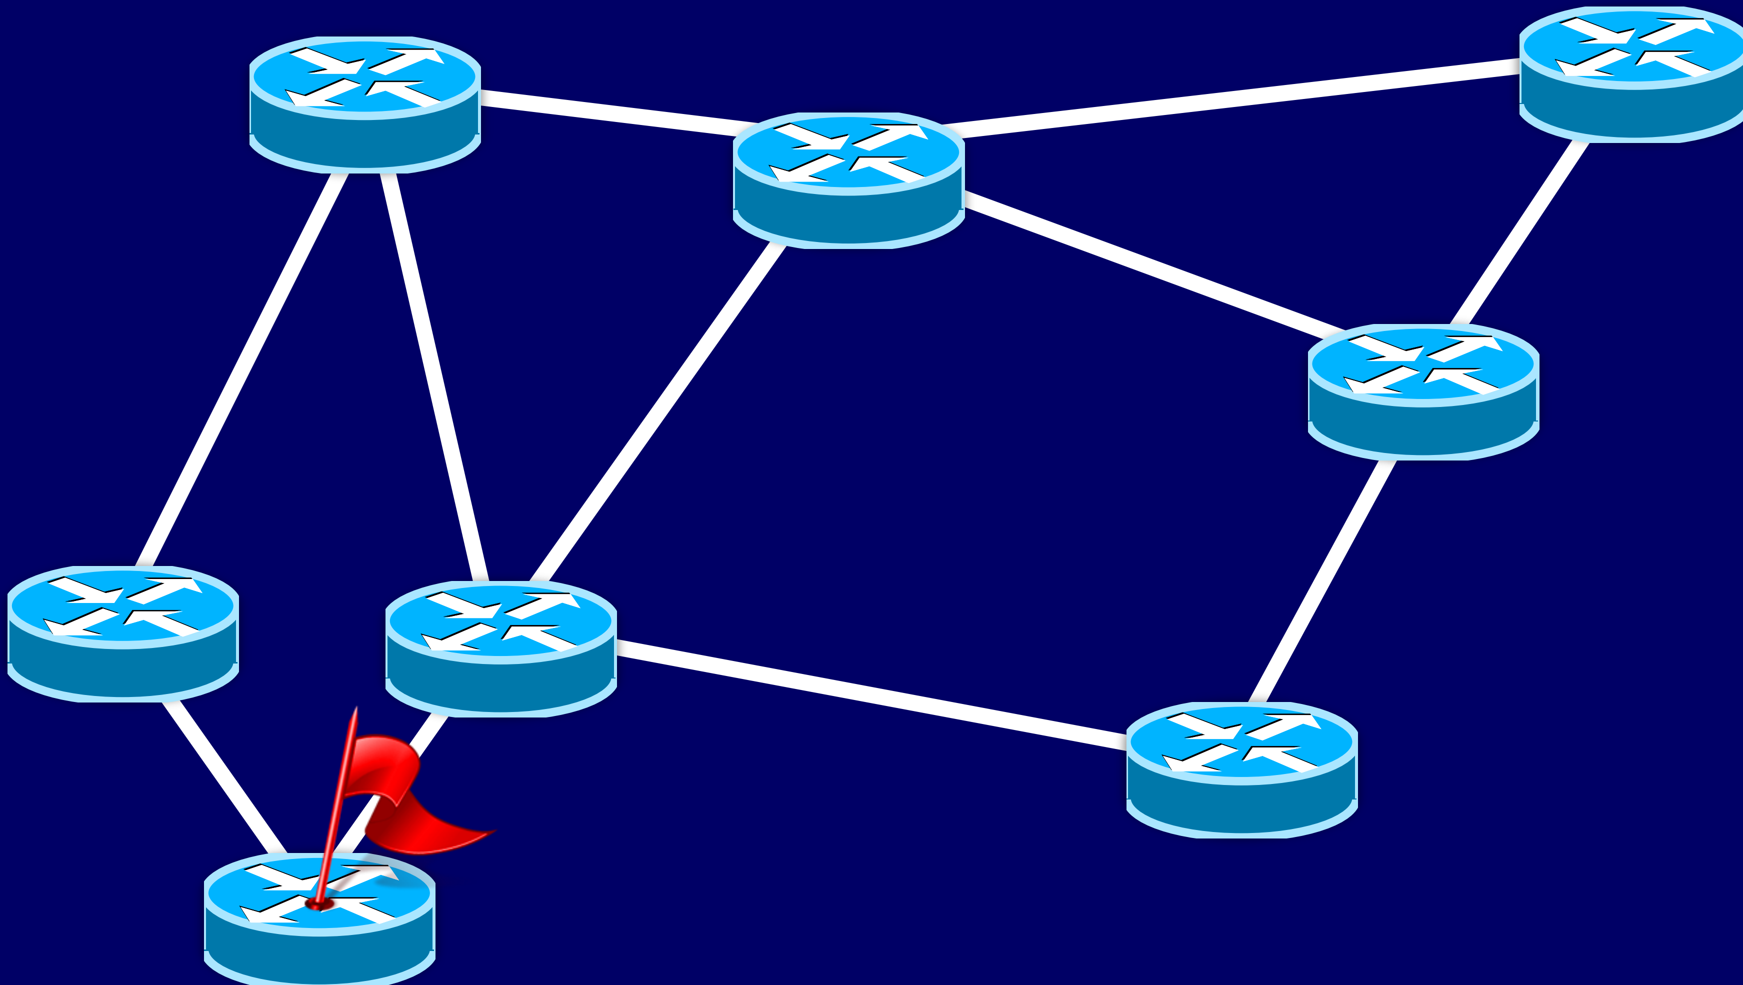

“Impact”

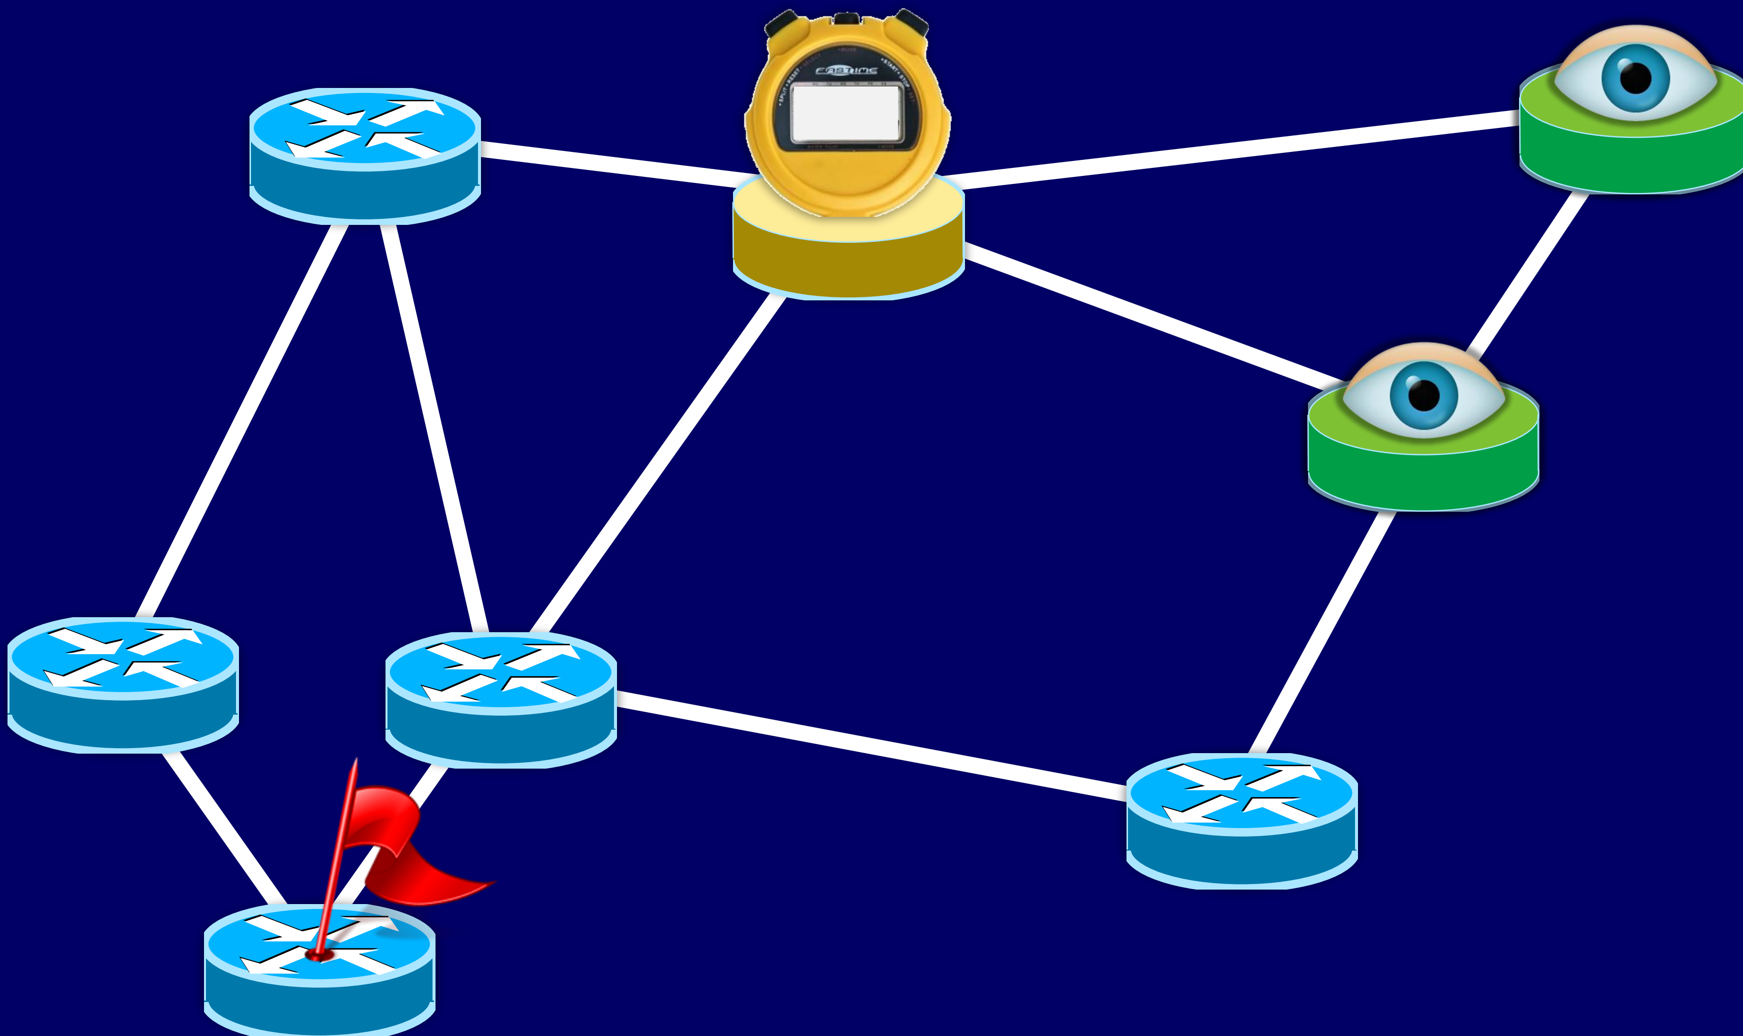

“Impact”

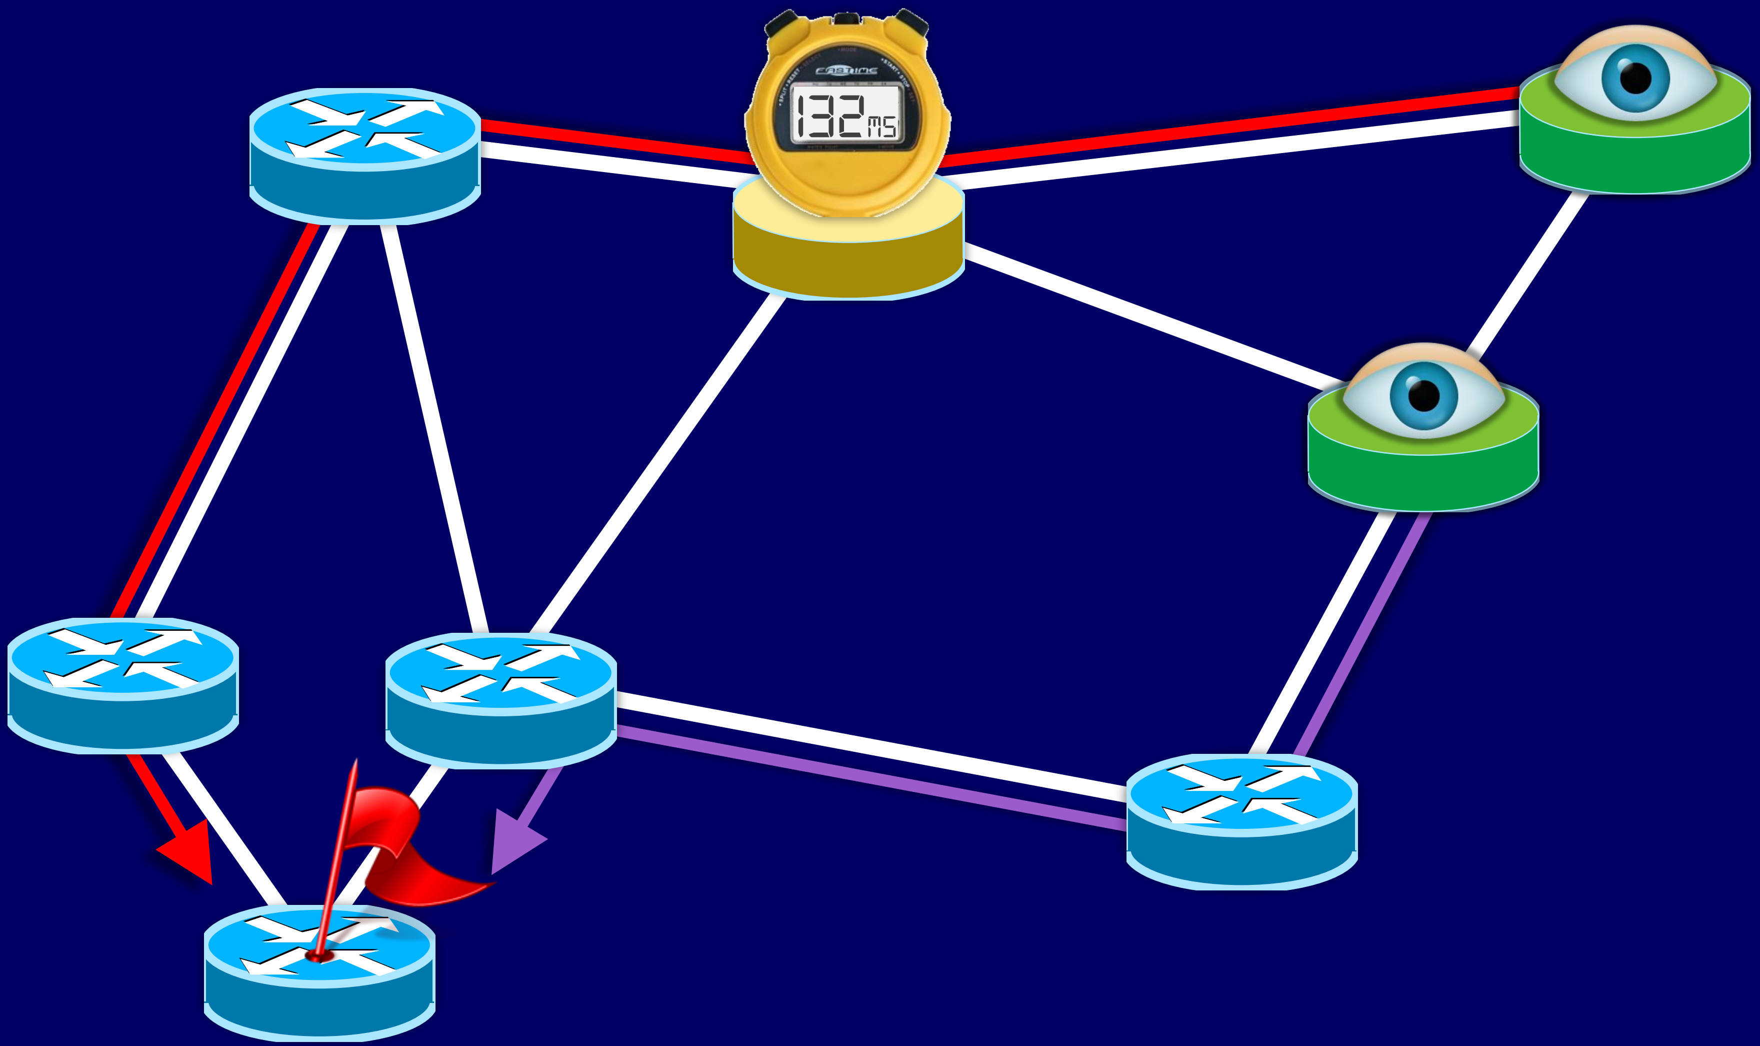

"Impact"

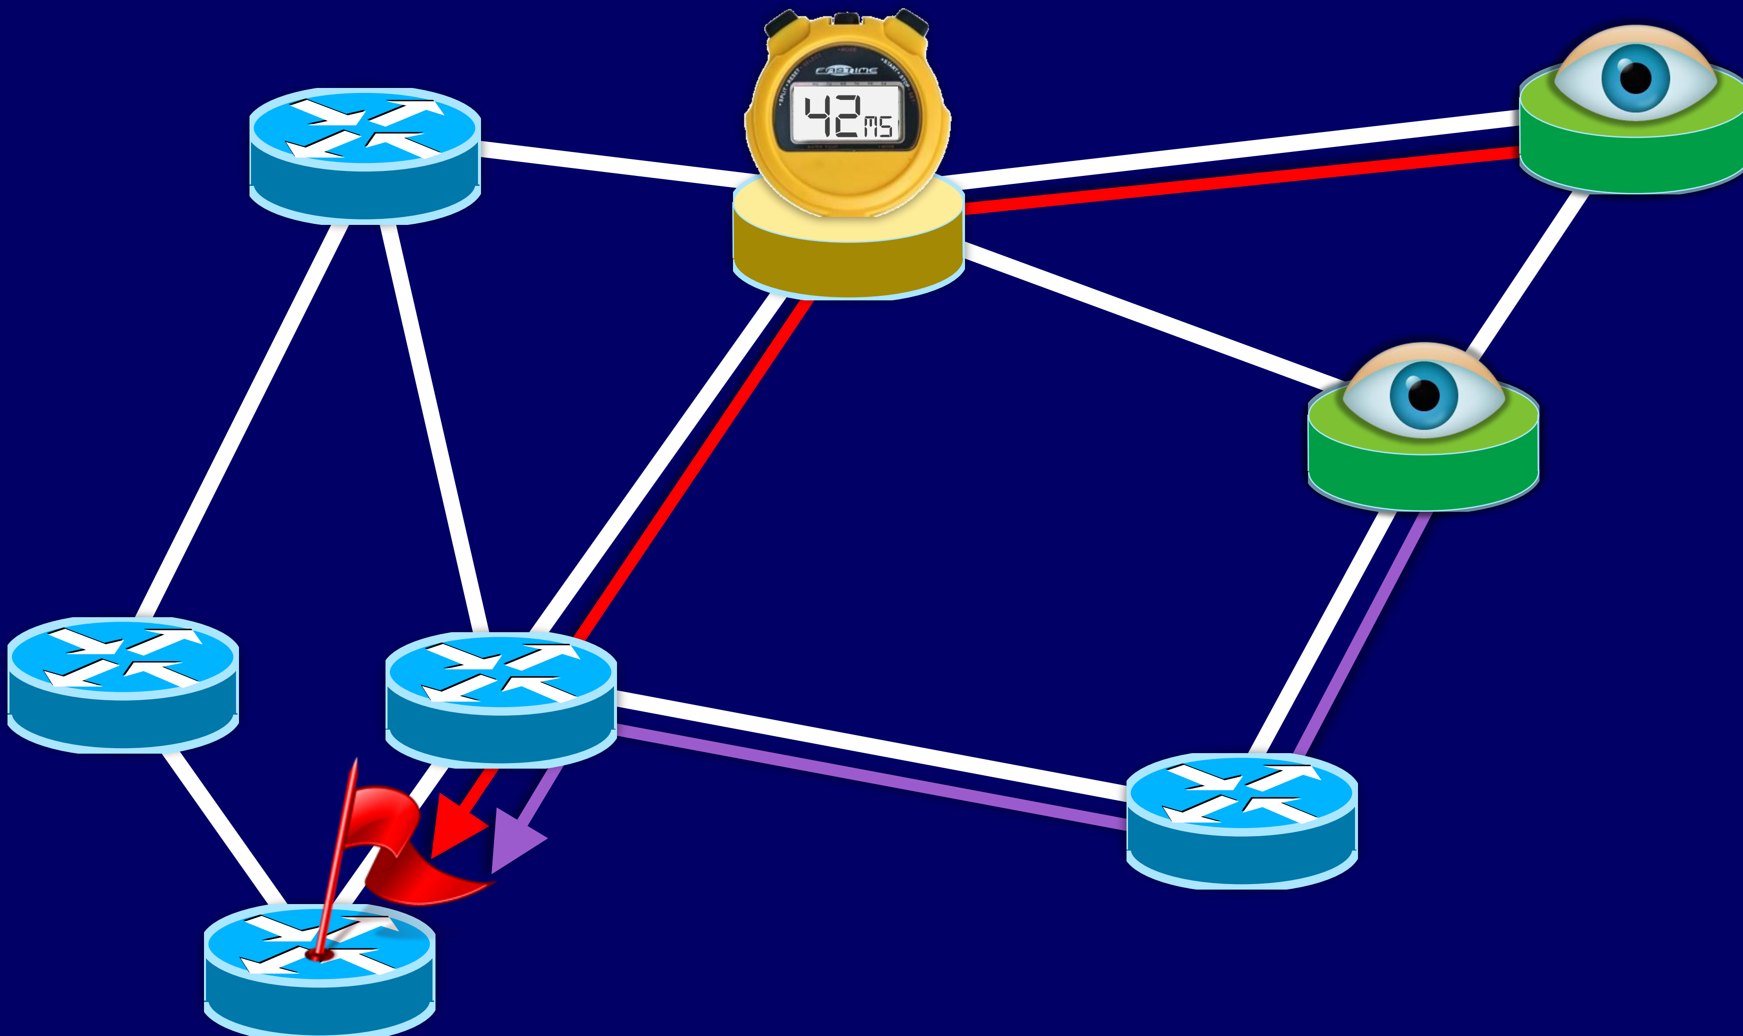

Motivations

Motivations

Motivations

SLAs

Motivations

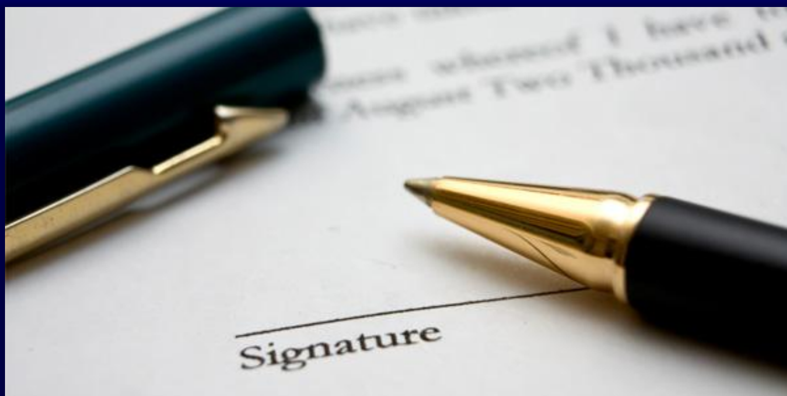

SLAs

QoE

Motivations

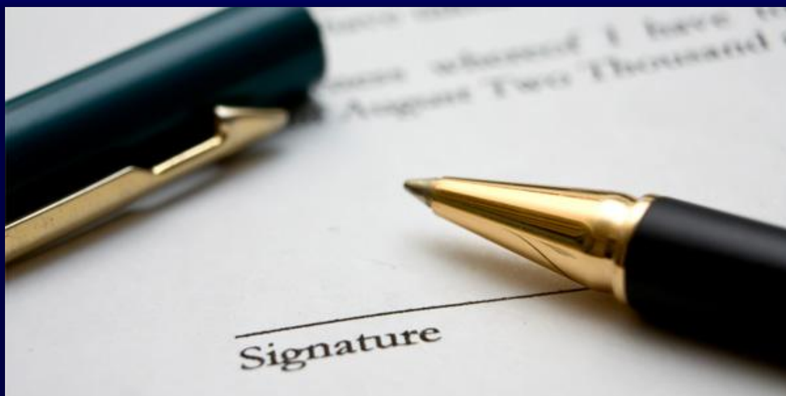

SLAs

$

QoE

Motivations

Motivations

Motivations

Motivations

Motivations

Motivations

Motivations

VS

Motivations

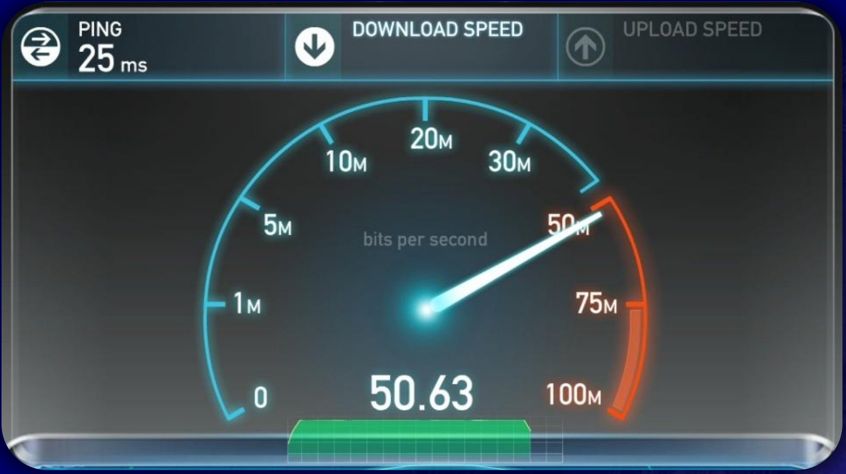

Motivations

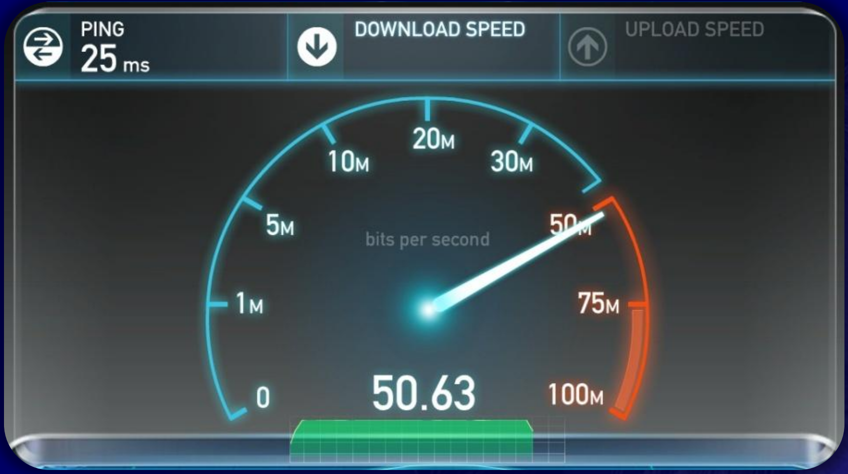

Motivations

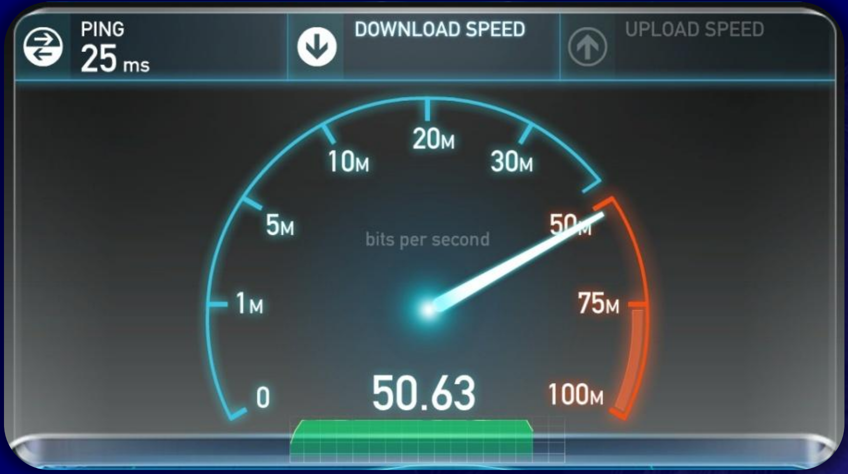

Motivations (and Applications)

Motivations (and Applications)

Motivations
(and Applications)

Motivations (and Applications)

Motivations (and Applications)

Motivations
(and Applications)

Oh, no! I'm sensitive to delay!

Motivations
(and Applications)

Oh, no! I'm sensitive to delay!

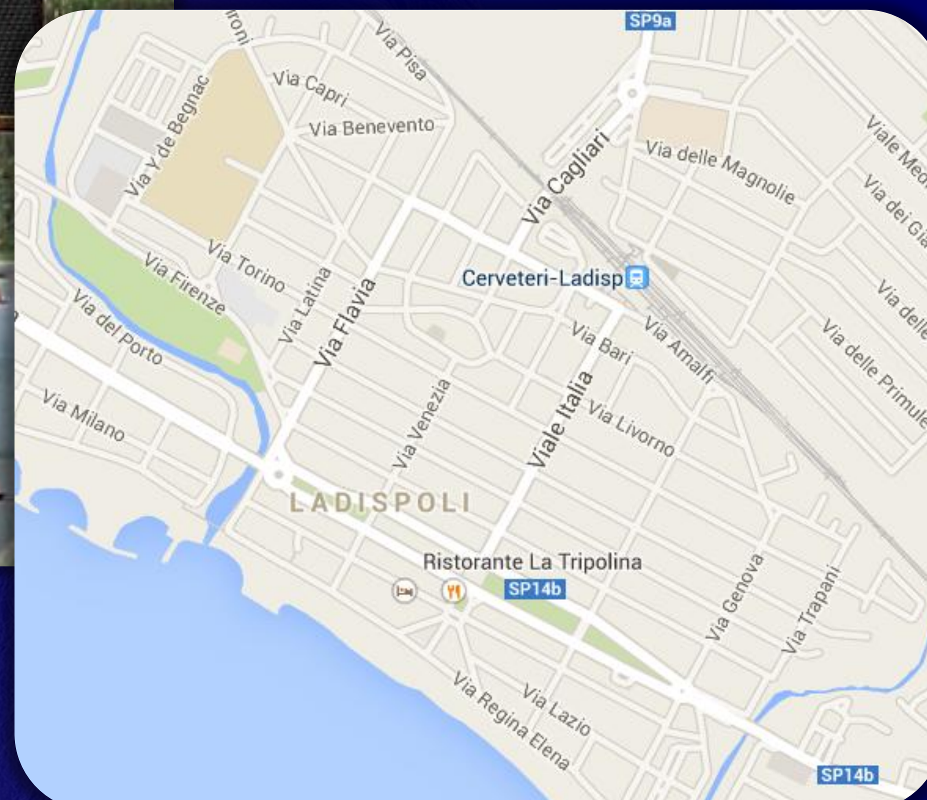

Motivations (and Applications)

Oh, no! I'm sensitive to delay!

Motivations (and Applications)

Motivations
(and Applications)

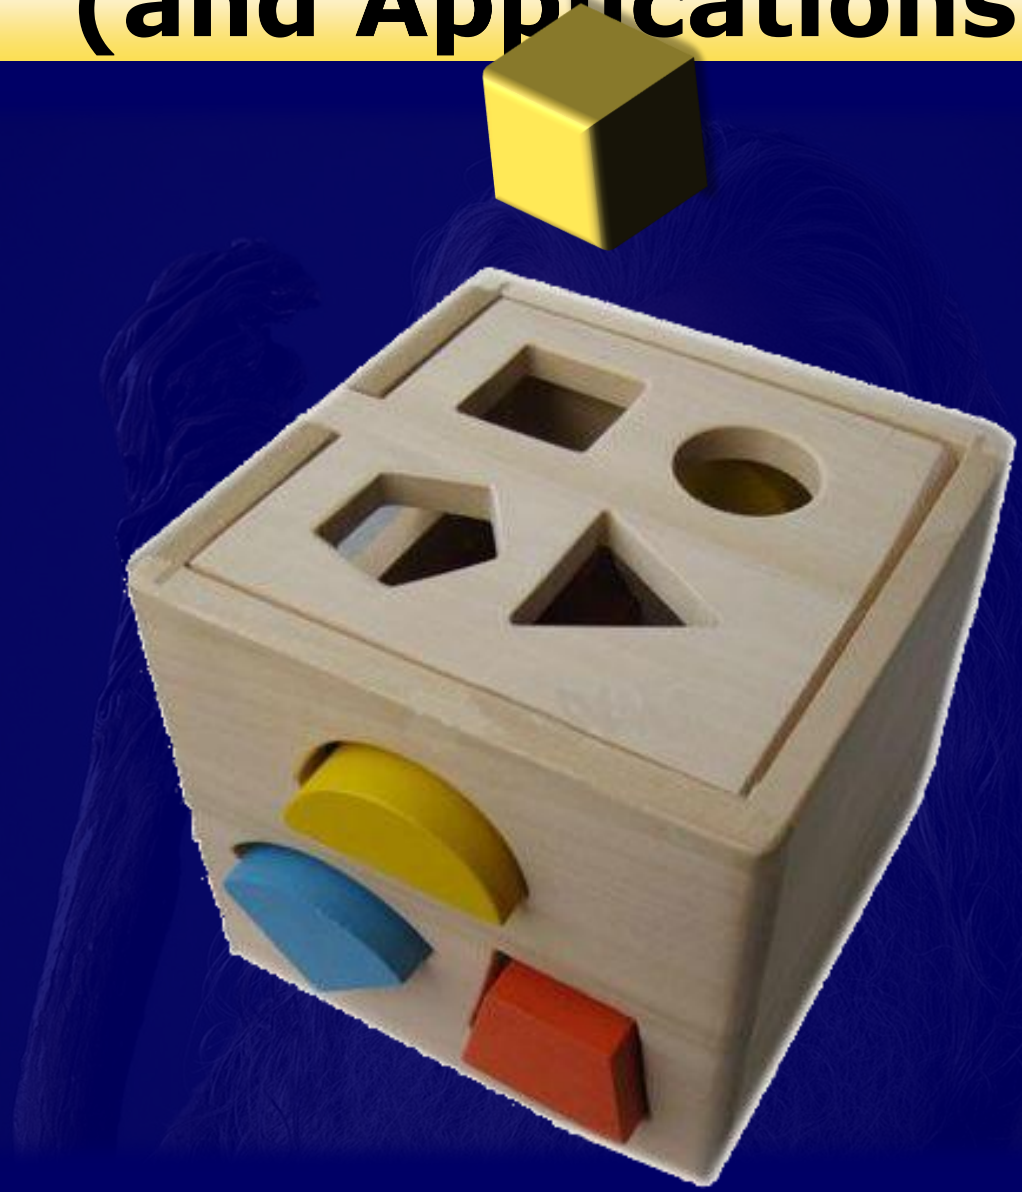

Motivations (and Applications)

Motivations (and Applications)

ROMA TRE UNIVERSITY
Department of Engineering

Impact
BGP Routing Changes
Network Delay Variations

Massimo Rimondini
Claudio Squarcella
Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

ROMA TRE UNIVERSITY
Department of Engineering

Automated Investigation of the Impact BGP Routing Changes Network Delay Variations

Massimo Rimondini
Claudio Squarcella
Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

ROMA TRE UNIVERSITY
Department of Engineering

Towards an Automated Investigation of the Impact of BGP Routing Changes on Network Delay Variations

Massimo Rimondini

Claudio Squarcella

Giuseppe Di Battista

Passive and Active Measurement Conference (PAM 2014)
March 11th, 2014

Our Contributions

Our Contributions

1. Methodology

Our Contributions

1. Methodology

Automated Investigation of the Impact of BGP Routing Changes on Network Delay Variations

Our Contributions

1. Methodology
   - Determines if a routing change caused a significant RTT variation
   - Statistical methods

Our Contributions

1. Methodology
   - Determines if a routing change caused a significant RTT variation
   - Statistical methods
2. Application
   - RIPE RIS + RIPE Atlas
   - Test *in the wild*

Our Contributions

1. Methodology
   - Determines if a routing change caused a significant RTT variation
   - Statistical methods
2. Application
   - RIPE RIS + RIPE Atlas
   - Test *in the wild*
3. A-posteriori statistics

State of the Art

State of the Art

- Pucha, H., Zhang, Y., Mao, Z., Hu, Y.: **Understanding network delay changes caused by routing events**. Proc. SIGMETRICS 2007

- Routing changes cause delay variations
  - as opposed to congestion
  - average delays mostly impacted by interdomain changes

State of the Art

State of the Art

- Chuah, C.N., Bhattacharyya, S., Diot, C.: **Measuring I-BGP updates and their impact on traffic**. Tech. Rep. TR02-ATL-051099, Sprint ATL 2002
- Wang, F., Mao, Z.M., Wang, J., Gao, L., Bush, R.: **A Measurement study on the impact of routing events on end-to-end internet path performance**. SIGCOMM Comput. Commun. Rev. 36(4), 375–386, 2006
- Zhang, Y., Mao, Z., Wang, J.: **A framework for measuring and predicting the impact of routing changes**. Proc. INFOCOM 2007

- Routing convergence can cause performance degradations

State of the Art

- Chuah, C.N., Bhattacharyya, S., Diot, C.: **Measuring I-BGP updates and their impact on traffic**. Tech. Rep. TR02-ATL-051099, Sprint ATL 2002
- Wang, F., Mao, Z.M., Wang, J., Gao, L., Bush, R.: **A Measurement study on the impact of routing events on end-to-end internet path performance**. SIGCOMM Comput. Commun. Rev. 36(4), 375–386, 2006
- Zhang, Y., Mao, Z., Wang, J.: **A framework for measuring and predicting the impact of routing changes**. Proc. INFOCOM 2007

- Routing convergence can cause performance degradations

State of the Art

State of the Art

- Da Lozzo, G., Di Battista, G., Squarcella, C.: **Visual discovery of the correlation between BGP routing and round-trip delay active measurements**. Computing, 1–11, 2013

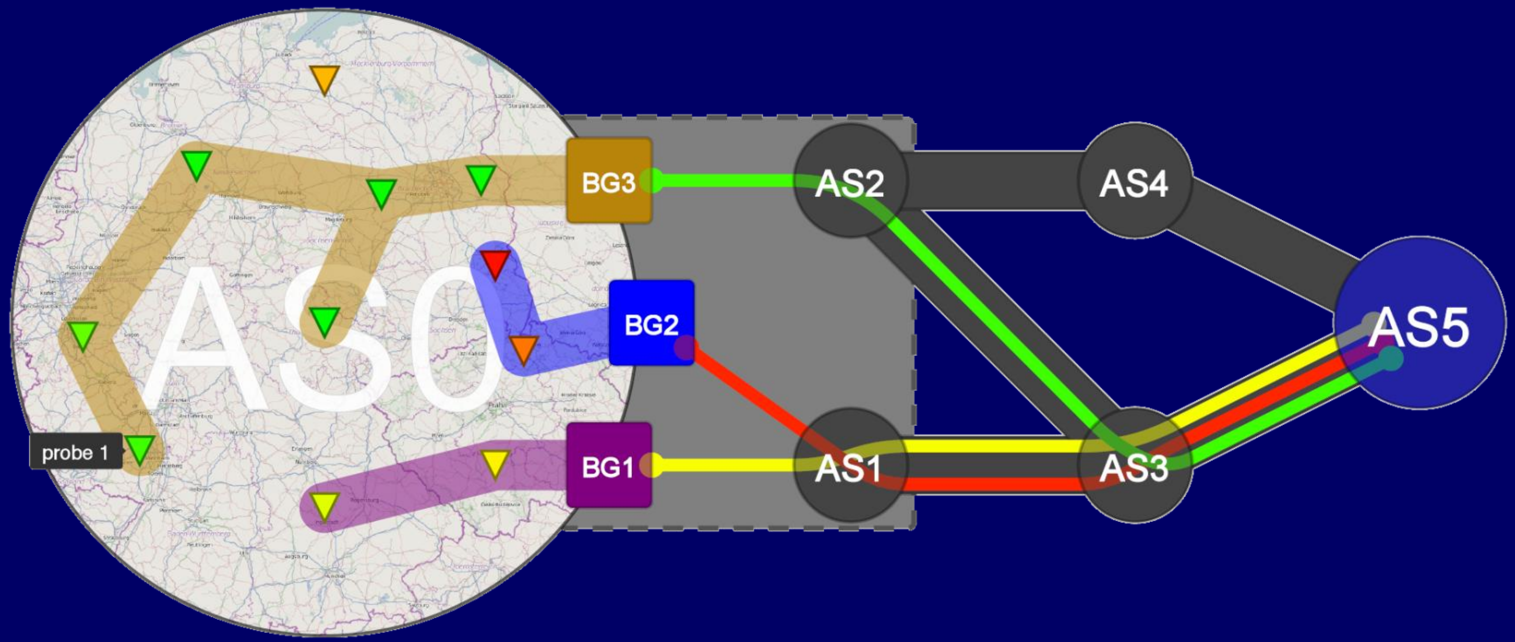

- Focus on the graphical metaphor

State of the Art

State of the Art

- Mahimkar, A., Ge, Z., Wang, J., Yates, J., Zhang, Y., Emmons, J., Huntley, B., Stockert, M.: **Rapid detection of maintenance induced changes in service performance**. Proc. CoNEXT 2011
- Mahimkar, A., Song, H., Ge, Z., Shaikh, A., Wang, J., Yates, J., Zhang, Y., Emmons, J.: **Detecting the performance impact of upgrades in large operational networks**. Proc. SIGCOMM 2010

- Identify patterns in performance changes
  - statistical rule mining
  - network configuration information

State of the Art

- Mahimkar, A., Ge, Z., Wang, J., Yates, J., Zhang, Y., Emmons, J., Huntley, B., Stockert, M.: **Rapid detection of maintenance induced changes in service performance**. Proc. CoNEXT 2011
- Mahimkar, A., Song, H., Ge, Z., Shaikh, A., Wang, J., Yates, J., Zhang, Y., Emmons, J.: **Detecting the performance impact of upgrades in large operational networks**. Proc. SIGCOMM 2010

- Identify patterns in performance changes
  - statistical rule mining
  - network configuration information

State of the Art

State of the Art

- Tsamoura, E., Gounaris, A.: **Incorporating change detection in network coordinate systems for large data transfers**. Proc. PCI 2013

- Predict network delay between host pairs
  - change detection algorithms

State of the Art

Scenario

Scenario

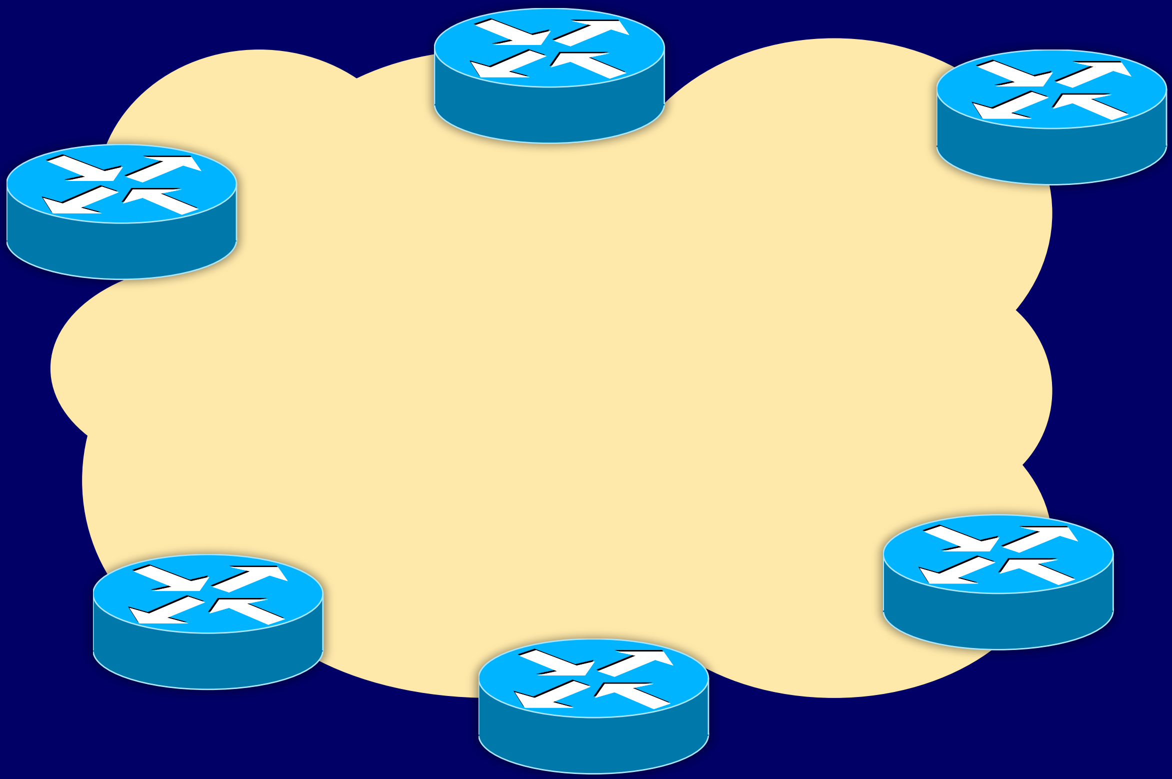

Scenario

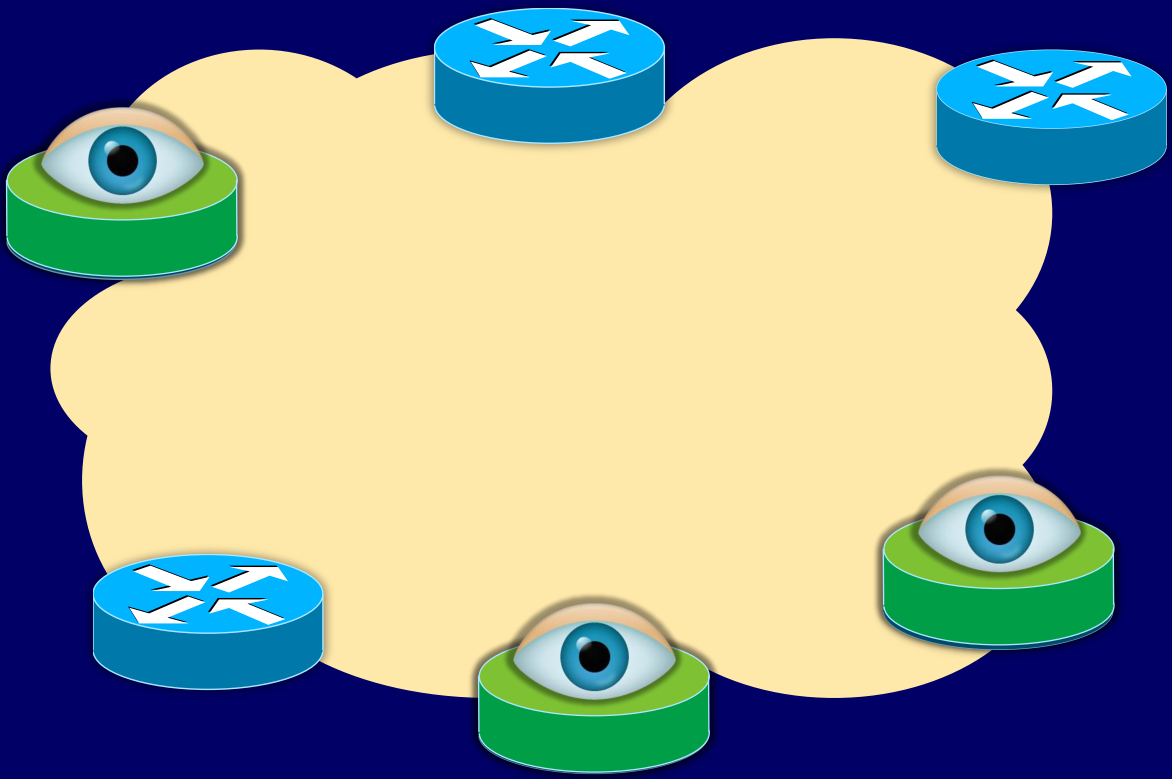

Scenario

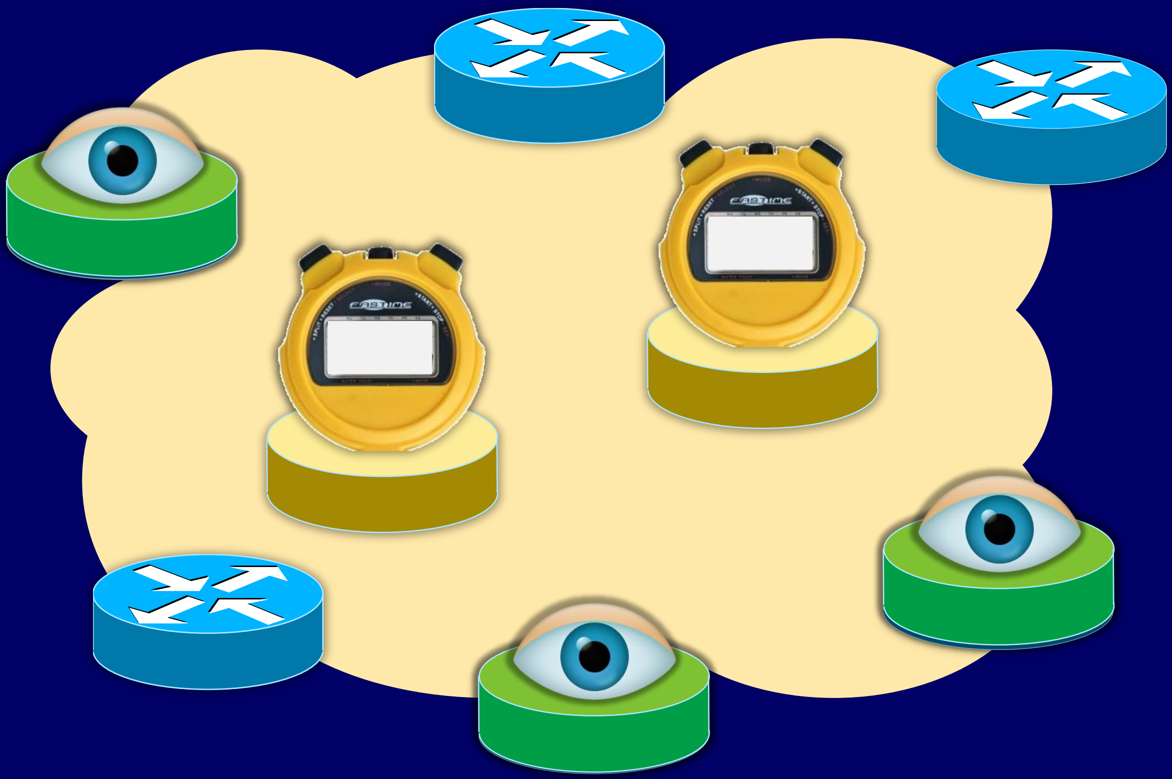

Scenario

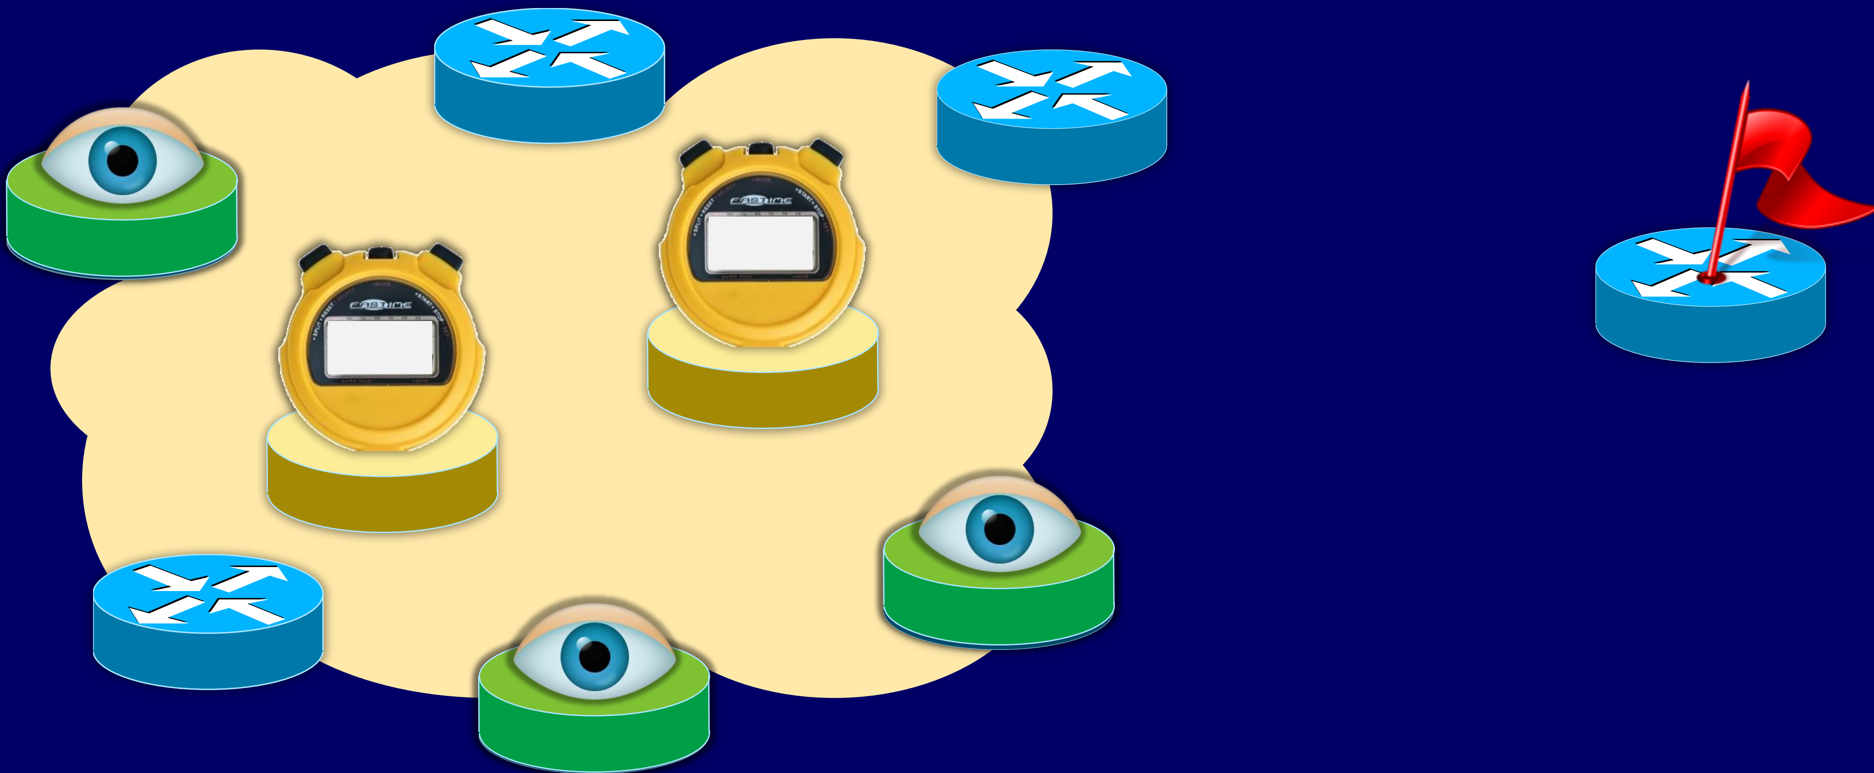

Scenario

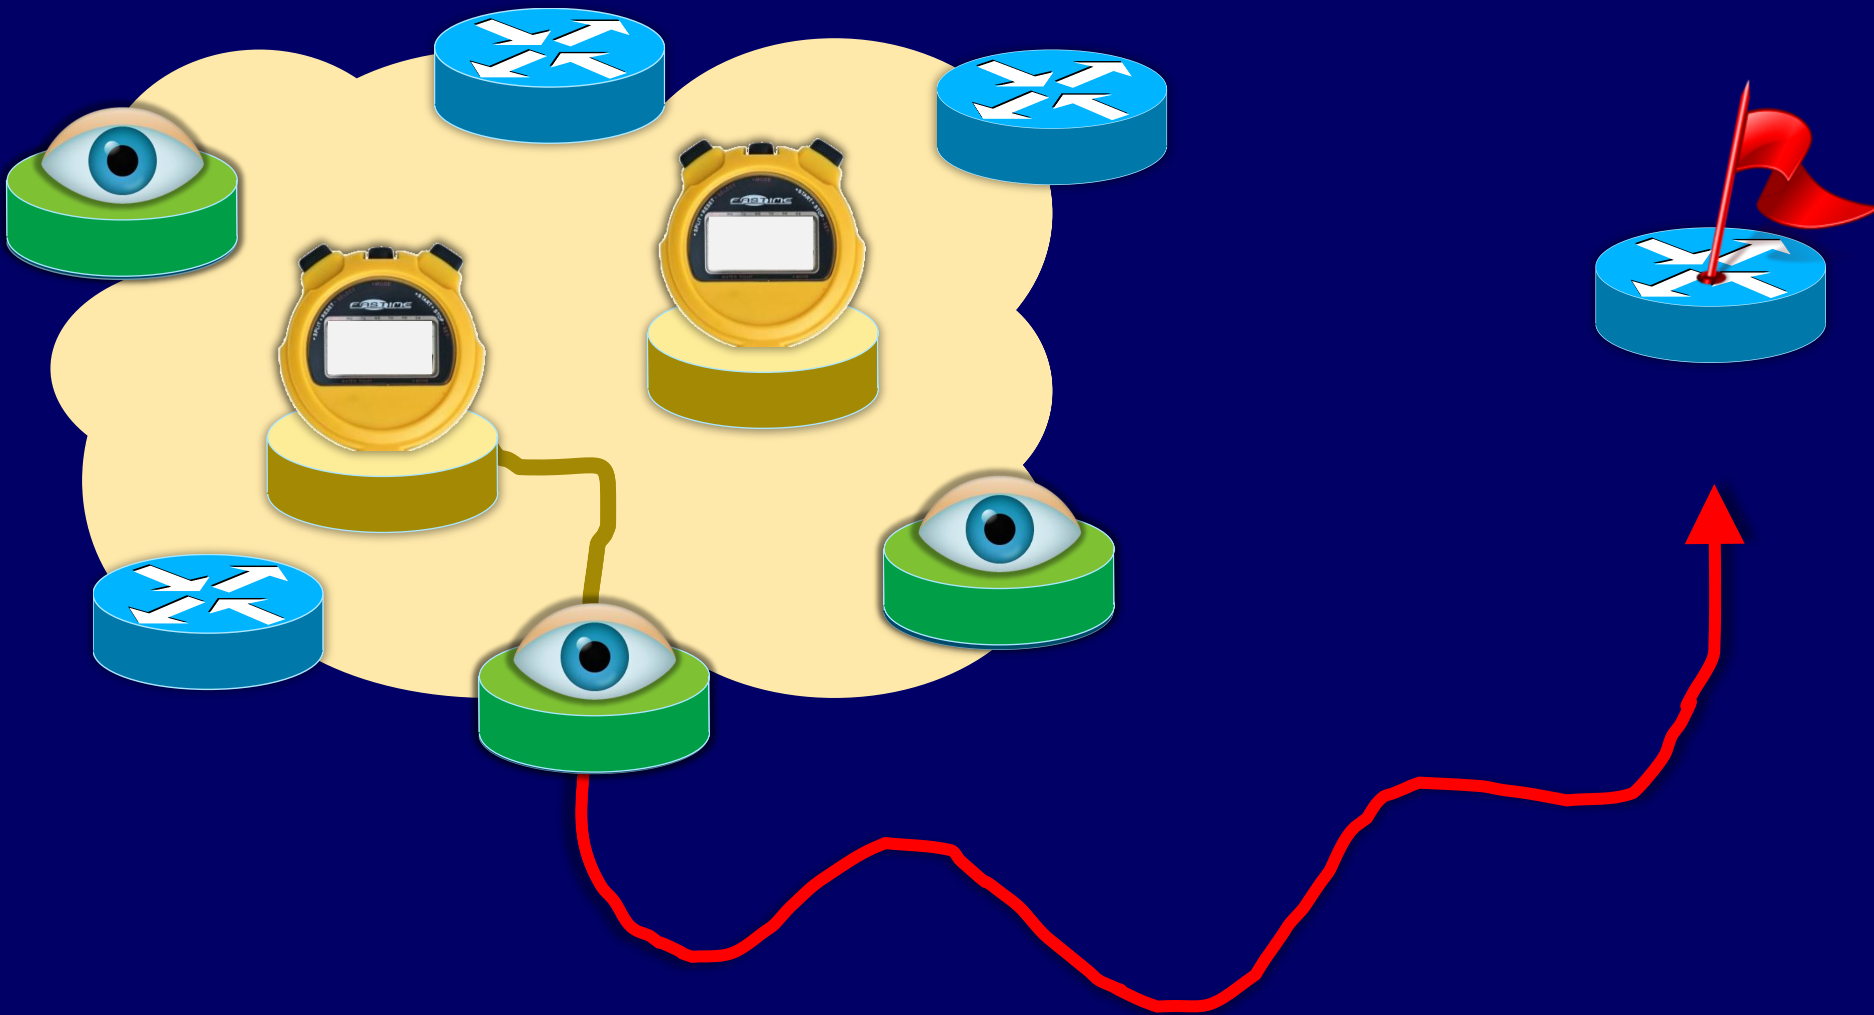

Scenario

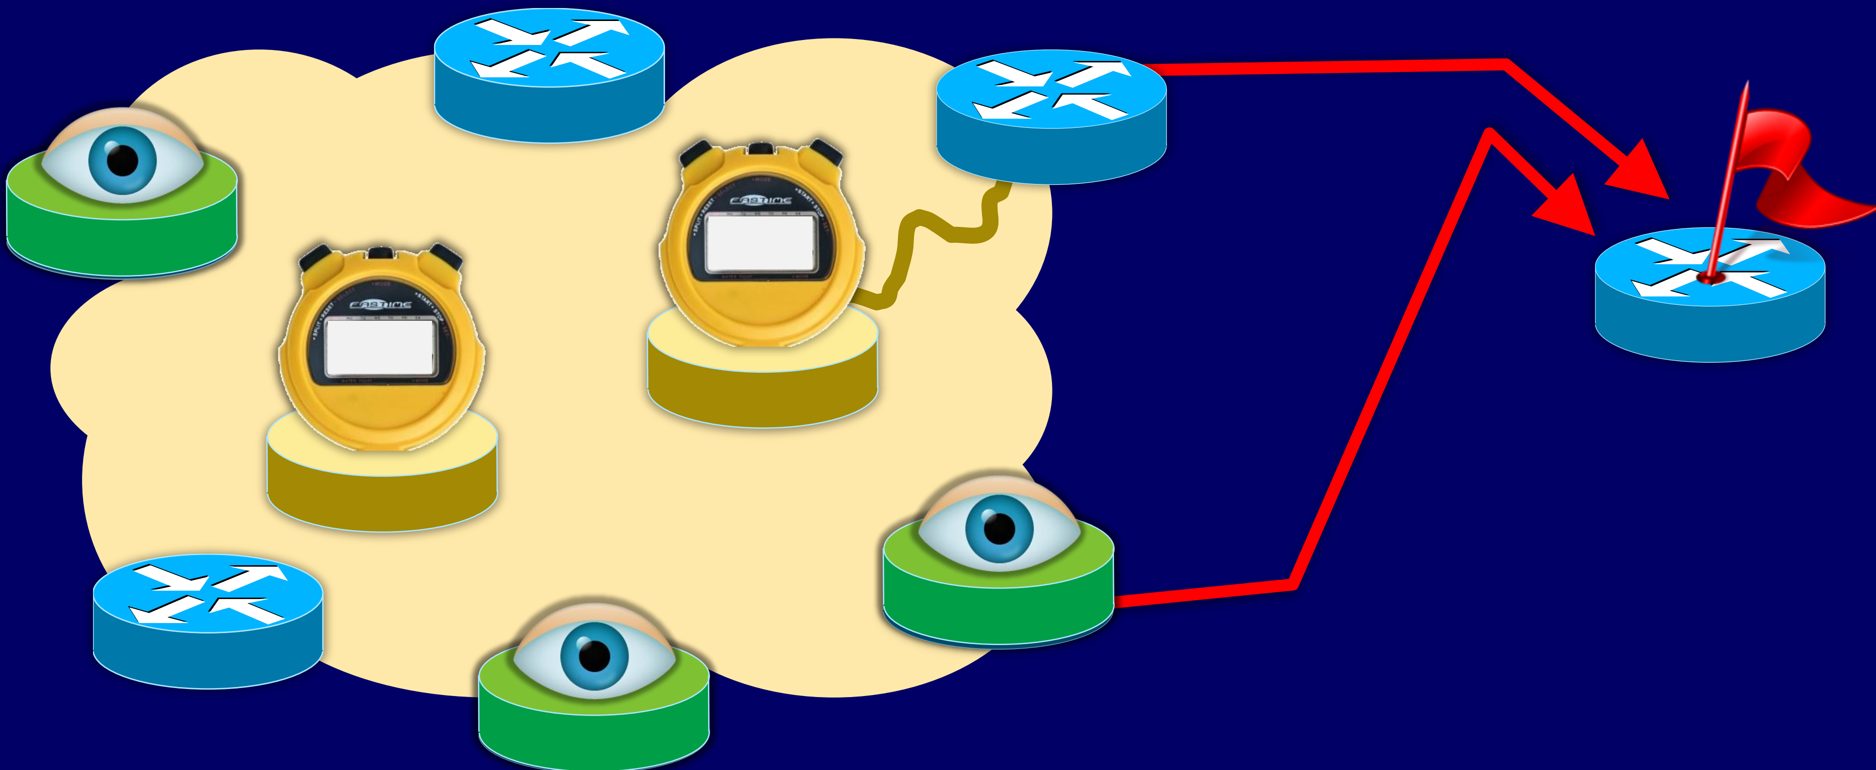

Scenario

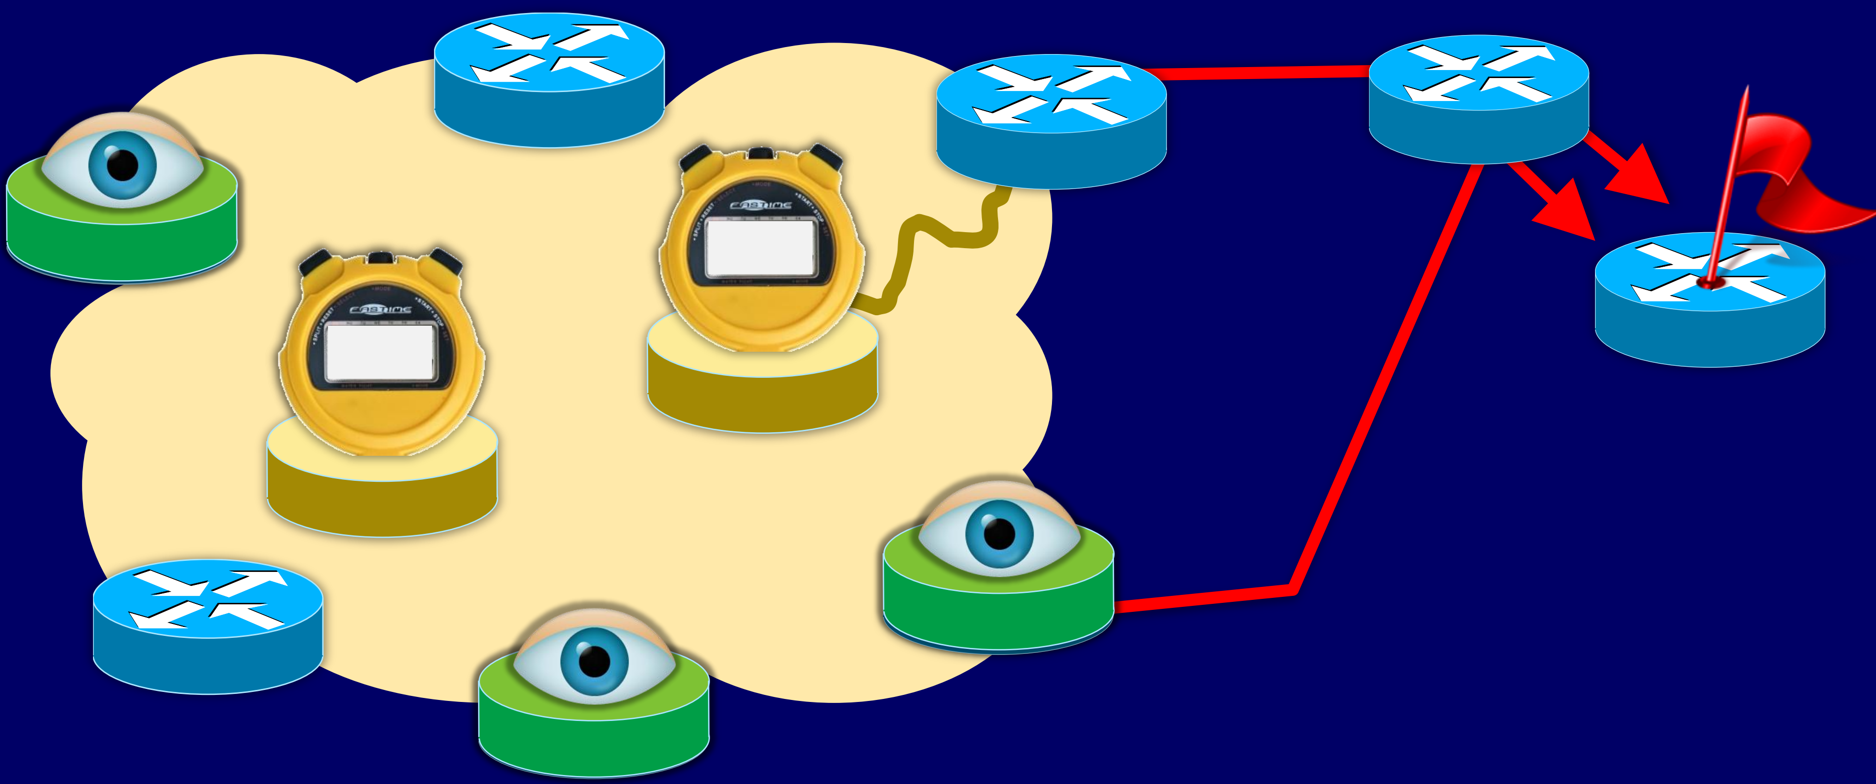

Scenario

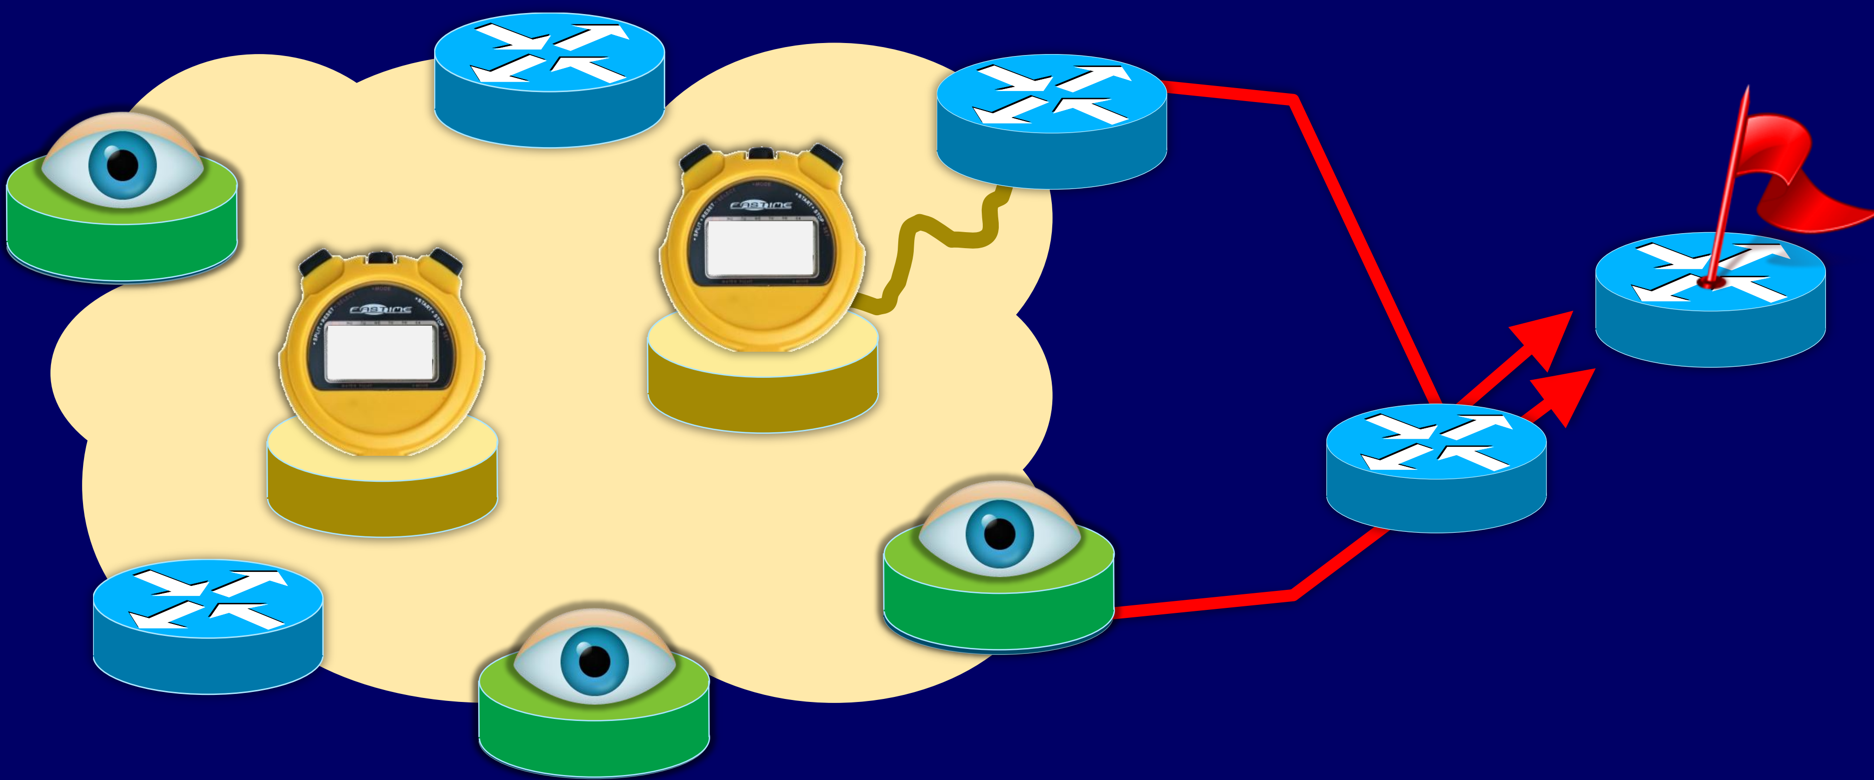

Methodology – Overview

Methodology – Overview

Raw RTT measurements

BGP updates

Methodology – Overview

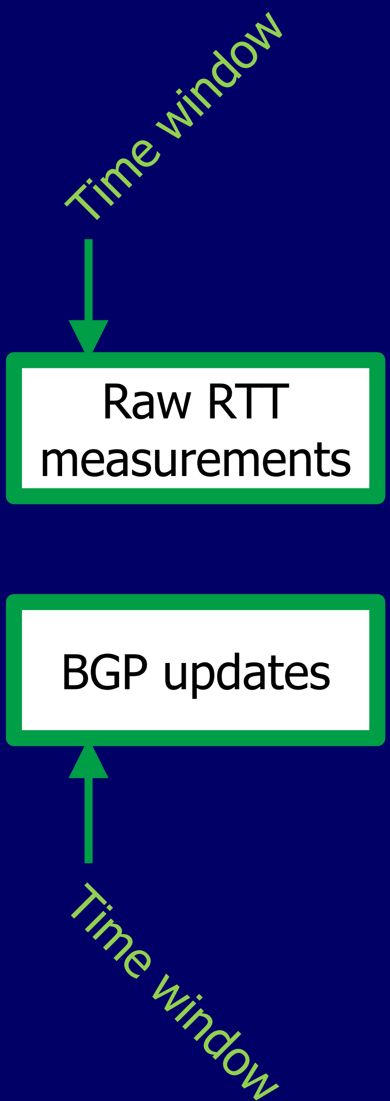

Methodology – Overview

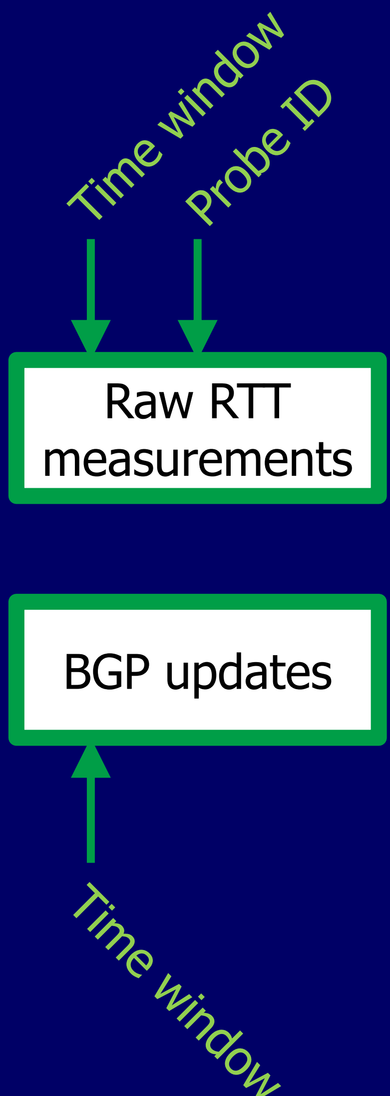

Methodology – Overview

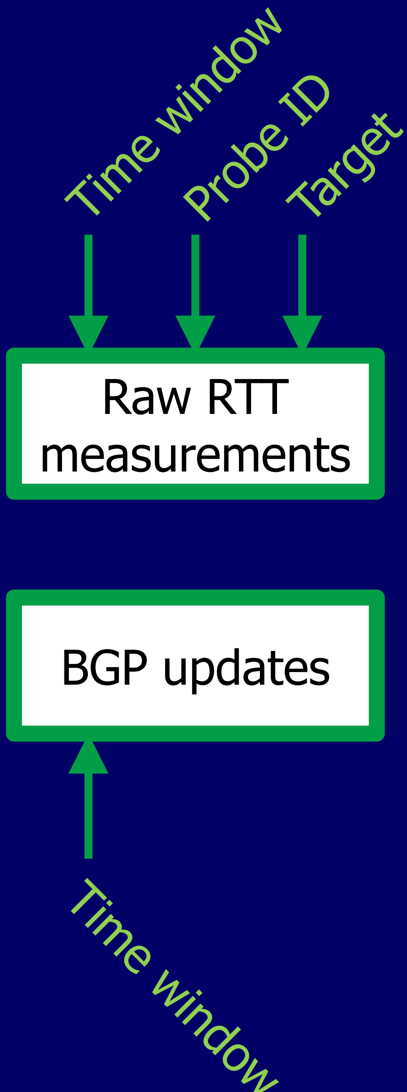

Methodology – Overview

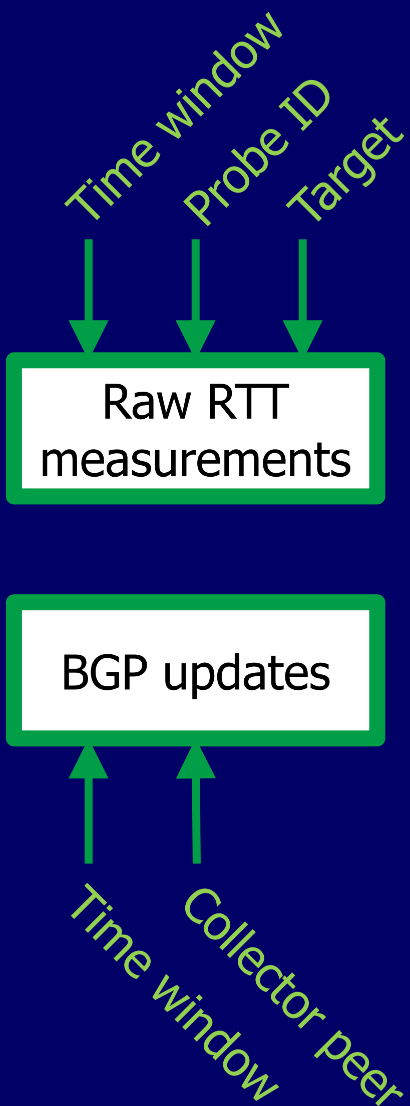

Methodology – Overview

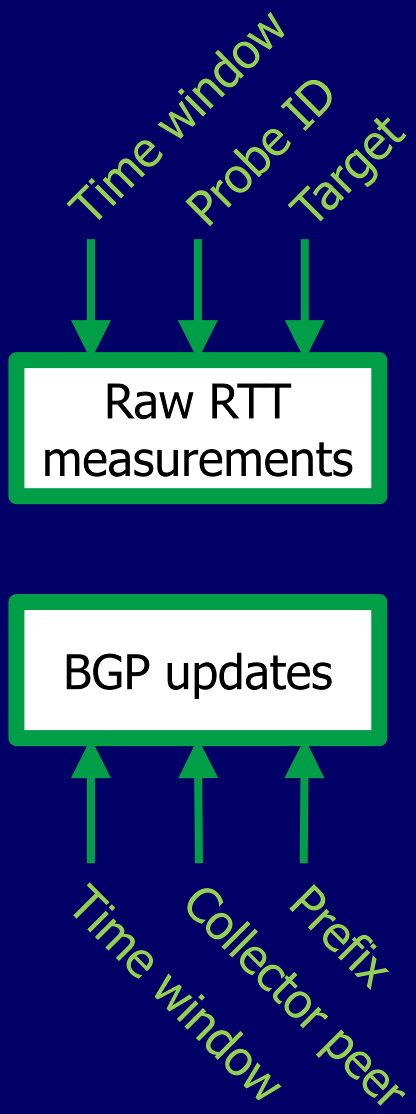

Methodology – Overview

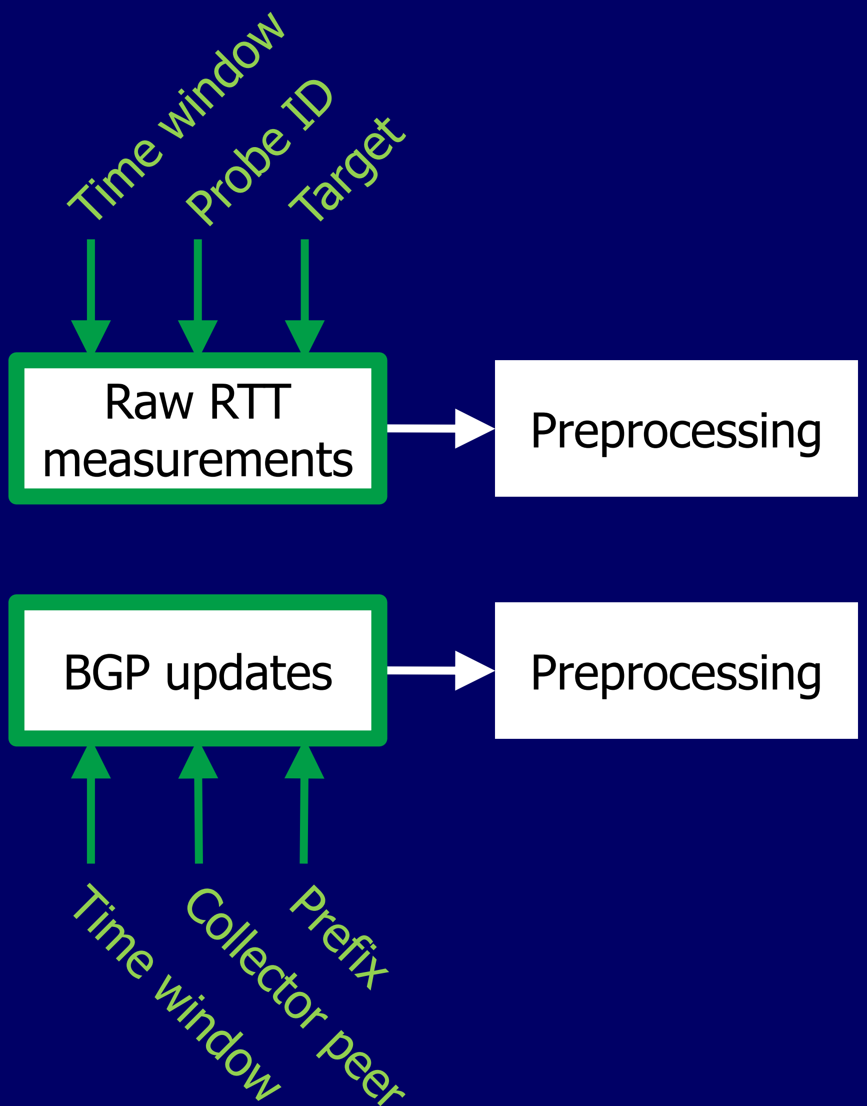

Methodology – Overview

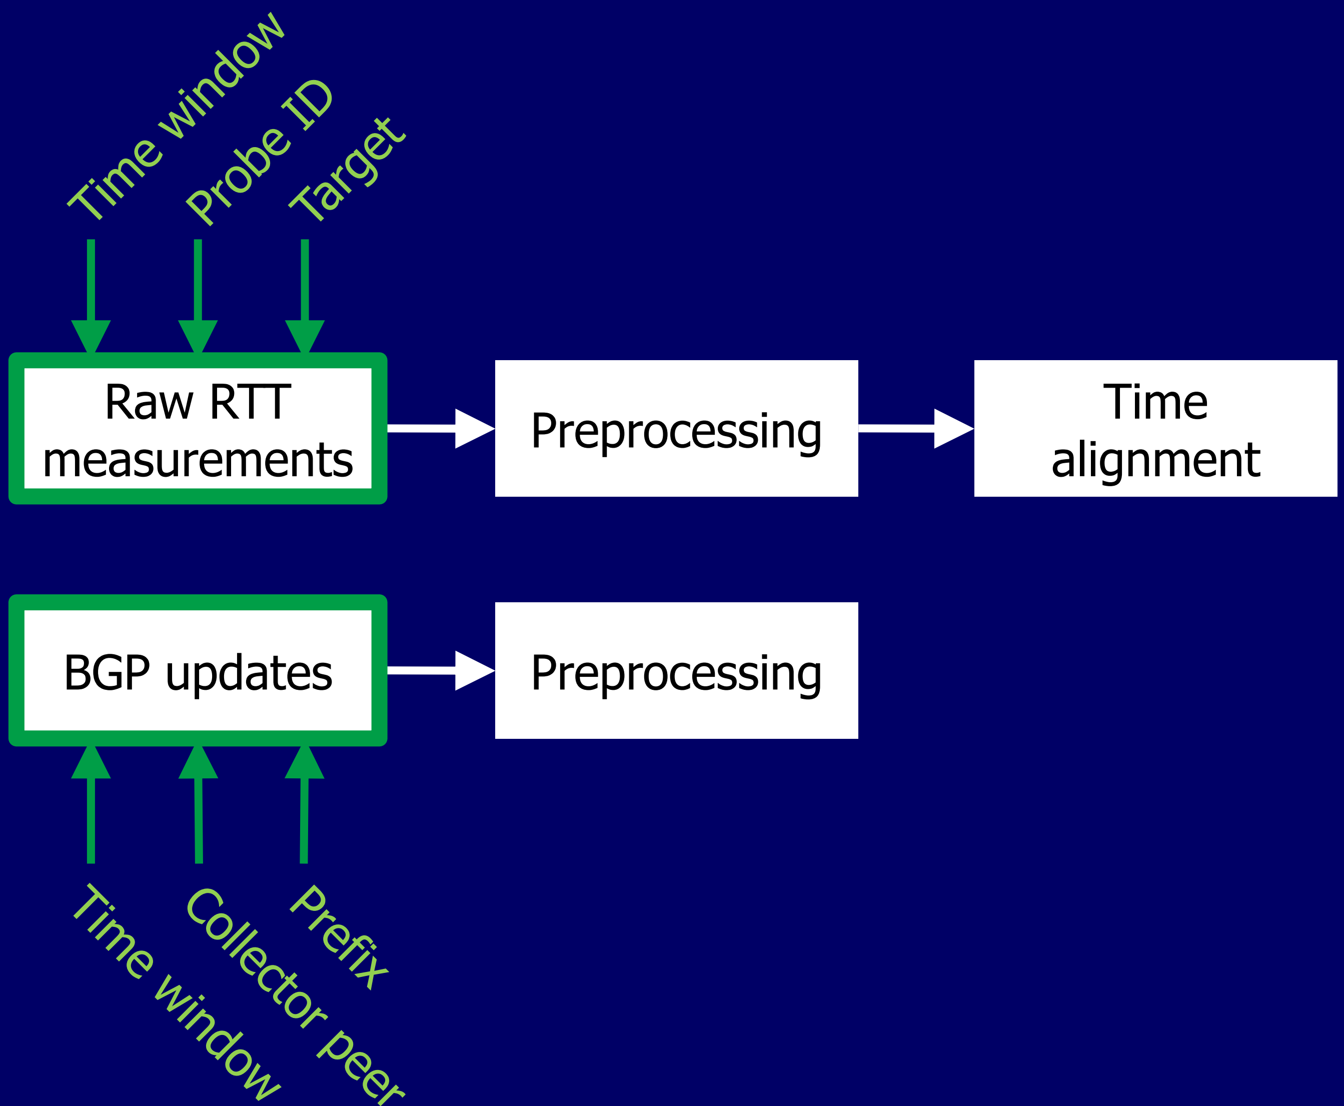

Methodology – Overview

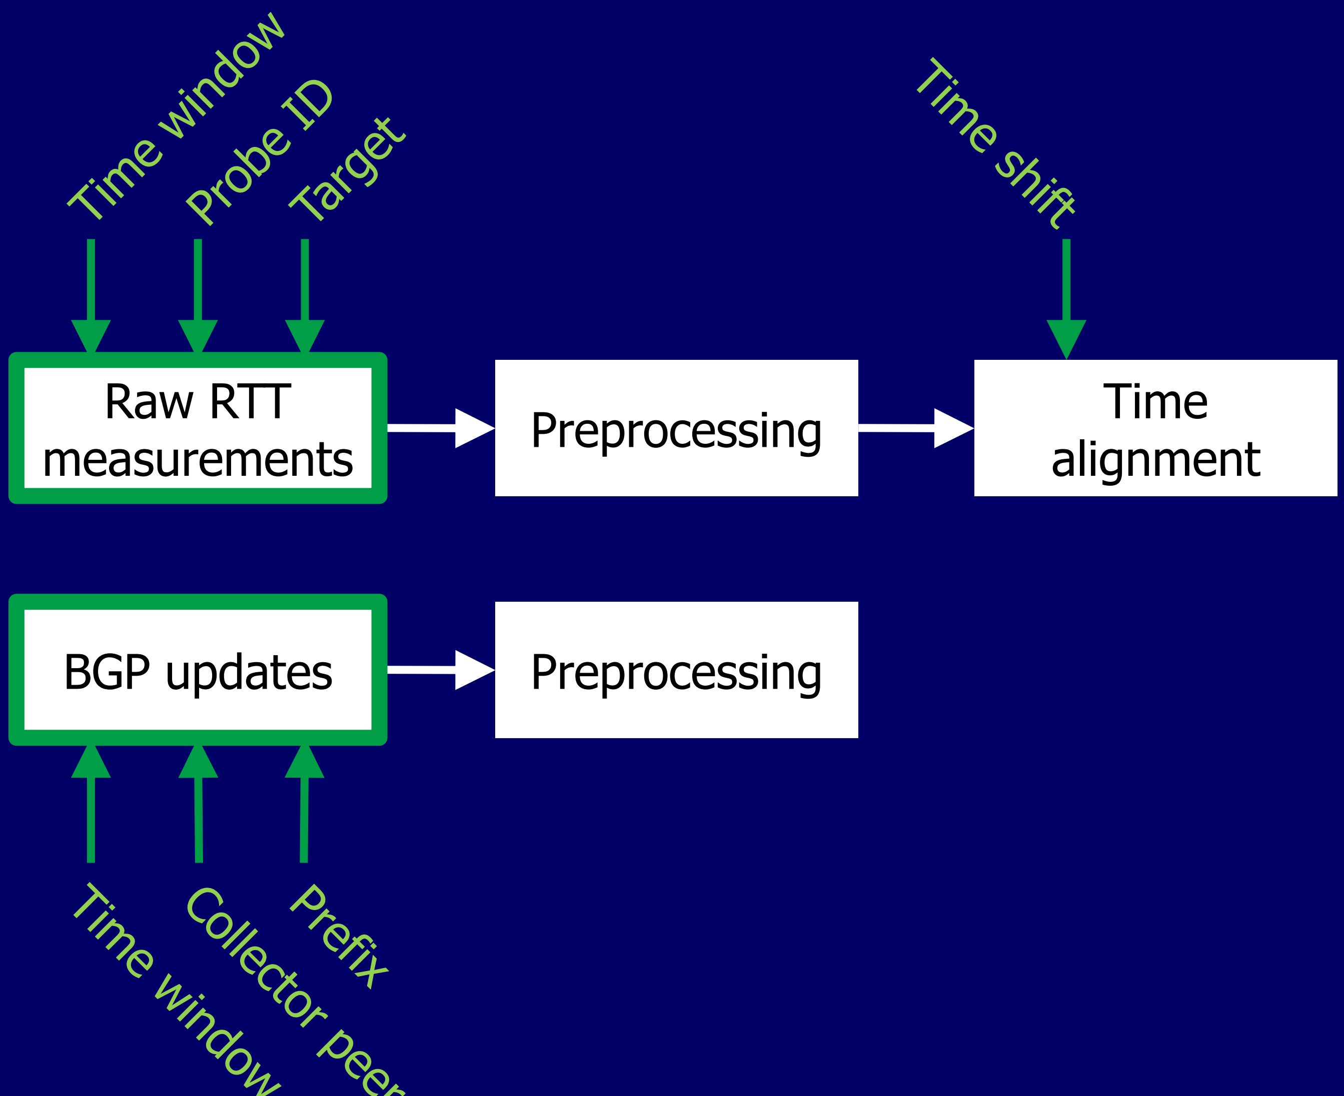

Methodology – Overview

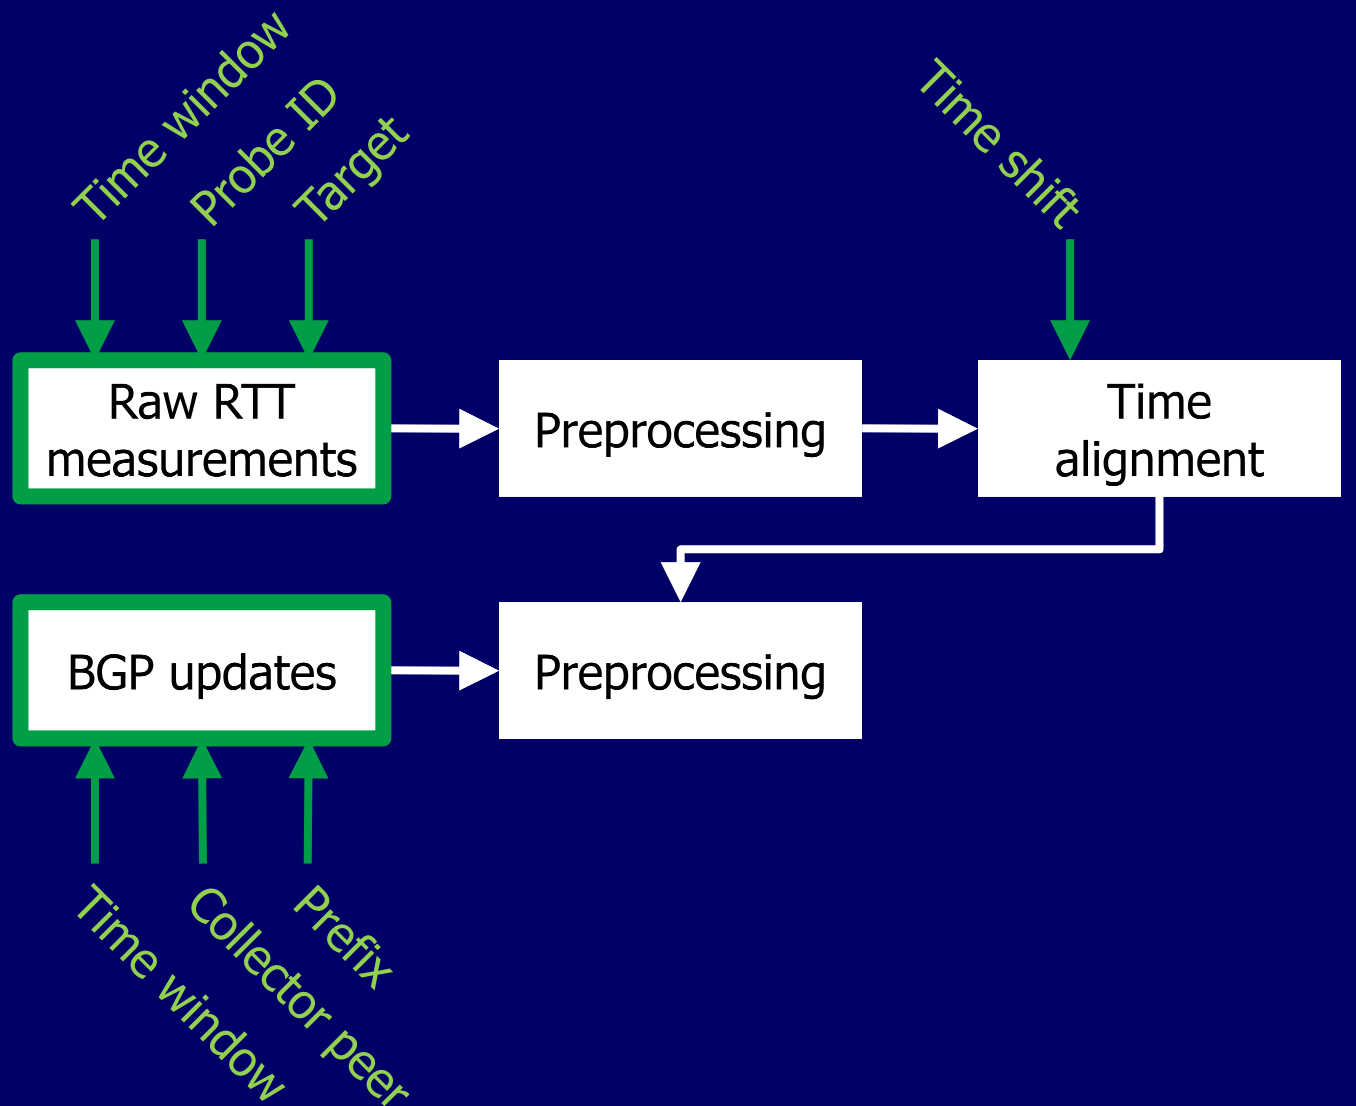

Methodology – Overview

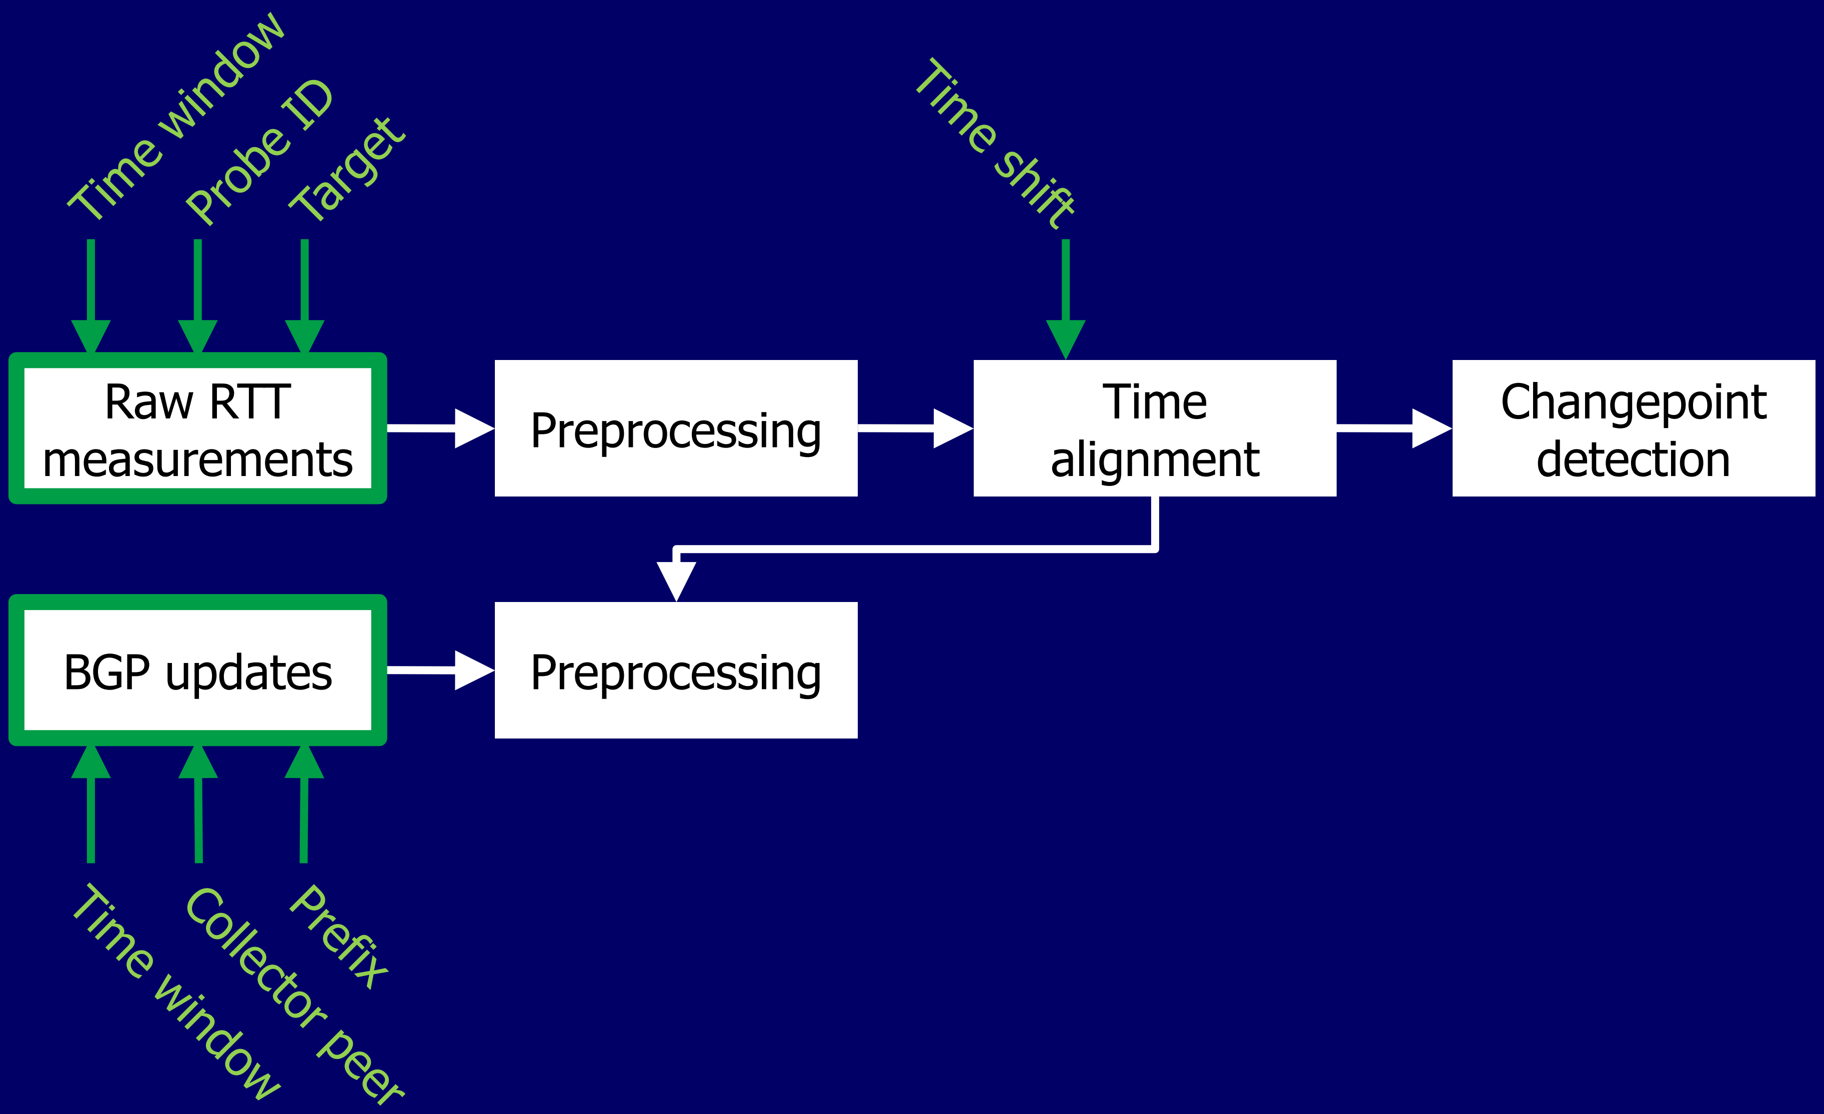

Methodology – Overview

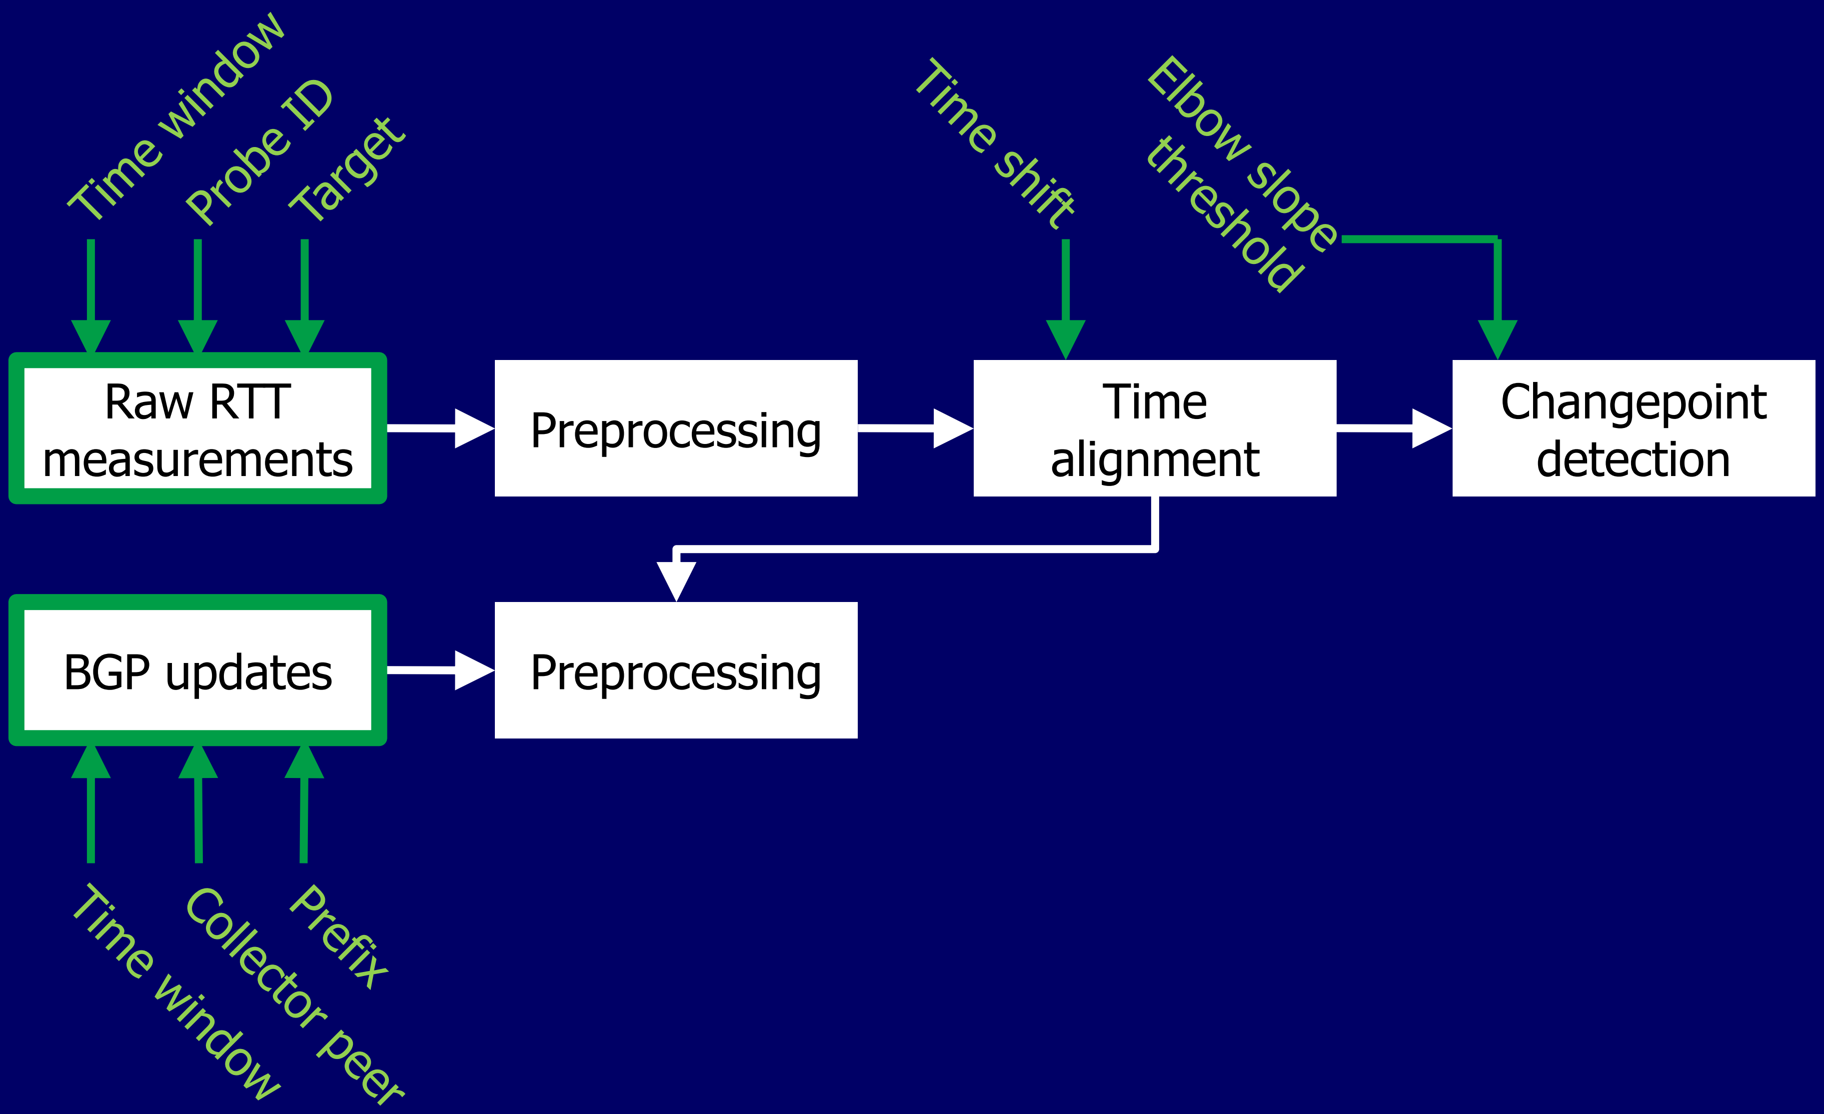

Methodology – Overview

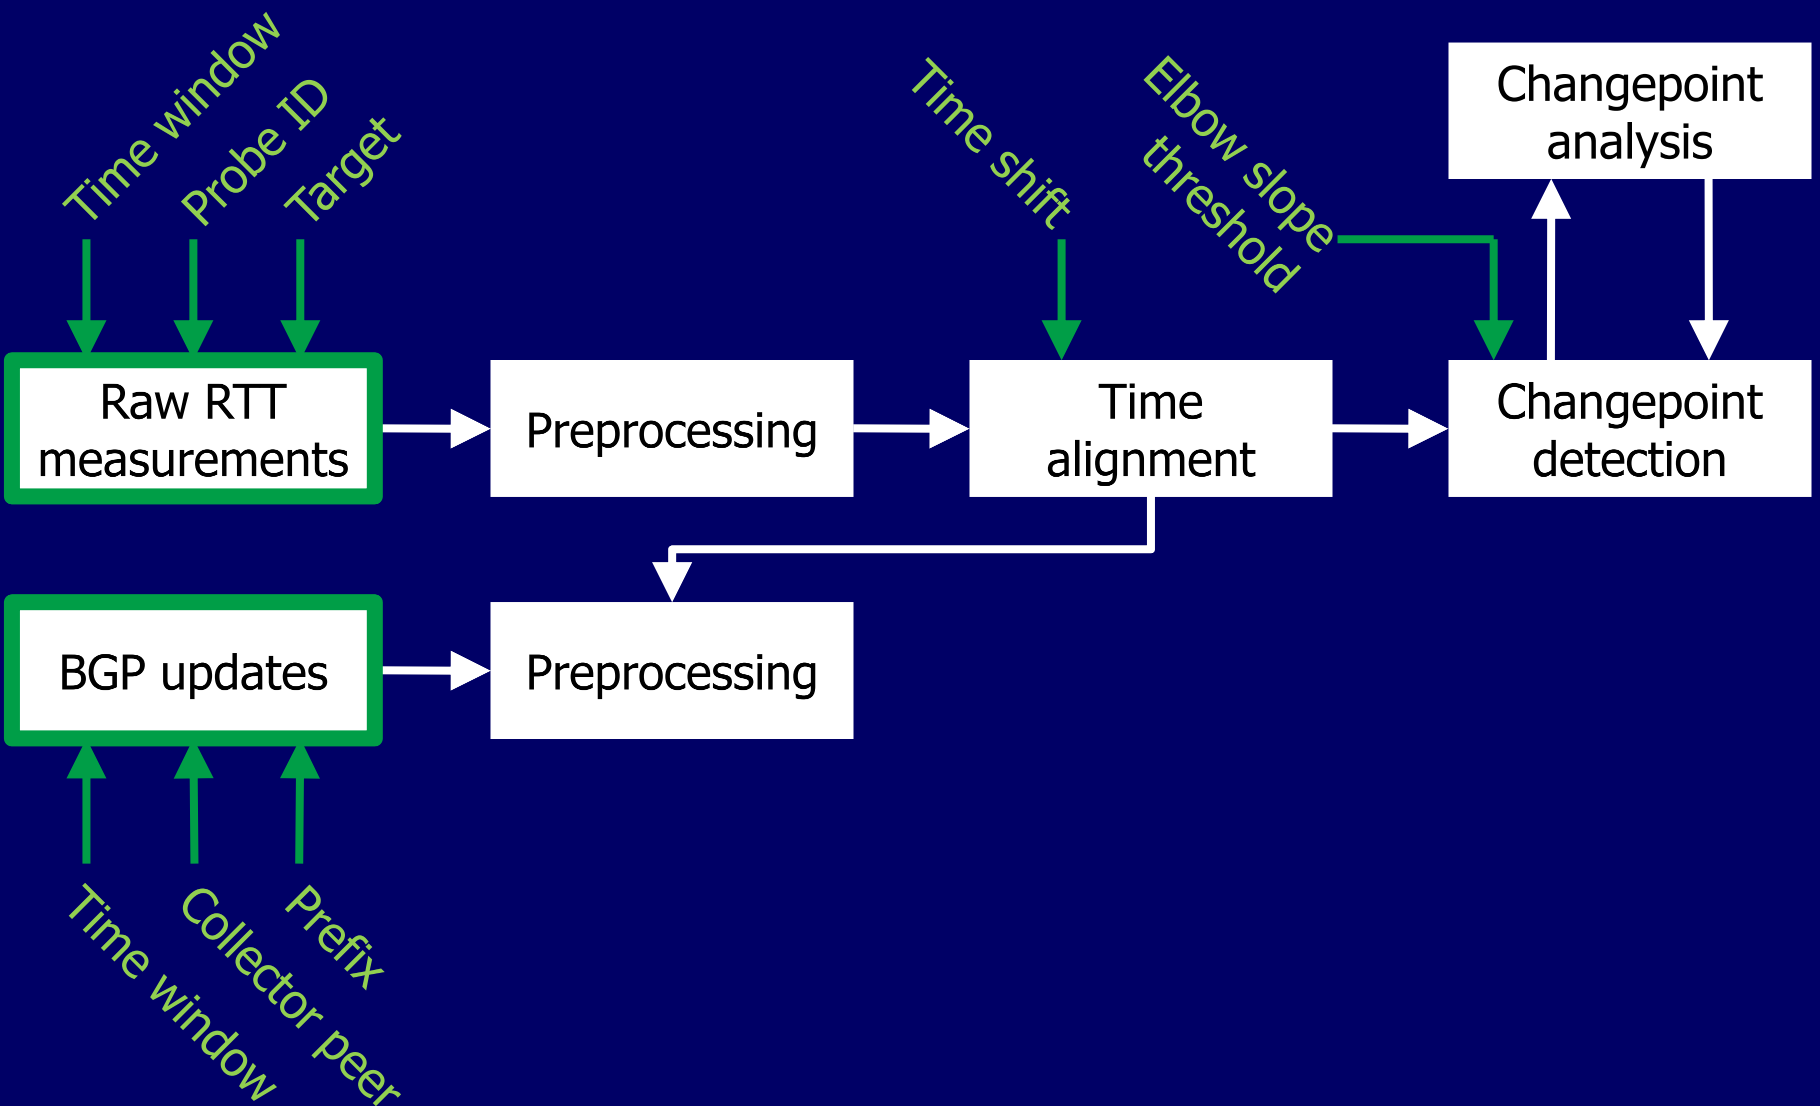

Methodology – Overview

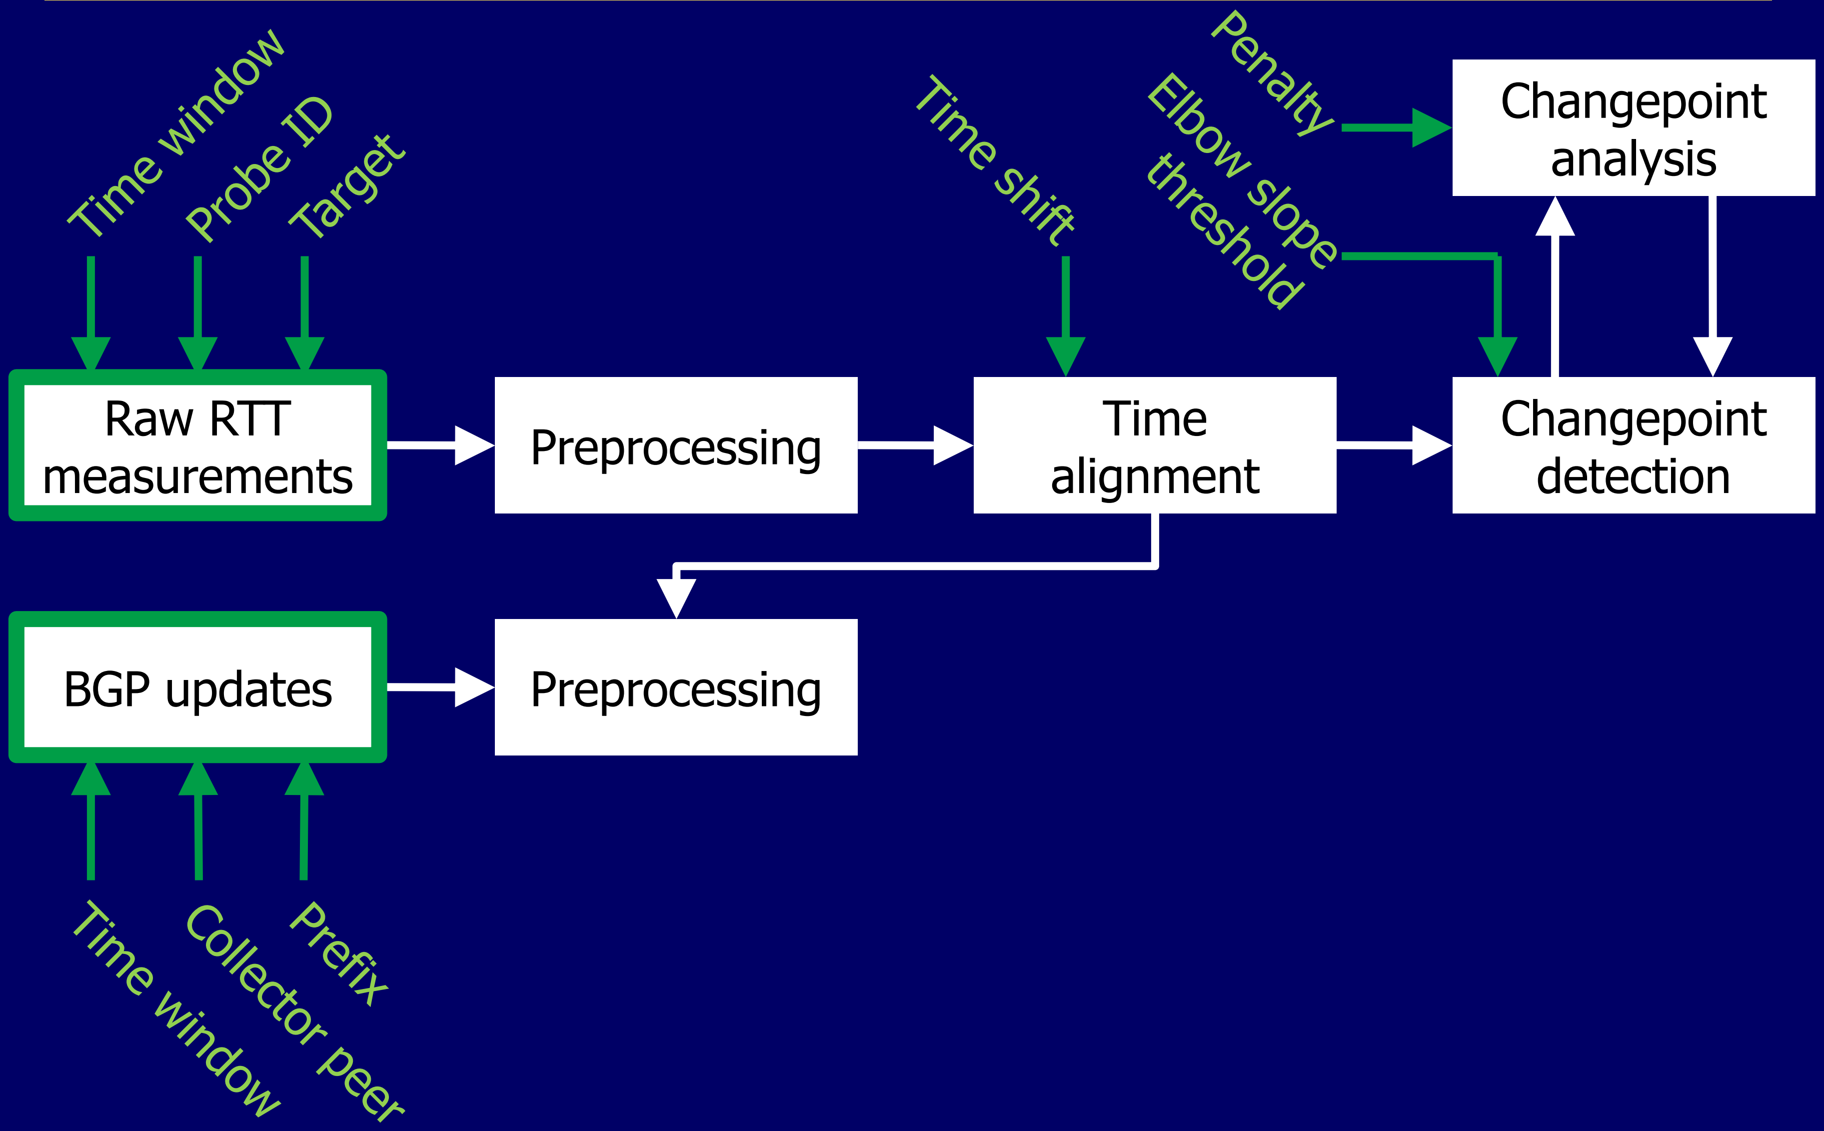

Methodology – Overview

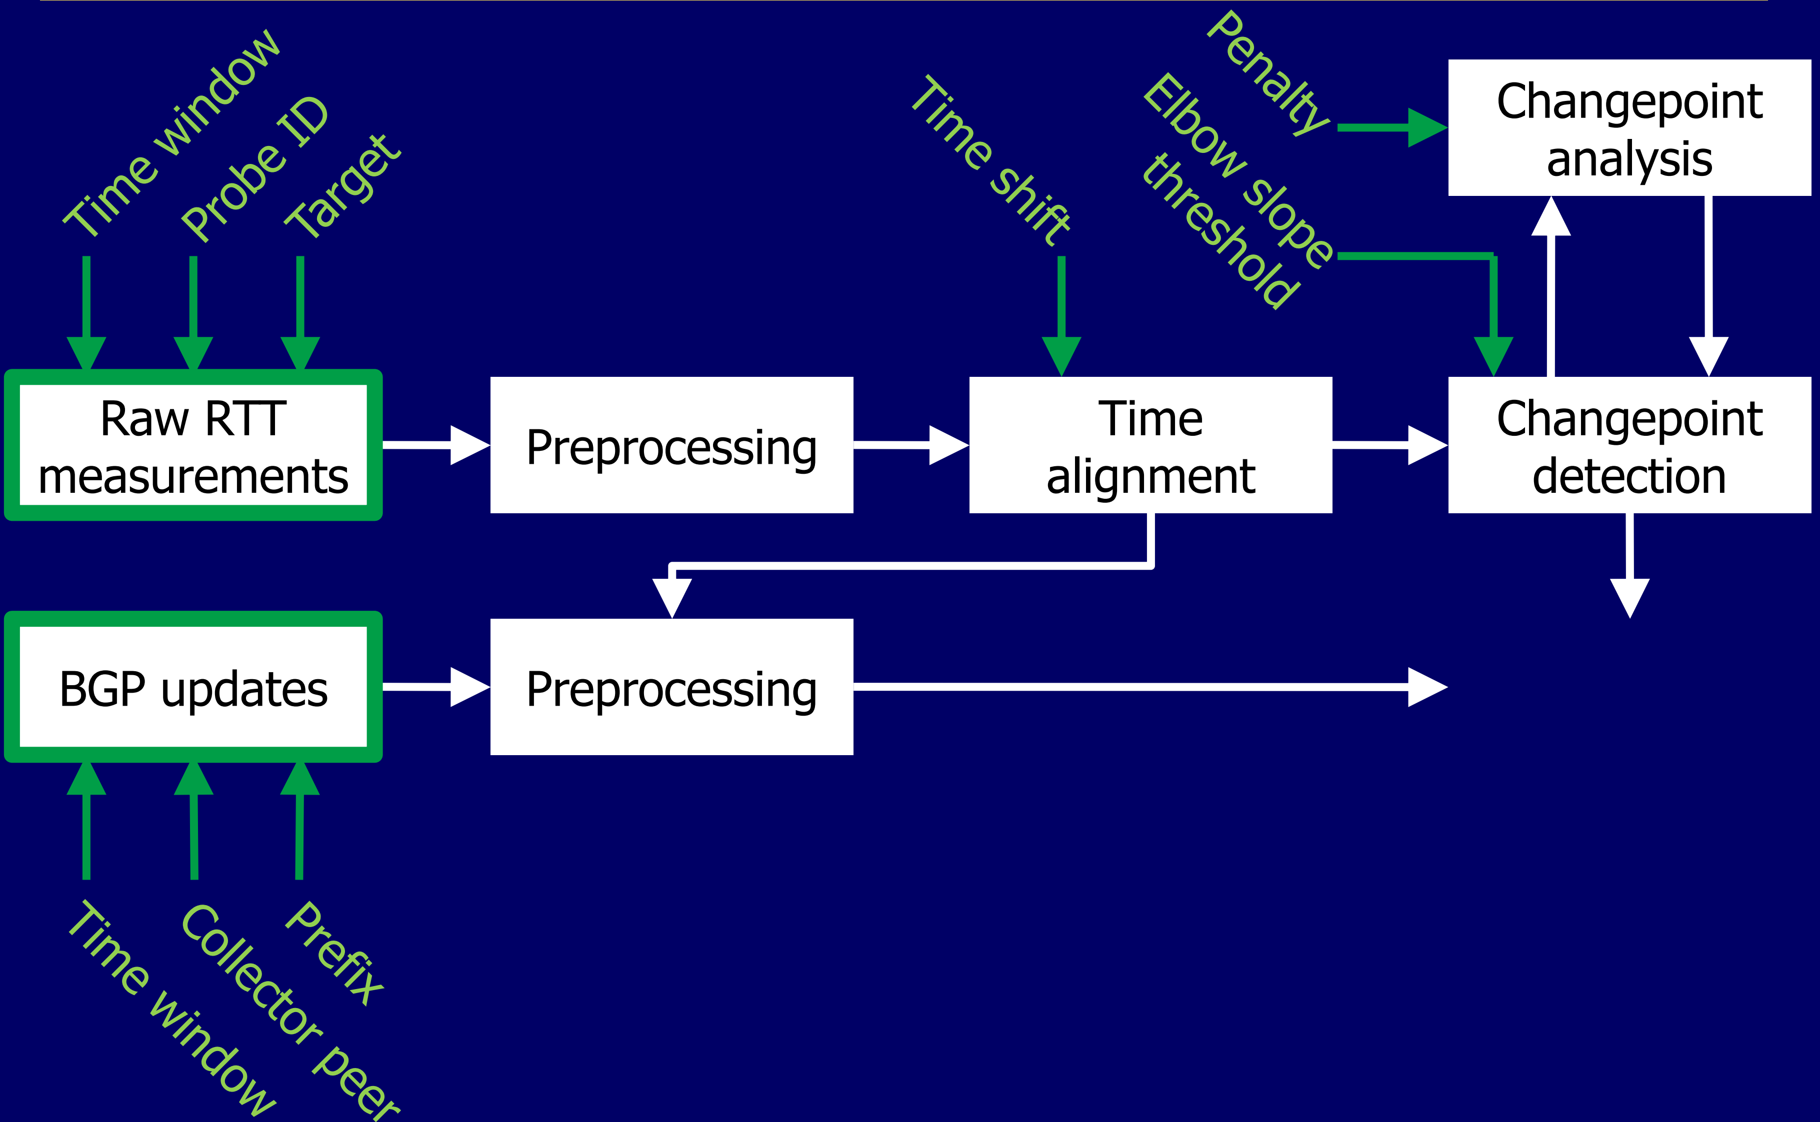

Methodology – Overview

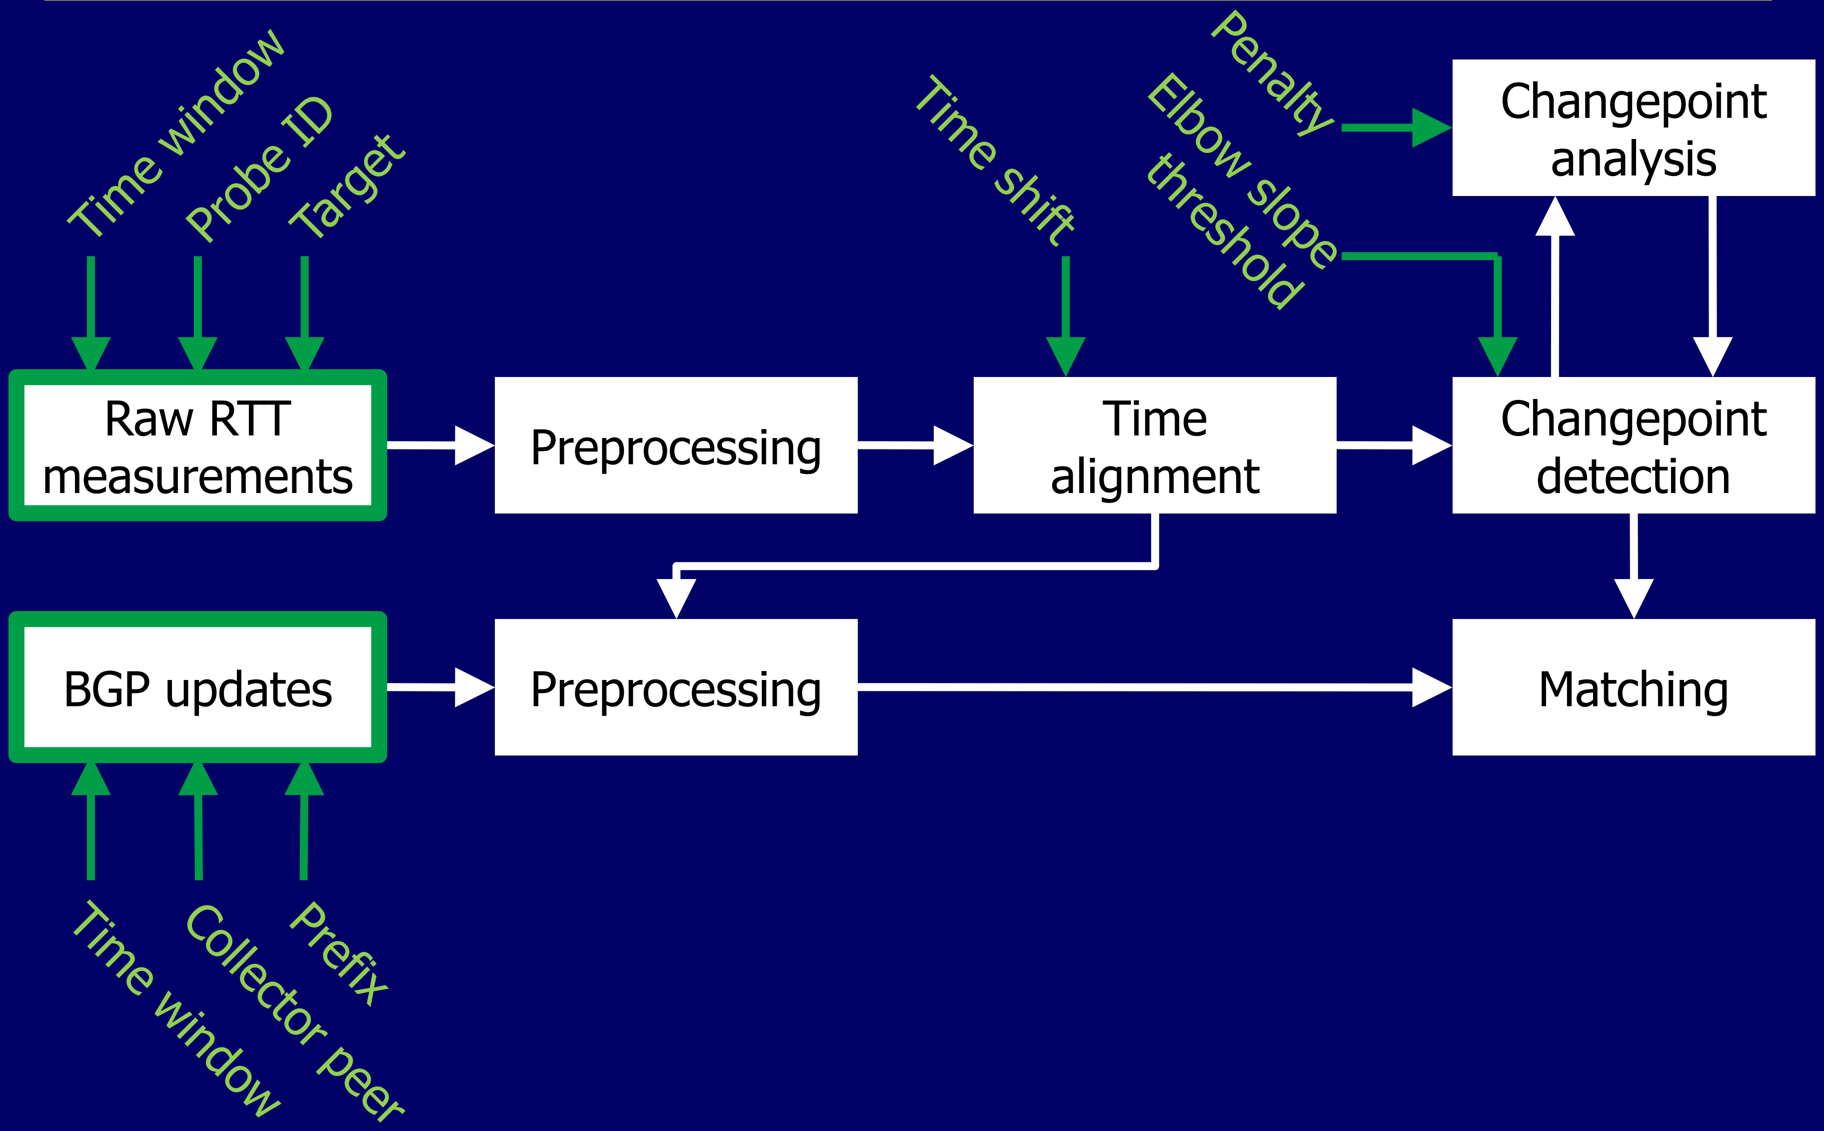

Methodology – Overview

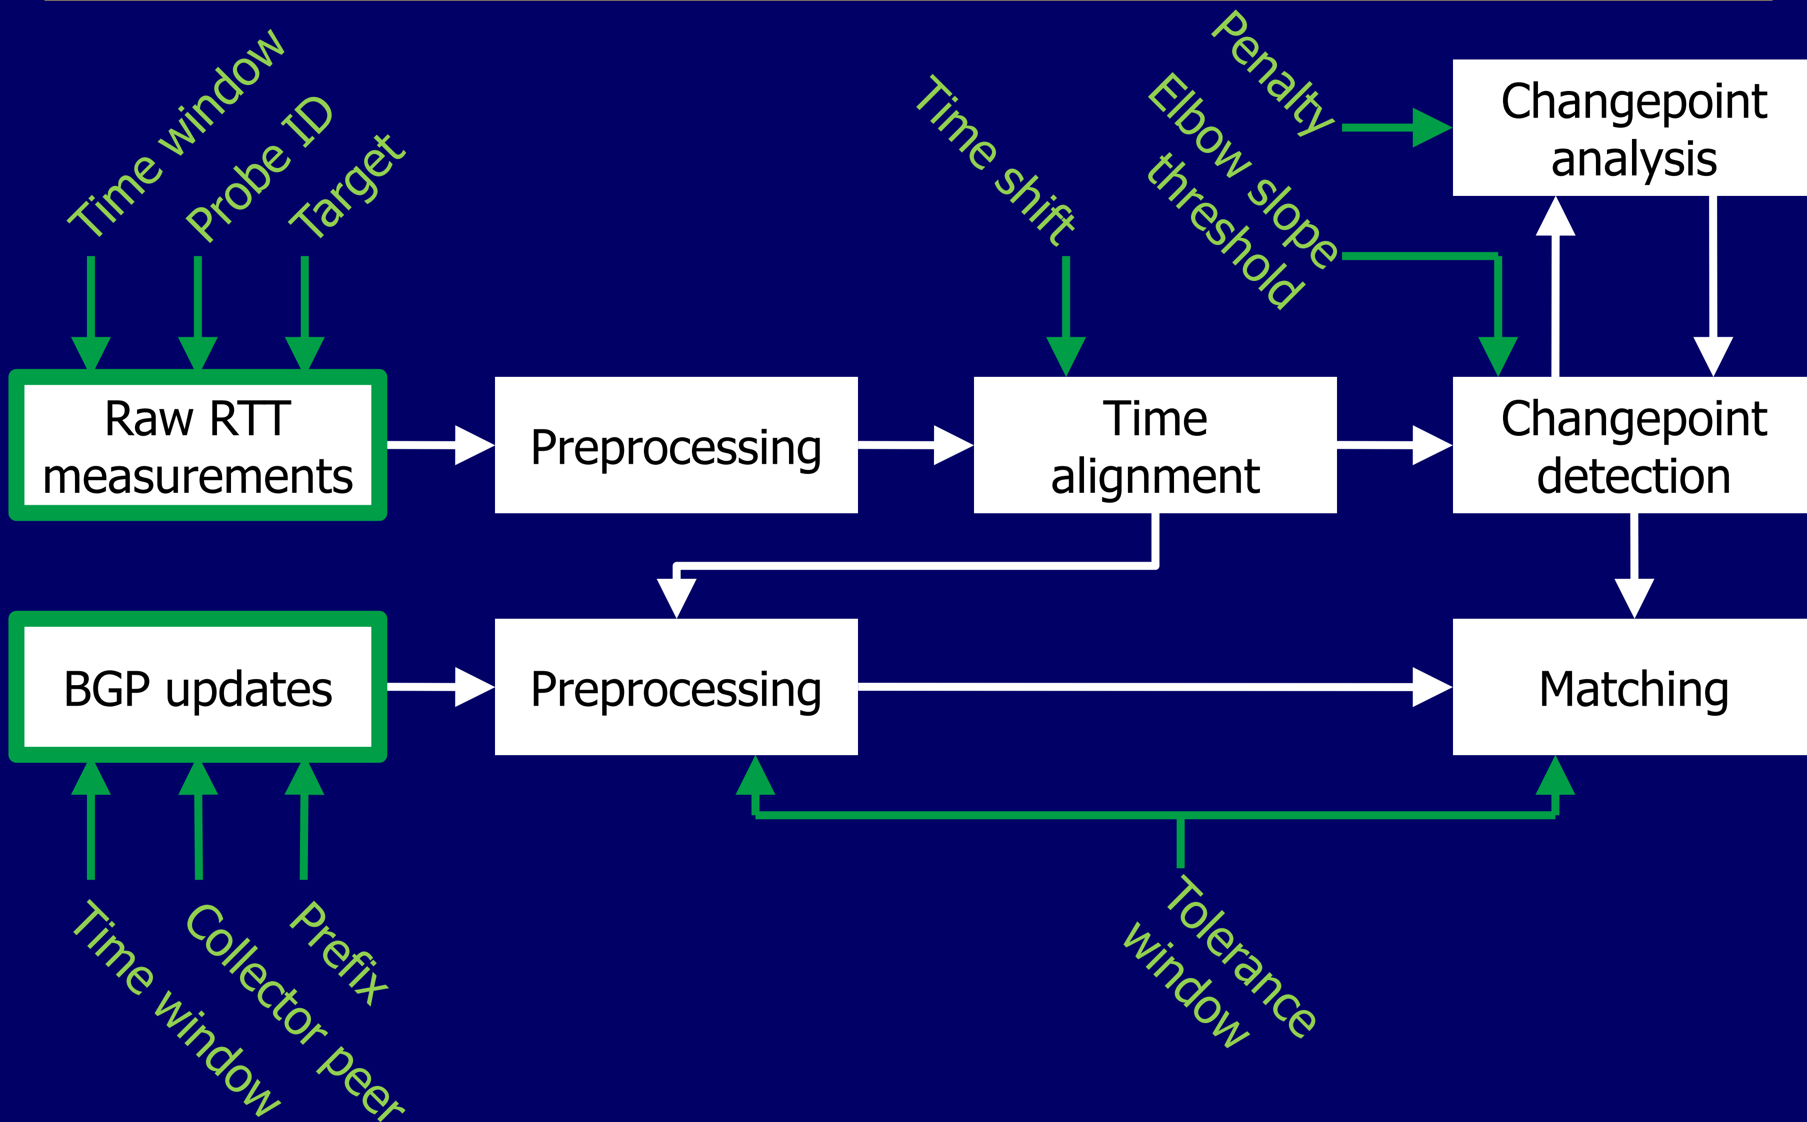

Methodology – Overview

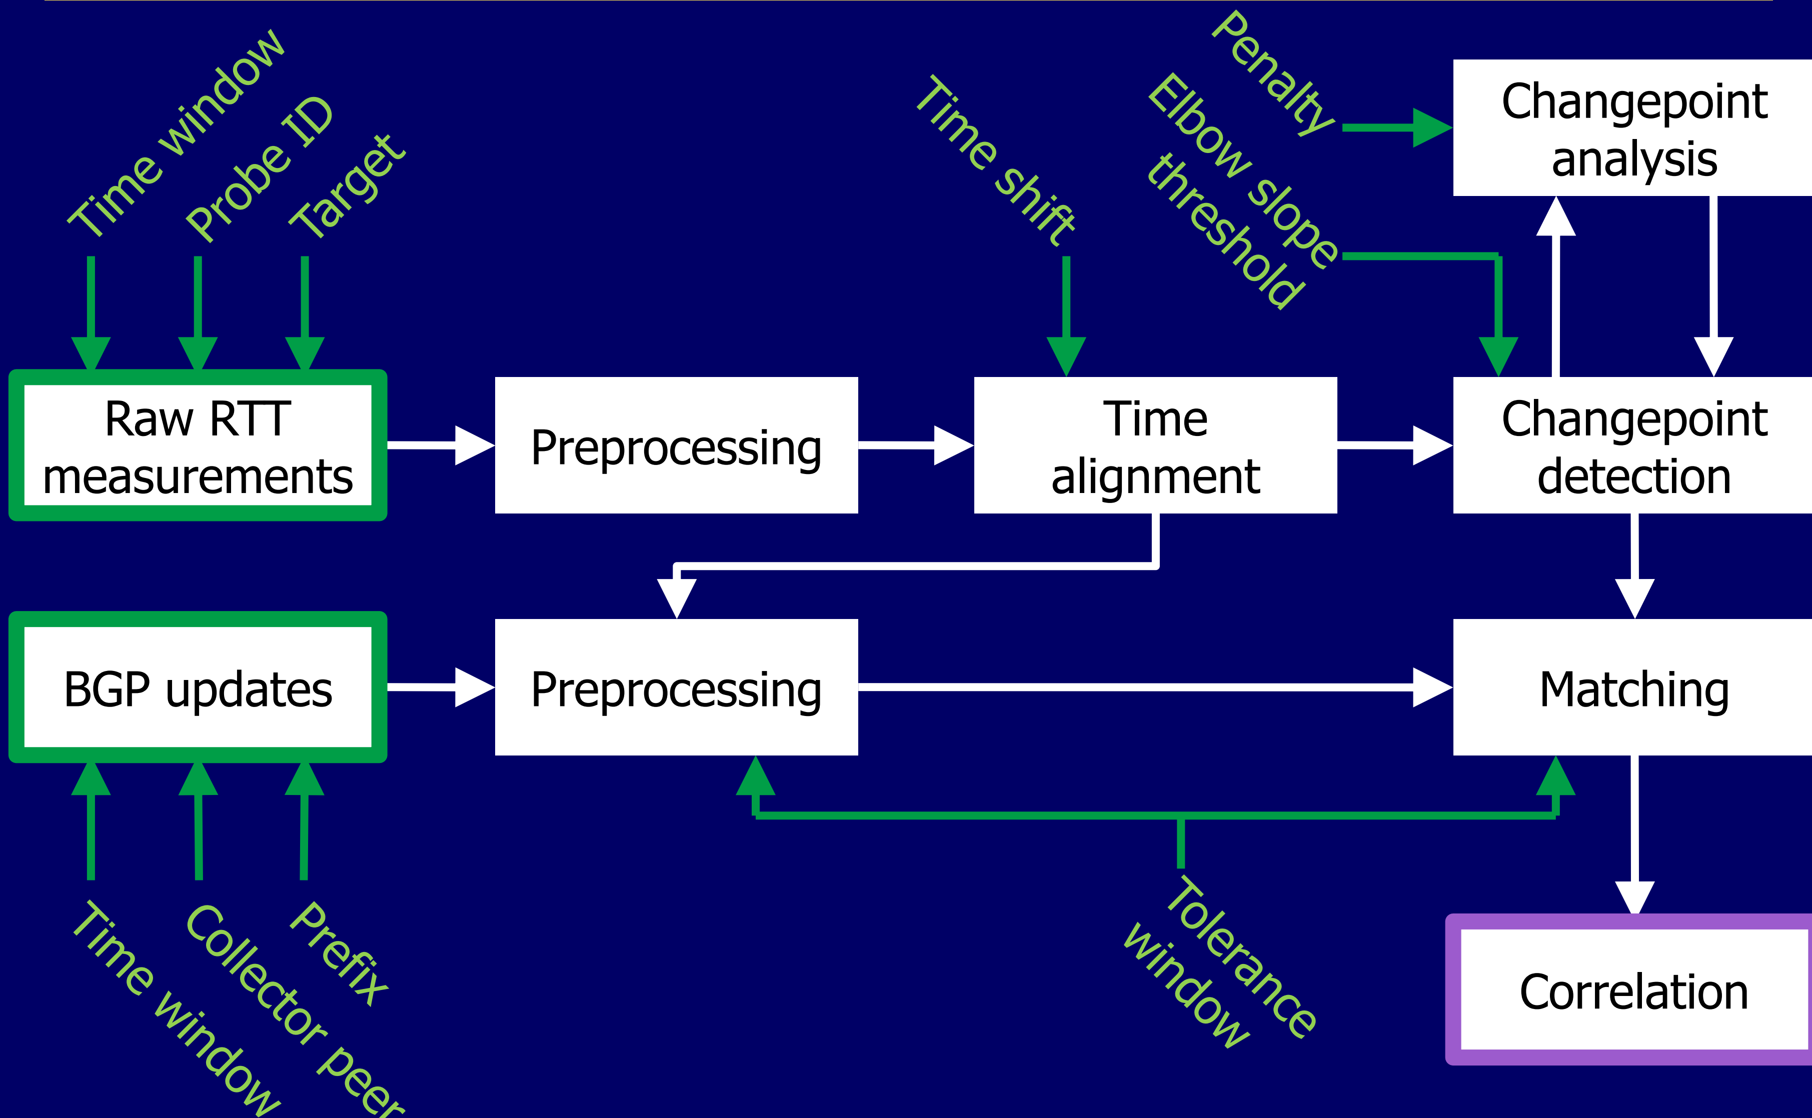

Methodology – Parameters

Methodology – Parameters

Data set	Correlation tuning
Time window	Time shift
Probe ID	Elbow slope threshold
Target	Penalty
(BGP) Prefix	Tolerance window
Collector peer	

Methodology – Parameters

Data set	Correlation tuning
Time window	Time shift
Probe ID	Elbow slope threshold
Target	Penalty
(BGP) Prefix	Tolerance window
Collector peer	

Methodology – Preprocessing

Methodology – Preprocessing

Methodology – Preprocessing

RTT Measurements **BGP updates**

Methodology – Preprocessing

RTT Measurements	BGP updates
Periodical	On-change

Methodology – Preprocessing

RTT Measurements	BGP updates
Periodical	On-change
Constant rate	Possibly bursty

Methodology – Preprocessing

RTT Measurements	BGP updates
Periodical	On-change
Constant rate	Possibly bursty
Highly variable	Few paths

Methodology – Preprocessing

RTT Measurements	BGP updates
Periodical	On-change
Constant rate	Possibly bursty
Highly variable	Few paths

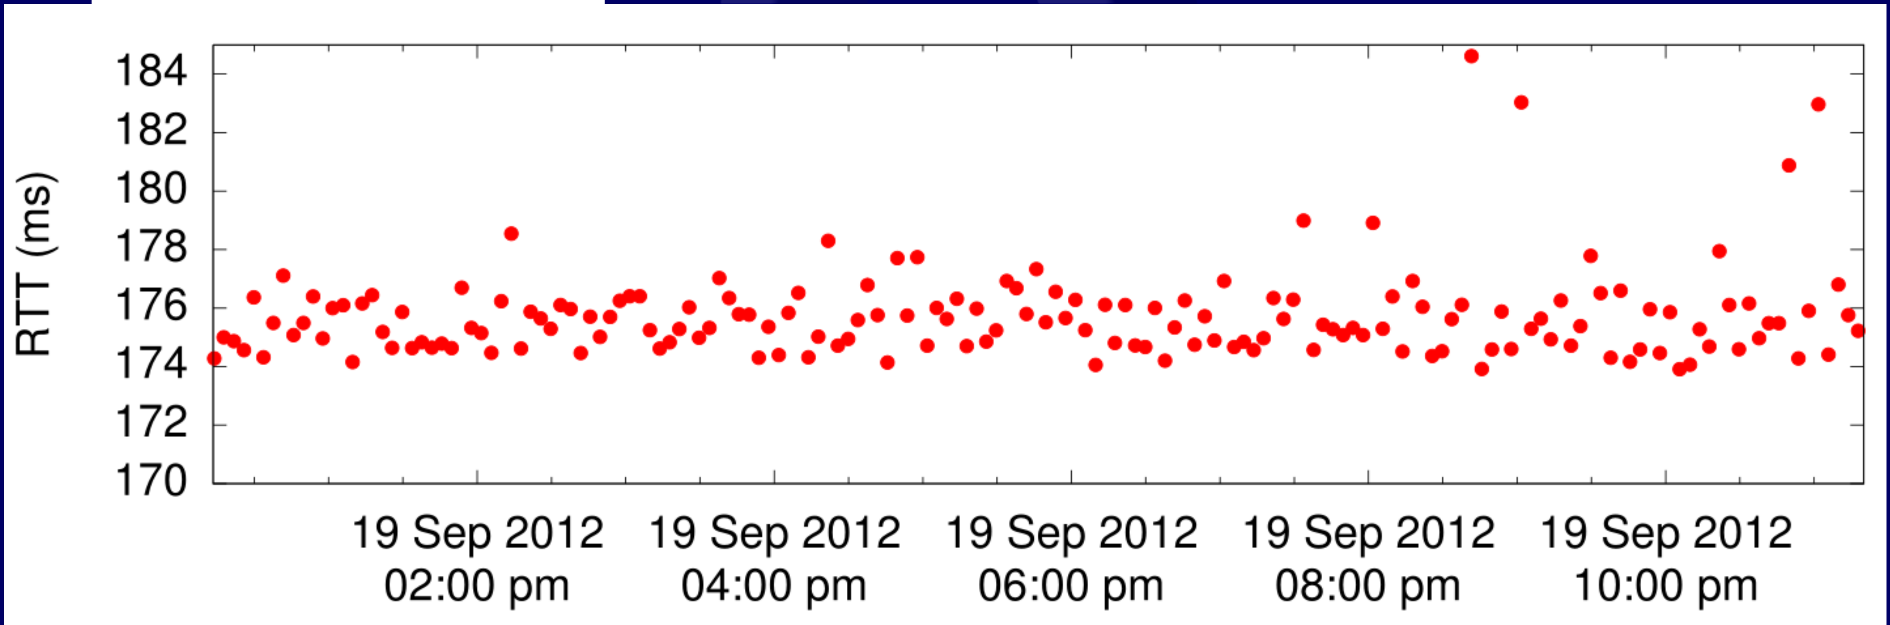

Methodology – Preprocessing

RTT Measurements	BGP updates
Periodical	On-change
Constant rate	Possibly bursty
Highly variable	Few paths

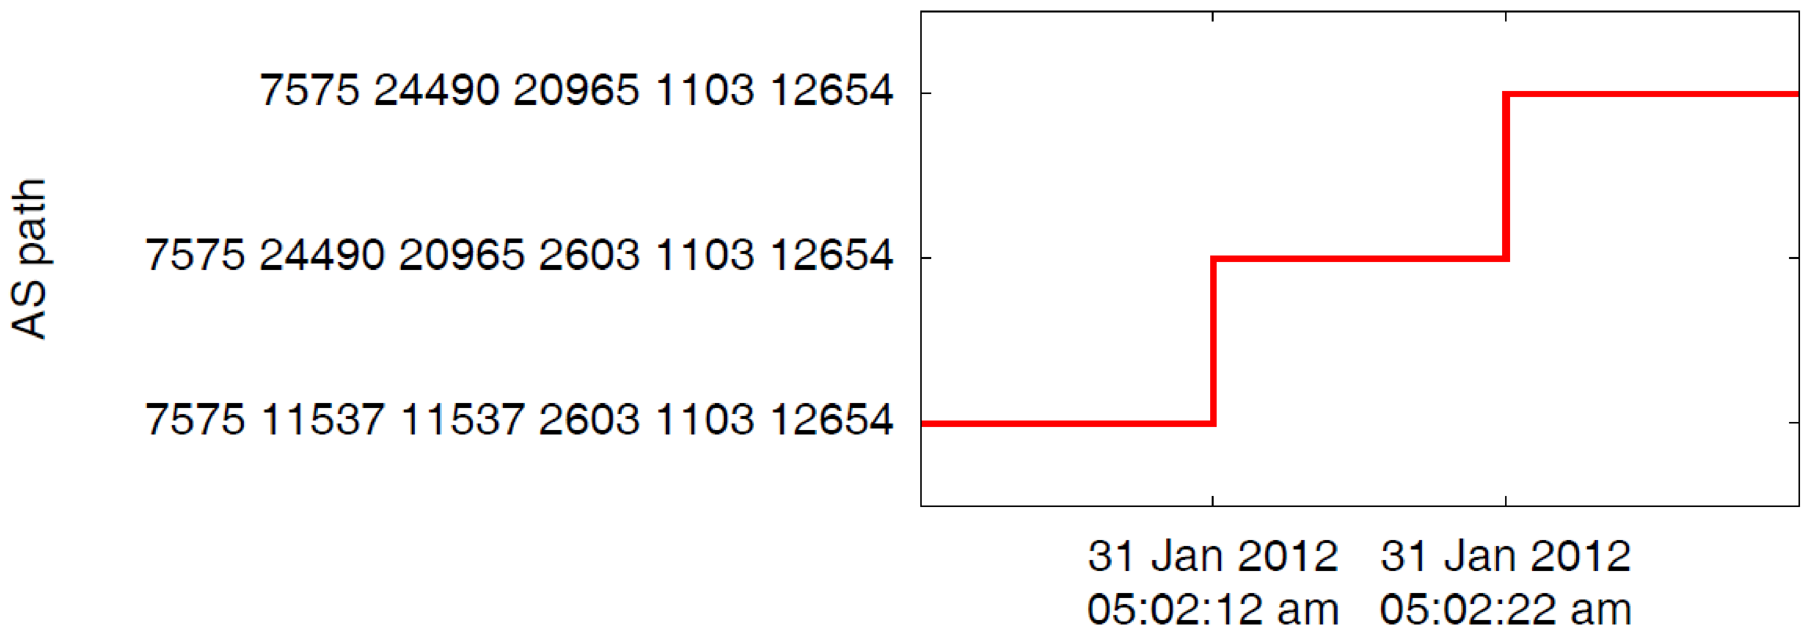

Methodology – Preprocessing

Methodology – Preprocessing

Methodology – Time shift

- Account for
  - clock offsets
  - BGP update propagation delays
    - MRAI
    - relative position of devices

Methodology – Changepoint detection

- Goal

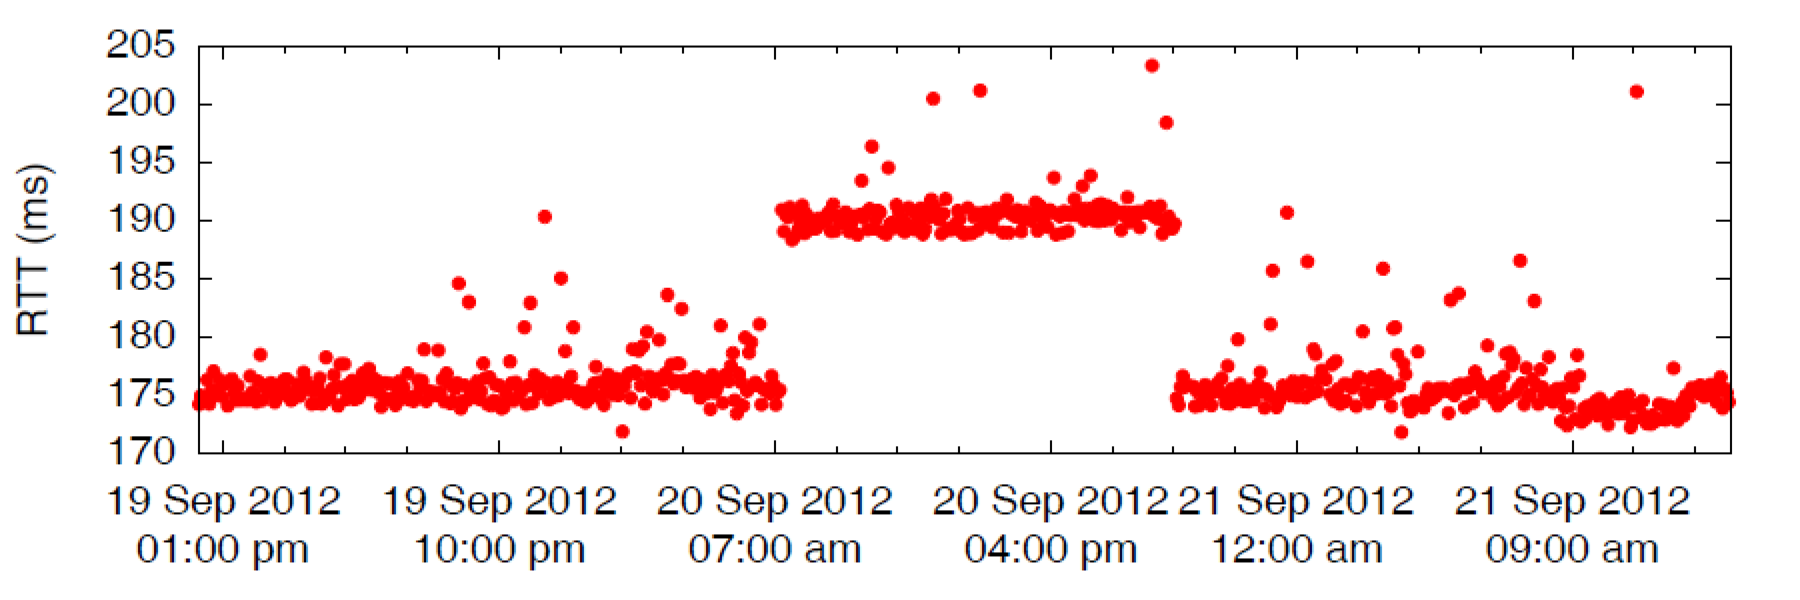

Methodology – Changepoint detection

- Goal

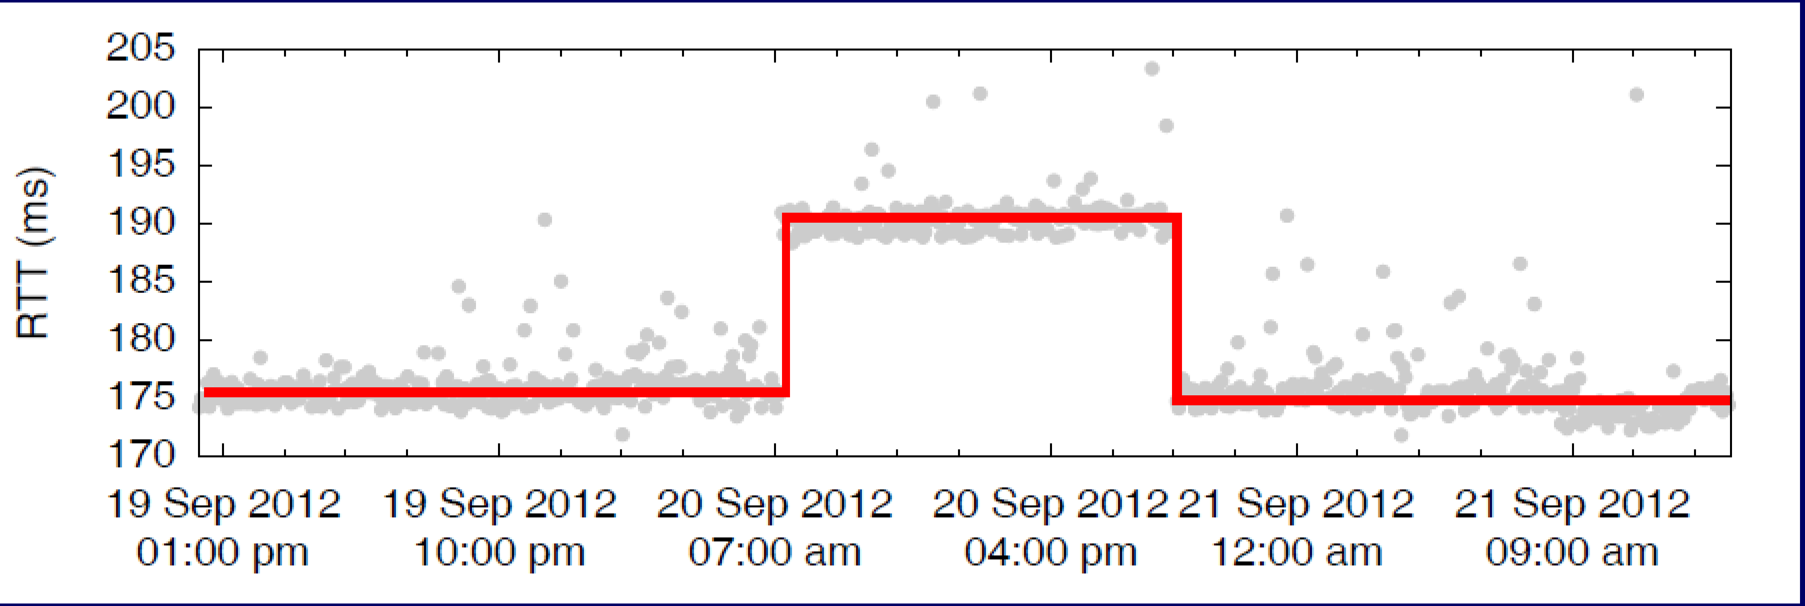

Methodology – Changepoint detection

- Technique
  - Changepoint analysis statistical method(s)

Methodology – Changepoint detection

- Technique
  - Changepoint analysis statistical method(s)

Pruned Exact Linear Time (PELT)

- Killick, R., Fearnhead, P., Eckley, I.: **Optimal detection of changepoints with a linear computational cost**. Jour. Amer. Stat. Assoc. 107(500), 1590–1598, 2012

Methodology – Changepoint detection

- Technique
  - Changepoint analysis statistical method(s)

    Pruned Exact Linear Time (PELT)

  - Detect mean & variance shifts in time series data

Methodology – Changepoint detection

- Technique
  - Changepoint analysis statistical method(s)

    Pruned Exact Linear Time (PELT)

  - Detect mean & variance shifts in time series data
    - long-lasting small changes
    - short-lived significant changes

Methodology – Changepoint detection

- Technique
  - Changepoint analysis statistical method(s)

    Pruned Exact Linear Time (PELT)

  - Detect mean & variance shifts in time series data
    - long-lasting small changes
    - short-lived significant changes
  - Tunable

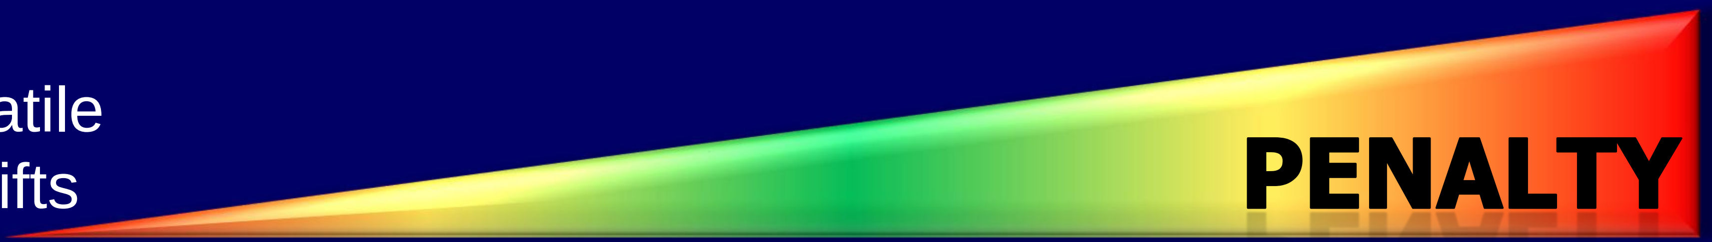

Methodology – Changepoint detection

- Technique
  - Changepoint analysis statistical method(s)

    Pruned Exact Linear Time (PELT)

  - Detect mean & variance shifts in time series data
    - long-lasting small changes
    - short-lived significant changes
  - Tunable

Methodology – Changepoint detection

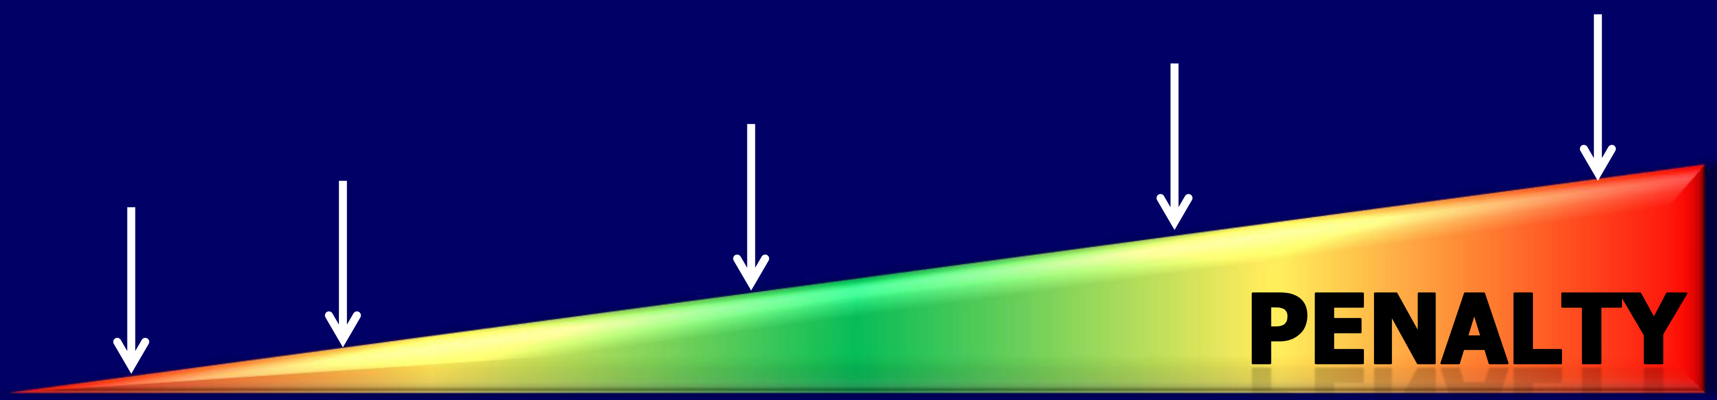

Methodology – Changepoint detection

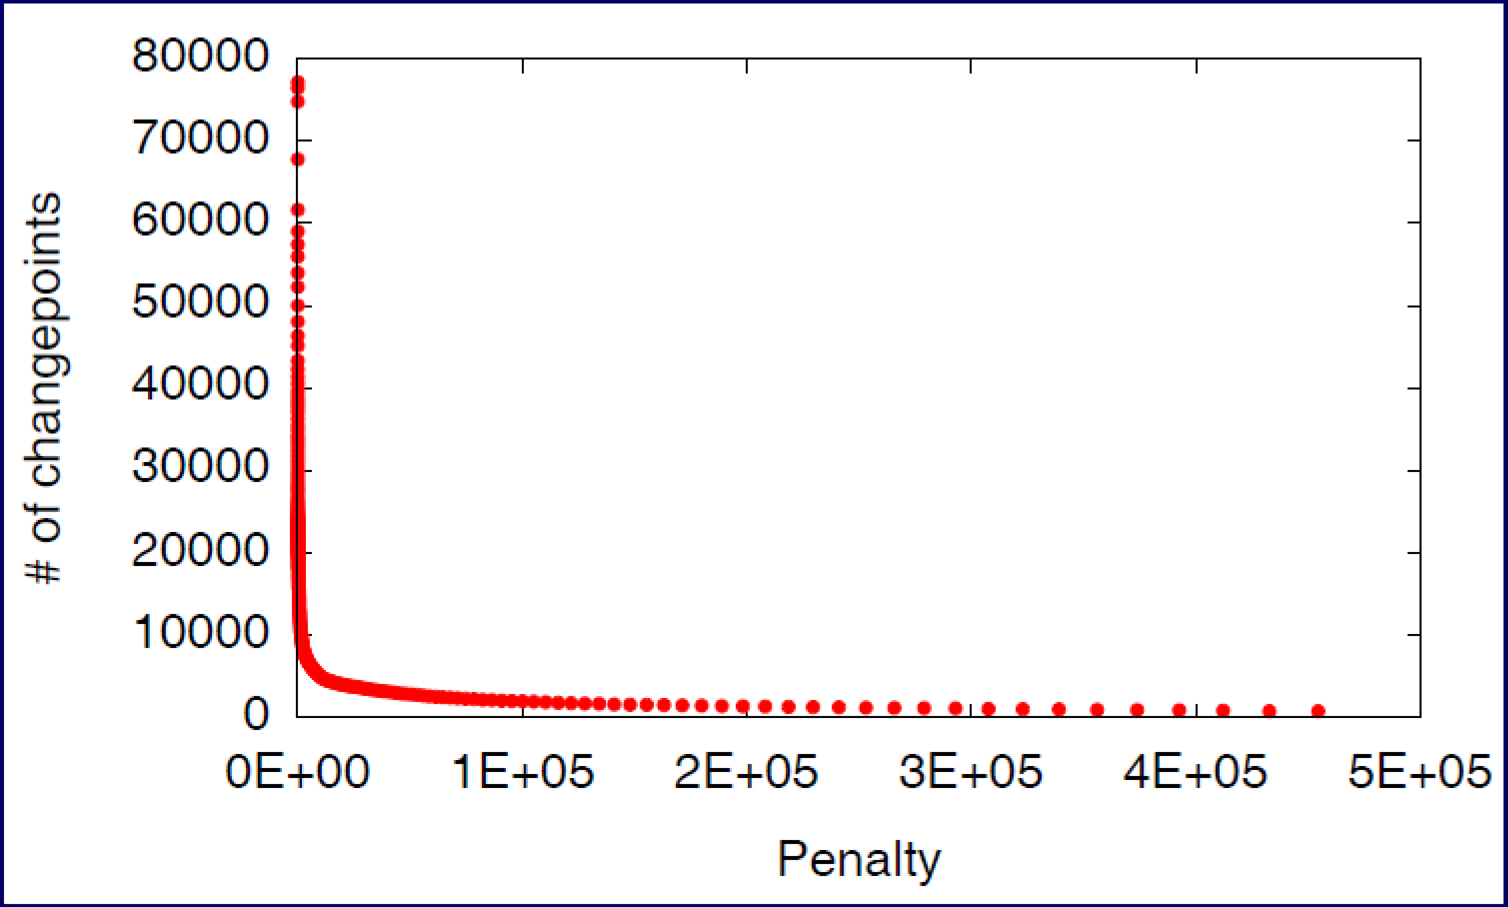

Methodology – Changepoint detection

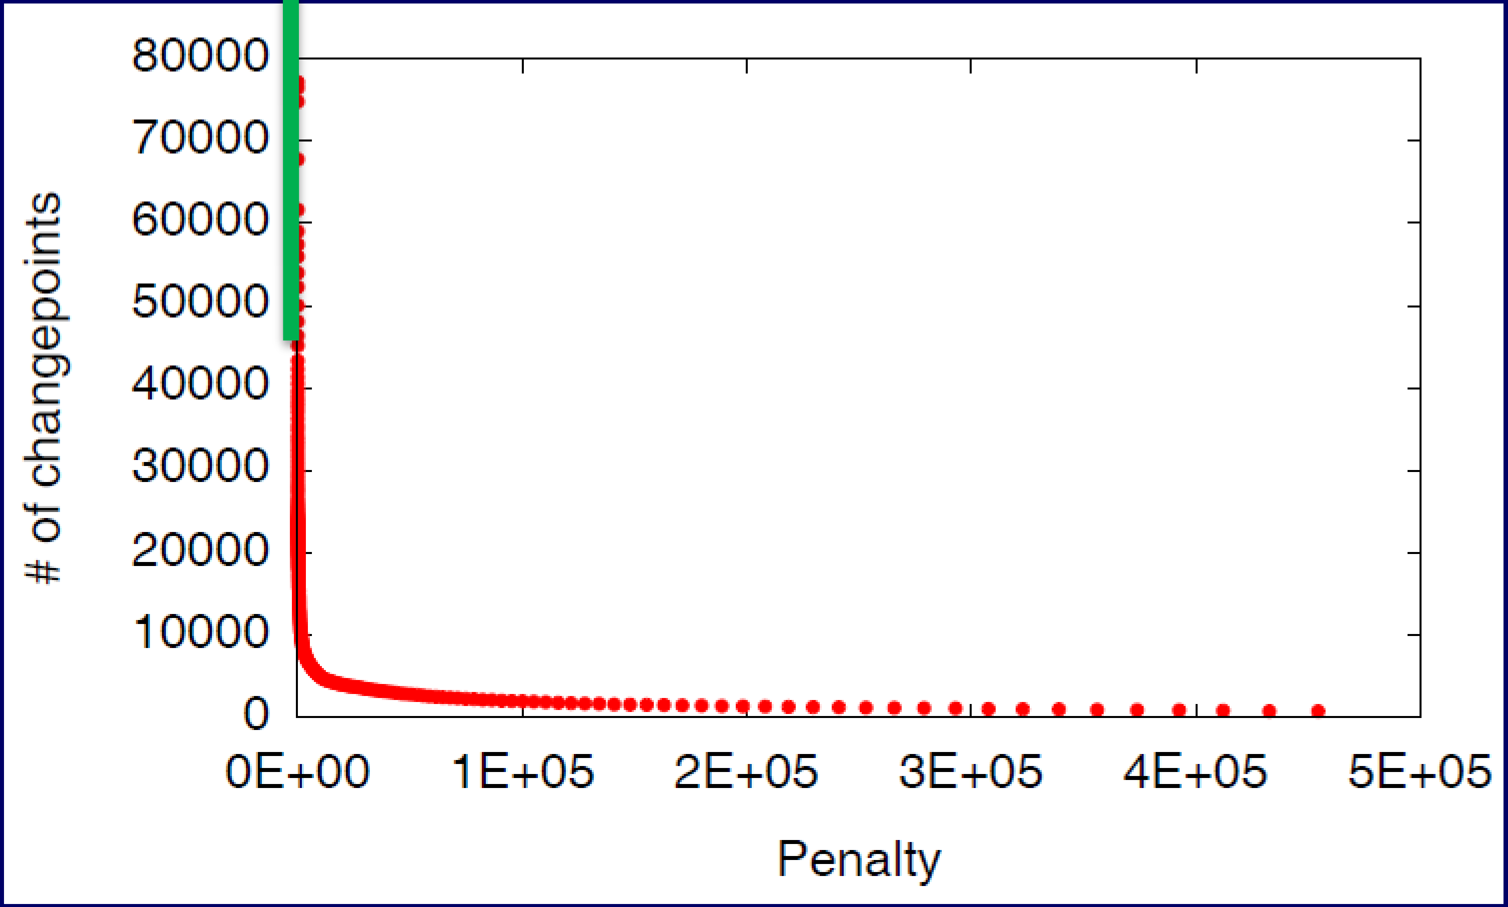

Methodology – Changepoint detection

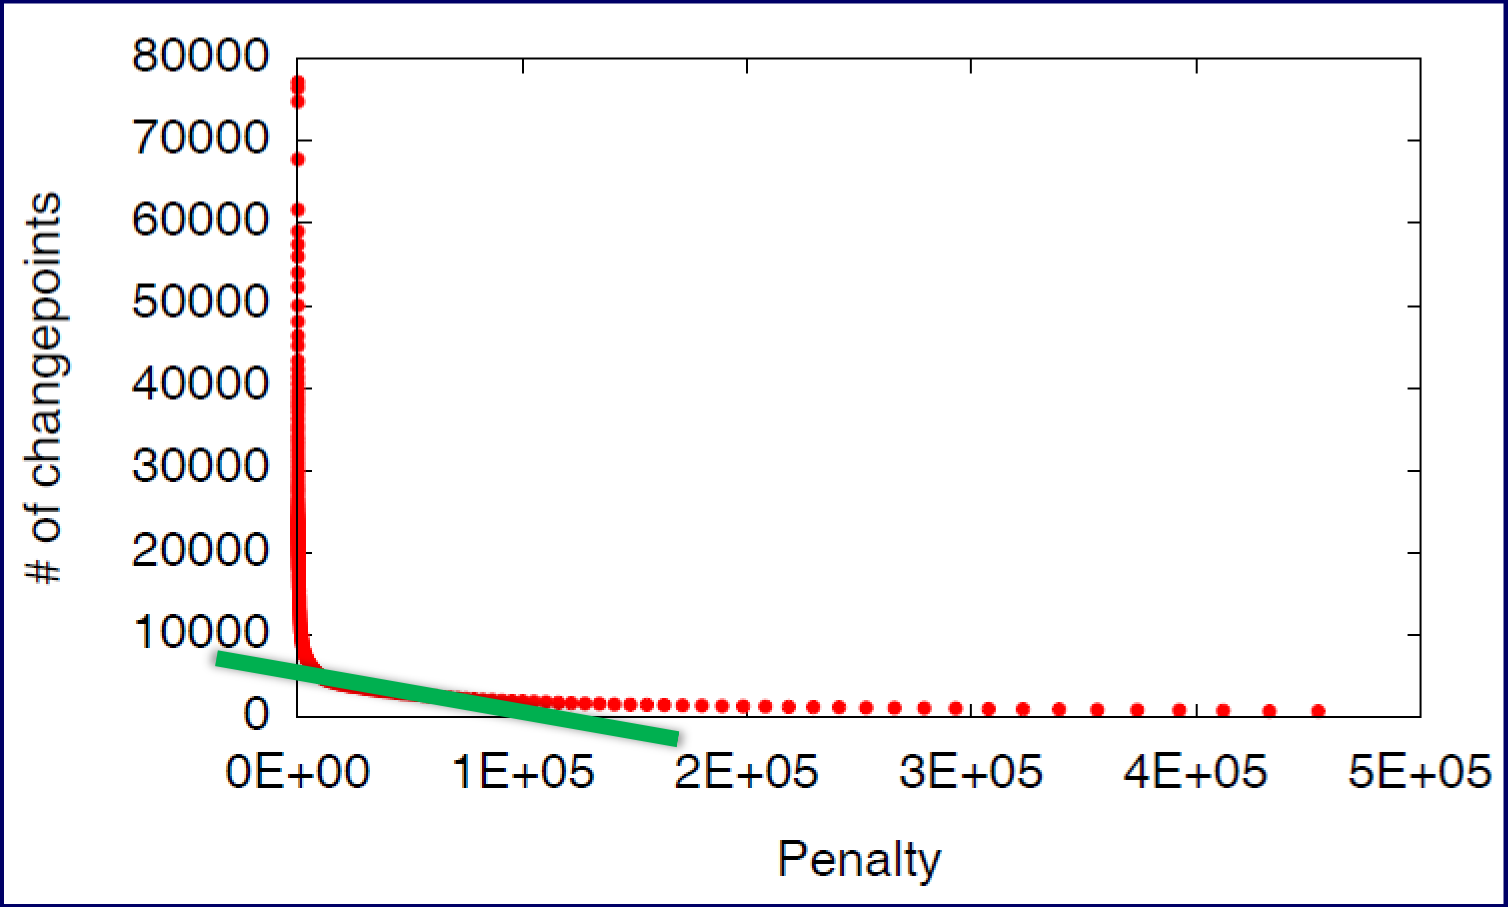

Methodology – Changepoint detection

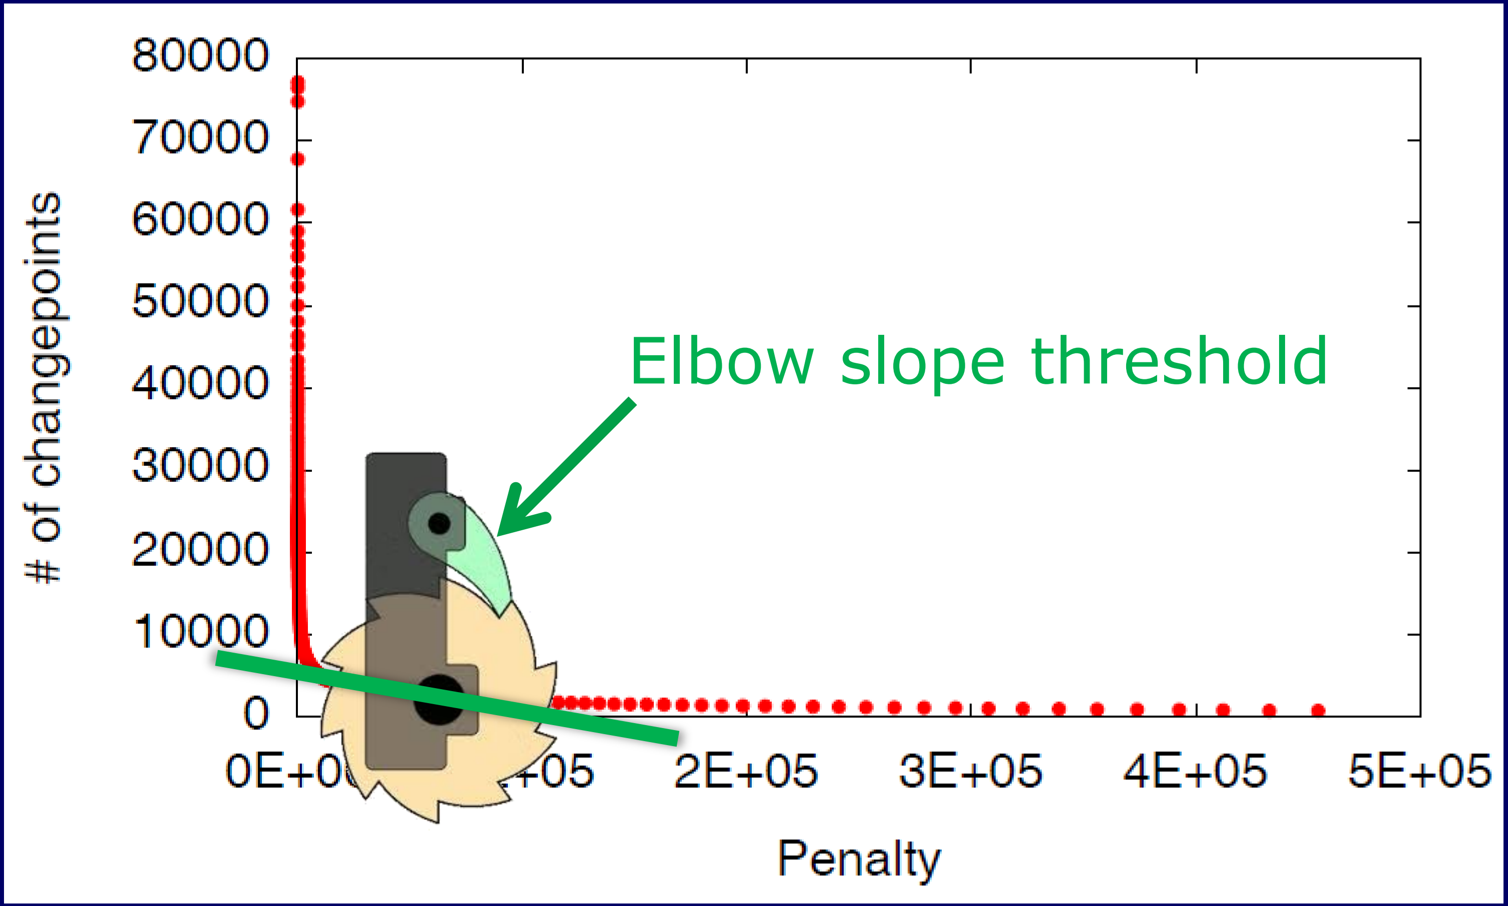

Methodology – Changepoint detection

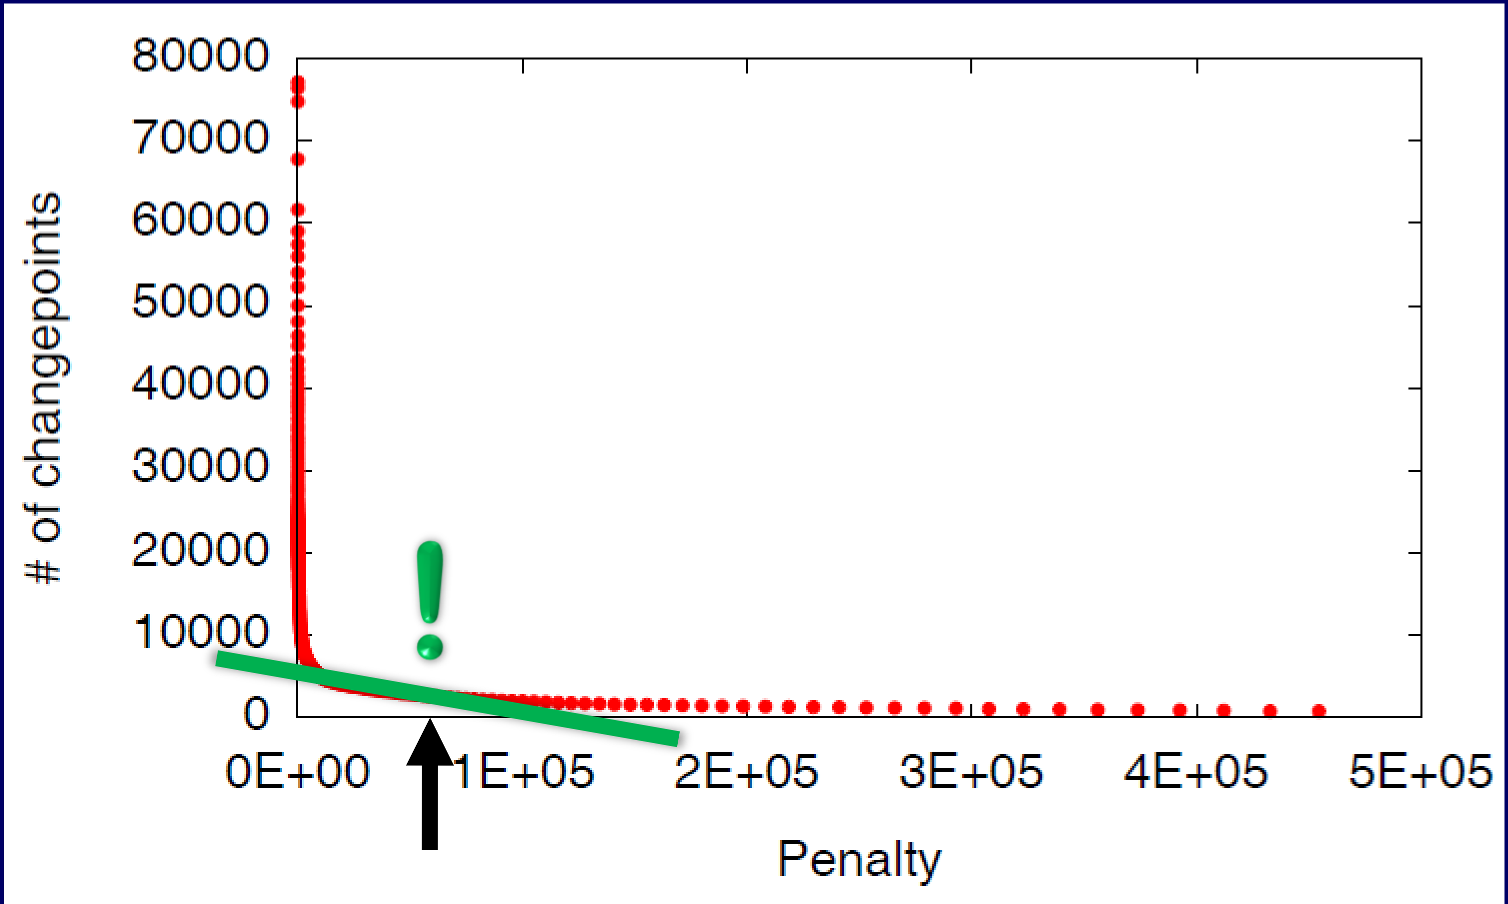

Methodology – Matching

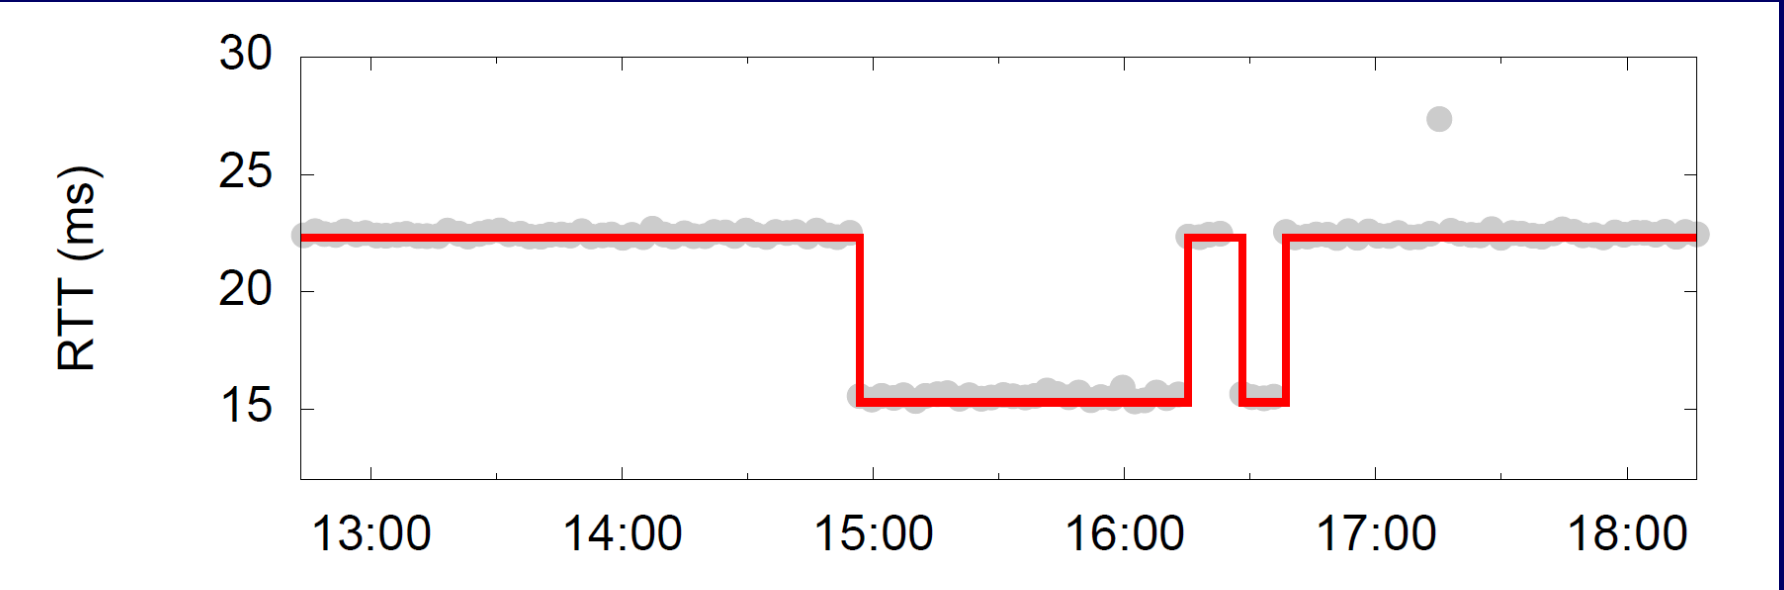

Methodology – Matching

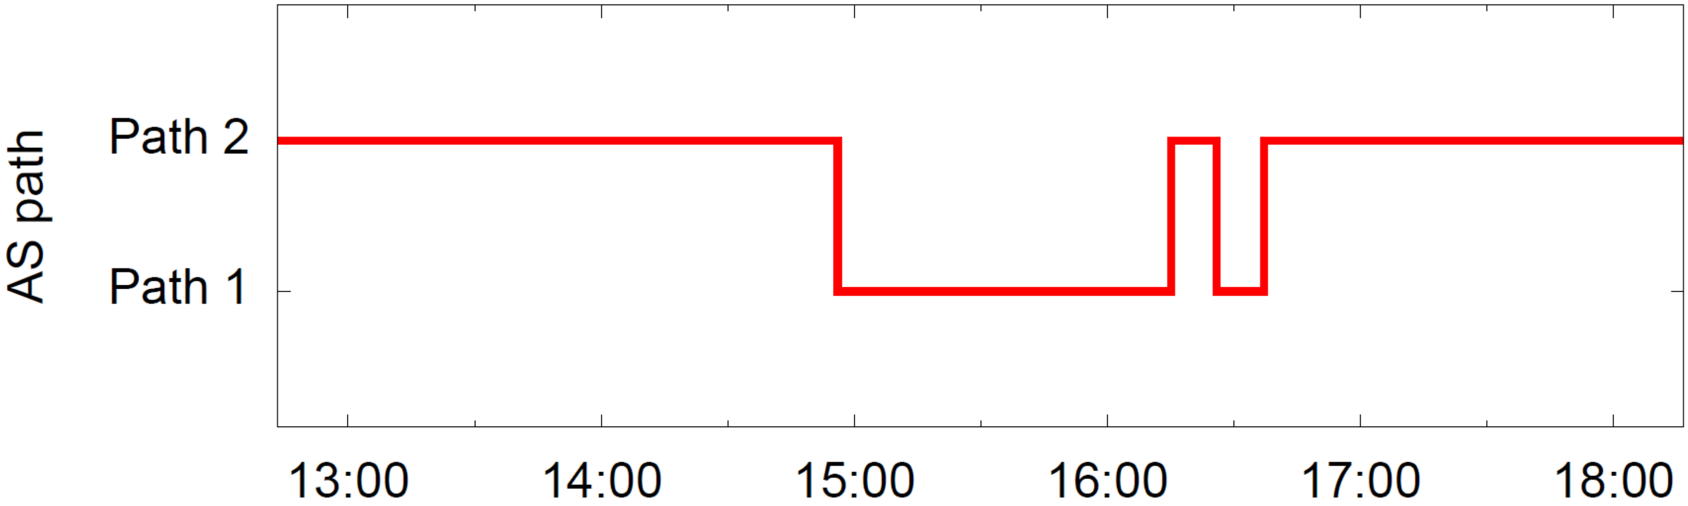

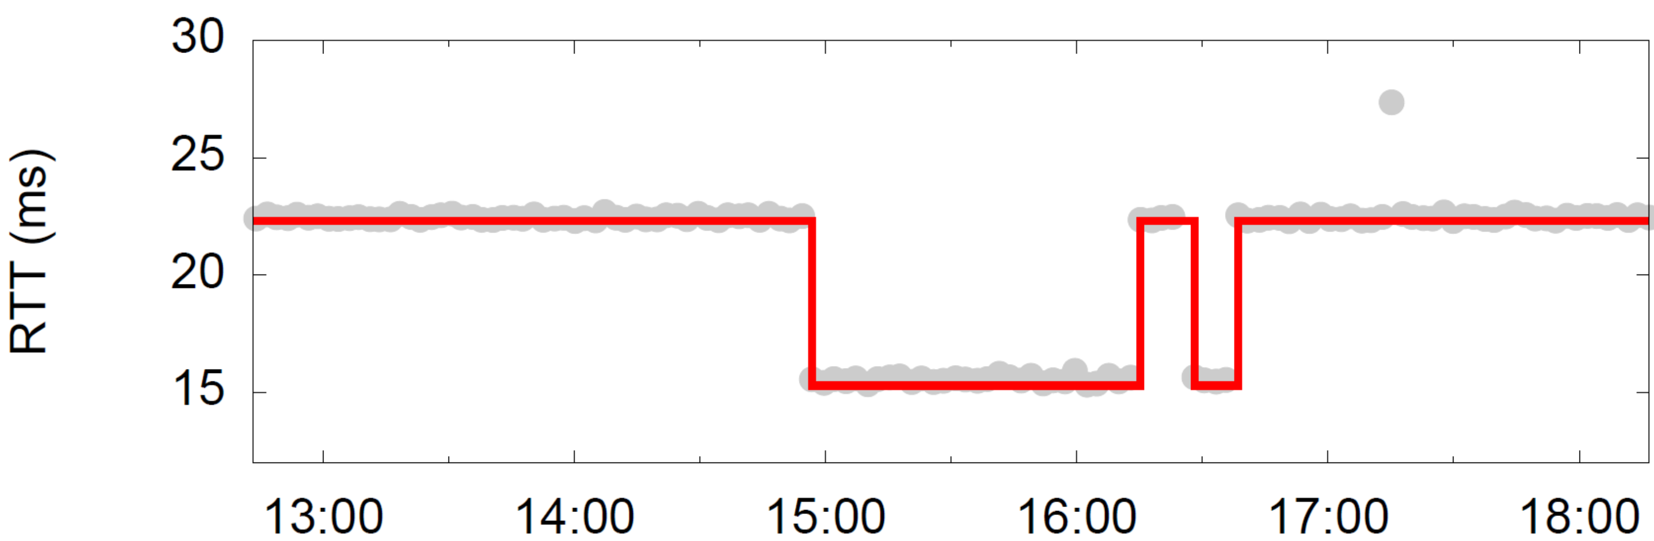

Methodology – Matching

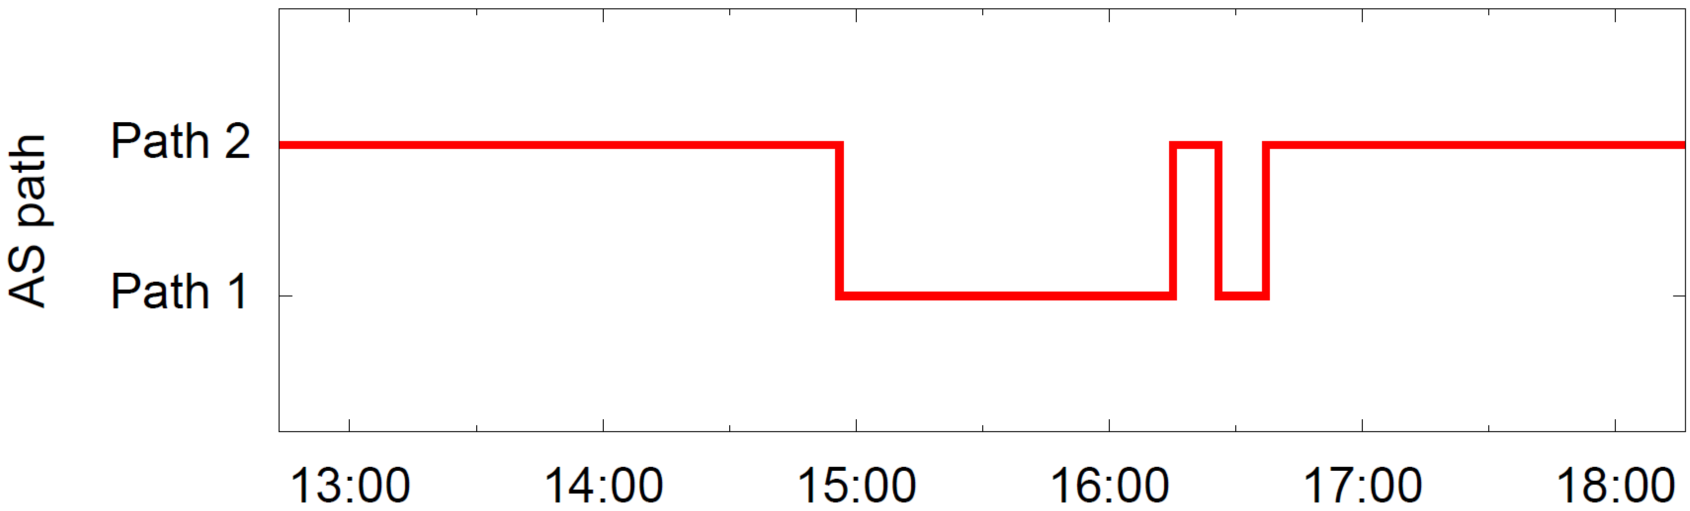

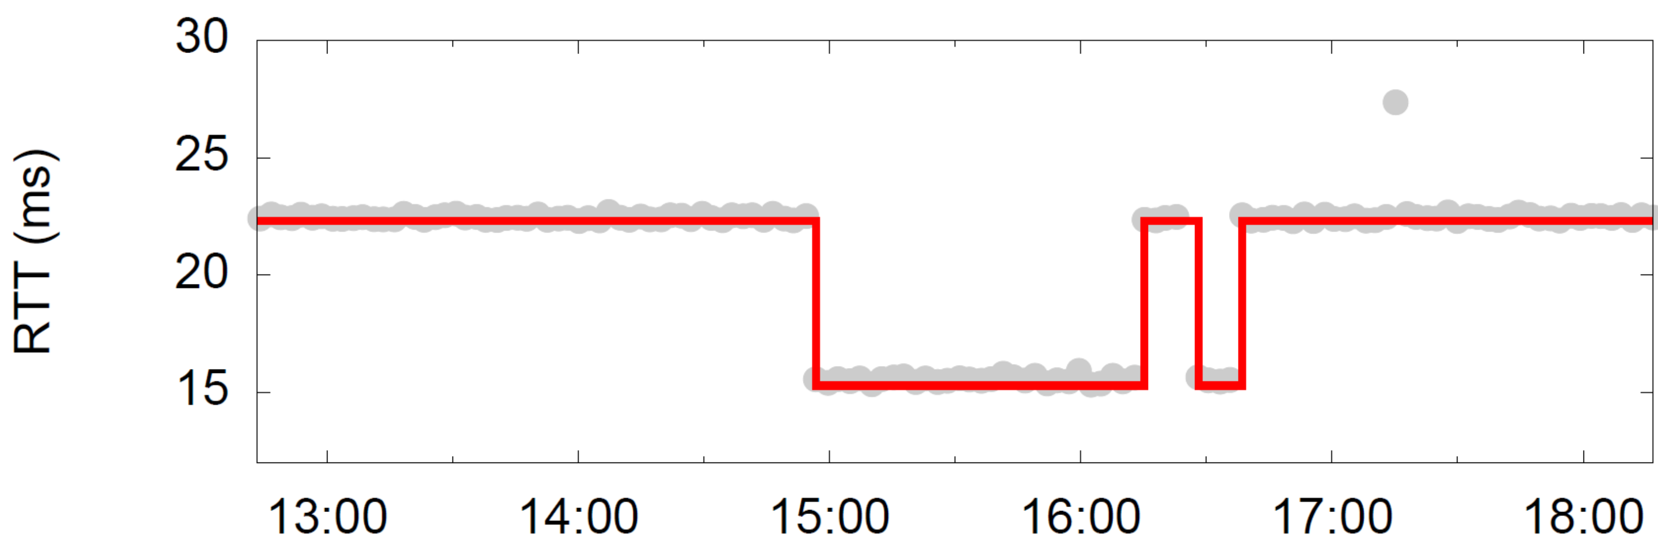

Methodology – Matching

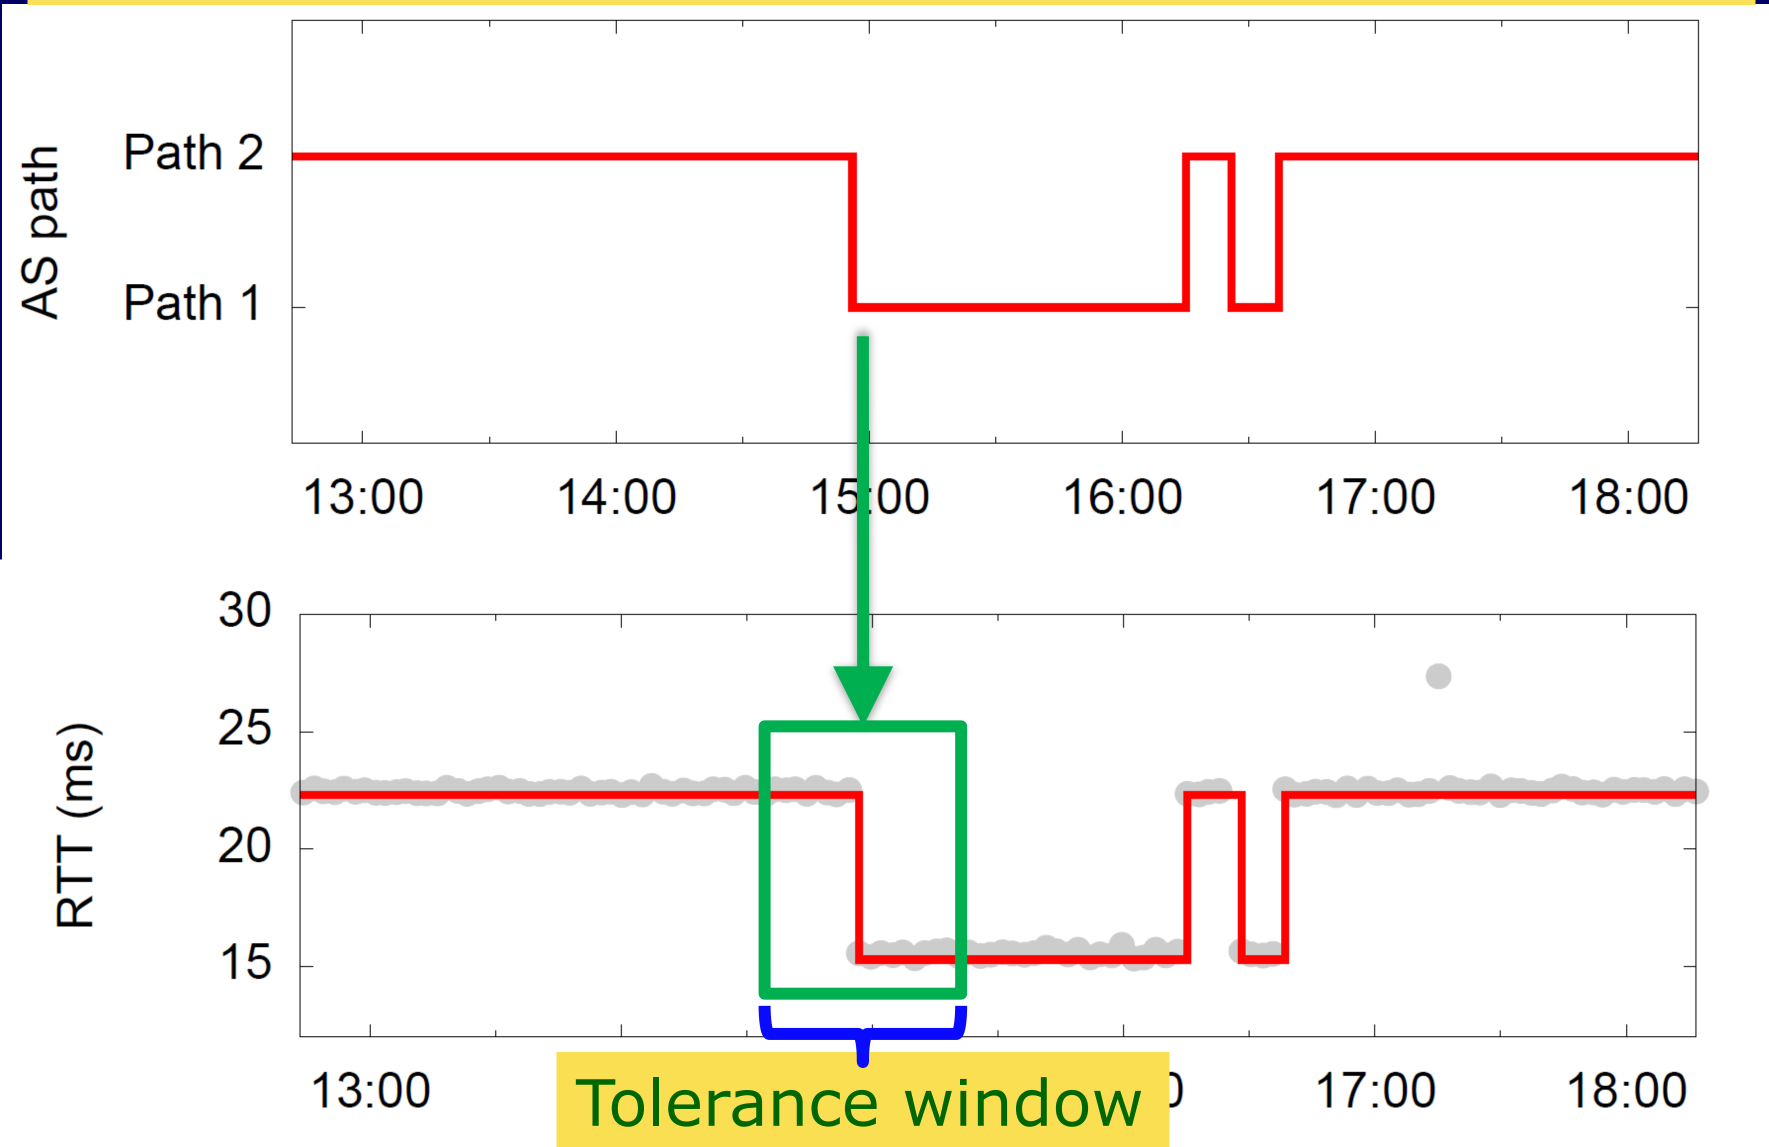

Methodology – Matching

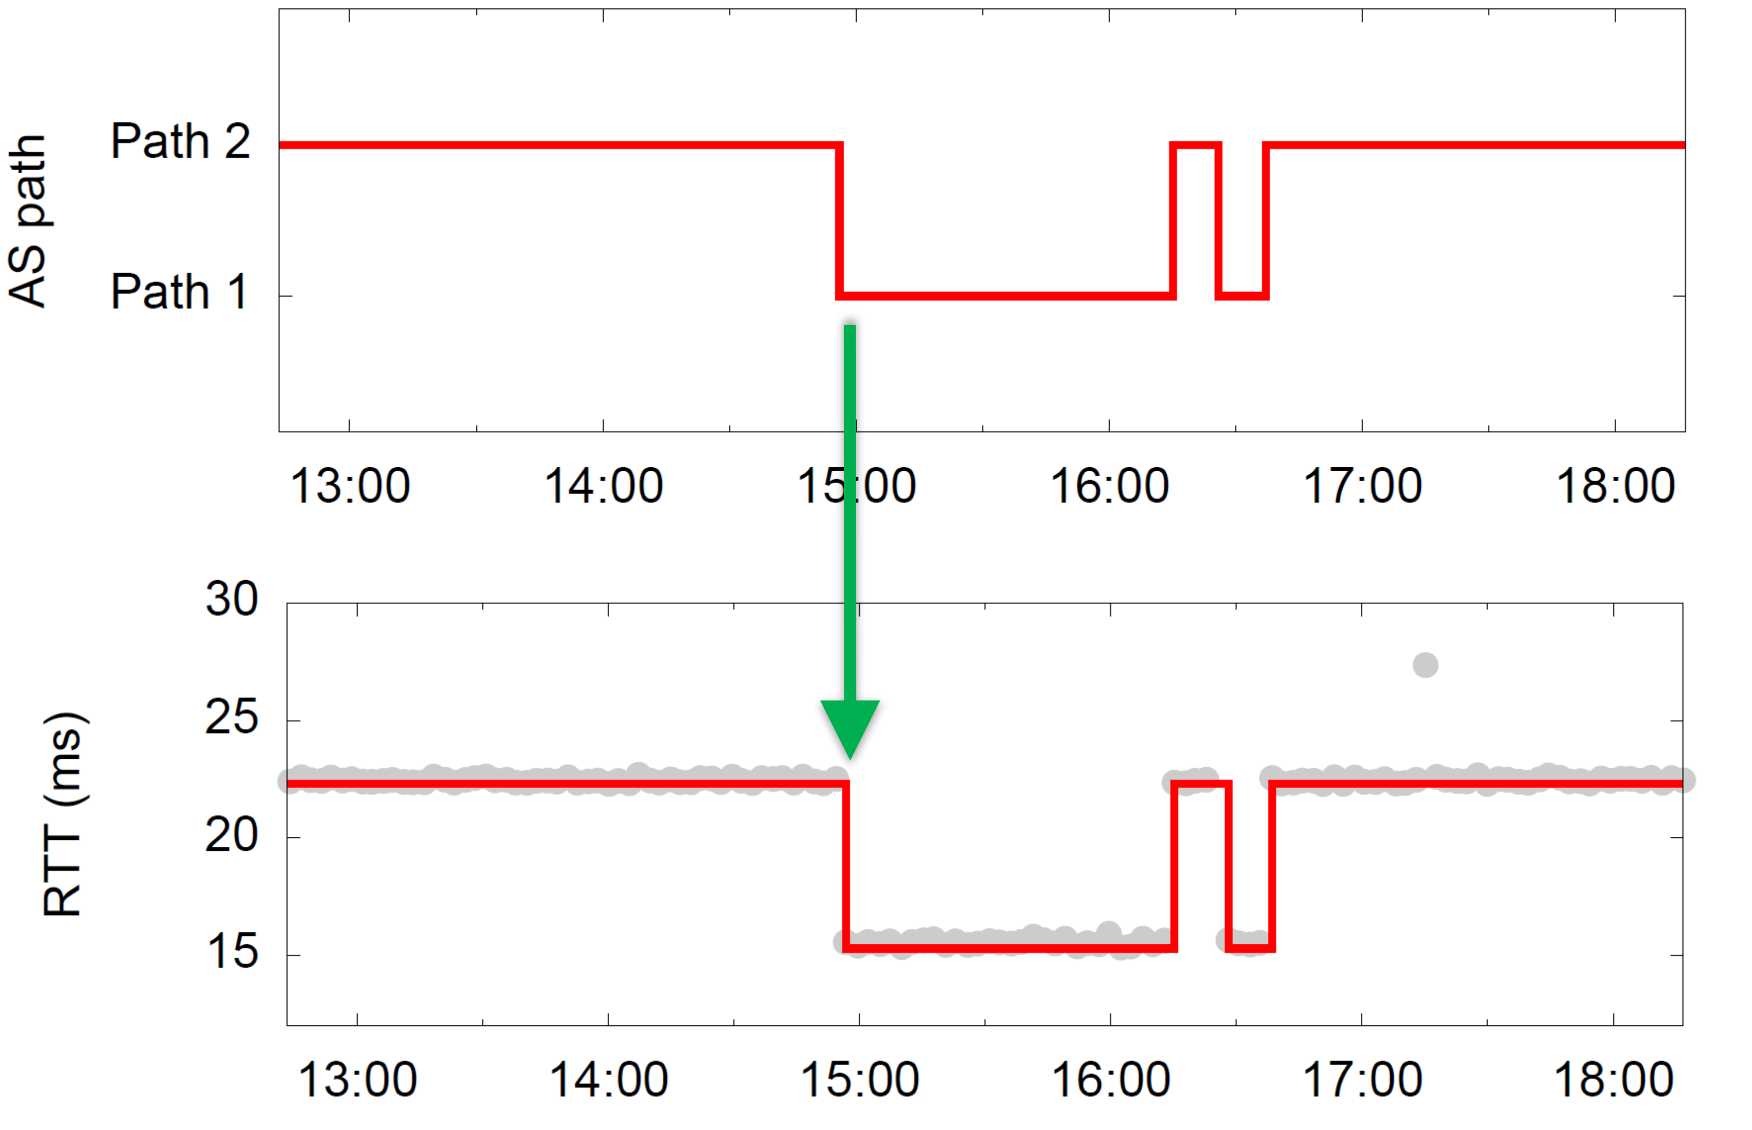

Methodology – Matching

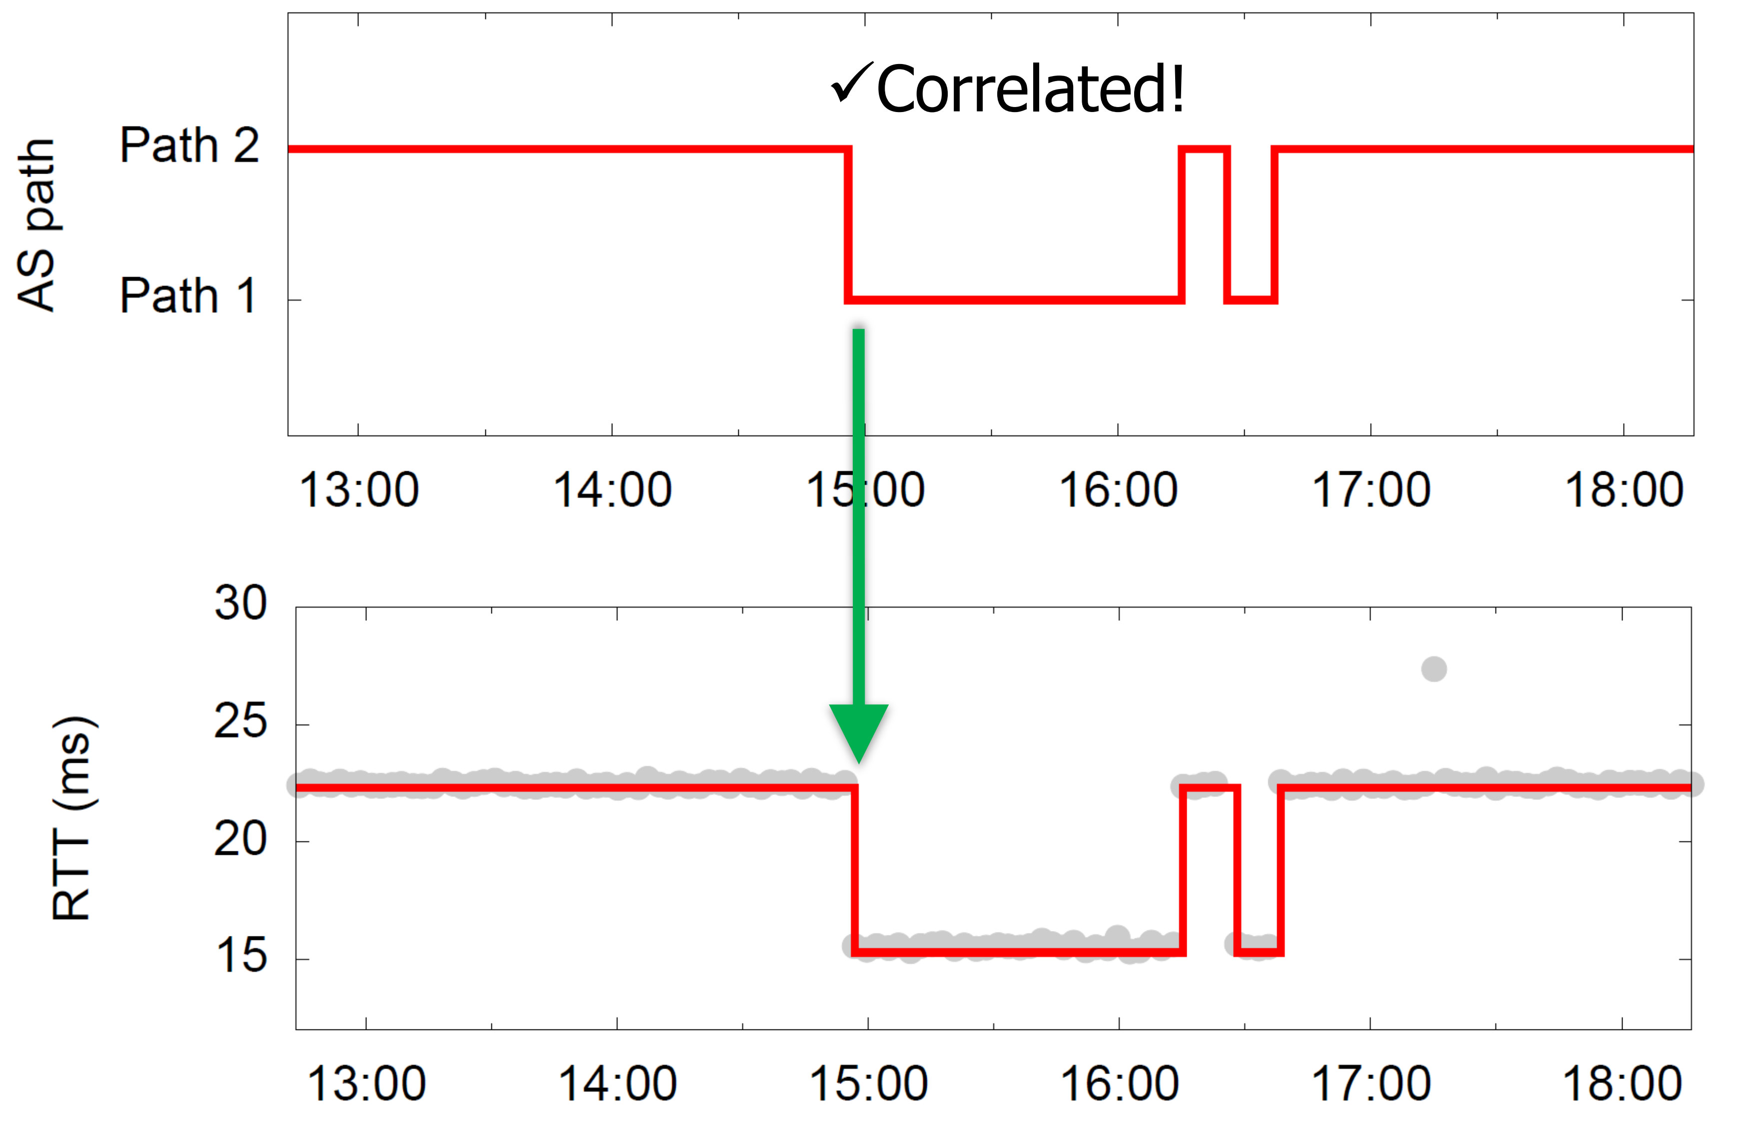

Methodology – Matching

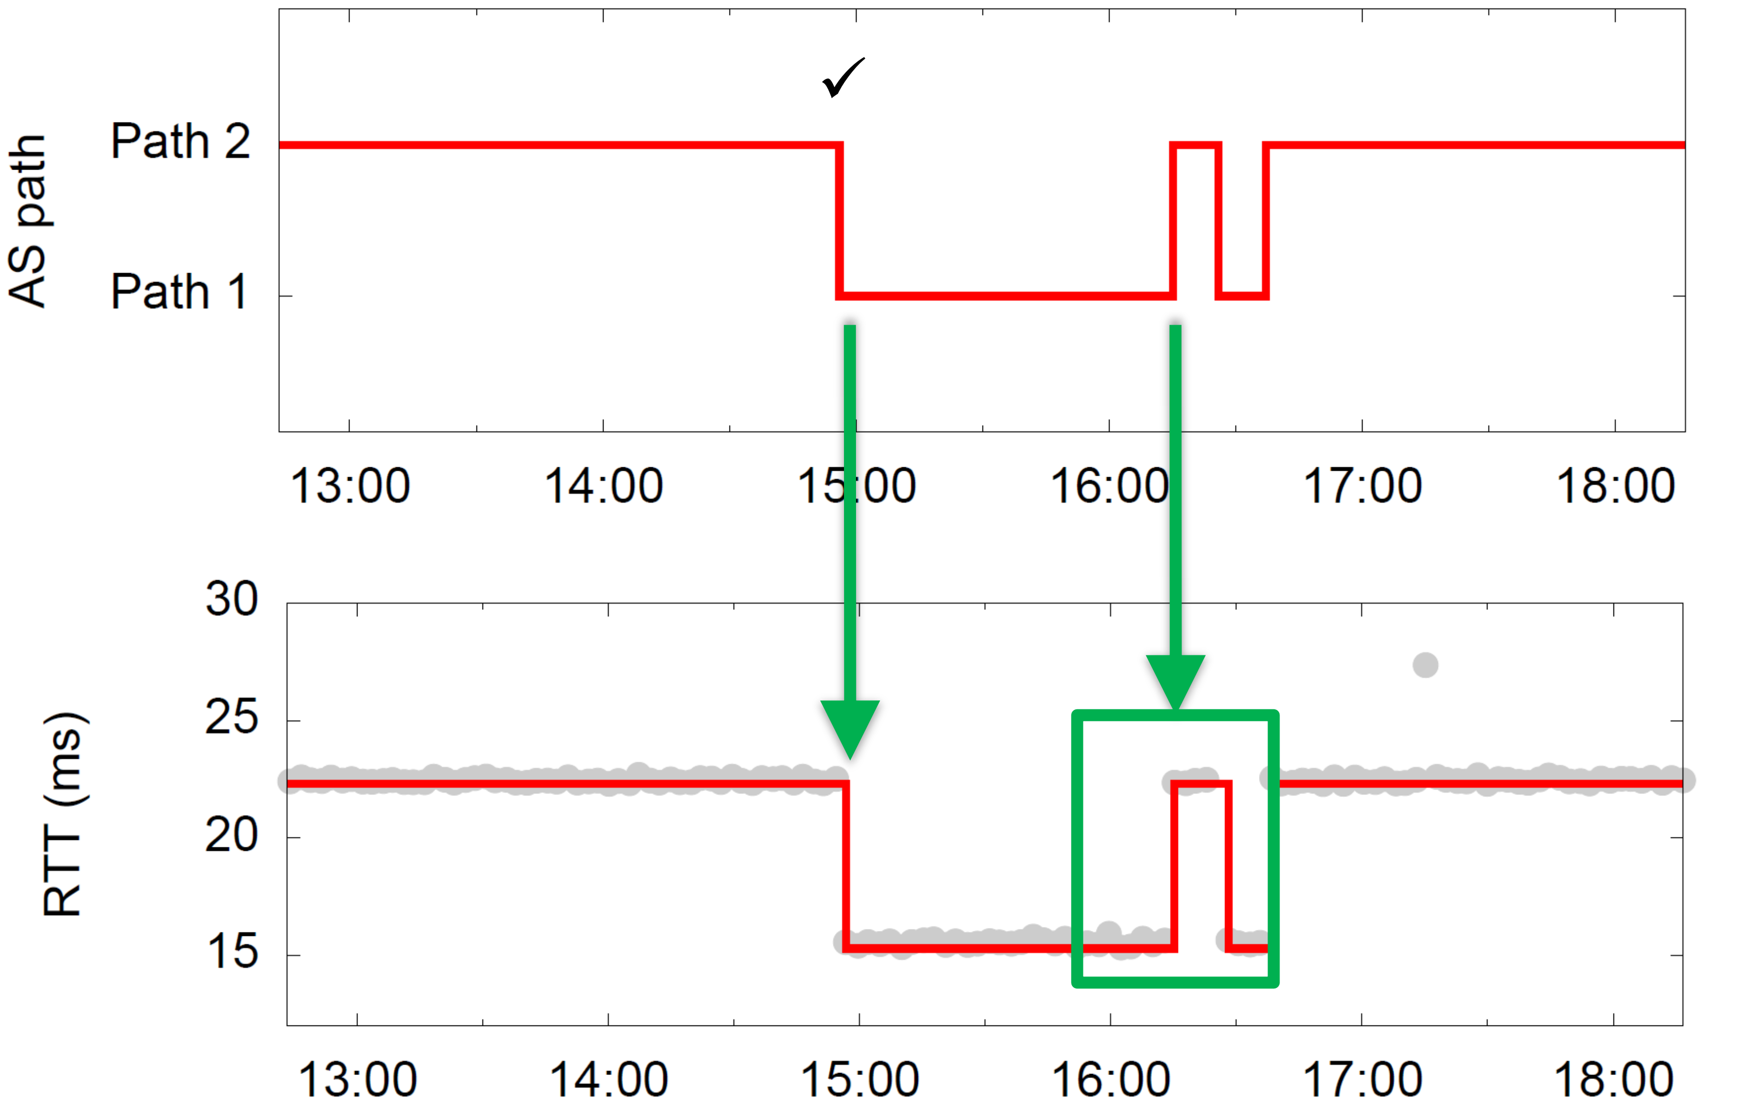

Methodology – Matching

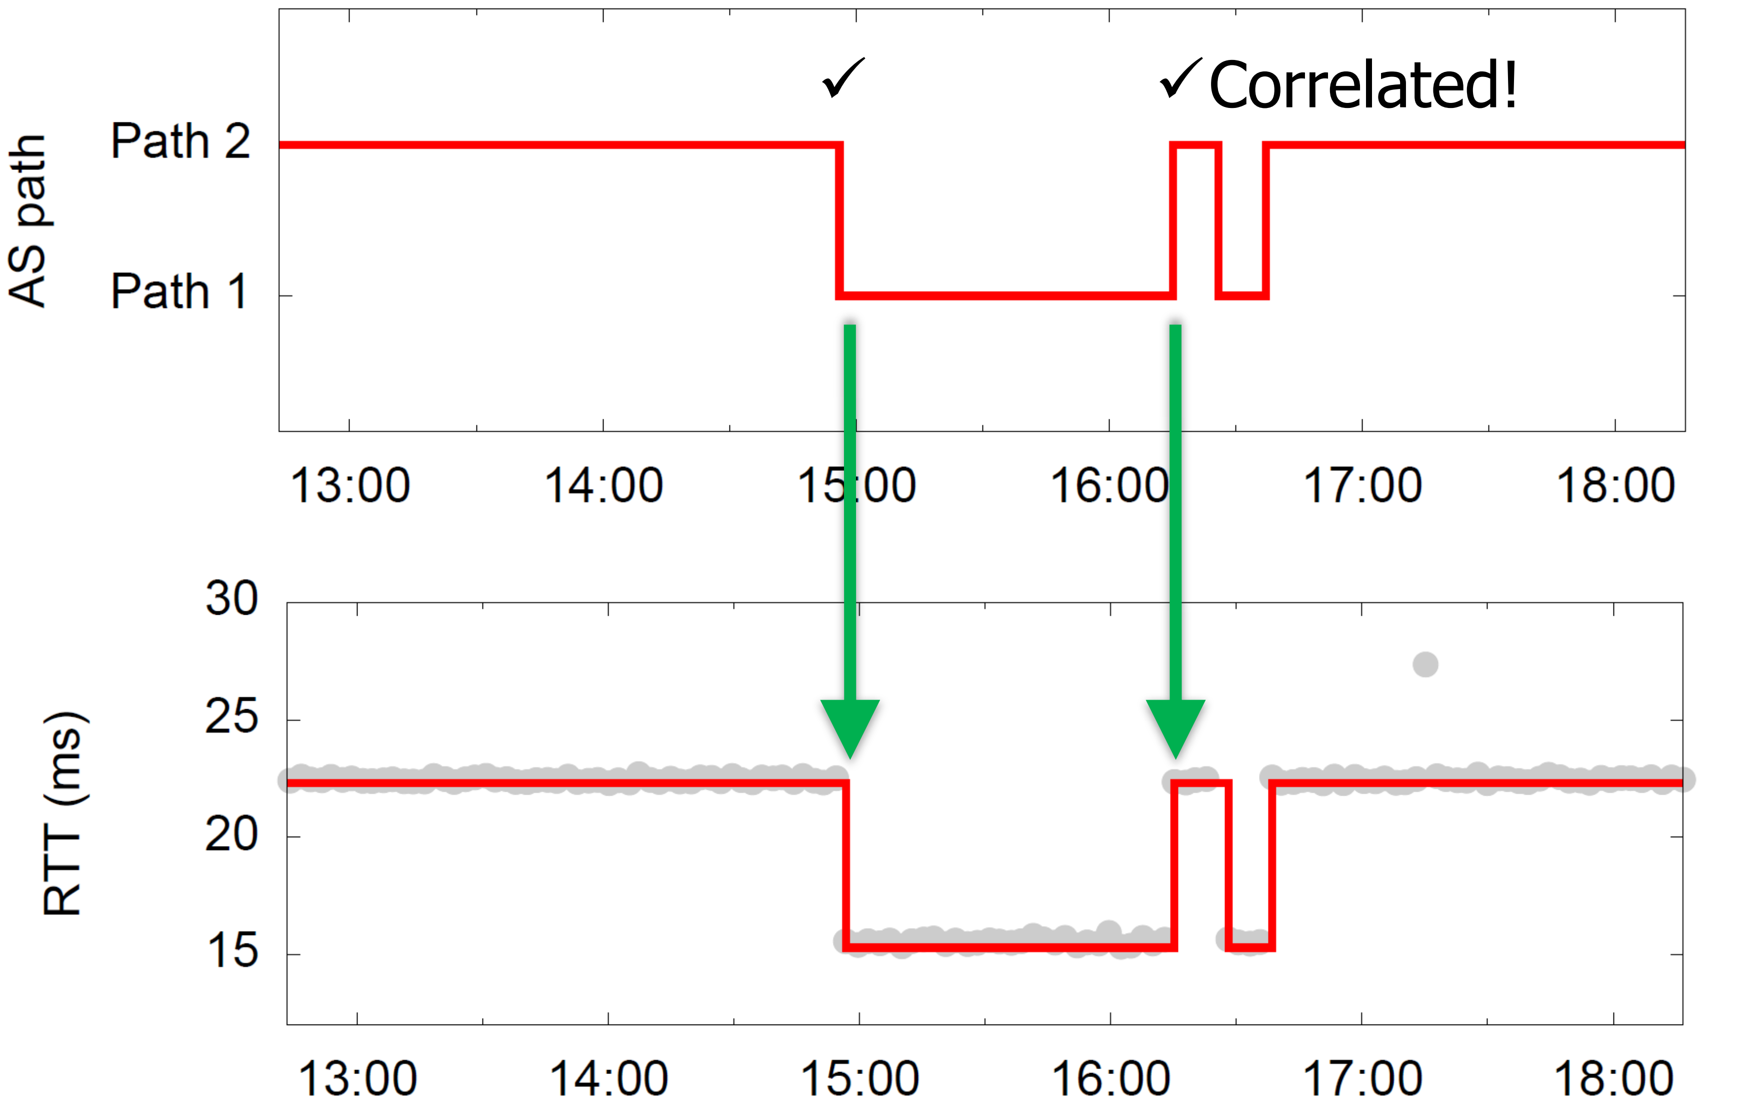

Methodology – Matching

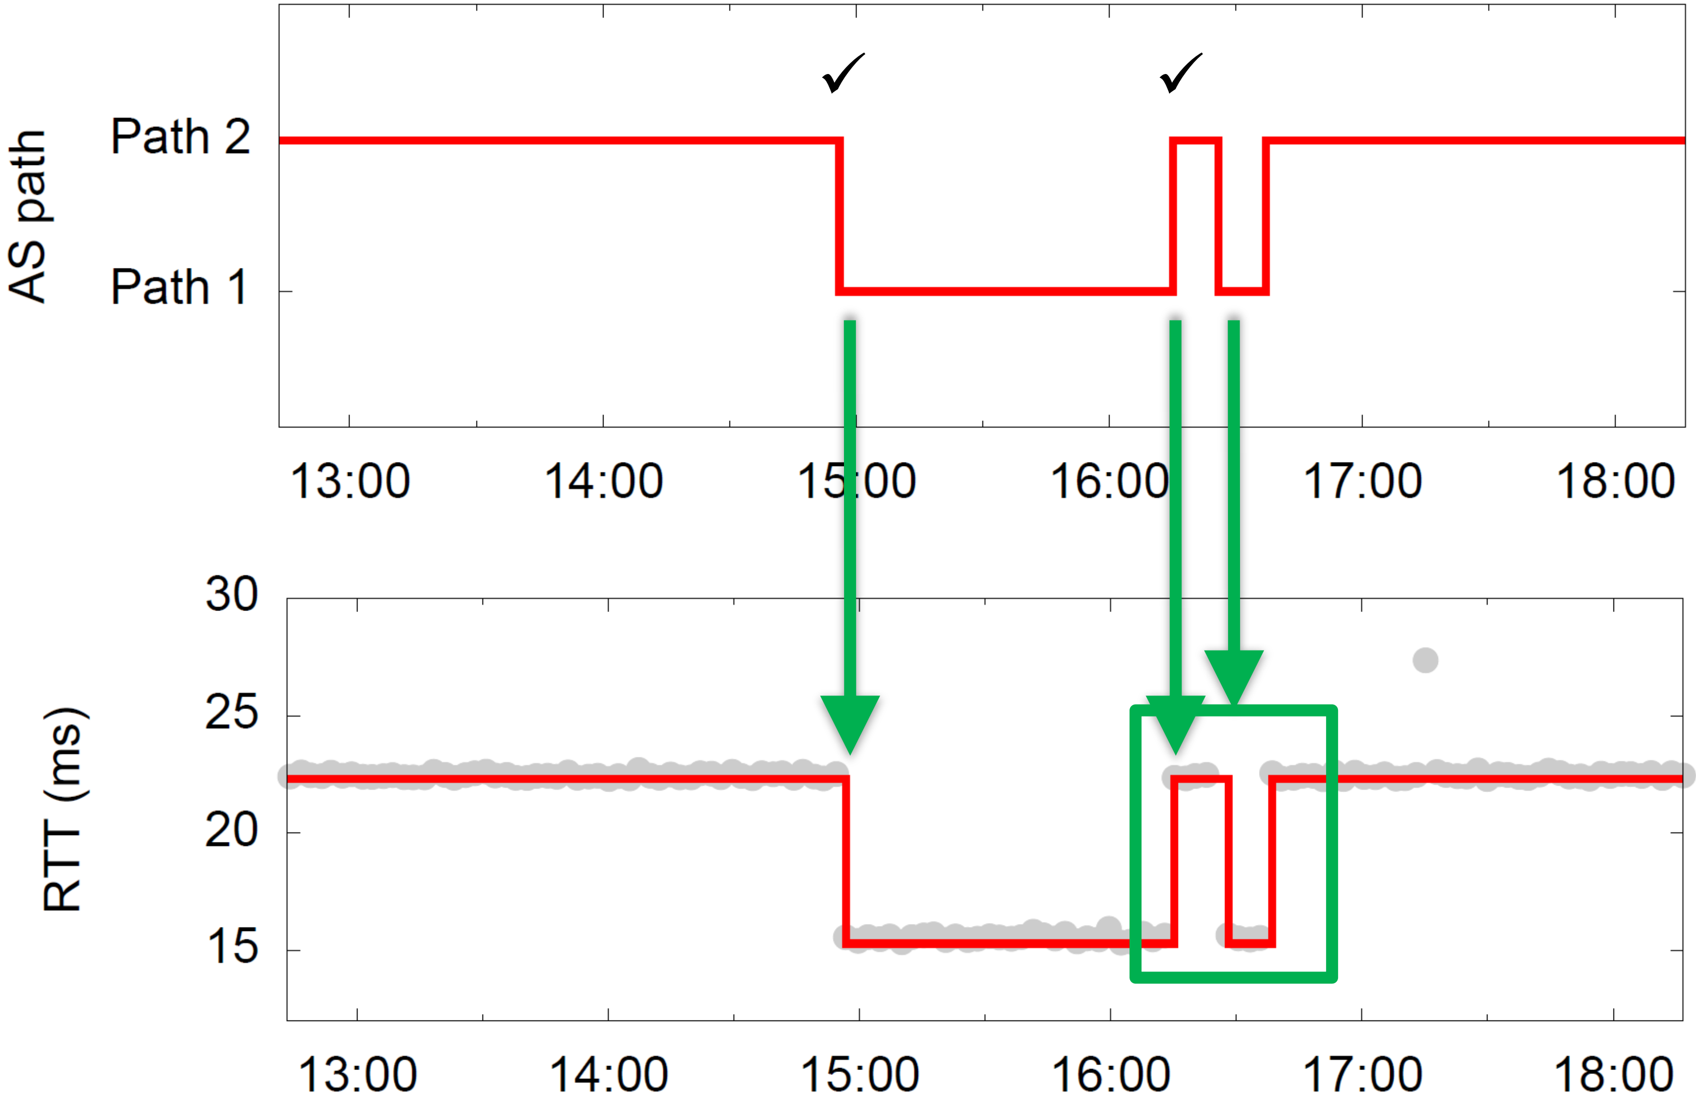

Methodology – Matching

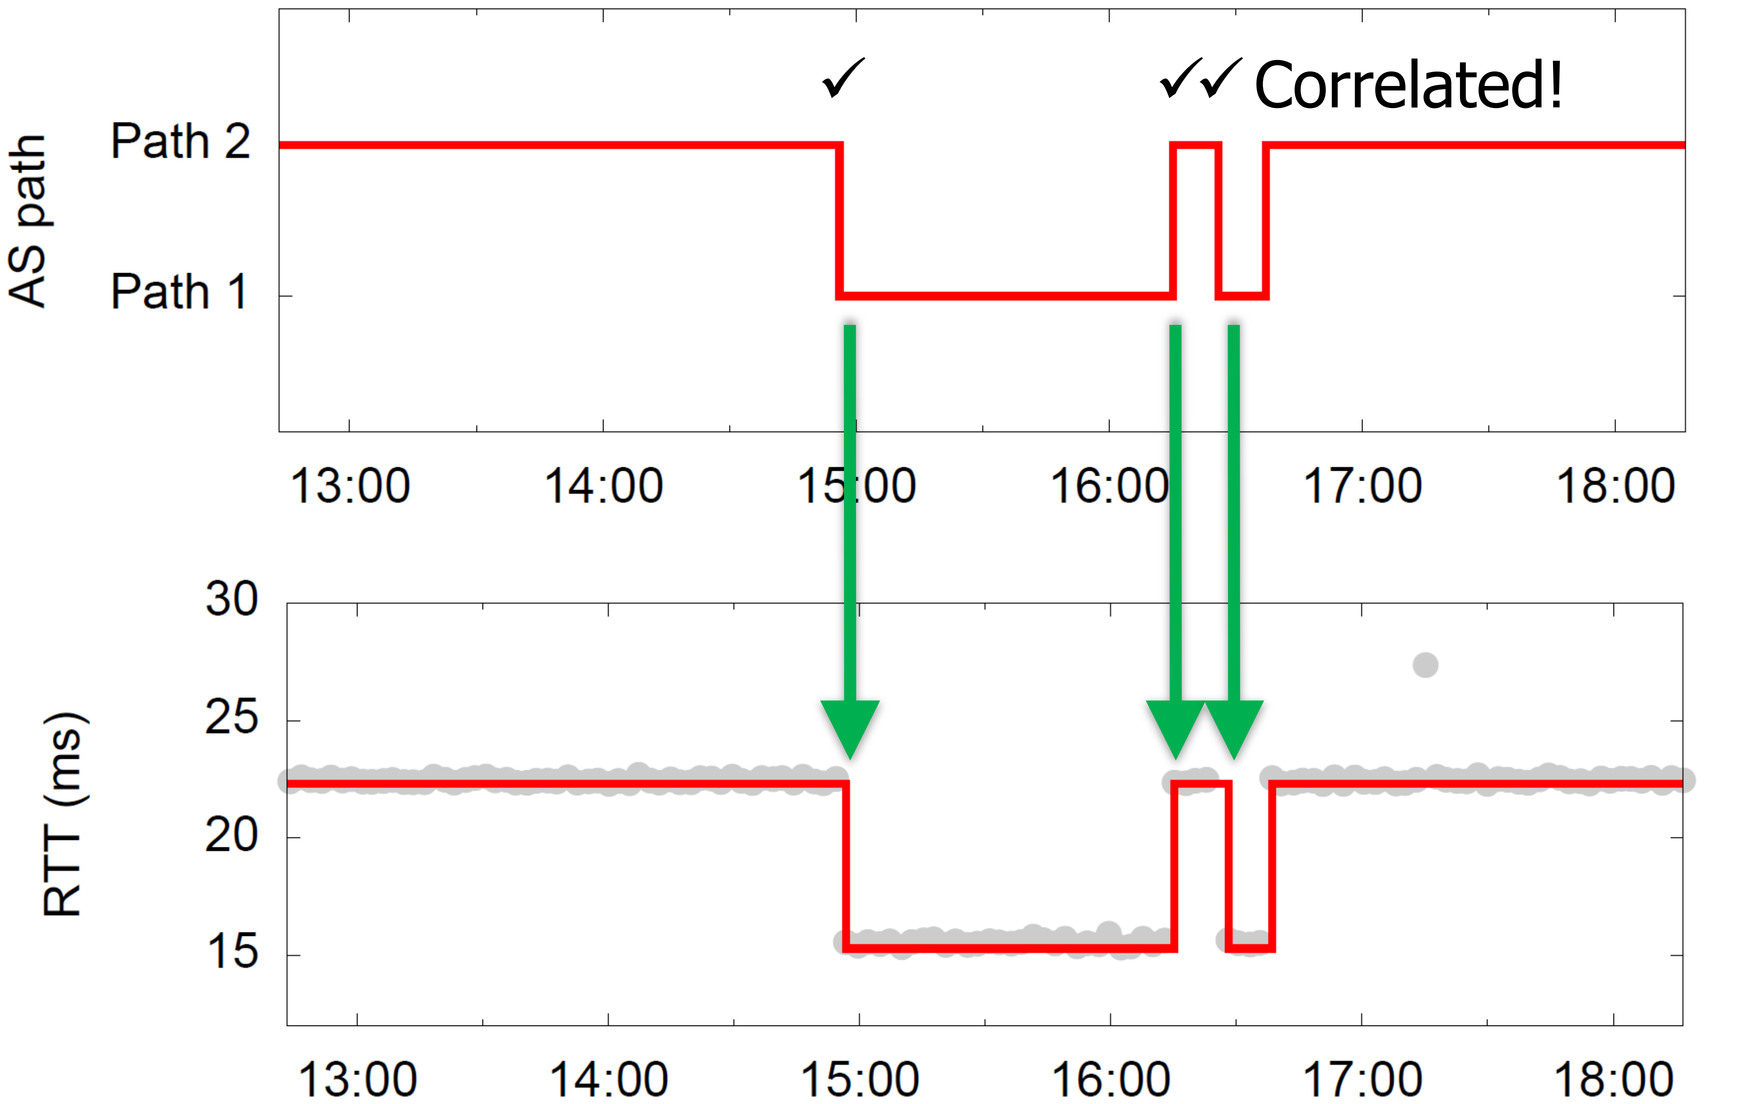

Methodology – Matching

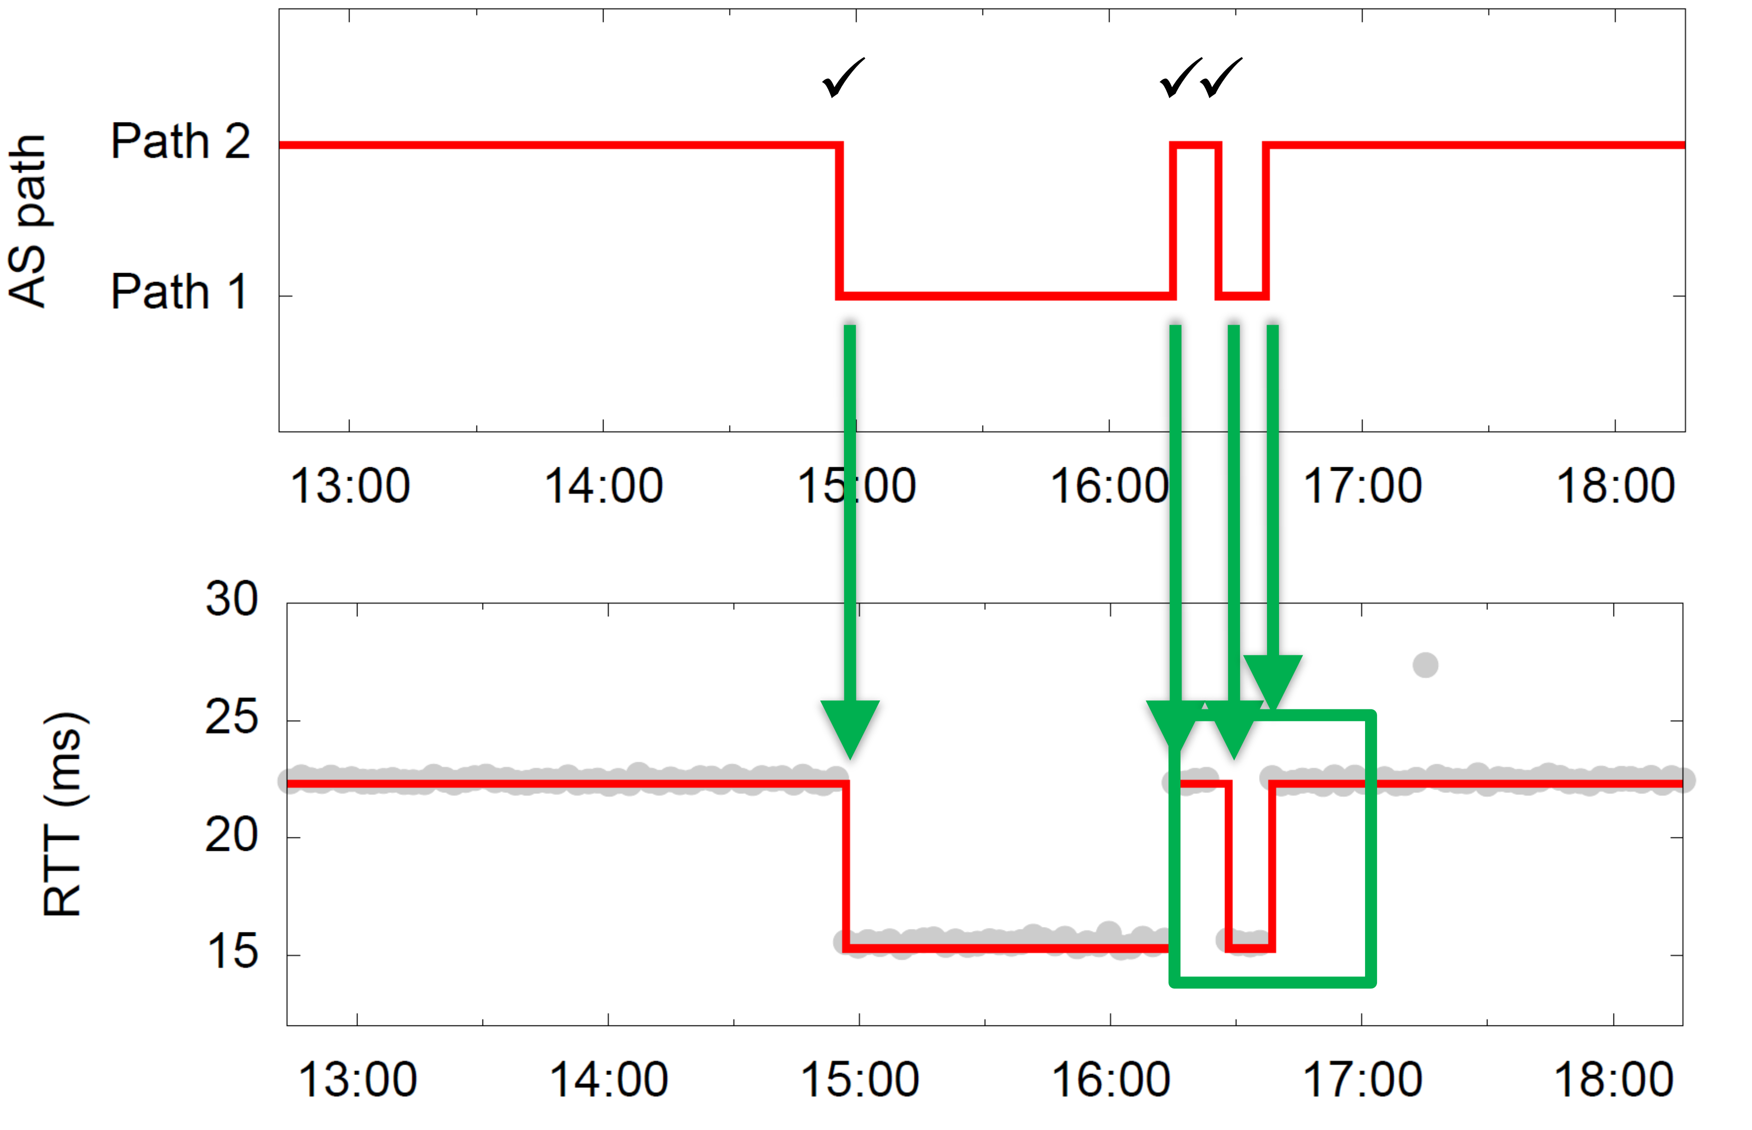

Methodology – Matching

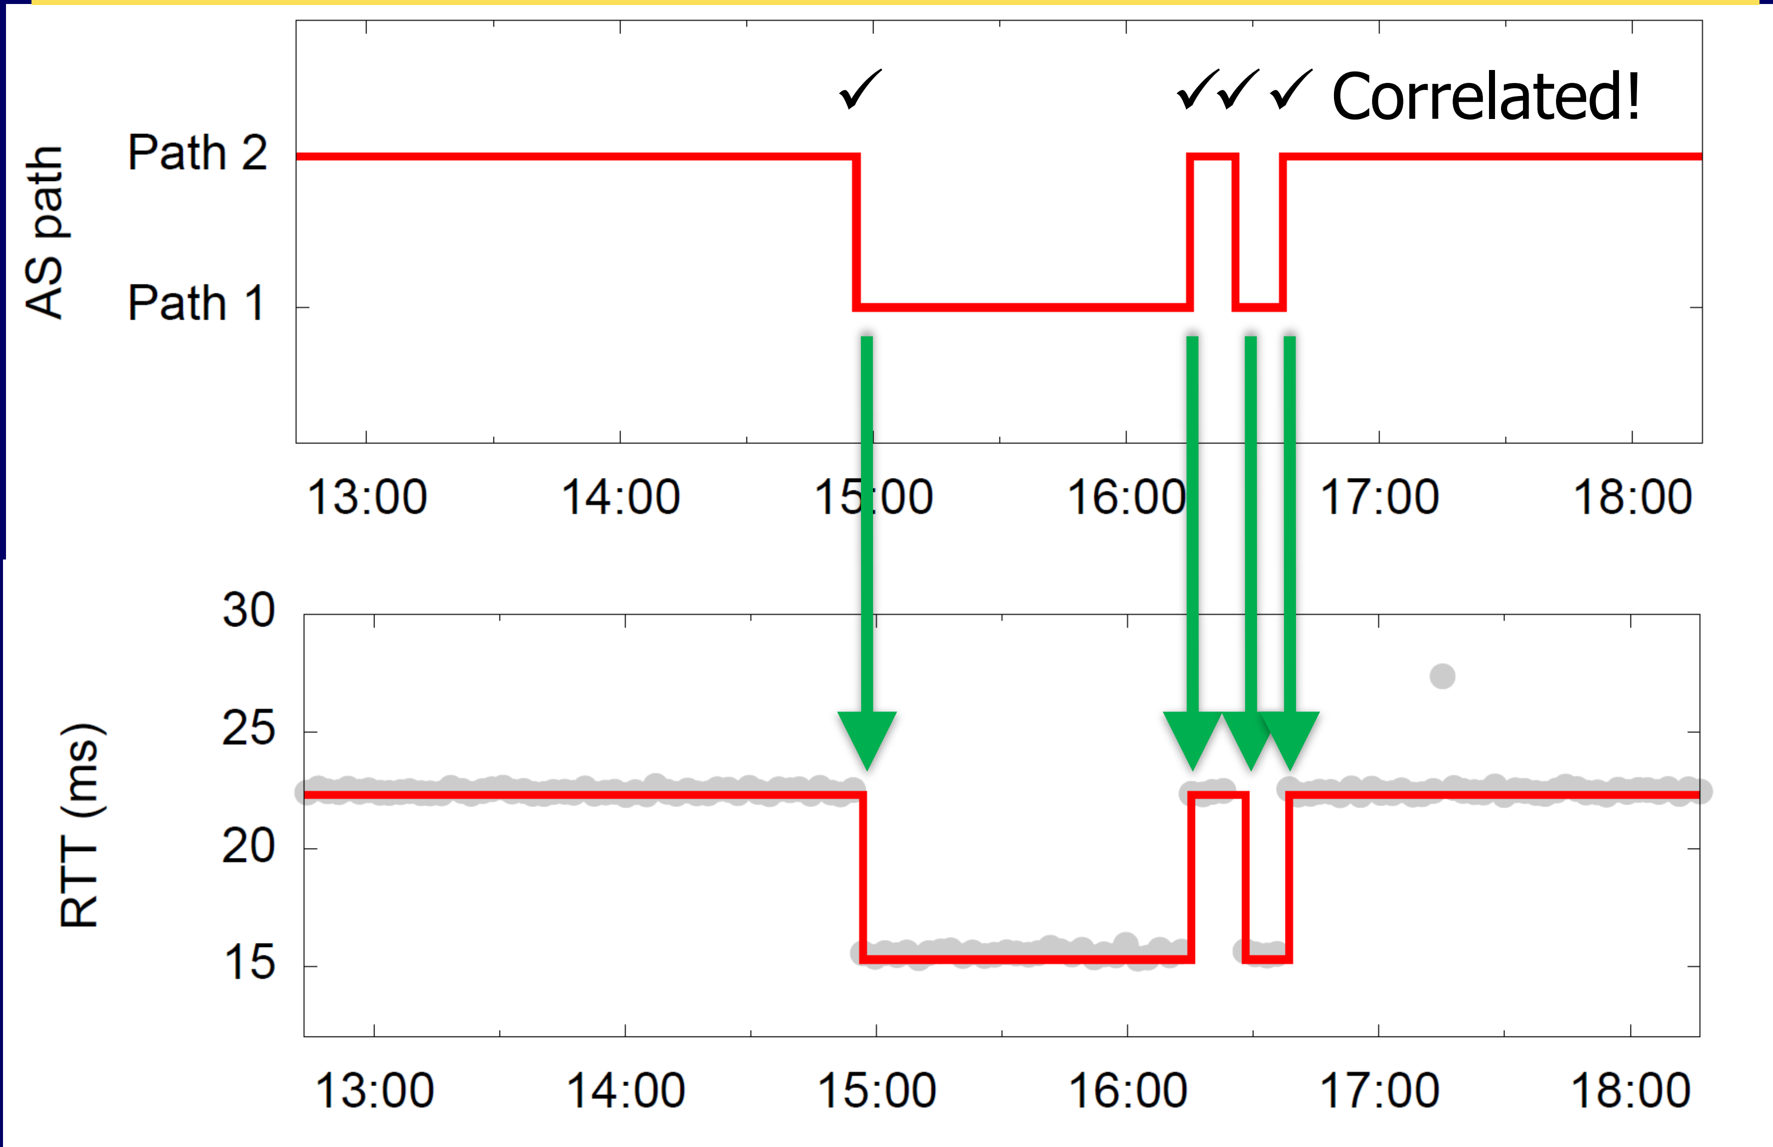

Methodology – Correlation

Correlation factor

$$cf := \frac{\#\text{ of correlated BGP updates}}{\#\text{ of BGP updates}}$$

Methodology – Correlation

Correlation factor

$$cf := \frac{\#\text{ of correlated BGP updates}}{\#\text{ of BGP updates}}$$

Methodology – Correlation

Correlation factor

$$cf := \frac{\#\text{ of correlated BGP updates}}{\#\text{ of BGP updates}}$$

Time window
Probe ID
Target

Raw RTT measurements

How much are they correlated?

Time window
Collector peer
Prefix

BGP updates

Methodology – Correlation

Correlation factor

$$cf := \frac{\text{\# of correlated BGP updates}}{\text{\# of BGP updates}}$$

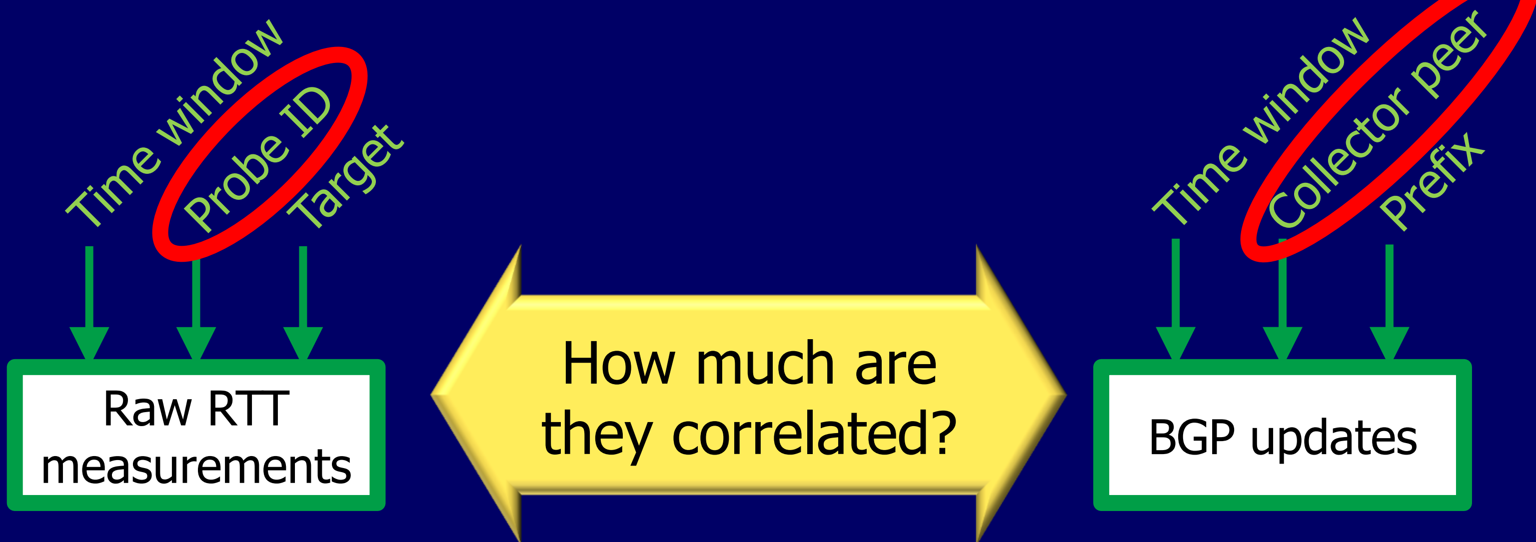

Methodology – Shortcomings

- We do not account for:
  - Routing changes on the reverse path
  - RTT biases
    - Clock sync
    - Load balancers
    - …

Methodology – Shortcomings

- We do not account for:
  - Routing changes on the reverse path
  - RTT biases
    - Clock sync
    - Load balancers
    - …

- Yet obtain interesting results!

Interesting Results

Experimental Results

Experimental Results

- Data sources

Experimental Results

- Data sources

Ripe Atlas

Experimental Results

- Data sources

SamKnows™

Experimental Results

- Data sources

FCC Measuring Broadband America

Experimental Results

- Data sources

MisuraInternet

Experimental Results

- Data sources

BISMark

Experimental Results

- Data sources

CAIDA Archipelago

Experimental Results

- Data sources

M-Lab

Experimental Results

- Data sources

Route Views

Experimental Results

- Data sources

RIPE RIS

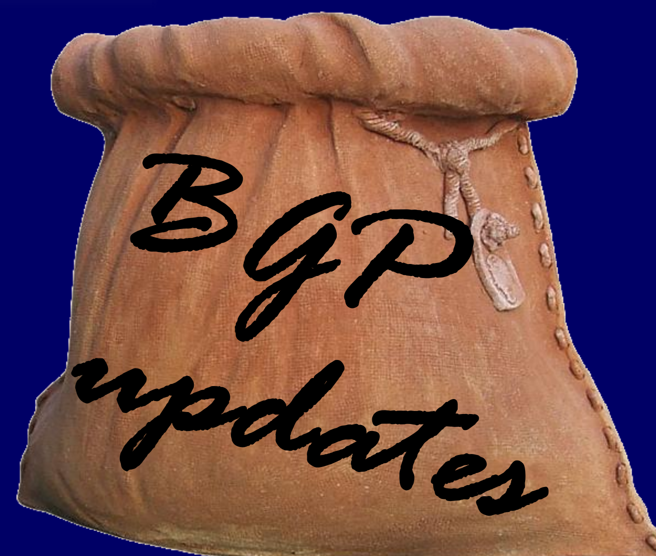

Experimental Results

- Data sources

Experimental Results

- Data sources

Ripe Atlas

Experimental Results

- Data sources

Ripe Atlas

RIPE NCC

RIPE RIS

Experimental Results

- Data sources
  - Jan 2013 footprint:
    - 55 ASes
    - 126 CPs
    - 200 probes

Ripe Atlas

RIPE RIS

Experimental Results

- Data sources
  - Jan 2013 footprint:
    - 55 ASes
    - 126 CPs
    - 200 probes

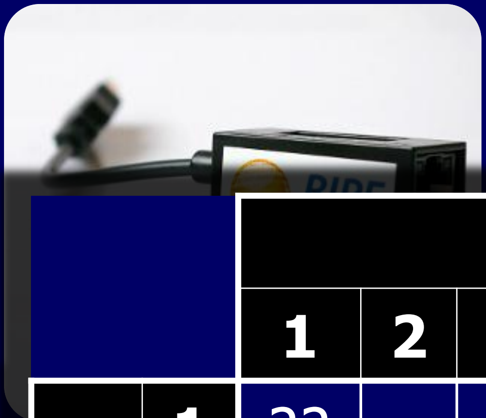

# of CPs		# of probes							
		1	2	3	4	5	7	13	22
	1	22		1	1				
	2	12	3	3	2				1
	3	1			1	1			
	4	1	1						
	5	1						1	
	6	1					1		
	7			1					

Experimental Results

- Data sources
  - Jan 2013 footprint:
    - 55 ASes
    - 126 CPs
    - 200 probes

Ripe Atlas

RIPE RIS

Experimental Results

- Data sources
  - Jan 2013 footprint:
    - 55 ASes
    - 126 CPs
    - 200 probes

  - Time window = 2 years (Jan 2011-Dec 2012)

Experimental Results

- Data sources
  - Jan 2013 footprint:
    - 55 ASes
    - 126 CPs
    - 200 probes

  - Time window = 2 years (Jan 2011-Dec 2012)
    - 23 Targets
    - one RTT every 4 minutes

Experimental Results

- Targets

ID	Target IP	BGP prefix	% of AS paths of len ≤5 (global avg.)	% of probes with avg. RTT ≤300ms
1001	193.0.14.129 `k.root-servers.net`	193.0.14.0/24 (Anycast)	87.5%	99.5%
1003	193.0.0.193 `ns.ripe.net`	193.0.0.0/21 (Unicast)	87.2%	97.3%
1004	192.5.5.241 `f.root-servers.net`	192.5.5.0/24 (Anycast)	57.8%	100%
1005	192.36.148.17 `i.root-servers.net`	192.36.148.0/24 (Anycast)	55.5%	99.1%

Experimental Results

- Targets

ID	Target IP	BGP prefix	% of AS paths of len ≤5 (global avg.)	% of probes with avg. RTT ≤300ms
1001	193.0.14.129 `k.root-servers.net`	193.0.14.0/24 (Anycast)	87.5%	99.5%
1003	193.0.0.193 `ns.ripe.net`	193.0.0.0/21 (Unicast)	87.2%	97.3%
1004	192.5.5.241 `f.root-servers.net`	192.5.5.0/24 (Anycast)	57.8%	100%
1005	192.36.148.17 `i.root-servers.net`	192.36.148.0/24 (Anycast)	55.5%	99.1%

- Choice driven by... data availability!

Back to Correlation

Back to Correlation

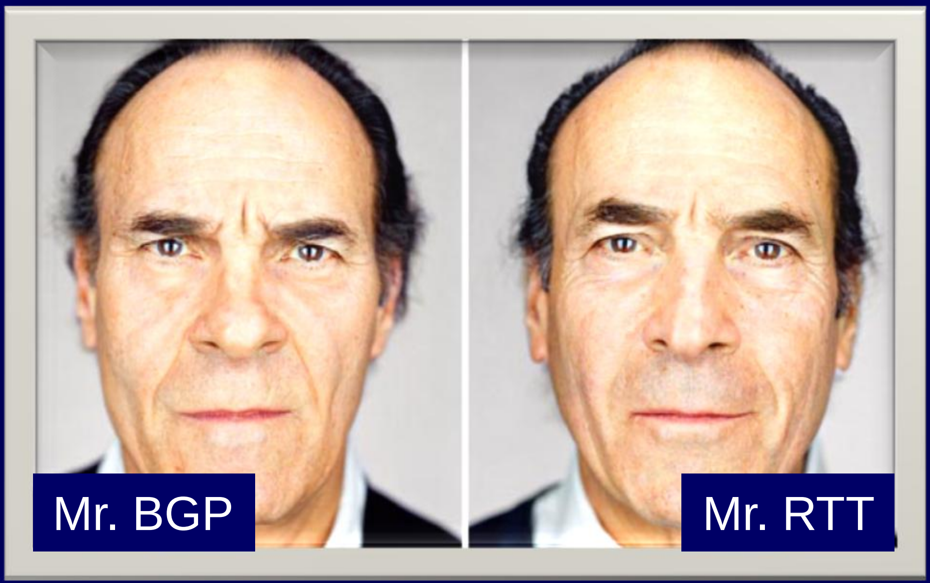

Photo ©Martin Schoeller

Back to Correlation

Photo ©Martin Schoeller

Back to Correlation

Photo ©Martin Schoeller

Back to Correlation

- Find the combination of
  - Time shift
  - Elbow slope threshold ($\Rightarrow$ Penalty)
  - Tolerance window

  that maximizes distinction between well correlated and badly correlated data

Back to Correlation

- Observation: for arbitrary
    - Time shift
    - Elbow slope threshold
    - Tolerance window
  - considering a fixed Target and a few Prefixes (one comprising the Target)
  - for all probe/CP pairs (in the same AS)

we get...

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
- considering a fixed Target and a few Prefixes (one comprising the Target)
- for all probe/CP pairs (in the same AS)

we get...

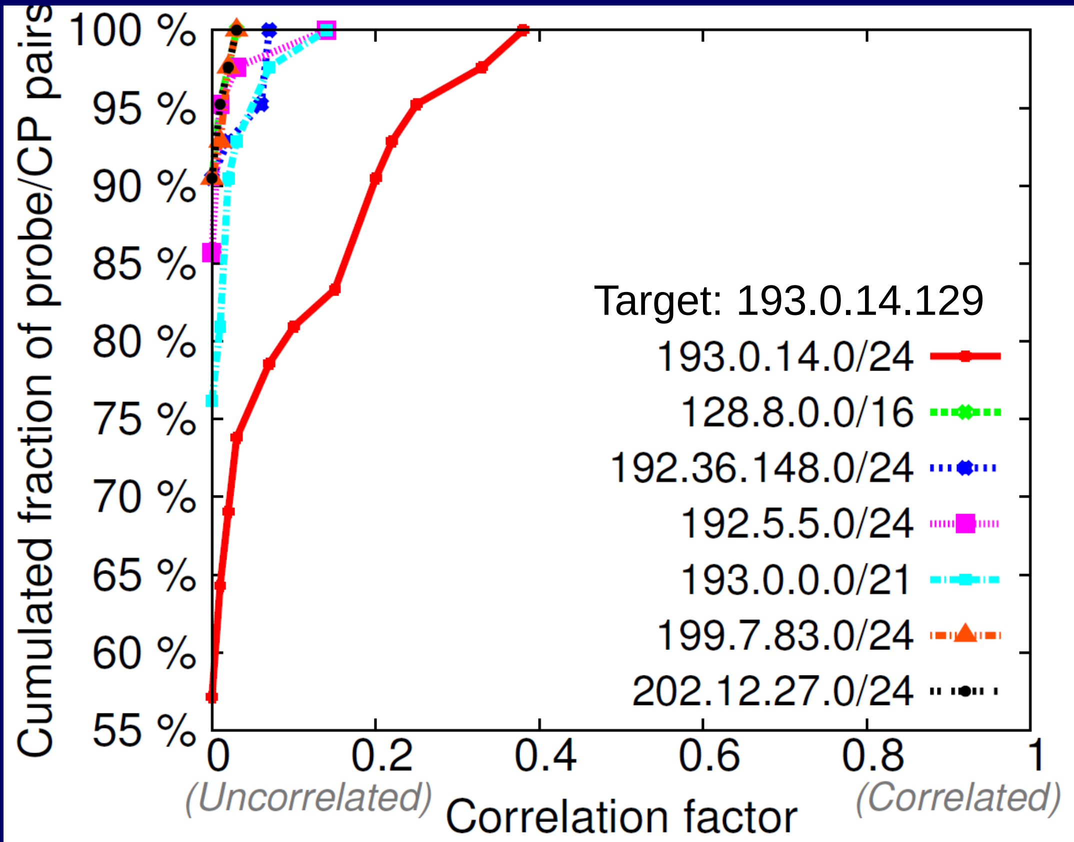

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
  - considering a fixed Target and a few Prefixes (one comprising the Target)
  - for all probe/CP pairs (in the same AS)

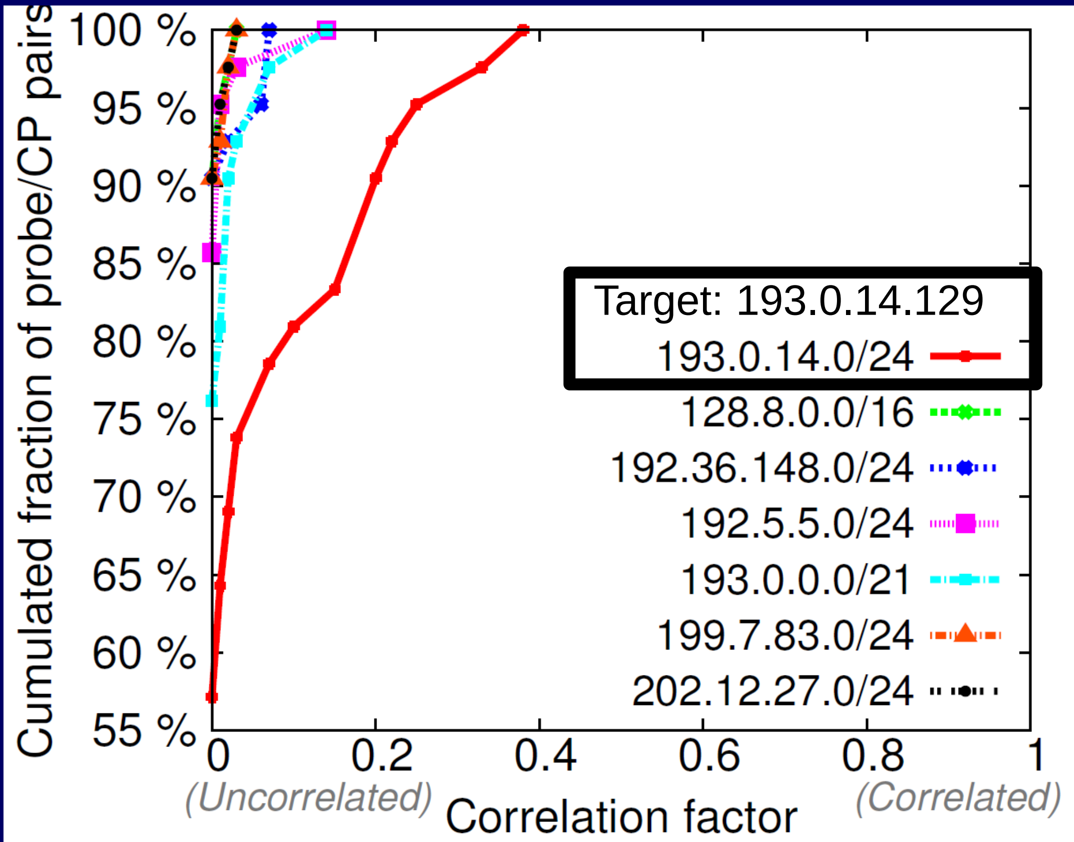

we get...

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
  - considering a fixed Target and a few Prefixes
    (one comprising the Target)
  - for all probe/CP pairs
    (in the same AS)

we get...

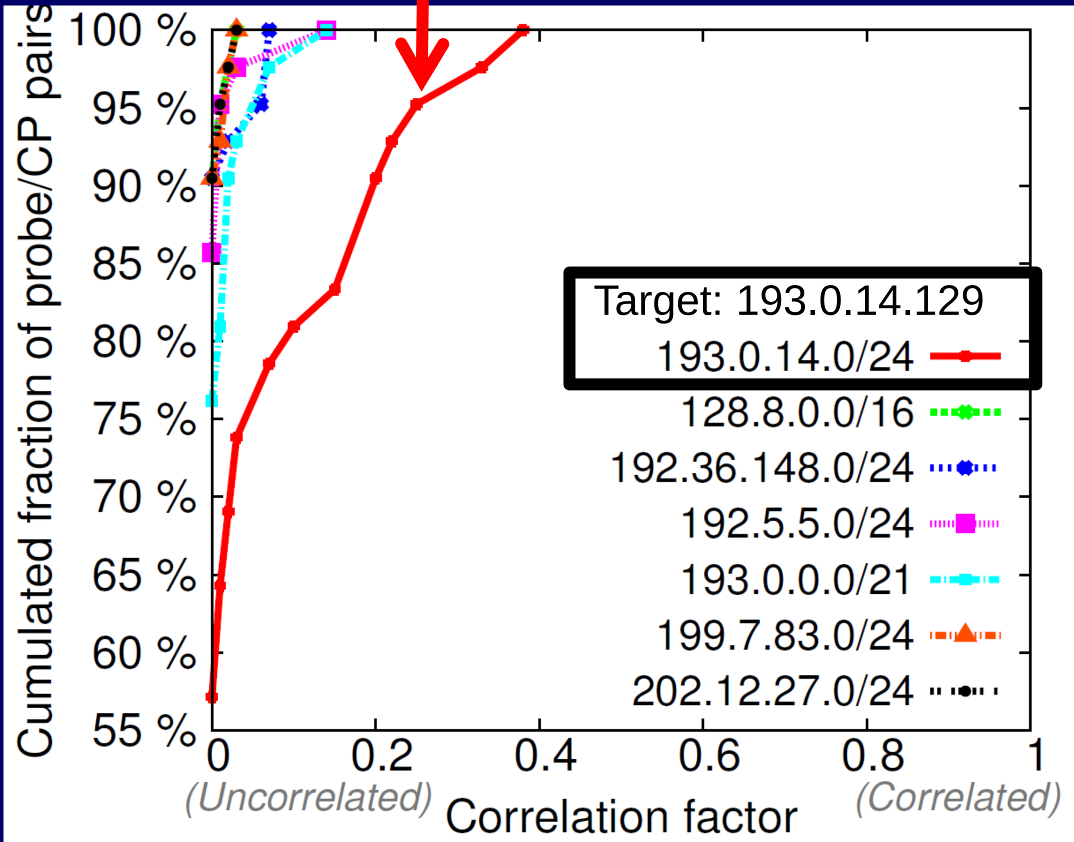

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
  - considering a fixed Target and a few Prefixes (one comprising the Target)
  - for all probe/CP pairs (in the same AS)

we get...

Higher *cf* values

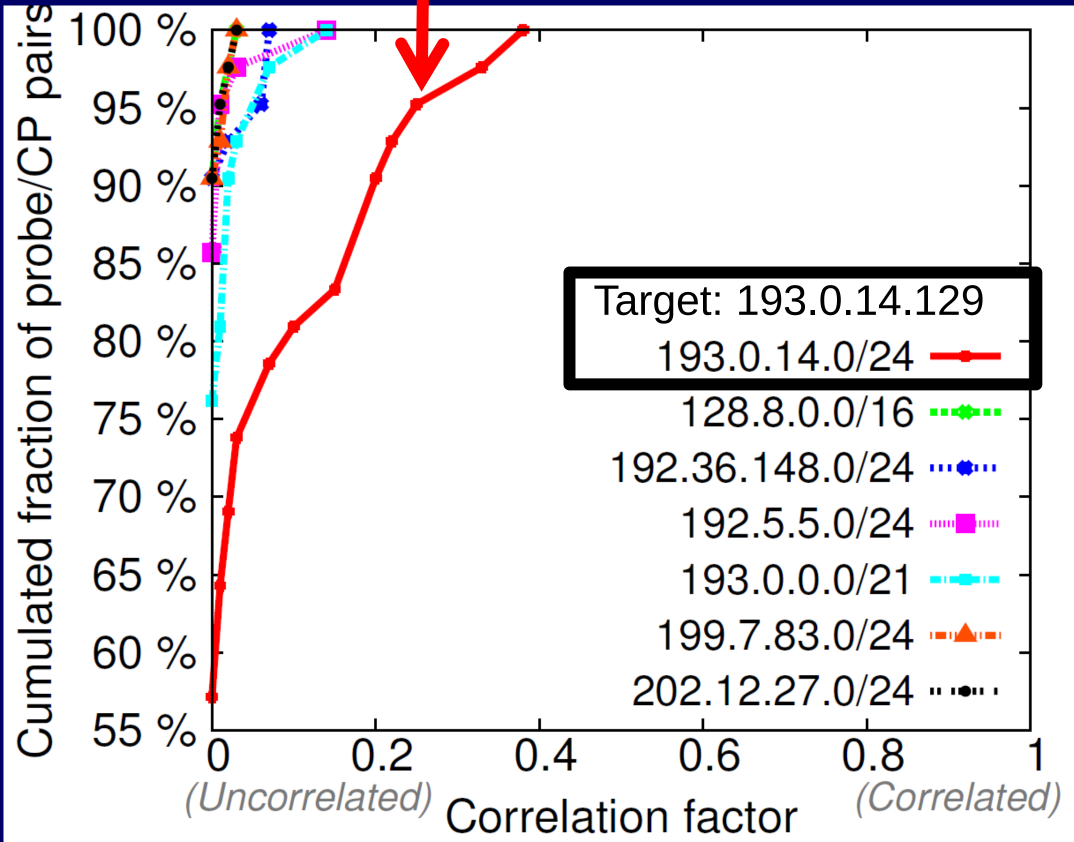

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
  - considering a fixed Target and a few Prefixes (one comprising the Target)
  - for all probe/CP pairs (in the same AS)

we get...

Higher *cf* values

Better correlation

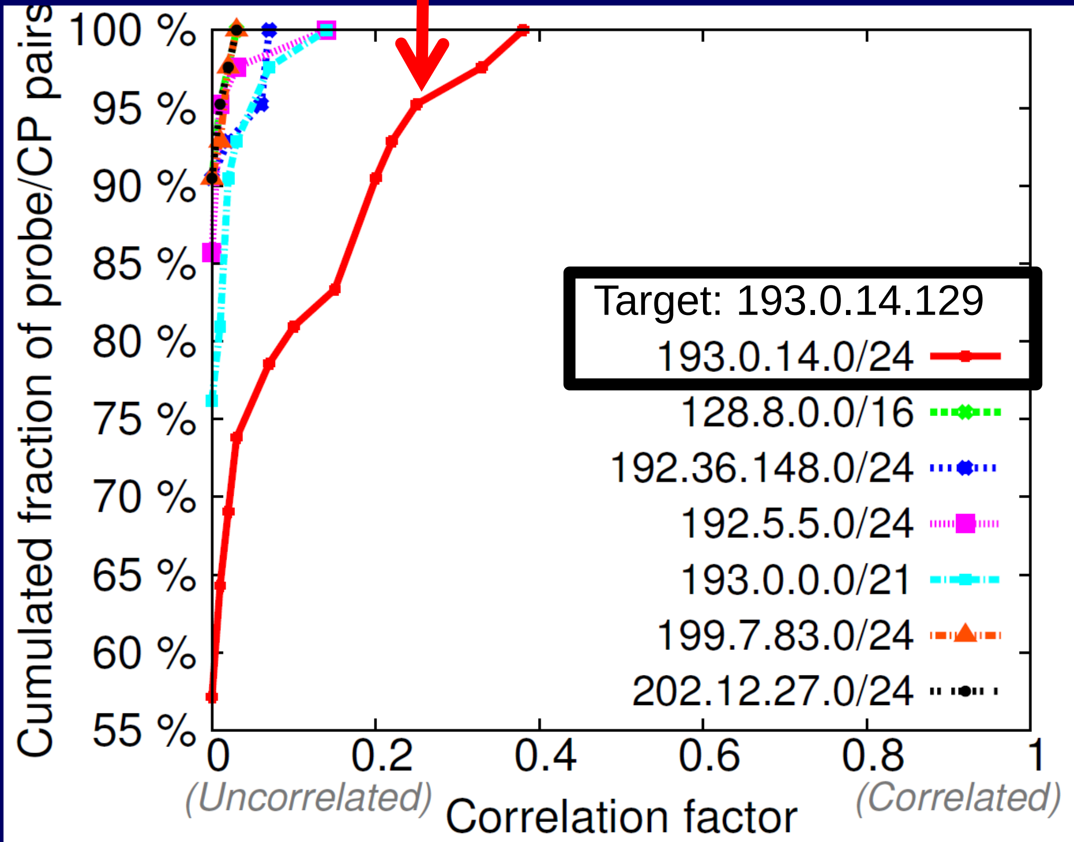

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
  - considering a fixed Target and a few Prefixes
    (one comprising the Target)
  - for all probe/CP pairs
    (in the same AS)

we get...

Higher *cf* values

Better correlation

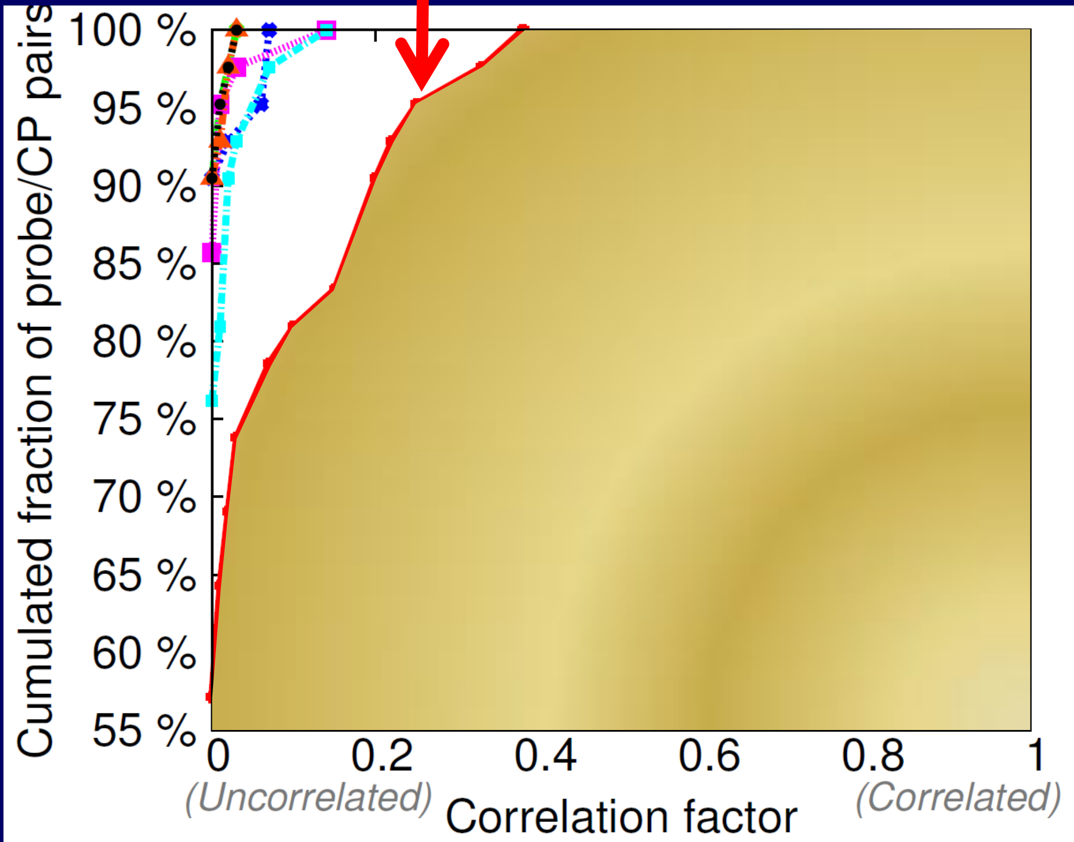

Back to Correlation

- Observation: for arbitrary
  - Time shift
  - Elbow slope threshold
  - Tolerance window
  - considering a fixed Target and a few Prefixes
    (one comprising the Target)
  - for all probe/CP pairs
    (in the same AS)

we get...

Higher *cf* values

Better correlation

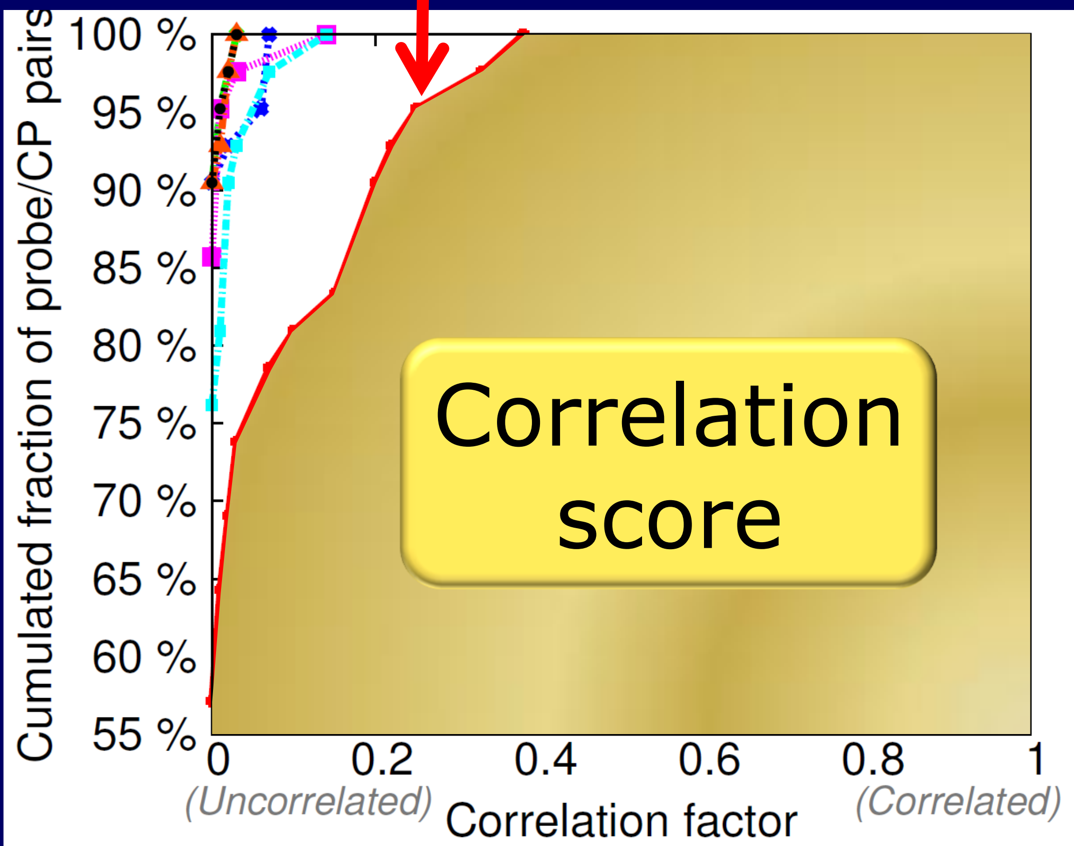

Back to Correlation

Correlation score

- Independent of the specific probe and CP
- Higher score $\Rightarrow$ lower correlation

Back to Correlation

Correlation score

- Independent of the specific probe and CP
- Higher score $\Rightarrow$ lower correlation

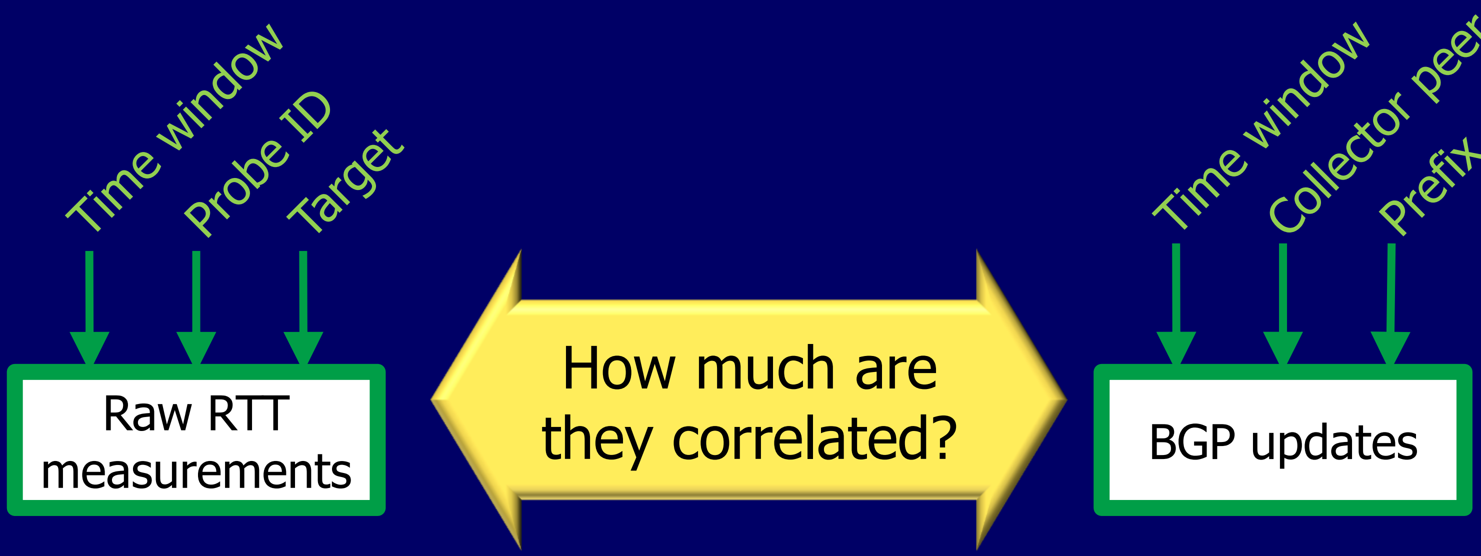

Back to Correlation

Correlation score

- Independent of the specific probe and CP
- Higher score $\Rightarrow$ lower correlation

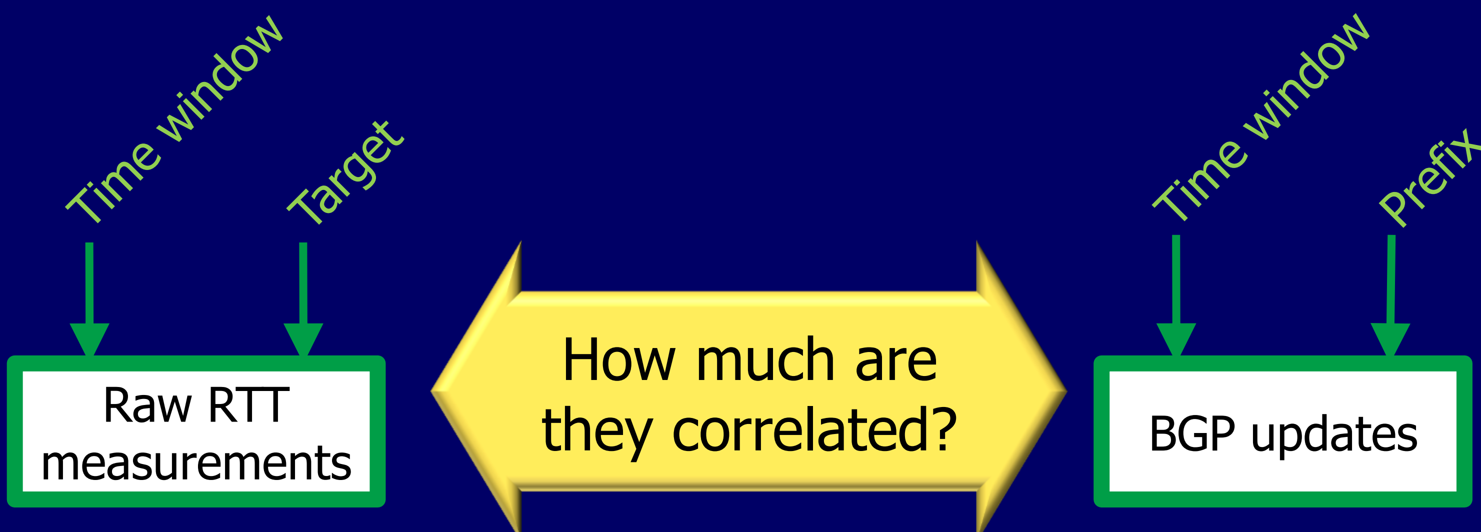

Parameter Tuning

- Time shift, Elbow slope threshold

Parameter Tuning

- Time shift, Elbow slope threshold

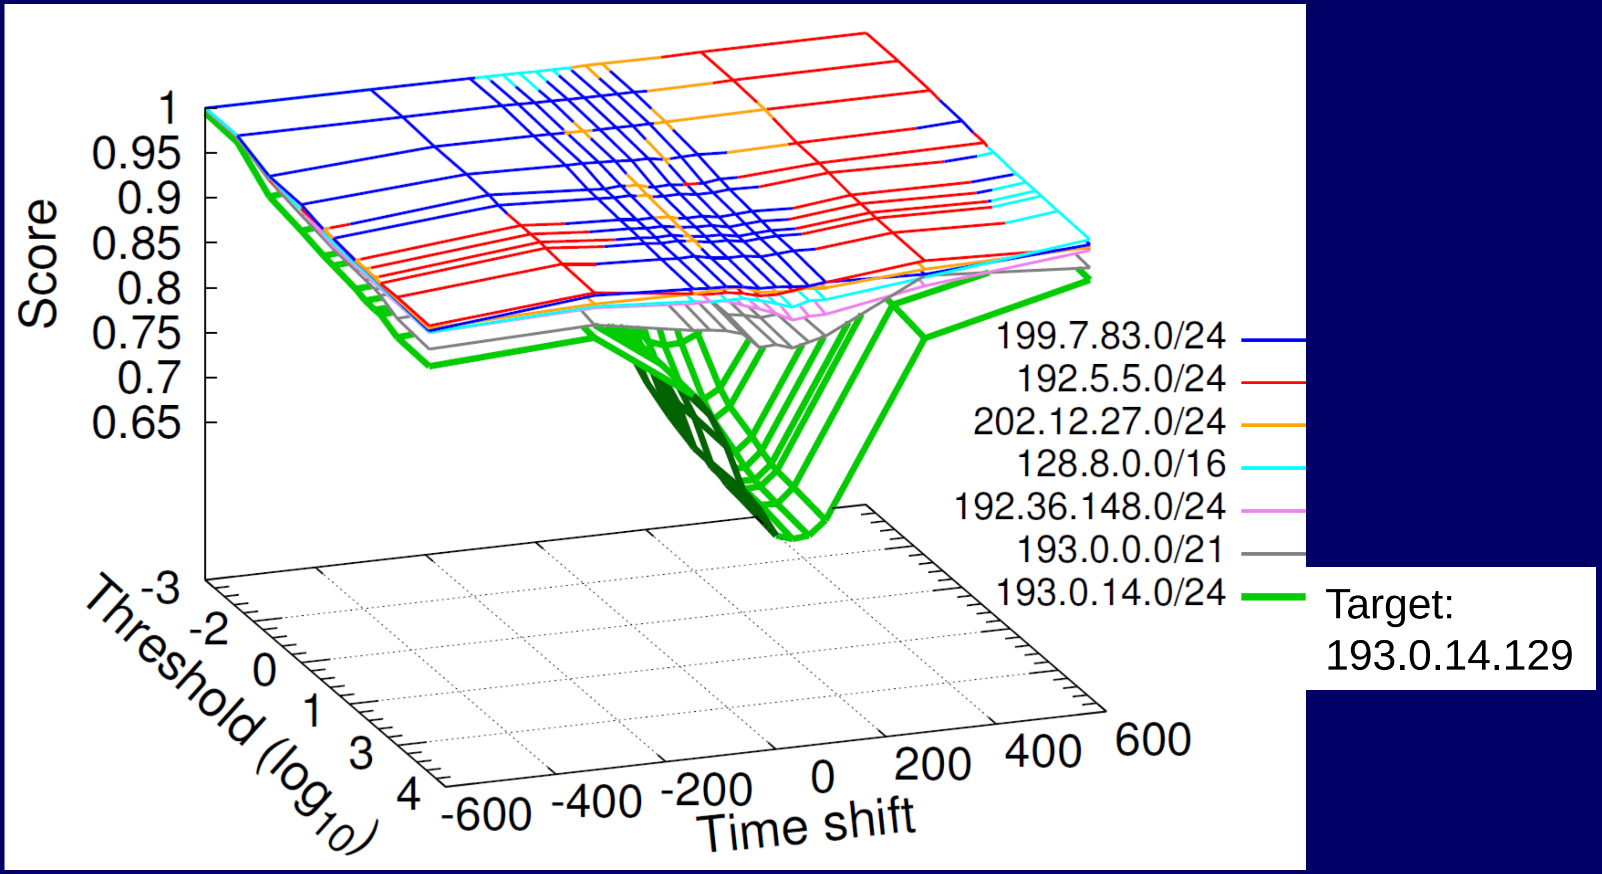

Parameter Tuning

- Time shift, Elbow slope threshold

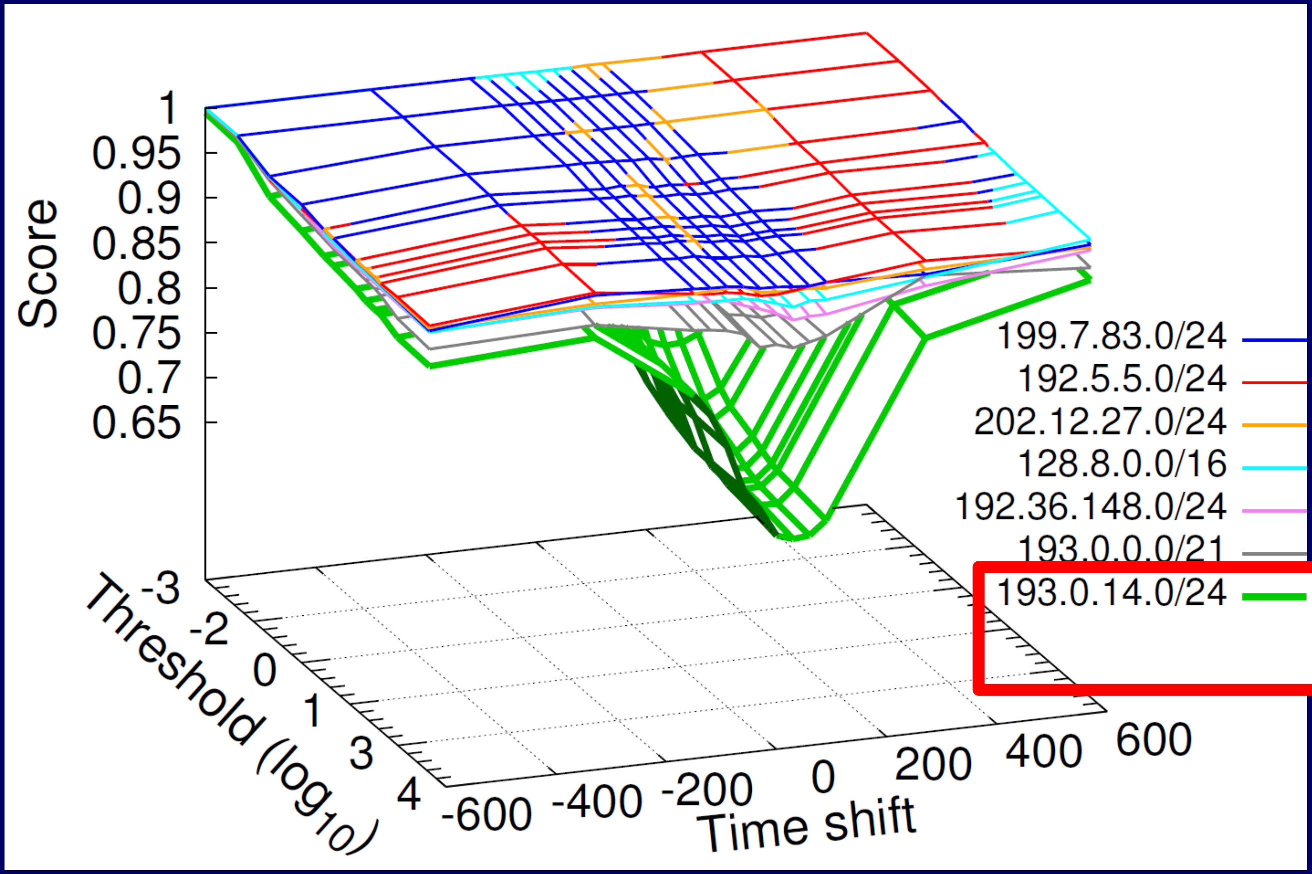

Parameter Tuning

- Time shift, Elbow slope threshold

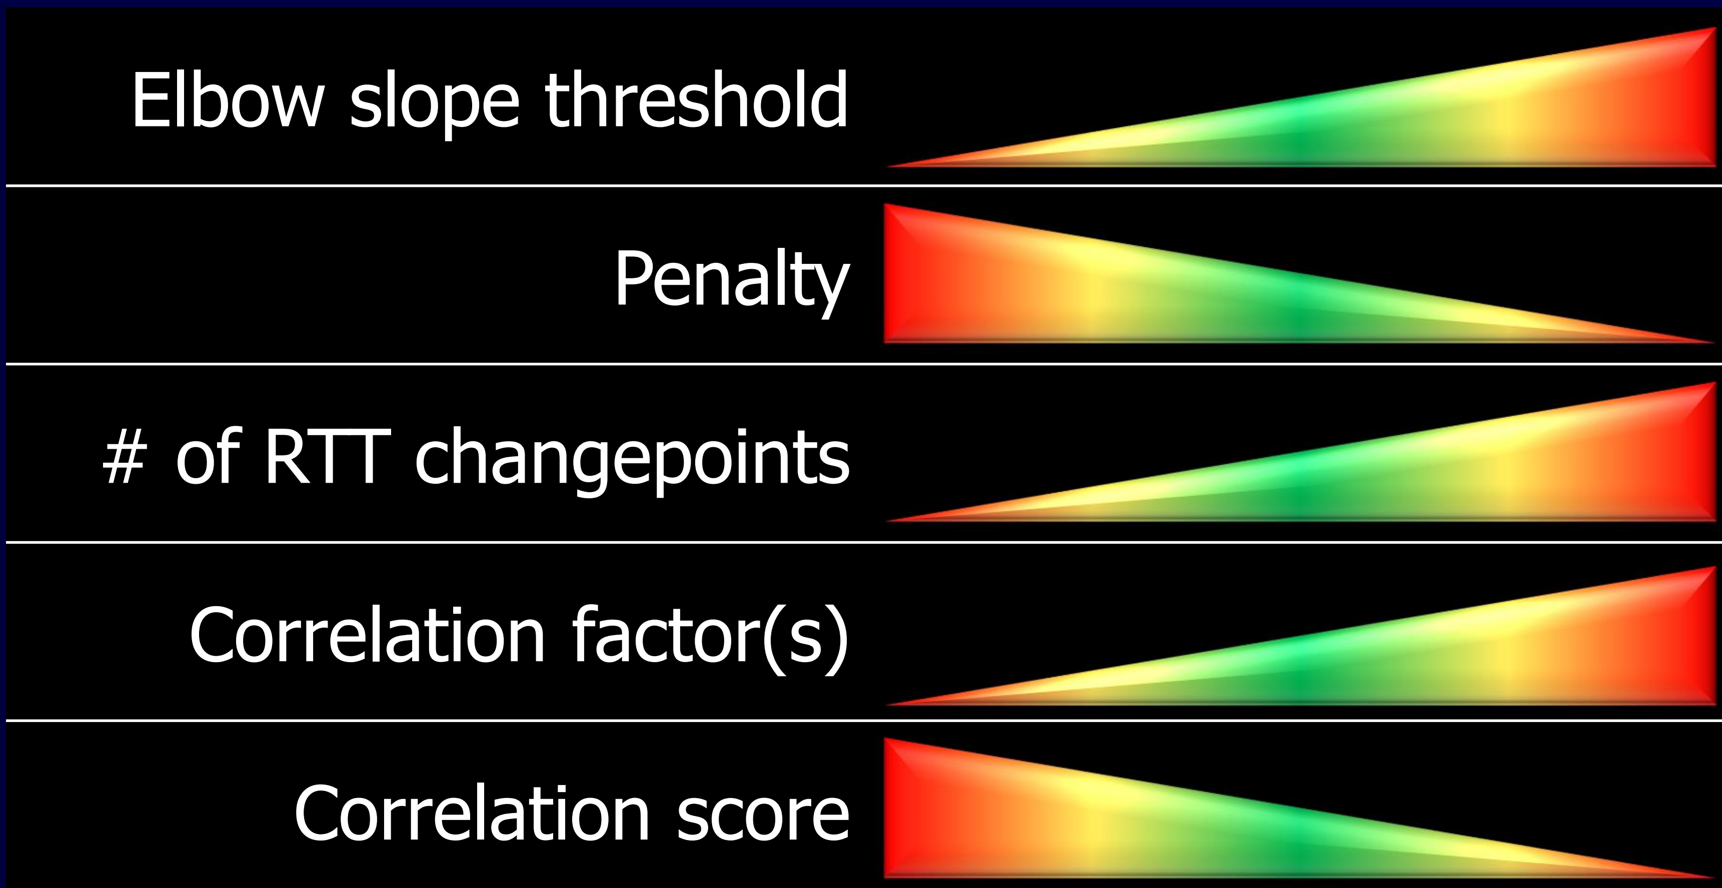

Parameter Tuning

- Time shift, Elbow slope threshold

Elbow slope threshold	
Penalty	
# of RTT changepoints	
Correlation factor(s)	
Correlation score	

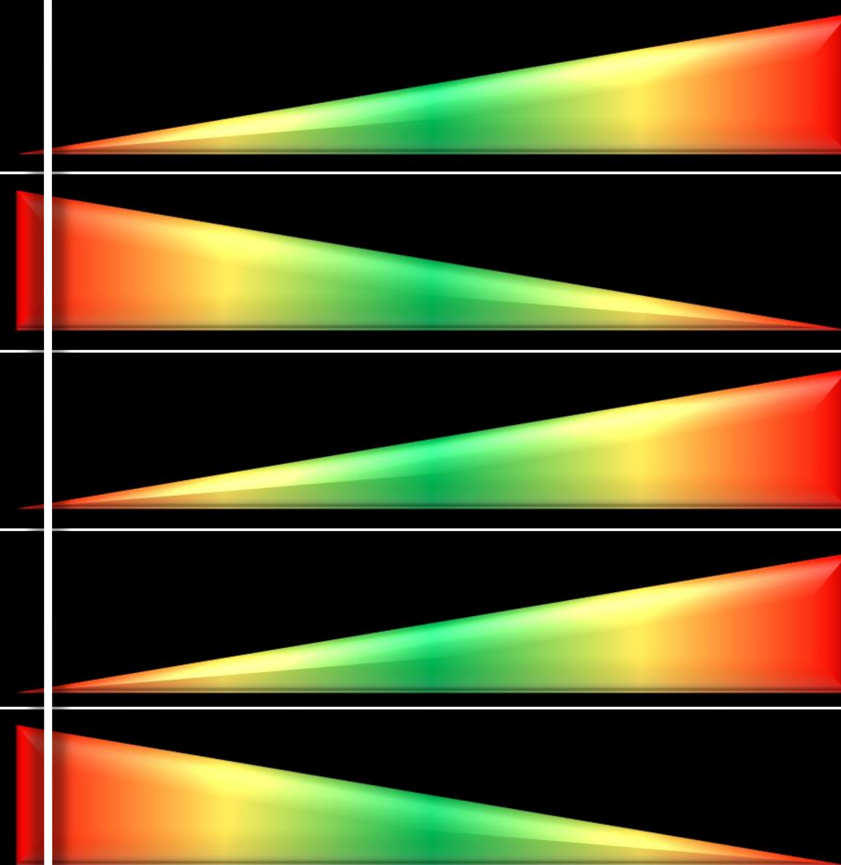

Parameter Tuning

- Time shift, Elbow slope threshold

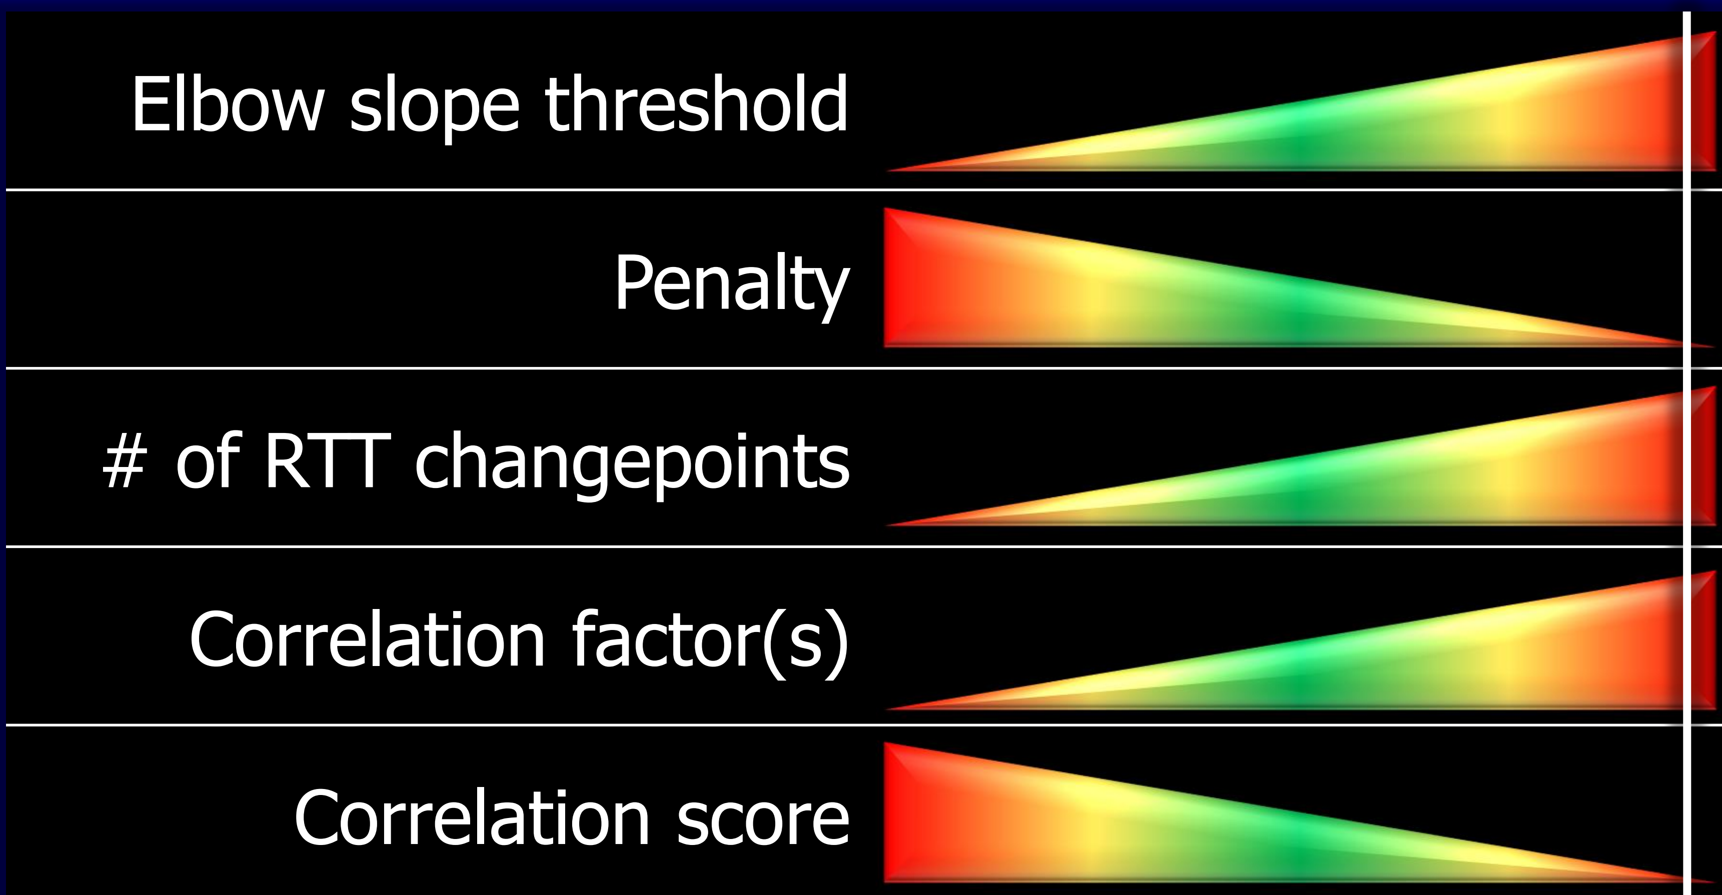

Parameter Tuning

- Time shift, Elbow slope threshold

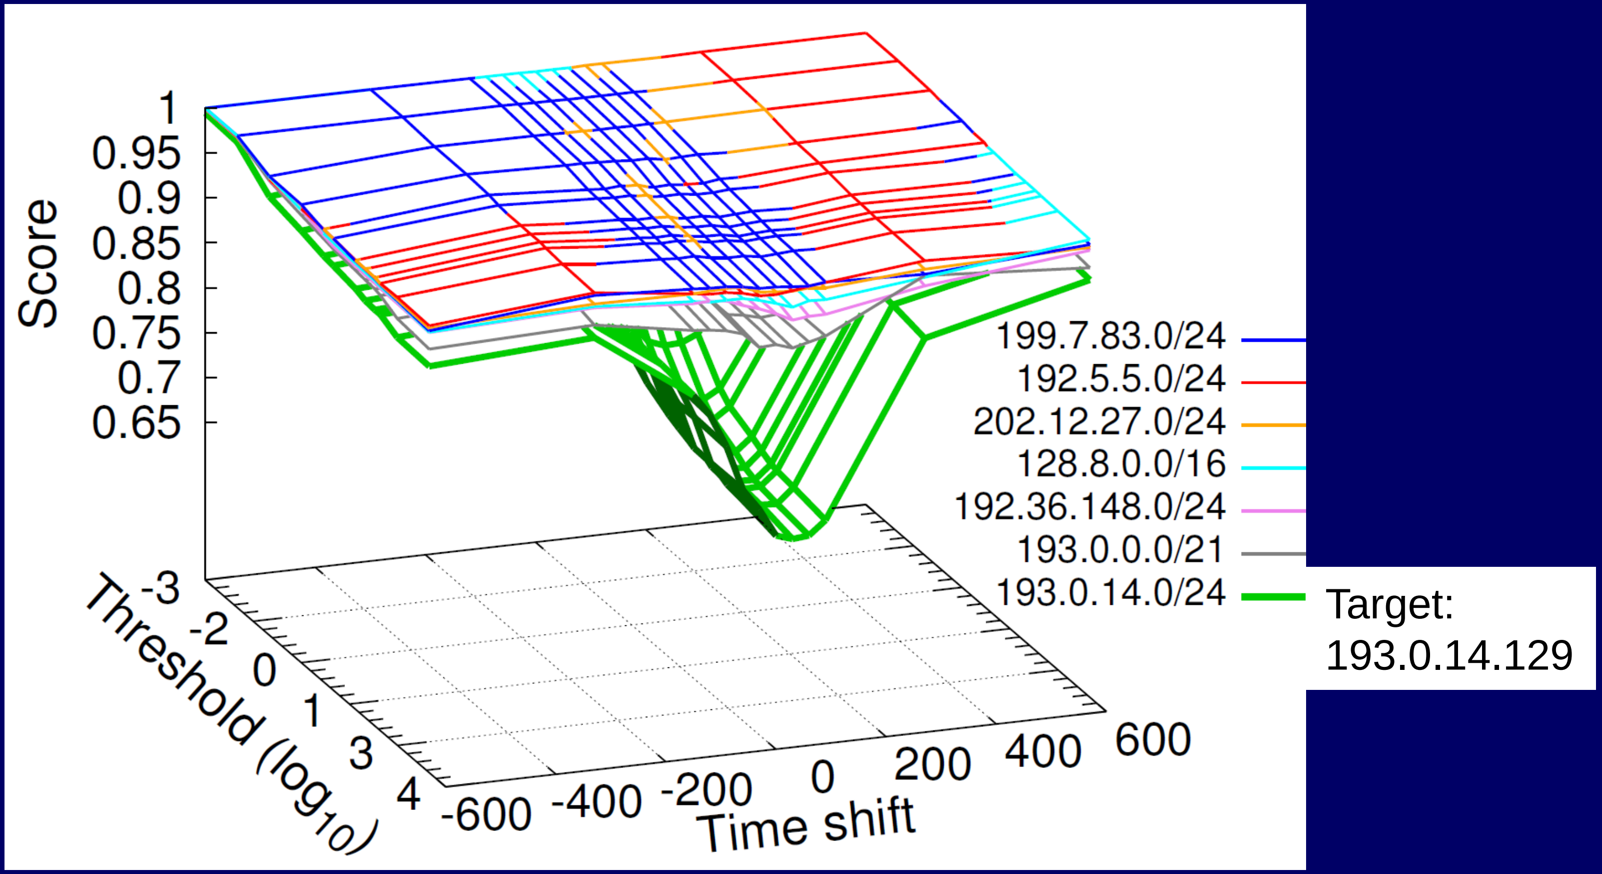

Parameter Tuning

- Time shift, Elbow slope threshold

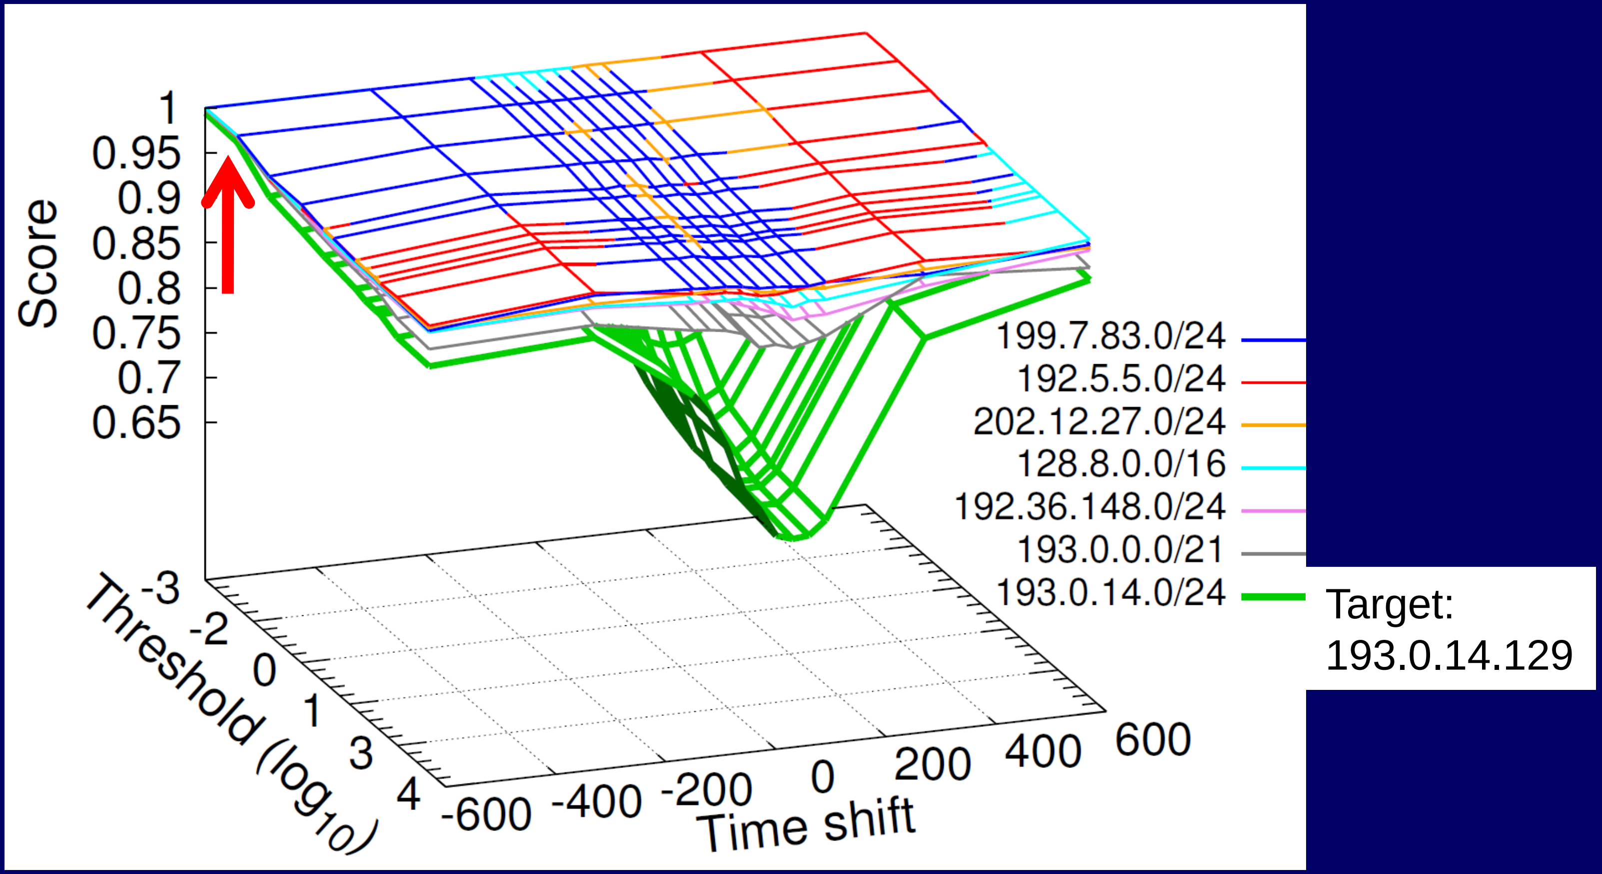

Parameter Tuning

- Time shift, Elbow slope threshold

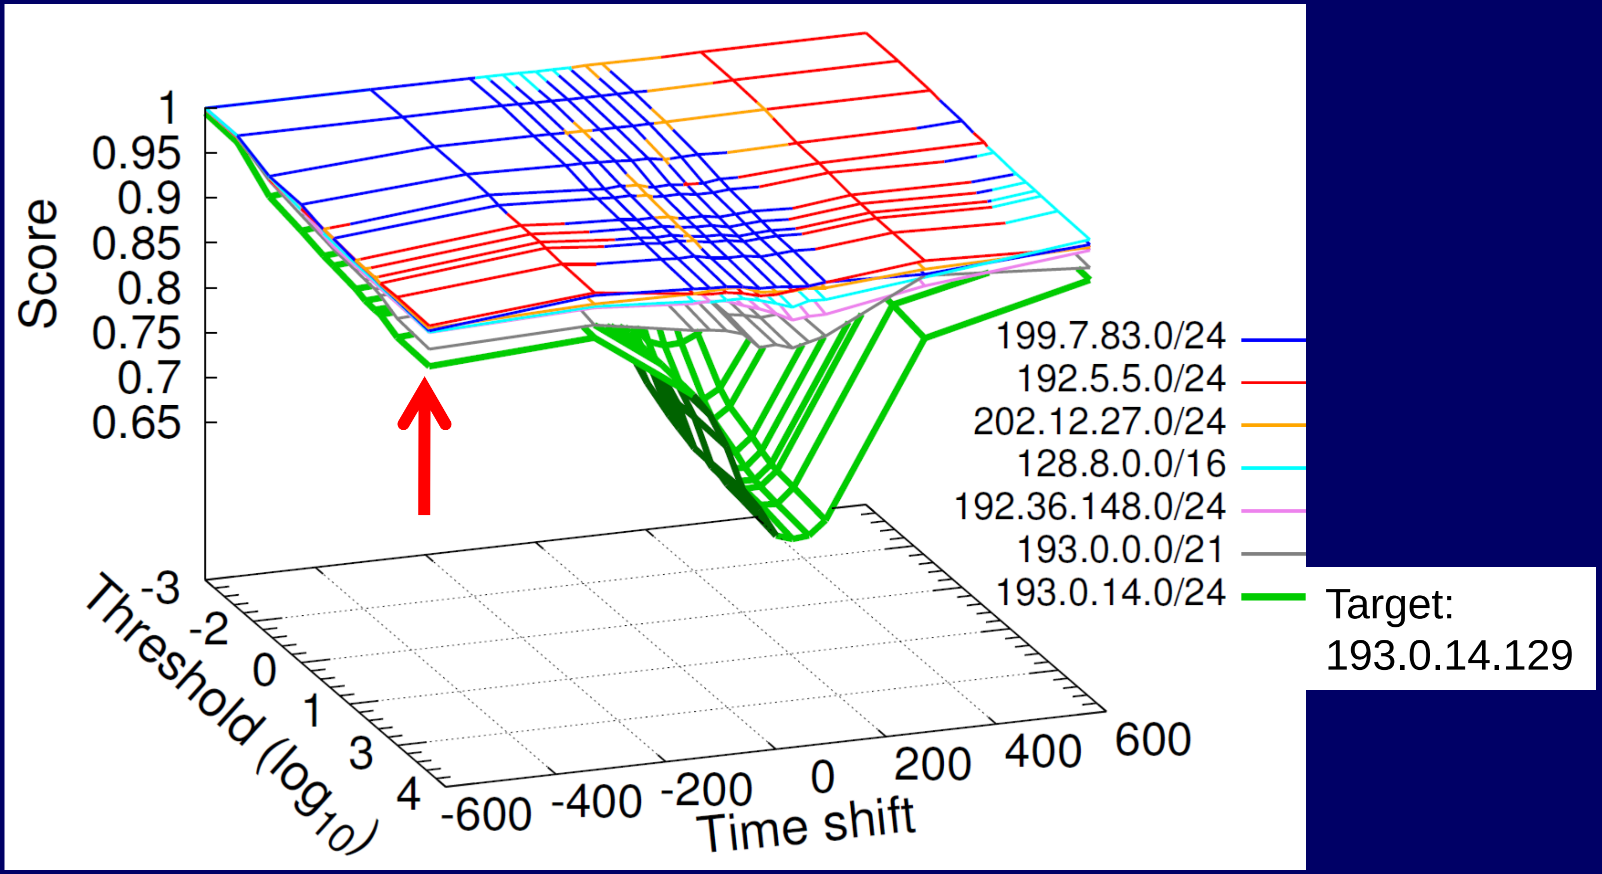

Parameter Tuning

- Time shift, Elbow slope thr

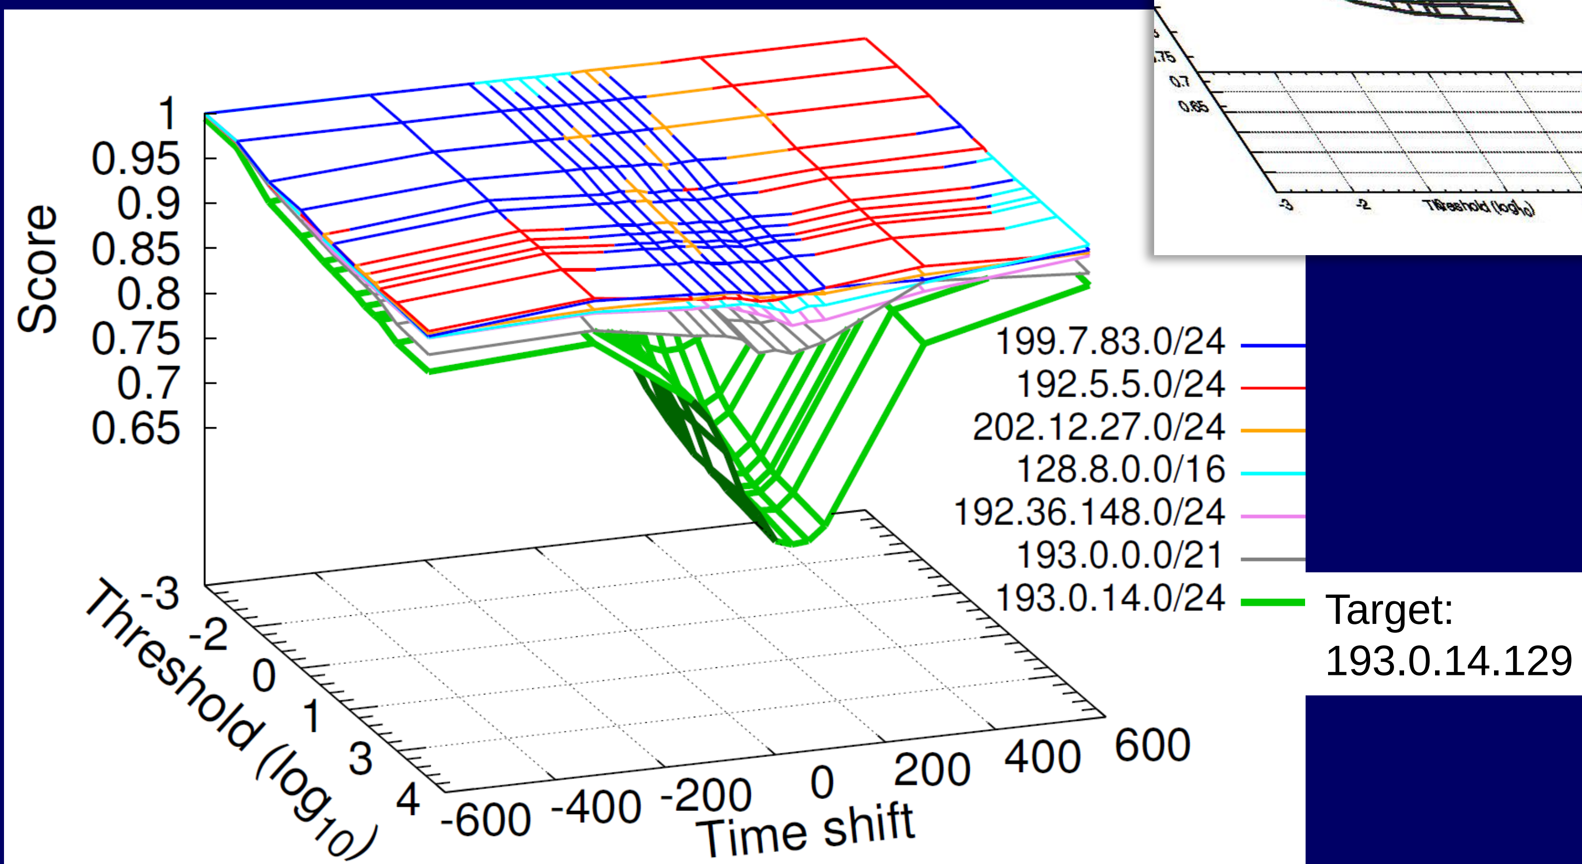

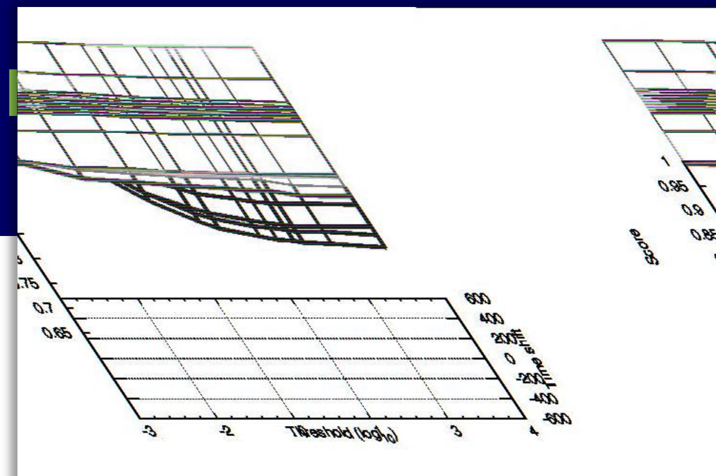

Parameter Tuning

- Time shift, Elbow slope thr

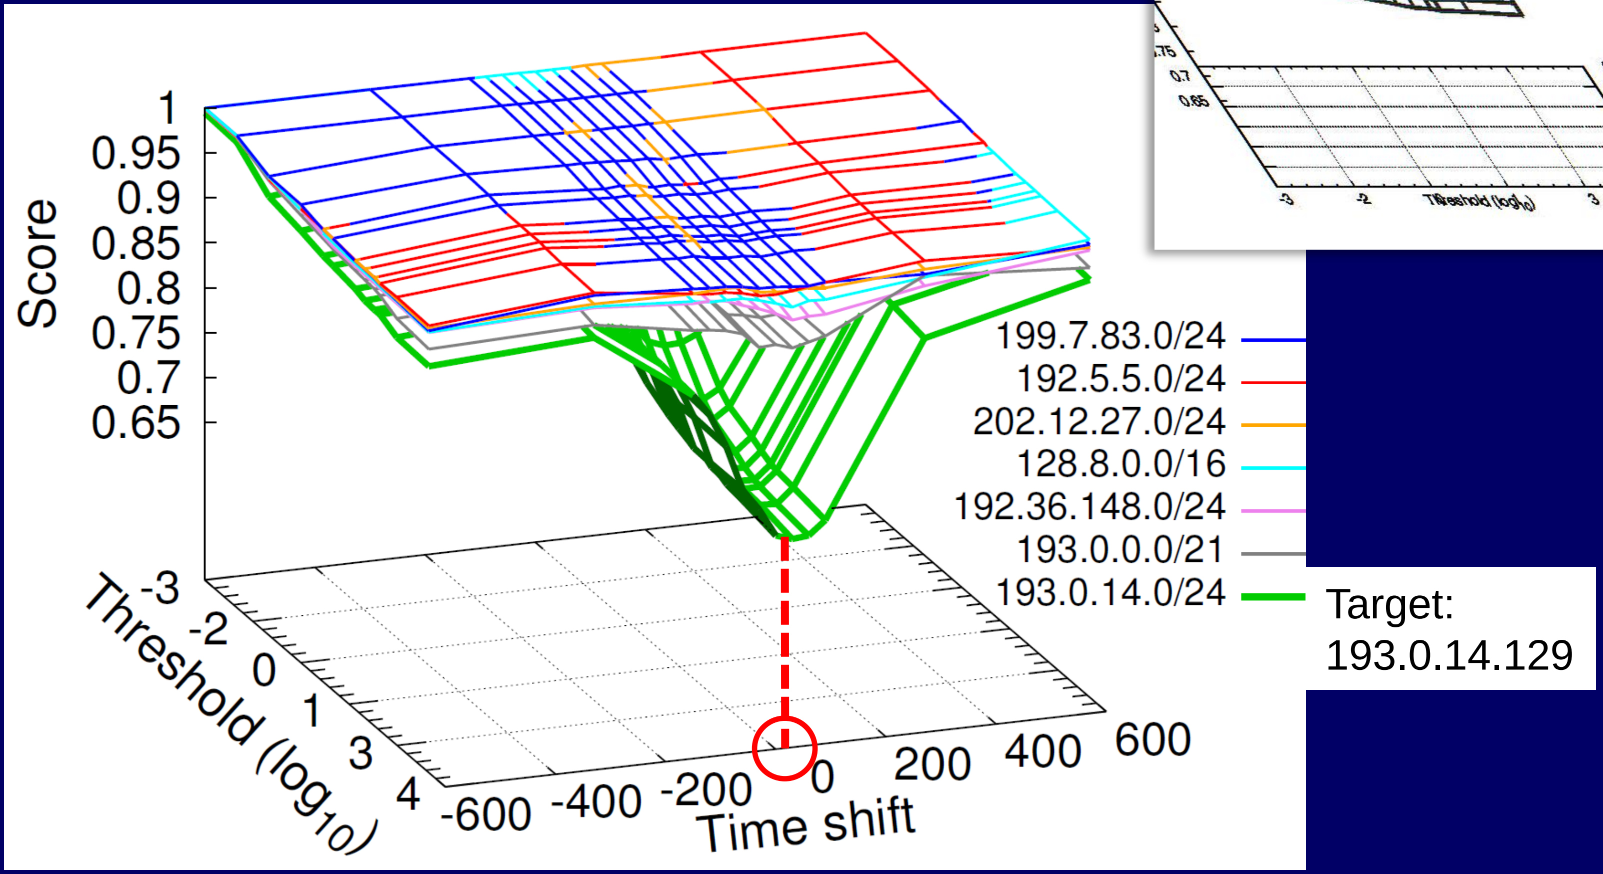

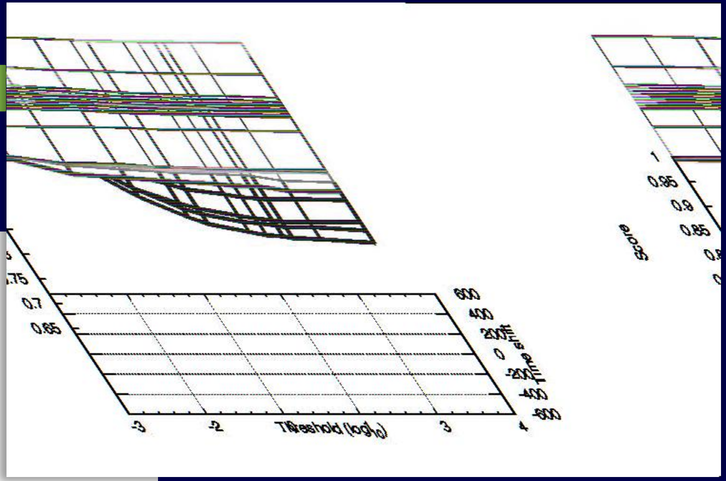

Parameter Tuning

- Time shift, Elbow slope thr

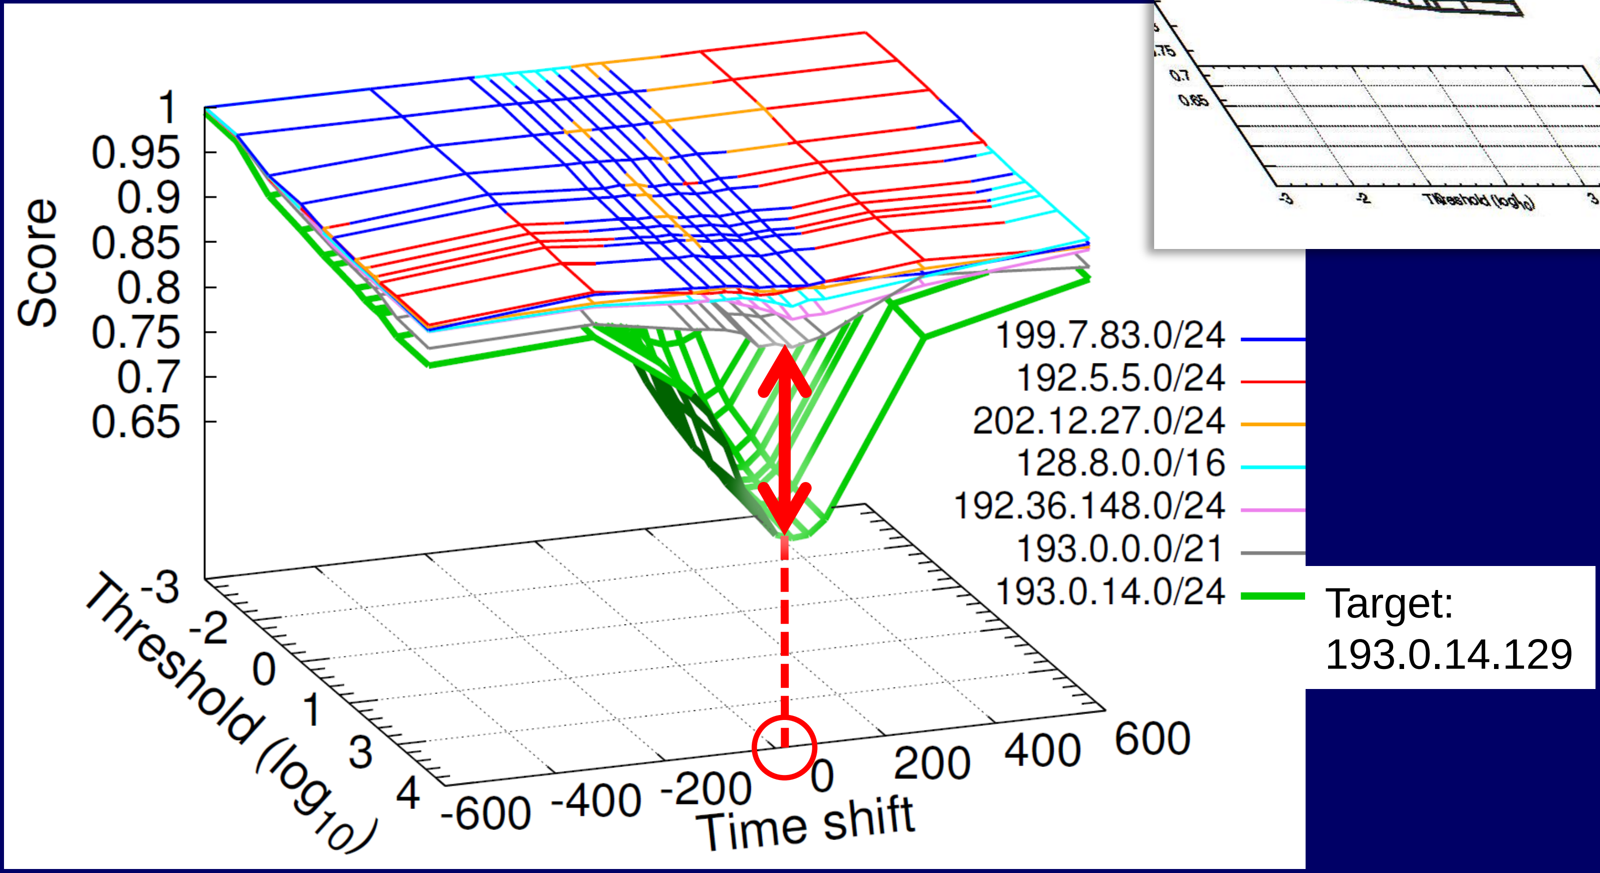

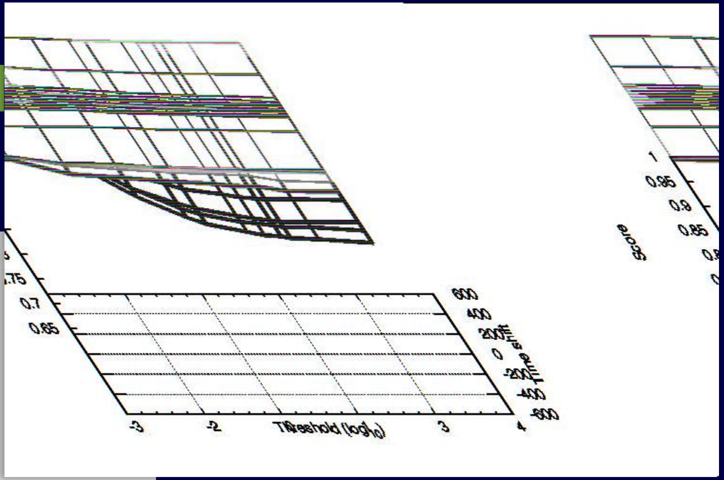

Parameter Tuning

- Time shift = 60s
- Elbow slope threshold = 10^4

Parameter Tuning

- Time shift = 60s
- Elbow slope threshold = 10^4
- Tolerance window = 5 mins
  - Based on the rate of RTT measurements

Parameter Tuning

- Time shift = 60s
- Elbow slope threshold = 10^4
- Tolerance window = 5 mins
  - Based on the rate of RTT measurements

- Good for all Targets and Prefixes!

Statistical Analyses

Measurement ID			
1001	1003	1004	1005

Statistical Analyses

Measurement ID			
1001	1003	1004	1005

Path-change: occurrence of $P_1 \rightarrow P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\mathrm{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%

Path-change: occurrence of $P_1 \rightarrow P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\mathrm{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%

Path-change: occurrence of $P_1 \rightarrow P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\mathrm{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%

Path-change: occurrence of $P_1 \rightarrow P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \rightarrow P_2$ and $P_2 \rightarrow P_1$ seen by the same probe/CP pair

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\mathrm{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%
path-change-pairs with $\mathrm{sign}(\Delta RTT_{P_1 \to P_2}) = -\mathrm{sign}(\Delta RTT_{P_2 \to P_1})$	64.8%	52.1%	43.3%	68.8%

Path-change: occurrence of $P_1 \to P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \to P_2$ and $P_2 \to P_1$ seen by the same probe/CP pair

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\operatorname{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%
path-change-pairs with $\operatorname{sign}(\Delta RTT_{P_1 \to P_2}) = -\operatorname{sign}(\Delta RTT_{P_2 \to P_1})$	64.8%	52.1%	43.3%	68.8%

Path-change: occurrence of $P_1 \to P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \to P_2$ and $P_2 \to P_1$ seen by the same probe/CP pair

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\operatorname{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%
path-change-pairs with $\operatorname{sign}(\Delta RTT_{P_1 \to P_2}) = -\operatorname{sign}(\Delta RTT_{P_2 \to P_1})$	64.8%	52.1%	43.3%	68.8%
path-changes with $\operatorname{sign}(\Delta ASpathlen) = \operatorname{sign}(\Delta RTT)$	76.4%	57.4%	64%	80.6%

Path-change: occurrence of $P_1 \to P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \to P_2$ and $P_2 \to P_1$ seen by the same probe/CP pair

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\mathrm{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%
path-change-pairs with $\mathrm{sign}(\Delta RTT_{P_1 \to P_2}) = -\mathrm{sign}(\Delta RTT_{P_2 \to P_1})$	64.8%	52.1%	43.3%	68.8%
path-changes with $\mathrm{sign}(\Delta ASpathlen) = \mathrm{sign}(\Delta RTT)$	76.4%	57.4%	64%	80.6%

Path-change: occurrence of $P_1 \to P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \to P_2$ and $P_2 \to P_1$ seen by the same probe/CP pair

Statistical Analyses

	Measurement ID			
	1001	1003	1004	1005
path-changes with consistent $\mathrm{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%
path-change-pairs with $\mathrm{sign}(\Delta RTT_{P_1 \to P_2}) = -\mathrm{sign}(\Delta RTT_{P_2 \to P_1})$	64.8%	52.1%	43.3%	68.8%
path-changes with $\mathrm{sign}(\Delta ASpathlen) = \mathrm{sign}(\Delta RTT)$	76.4%	57.4%	64%	80.6%
path-changes with $\sigma_{\Delta RTT} / \overline{\Delta RTT} < 0.25$ (~same BGP change $\Rightarrow$ same RTT change)	73.6%	75.5%	95.5%	93.1%

Path-change: occurrence of $P_1 \to P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \to P_2$ and $P_2 \to P_1$ seen by the same probe/CP pair

Statistical Analyses

	Measurement ID			
	1001	**1003**	**1004**	**1005**
path-changes with consistent $\text{sign}(\Delta RTT)$	87.5%	78.6%	72.5%	86.4%
path-change-pairs with $\text{sign}(\Delta RTT_{P_1 \to P_2}) = -\text{sign}(\Delta RTT_{P_2 \to P_1})$	64.8%	52.1%	43.3%	68.8%
path-changes with $\text{sign}(\Delta ASpathlen) = \text{sign}(\Delta RTT)$	76.4%	57.4%	64%	80.6%
path-changes with $\sigma_{\Delta RTT} / \overline{\Delta RTT} < 0.25$ (~same BGP change $\Rightarrow$ same RTT change)	73.6%	75.5%	95.5%	93.1%

Path-change: occurrence of $P_1 \to P_2$ recorded by a CP and matched with an RTT variation seen by a probe in the same AS

Path-change-pair: occurrence of $P_1 \to P_2$ and $P_2 \to P_1$ seen by the same probe/CP pair

Traceroute-based Validation

Traceroute-based Validation

- Traceroutes from RIPE Atlas
  - One every 20 mins

Traceroute-based Validation

- Traceroutes from RIPE Atlas
  - One every 20 mins
- IP$\rightarrow$AS mapping
  - Based on RIPE RIS BGP tables

Traceroute-based Validation

- Traceroutes from RIPE Atlas
  - One every 20 mins
- IP→AS mapping
  - Based on RIPE RIS BGP tables
- BGP-RTT → BGP-traceroute

Traceroute-based Validation

- Traceroutes from RIPE Atlas
  - One every 20 mins
- IP$\rightarrow$AS mapping
  - Based on RIPE RIS BGP tables
- BGP-RTT $\rightarrow$ BGP-traceroute
- Results:
  - For well correlated BGP-RTT data (high *cf*) there is evidence of correlation even with traceroute data

Traceroute-based Validation

- Traceroutes from RIPE Atlas
  - One every 20 mins
- IP→AS mapping
  - Based on RIPE RIS BGP tables
- BGP-RTT → BGP-traceroute
- Results:
  - For badly correlated BGP-RTT data (low *cf*) there is no evidence of correlation even with traceroute data

Further Possible Analyses

Further Possible Analyses

- Recall from the motivations

Further Possible Analyses

- Recall from the motivations
  - Equivalence classes

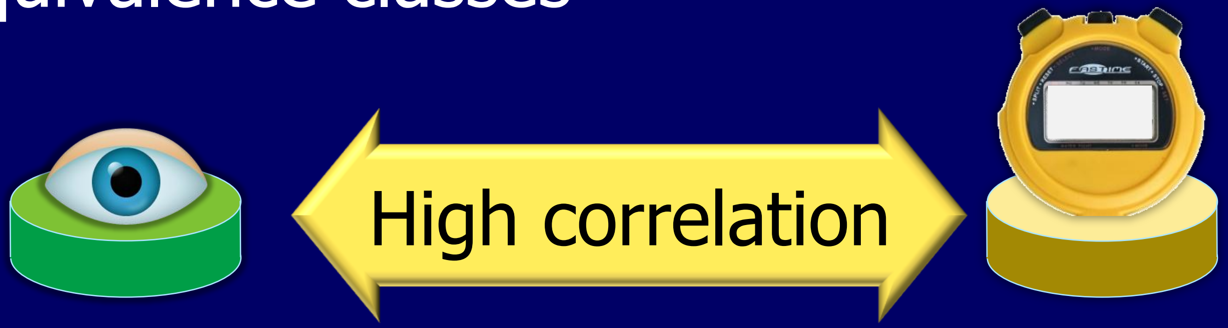

Further Possible Analyses

- Recall from the motivations
  - Equivalence classes

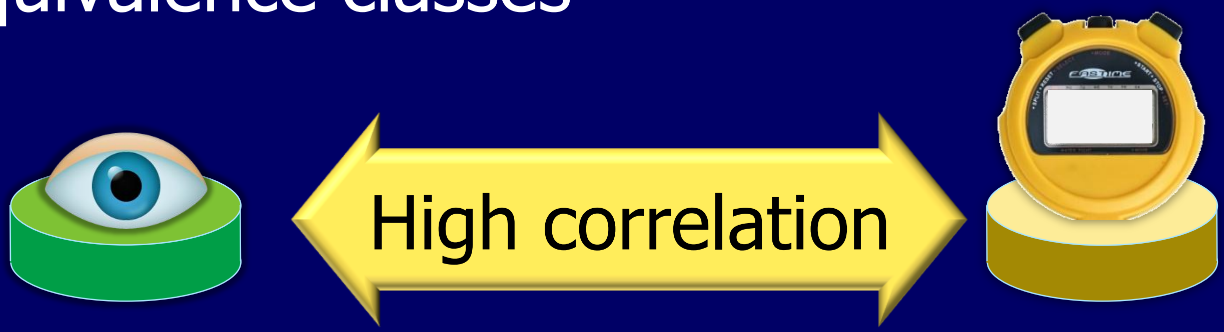

  - Placement of delay-sensitive services

Further Possible Analyses

- Recall from the motivations
  - Equivalence classes

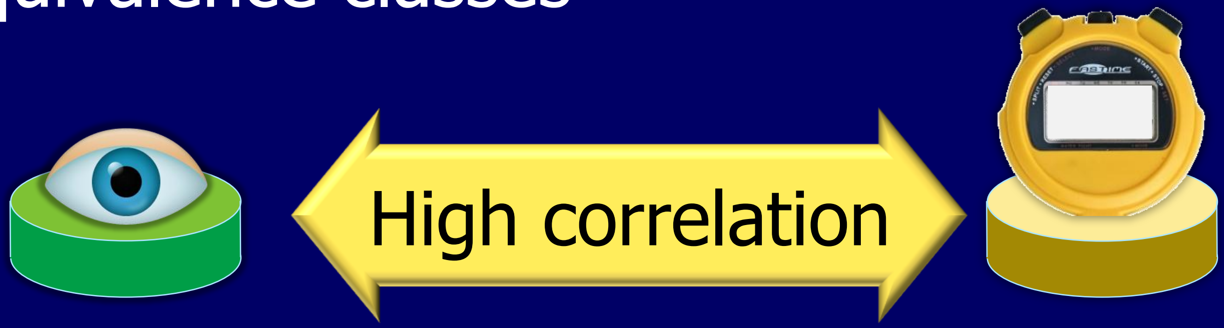

  - Placement of delay-sensitive services
  - Troubleshooting

Conclusions

Conclusions

- **Automated Investigation of the Impact of BGP Routing Changes on Network Delay Variations**

Conclusions

- **Automated Investigation of the Impact of BGP Routing Changes on Network Delay Variations**

- Application (+validation) & statistics

Future Work

Future Work

- Other data sources/vantage points/targets

Future Work

- Other data sources/vantage points/targets
- Other statistical methods
  - Consider noise, gaps, patterns

Future Work

- Other data sources/vantage points/targets
- Other statistical methods
  - Consider noise, gaps, patterns
    - Unreachability in BGP, gap in RTT: ok! ✓

Future Work

- Other data sources/vantage points/targets
- Other statistical methods
  - Consider noise, gaps, patterns
    - Unreachability in BGP, gap in RTT: ok! ✓
- 1-way performance indicators in a controlled environment
  - Wi-Fi

Future Work

- Other data sources/vantage points/targets
- Other statistical methods
  - Consider noise, gaps, patterns
    - Unreachability in BGP, gap in RTT: ok! ✓
- 1-way performance indicators in a controlled environment
  - Wi-Fi
- Further analyses

hank you!