Network neutrality inference

Zhiyong Zhang, Ovidiu Mara,
Katerina Argyraki

EPFL, UESTC

Neutrality

- Neutrality violation = a link subjects different flows to different policies

Neutrality

- Neutrality violation = a link subjects different flows to different policies

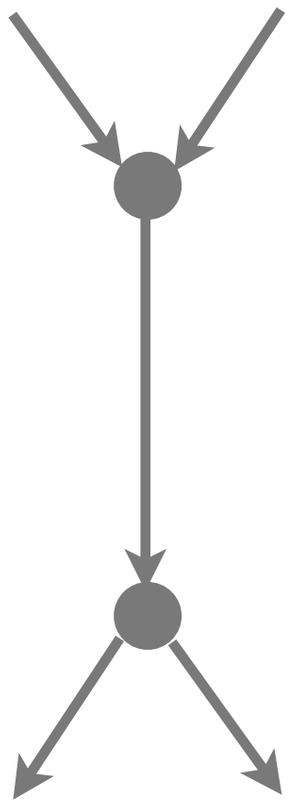

Neutrality

- Neutrality violation = a link subjects different flows to different policies

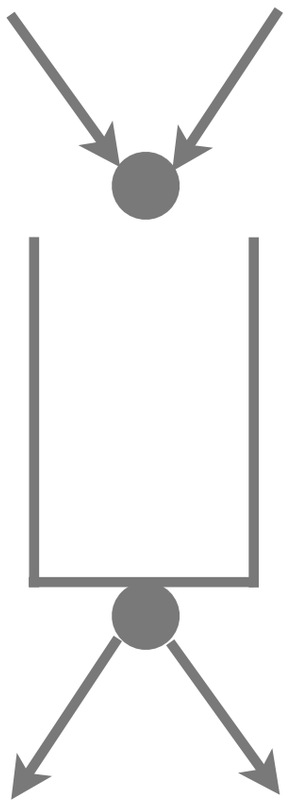

Neutrality

- Neutrality violation = a link subjects different flows to different policies

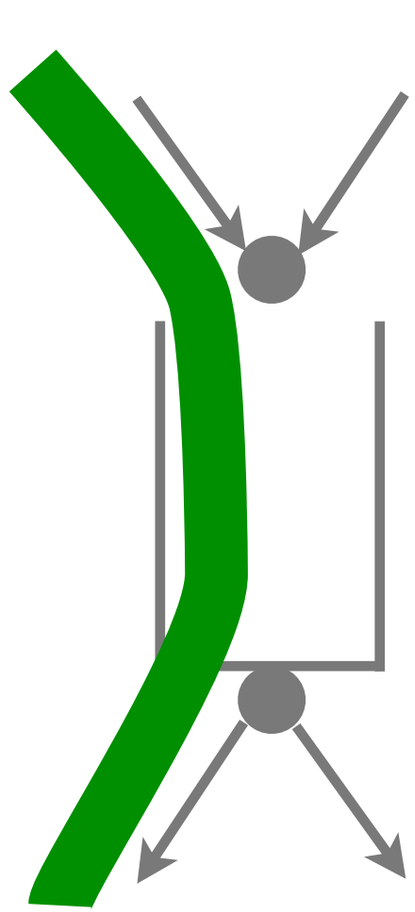

Neutrality

- Neutrality violation = a link subjects different flows to different policies

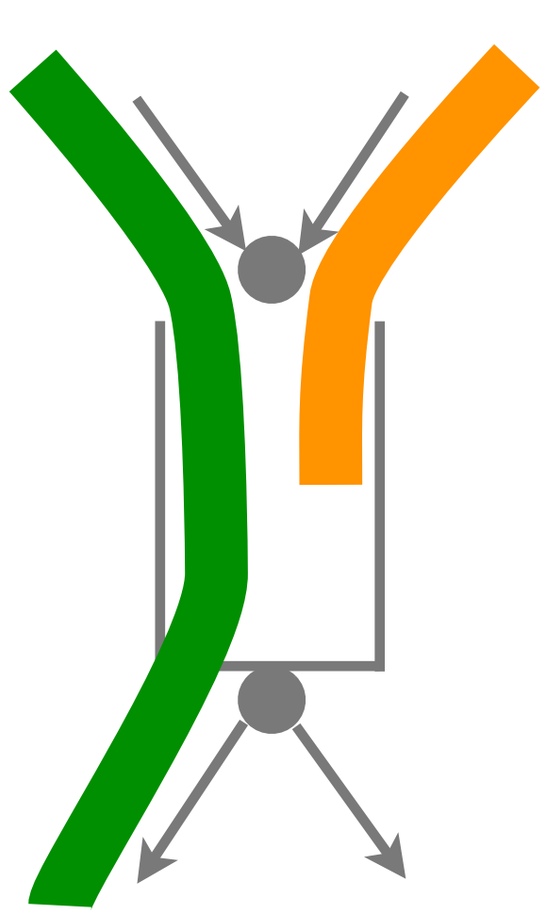

Neutrality

- Neutrality violation = a link subjects different flows to different policies

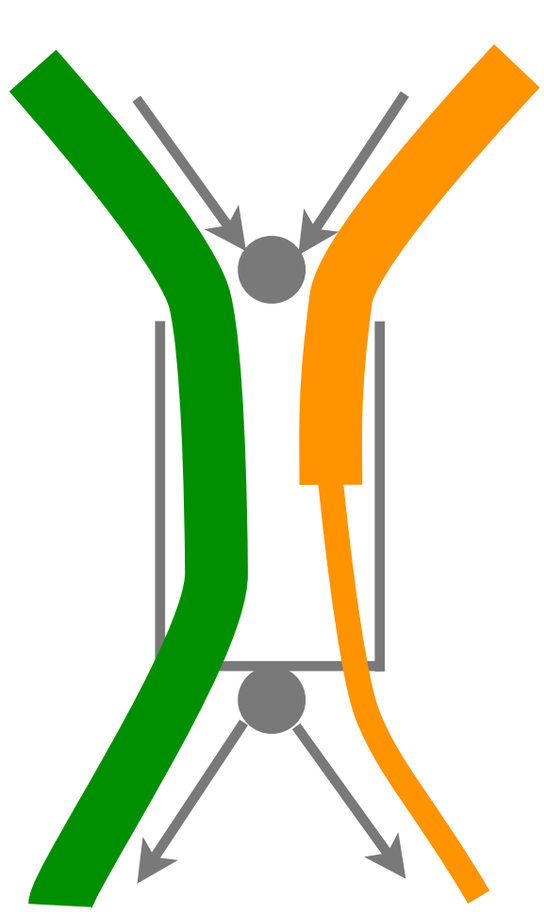

Transparency

- Neutrality violation should be transparent
- Is it feasible to externally observe neutrality violations?
- Is it feasible to localize them to specific links?

Inconsistent observations

- When the network is non-neutral,
  if we make diverse external observations,
  these will be inconsistent with each other.

Network tomography

- Input: network topology,
  path measurements

- Output: link properties
  - *average loss rate*
  - *average latency*
  - *congestion status*
  - *congestion probability*

Network tomography

- Input: network topology, path measurements

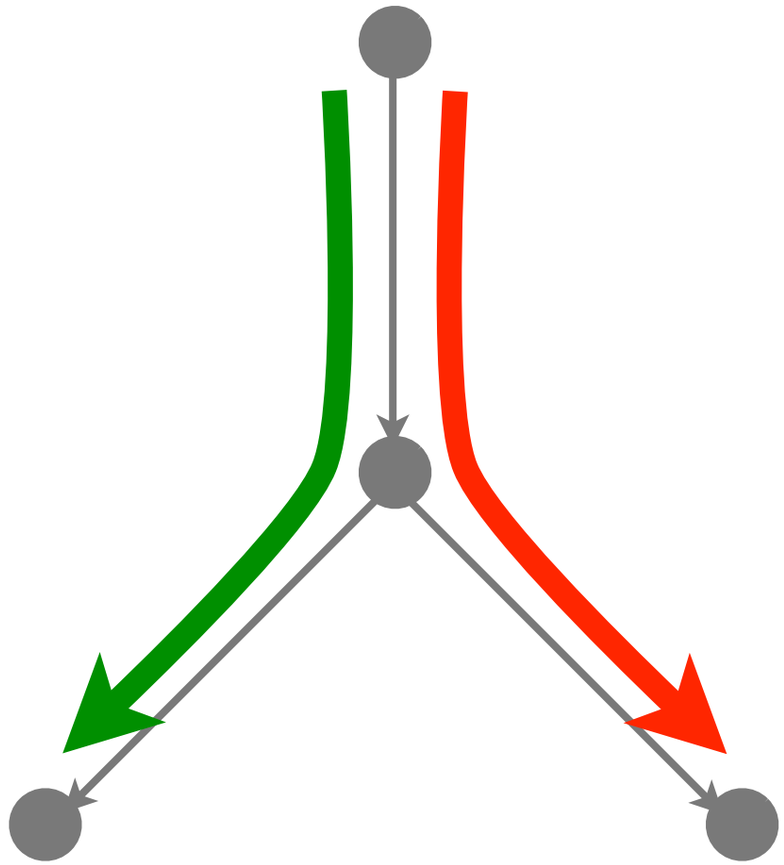

- Output: link properties
  - *average loss rate*
  - *average latency*
  - *congestion status*
  - *congestion probability*

Network tomography

- Input: network topology, path measurements

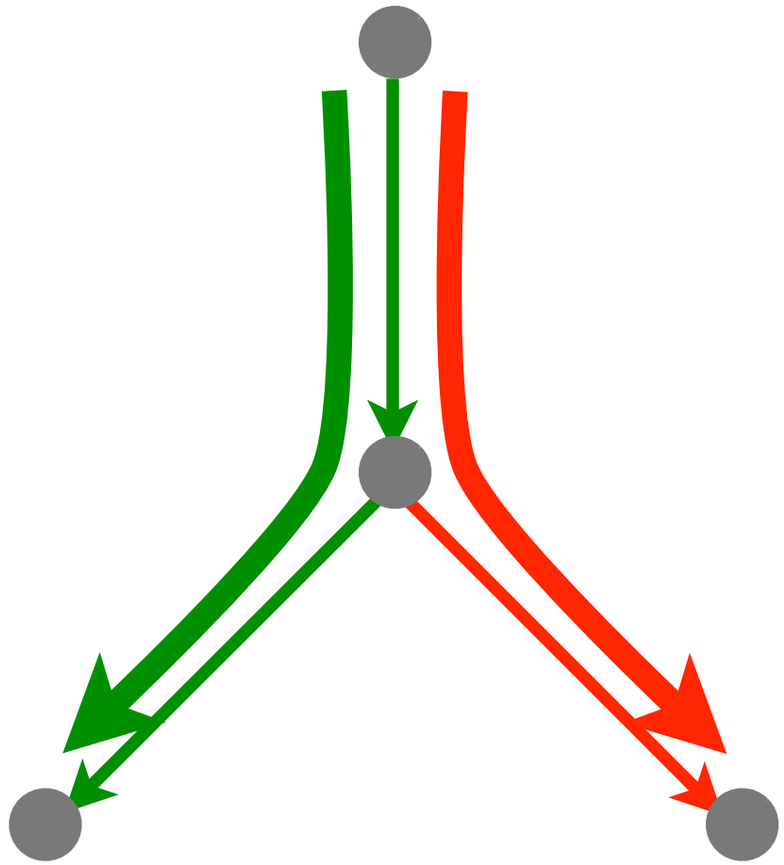

- Output: link properties
  - *average loss rate*
  - *average latency*
  - *congestion status*
  - *congestion probability*

Network tomography

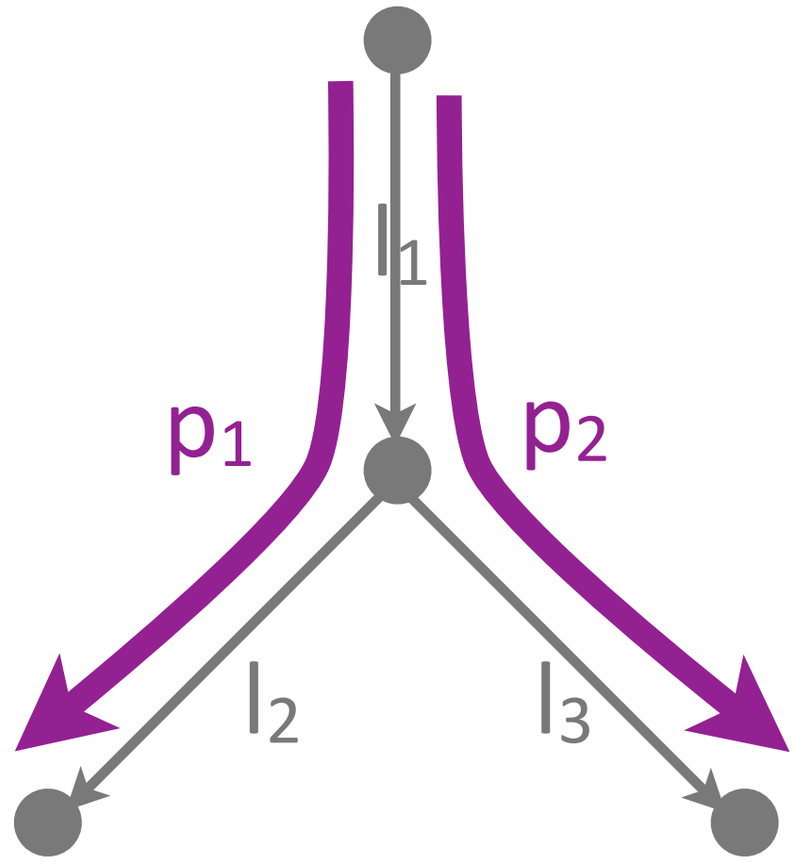

Network tomography

$P\{p_1 \text{ good}\} = P\{l_1 \text{ good}\} * P\{l_2 \text{ good}\}$

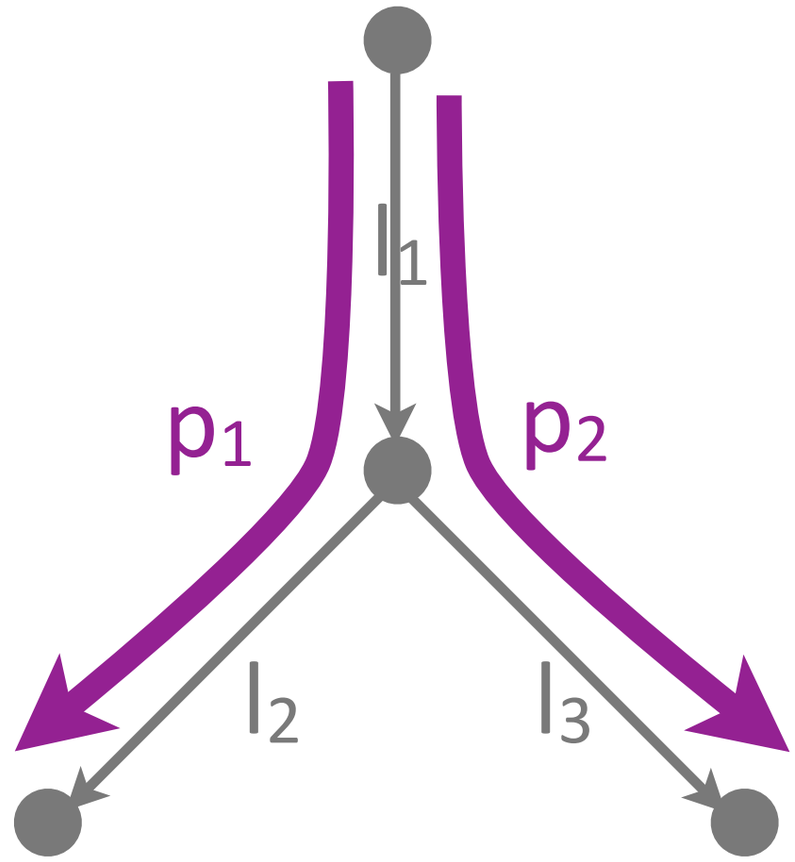

Network tomography

$$P\{p_1 \text{ good}\} = P\{l_1 \text{ good}\} * P\{l_2 \text{ good}\}$$

$$\log(P\{p_1 \text{ good}\}) = \log(P\{l_1 \text{ good}\}) + \log(P\{l_2 \text{ good}\})$$

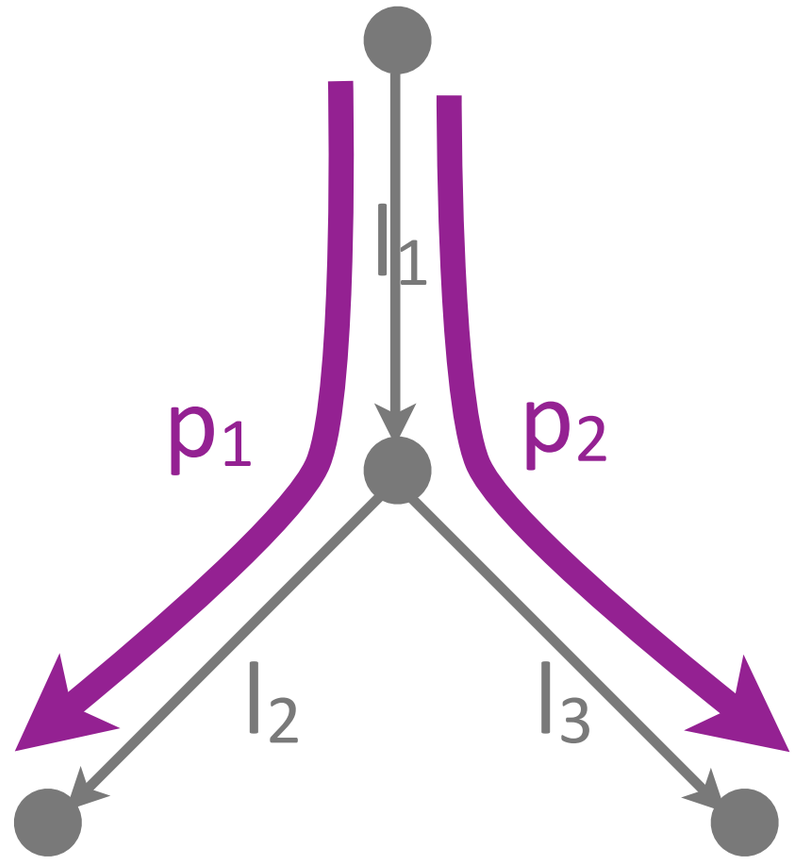

Network tomography

$$P\{p_1 \text{ good}\} = P\{l_1 \text{ good}\} * P\{l_2 \text{ good}\}$$

$$\log(P\{p_1 \text{ good}\}) = \log(P\{l_1 \text{ good}\}) + \log(P\{l_2 \text{ good}\})$$

$$Y_1 = X_1 + X_2$$

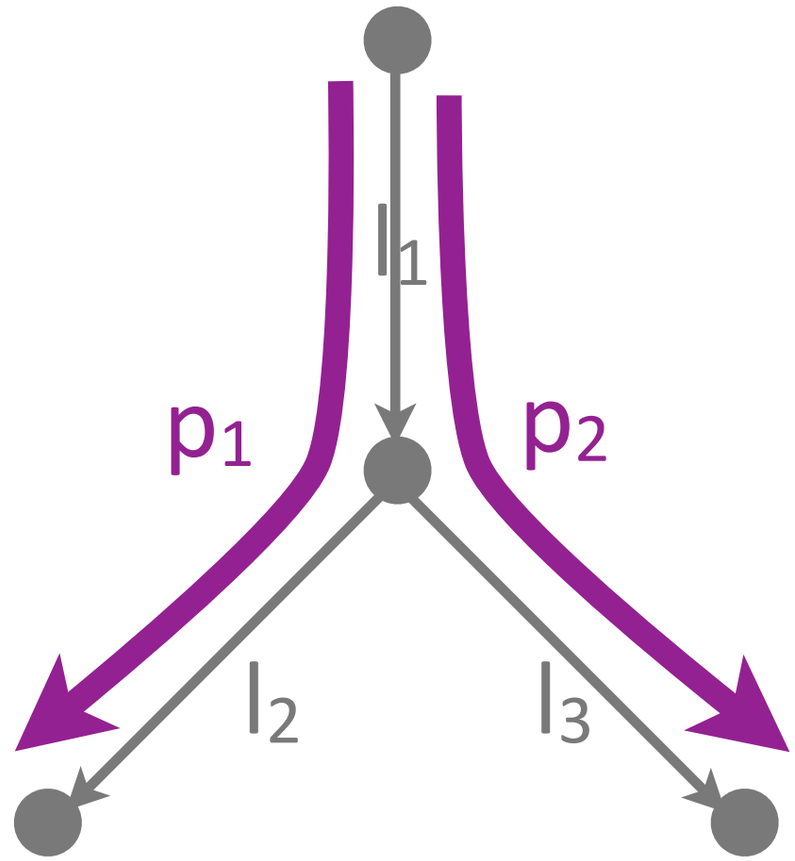

Network tomography

$Y_1 = X_1 + X_2$

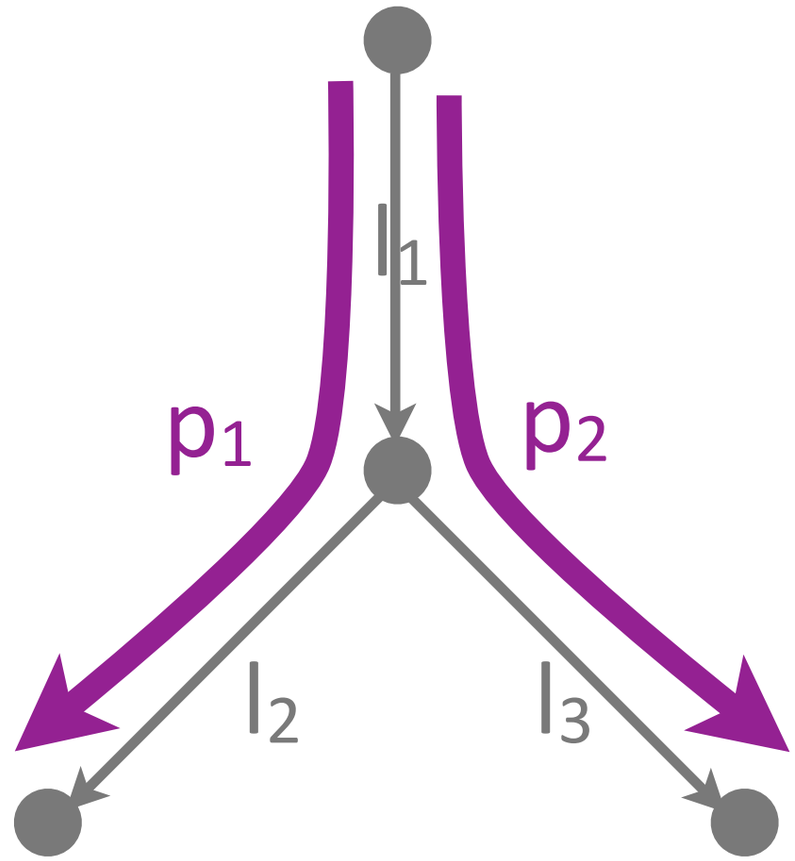

Network tomography

link properties (unknowns)

$Y_1 = X_1 + X_2$

$Y_2 = X_1 + X_3$

$Y_{12} = X_1 + X_2 + X_3$

e2e path measurements

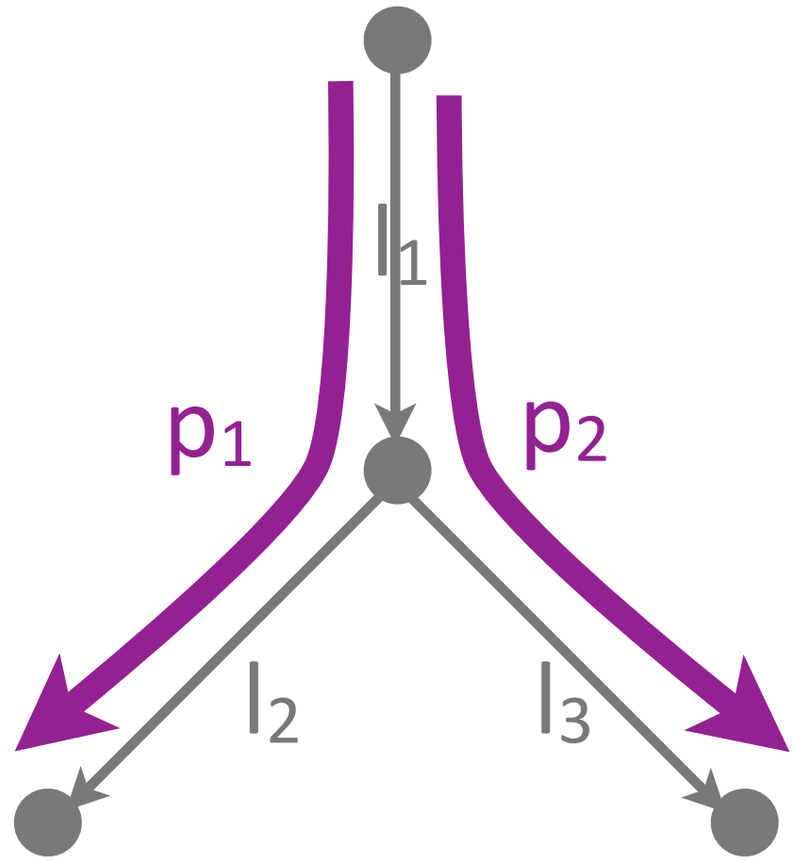

Tomography and neutrality

- Tomography fundamentally assumes network neutrality

- A link has the same property for all paths that traverse that link

Neutrality and tomography

- In a neutral network, any tomographic system of equations has a solution

- If we find an unsolvable system, this indicates neutrality violation

Outline

- Observability of neutrality violations
- Localization of neutrality violations
- Algorithm + results

Outline

- Observability of neutrality violations
- Localization of neutrality violations
- Algorithm + results

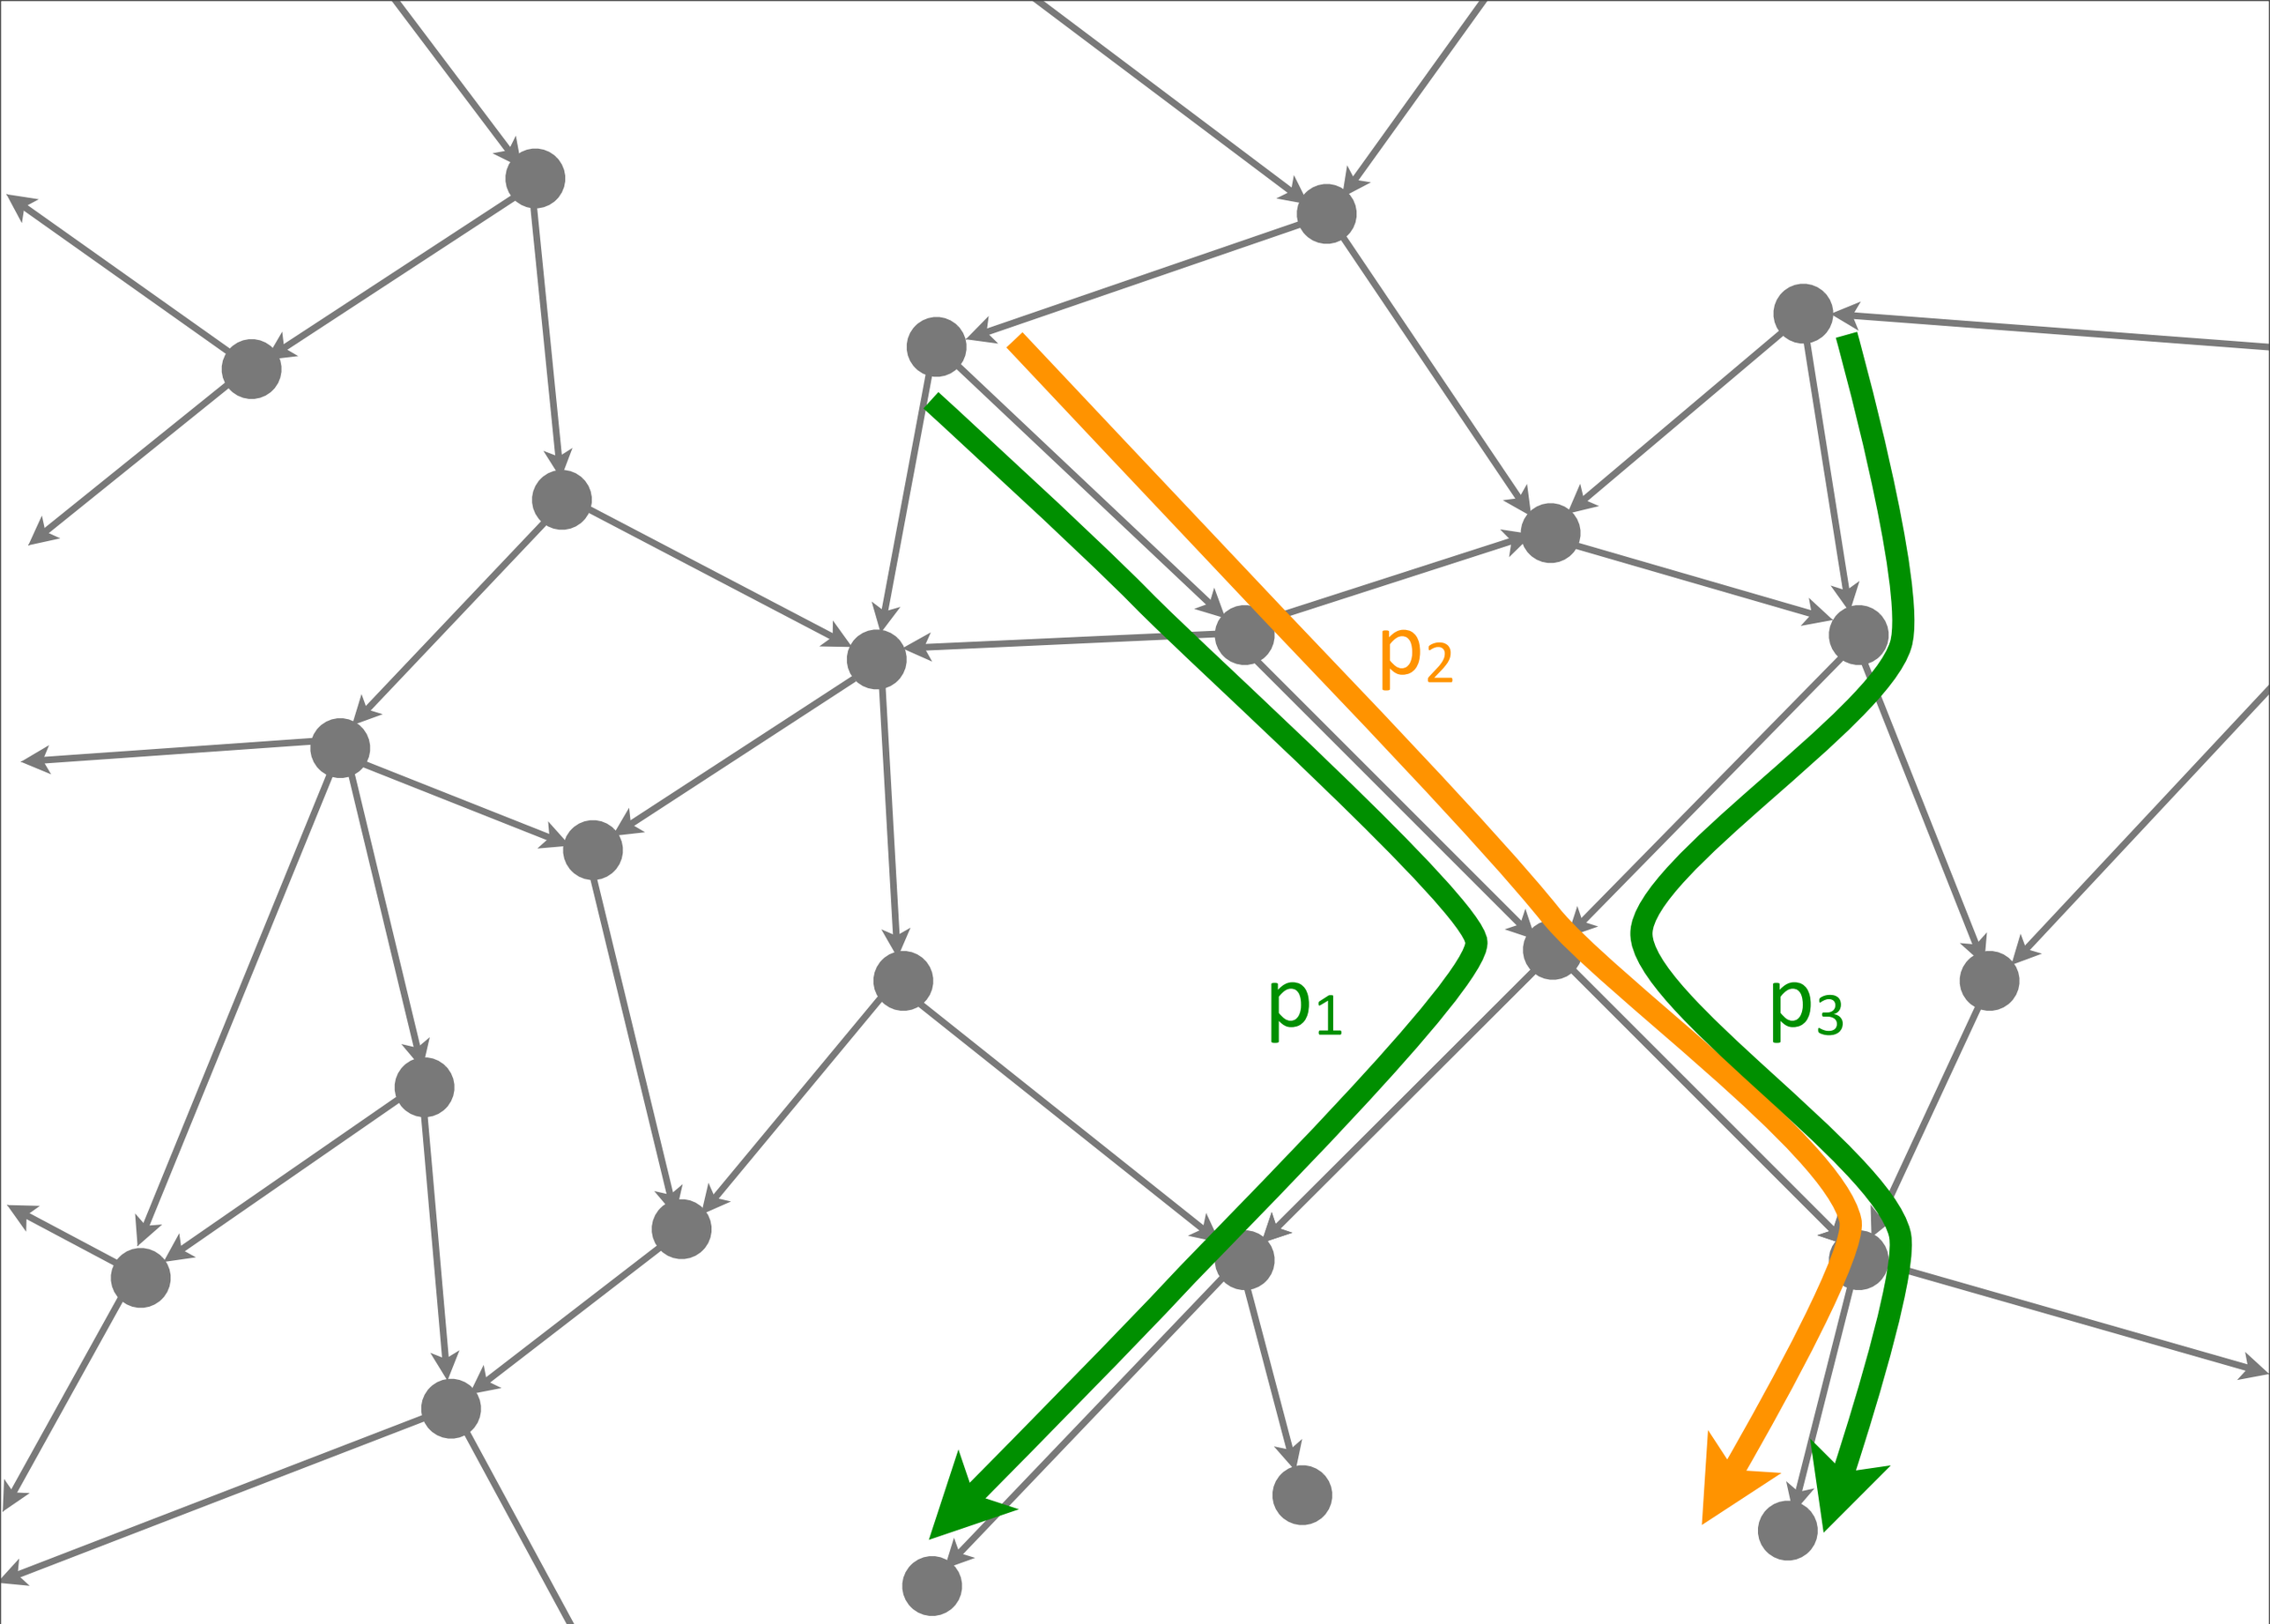
p_2
p_1
p_3

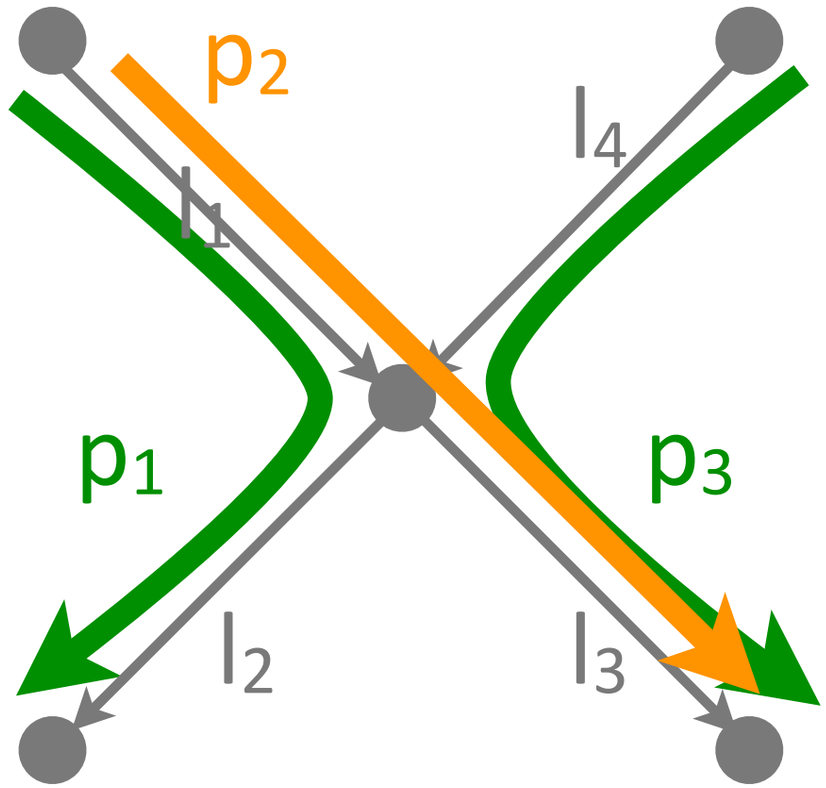
p2
l4
l1
p1
p3
l2
l3

$$Y_1 = X_1 + X_2$$

$$Y_2 = X_1 + X_3$$

$$Y_3 = X_3 + X_4$$

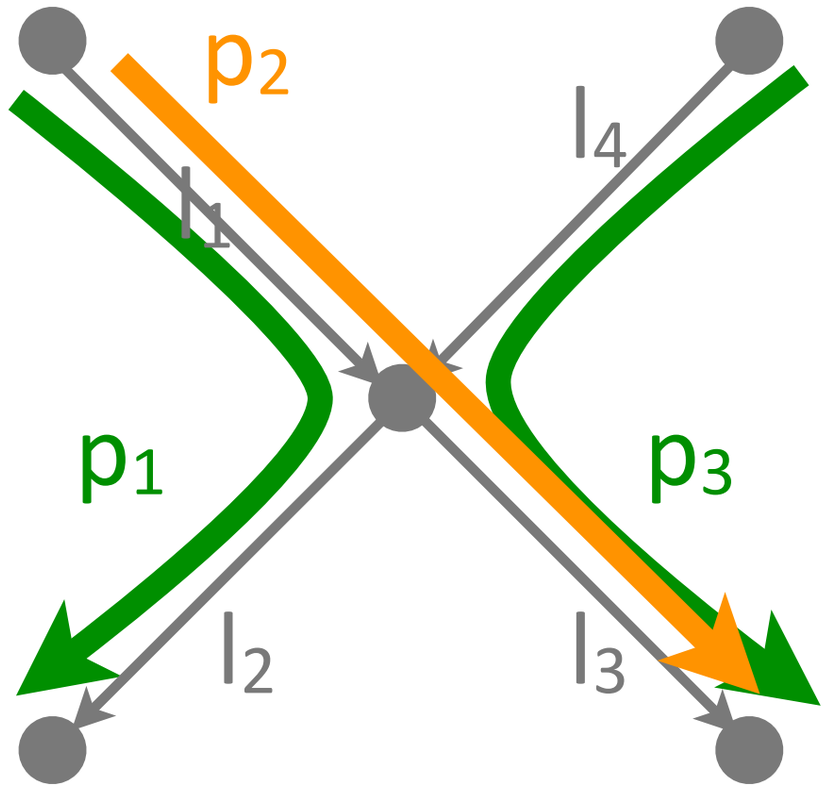

$$0 = X_1 + X_2$$

$$0 > Y_2 = X_1 + X_3$$

$$0 = X_3 + X_4$$

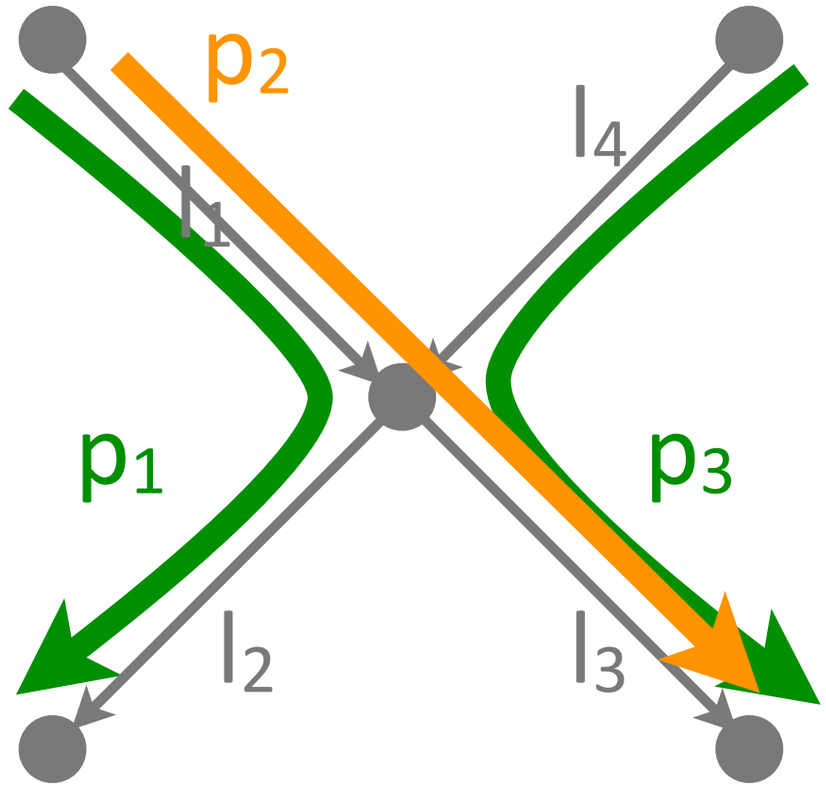

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + \cancel{X_3}^{0}$$

$$0 = \cancel{X_3}^{0} + \cancel{X_4}^{0}$$

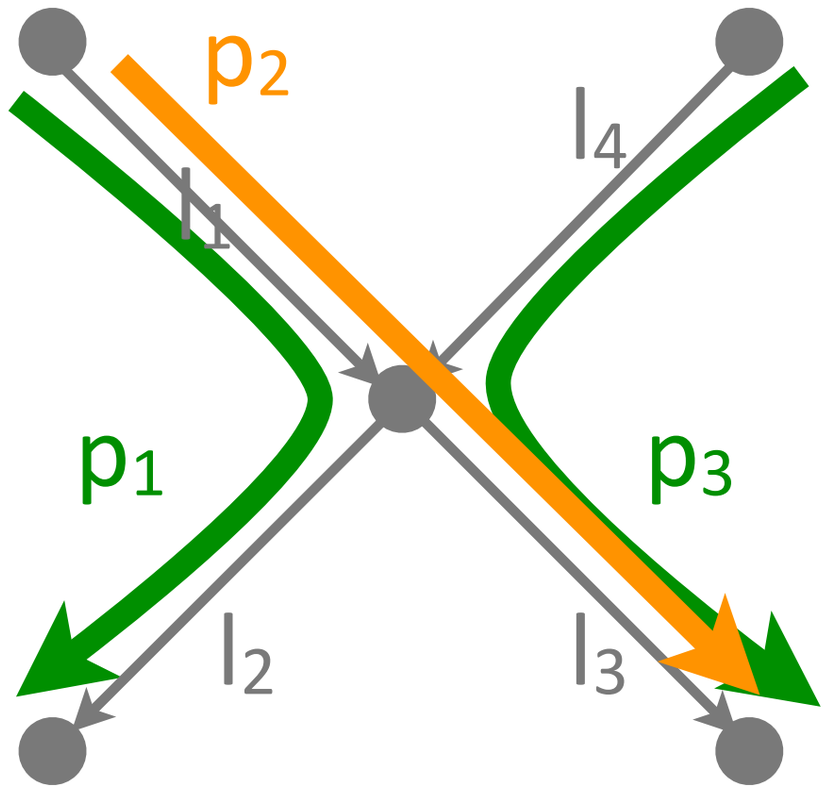

Observable neutrality violation

$$0 = \cancel{x_1}^{0} + \cancel{x_2}^{0}$$

$$0 > Y_2 = \cancel{x_1}^{0} + \cancel{x_3}^{0}$$

$$0 = \cancel{x_3}^{0} + \cancel{x_4}^{0}$$

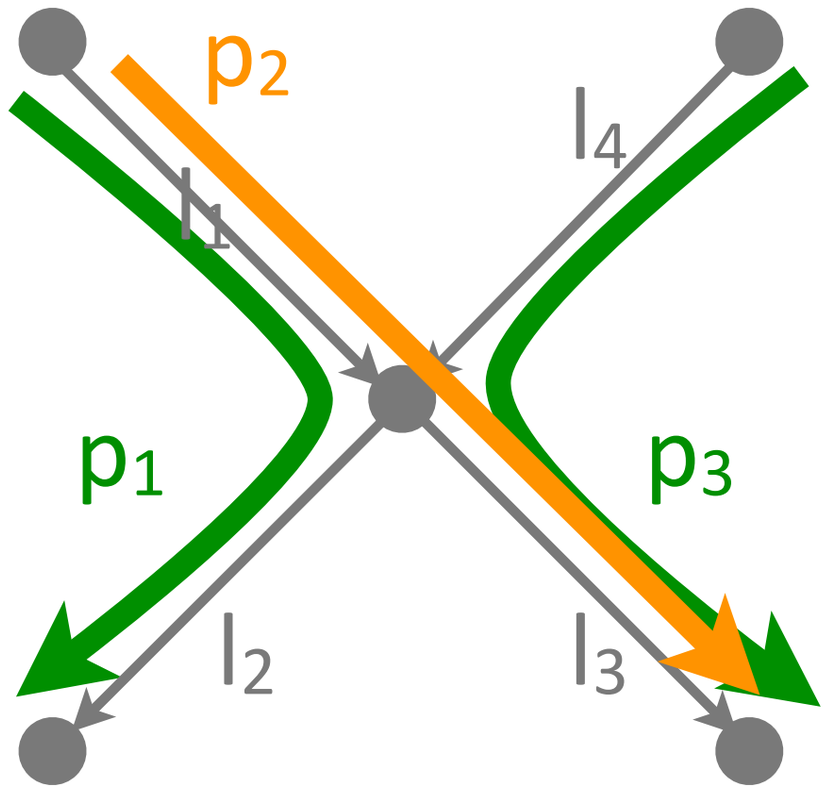

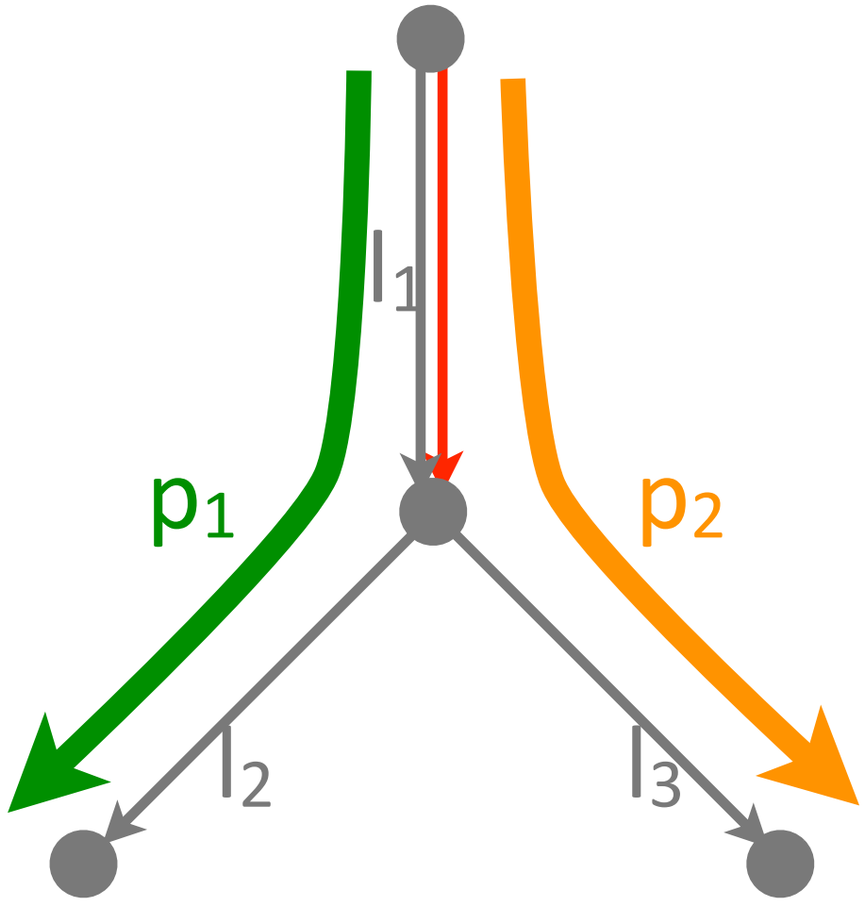
l1
p1
p2
l2
l3

$$Y_1 = X_1 + X_2$$

$$Y_2 = X_1 + X_3$$

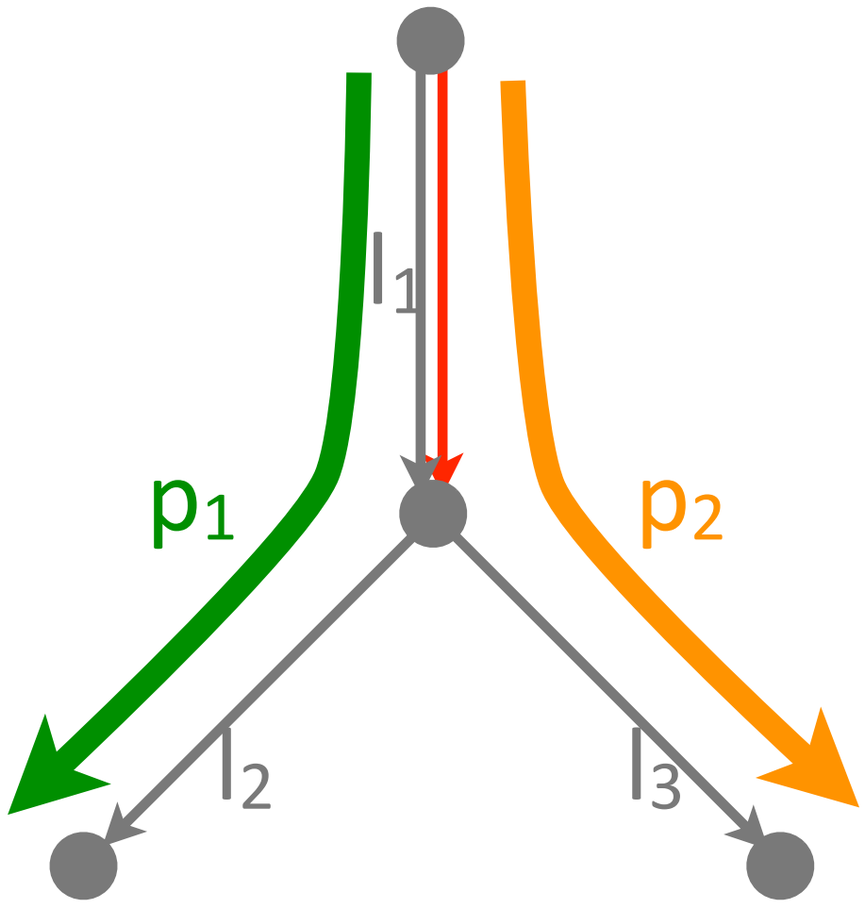

$$0 = X_1 + X_2$$

$$0 > Y_2 = X_1 + X_3$$

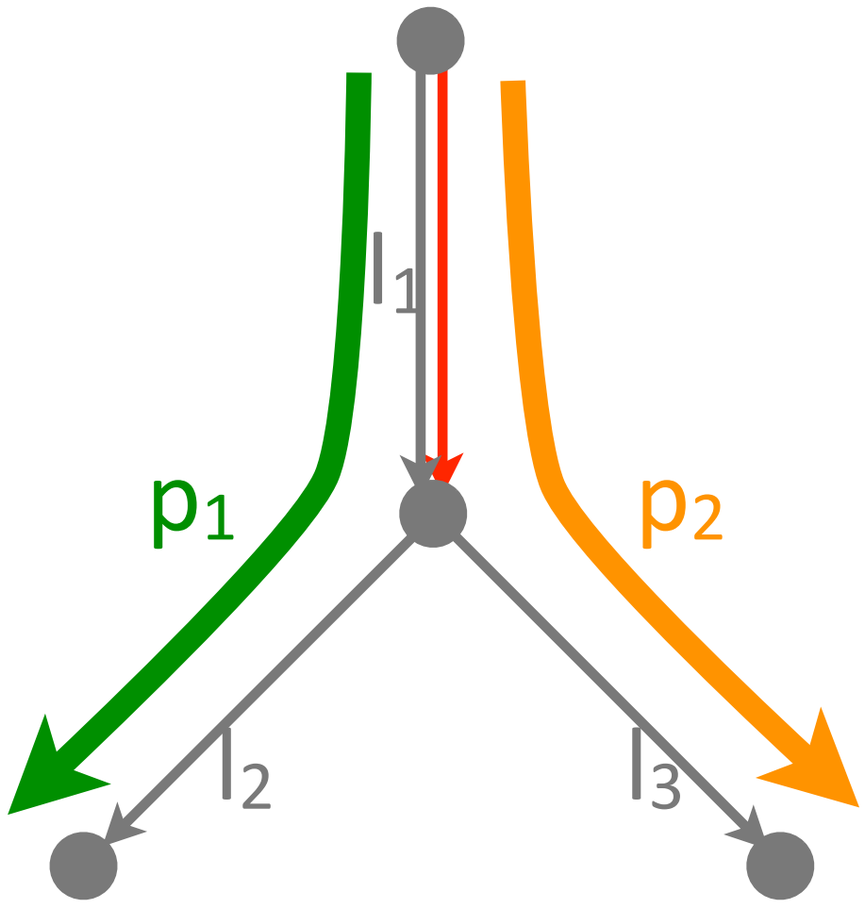

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

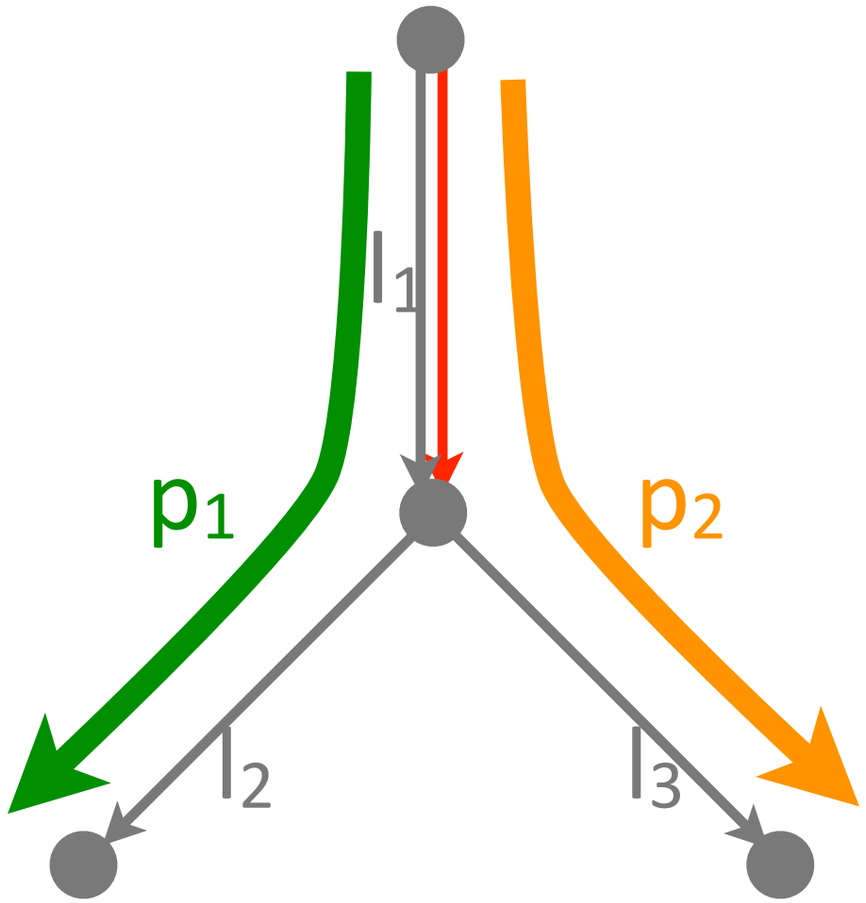

Non-observable neutrality violation

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

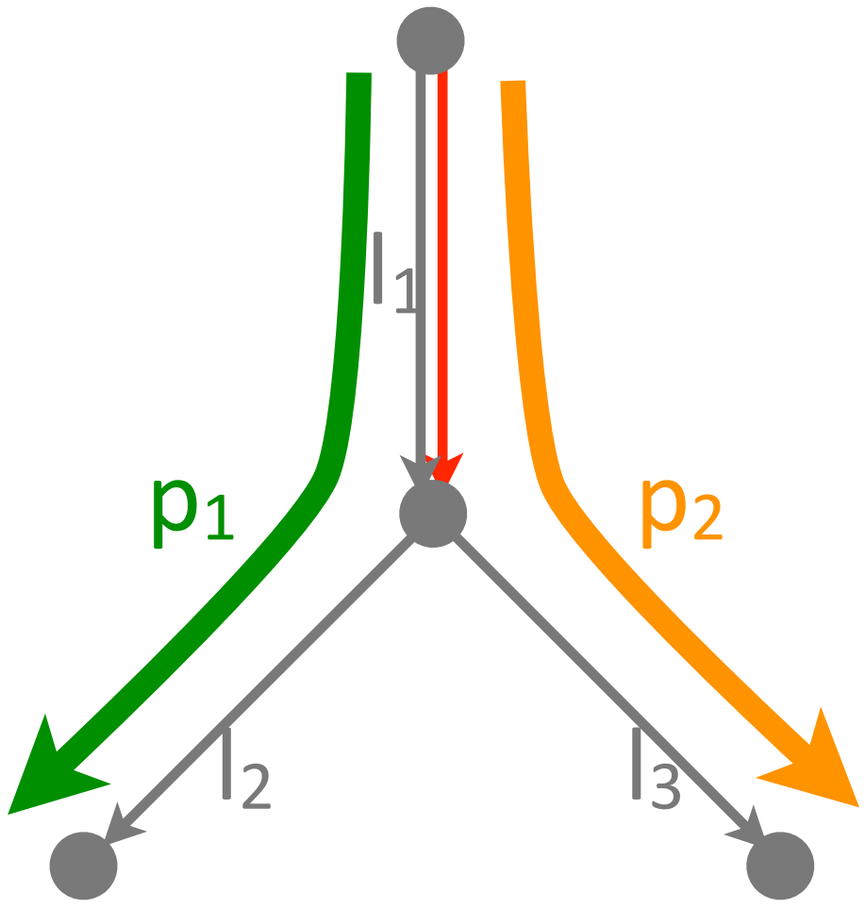

Equivalent neutral network

- Any non-neutral network has a neutral equivalent
- A neutral network that:

  given the same traffic input

  would produce the same external observations

  as the non-neutral network.

Equivalent neutral network

- A purely theoretical construct

- Enables to formally specify which neutrality violations are externally observable

Non-observable neutrality violation

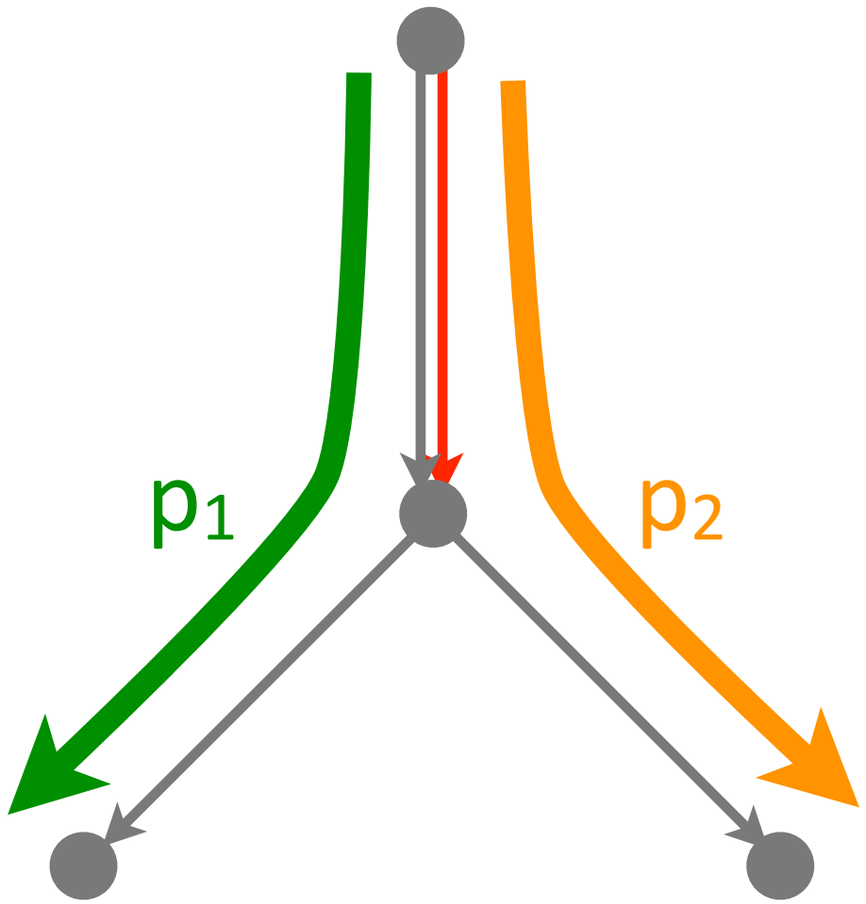

neutral equivalent

original

Non-observable neutrality violation

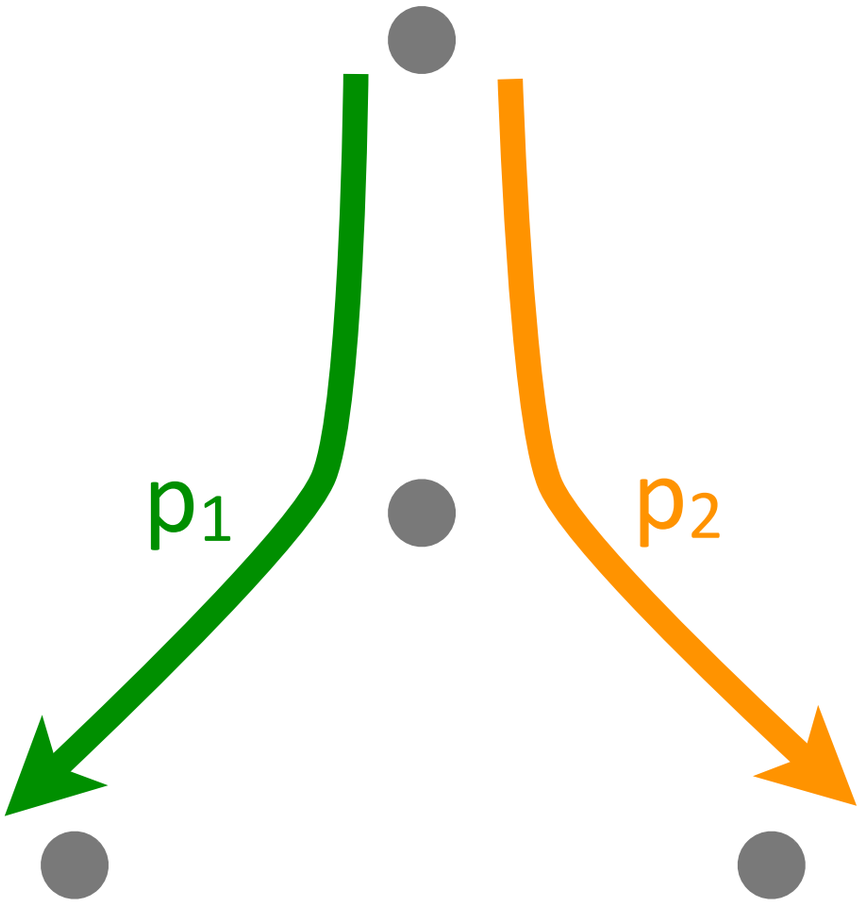

neutral equivalent

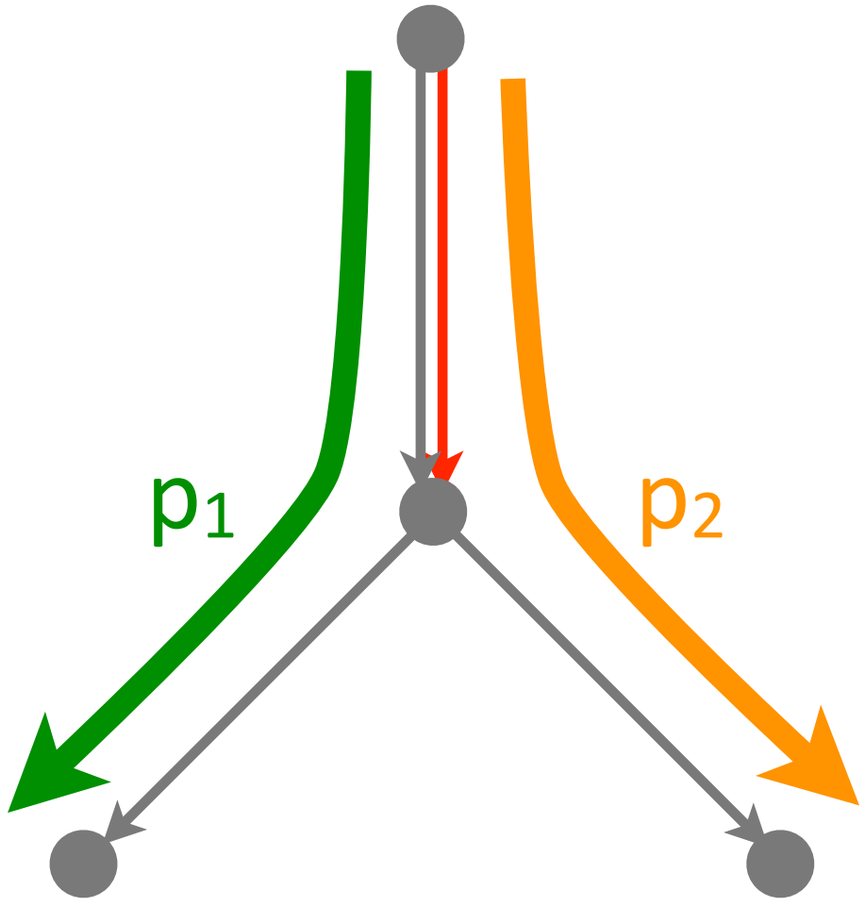

original

Non-observable neutrality violation

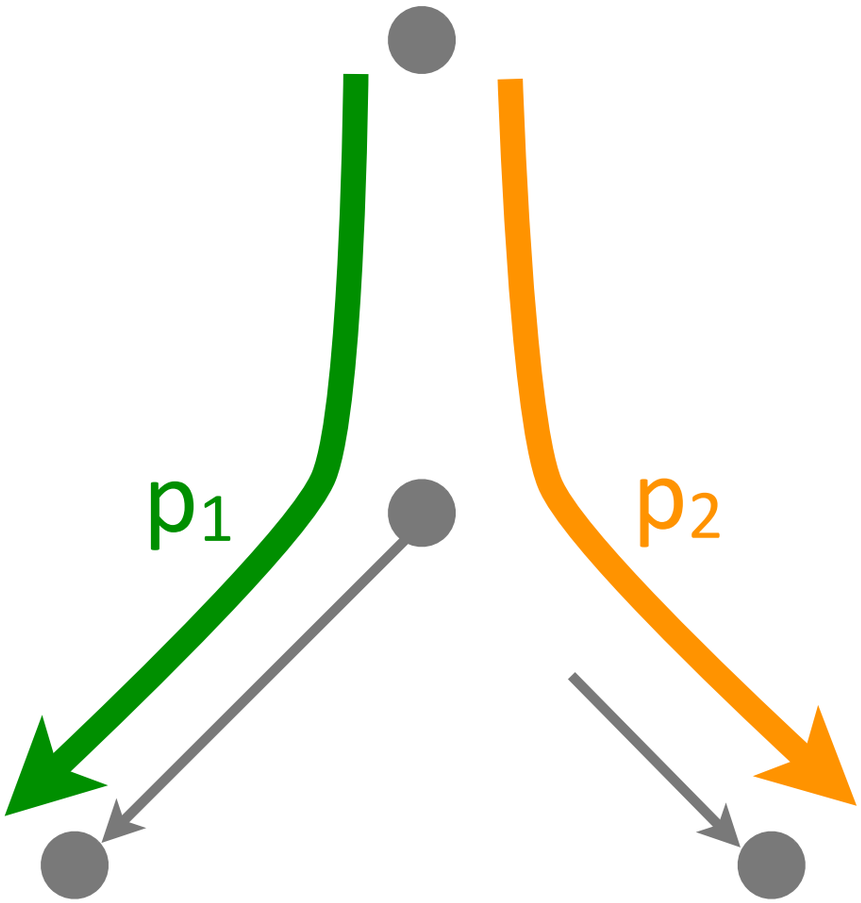

neutral equivalent

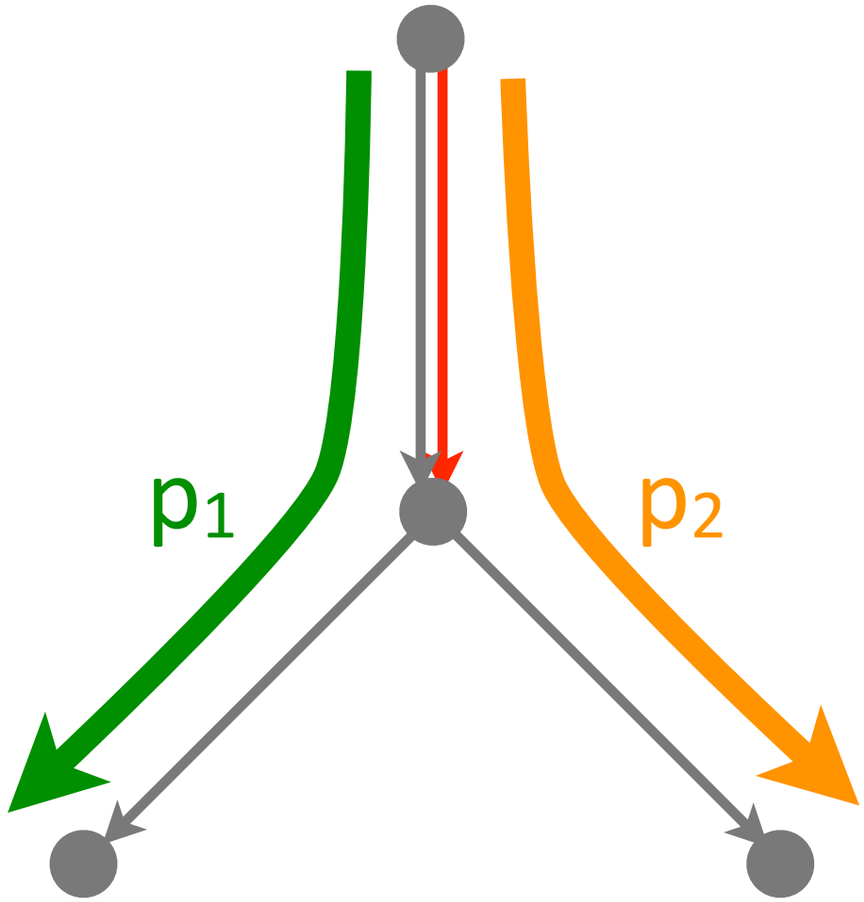

original

Non-observable neutrality violation

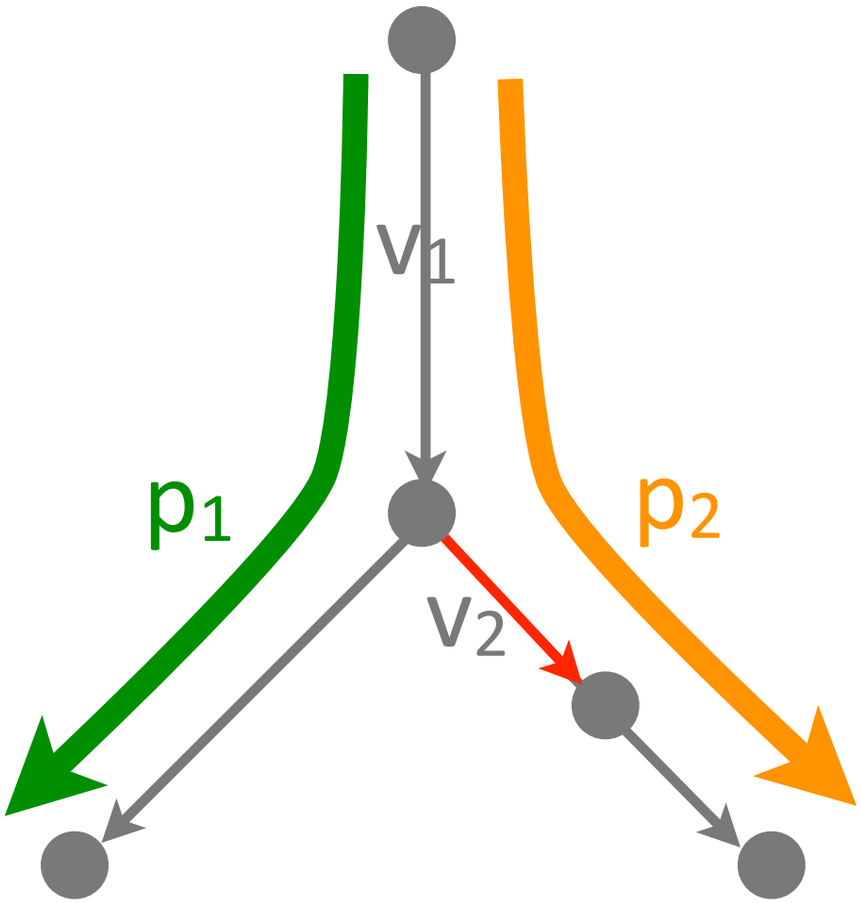

neutral equivalent

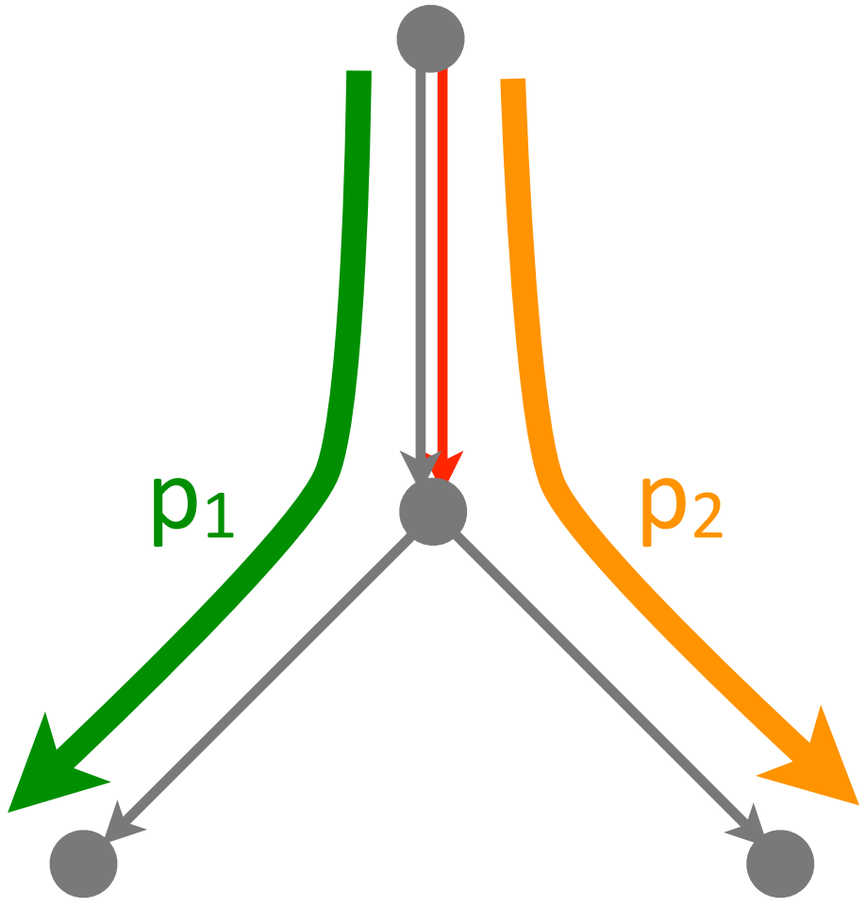

original

Observable neutrality violation

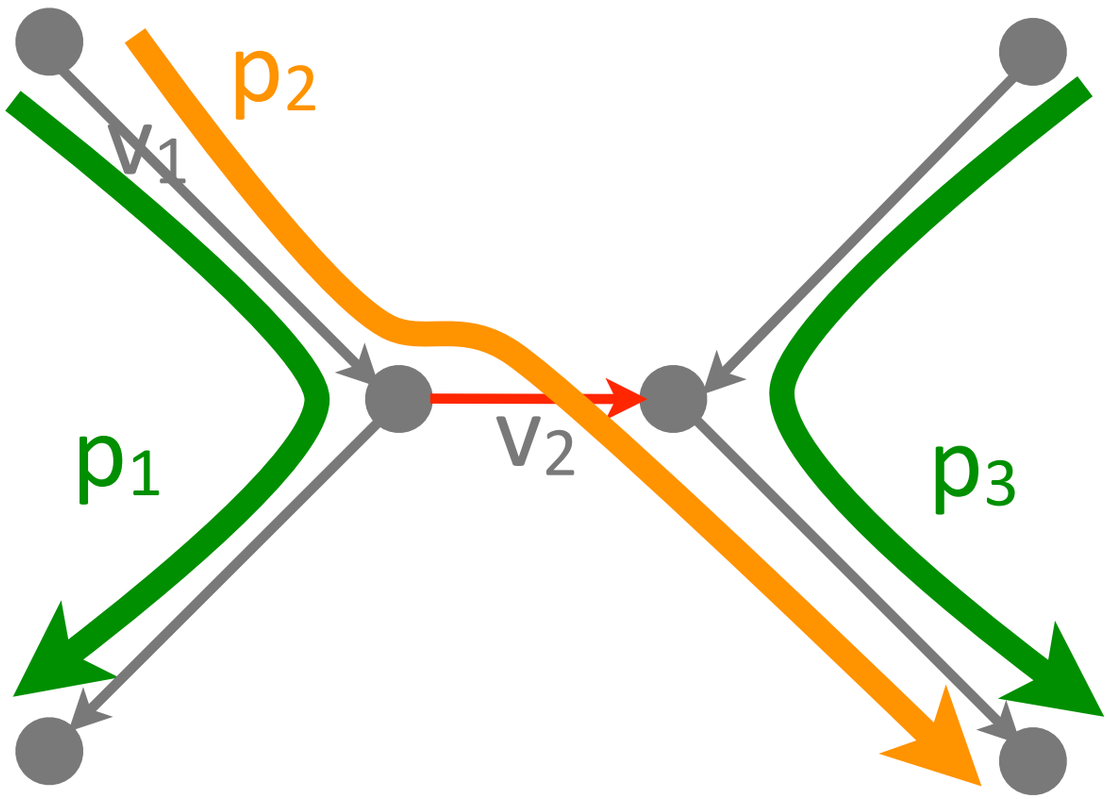

neutral equivalent

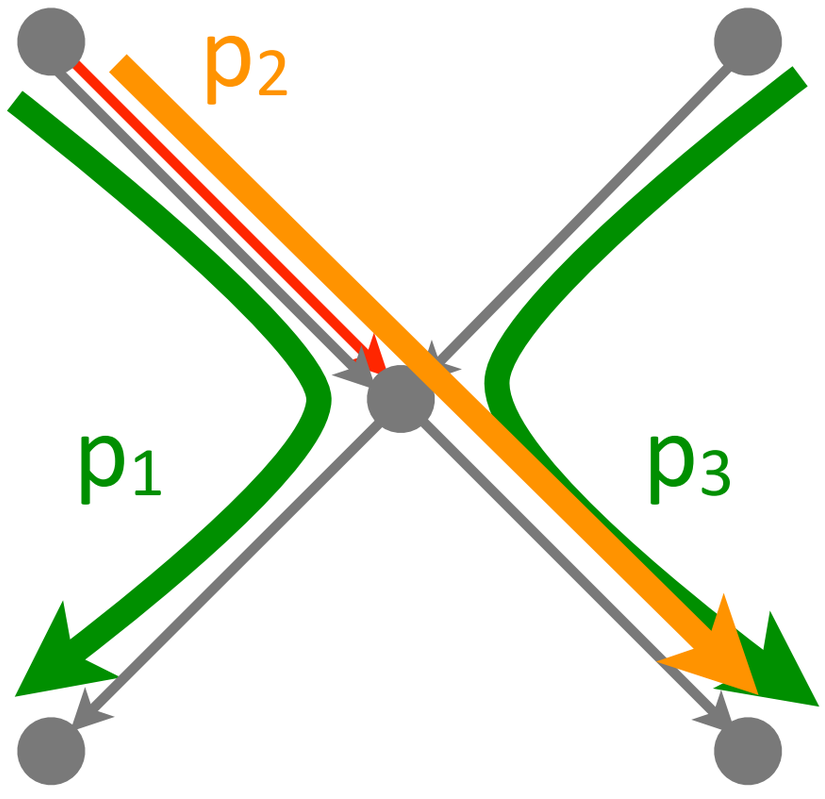

original

Observability

- Neutrality violation is externally observable,

  if and only if

  there exists at least one

  distinguishable virtual link

  in the equivalent neutral network.

Outline

- Observability of neutrality violation
- Localization of neutrality violations
- Algorithm + results

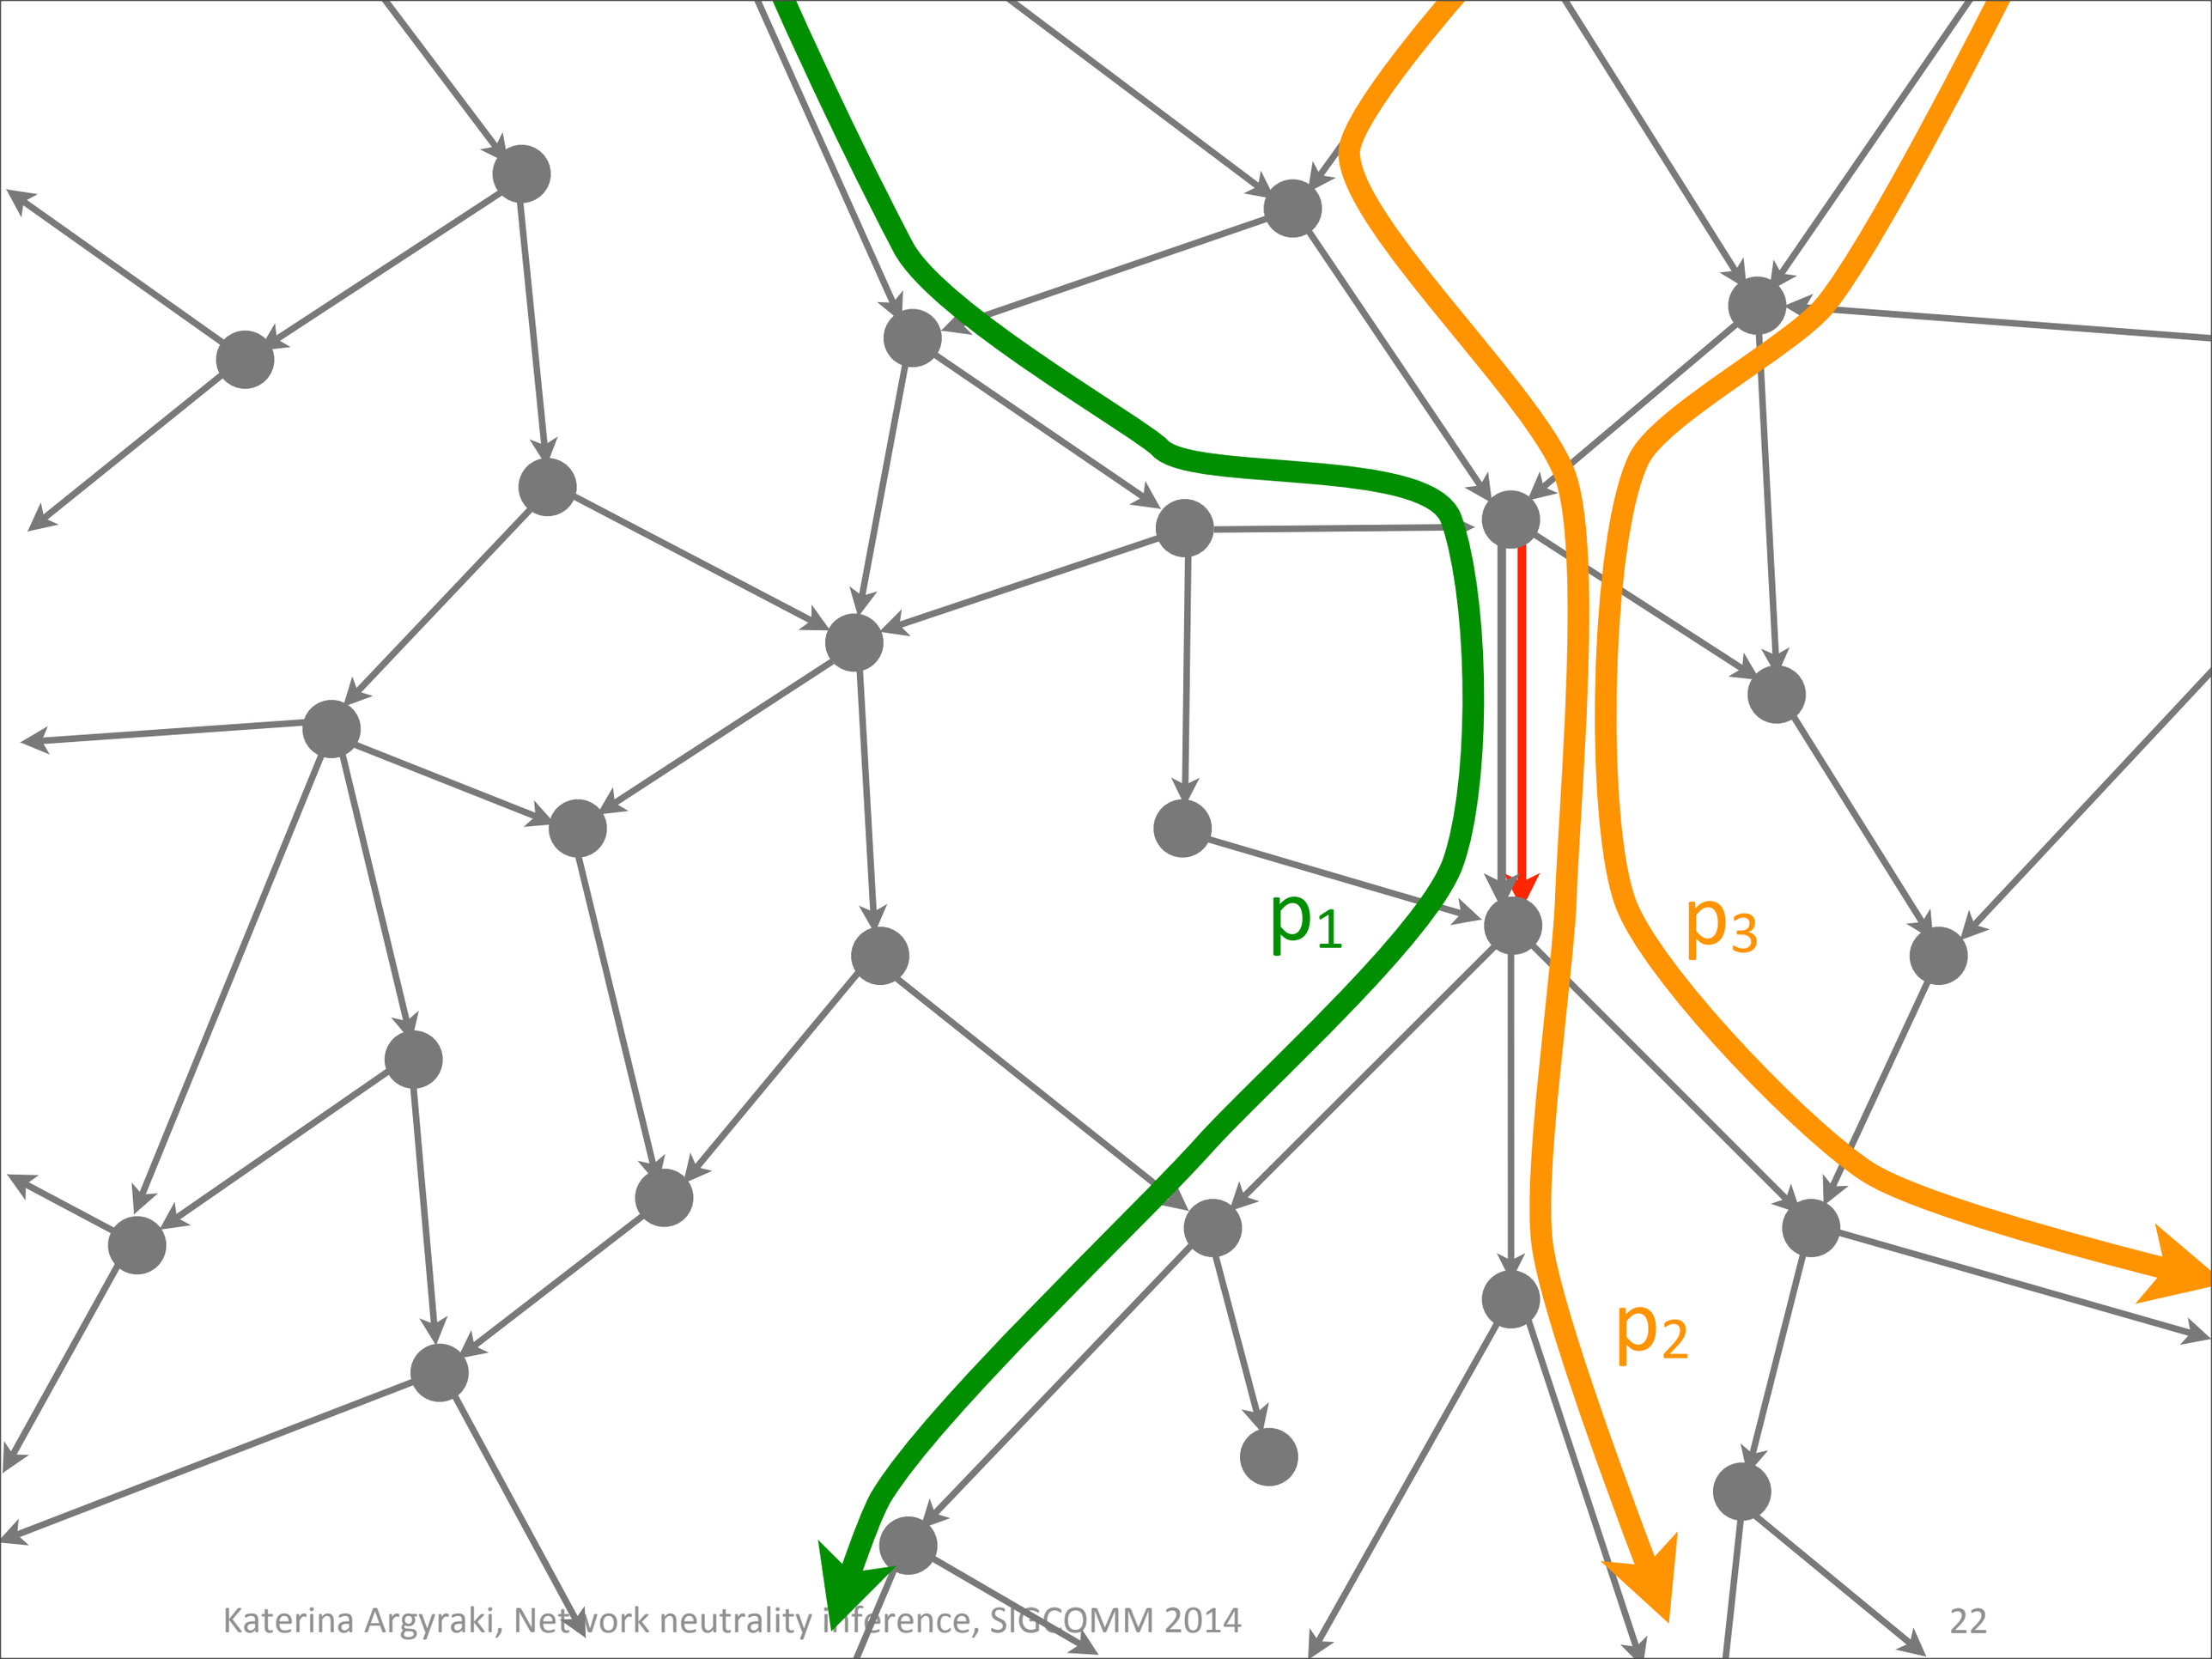
p1
p3
p2

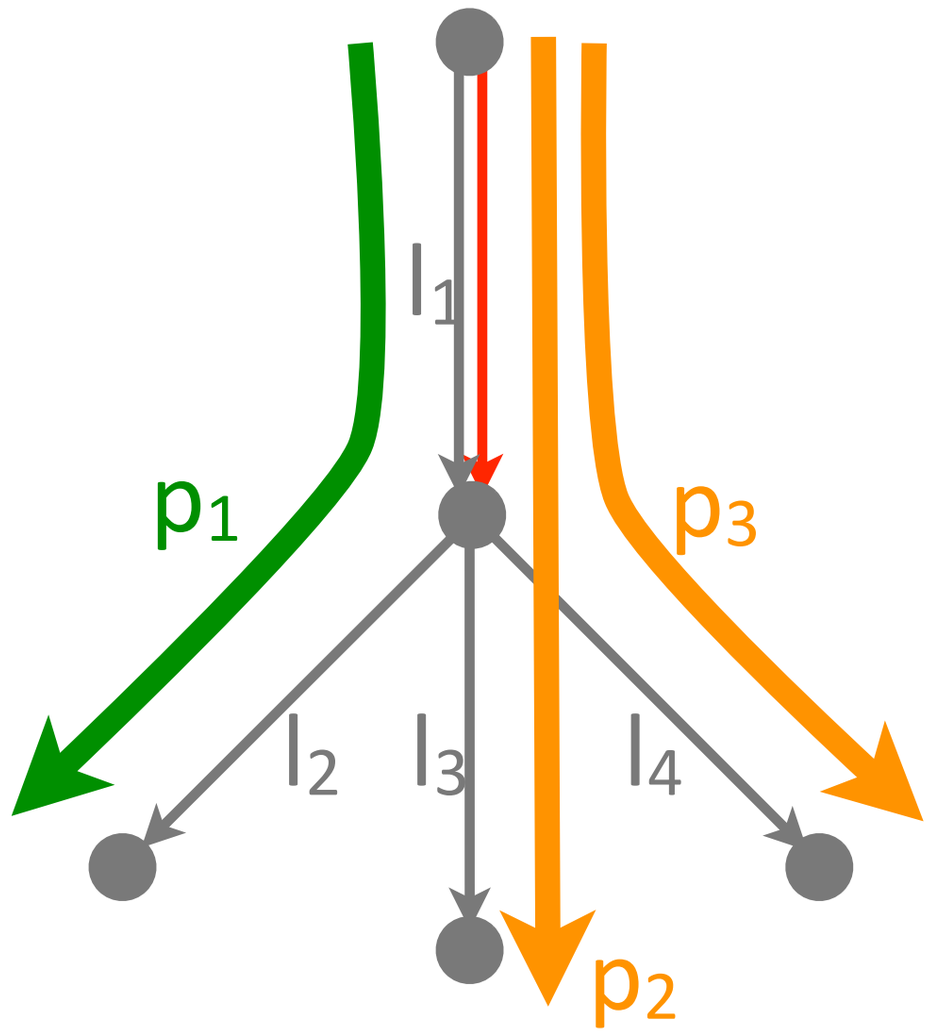
l1
p1
p3
l2
l3
l4
p2

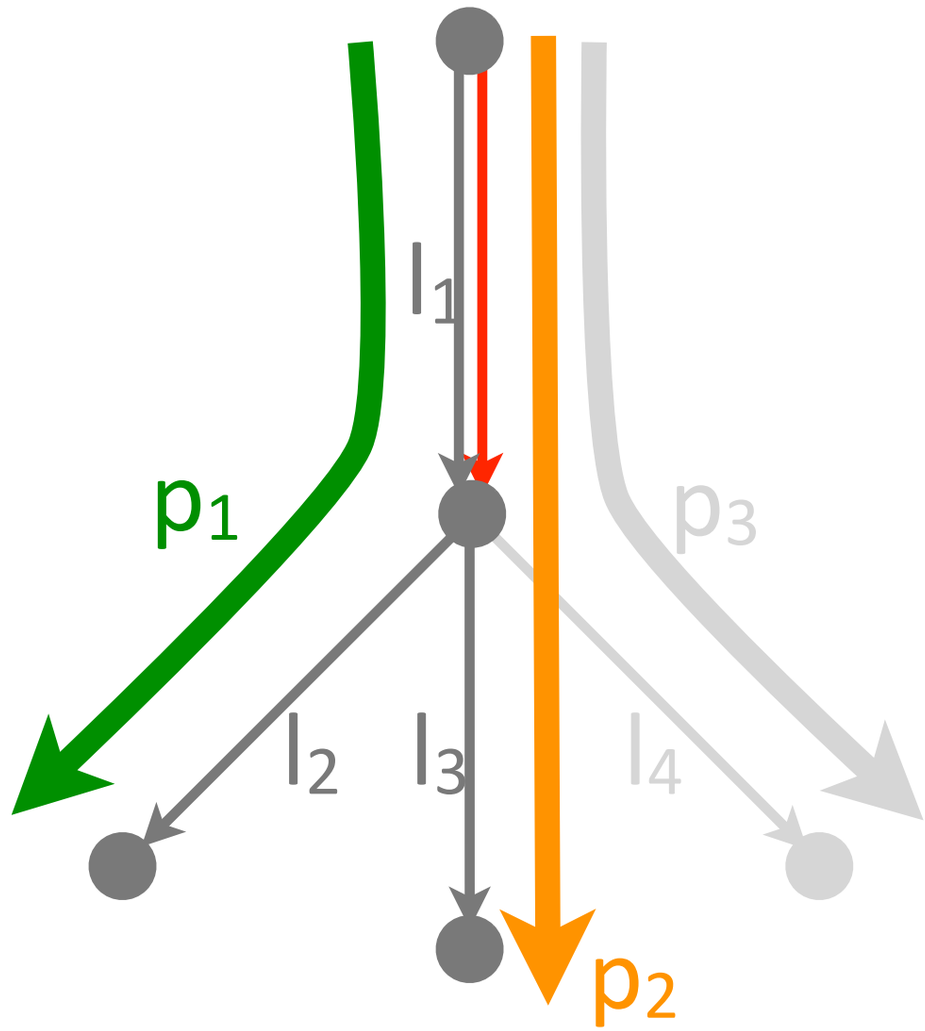
l1
p1
p3
l2
l3
l4
p2

$$0 = X_1 + X_2$$

$$0 > Y_2 = X_1 + X_3$$

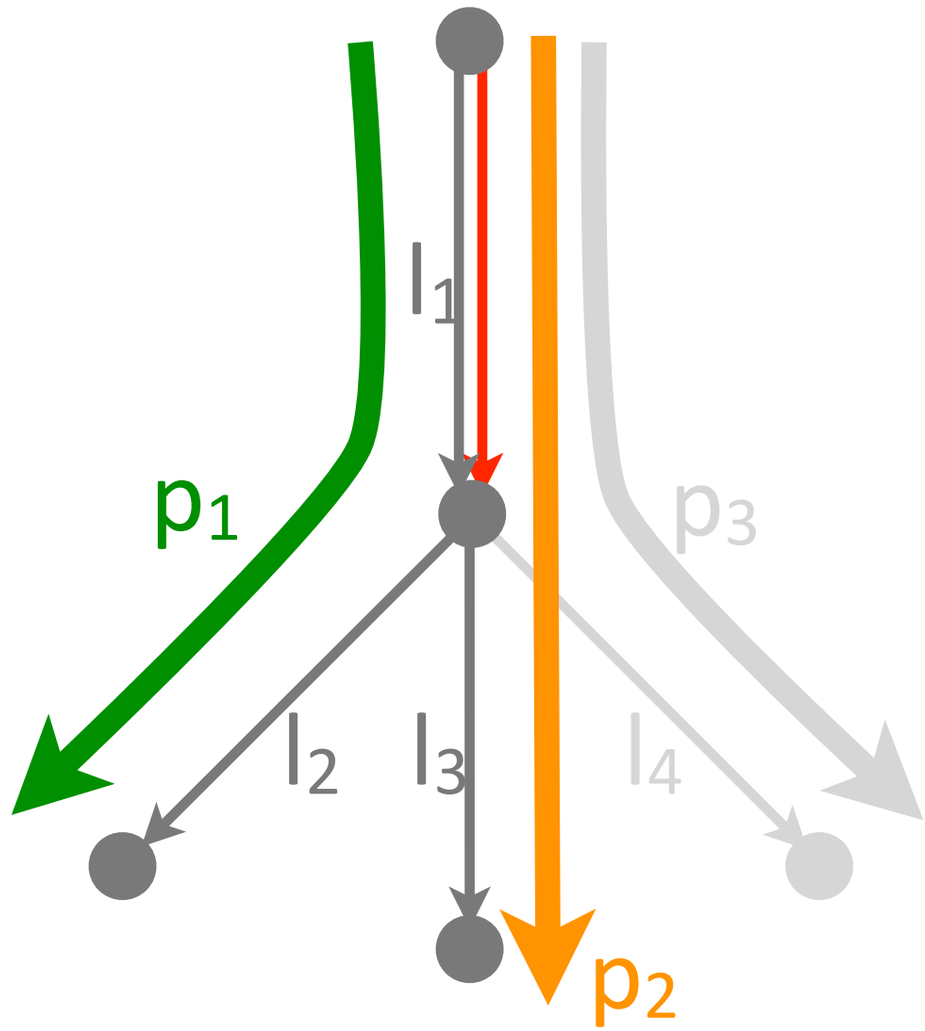

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

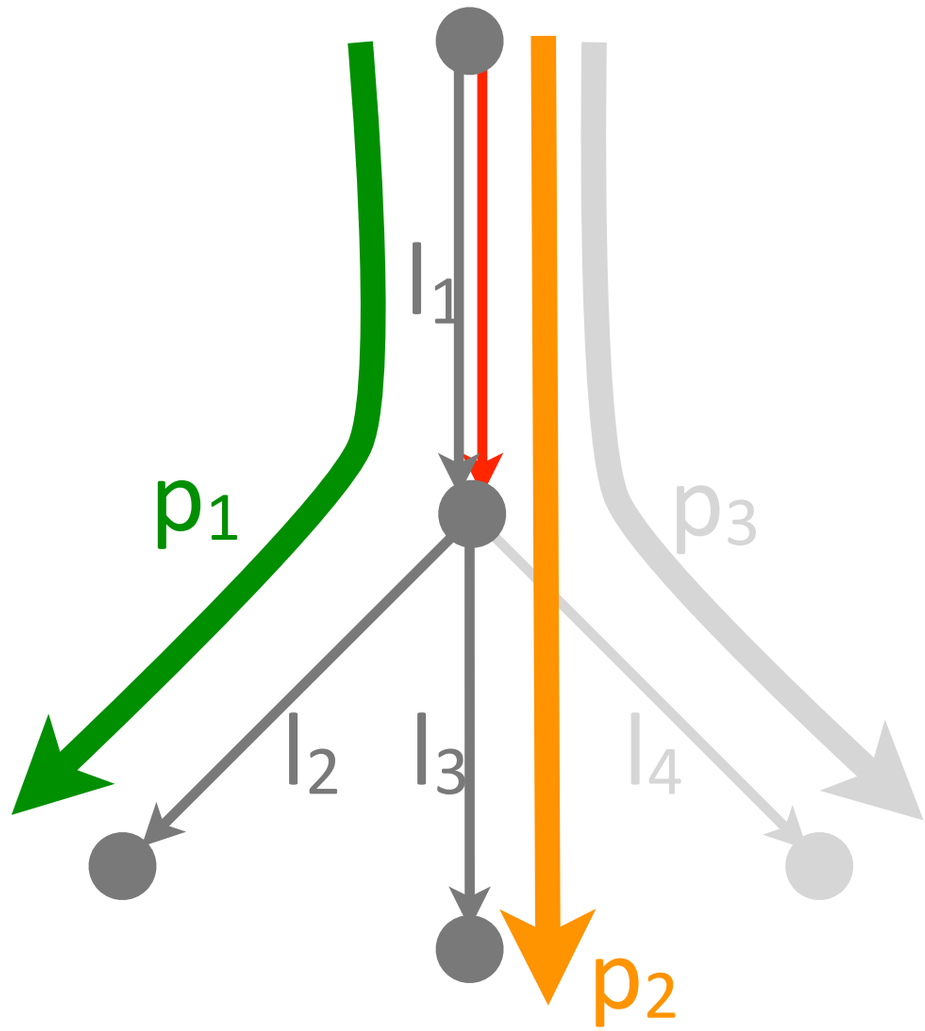

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

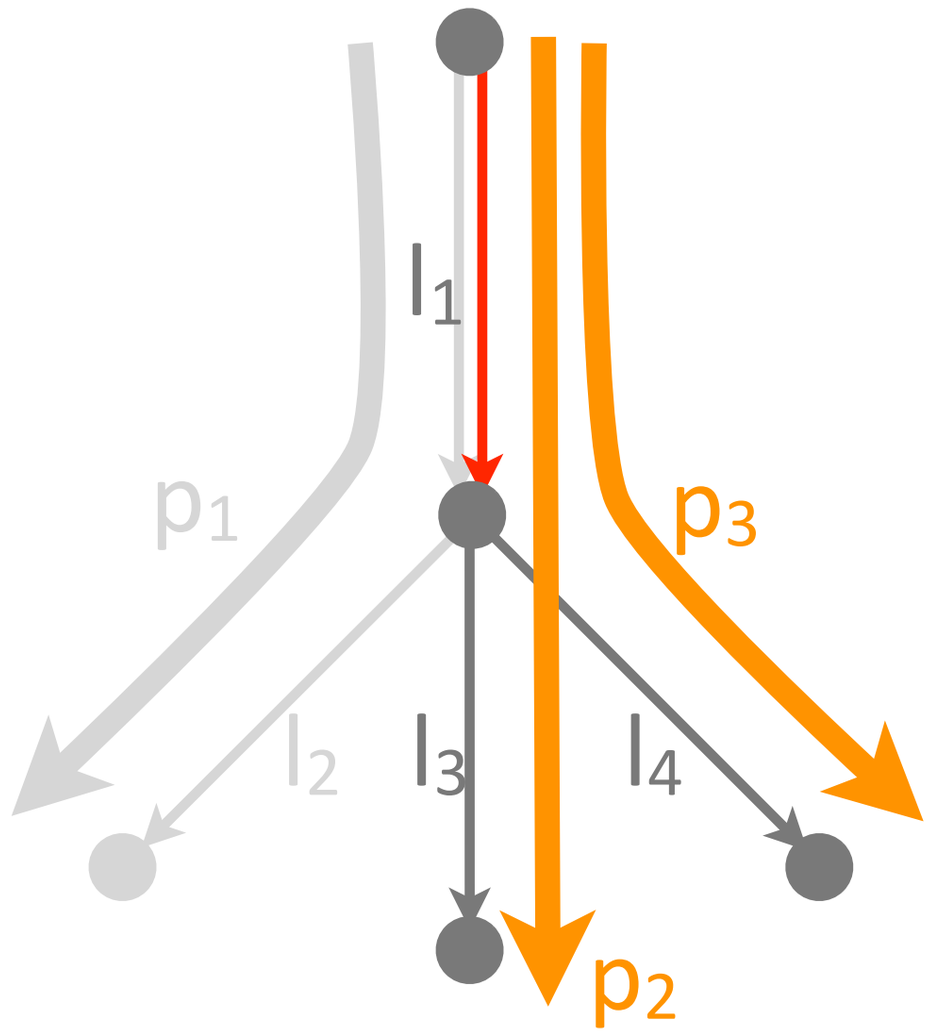

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

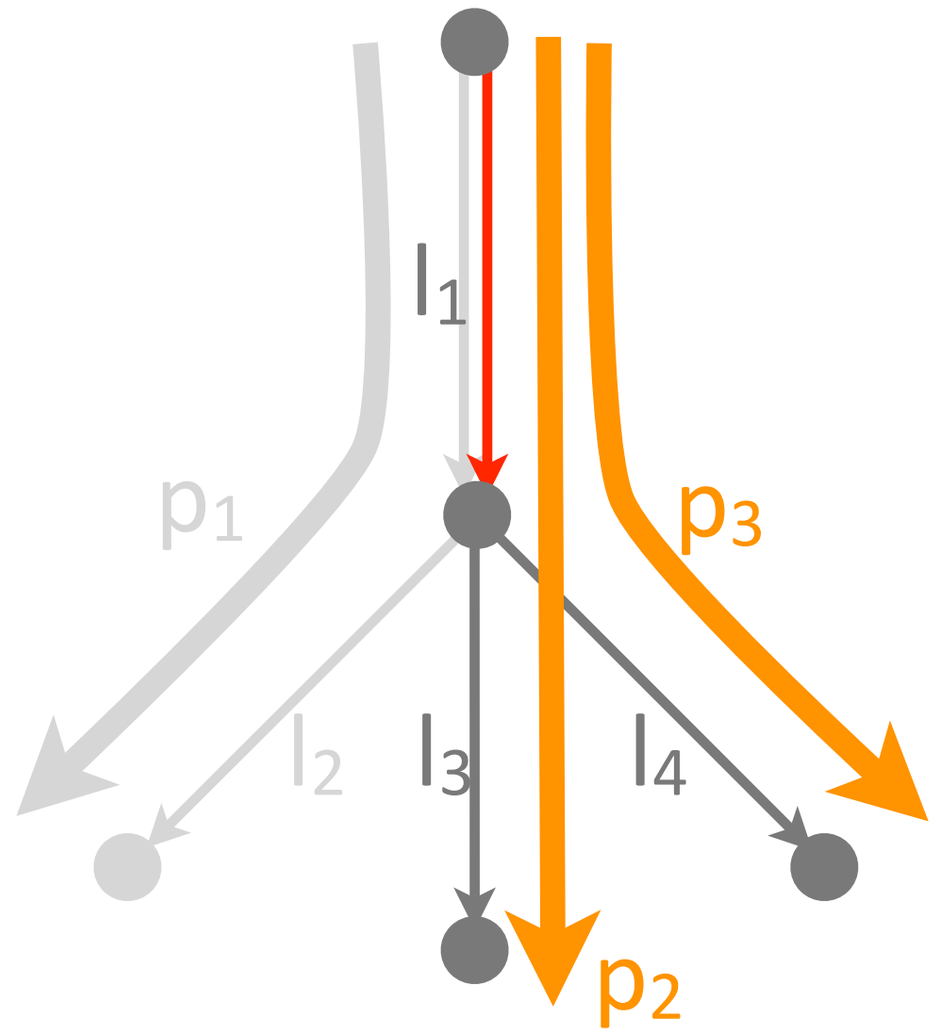

$$Y_2 = X_1 + X_3$$

$$Y_3 = X_1 + X_4$$

$$Y_{23} = X_1 + X_3 + X_4$$

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

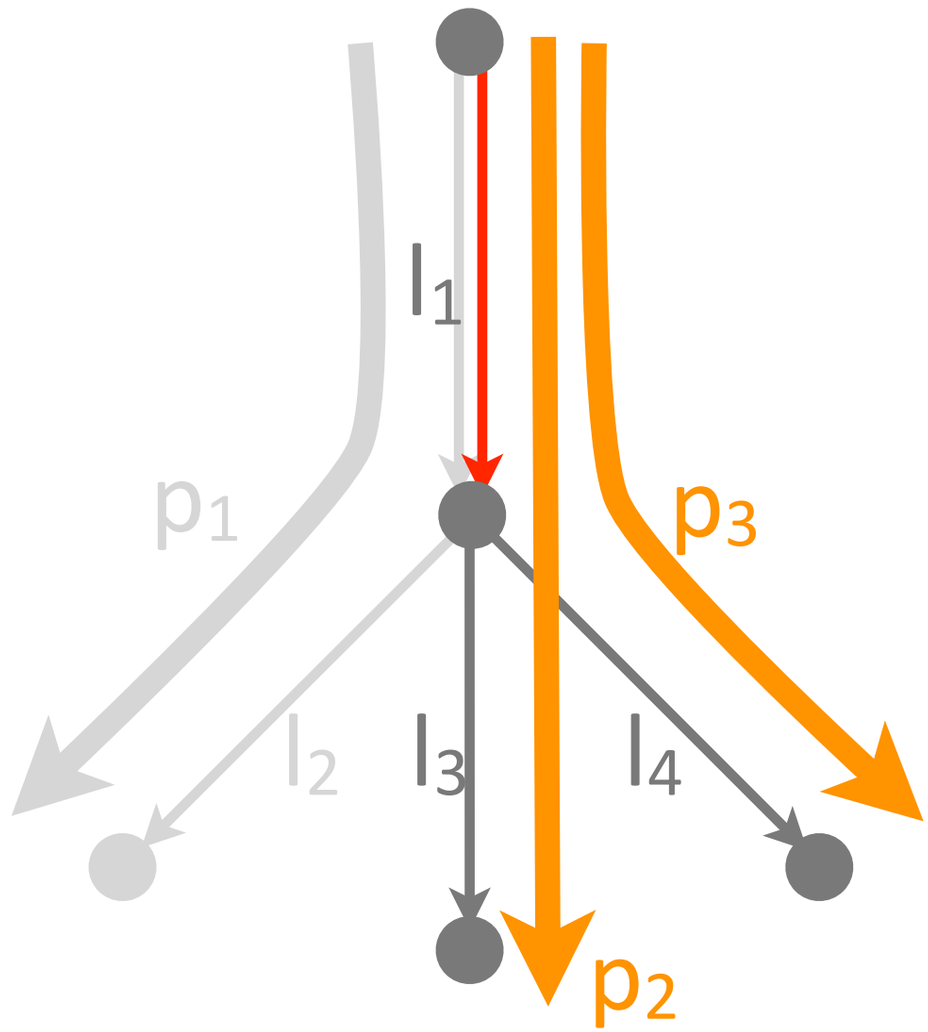

$$Y_2 = X_1 + X_3$$

$$Y_3 = X_1 + X_4$$

$$Y_{23} = X_1 + X_3 + X_4$$

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

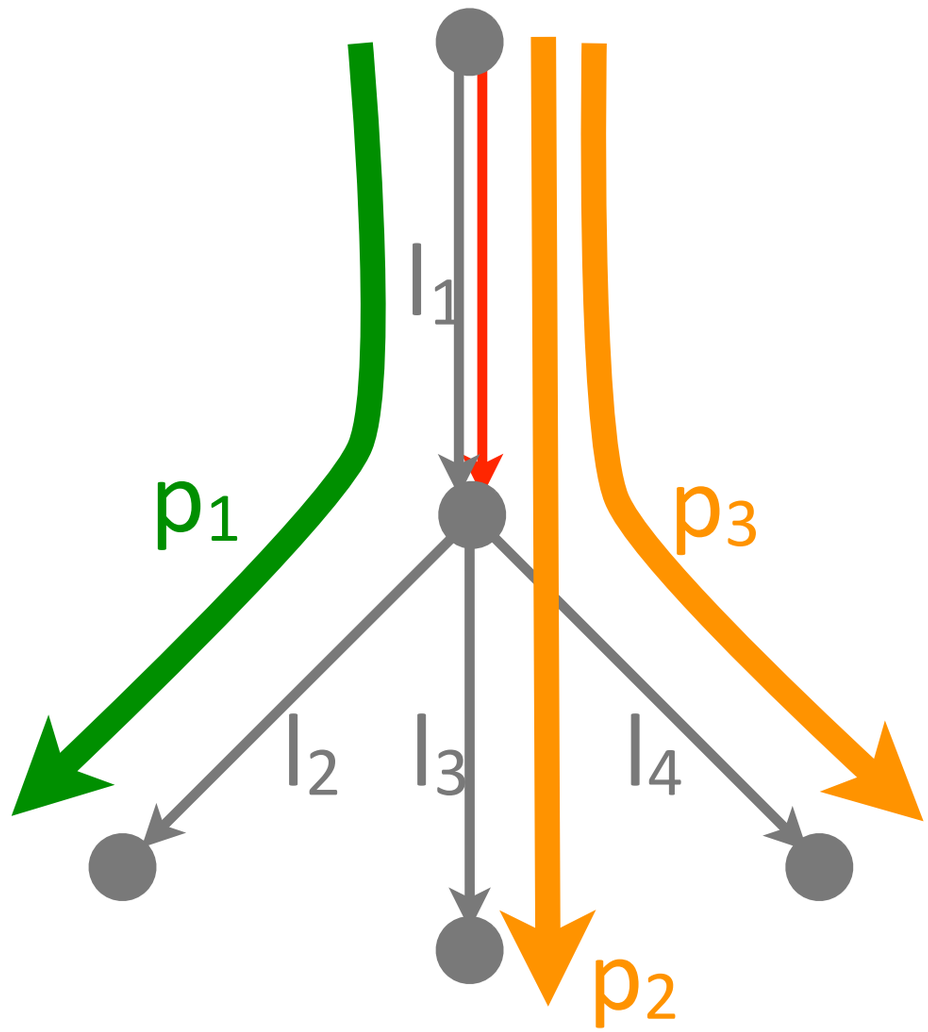

$$Y_2 = X_1 + X_3$$

$$Y_3 = X_1 + X_4$$

$$Y_{23} = X_1 + X_3 + X_4$$

Localizable violation

$$0 = \cancel{X_1}^{0} + \cancel{X_2}^{0}$$

$$0 > Y_2 = \cancel{X_1}^{0} + X_3$$

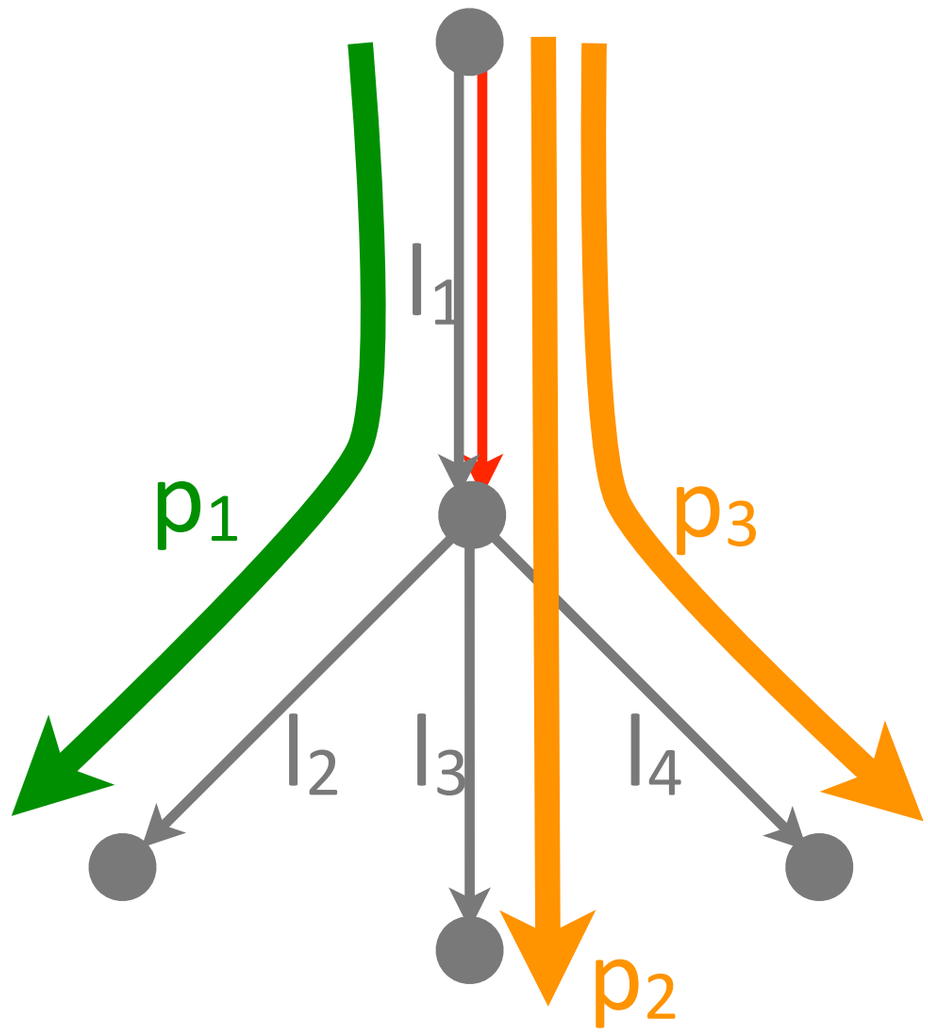

$$Y_2 = X_1 + X_3$$

$$Y_3 = X_1 + X_4$$

$$Y_{23} = X_1 + X_3 + X_4$$

Localizability

- We can localize a neutrality violation to a particular link sequence,

  if there exist at least two path pairs that intersect exactly at this link sequence (…)

Outline

- Observability of neutrality violation
- Localization of neutrality violations
- Algorithm + results

Algorithm

- Considers each link sequence
- Identifies path pairs that intersect exactly at this link sequence
- Checks whether these form an unsolvable tomographic system of equations

Measurements

- Divide time into intervals
- Measure "congestion probabilities"
  - *count the intervals in which a path or a pair of paths is congestion-free*

Output

- Non-neutral link sequences
  - *include at least one non-neutral link*
  - *satisfy our localization condition*

Metrics

- False positives
  - *fraction of neutral links that are in a link sequence that is misclassified as non-neutral*

- False negatives
  - *fraction of non-neutral links that are not in a link sequence that is classified as non-neutral*

- Granularity
  - *average size of link sequences classified as non-neutral*

Evaluation platform

- Network emulator (LINE)
  - *virtual interface per end-host*
  - *runs in user space*
- Real network stacks, TCP traffic
- Real packet queues, policers, shapers
  - *(in software)*
- Simulated transfer and propagation delay

Simplest setup: uniform traffic

- Two traffic classes: high and low priority
- All traffic has the same TCP flow properties
- All links implement drop-tail policy
- Non-neutral links police low-priority at 30%

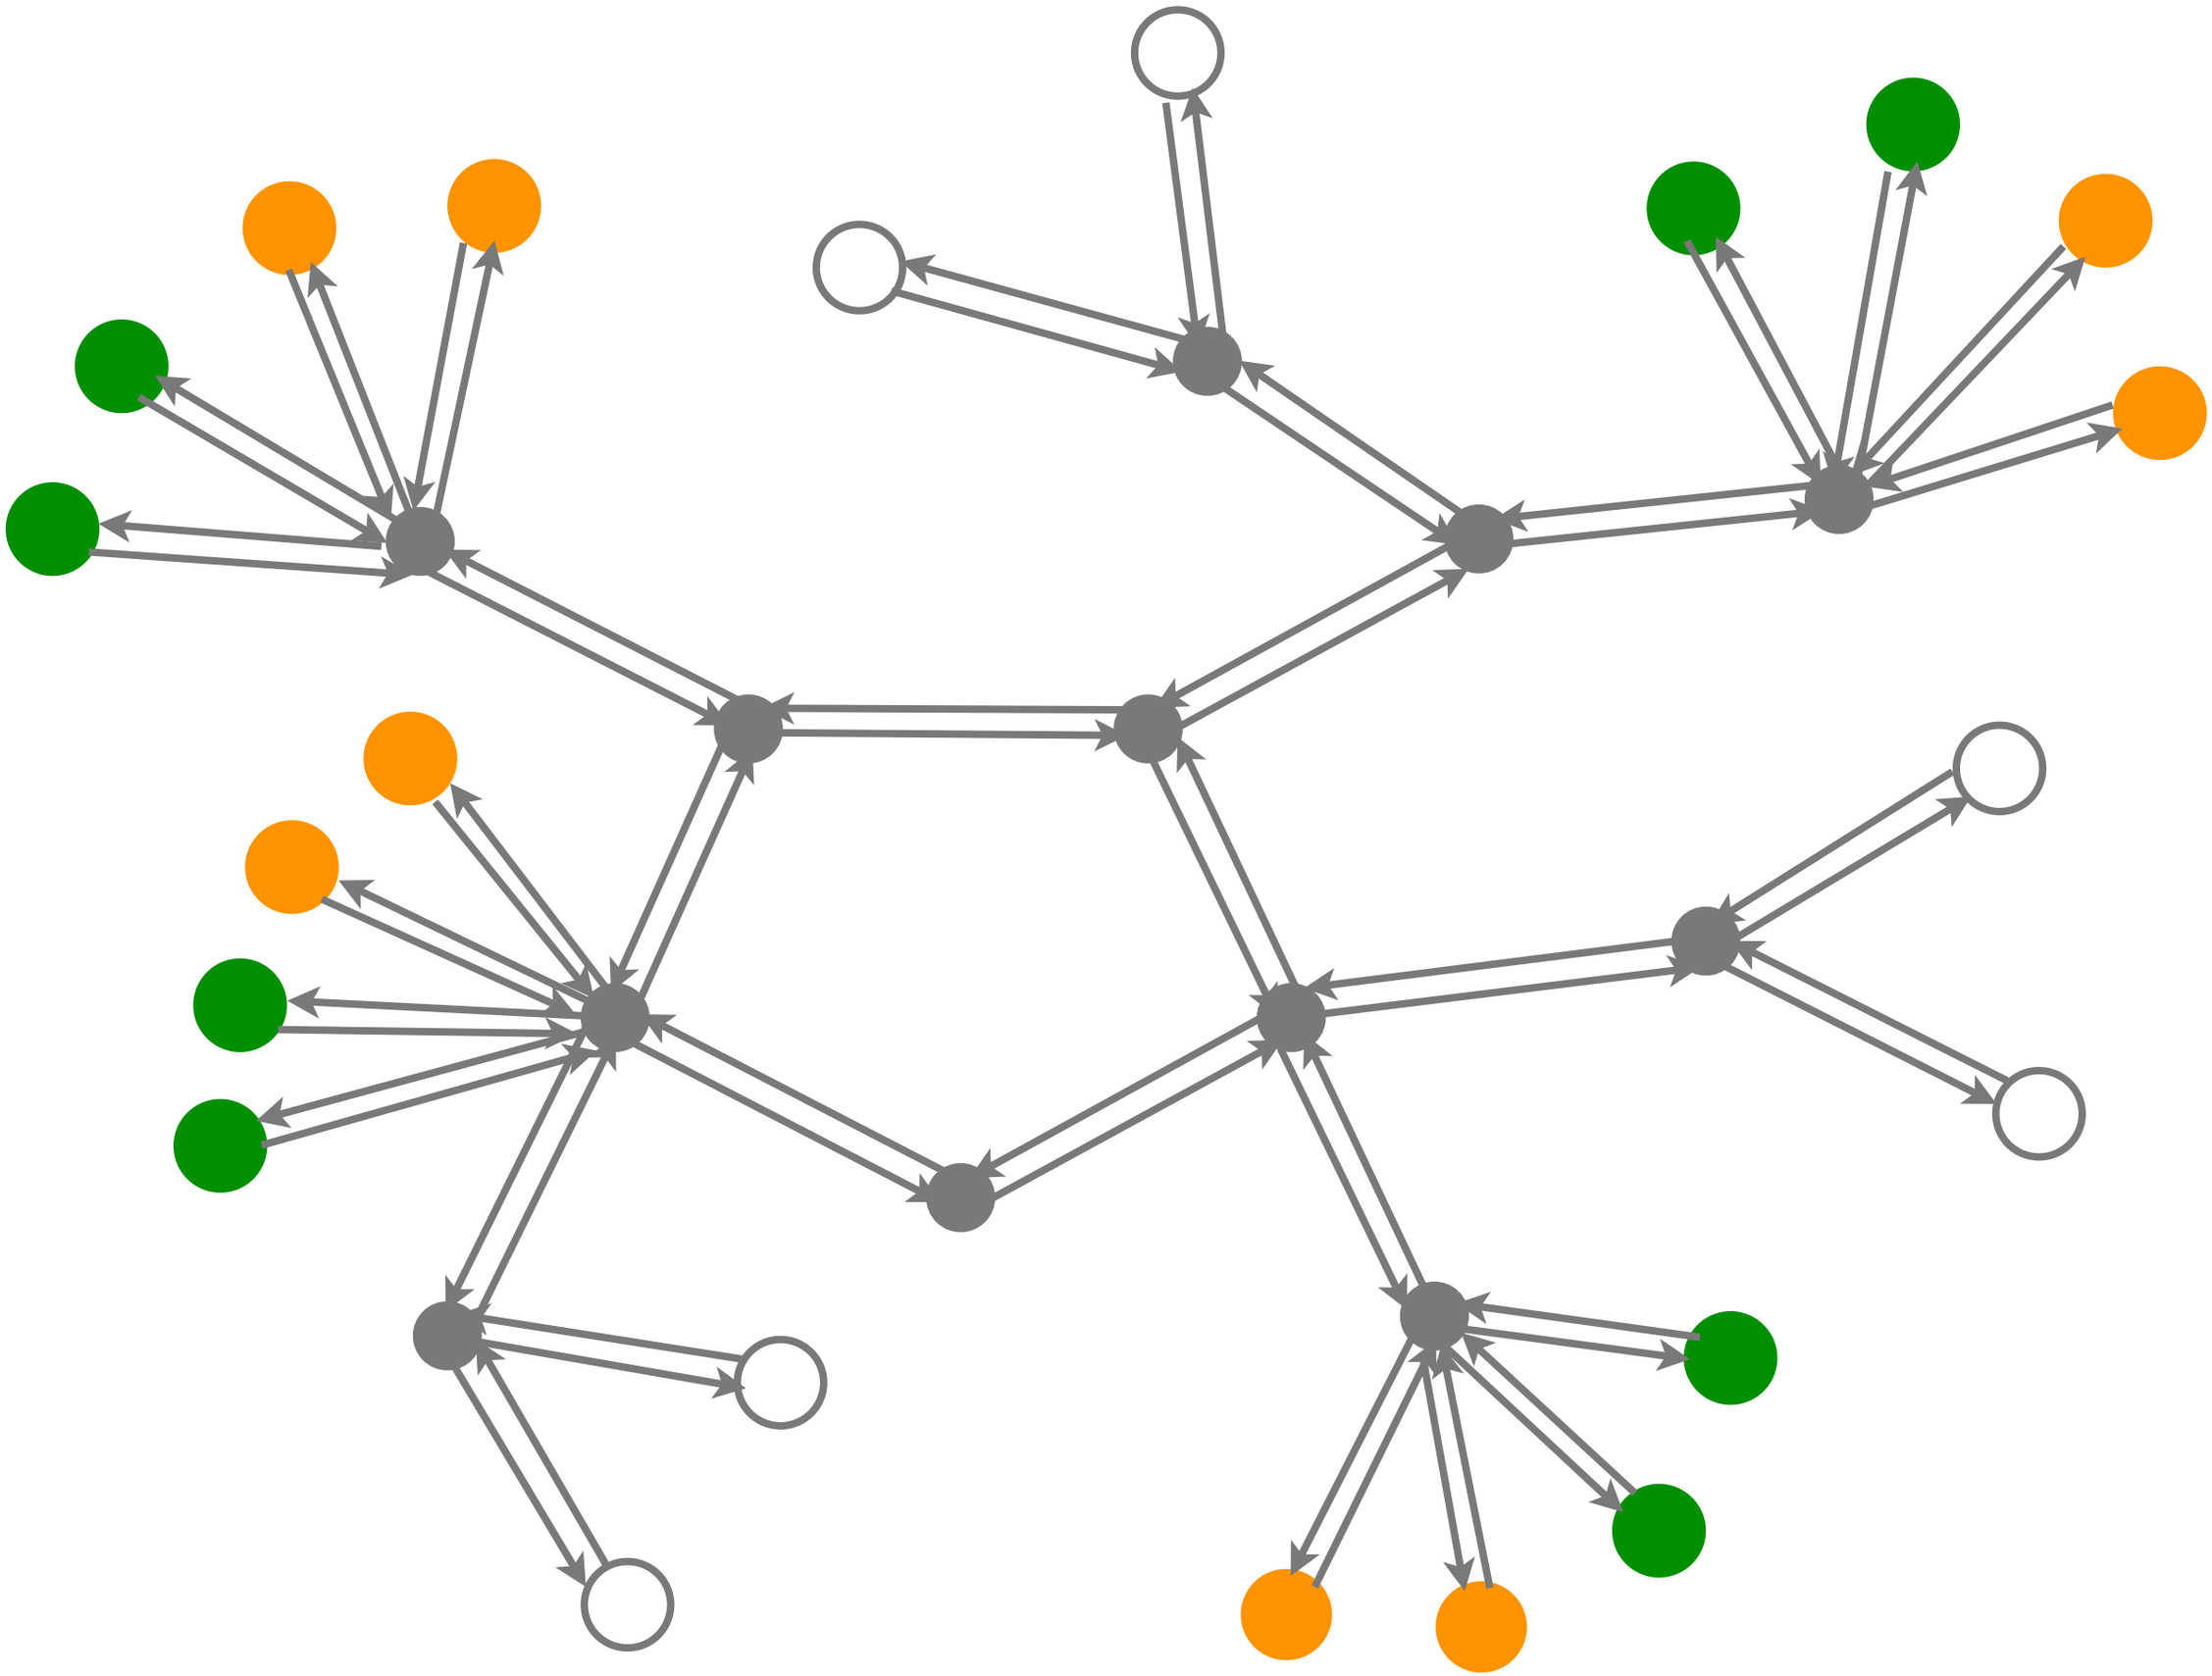

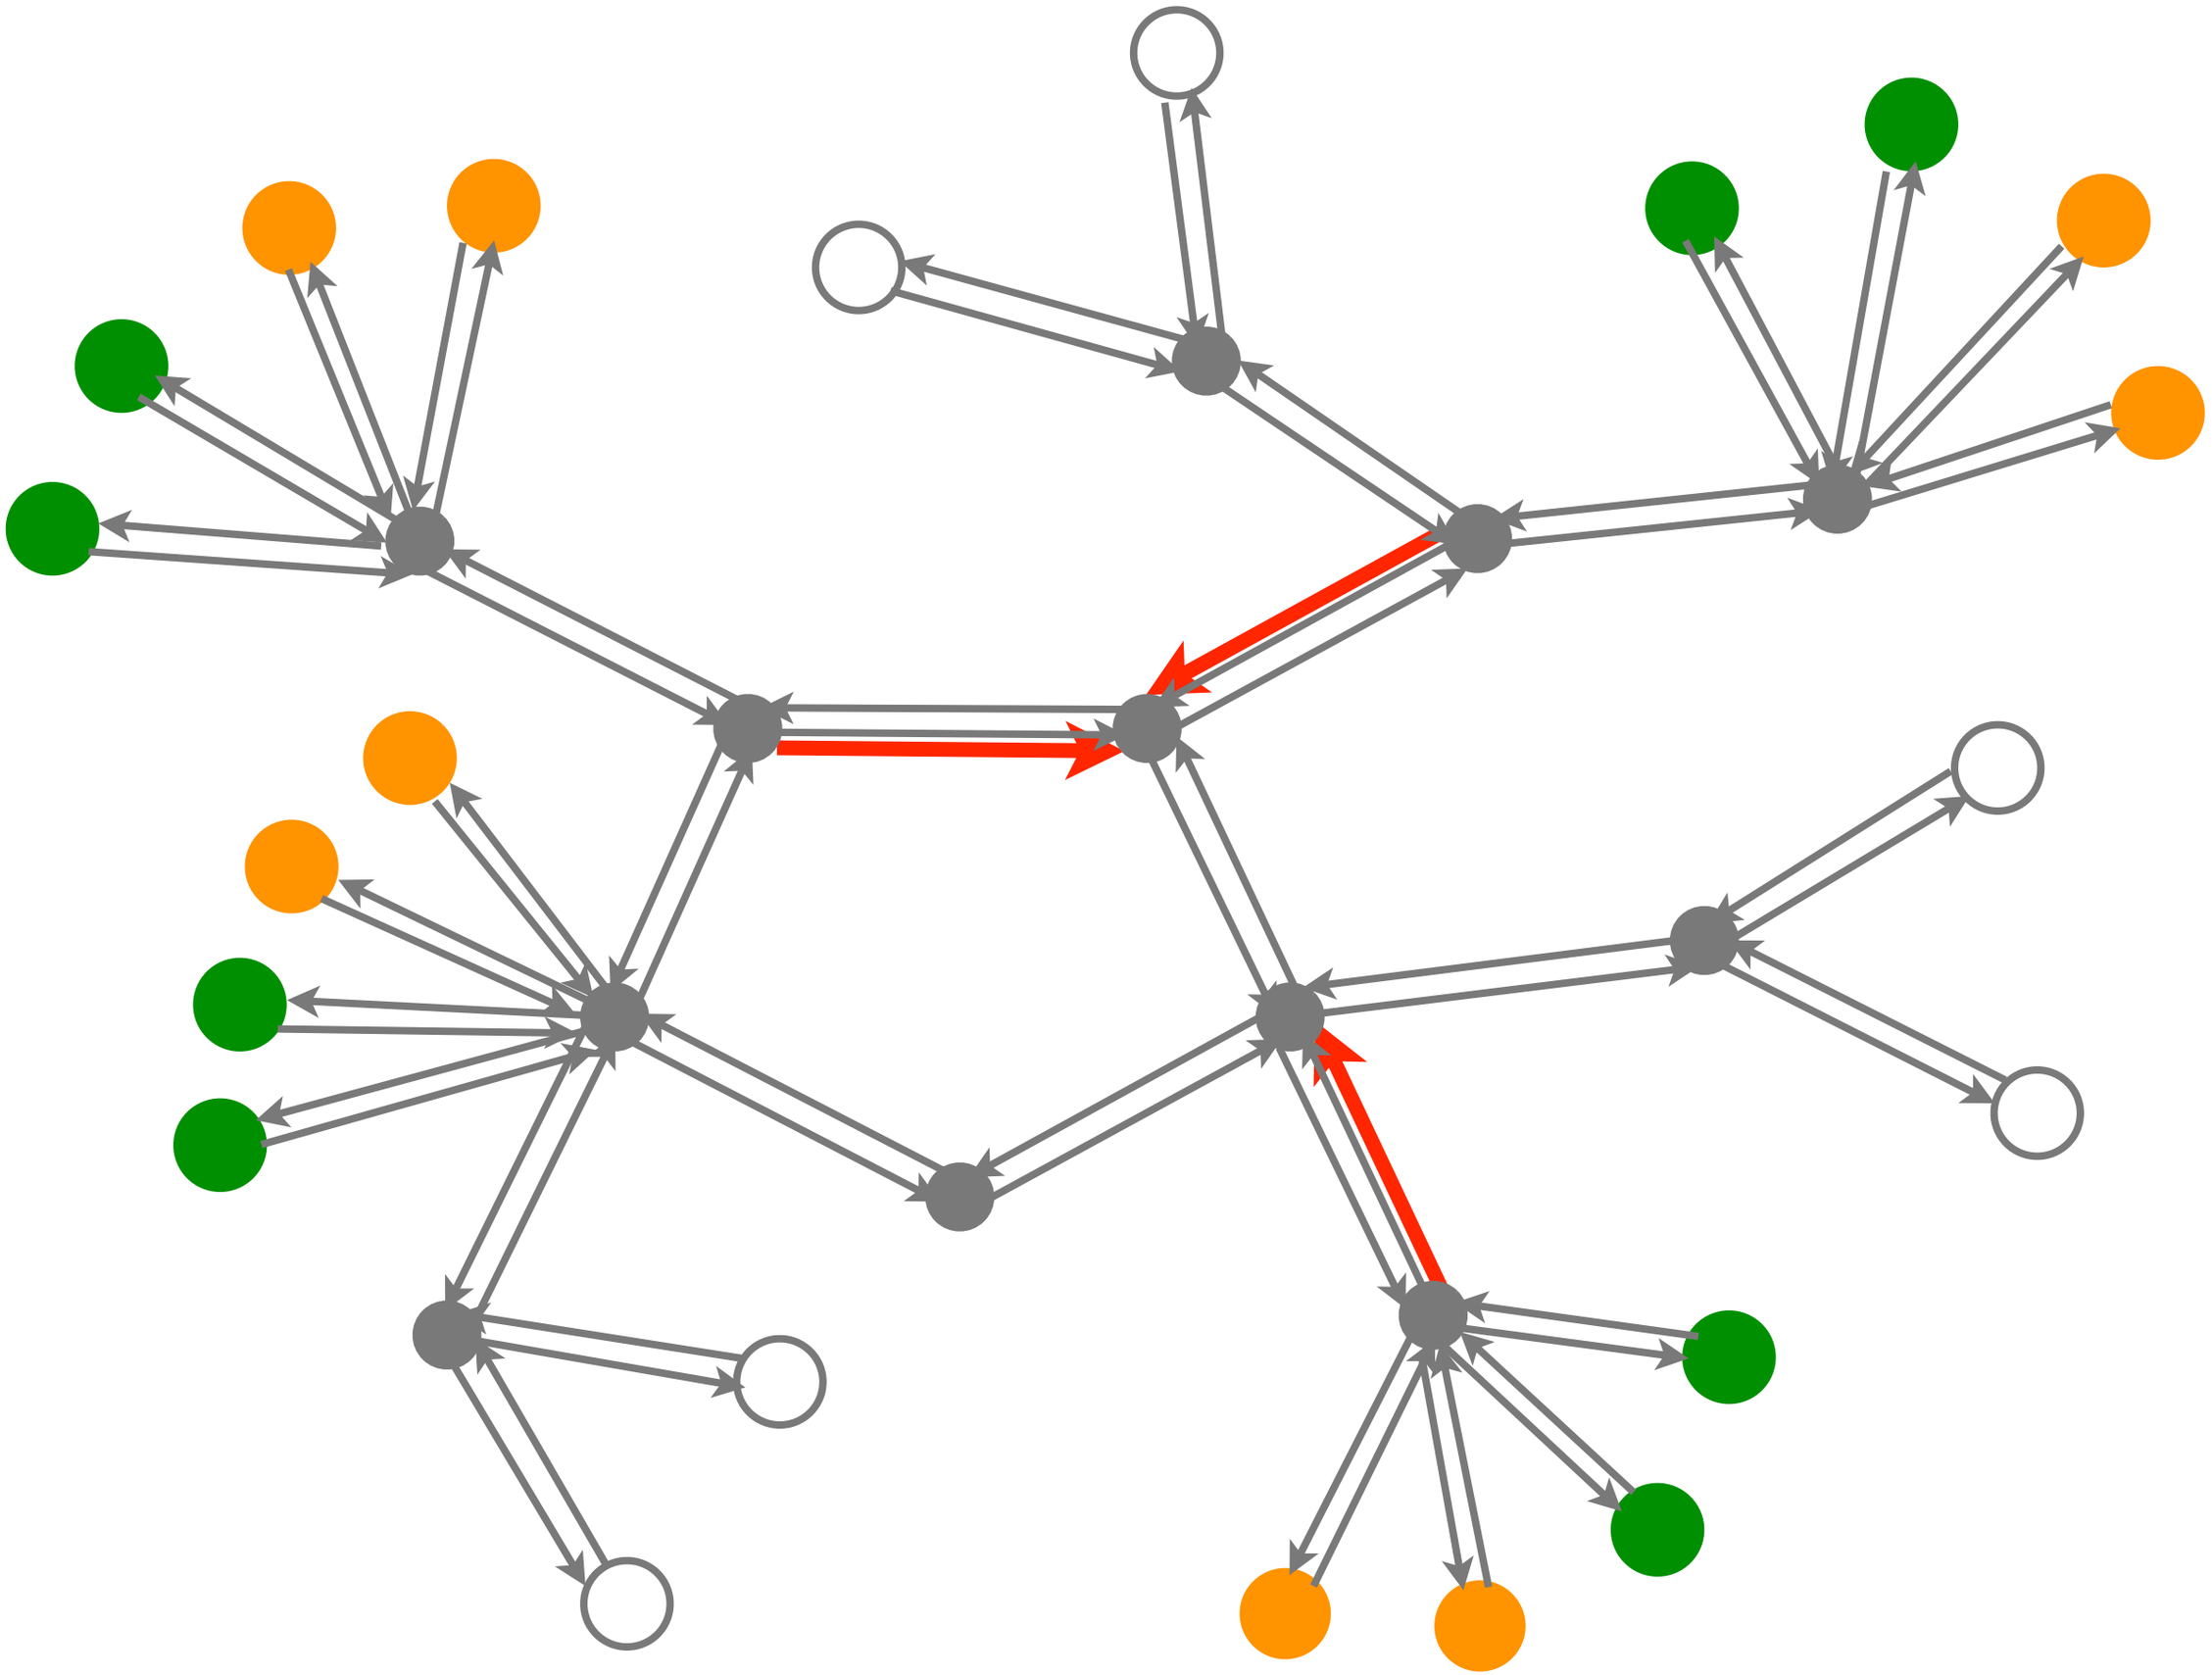

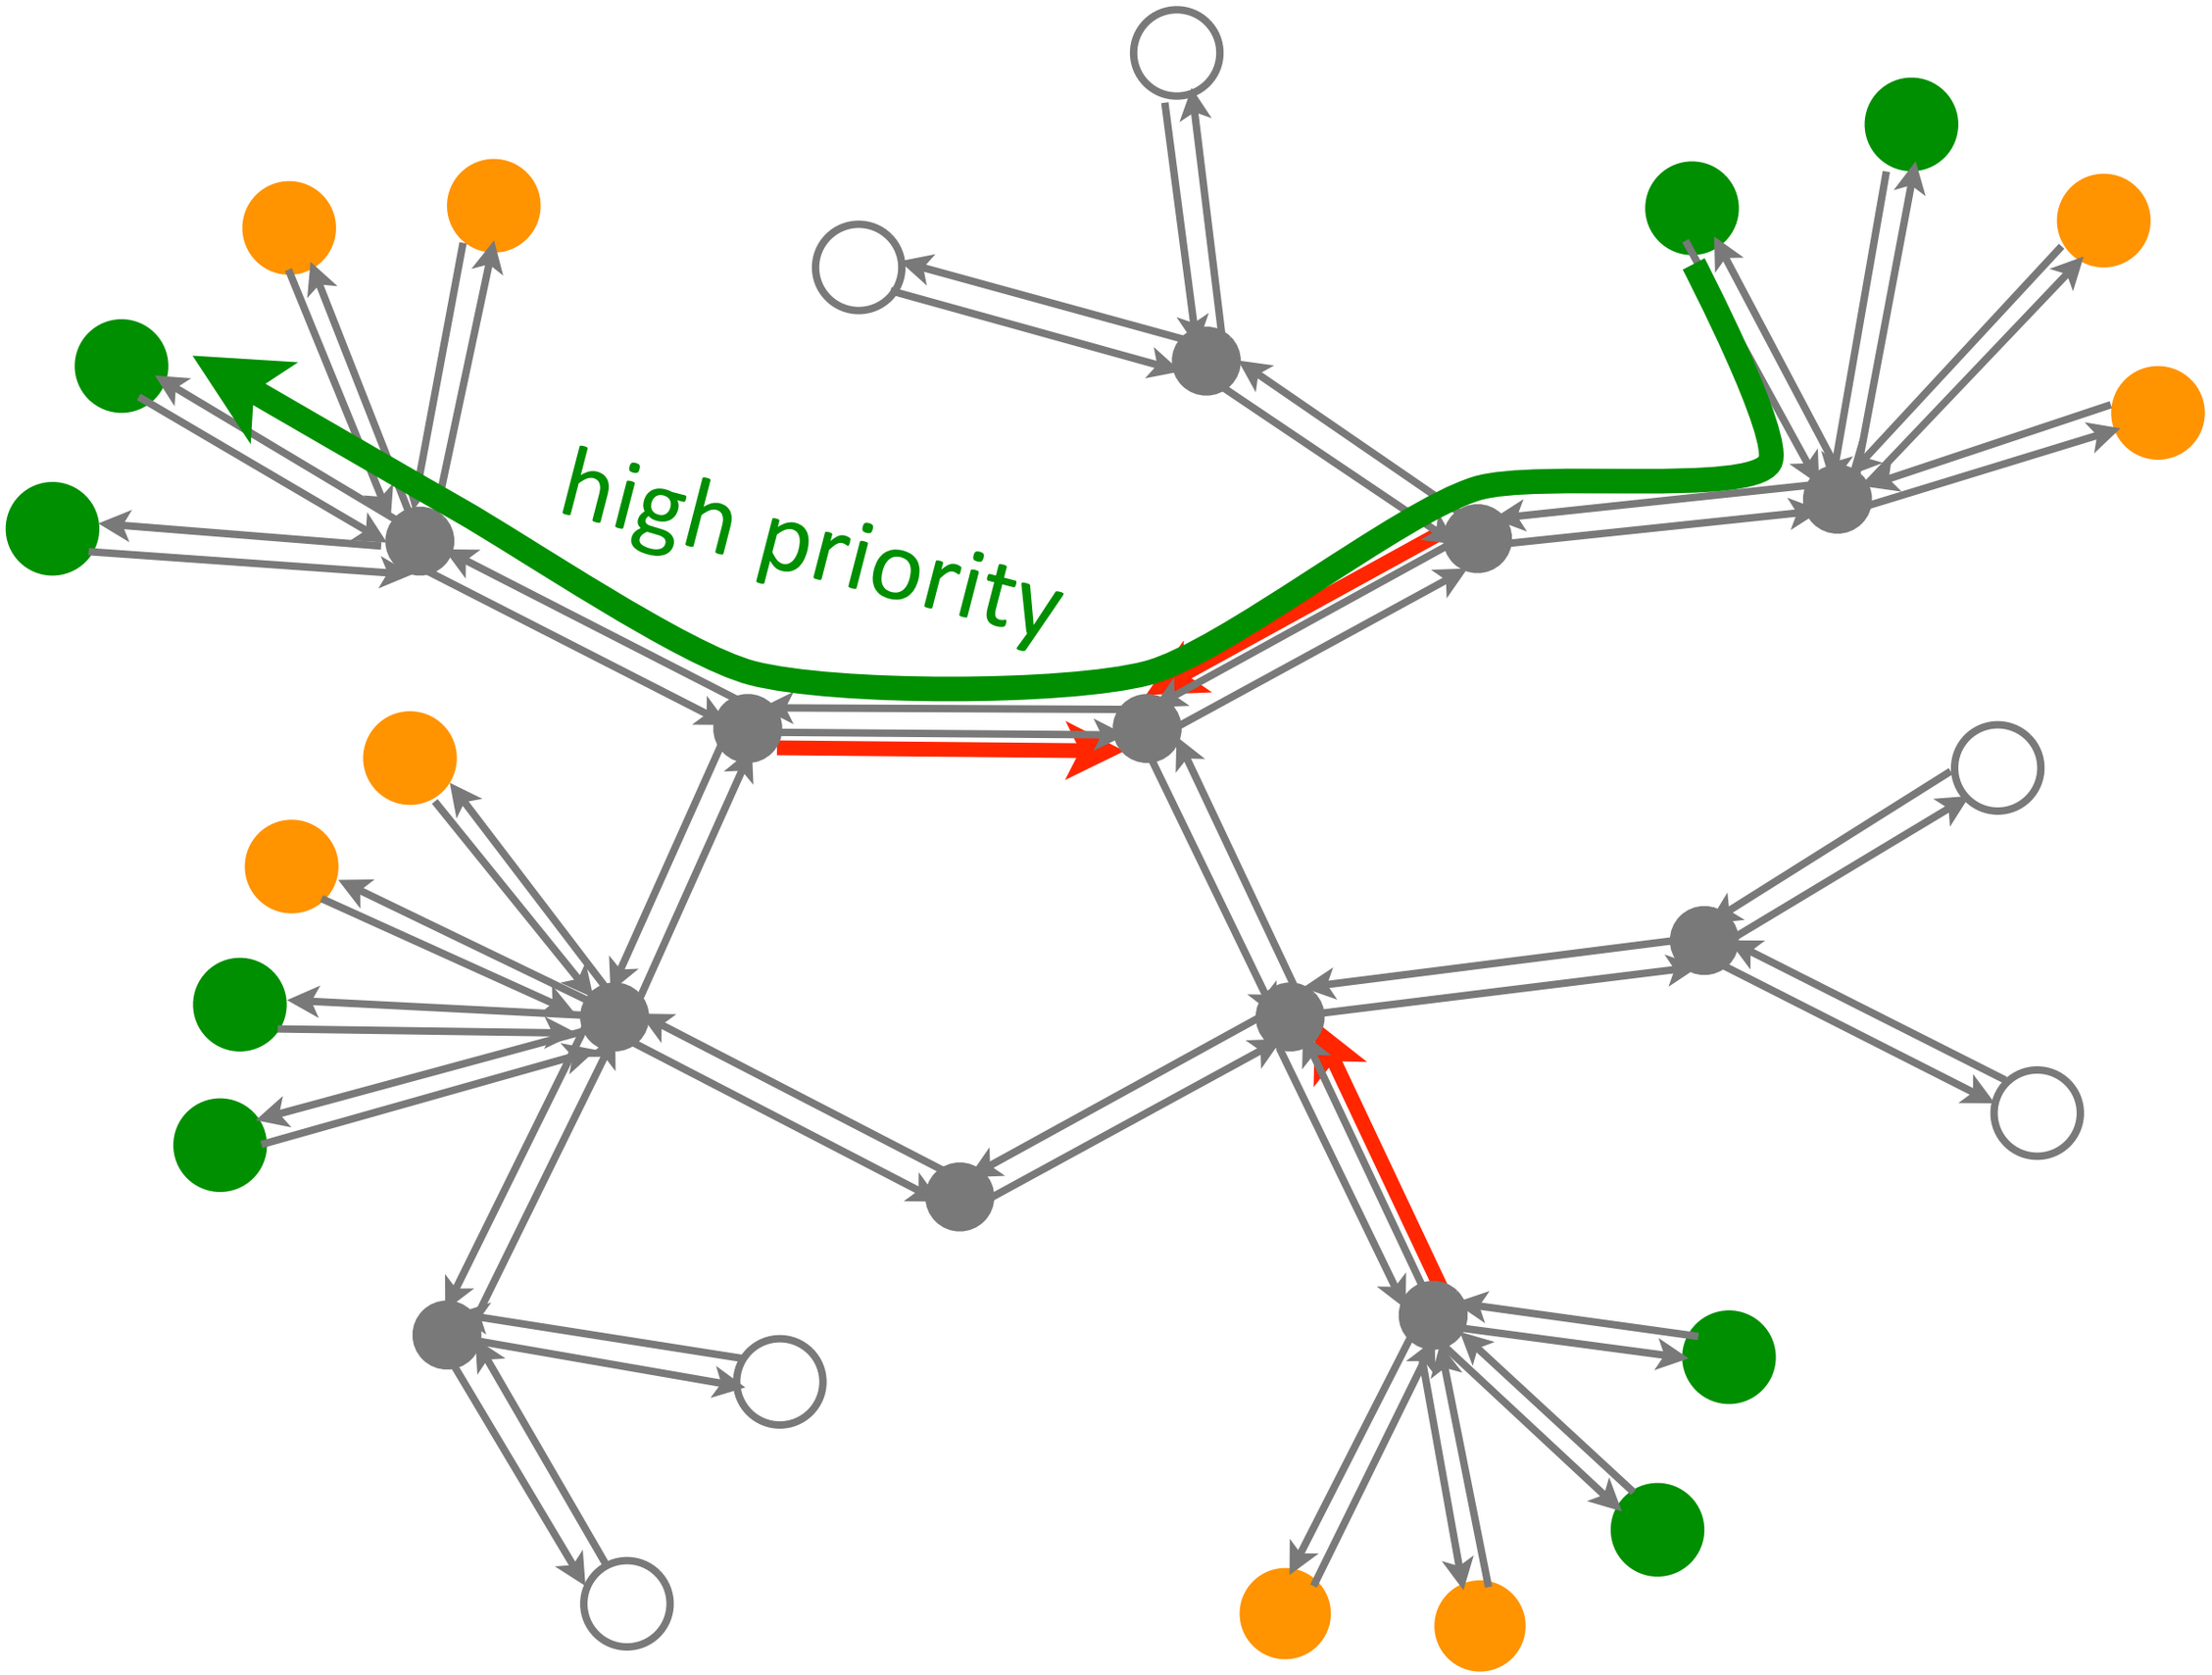
high priority

high priority

low priority

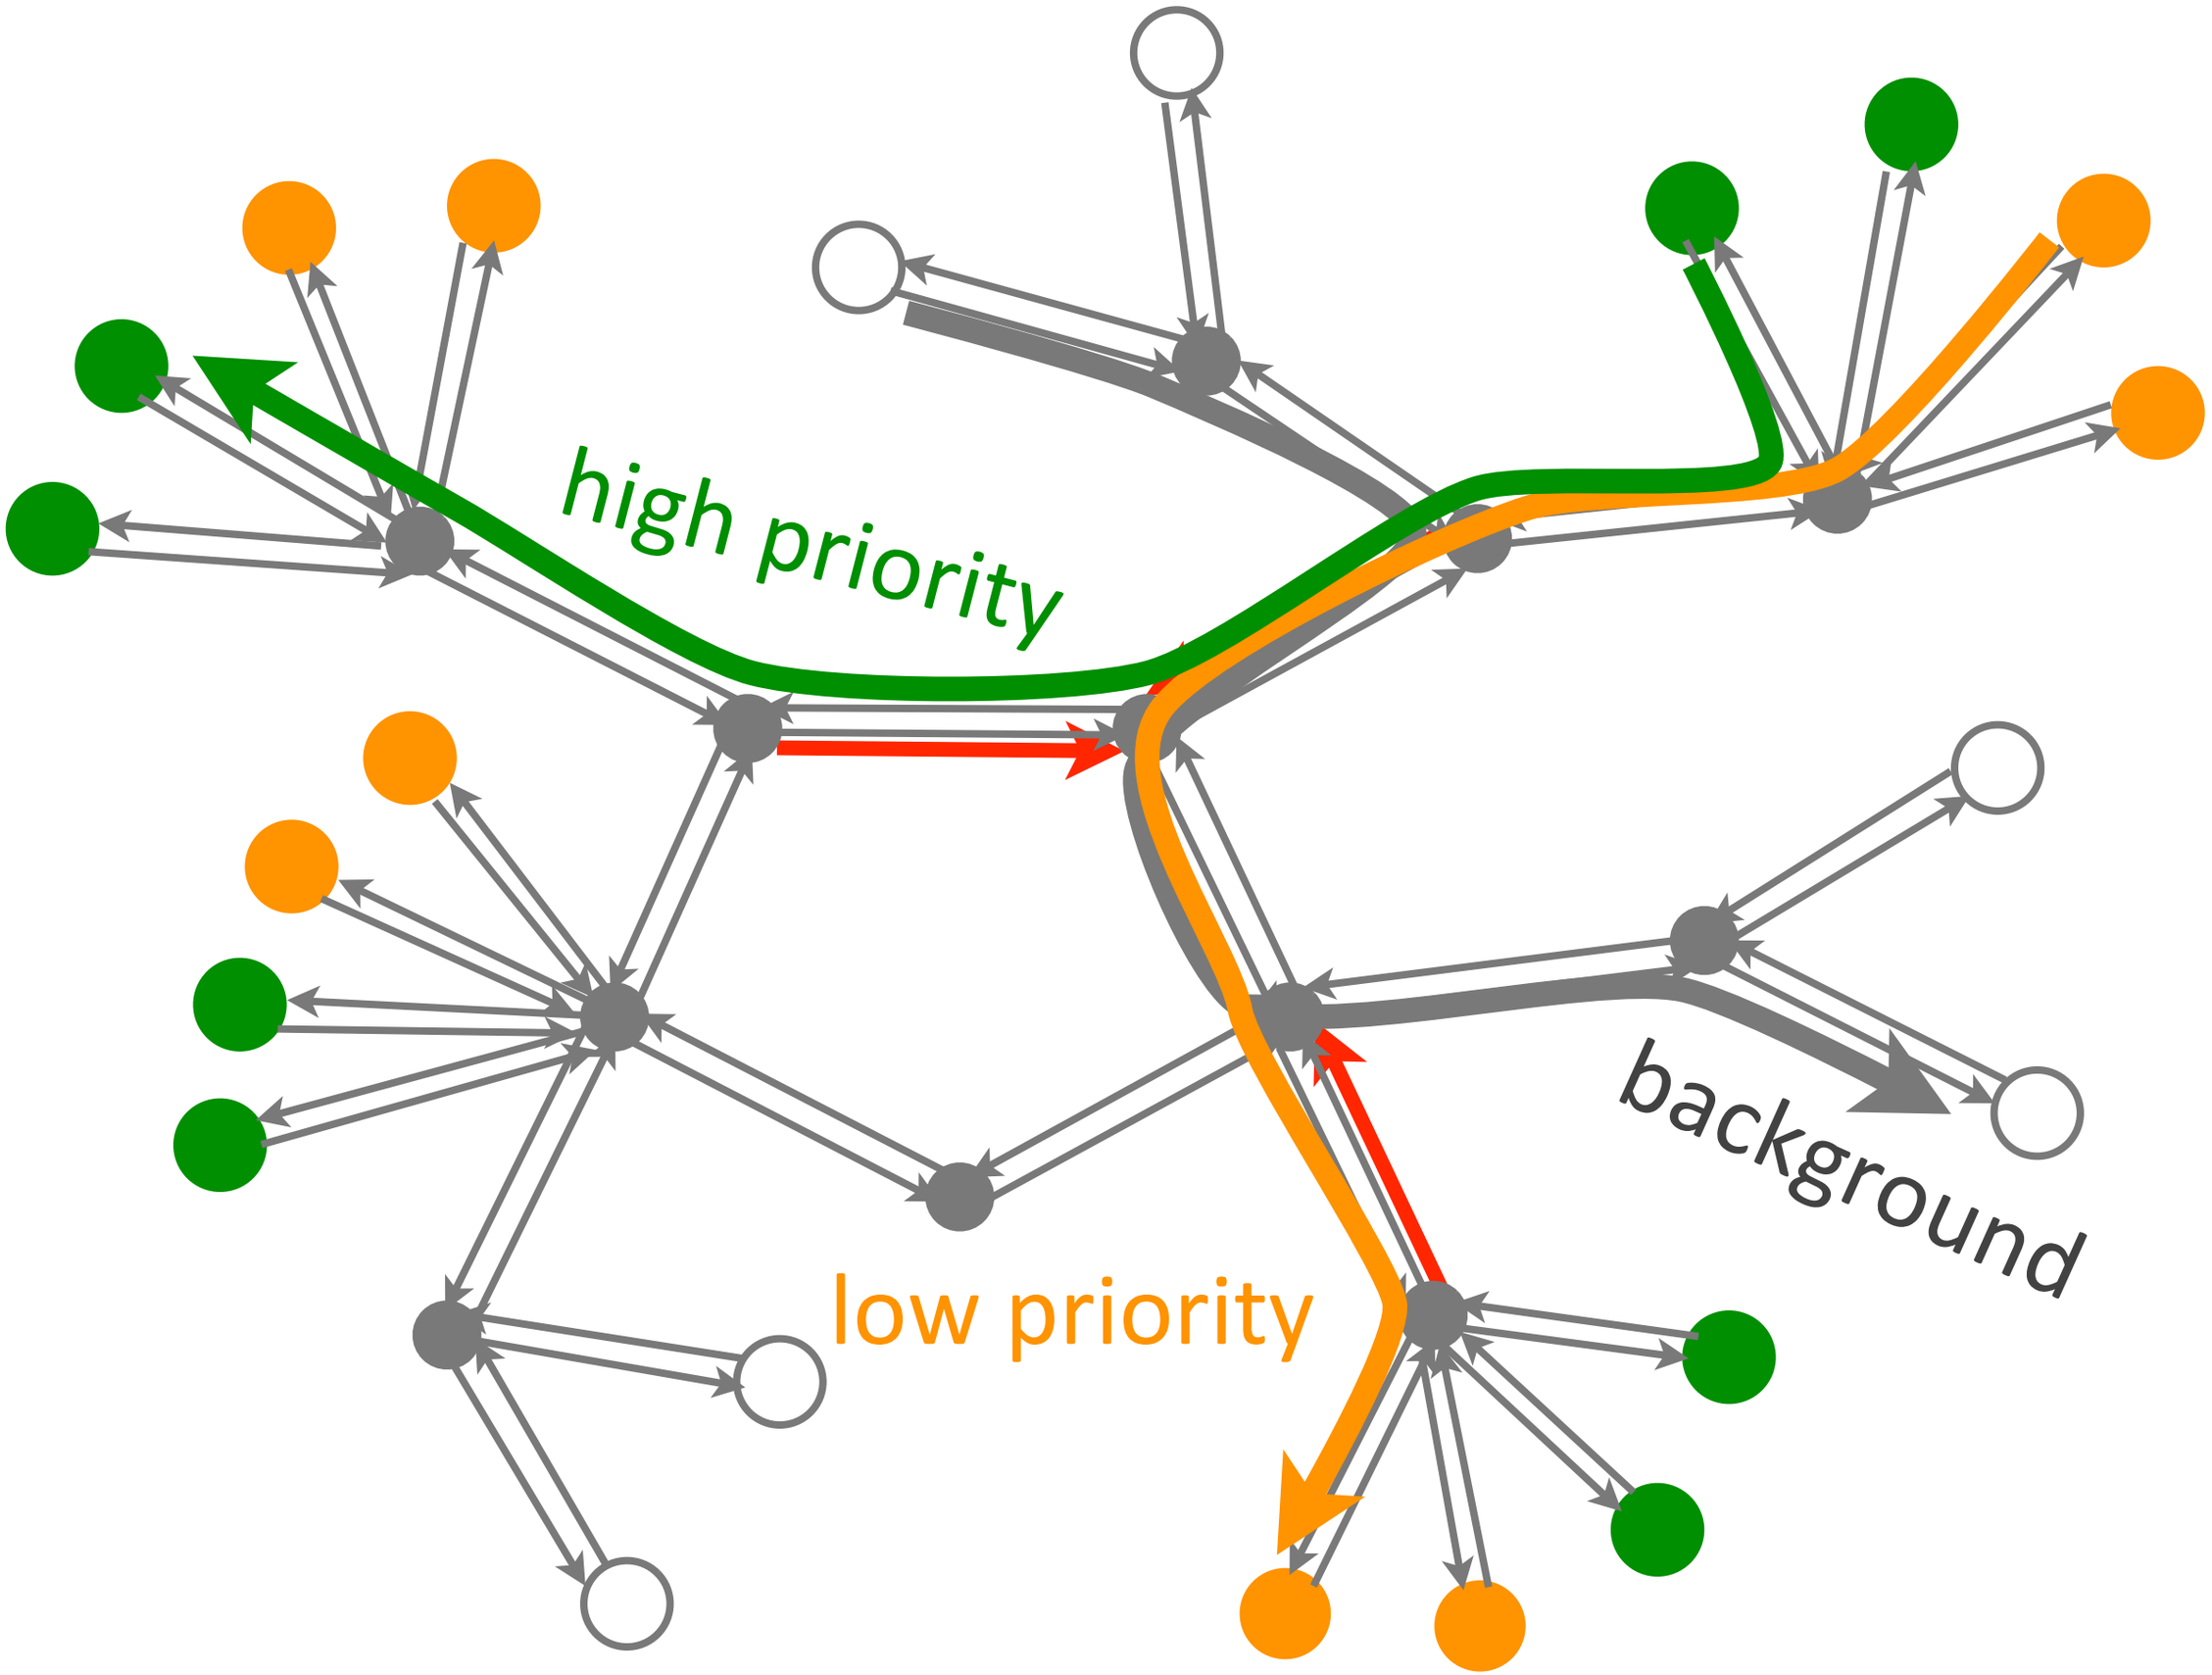
high priority
low priority
background

Results

- No false positives
- No false negatives
- Granularity: 1.8

Under the hood

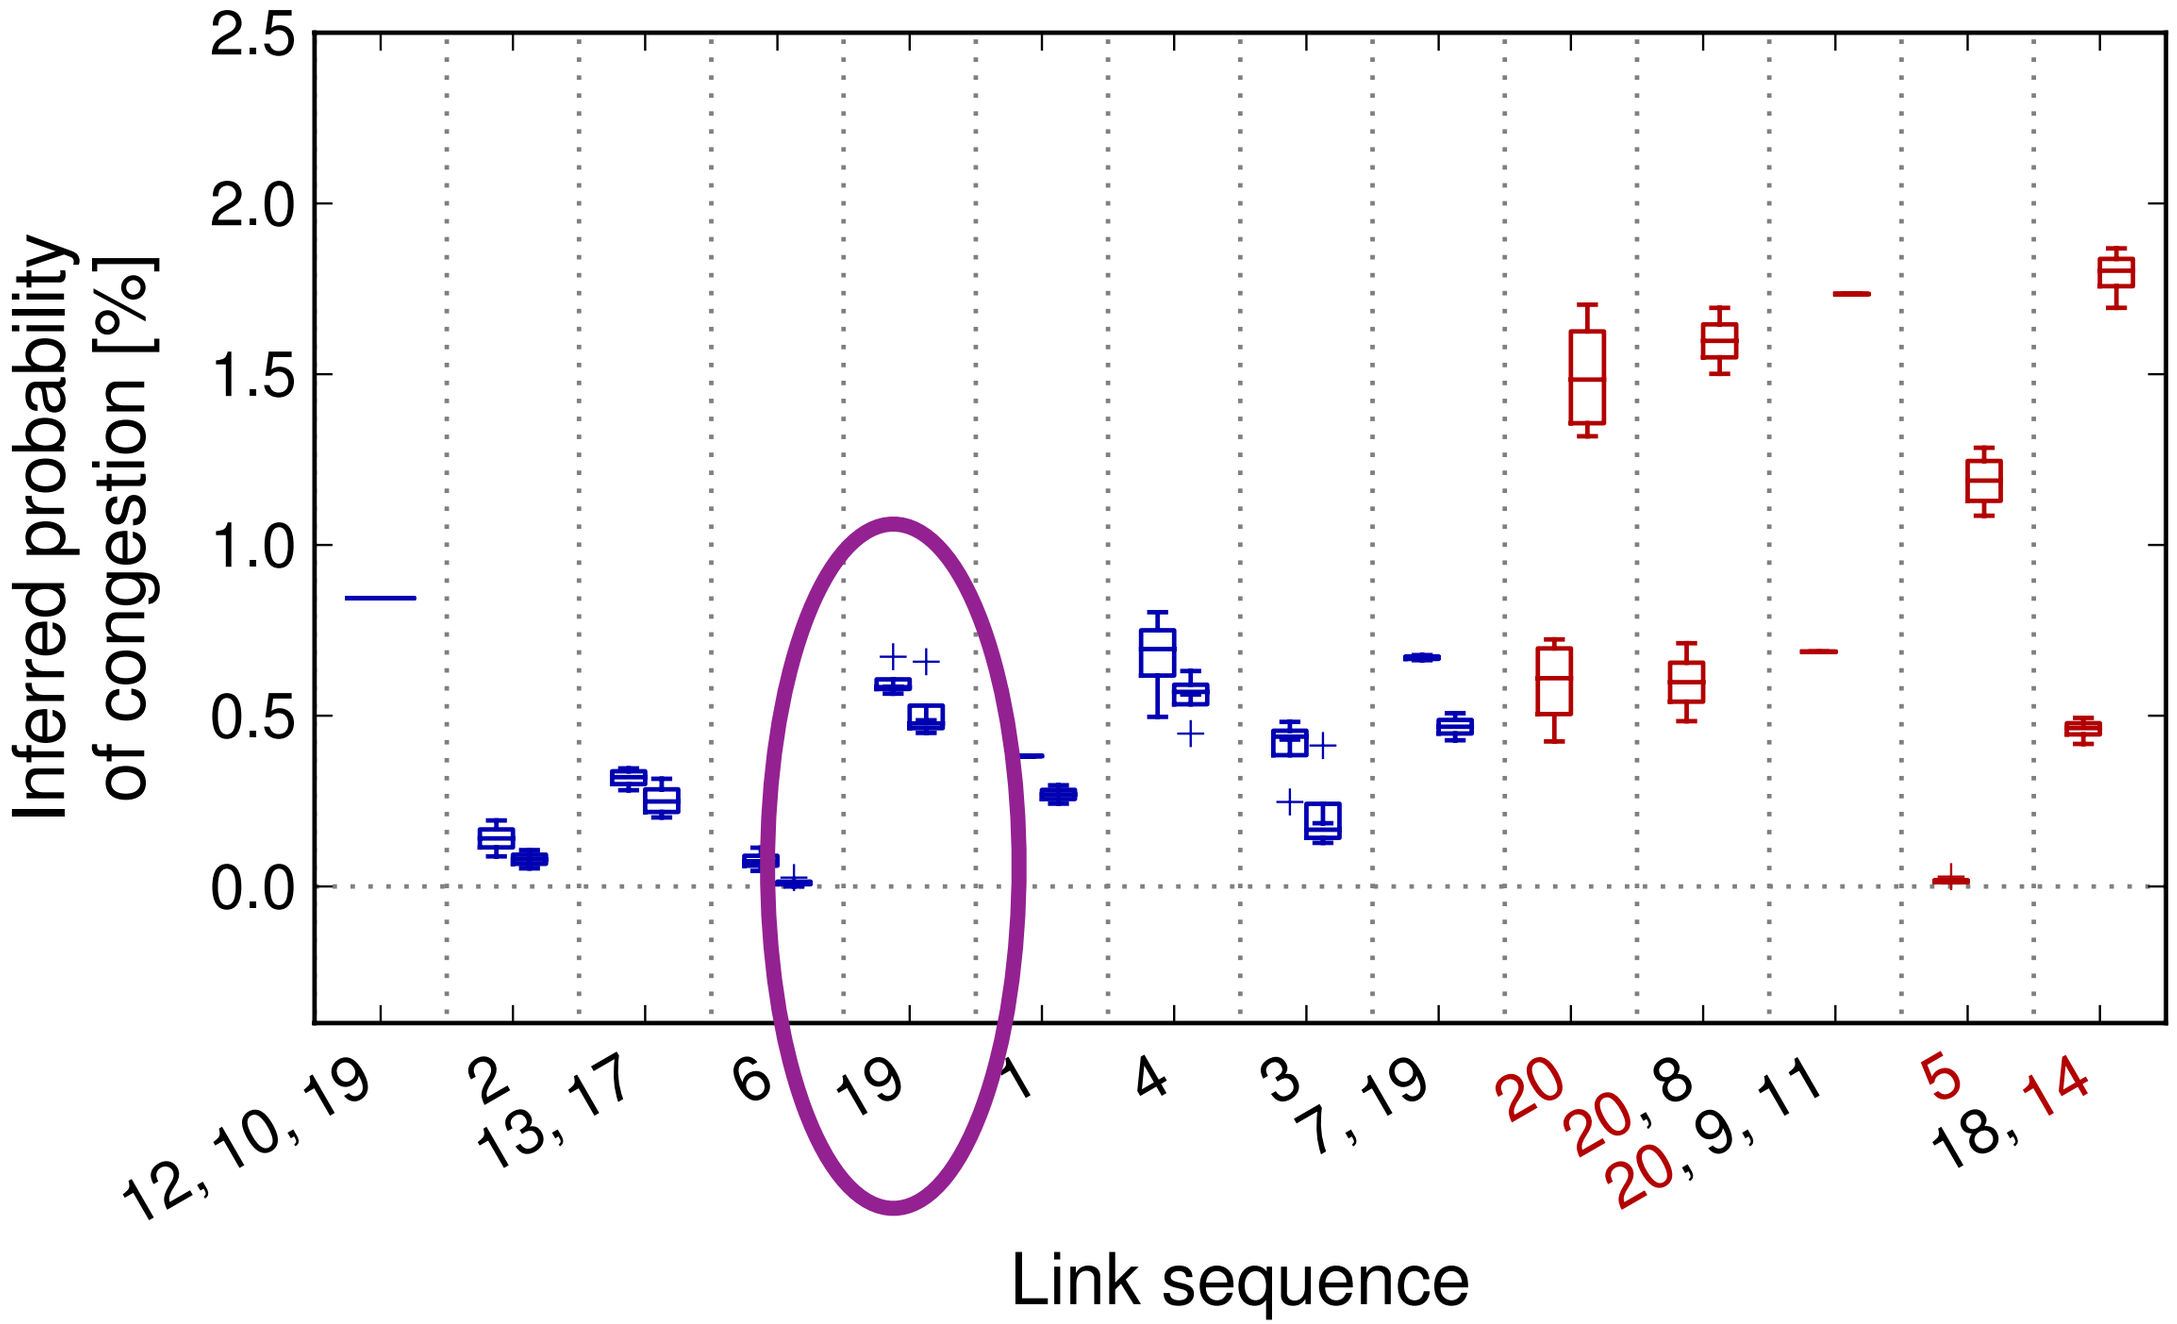

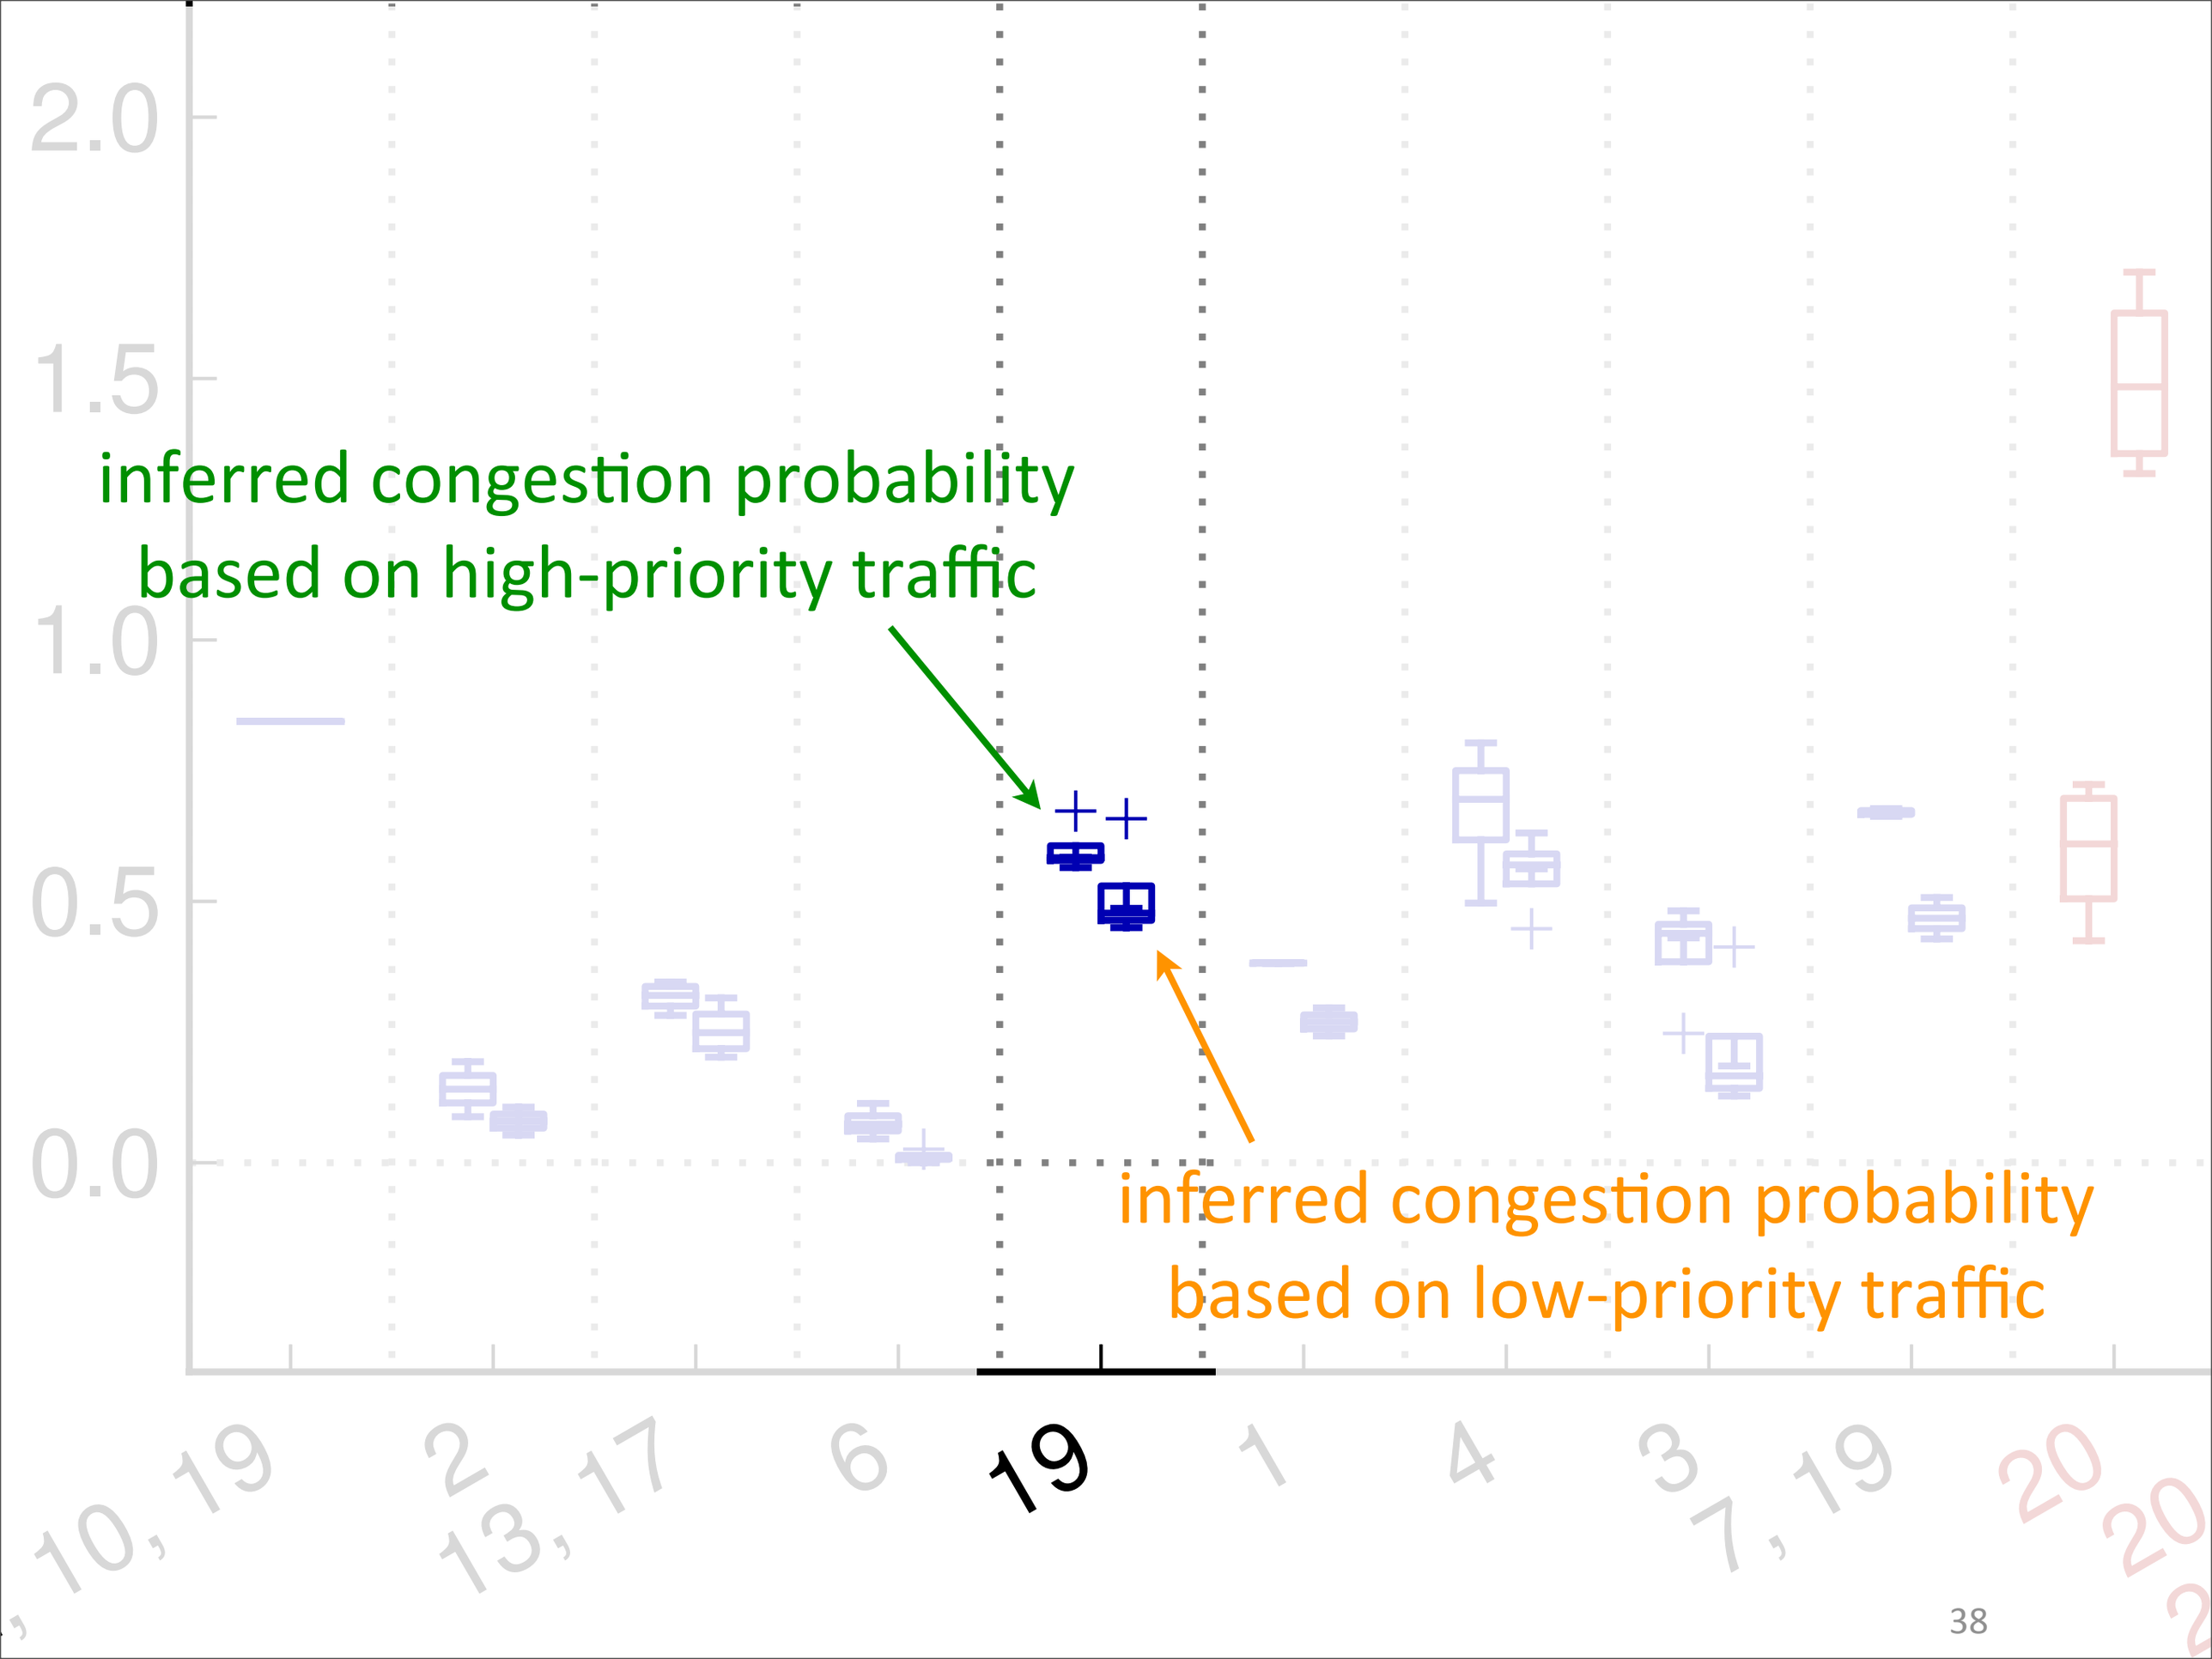
2.0
1.5
1.0
0.5
0.0
inferred congestion probability
based on high-priority traffic
inferred congestion probability
based on low-priority traffic
19

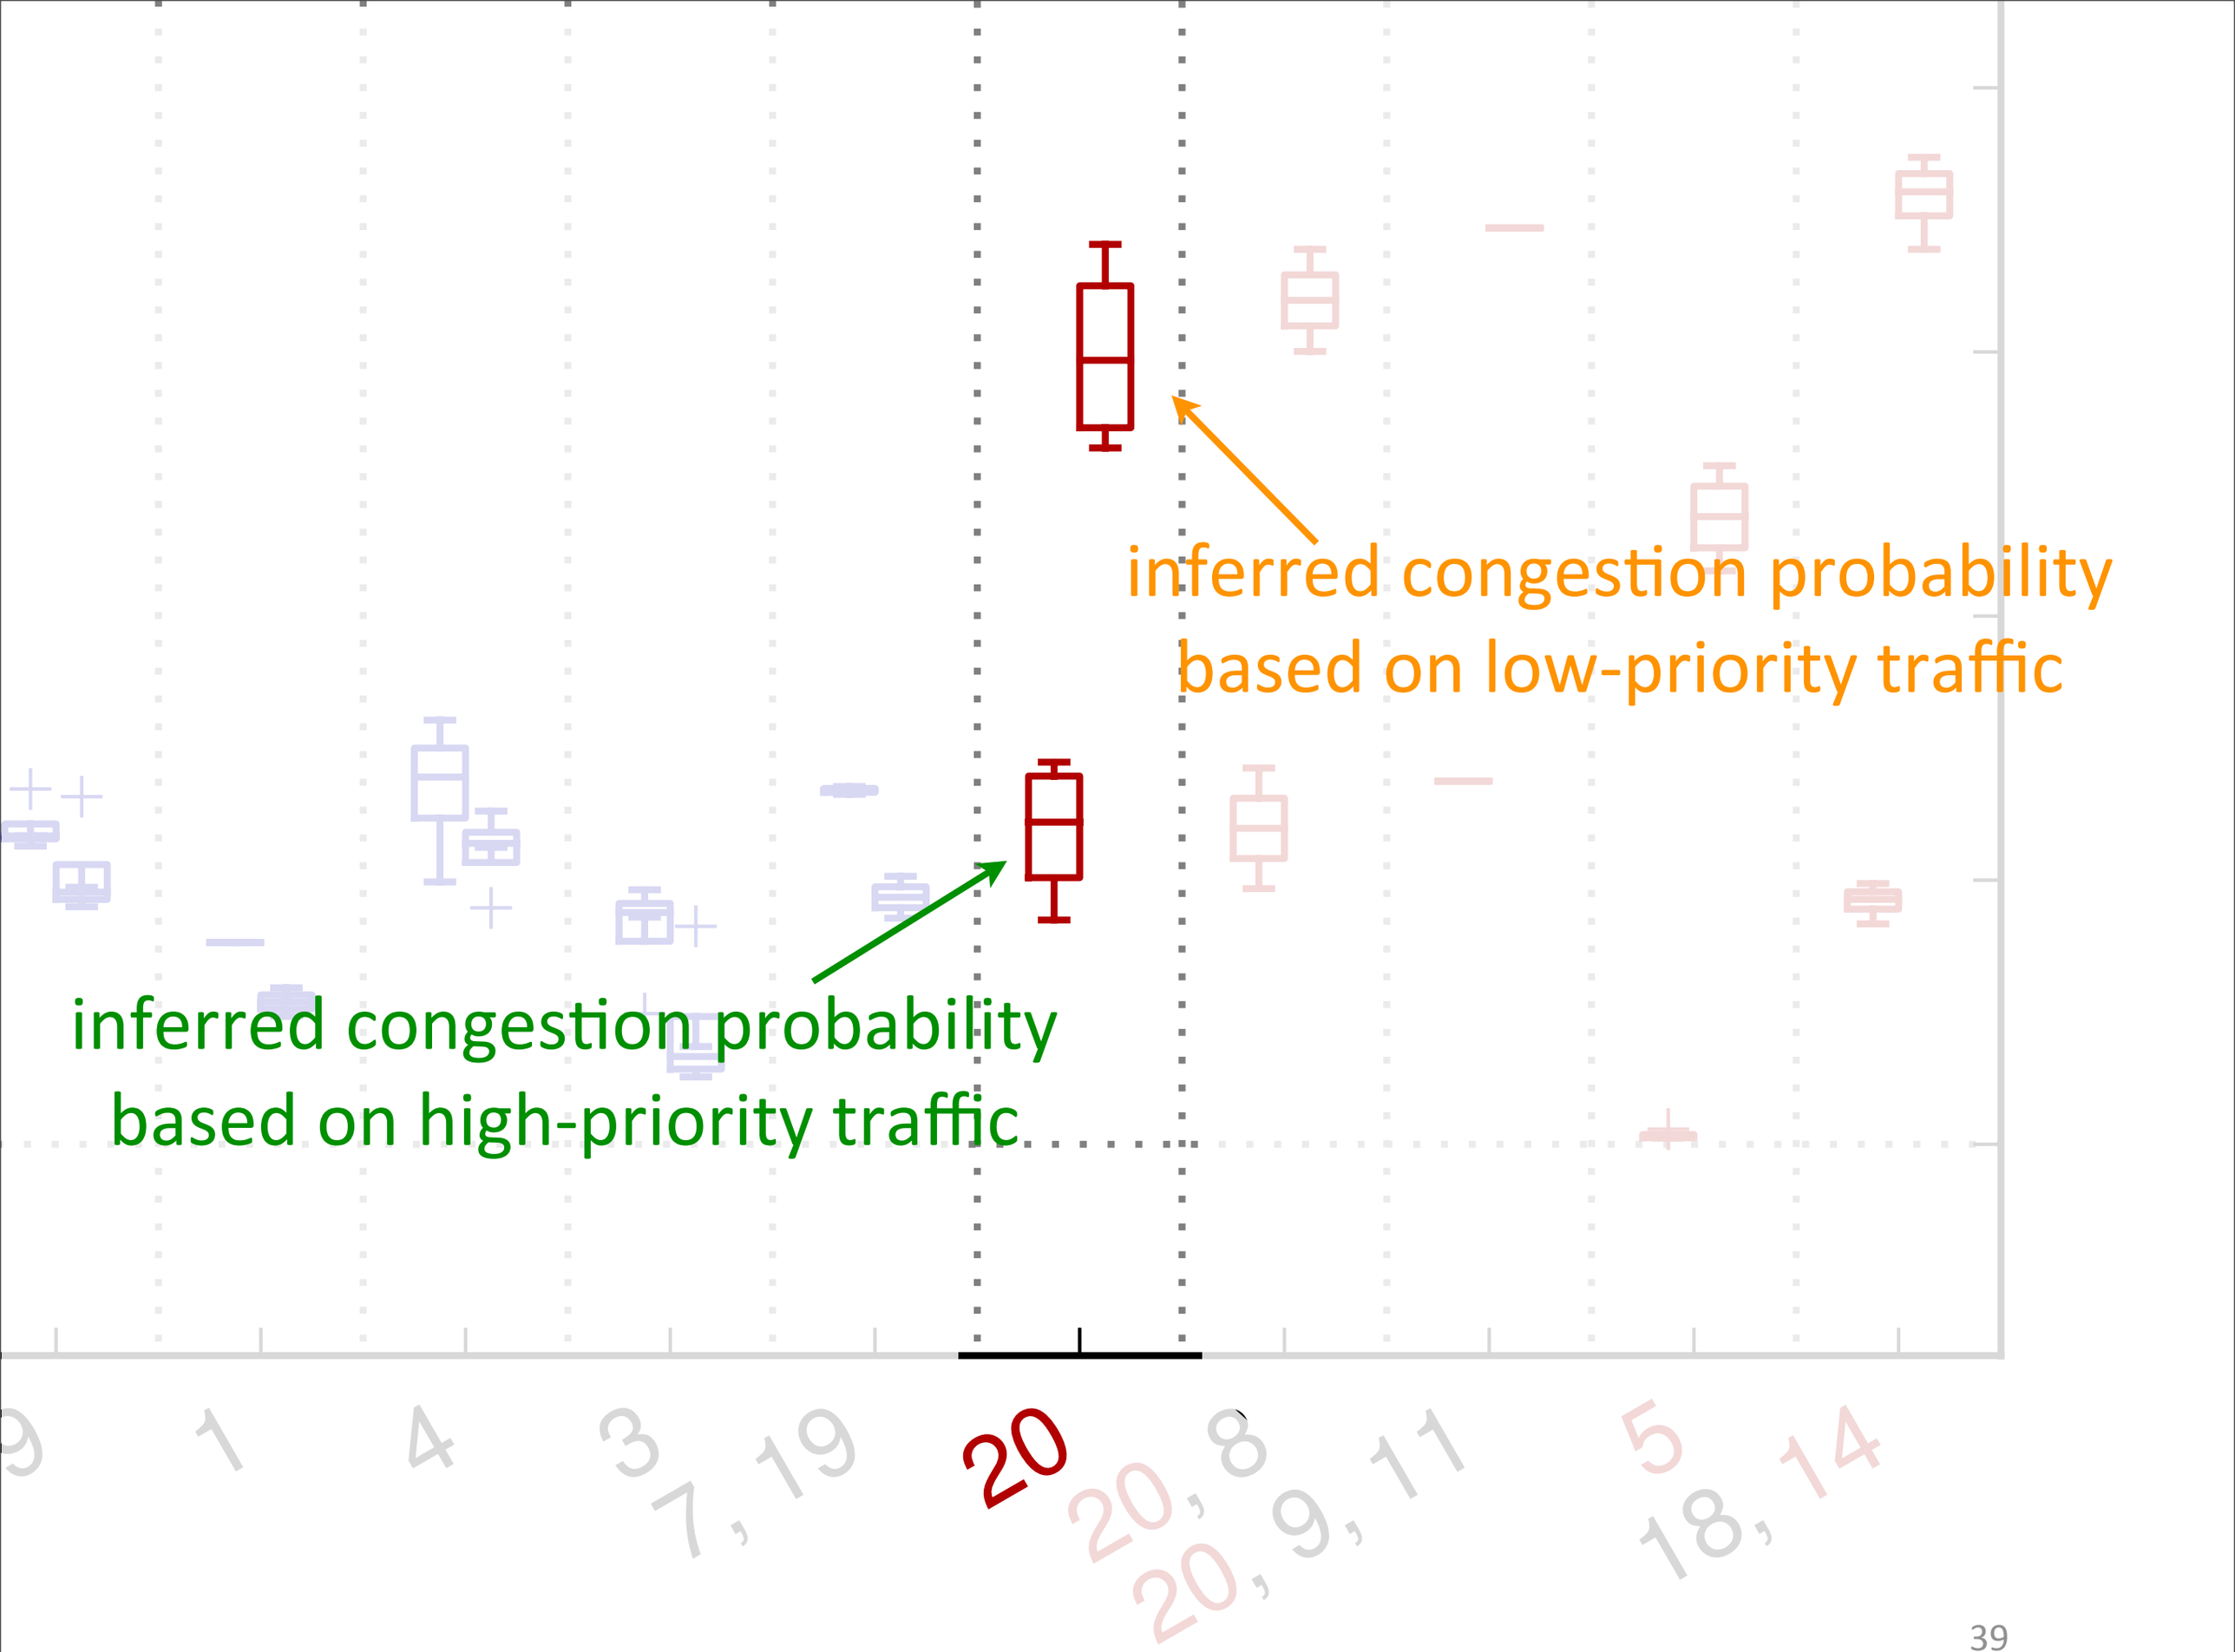
inferred congestion probability
based on low-priority traffic
inferred congestion probability
based on high-priority traffic
9
1
4
3
7, 19
20
20, 8
20, 9, 11
5
18, 14

Under the hood

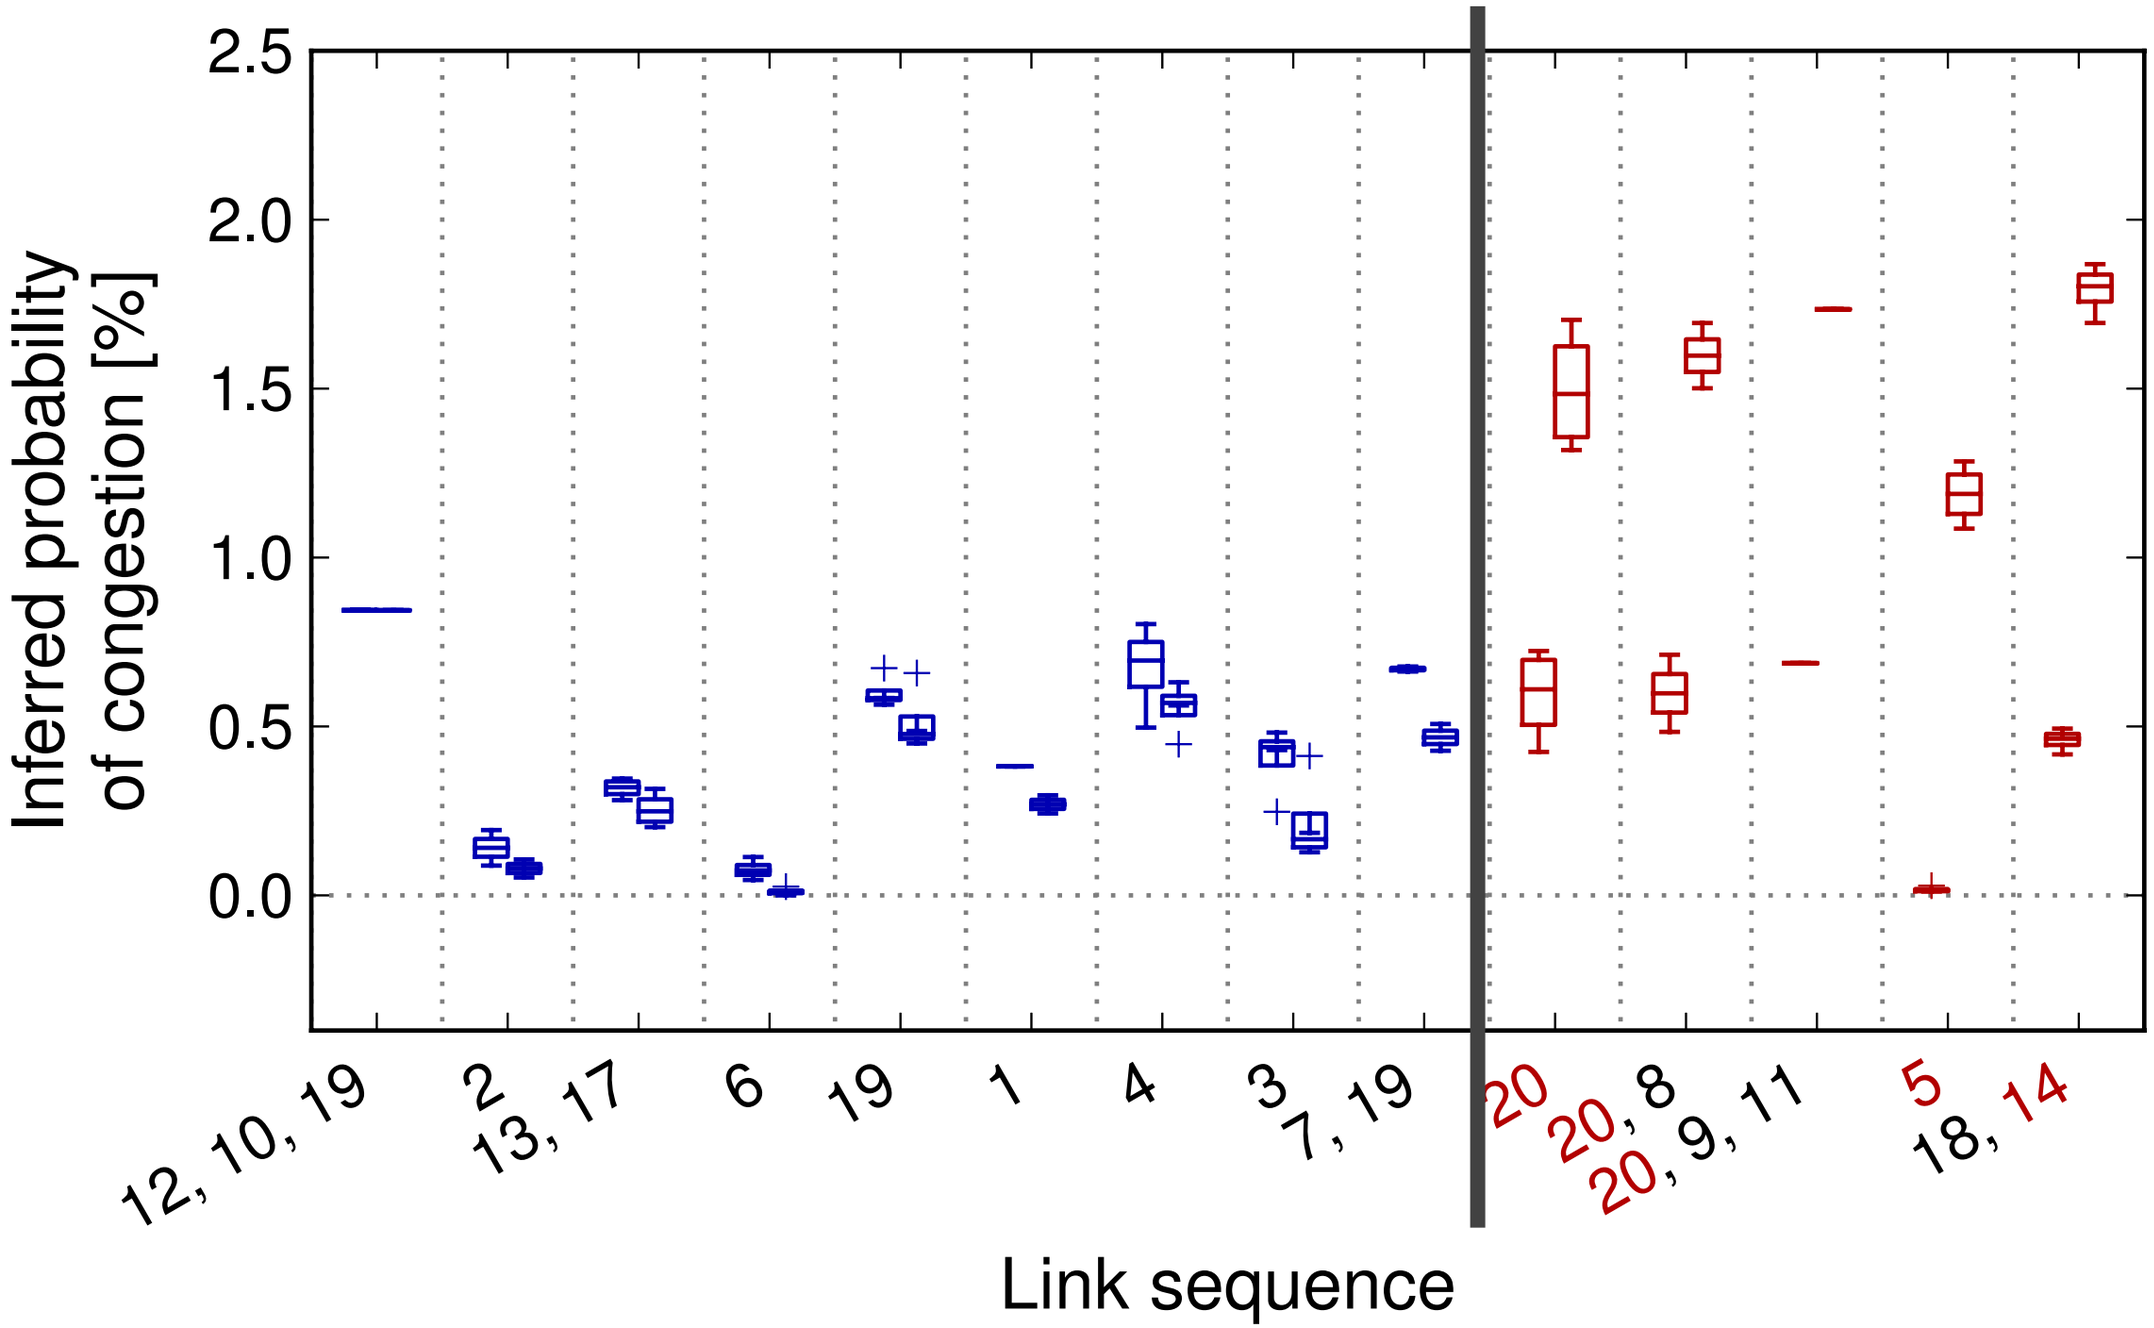

In the paper

- Theoretical assumptions
  - *link independence, class correlation*
- Solutions to practical problems
  - *TCP cautiousness, loss burstiness*
- Other experimental setups
  - *varying flow sizes, RTTs, congestion control algorithms, policing/shaping rates*

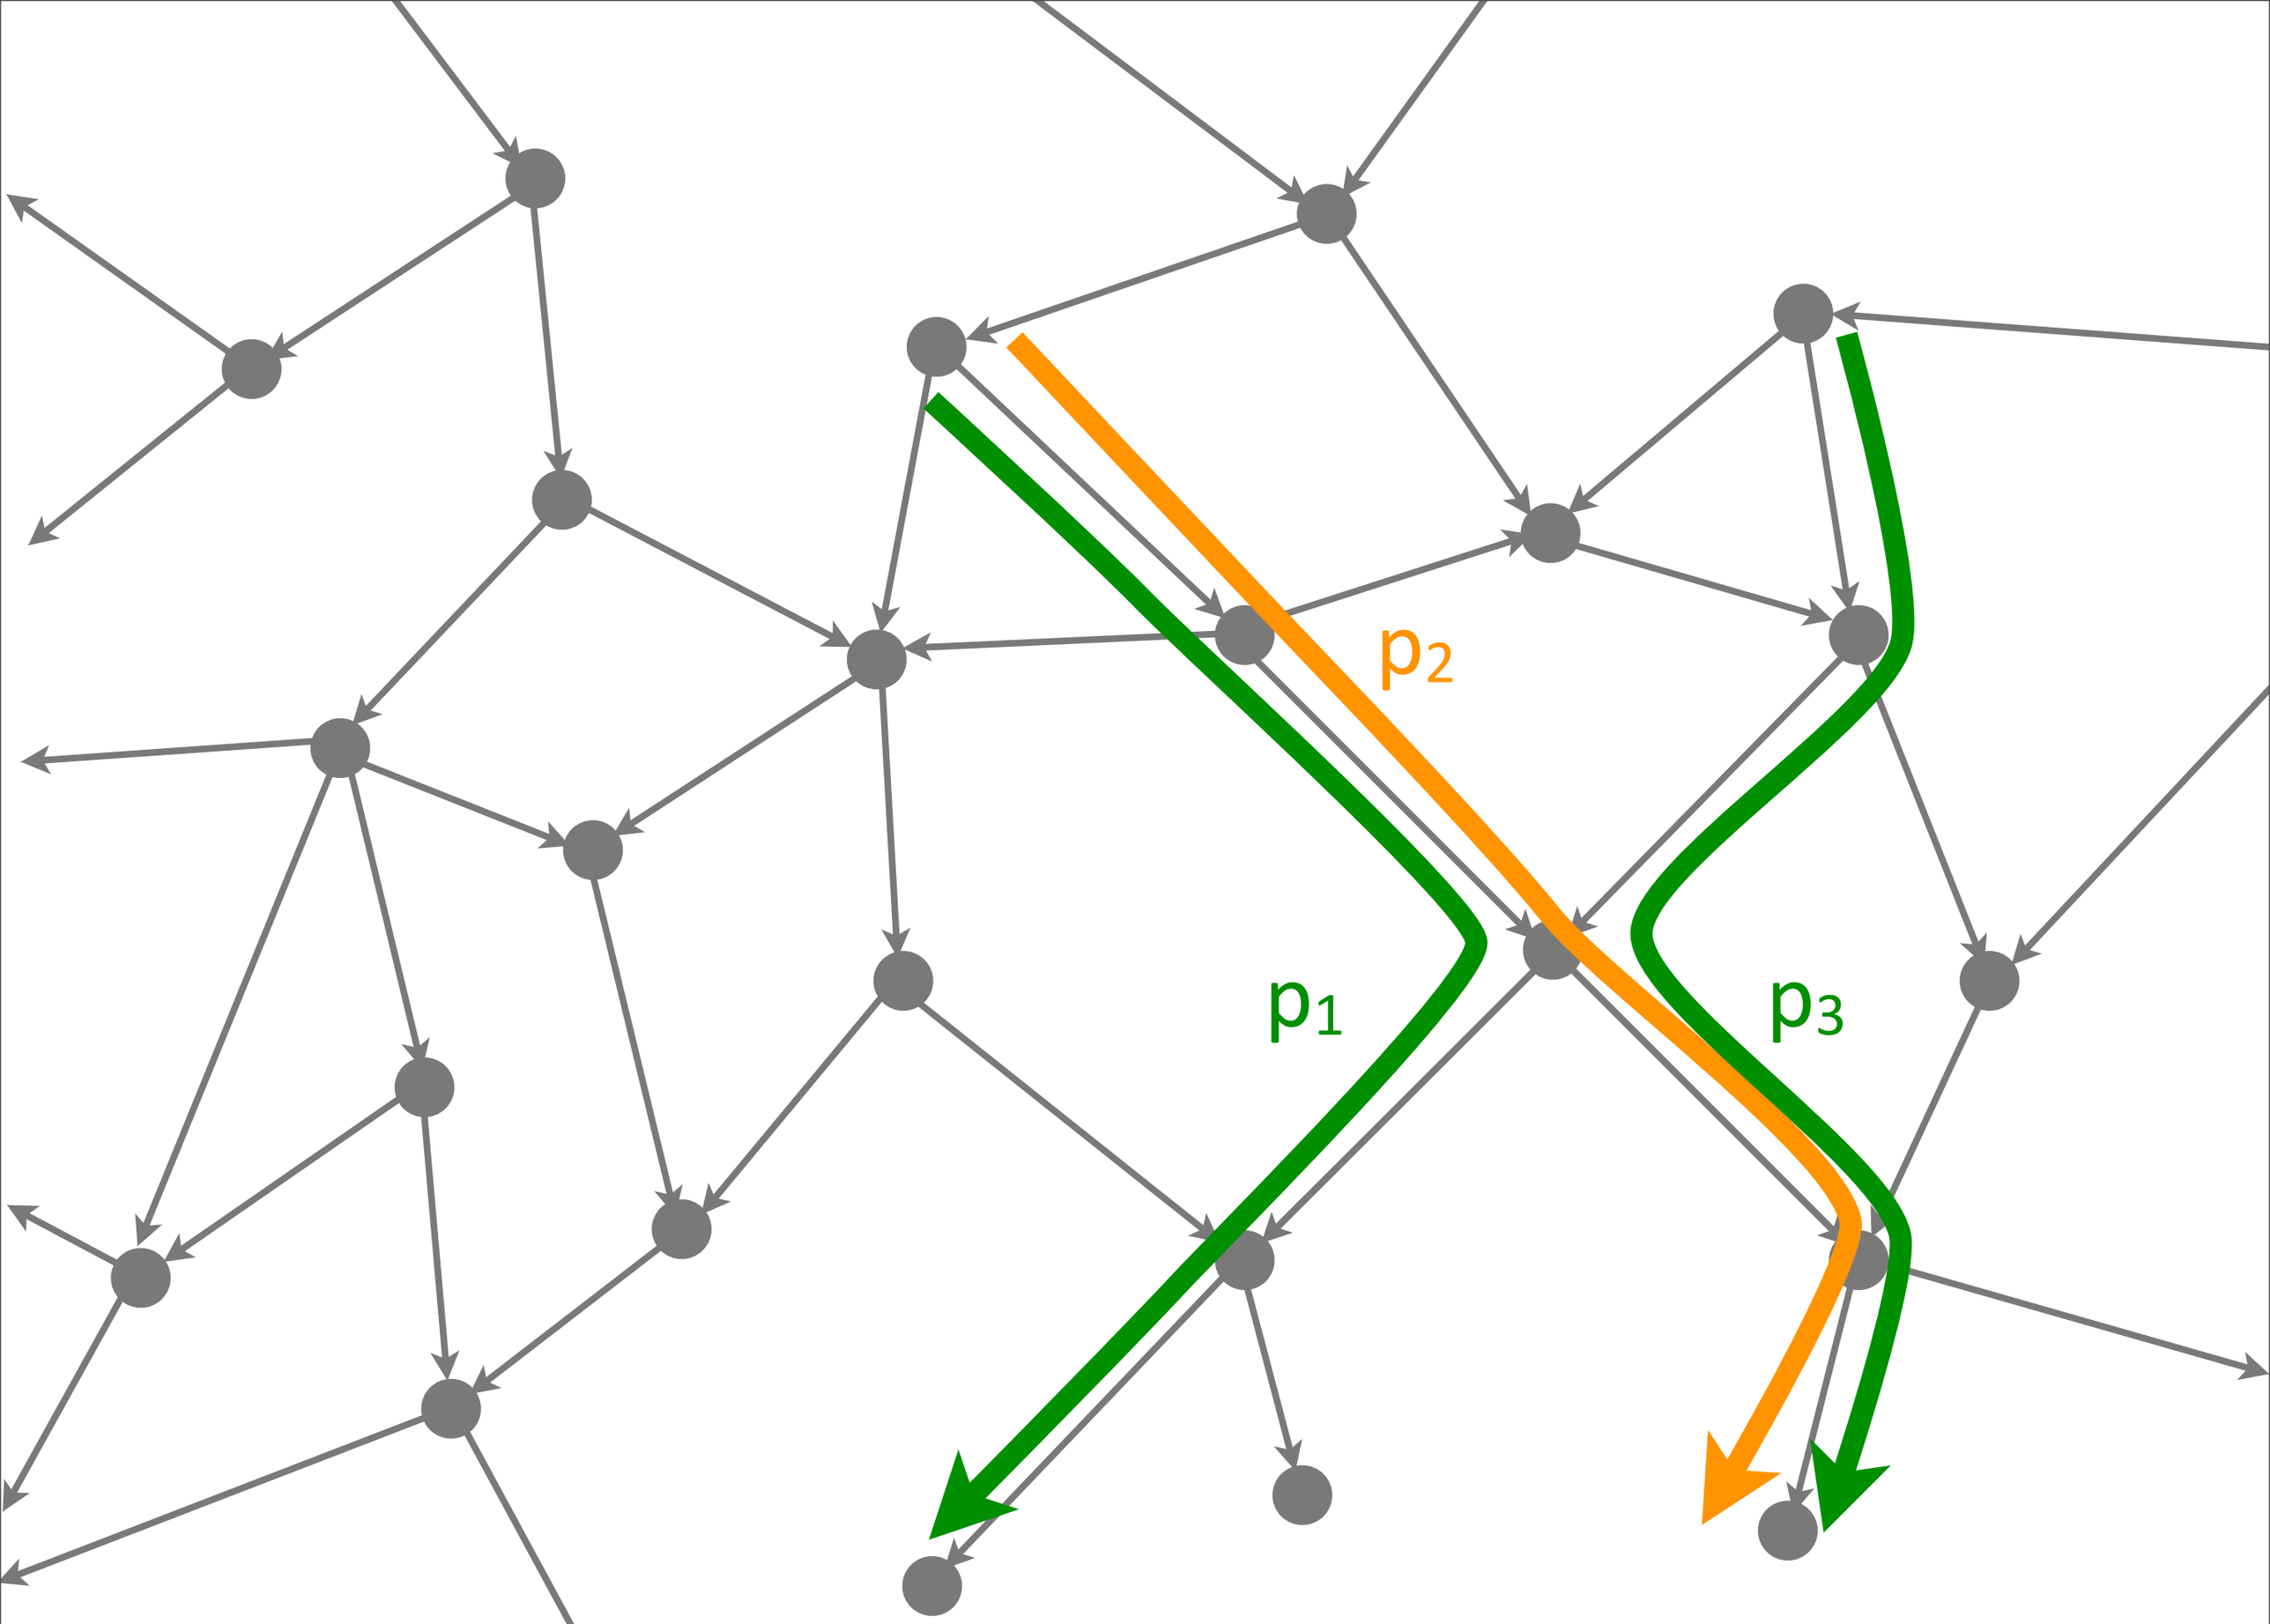
p2
p1
p3

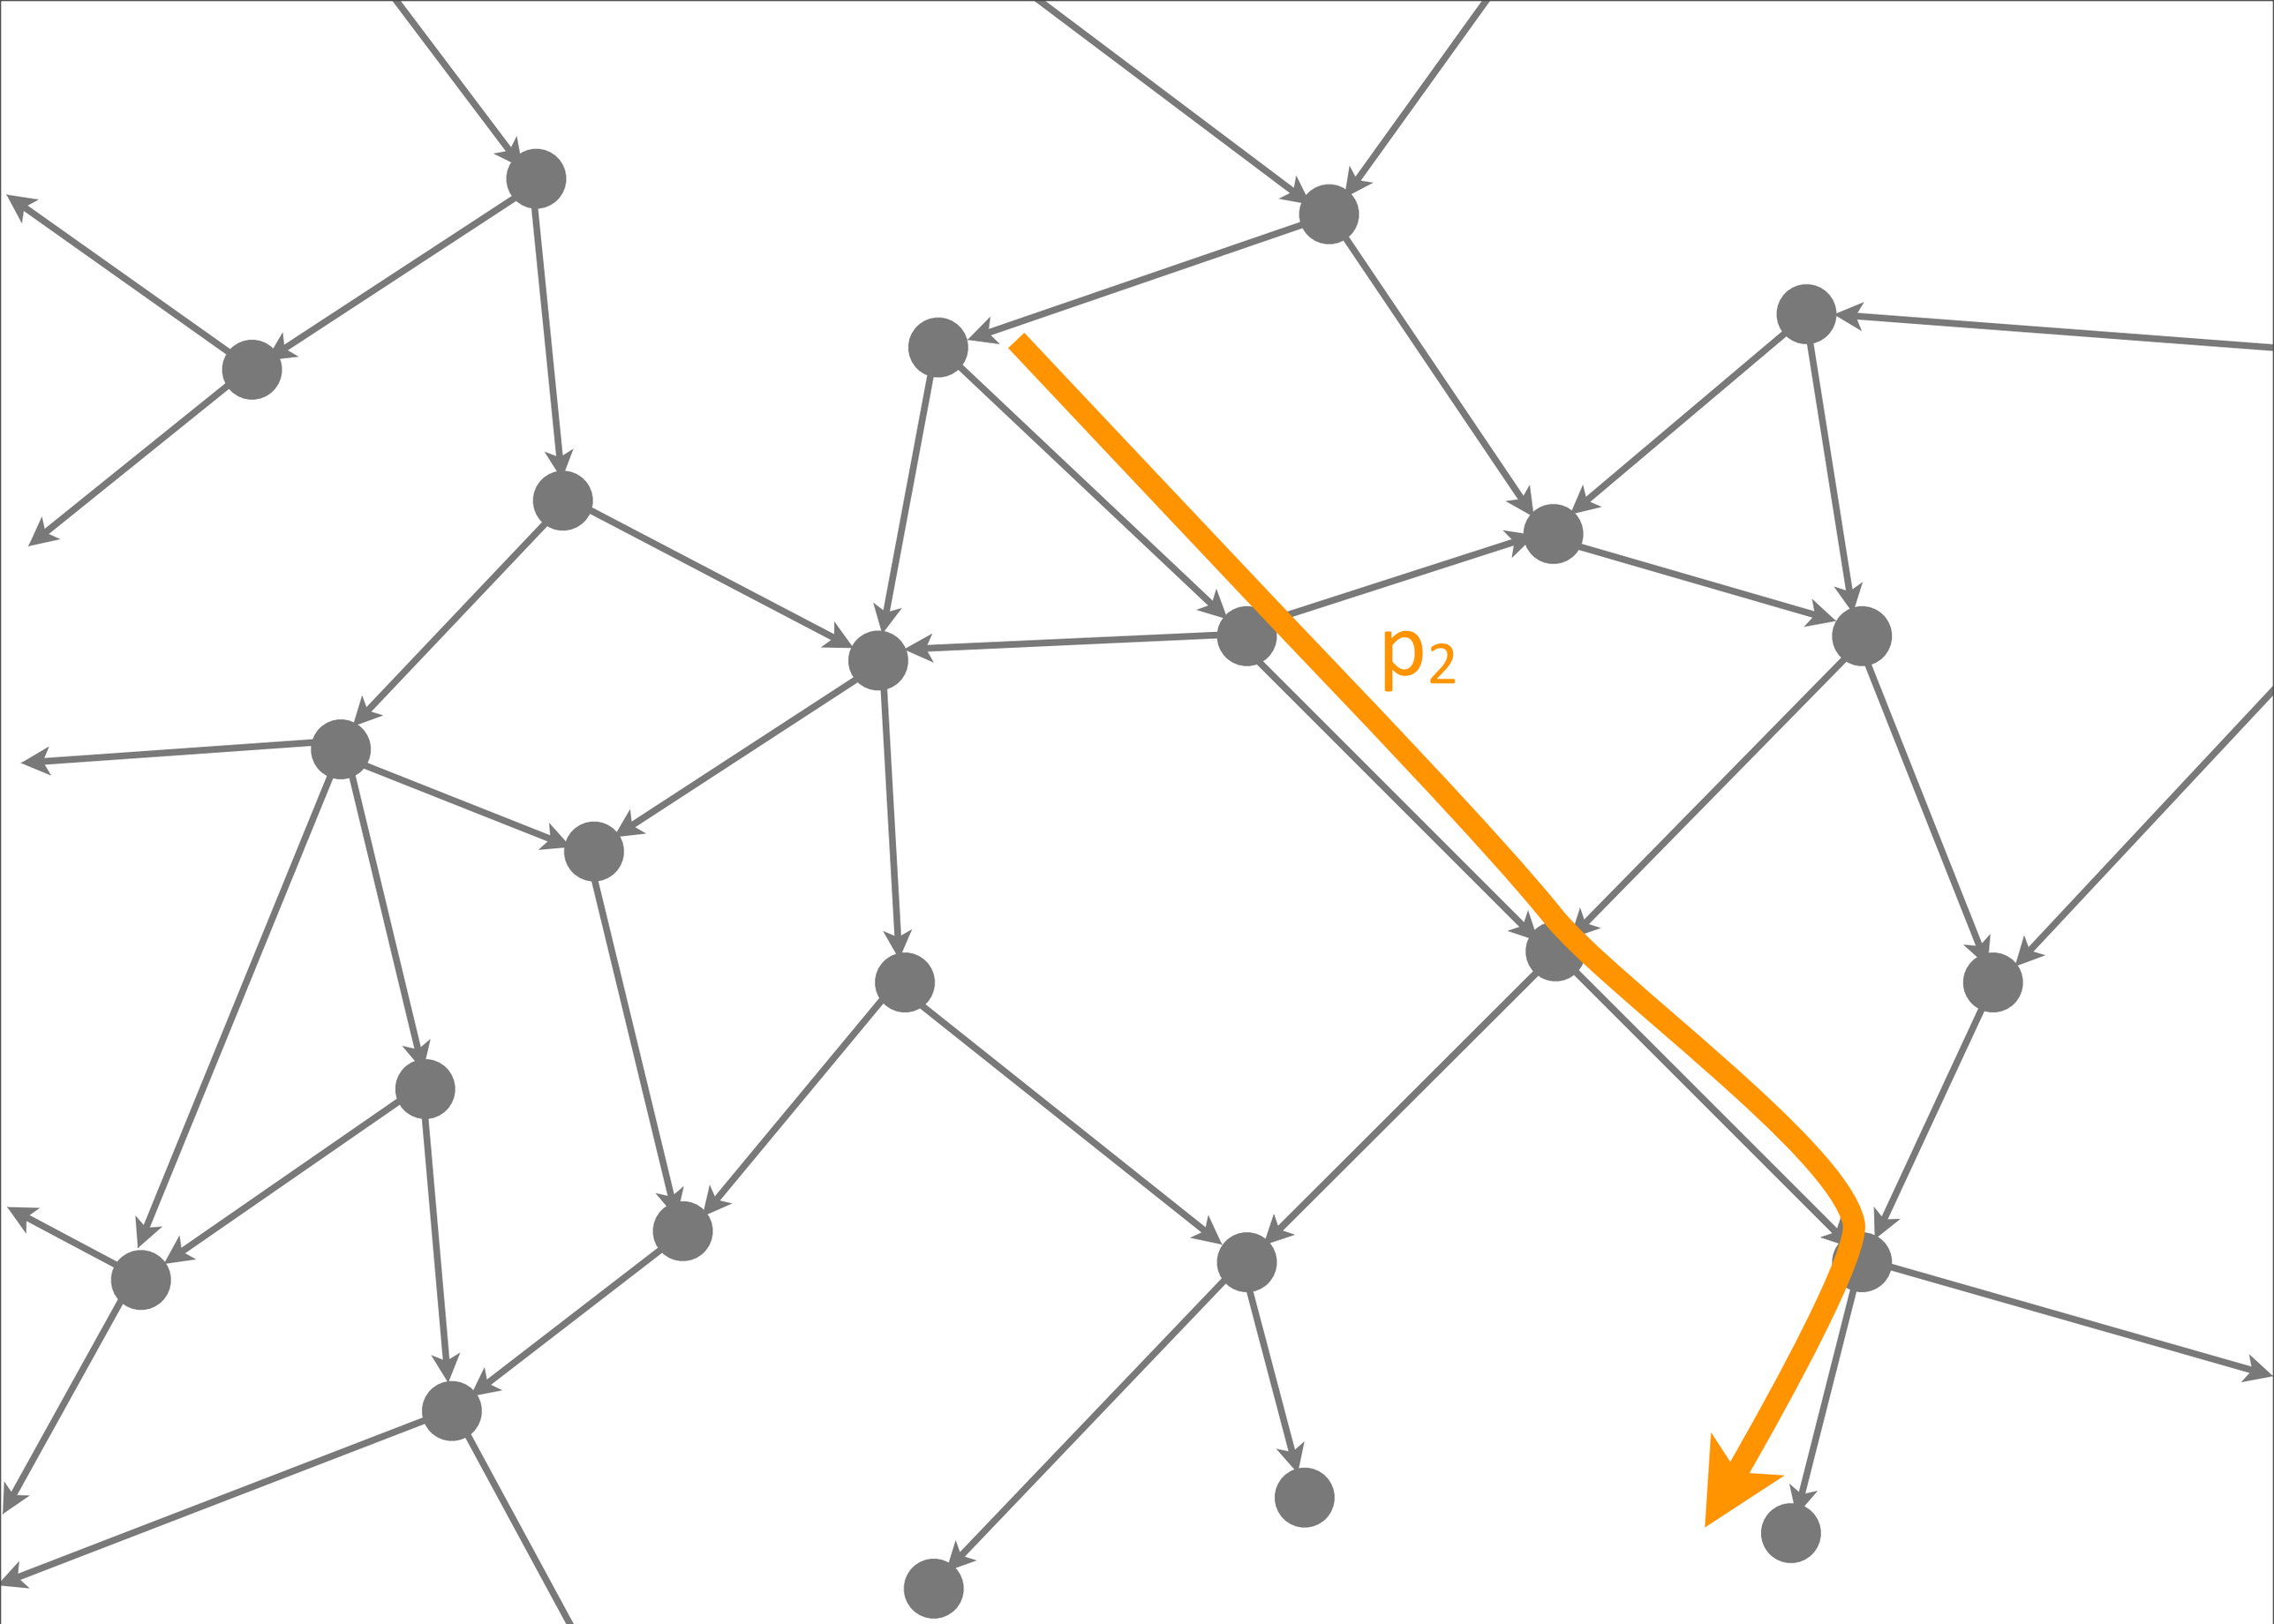
p_2

Related work

- Detection of application-layer and transport-layer differentiation
  - *DiffProbe, Kanuparthy and Dovrolis, 2010*
  - *Glasnost, Dischinger et al., 2010*

- Detection of traffic shaping
  - *ShaperProbe, Kanuparthy and Dovrolis, 2011*
  - *Packsen, Weinsberg et al., 2011*

Contributions

- Detection on neutrality violations independently from differentiation criteria
  - *iff there exists a distinguishable virtual link in the neutral equivalent*

- Localization of neutrality violations to specific link sequences
  - *if there exist two path pairs that intersect exactly at that link sequence (...)*