Improving network agility with seamless BGP reconfigurations

Laurent Vanbever, UCLouvain

laurent.vanbever@uclouvain.be

Princeton University

Monday, April 2 2012

Based on joint work with

Stefano Vissicchio, Luca Cittadini, Cristel Pelsser, Pierre François and Olivier Bonaventure

"When you are changing
the tires of a moving car,

--Vijay Gill

"When you are changing
the tires of a moving car,

make sure one wheel is
on the ground at all time"

--Vijay Gill

Why does seamless reconfigurations matter ?

BGP configuration is critical for ISPs

enforce business relationship

BGP configuration is often changed

On average, > 400 changes accounted per month in a Tier1

Changing a BGP configuration can impact availability

even if the initial and final configurations are safe

Seamless BGP reconfiguration

Progressively modify the BGP configuration of a running network with no impact on the services

Following which order ?

Seamless BGP reconfiguration

Progressively modify the BGP configuration of a running network with no impact on the services

Seamless BGP reconfiguration

Progressively modify the BGP configuration of a running network with no impact on the services

Approach

Provable reconfiguration strategies

Improving network agility with seamless BGP reconfigurations

BGP background

Finding an ordering

Reconfiguration framework

Improving network agility with seamless BGP reconfigurations

BGP background

Finding an ordering

Reconfiguration framework

BGP is the only inter-domain routing protocol used today

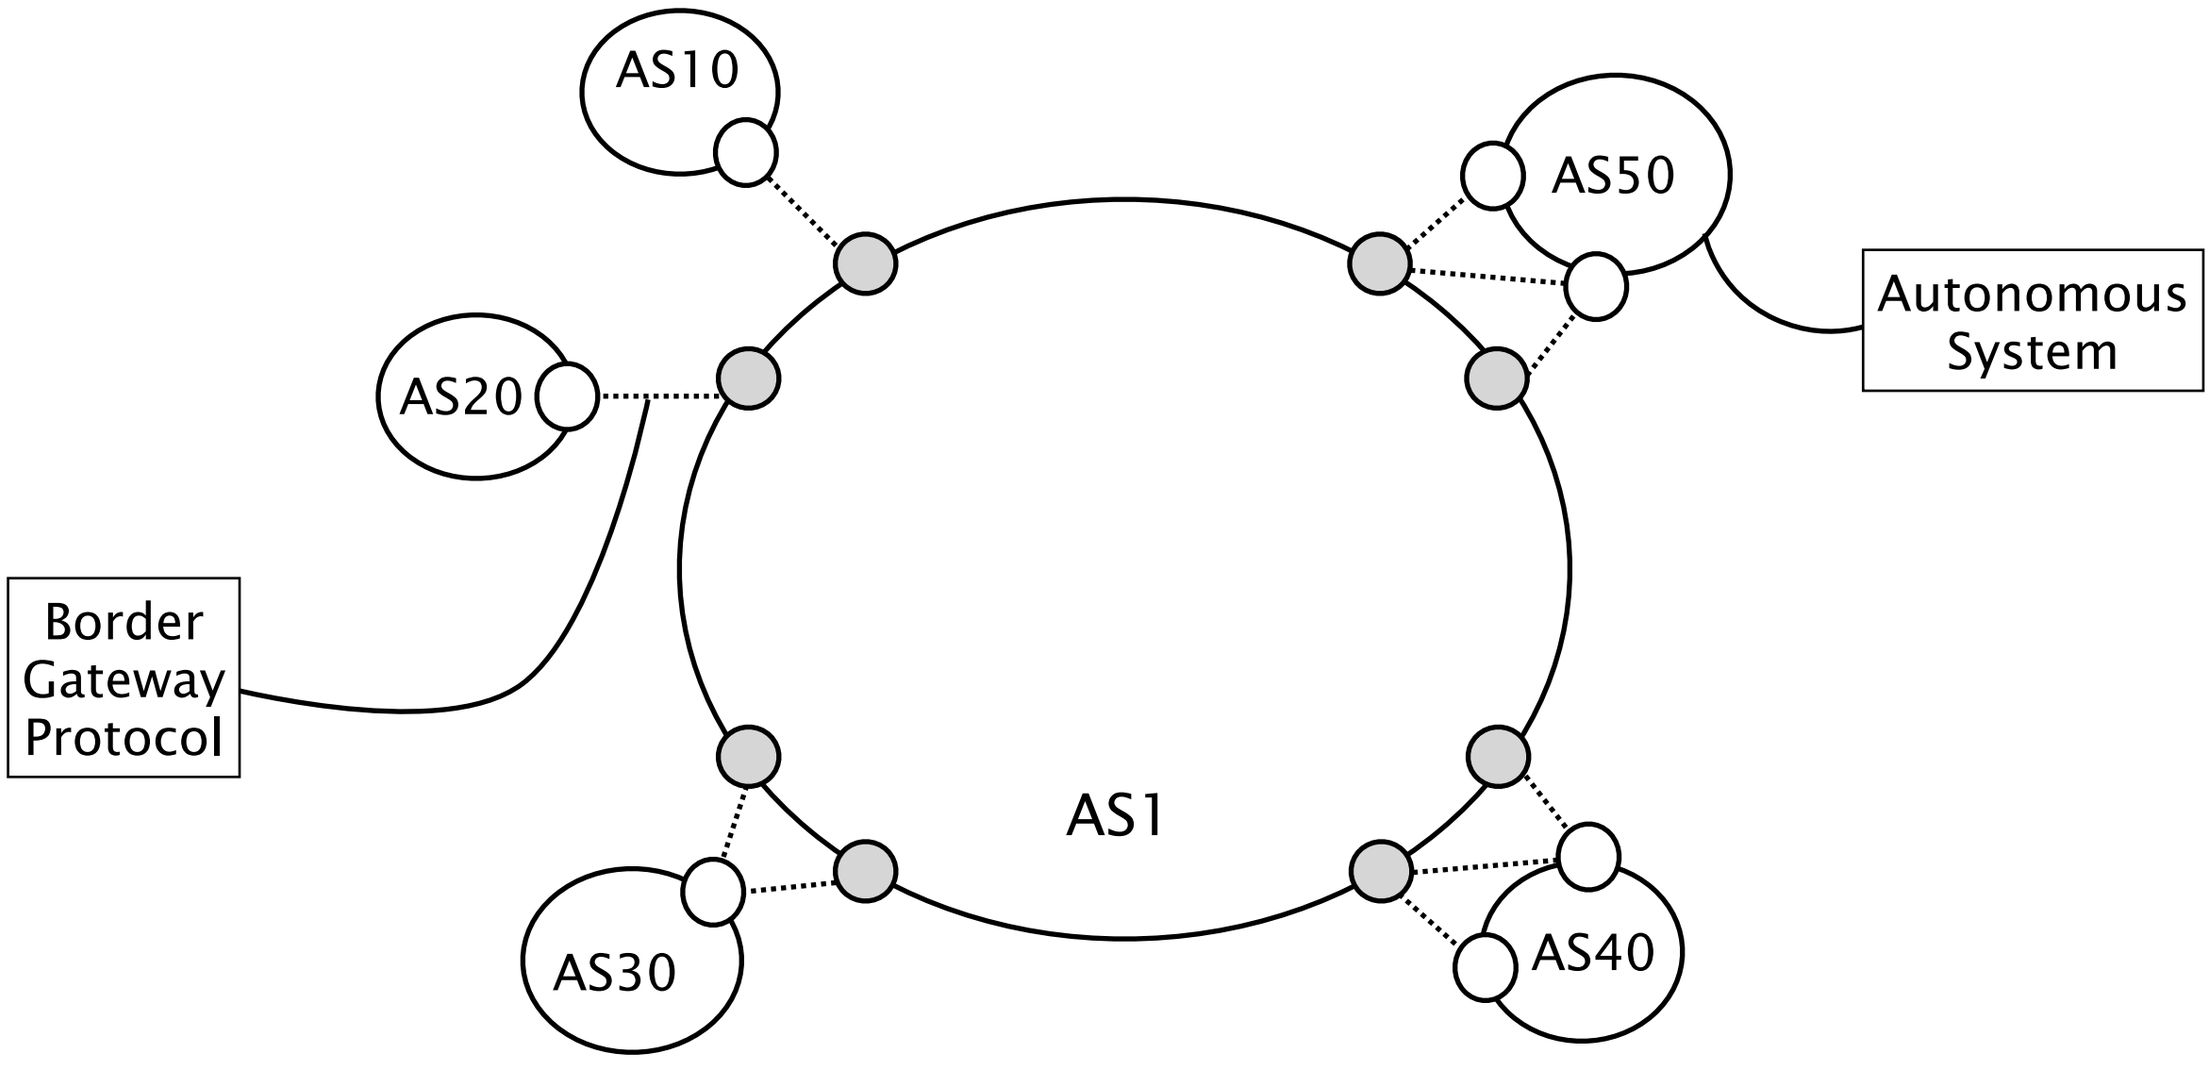

BGP comes in two flavors

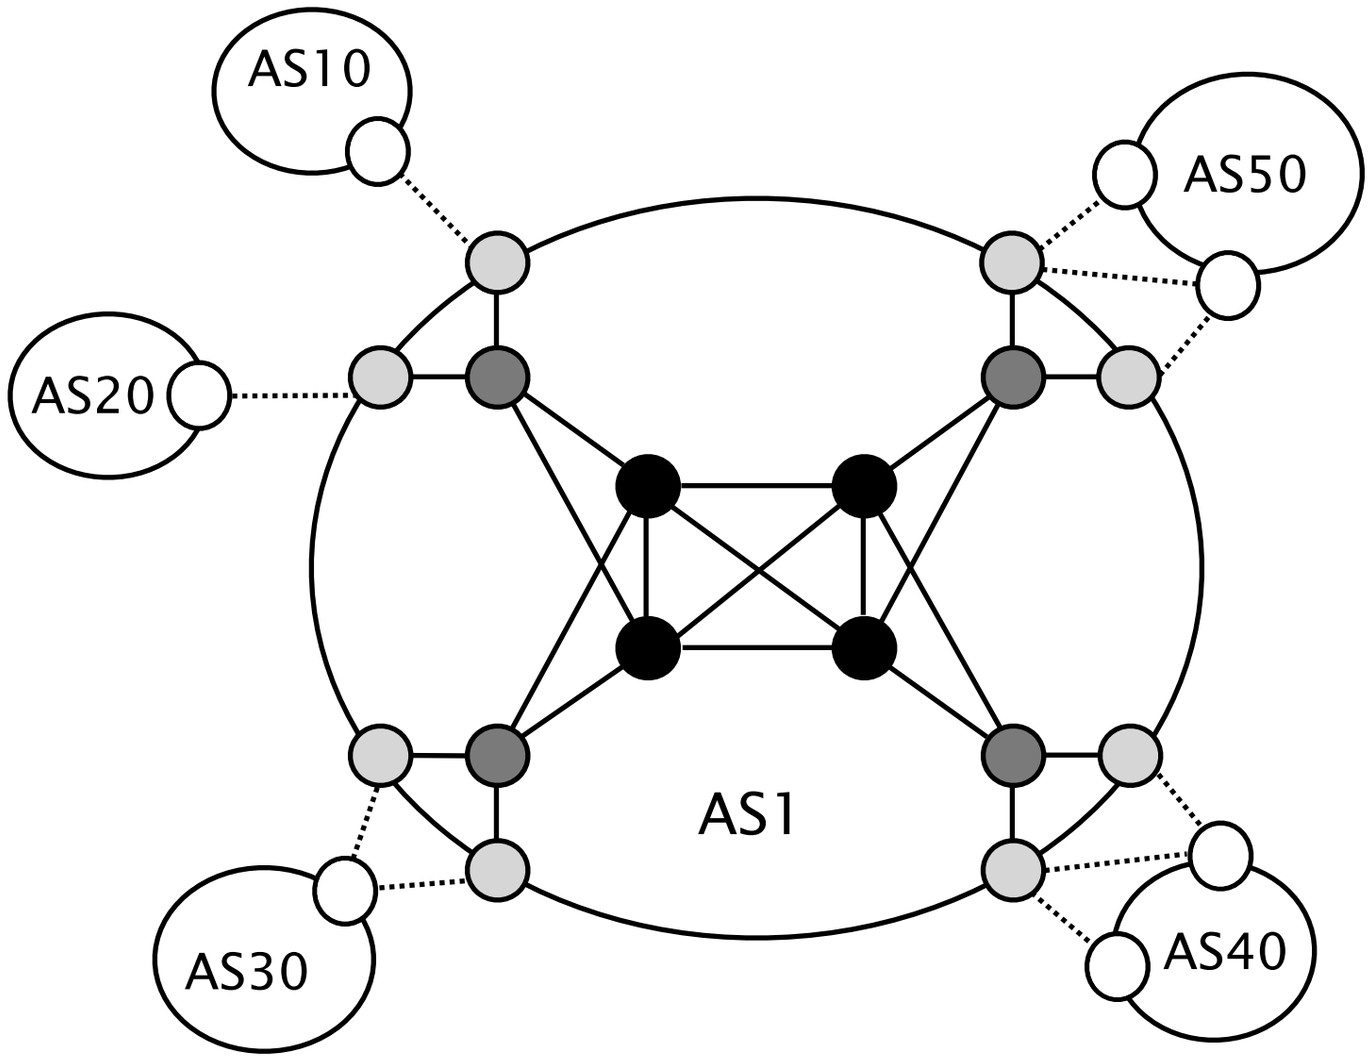

external BGP (eBGP) exchanges reachability information between ASes

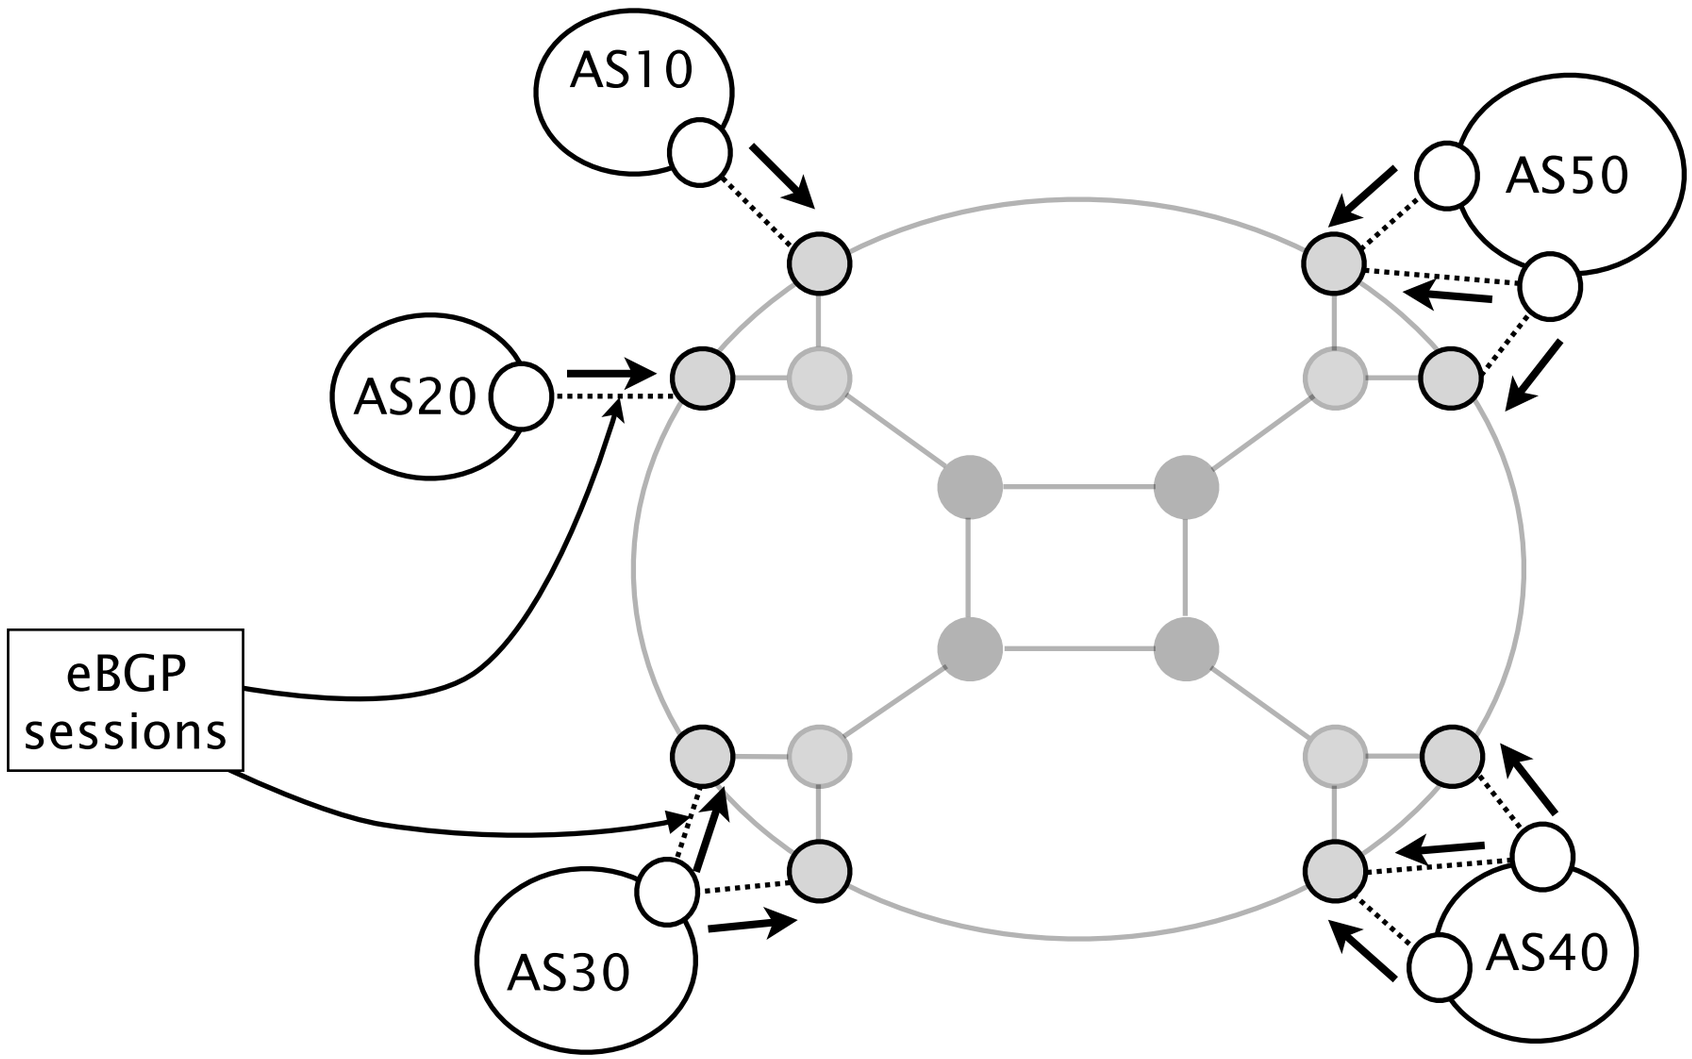

internal BGP (iBGP) distributes externally learned routes within the AS

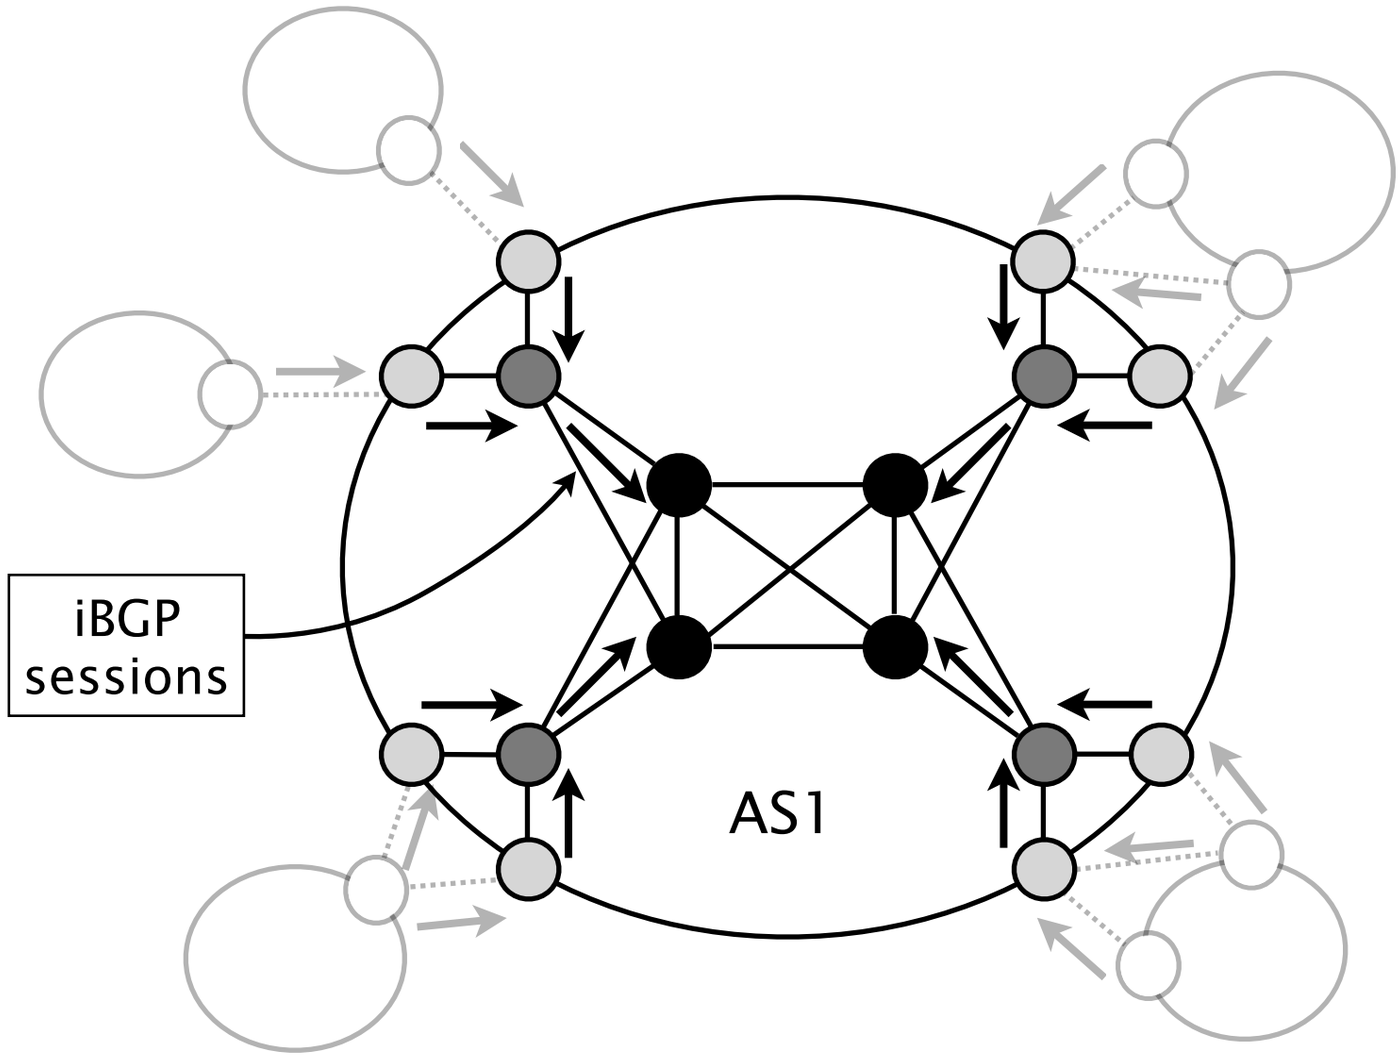

To scale, iBGP routers are often organized in a hierarchy

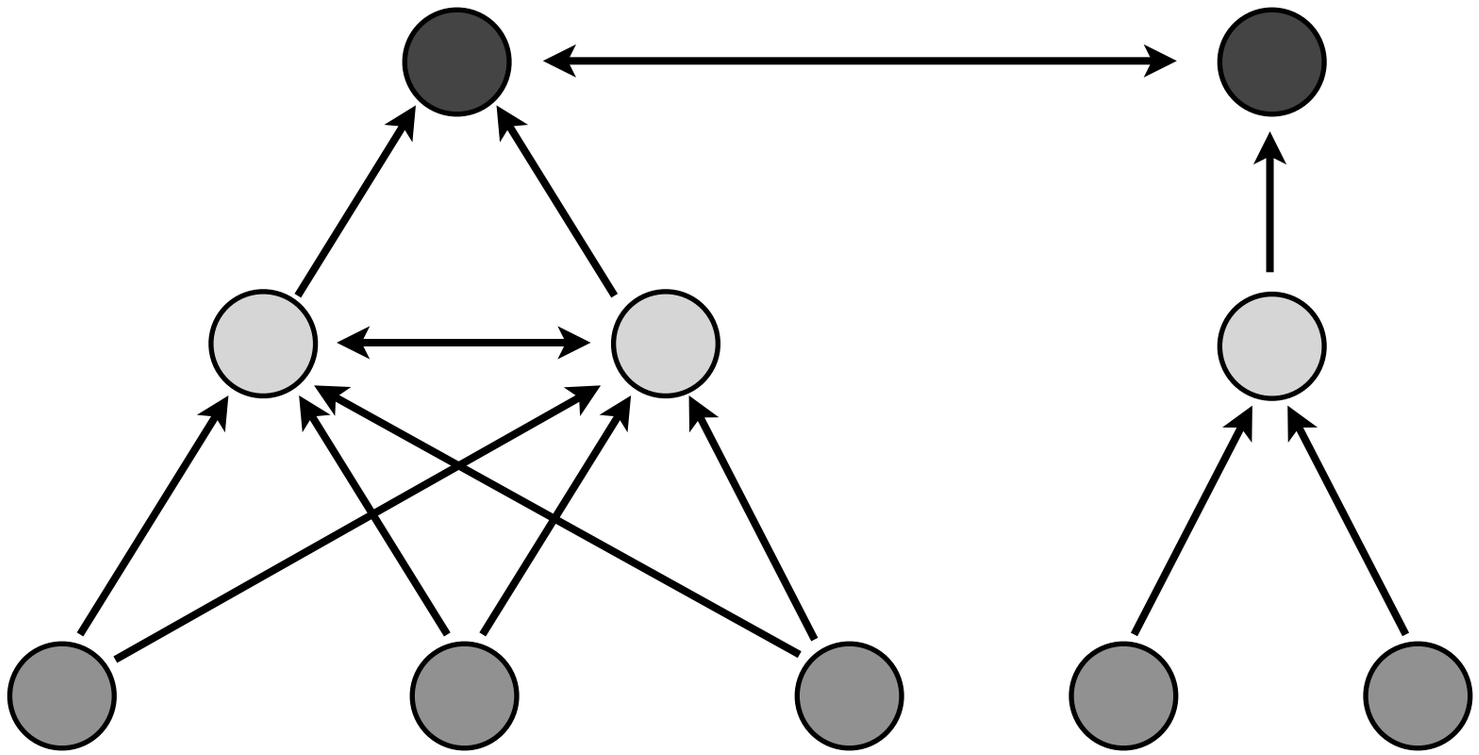

Lower layers rely on upper layers to learn and propagate routing informations

Reconfiguring BGP can be disruptive

BGP reconfigurations can create

- signaling anomalies [Griffin, SIGCOMM02]
- dissemination anomalies [Vissicchio, INFOCOM12]
- forwarding anomalies [Griffin, SIGCOMM02]

or any combination of those

Reconfiguring BGP can be disruptive

BGP reconfigurations can create

- signaling anomalies — routing oscillations
- dissemination anomalies — black holes
- forwarding anomalies — forwarding loops, traffic shifts

or any combination of those

Reconfiguring BGP can be disruptive

BGP reconfigurations can create

- signaling anomalies
- dissemination anomalies
- forwarding anomalies

or any combination of those

How much ?

Let’s migrate from a iBGP full-mesh to an iBGP hierarchy

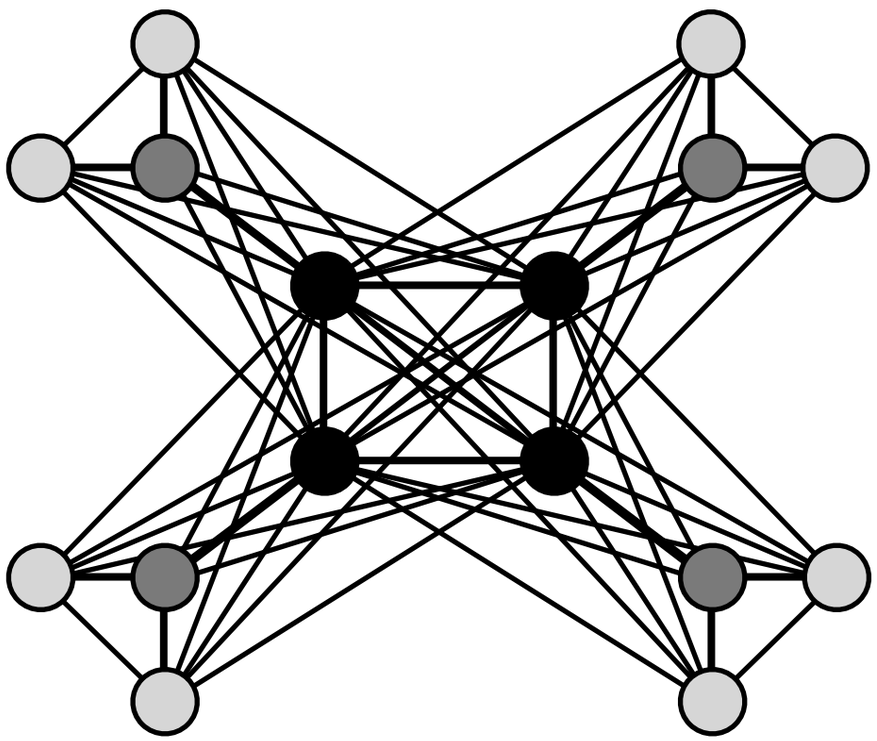

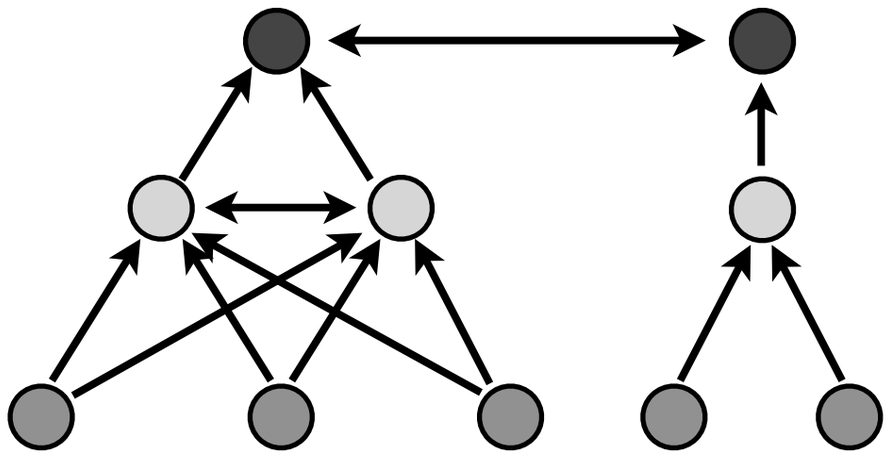

Let's migrate from a iBGP full-mesh to an iBGP hierarchy, following best practices

Establish the RR sessions in a bottom-up manner, then remove the full-mesh sessions

[Herrero10]

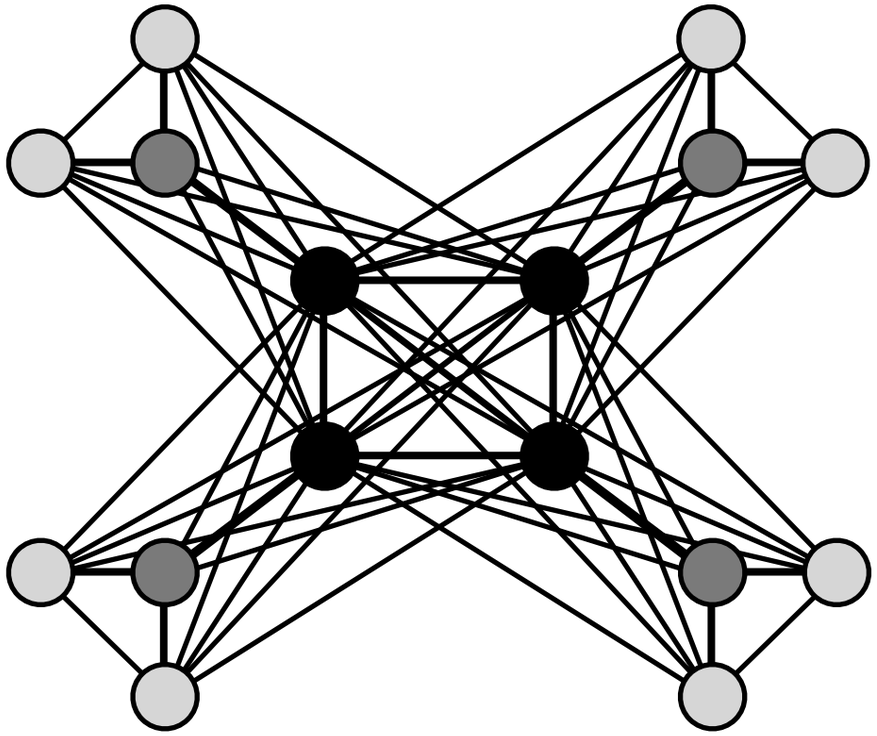

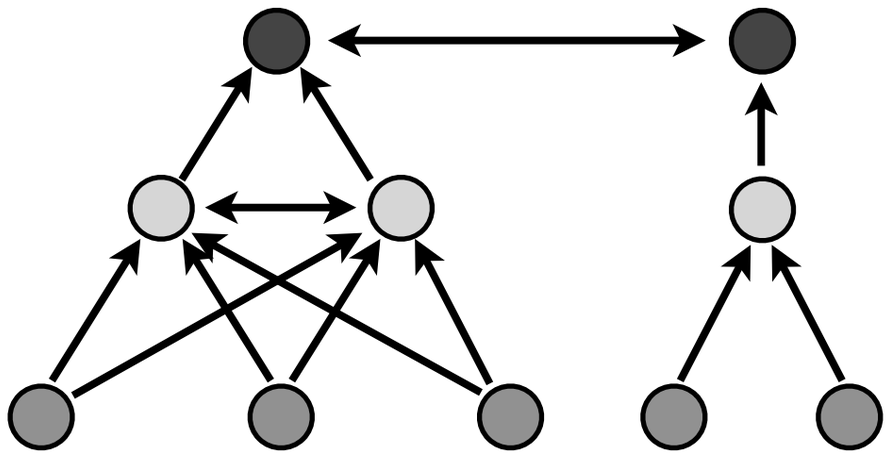

Best practices do not work

Tier1 (50) experiments
(cumul. frequency)

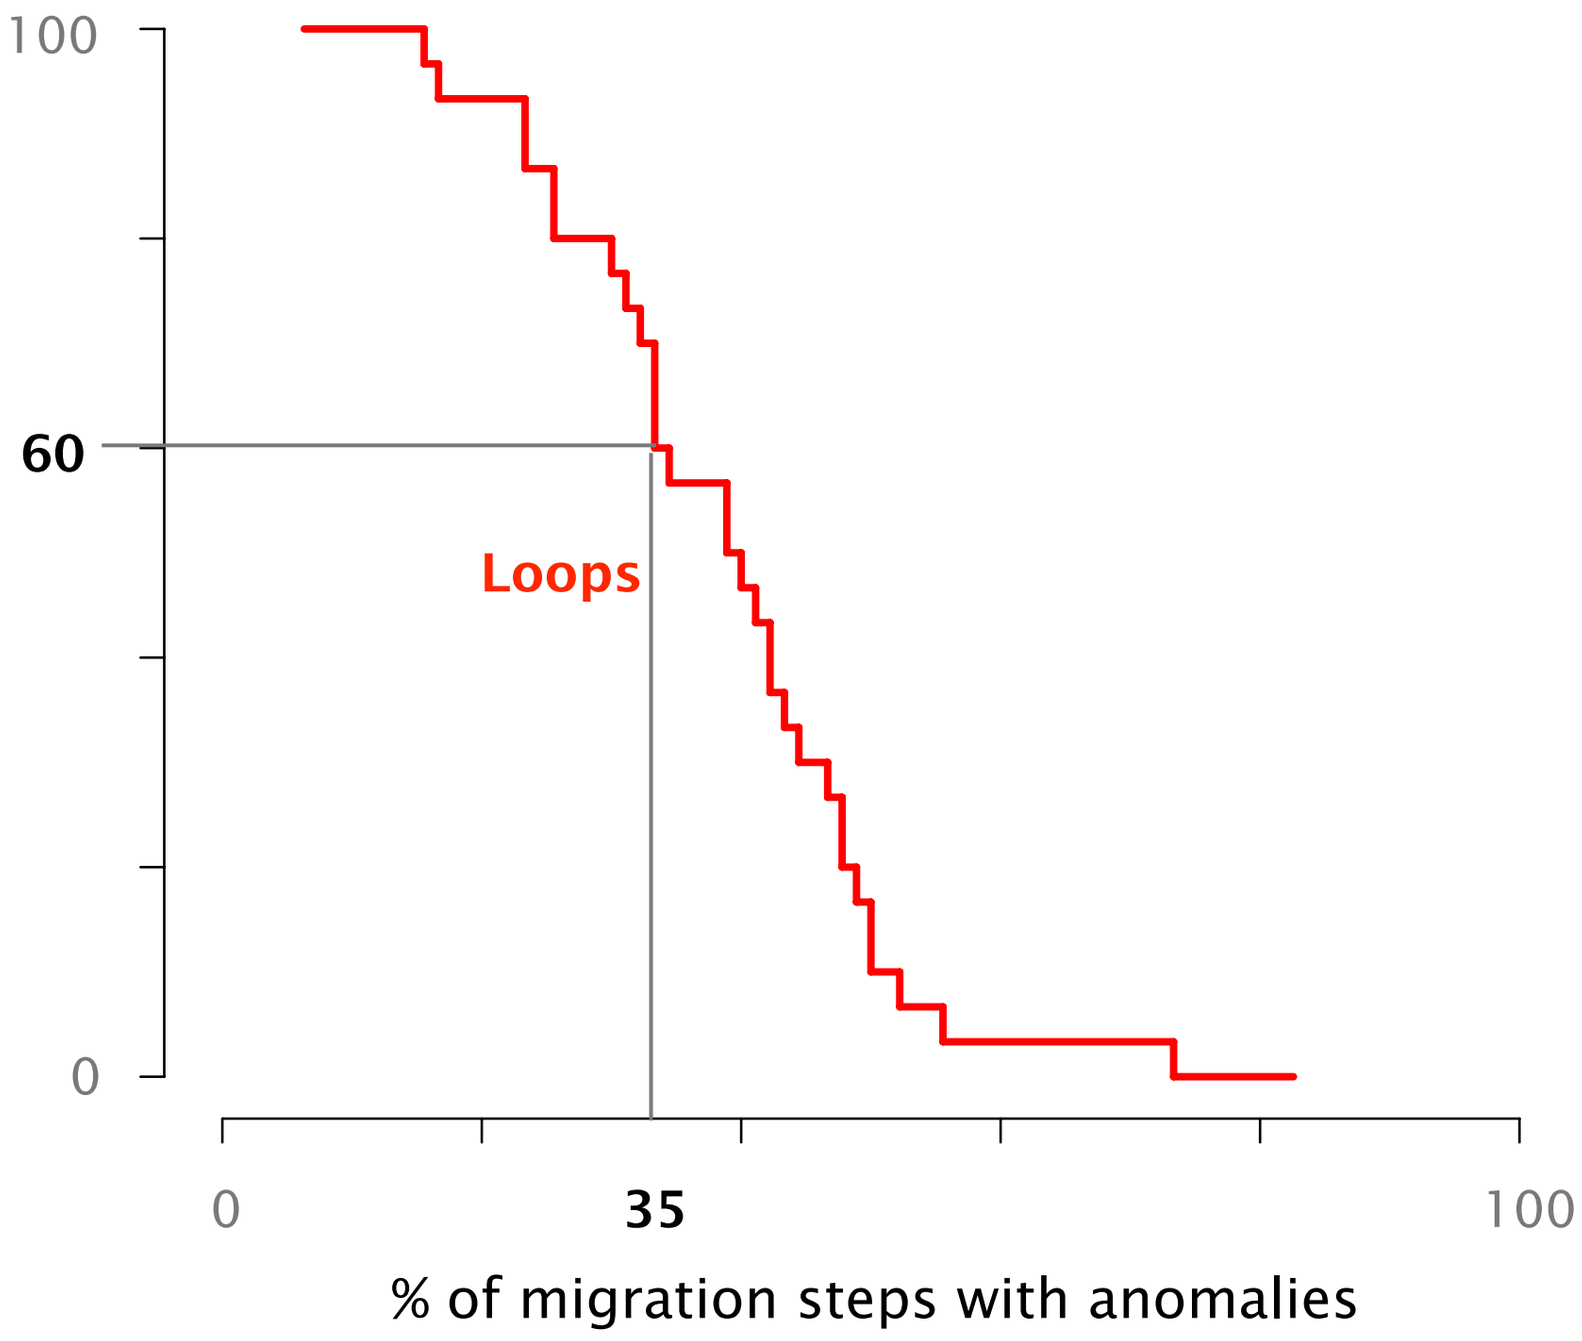

60% of the experiments
were subject to loops
for > **35%** of the steps

Best practices do not work

Tier1 (50) experiments
(cumul. frequency)

100% of the experiments were subject to traffic shifts for > **40%** of the steps

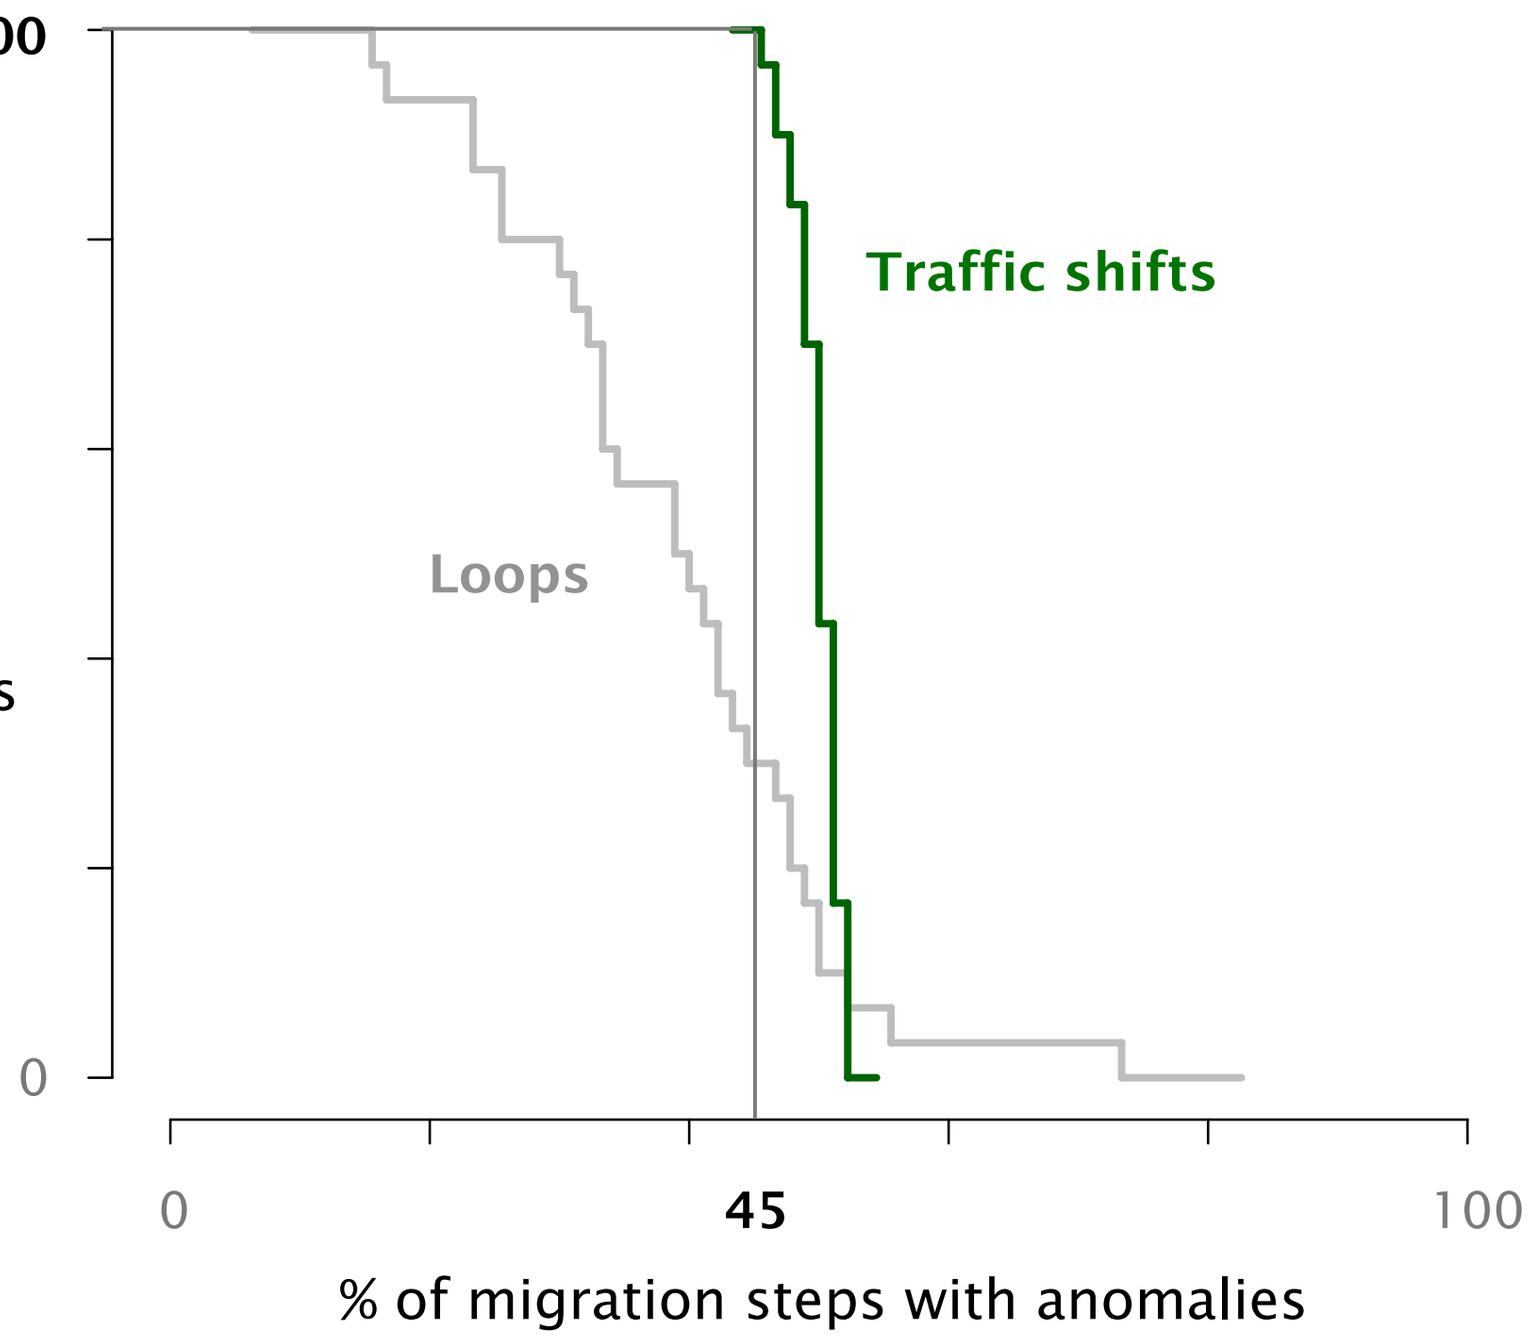

Improving network agility with seamless BGP reconfigurations

BGP background

Finding an ordering

Reconfiguration framework

To avoid reconfiguration problems, a proper operational ordering must be enforced

Given an initial & final, anomaly-free, BGP configuration.

Find a sequence of configuration changes such that

- signaling anomalies
- dissemination anomalies
- forwarding anomalies

never occur, during any migration step

Find a sequence of configuration changes

Find a sequence of configuration changes

Does it always exist ?

Find a sequence of configuration changes

Does it always exist ?

Is it easy to compute ?

We model iBGP configurations by using extended Stable Path Problem instances

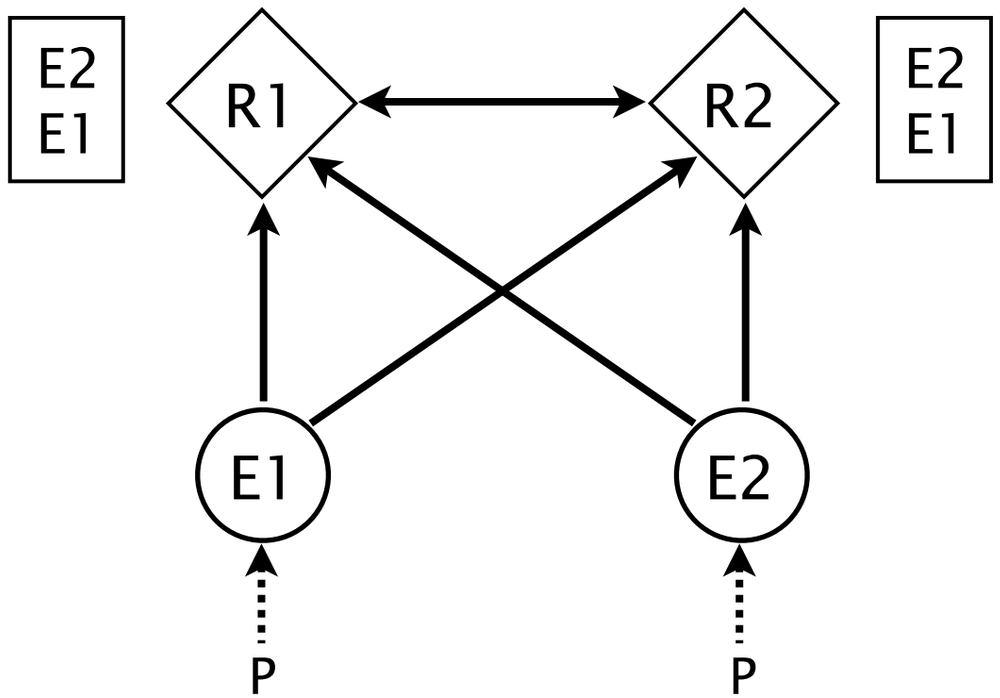

We model iBGP configurations by using extended Stable Path Problem instances

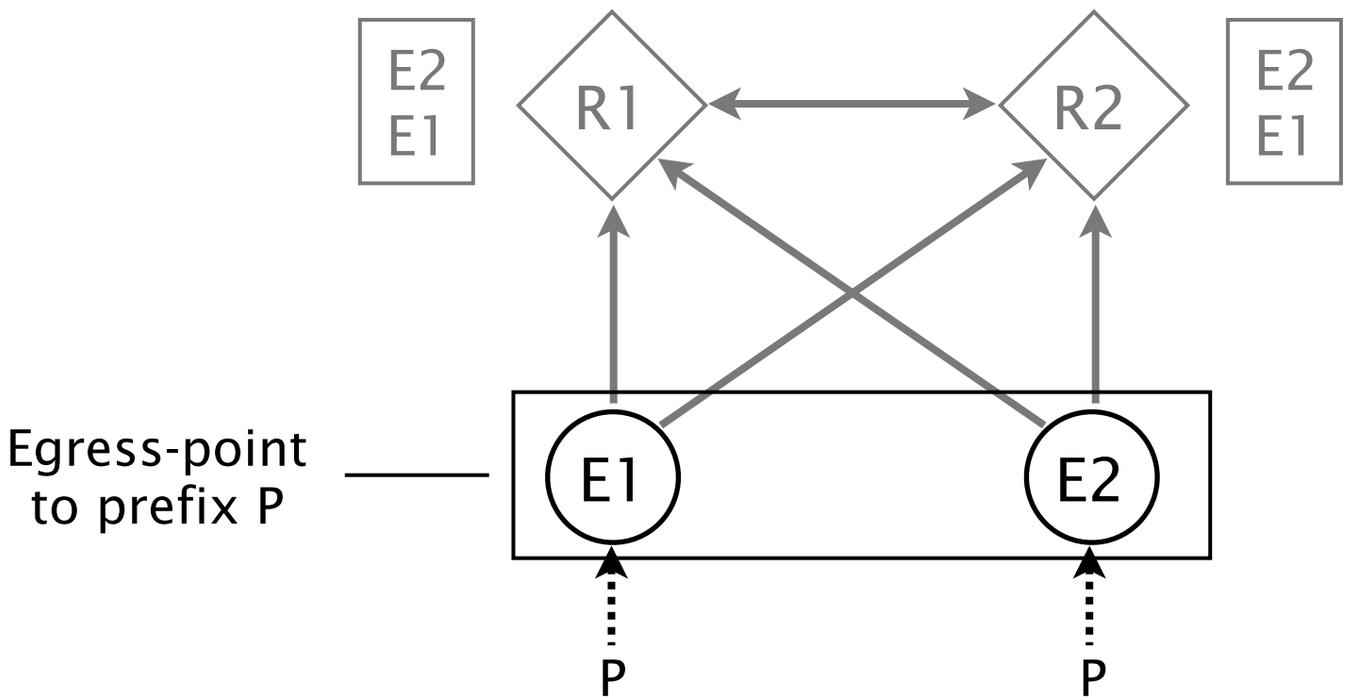

We model iBGP configurations by using extended Stable Path Problem instances

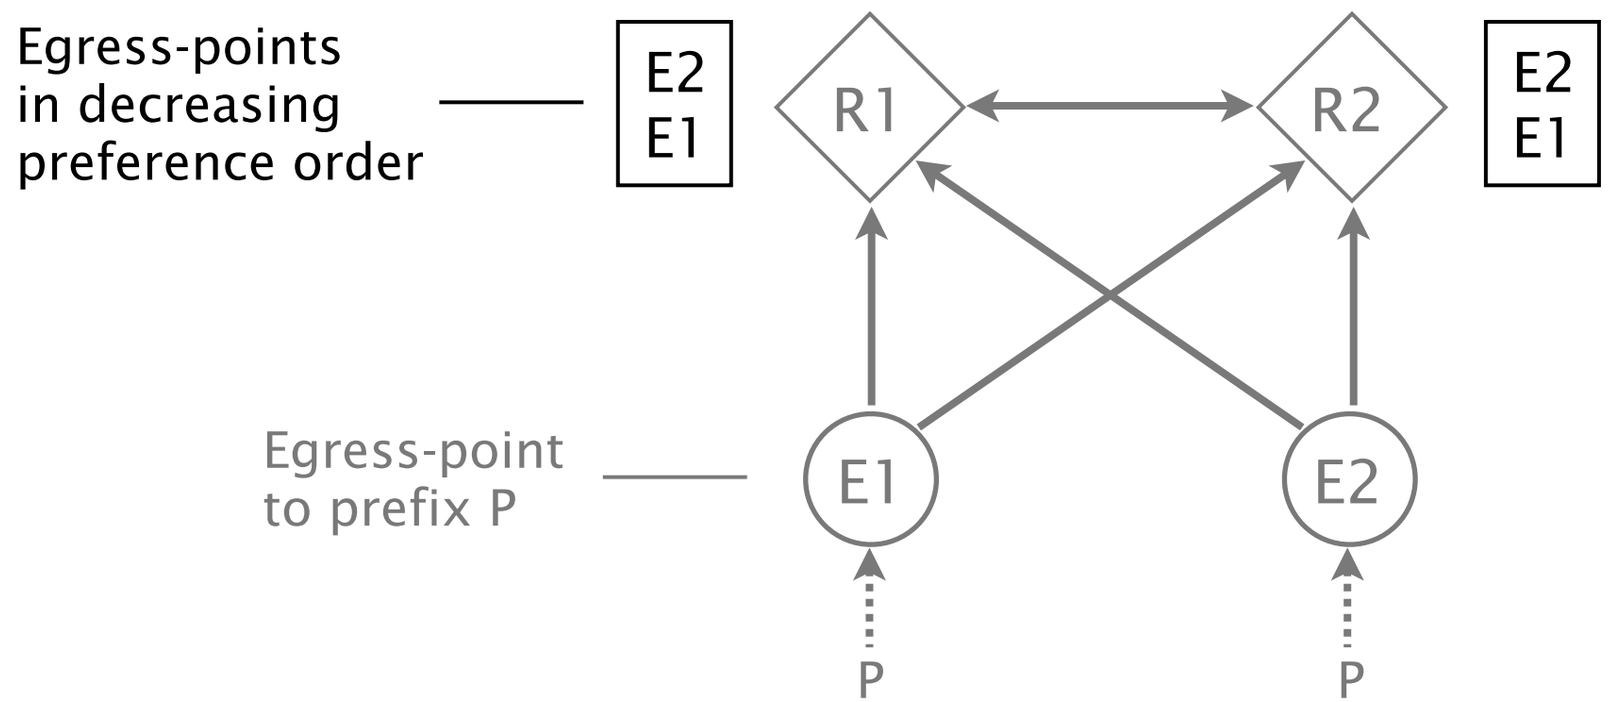

We model iBGP configurations by using extended Stable Path Problem instances

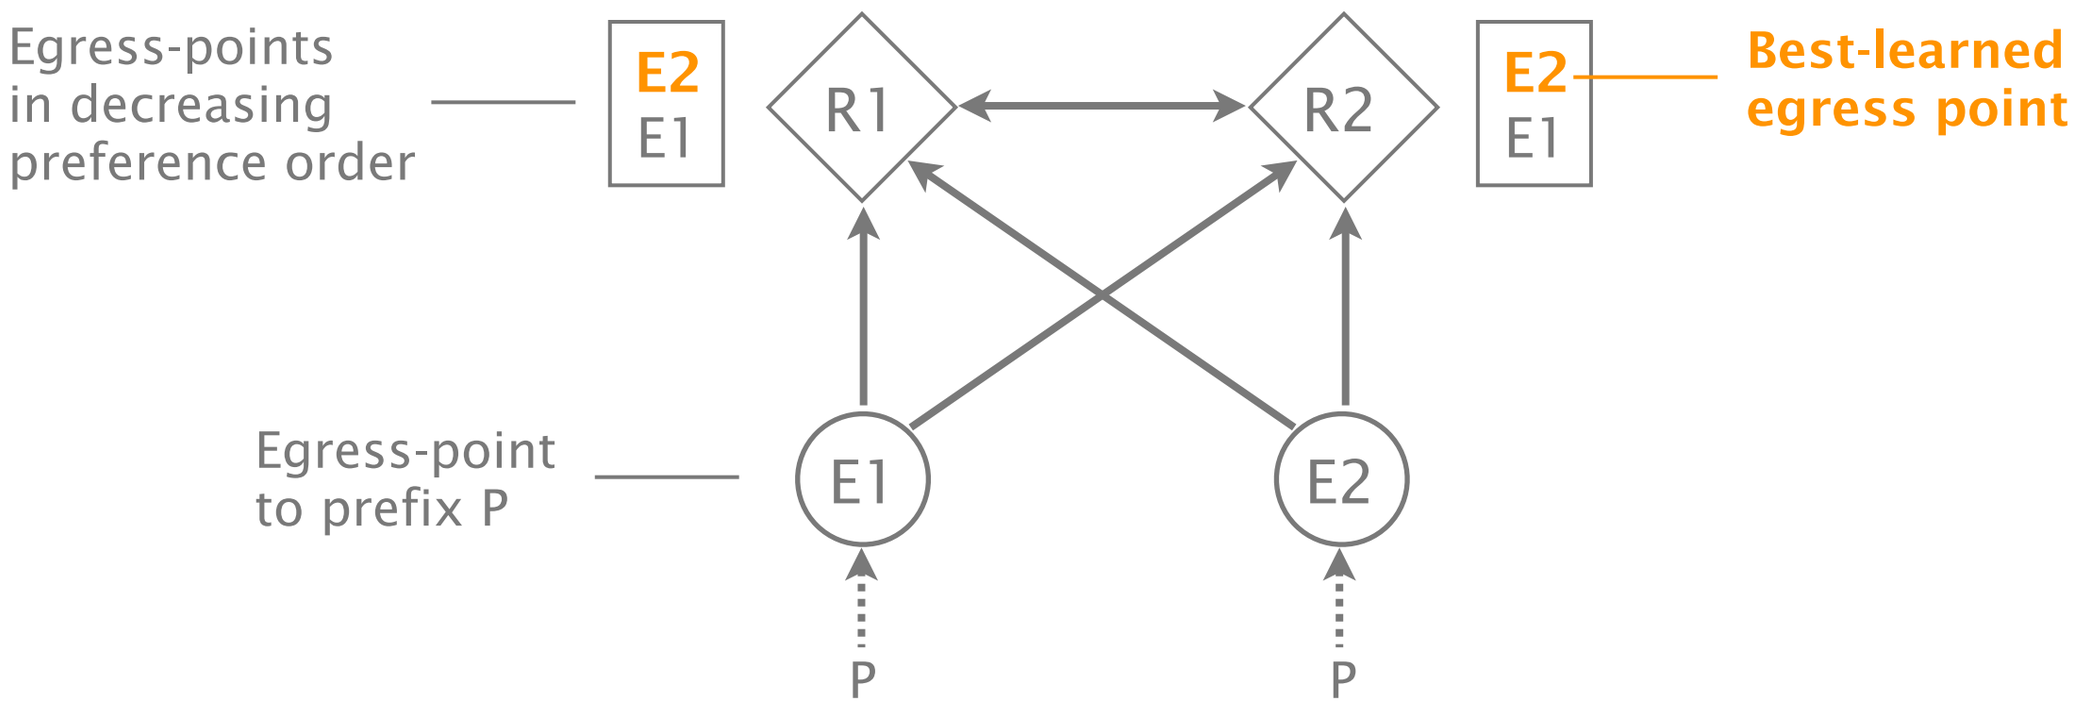

A stable BGP configuration determines the forwarding paths being used

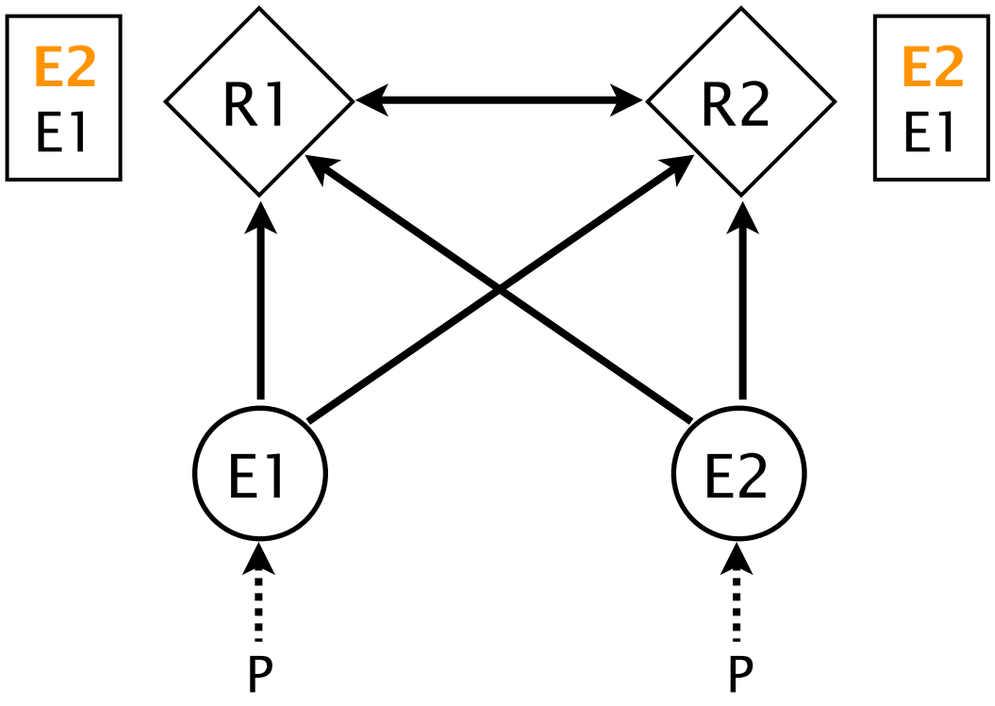

BGP configuration

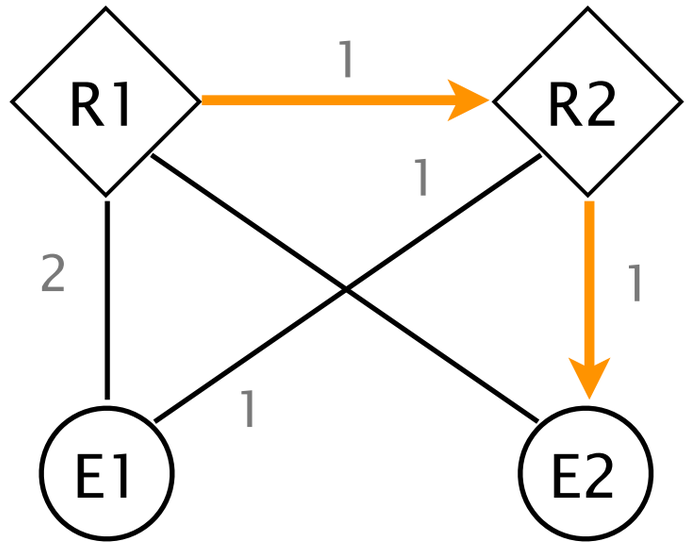

IGP configuration

A seamless migration ordering might not always exist

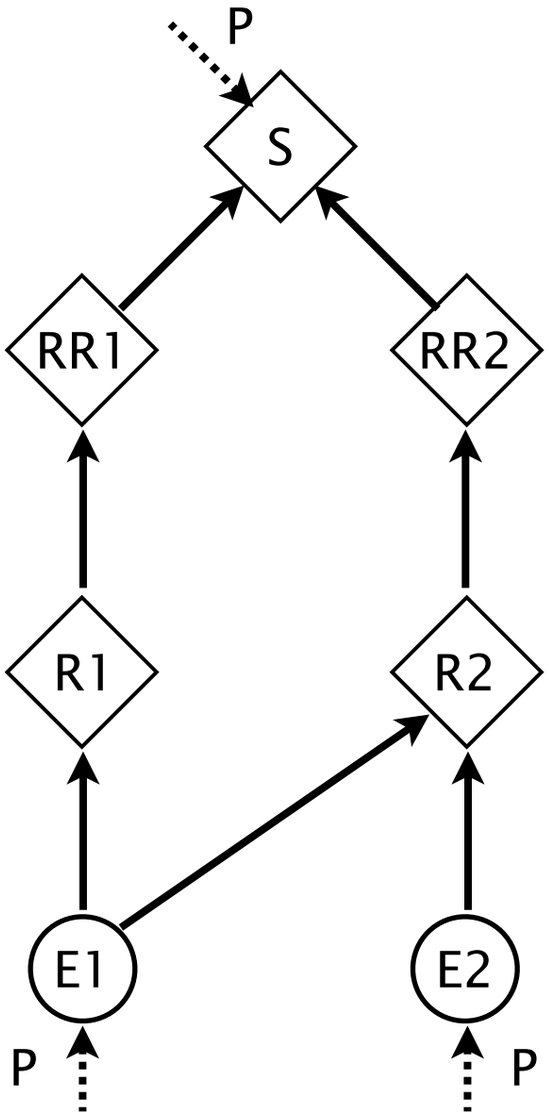

Initial BGP configuration

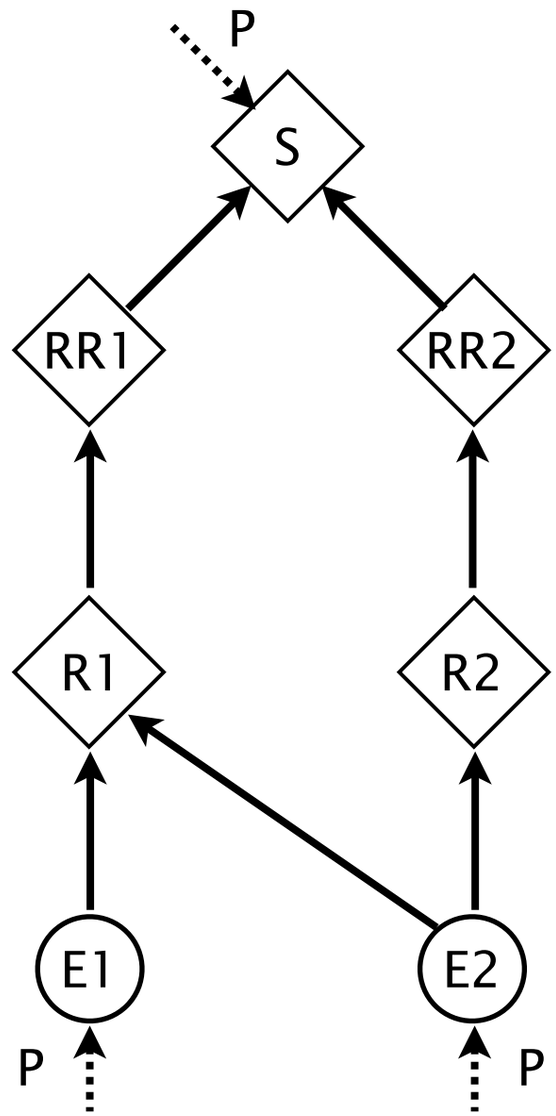

Final BGP configuration

A seamless migration ordering might not always exist

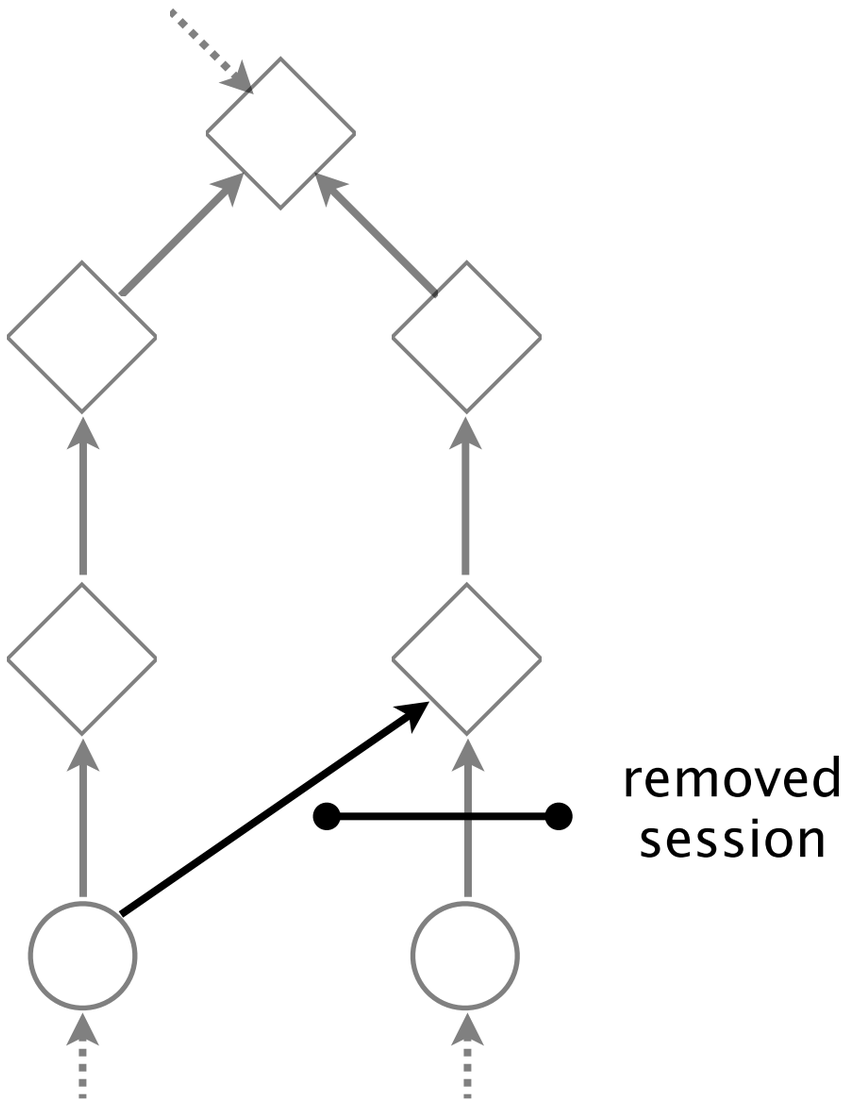

Initial BGP configuration

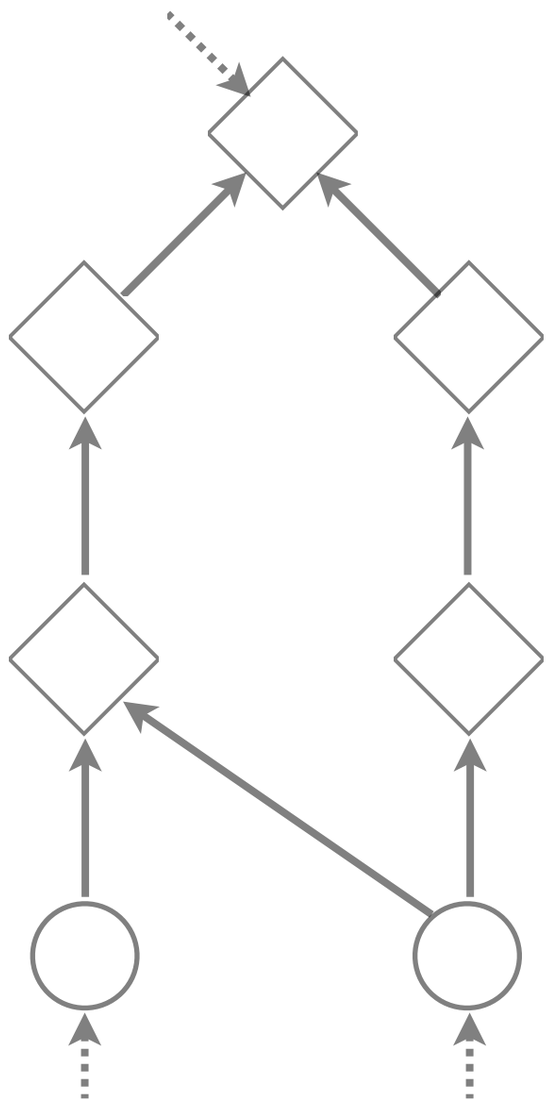

A seamless migration ordering might not always exist

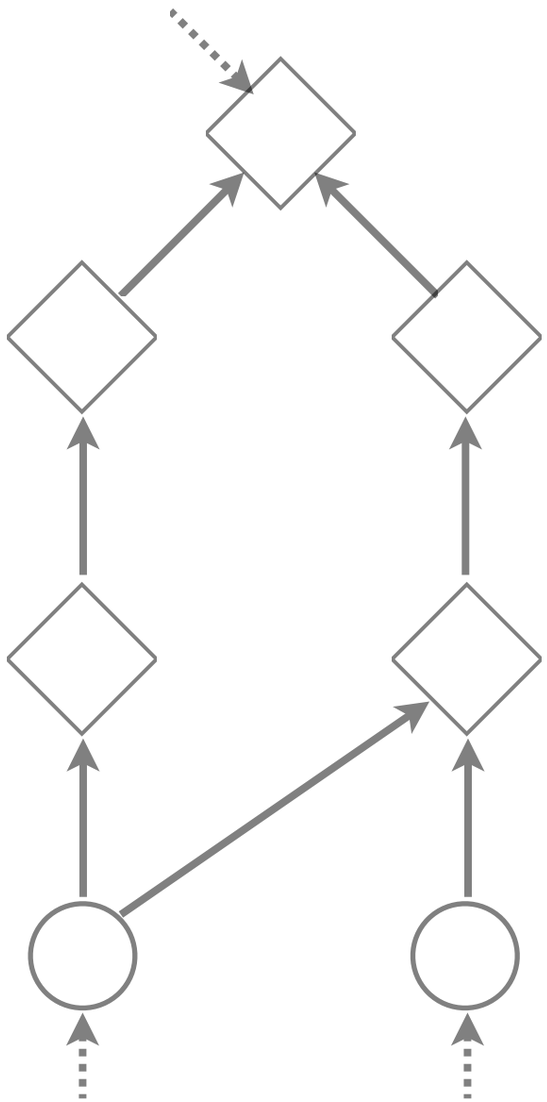

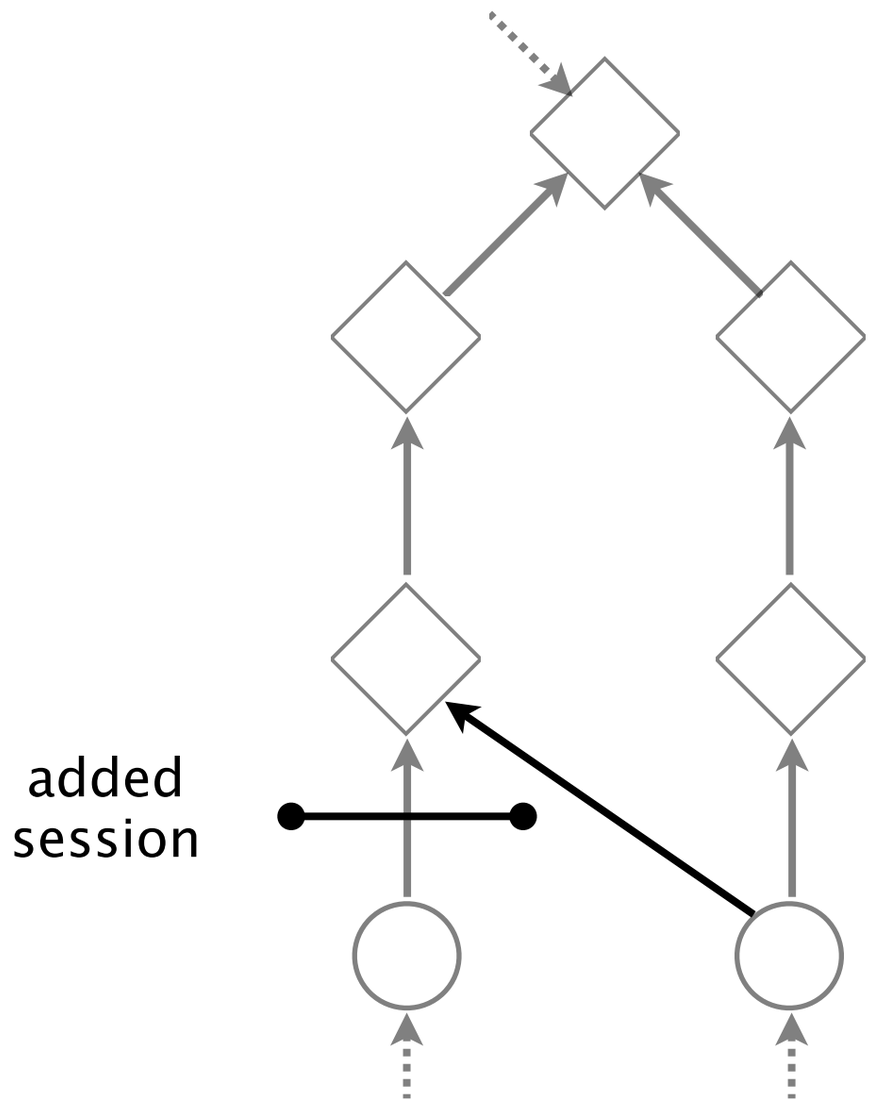

Final BGP configuration

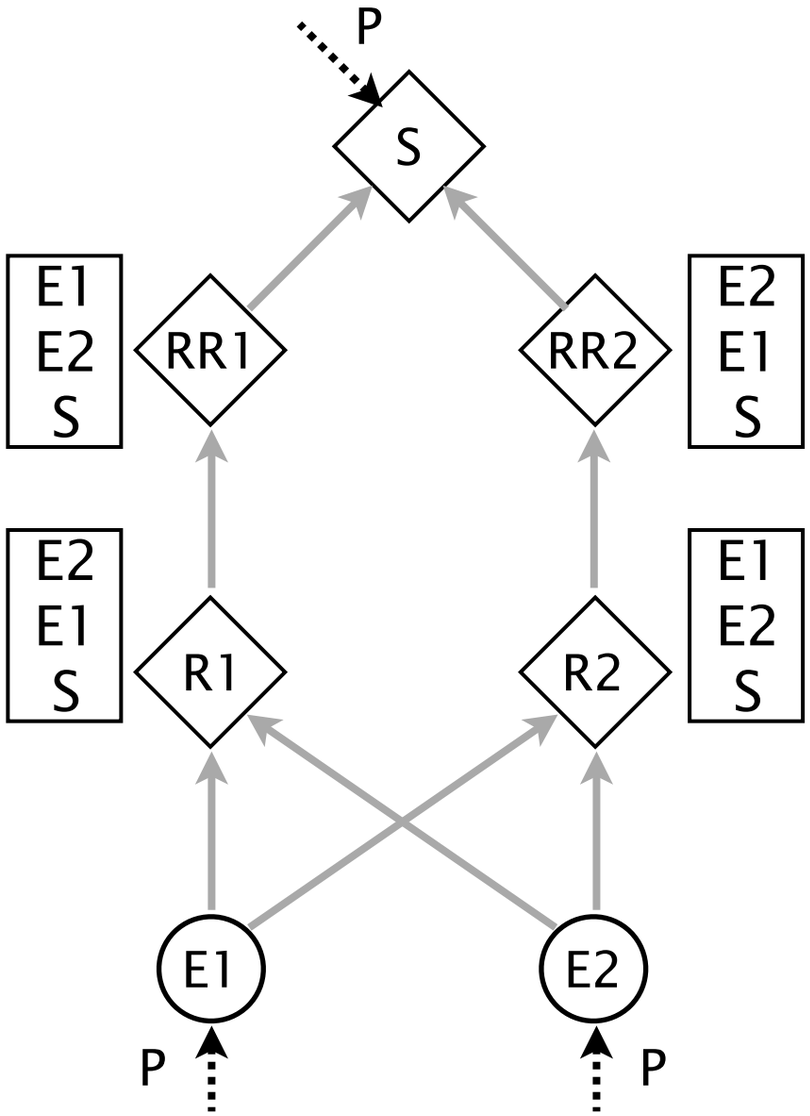

Path preferences

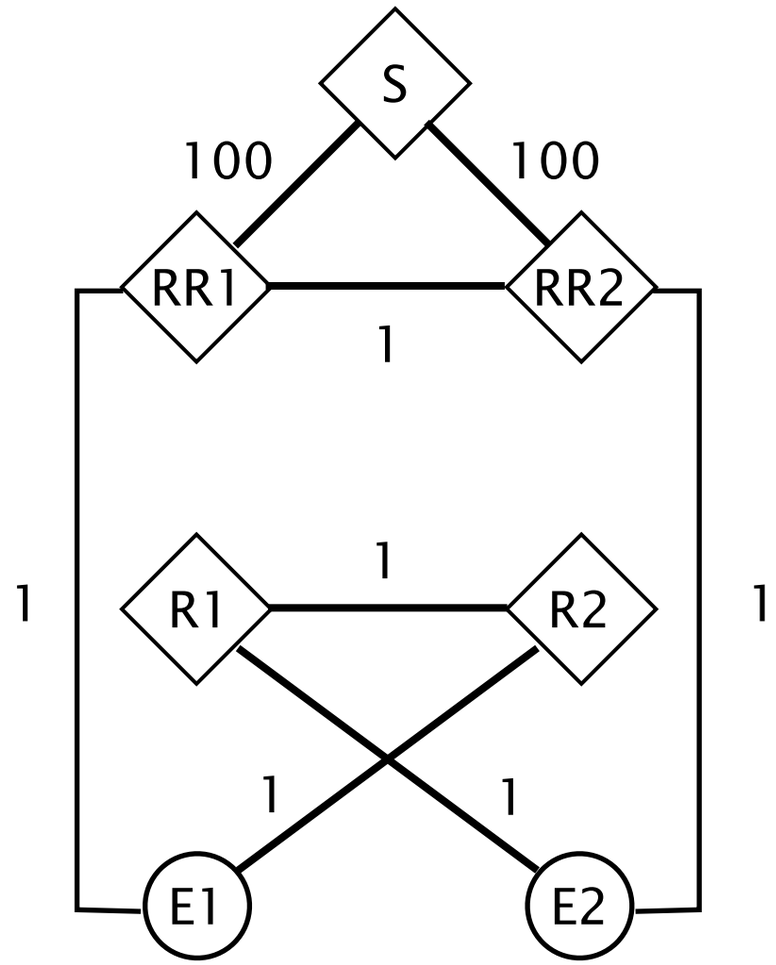

IGP configuration

The initial configuration is anomaly-free

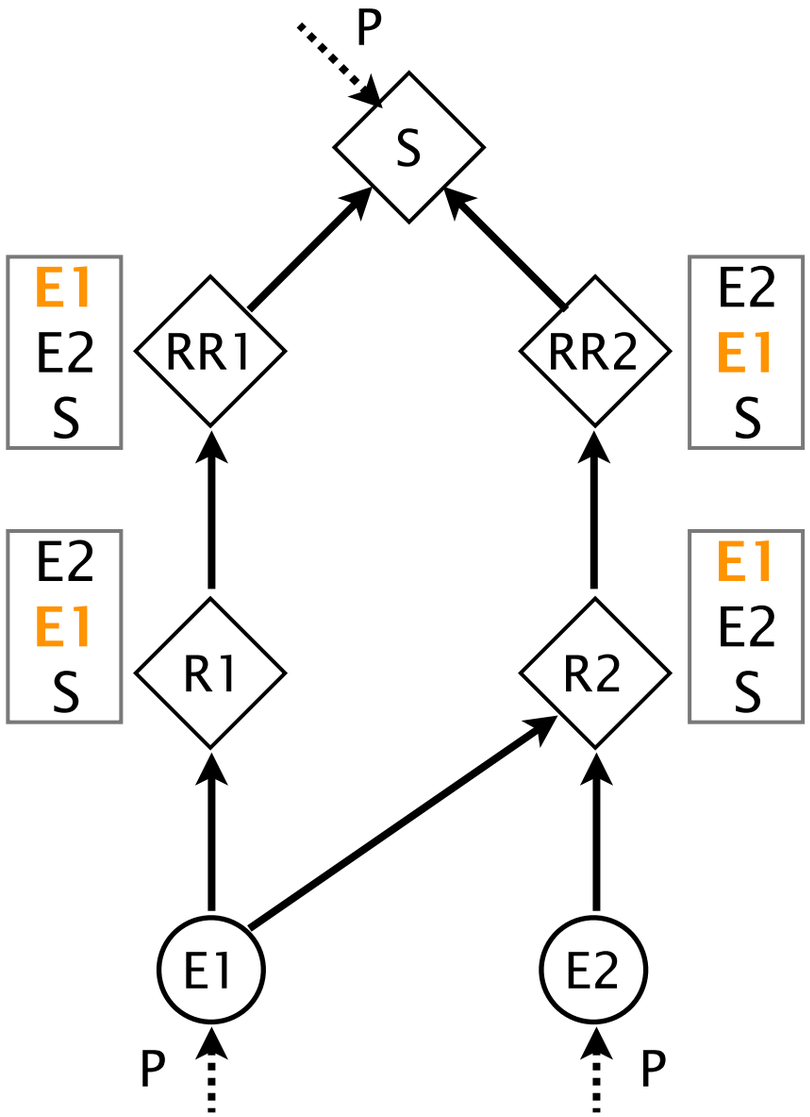

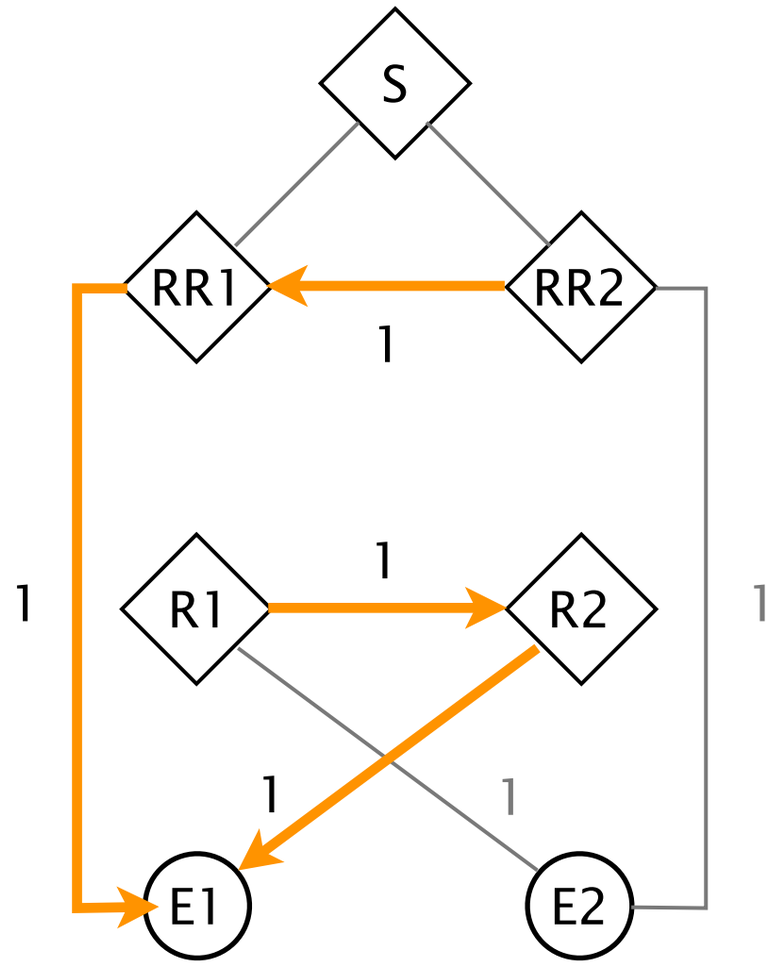

The final configuration is anomaly-free

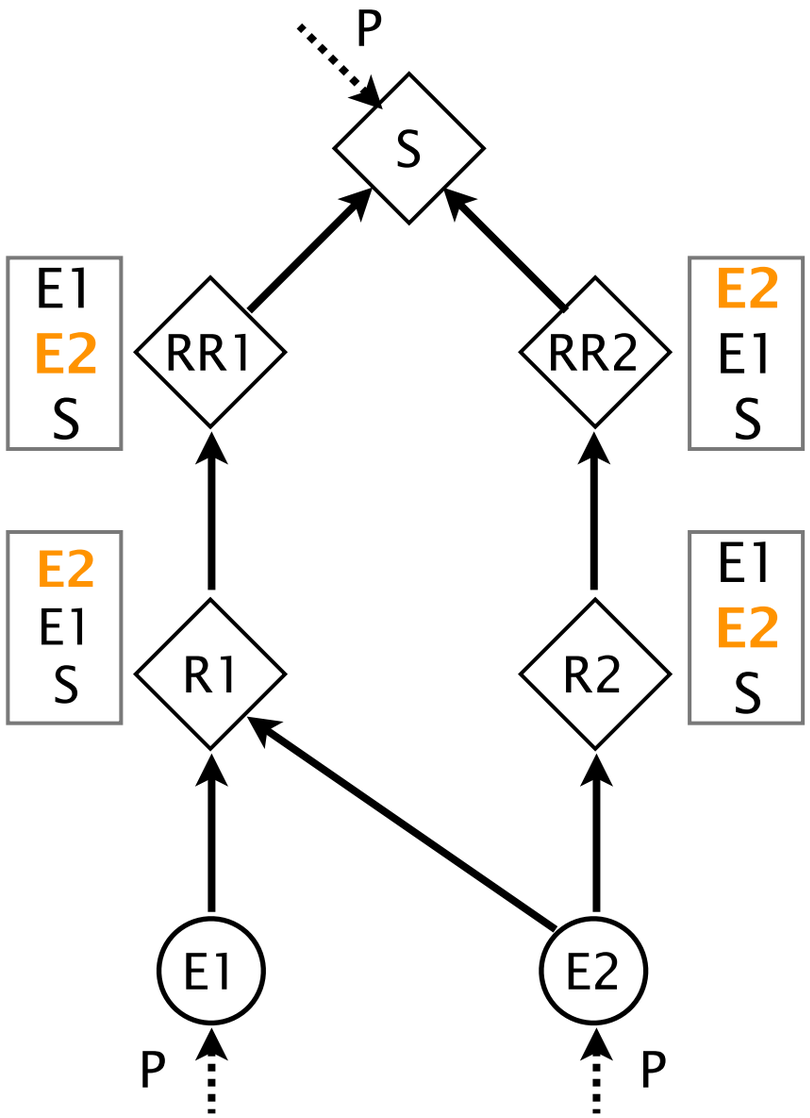

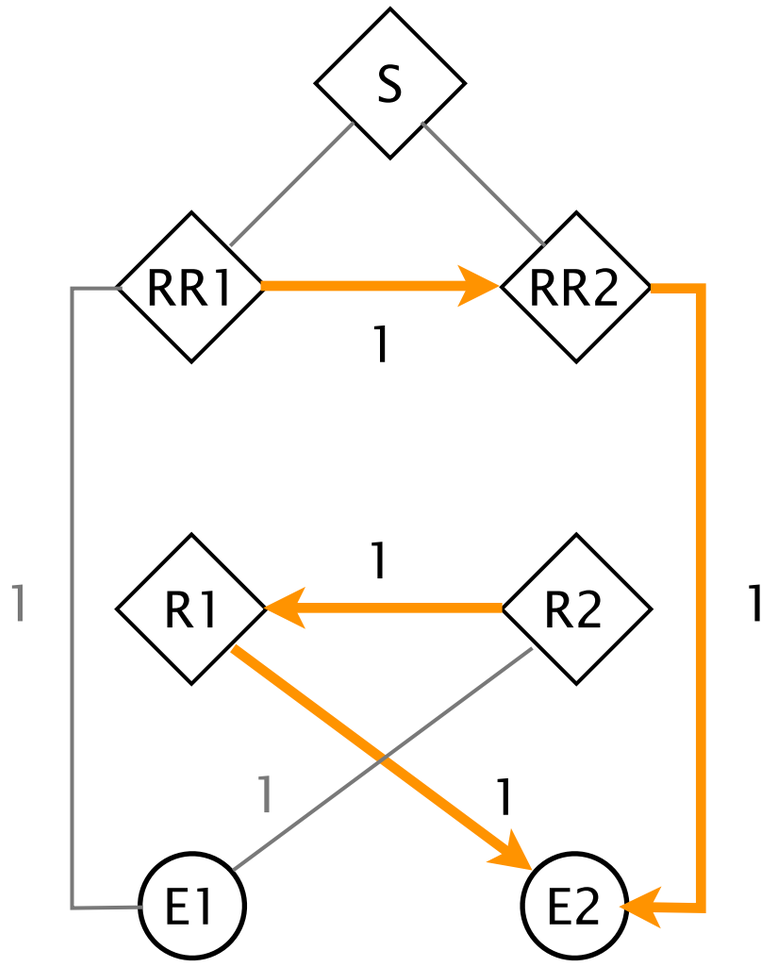

Let’s add the final session before removing the initial one

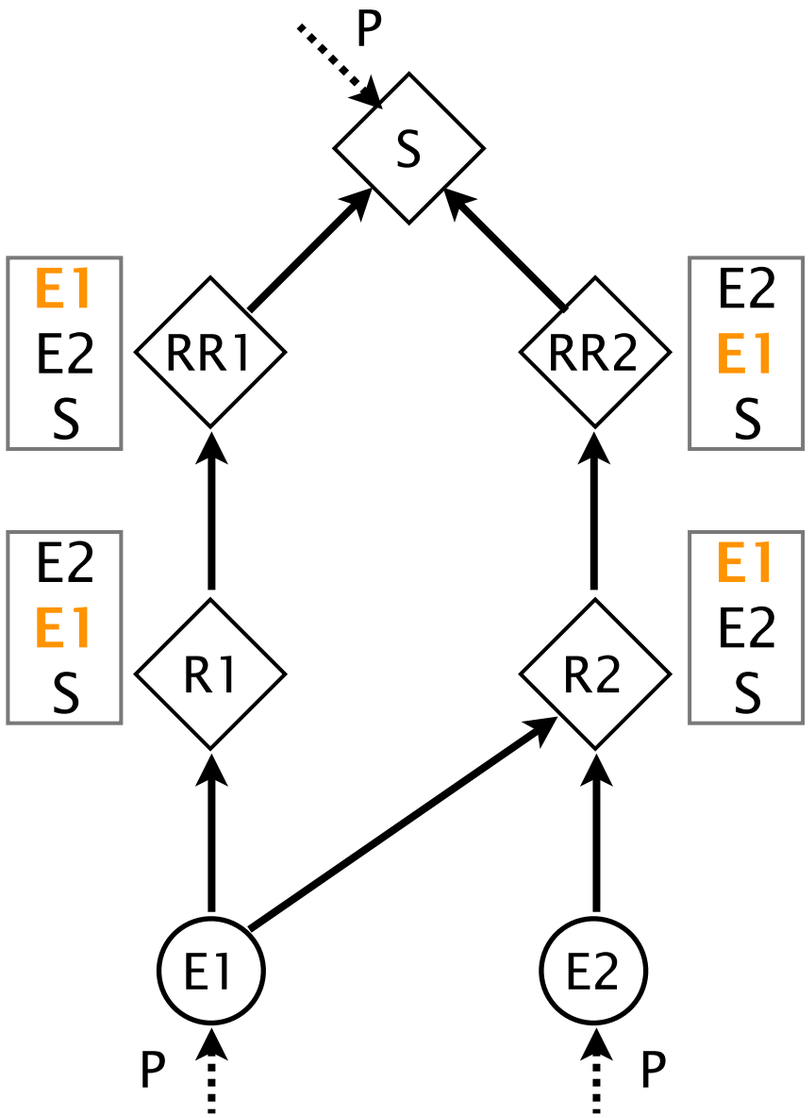

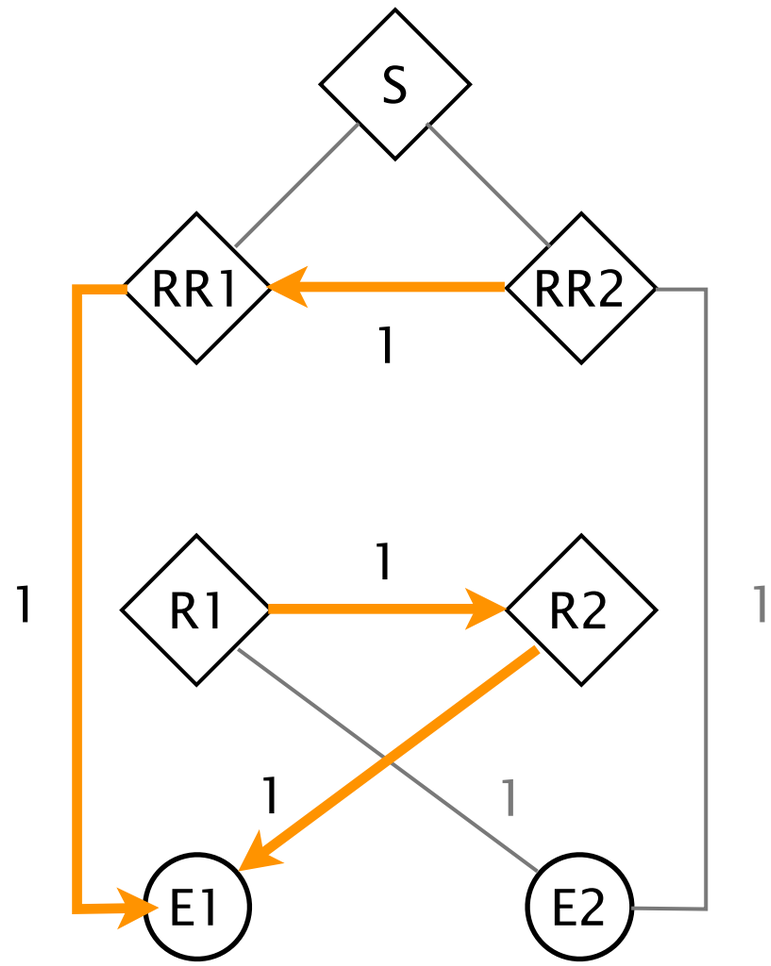

Let’s add the final session before removing the initial one

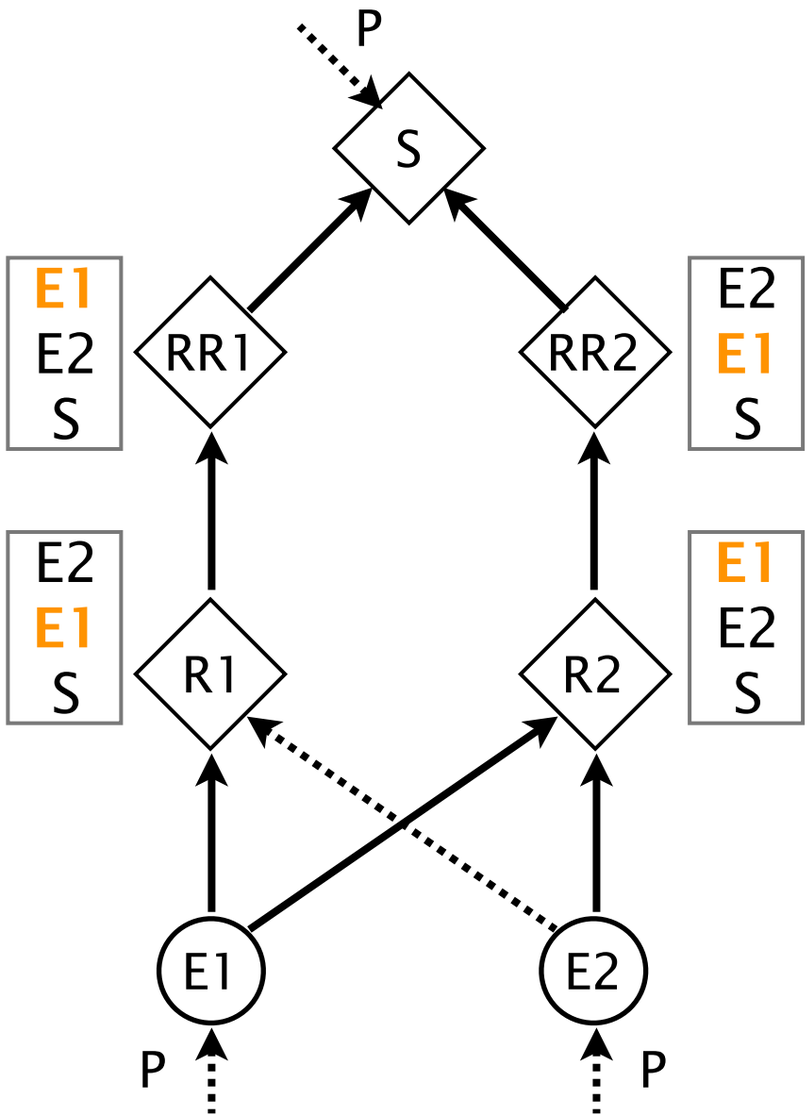

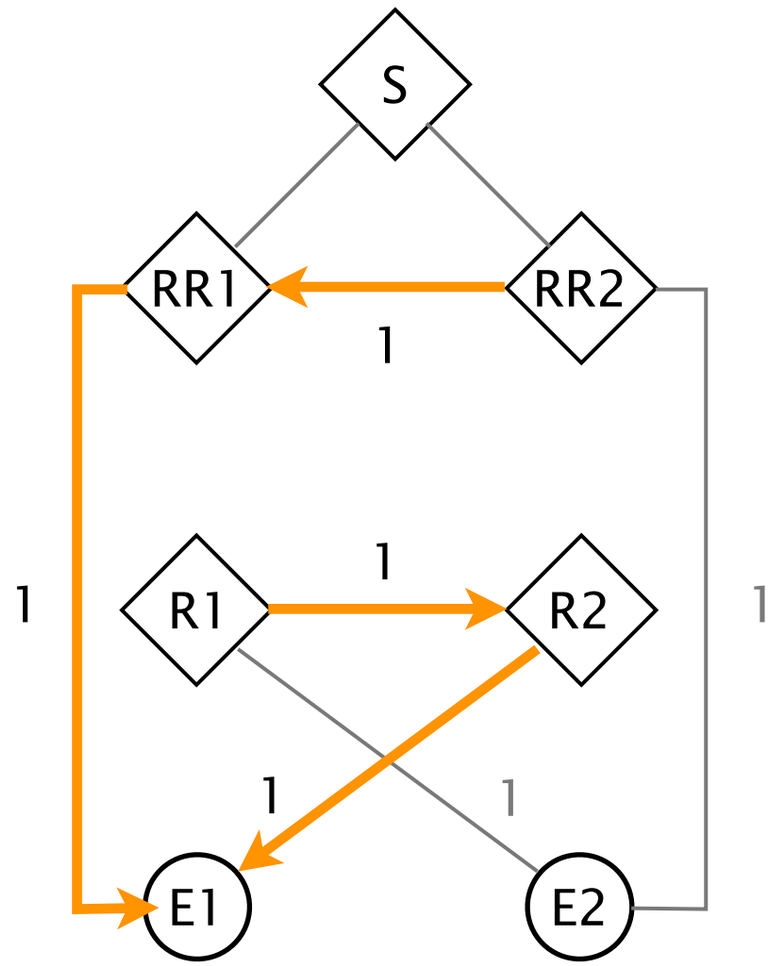

Let’s add the final session before removing the initial one

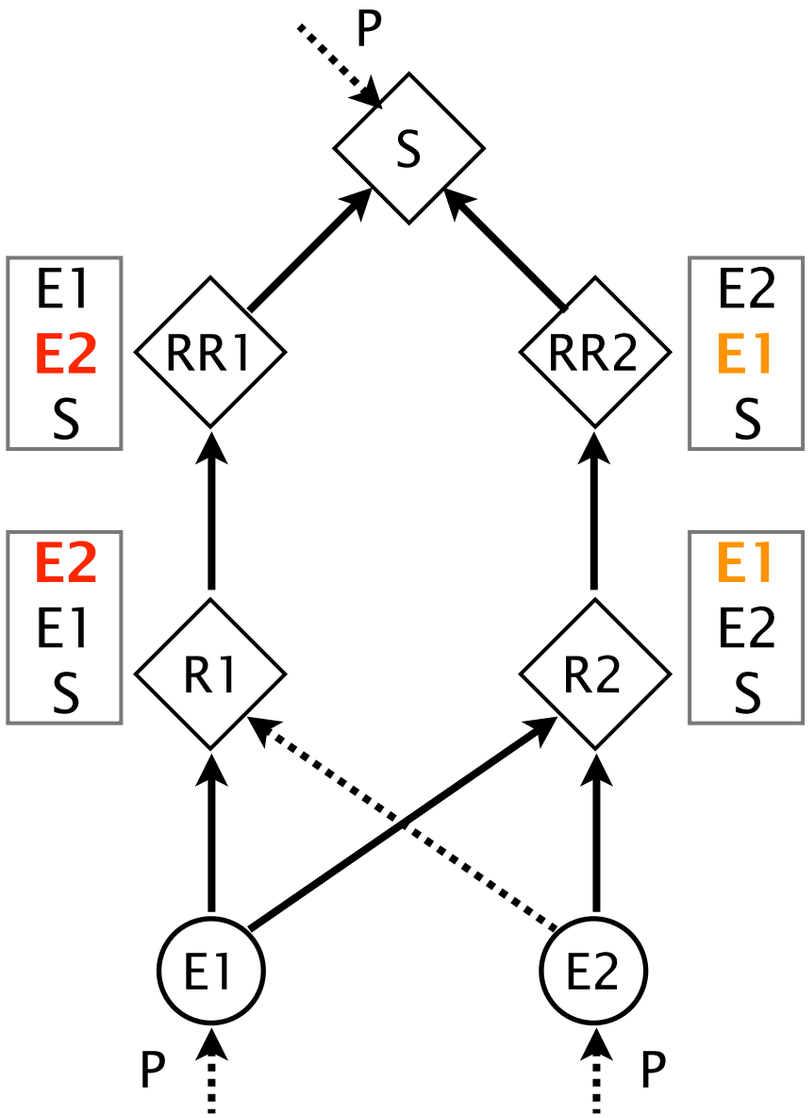

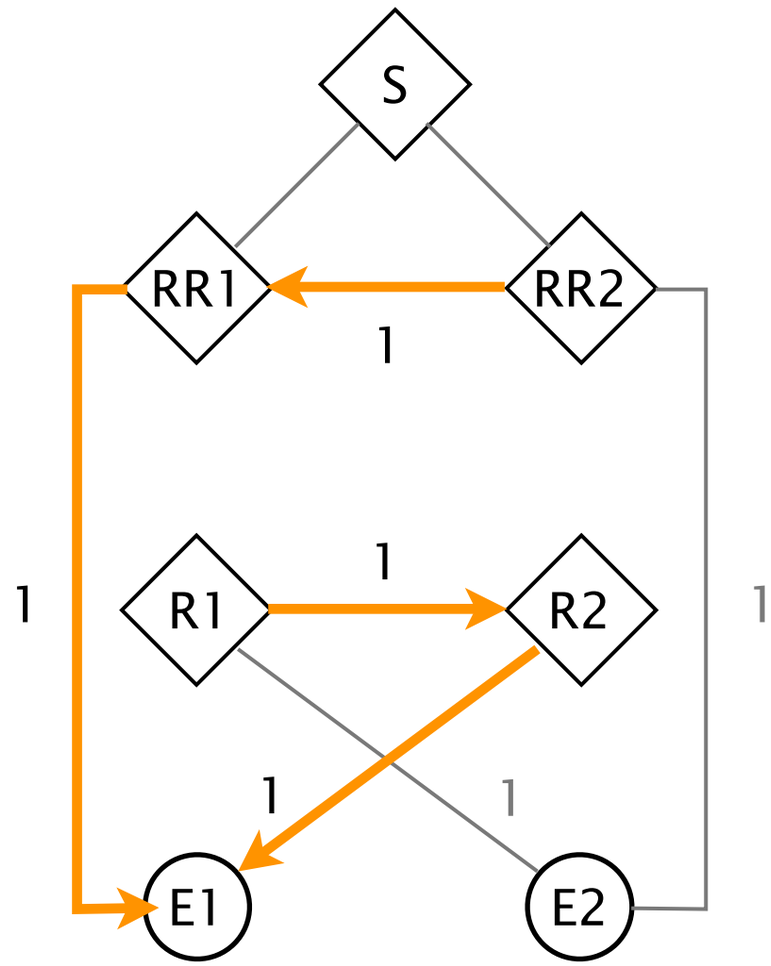

Let’s add the final session before removing the initial one

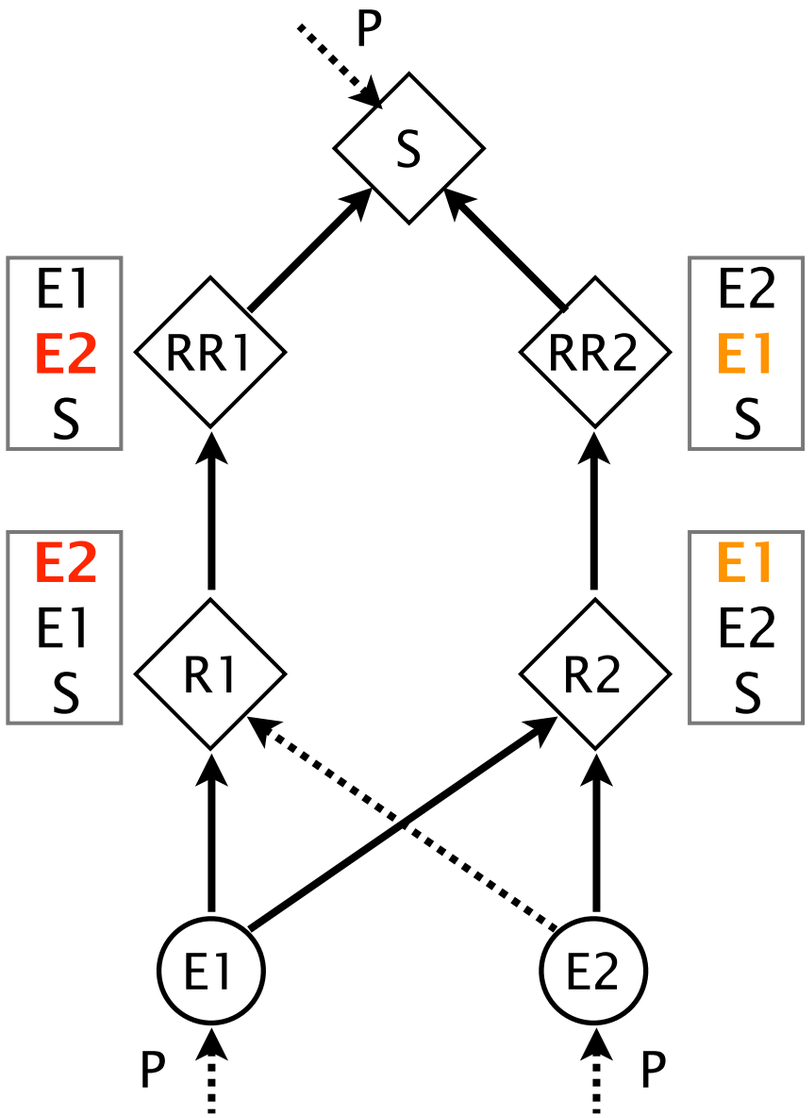

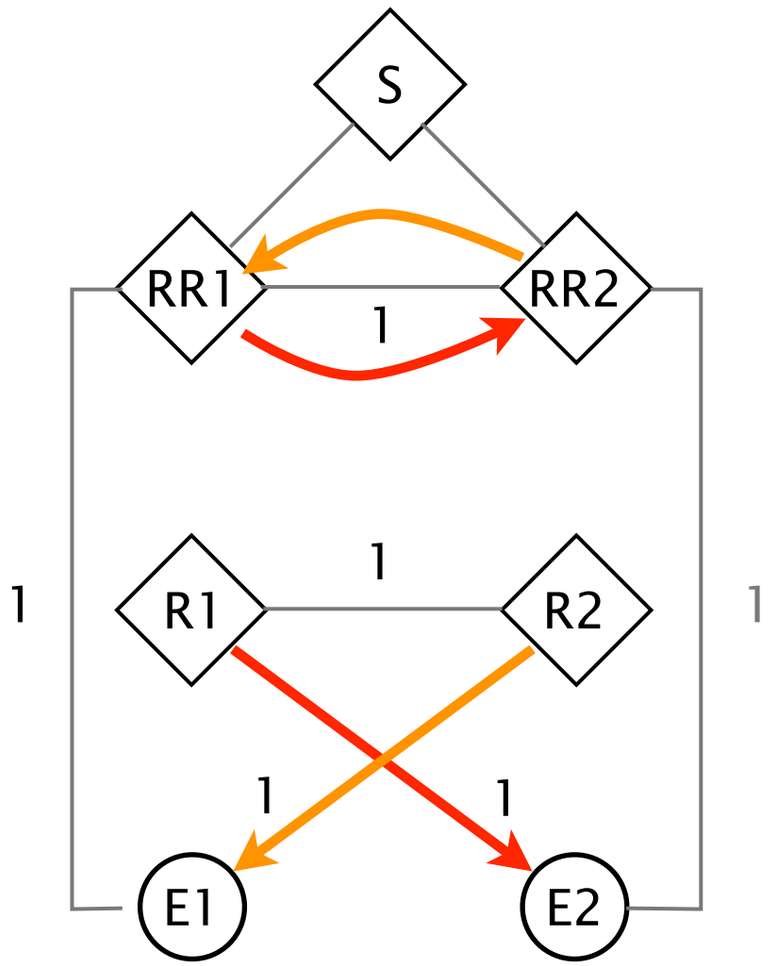

A forwarding loop is created

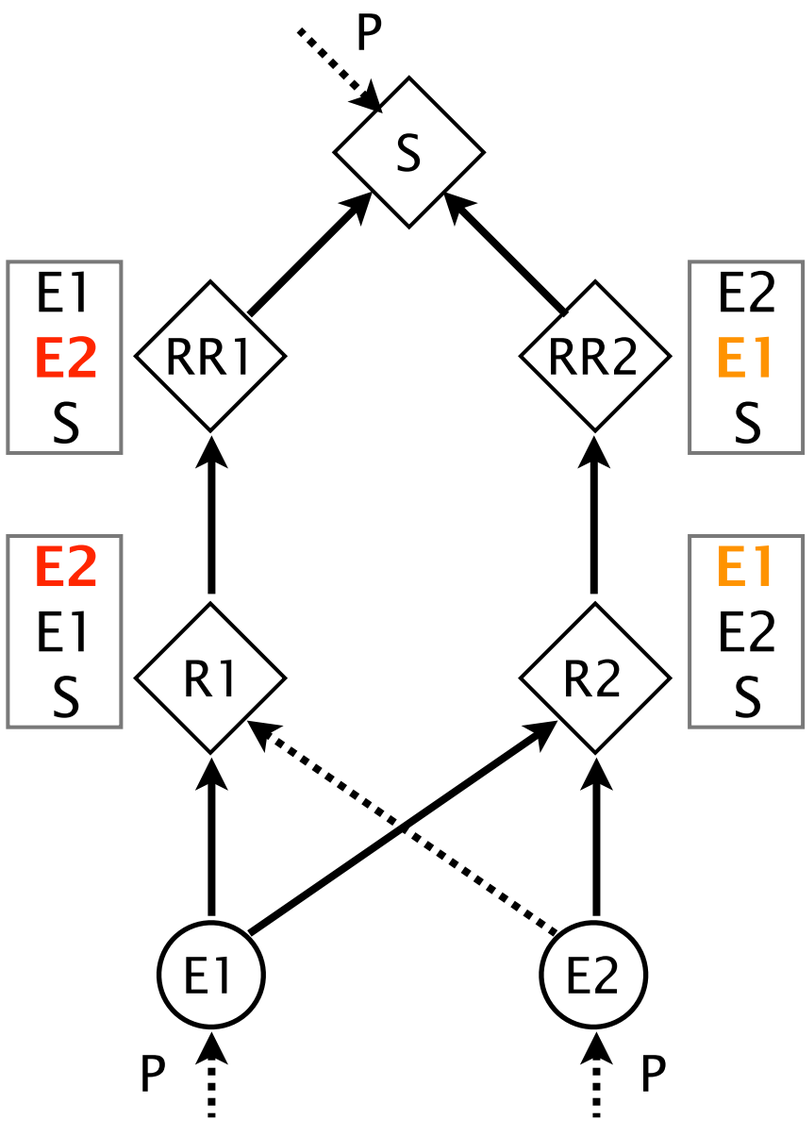

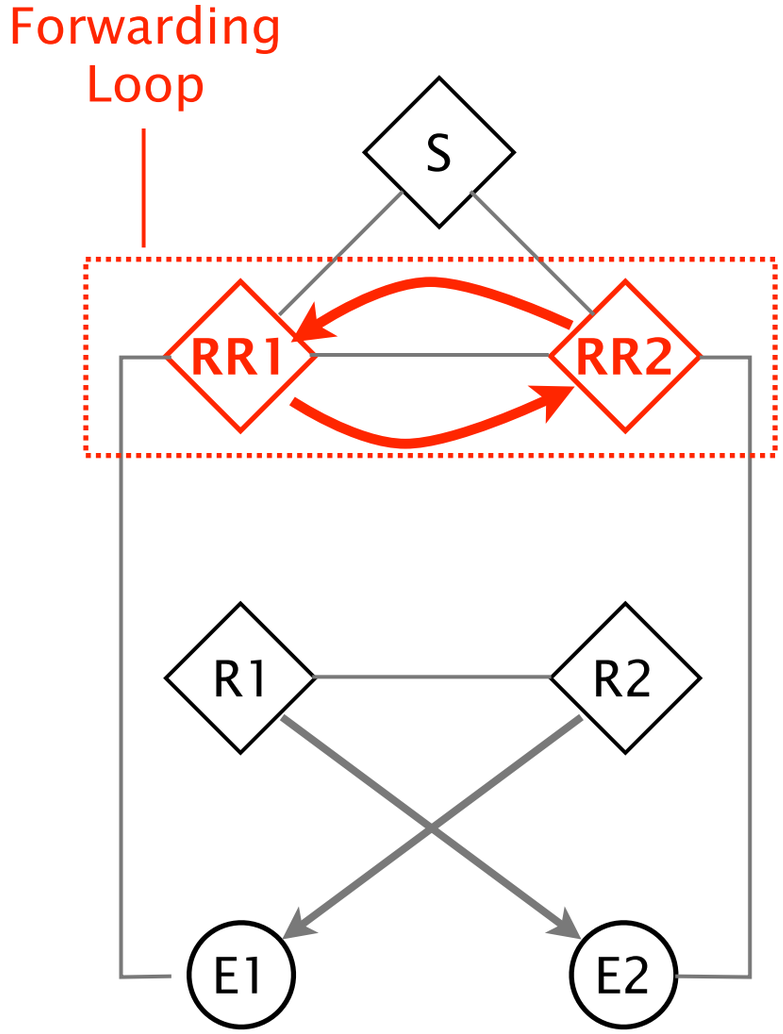

Let’s remove the initial session before adding the final one

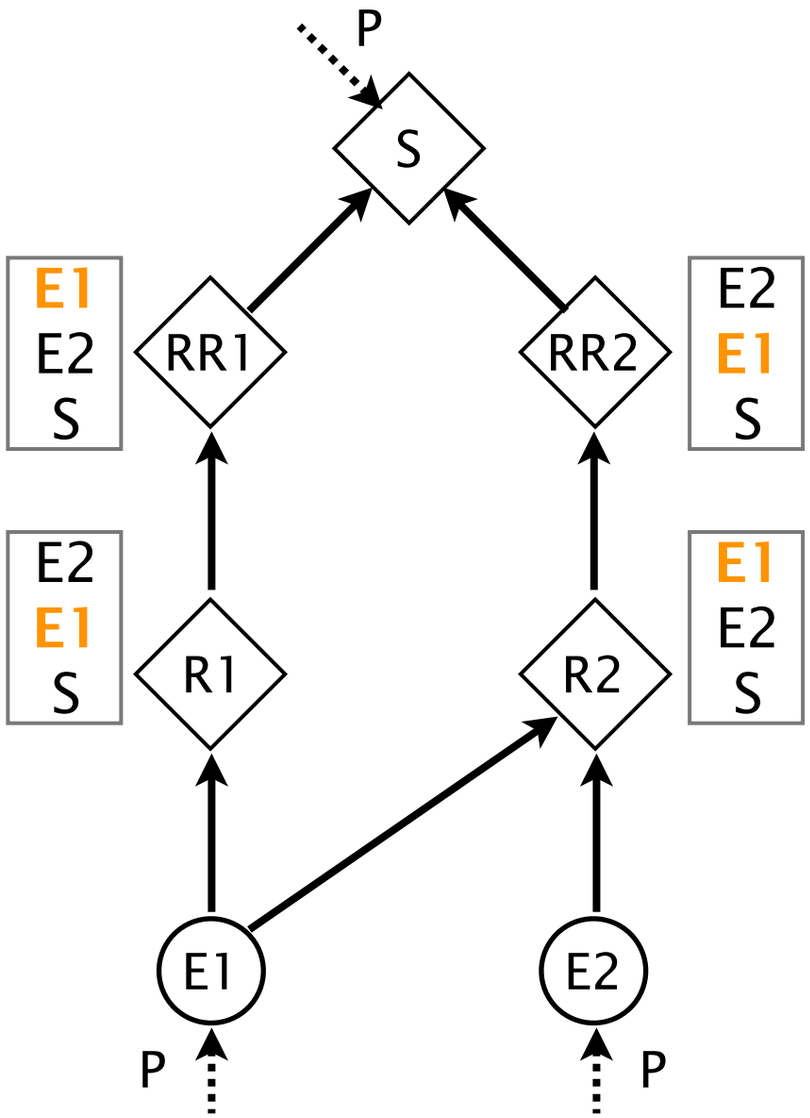

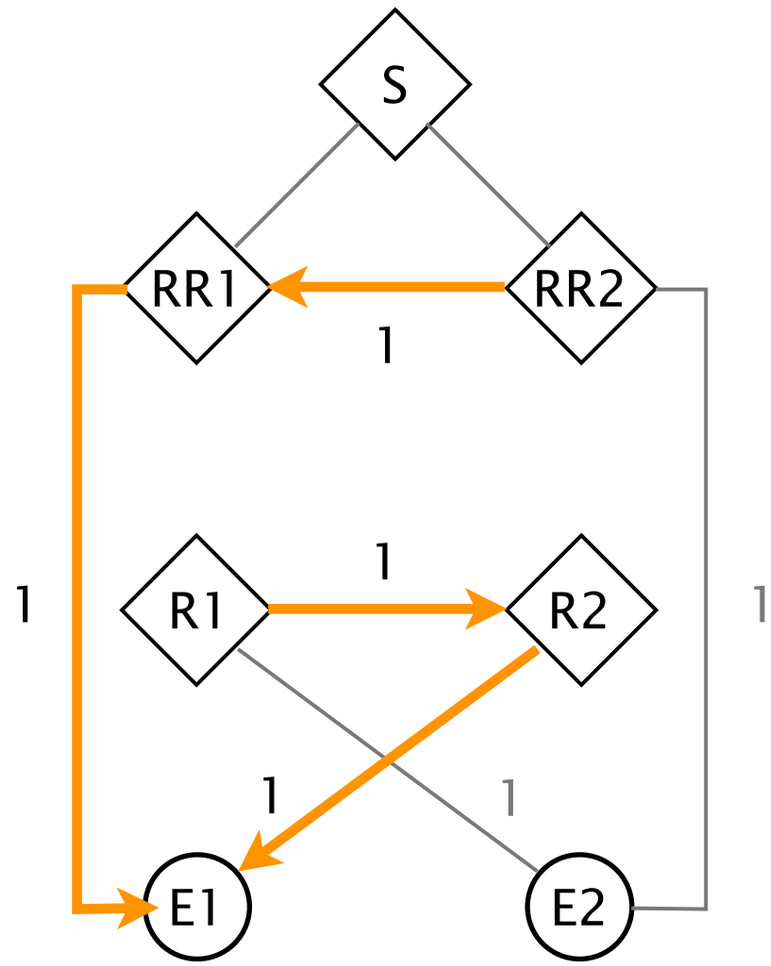

Let’s remove the initial session before adding the final one

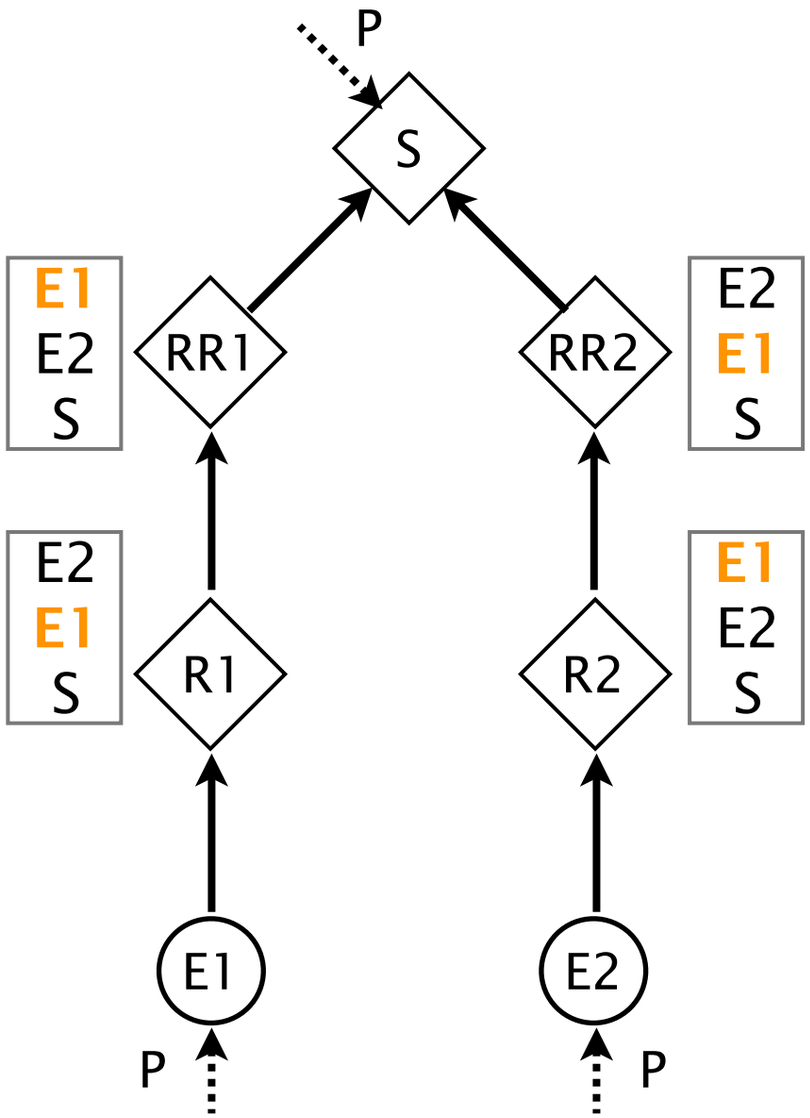

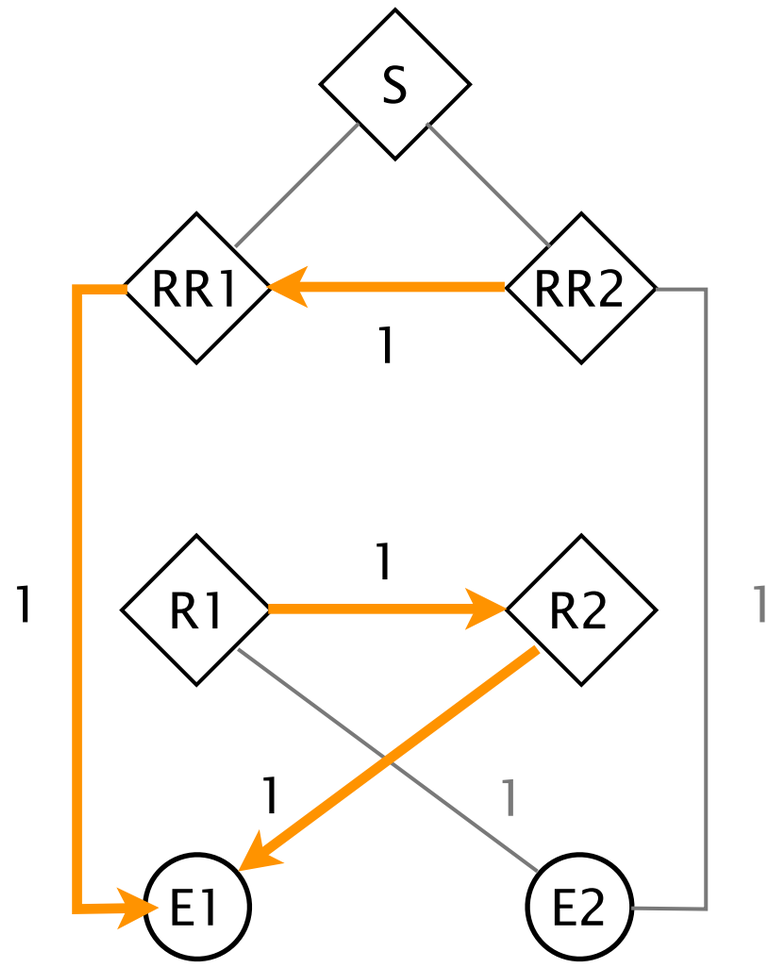

Let’s remove the initial session before adding the final one

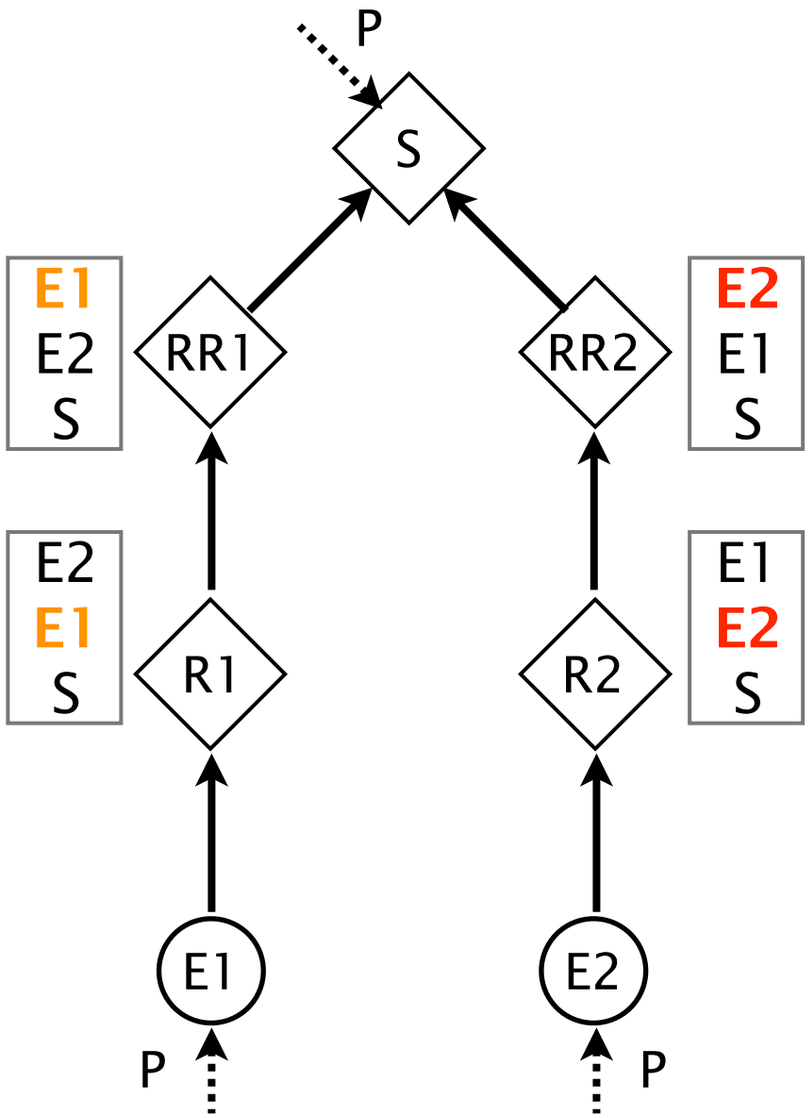

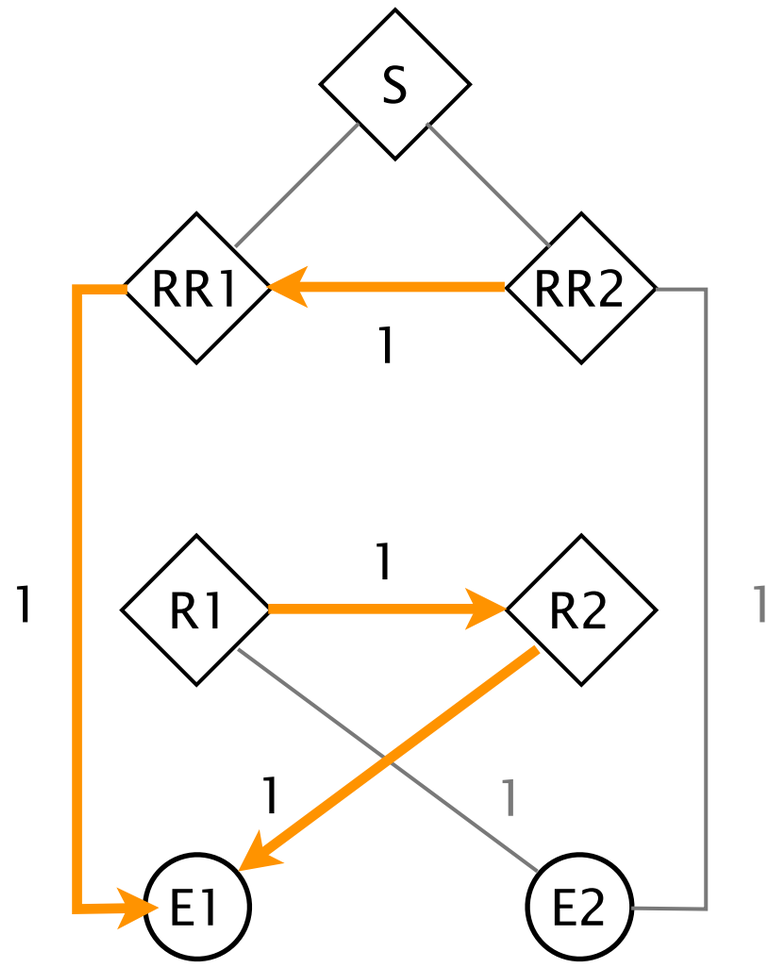

Let’s remove the initial session before adding the final one

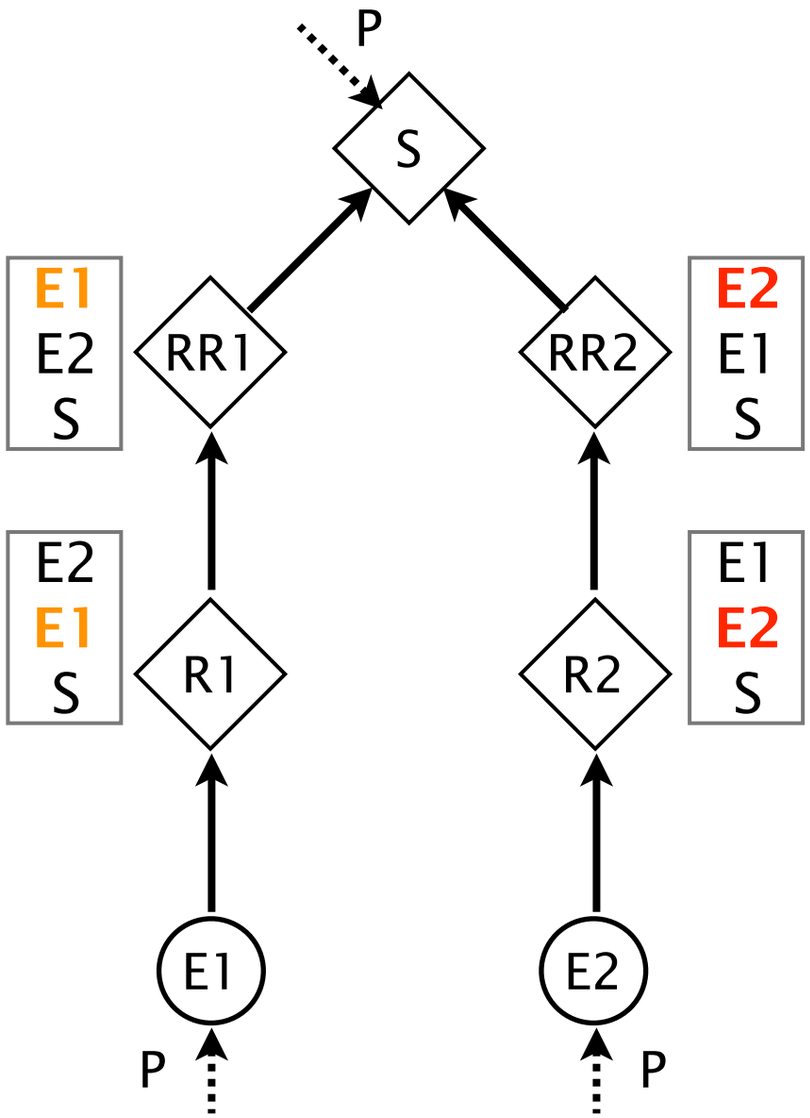

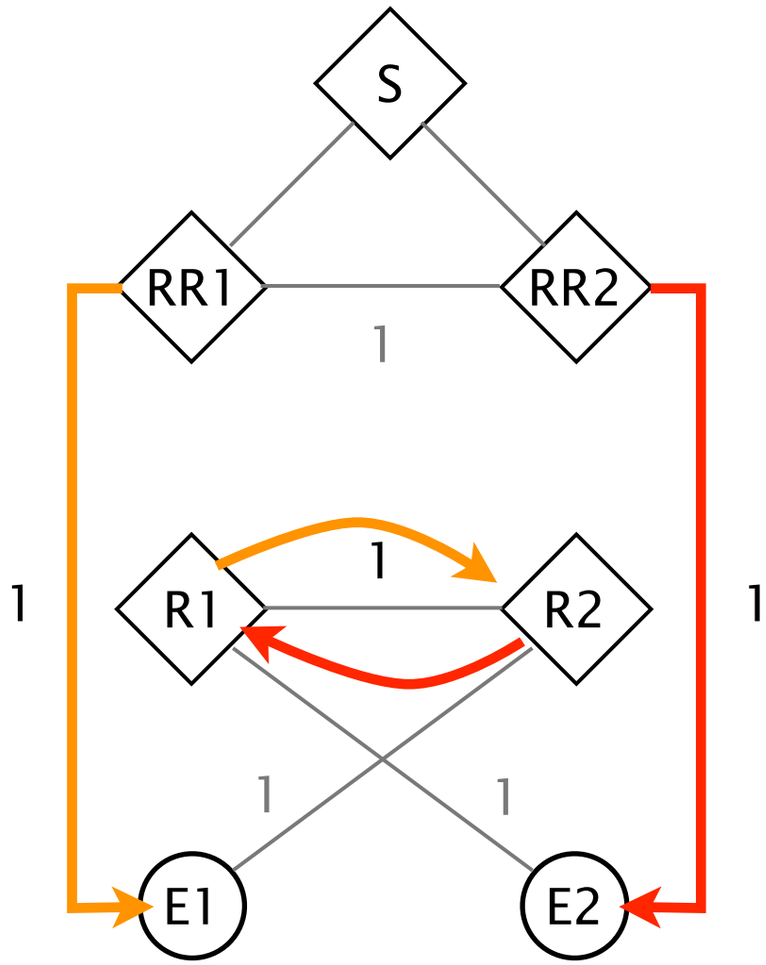

A forwarding loop is created as well

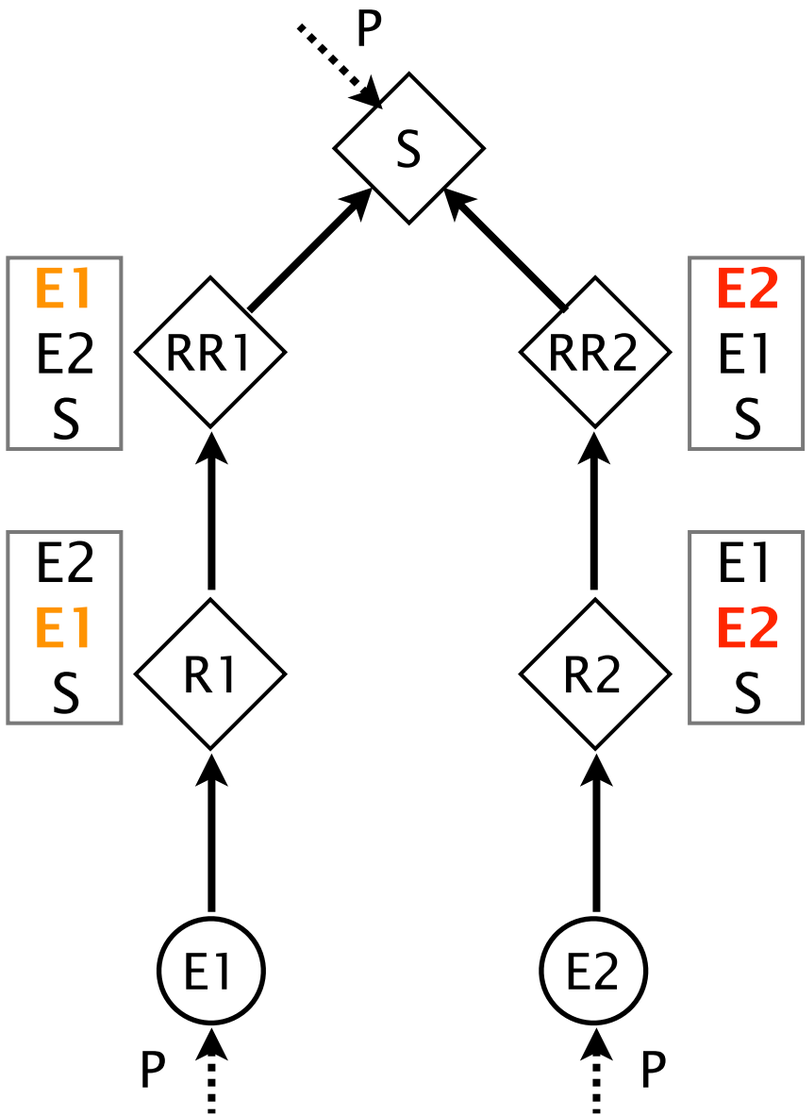

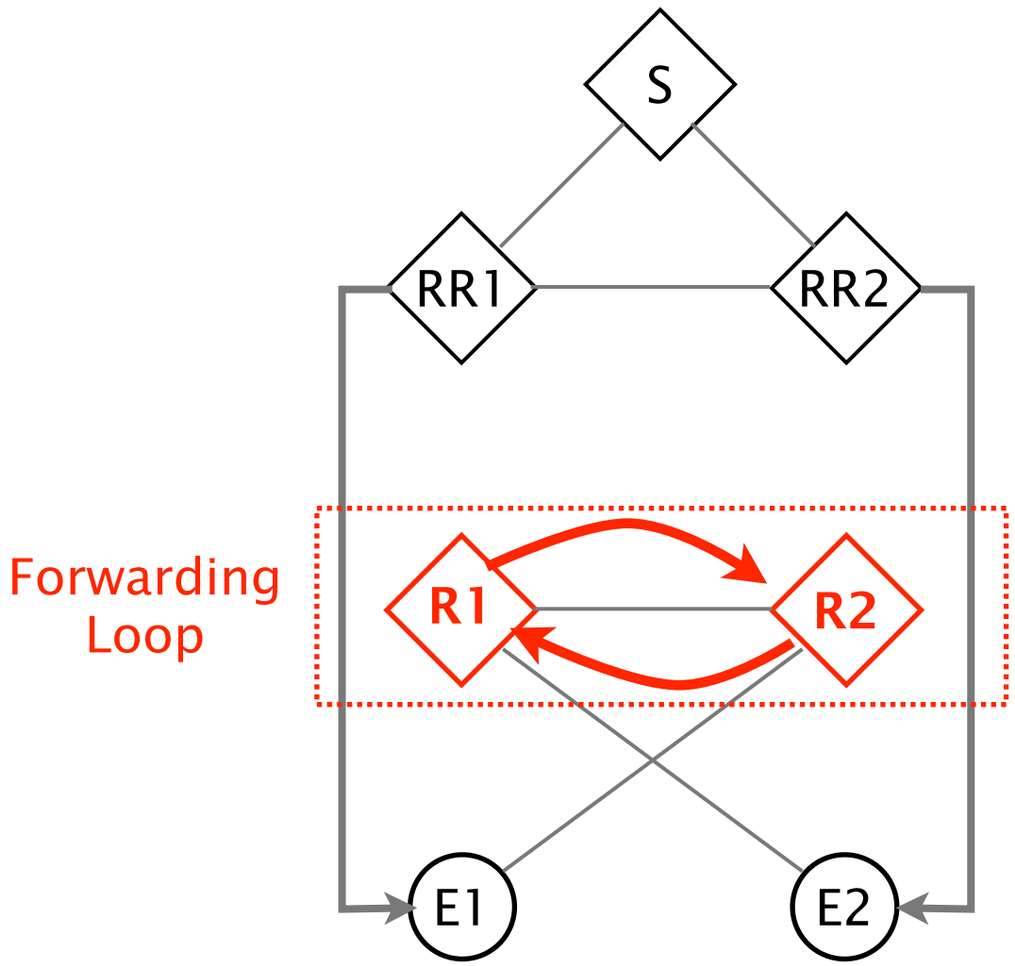

Find a sequence of configuration changes

Does it always exist ? **No.**

Find a sequence of configuration changes

Is it easy to compute ?

Finding a seamless migration ordering is computationally hard

Deciding if an ordering free from signaling anomalies exists is NP-hard

reduction in polynomial time from 3-SAT

Finding a seamless migration ordering is computationally hard

Deciding if an ordering free from signaling anomalies exists is NP-hard

reduction in polynomial time from 3-SAT

The same reduction applies for

- dissemination anomalies
- forwarding anomalies
- iBGP or eBGP reconfigurations

Find a sequence of configuration changes

Is it easy to compute ? **No.**

- **An algorithmic approach is not viable**

Improving network agility with seamless BGP reconfigurations

BGP background

Finding an ordering

Reconfiguration framework

Why is BGP reconfiguration so complex ?

Local reconfiguration can have global impact

in an unpredictable manner

Why is BGP reconfiguration so complex ?

Local reconfiguration can have global impact

in an unpredictable manner

To avoid that, we could run each configuration
in an independent routing plane

Similar to

- IGP reconfiguration [Vanbever, SIGCOMM11]
- Shadow configuration [Alimi, SIGCOMM08]

The reconfiguration framework leverages Ships-In-The-Night (SITN) migration for BGP

SITNs migrations consists in

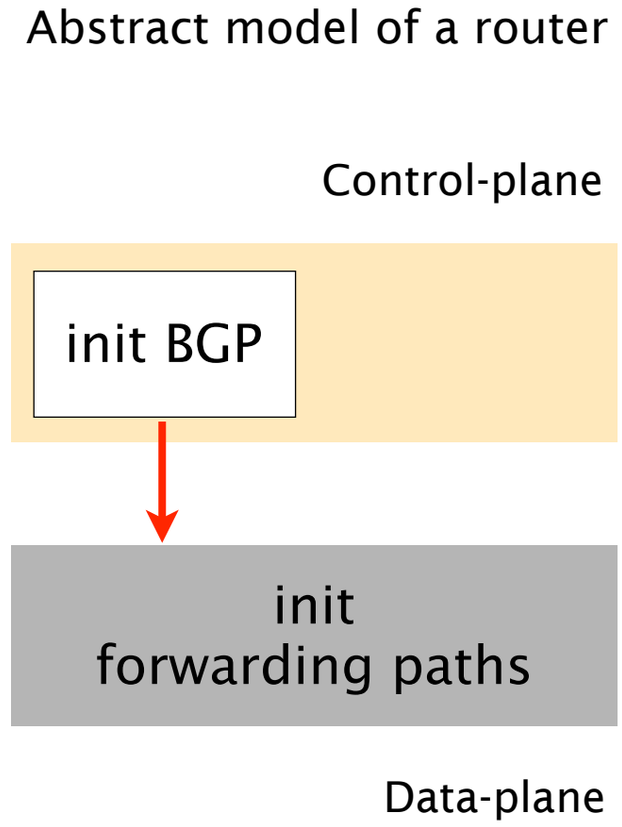

The reconfiguration framework leverages Ships-In-The-Night (SITN) migration for BGP

SITNs migrations consists in

1. running multiple BGP routing planes

Abstract model of a router

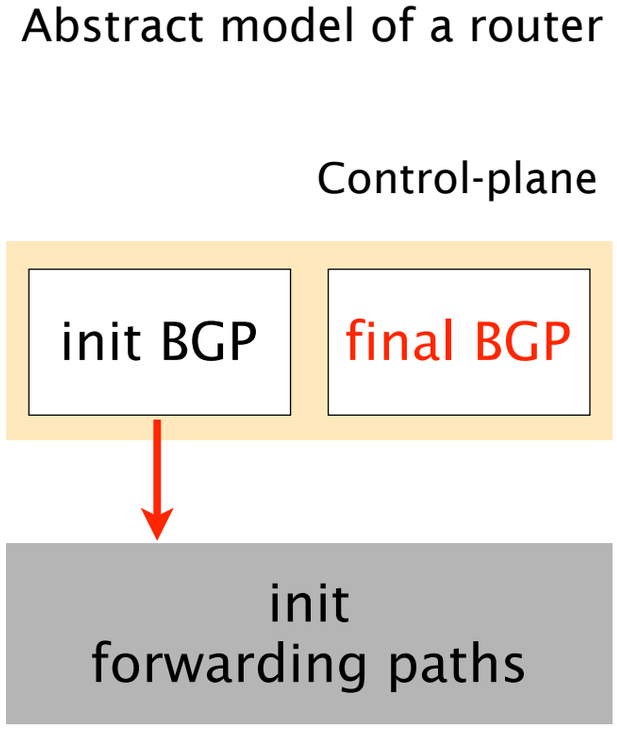

The reconfiguration framework leverages Ships-In-The-Night (SITN) migration for BGP

SITNs migrations consists in

2 waiting for each plane to converge

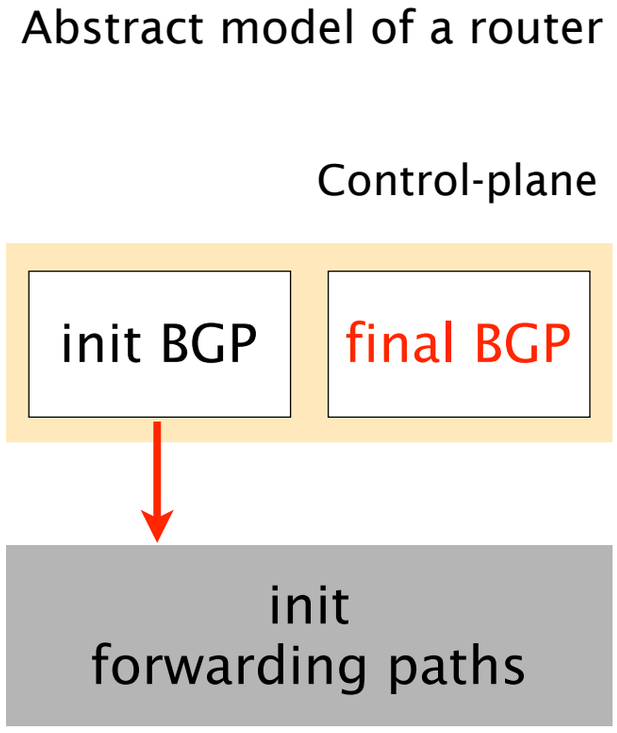

The reconfiguration framework leverages Ships-In-The-Night (SITN) migration for BGP

SITNs migrations consists in

3 modifying the process responsible for forwarding

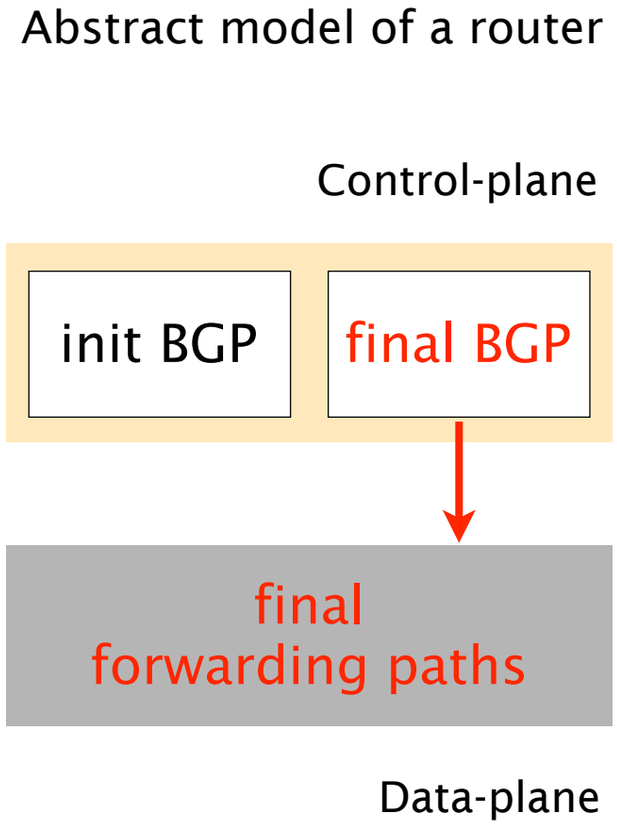

The reconfiguration framework leverages Ships-In-The-Night (SITN) migration for BGP

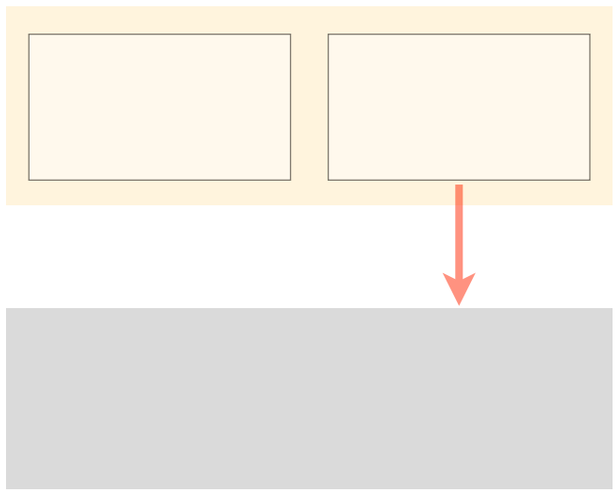

BGP SITN can be deployed on today's routers

Let’s reconfigure a network from an iBGP full-mesh ...

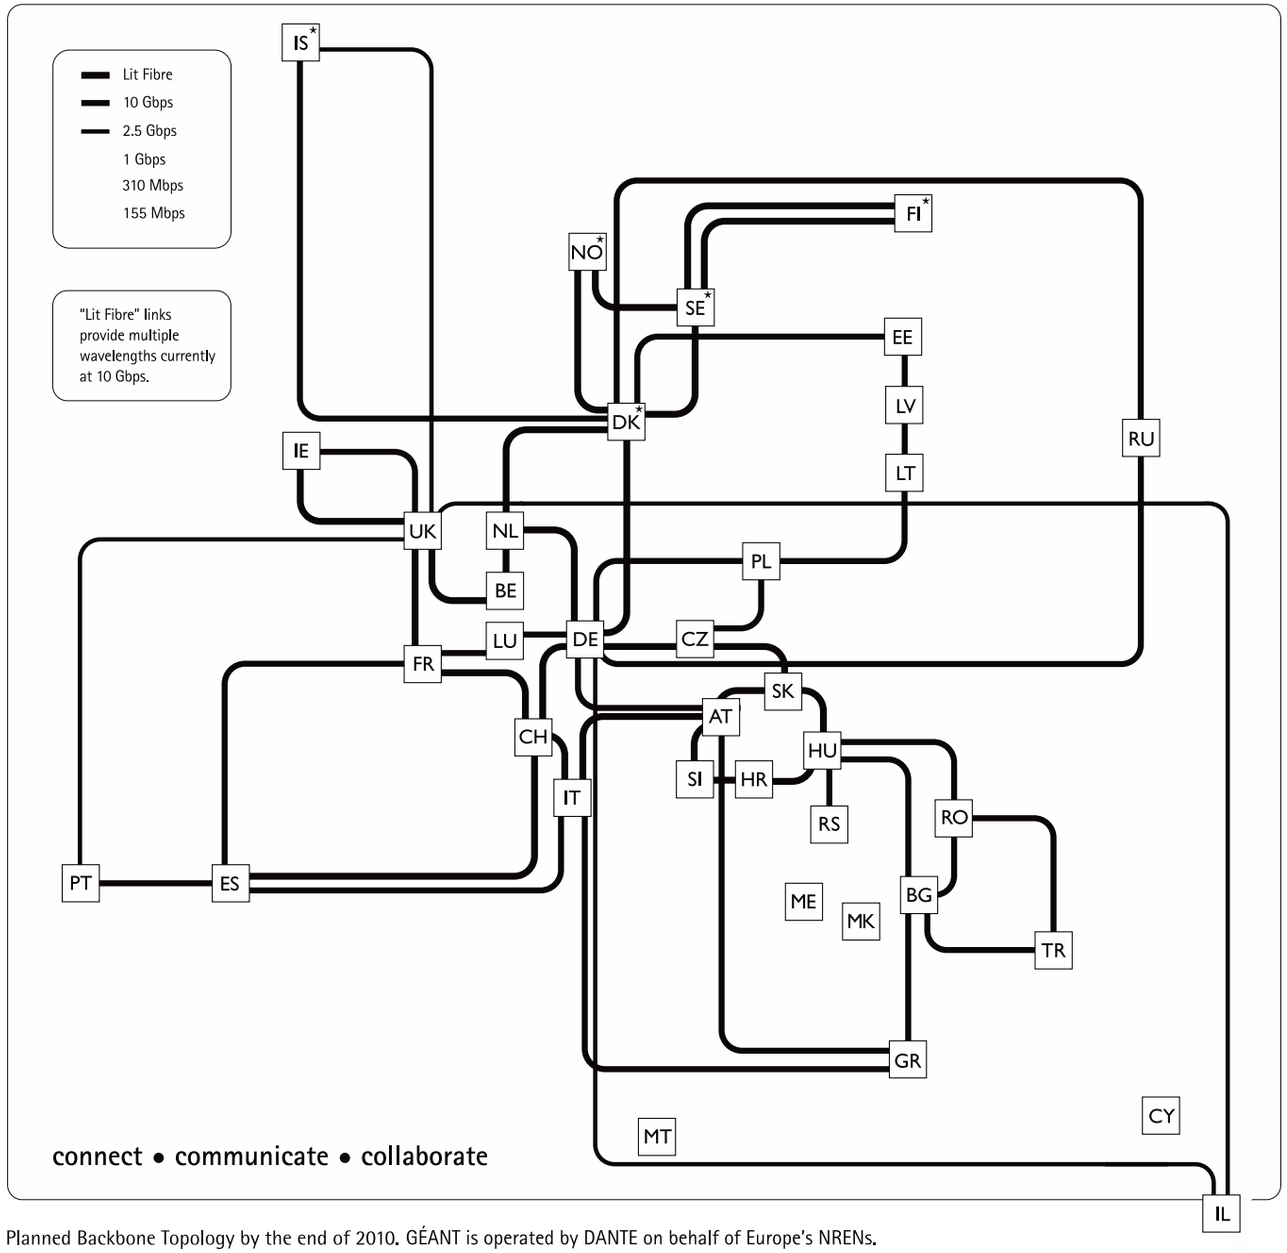

Planned Backbone Topology by the end of 2010. GÉANT is operated by DANTE on behalf of Europe's NRENs.

GEANT

European research network

36 routers (virtualized)

53 links

Let's reconfigure a network from an iBGP full-mesh to an iBGP hierarchy

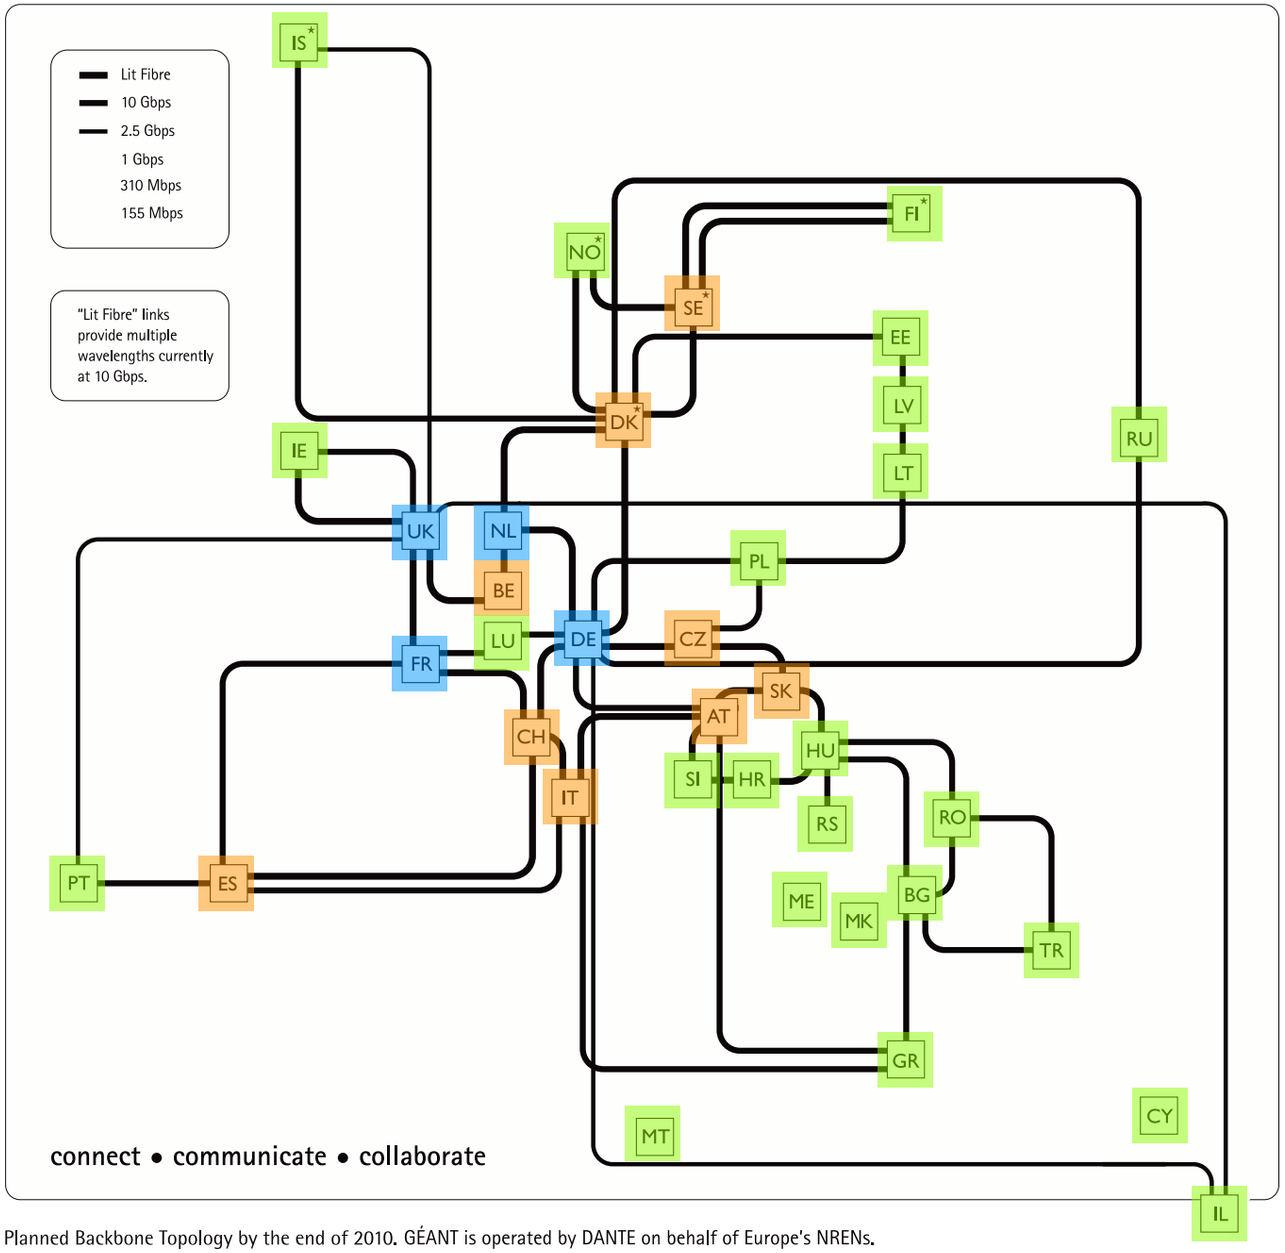

Planned Backbone Topology by the end of 2010. GÉANT is operated by DANTE on behalf of Europe's NRENs.

GEANT

European research network

36 routers (virtualized)

53 links

iBGP hierarchy

- Top
- Middle
- Bottom

Following best practices, traffic was lost for *30% of the process*

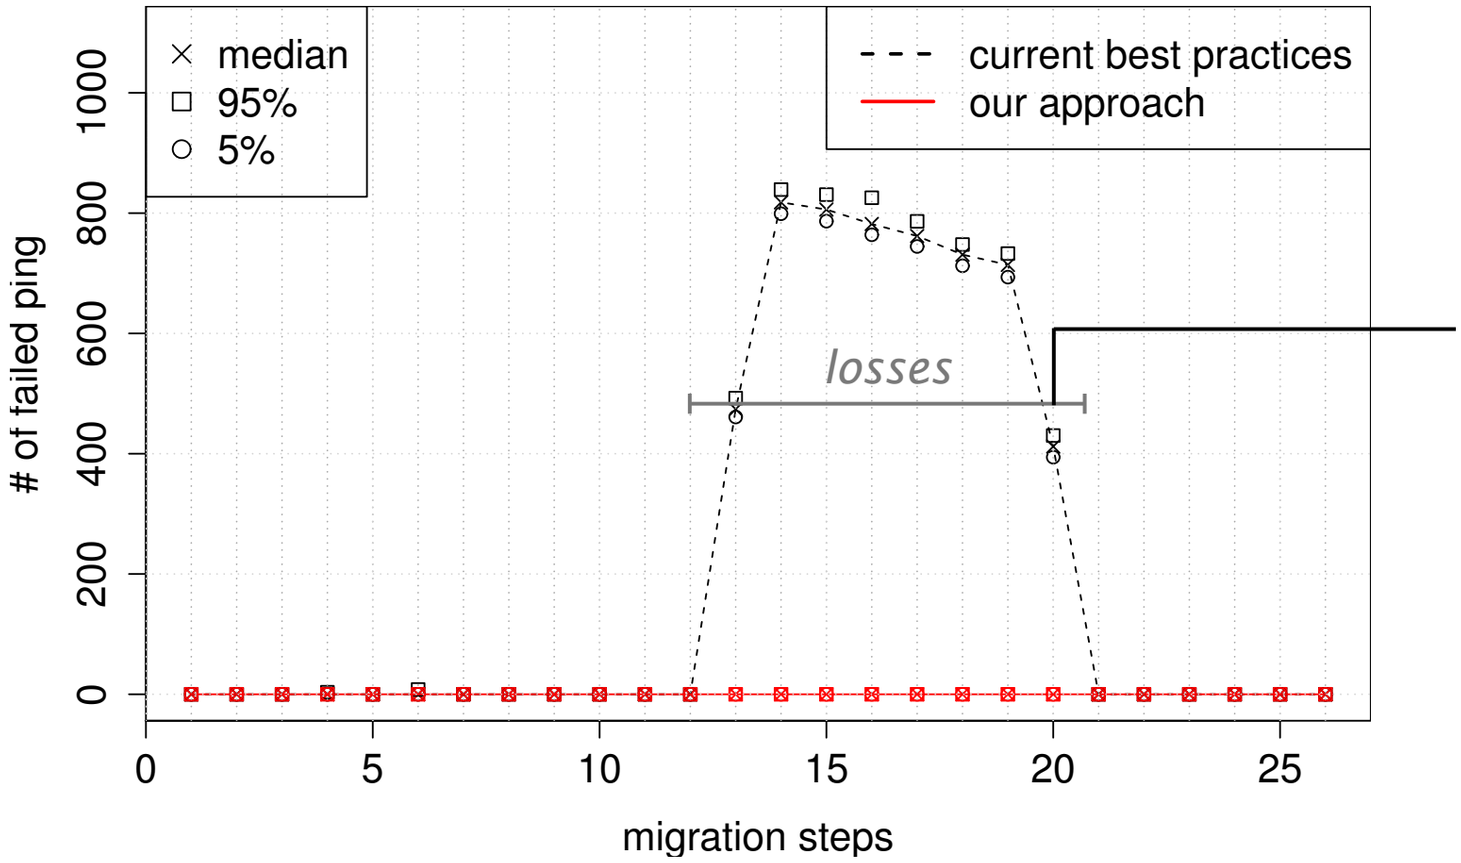

losses from 7 routers

60% of GEANT routing table is impacted !

Average results (30 repetitions) computed on 120+ pings per step from every router to 16 summary prefixes

Following our approach, lossless reconfiguration was achieved

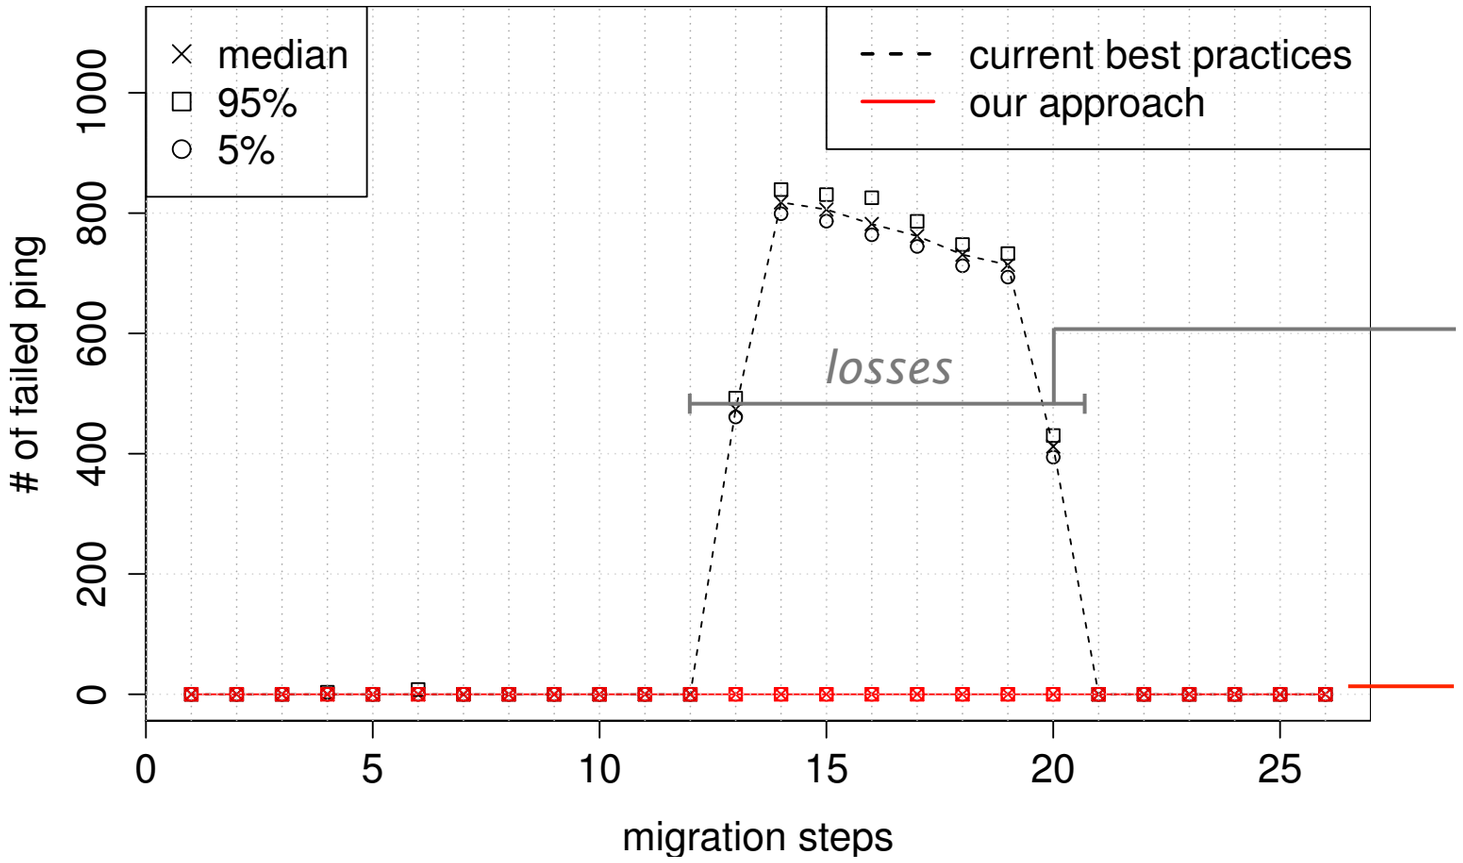

losses from 7 routers

60% of GEANT
routing table is impacted !

No loss occurred
with our approach

Average results (30 repetitions) computed on 120+ pings per step from every router to 16 summary prefixes

Improving network agility with seamless BGP reconfigurations

BGP background

Finding an ordering

Reconfiguration framework

Improving network agility with seamless BGP reconfigurations

BGP background

Finding an ordering

Reconfiguration framework

Contributions

1. Study BGP reconfiguration, both practically and theoretically

2. Show that a (seamless) operational ordering
   - might be needed
   - might not exist
   - is computationally hard to find

3. Implement and validate a BGP reconfiguration framework

Future works

1. Ensure seamless migrations when the initial and final configurations meet given properties (which ones ?)
2. Improve the scalability of the framework
3. Deployment of the framework

Improving network agility with seamless BGP reconfigurations

Laurent Vanbever

laurent.vanbever@uclouvain.be

Towards flexible networks

with seamless reconfiguration