Modelling Inter-Domain Routing

Olaf Maennel
University of Adelaide

Wolfgang Mühlbauer
Technical University Munich

Steve Uhlig
Université catholique de Louvain

Anja Feldmann
Technical University Munich

Matthew Roughan
University of Adelaide

Outline

1. Why modelling inter-domain routing?
2. A model of the Internet
3. Some results
4. Conclusion

Outline

1. Why modelling inter-domain routing?
2. A model of the Internet
3. Some results
4. Conclusion

Why modelling inter-domain routing?

- Create models which are able to answer what-if questions
- Understand "policies" and their impact

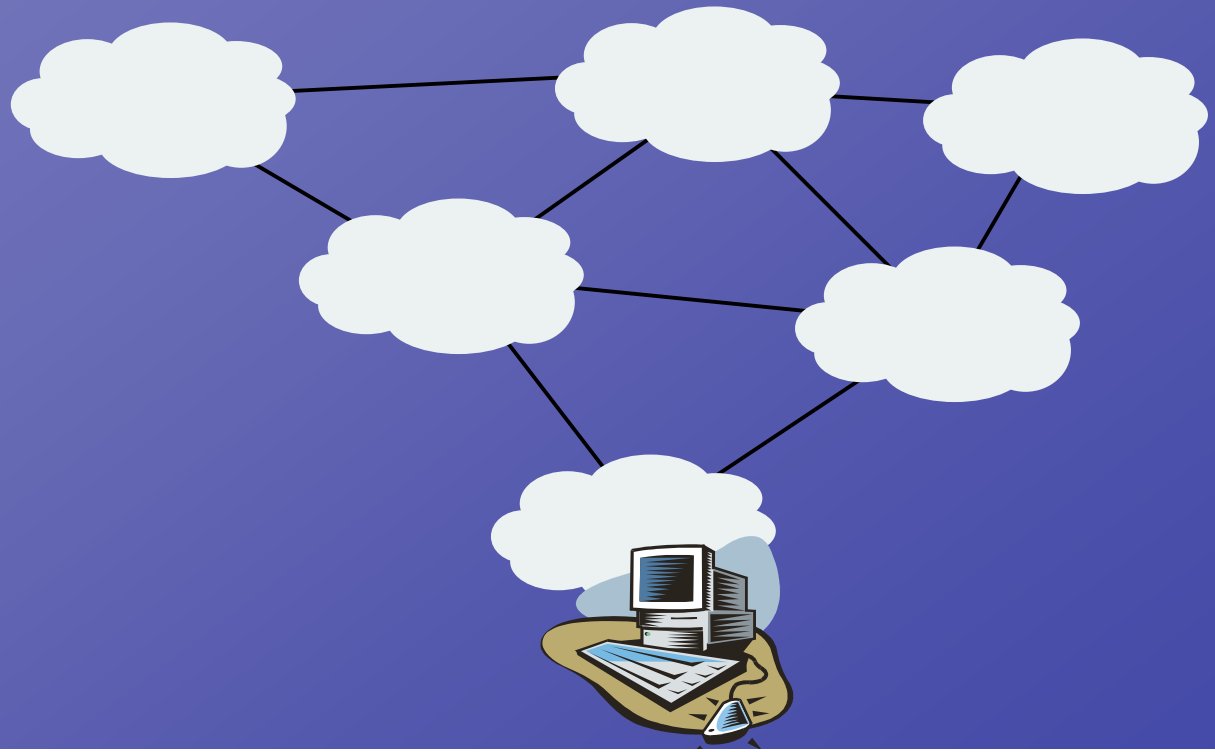

Typology of what-if questions

Goal of our inter-domain model is to potentially predict an outcome of changes

For that we need to know

- how routes propagate through the network
- which policies are applied on alternative paths

Knowing that helps

- to predict impact of policy changes
- to predict traffic flows
- or to debug the network

What-if: BGP policies

Understand potential impact of changes in policy:

- Impact of cancelled peering?
- How to identify “good” new peerings?
- Possibilities to improve policy config, e.g.,
  - check outcome of network configuration
  - formulate abstract AS-wide routing policy
- Poor path selection, e.g.,
  - identify highly asymmetric paths
  - identify long paths
  - changed connectivity between ASes

What-if: traffic flows

Predict traffic flow with the help of simulations:

- Inter-domain:
  - Predict where traffic enters/leaves the network?
  - How can I balance traffic among my neighbors?
- Intra-domain:
  - How traffic flows within my network?
    - by simulating end-to-end traceroutes
  - IGP/BGP interactions

What-if: debugging

- Optimize network performance
- Locate Internet routing instabilities
- Detect problematic routing conditions
  - Tim Griffin's BGP wedgies
- Checking what you are doing:
  - check if the current state is the desired one
  - emulate a planned configuration/policy change

Outline

1. Why modelling inter-domain routing?
2. A model of the Internet
3. Some results
4. Conclusion

Where do we start from?

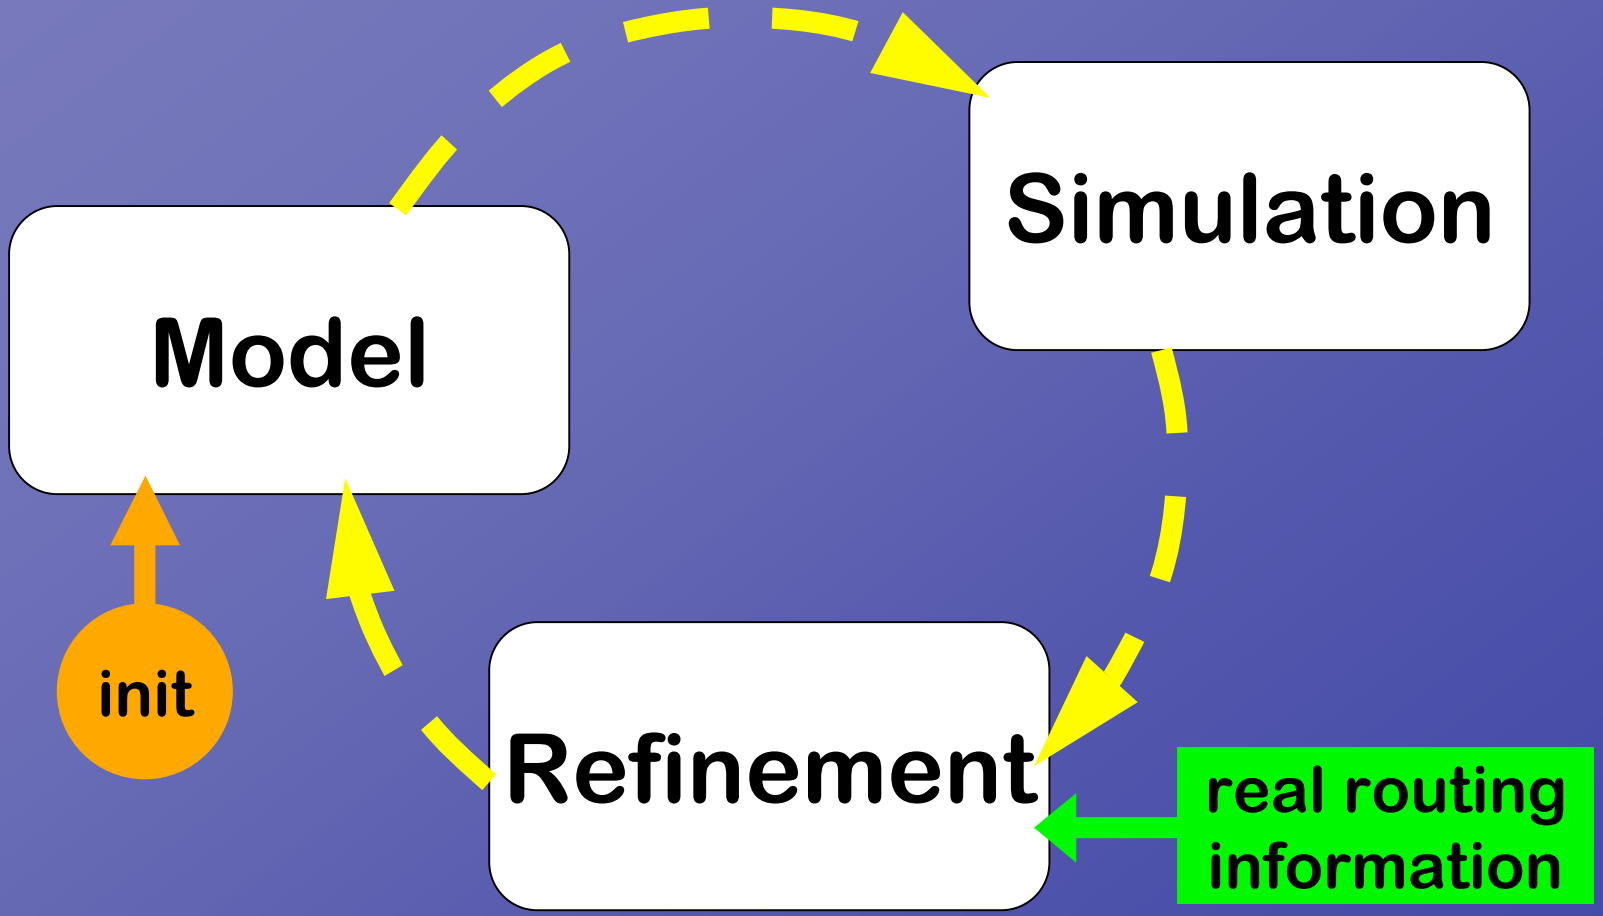

Where do we start from?

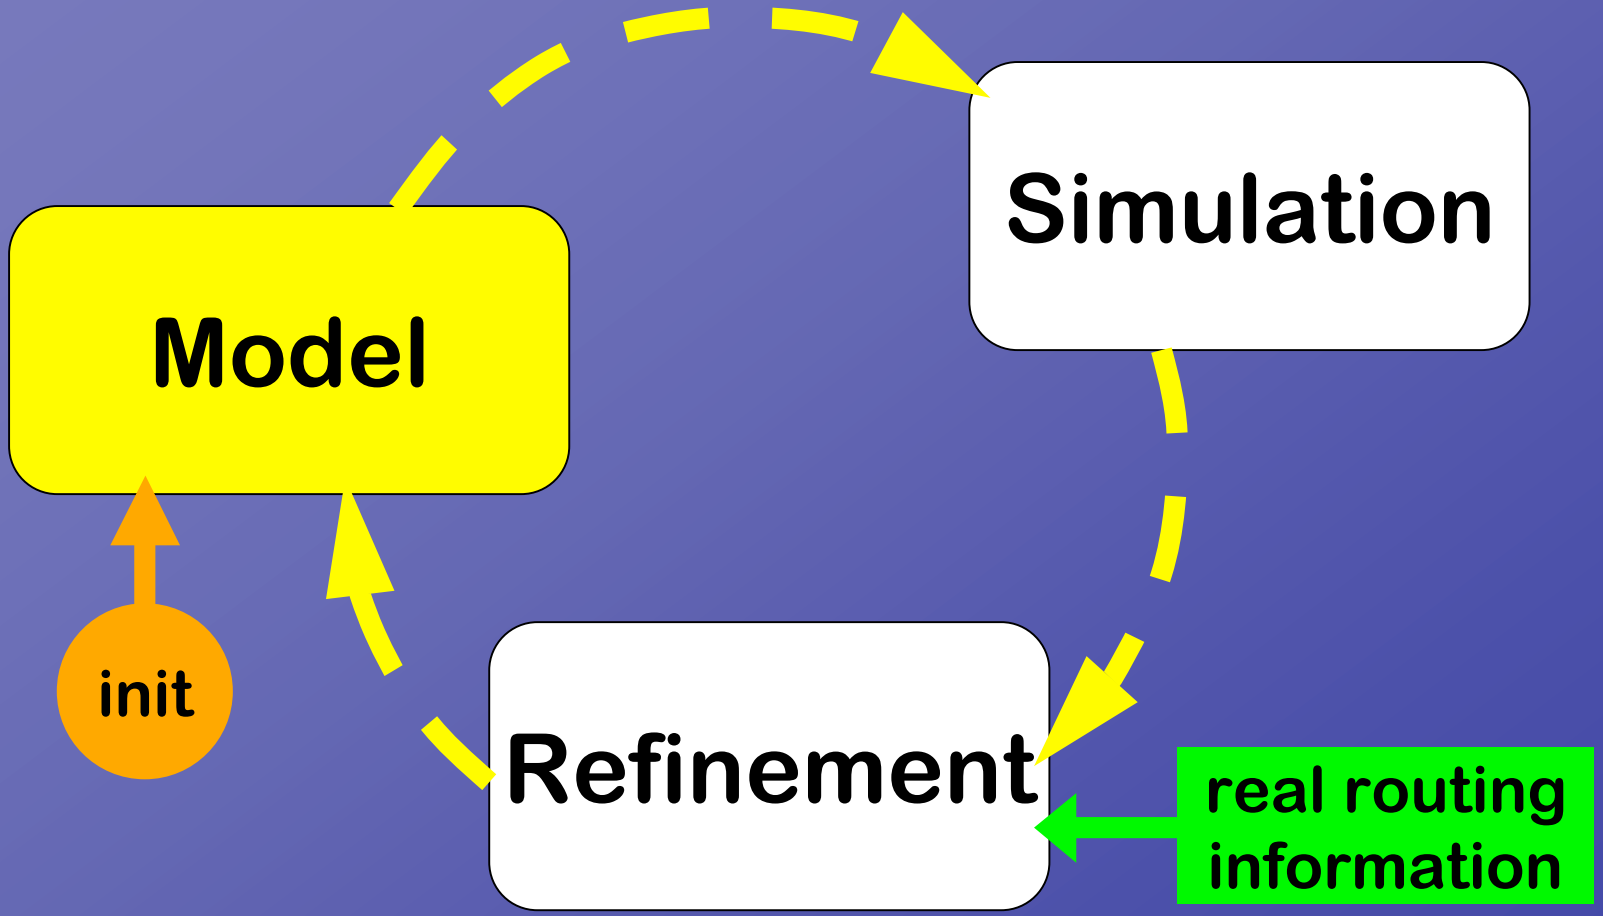

Build a model of the inter-domain routing system

Where do we start from?

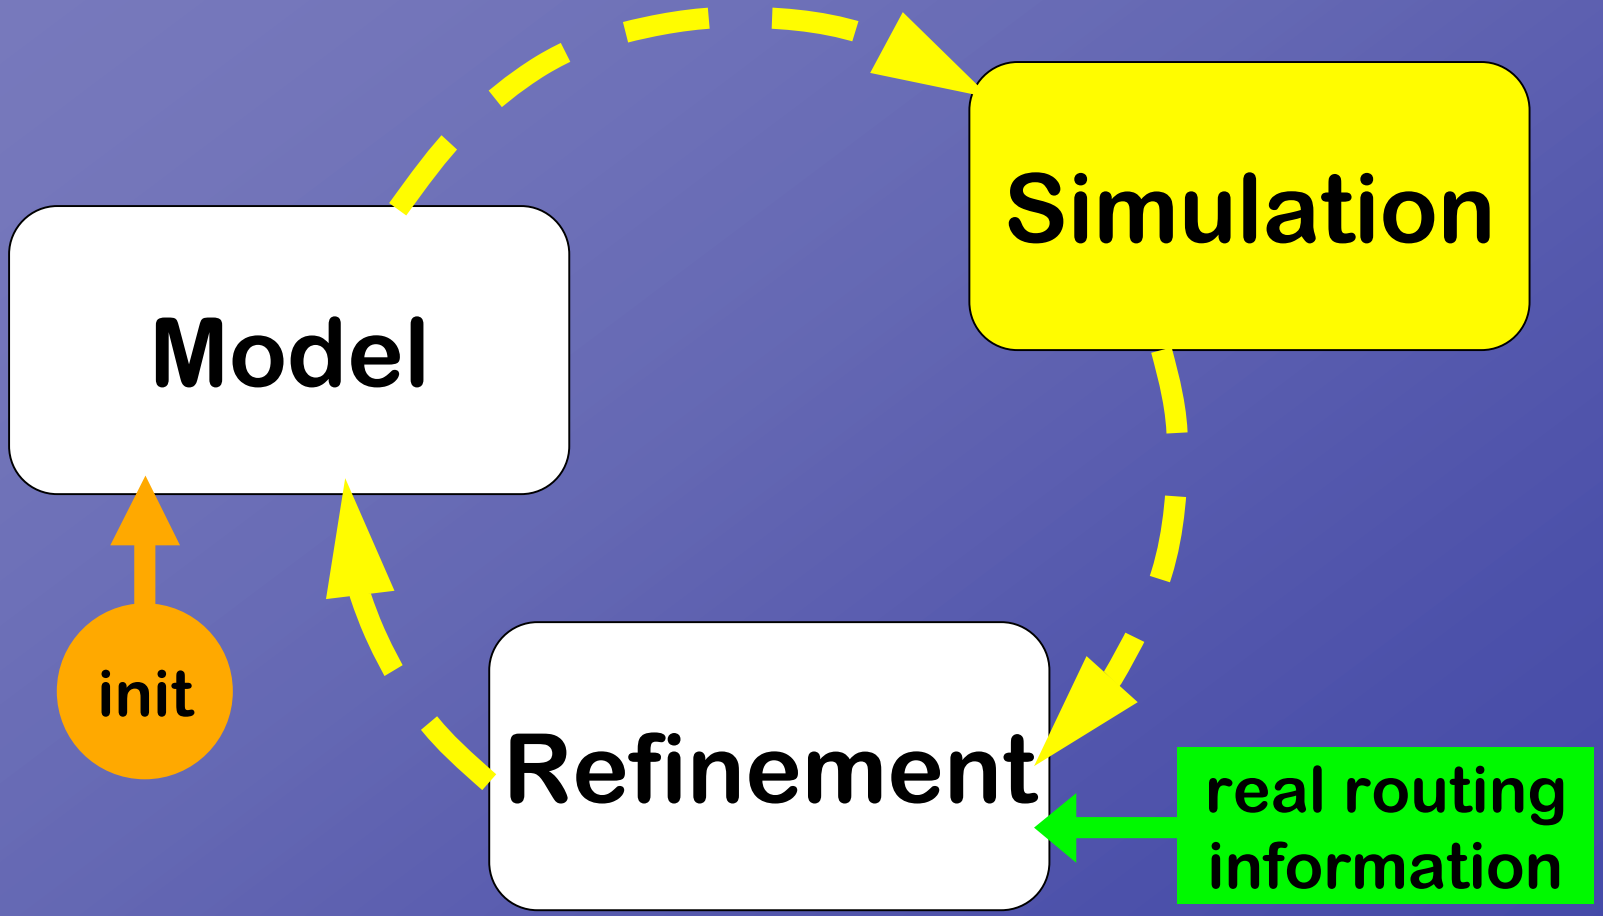

Simulate BGP path propagation…

Where do we start from?

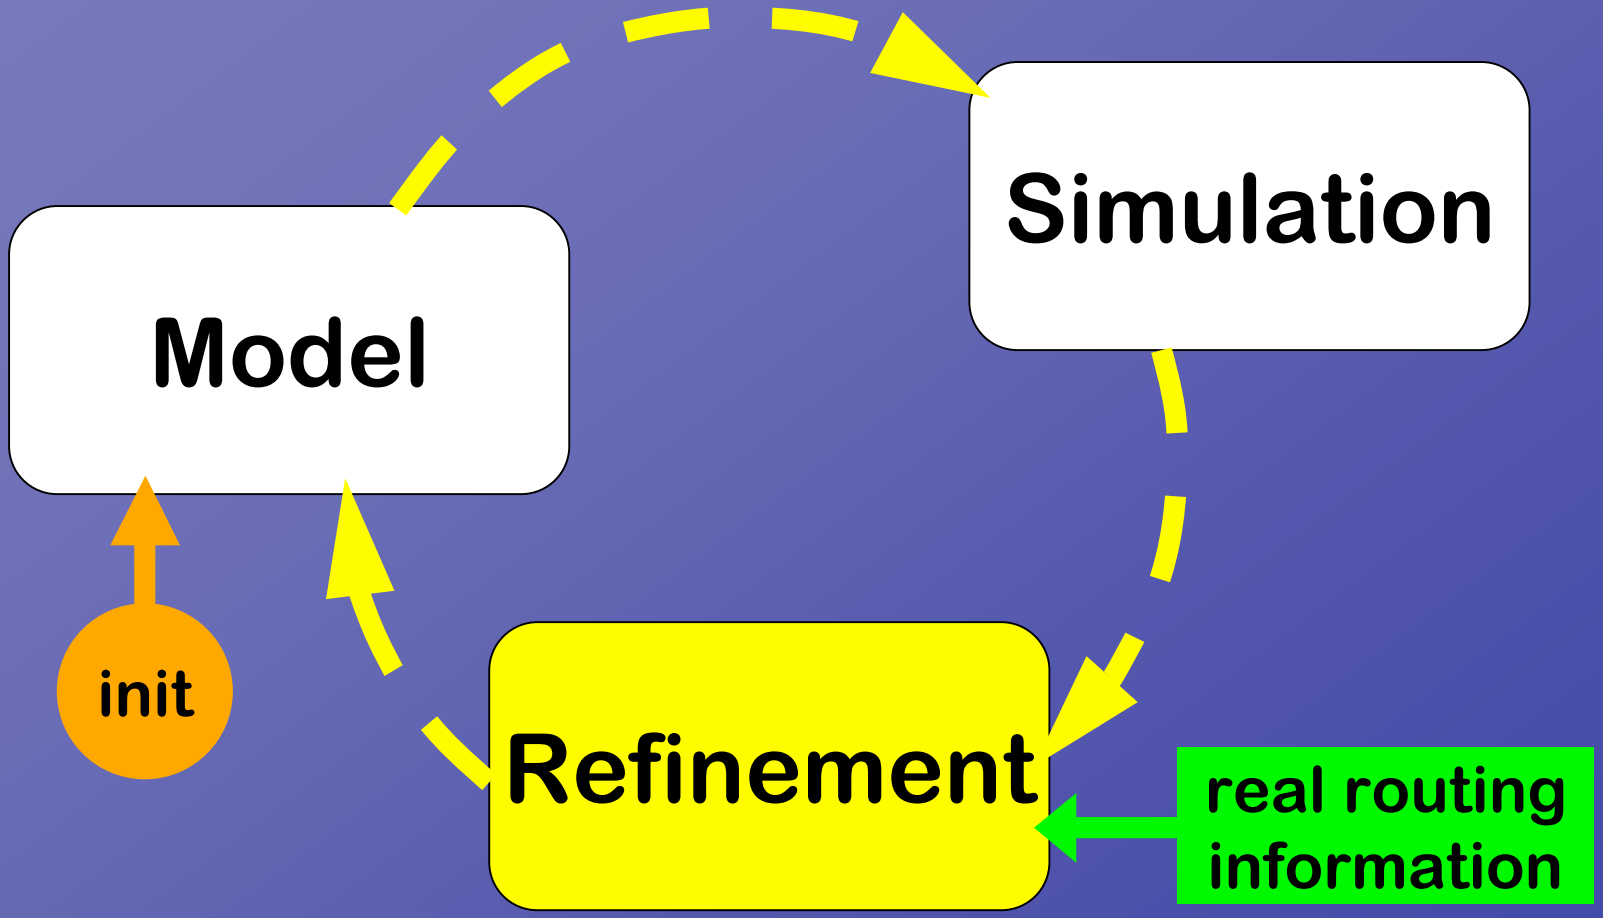

Compare observable (real) AS paths to simulation

“Simply” reproducing observable paths

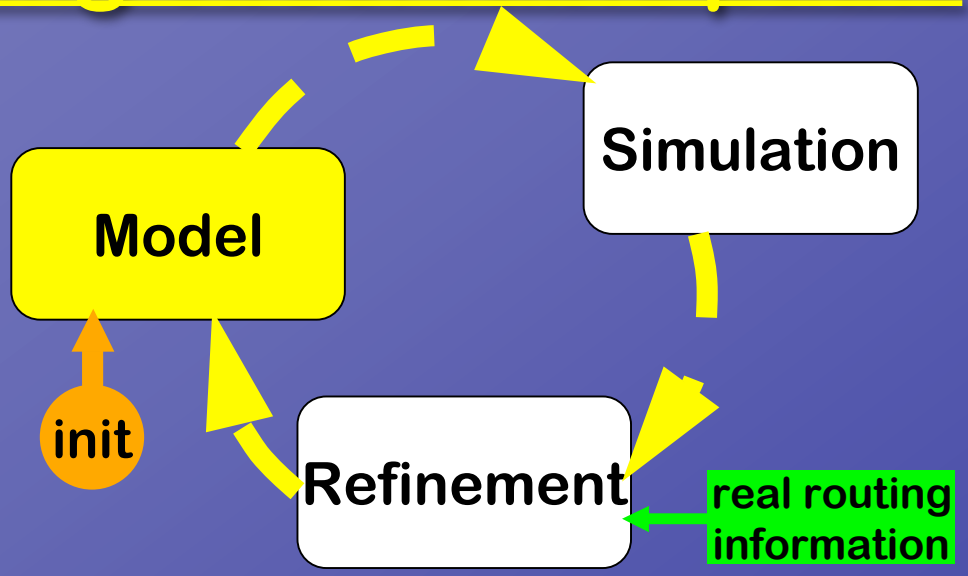

Premise:

- Only observable paths give us information about the AS-level topology and potential routing policies

Goal:

- Reproduce paths in C-BGP simulator

The C-BGP simulator

- Full BGP policies and decision process
- Parse Cisco and Juniper configs

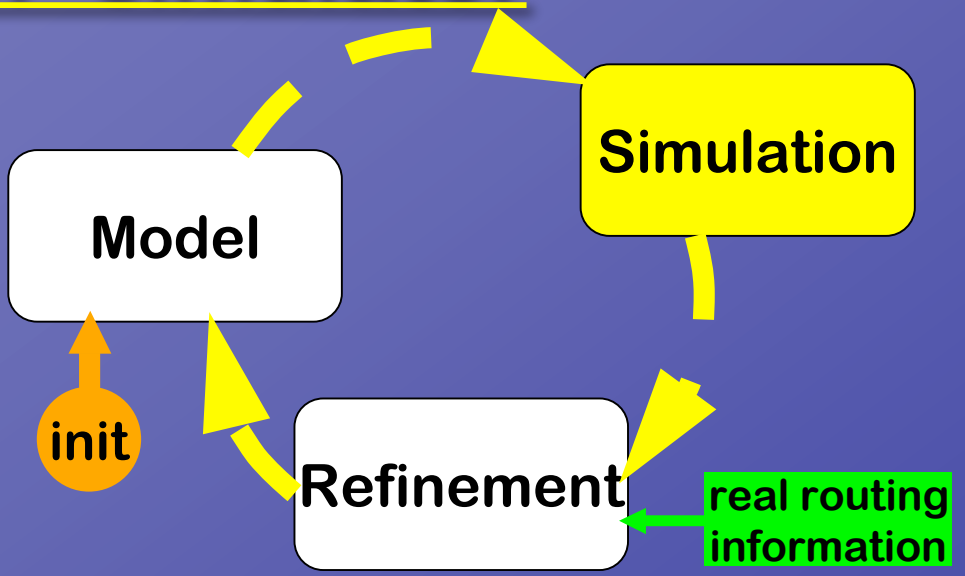

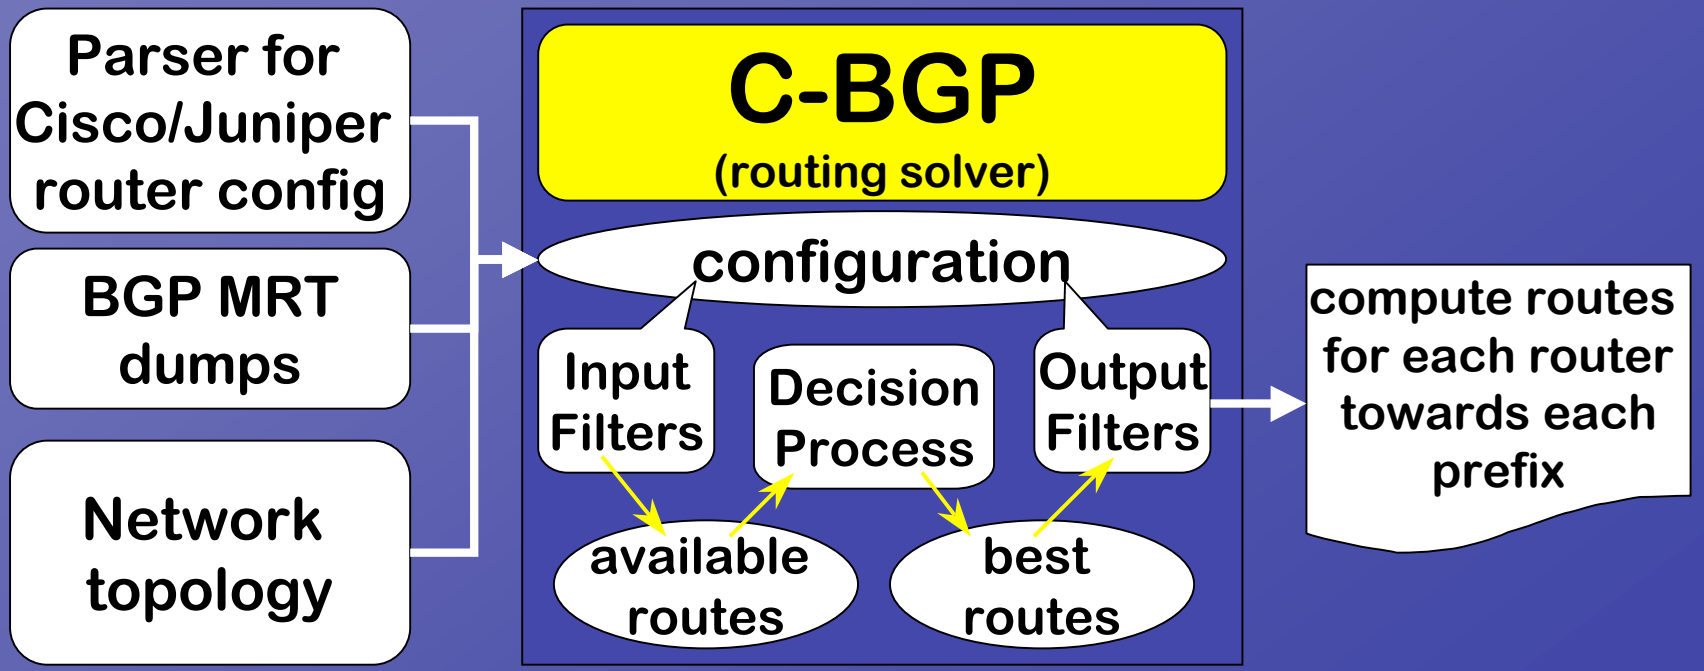

What to do?

Recall best path selection process:

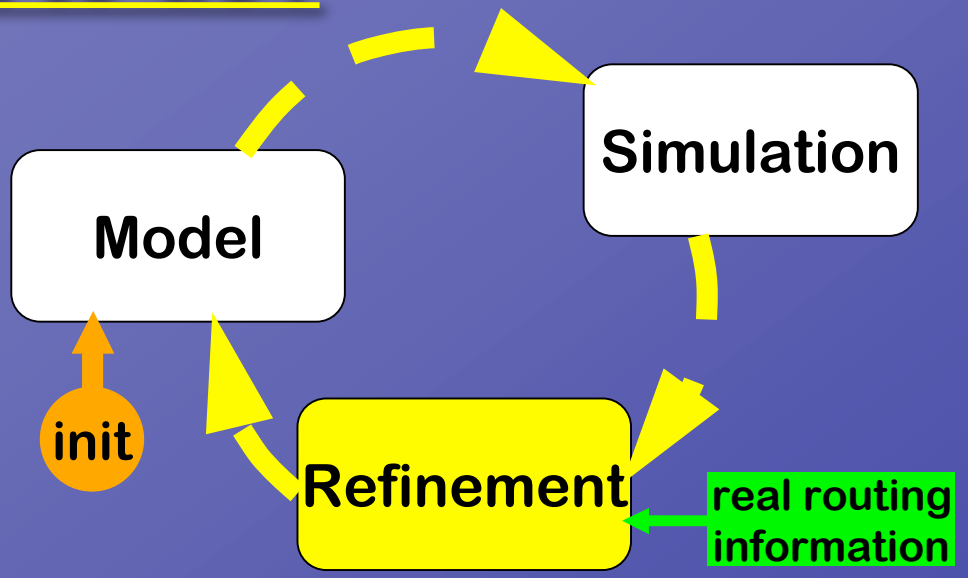

1. Local-pref
2. Shortest AS-path
3. Origin type
4. Lowest MED value
5. eBGP over iBGP
6. Lowest IGP cost ("hot-potato" routing)
7. Tie-break (e.g., Lowest router-ID)

What to do?

- **Split AS**, if multiple paths must be propagated
- **Filter**, if longer paths must be propagated
- Get rid of random decisions (**lowest router-ID**), when supporting information is available

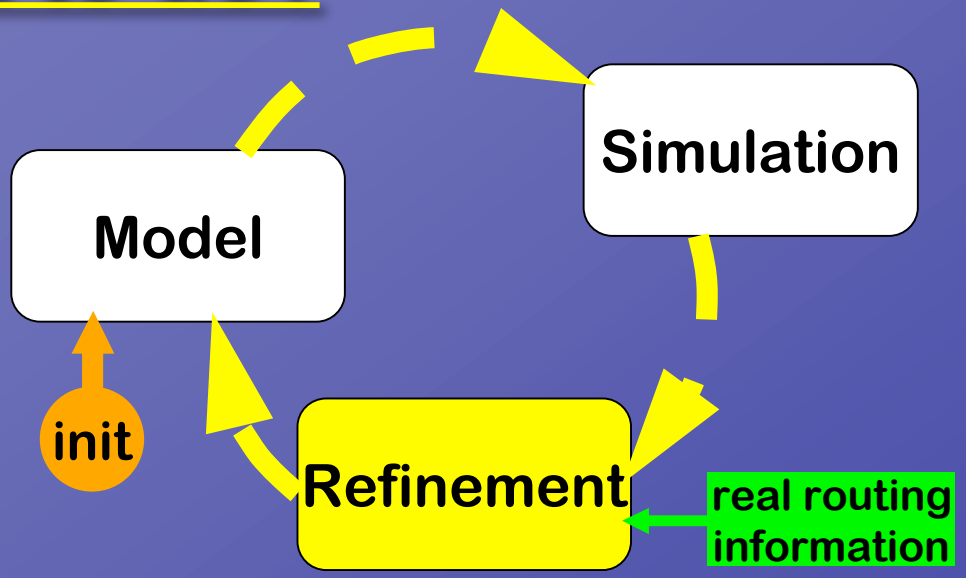

Splitting ASes

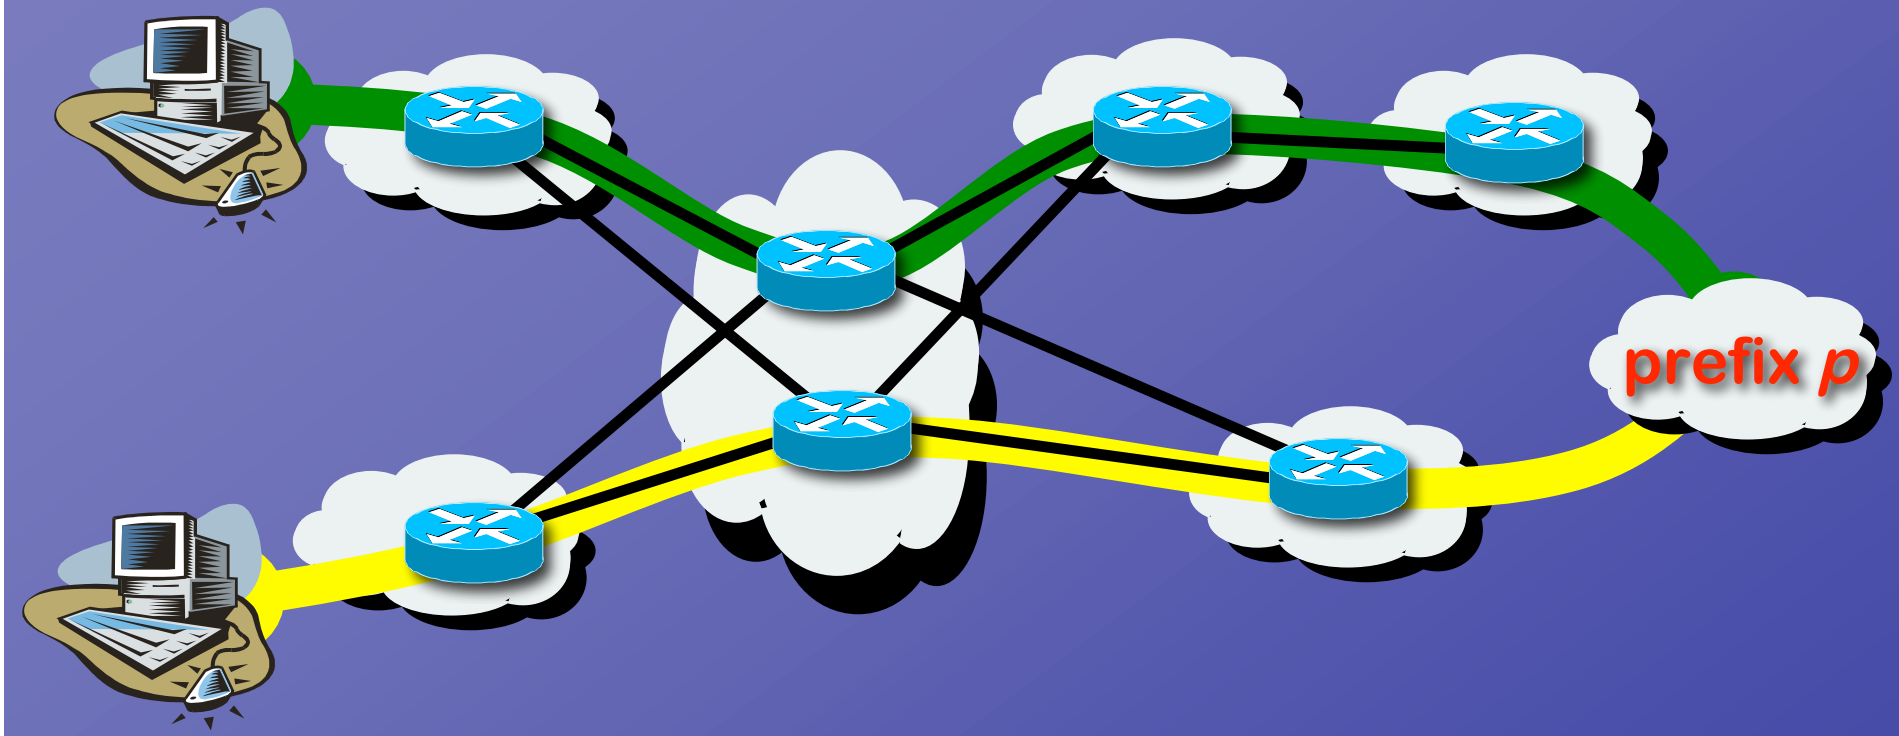

Multiple paths must be propagated …
=> need multiple routers in model!

How to propagate longer paths?

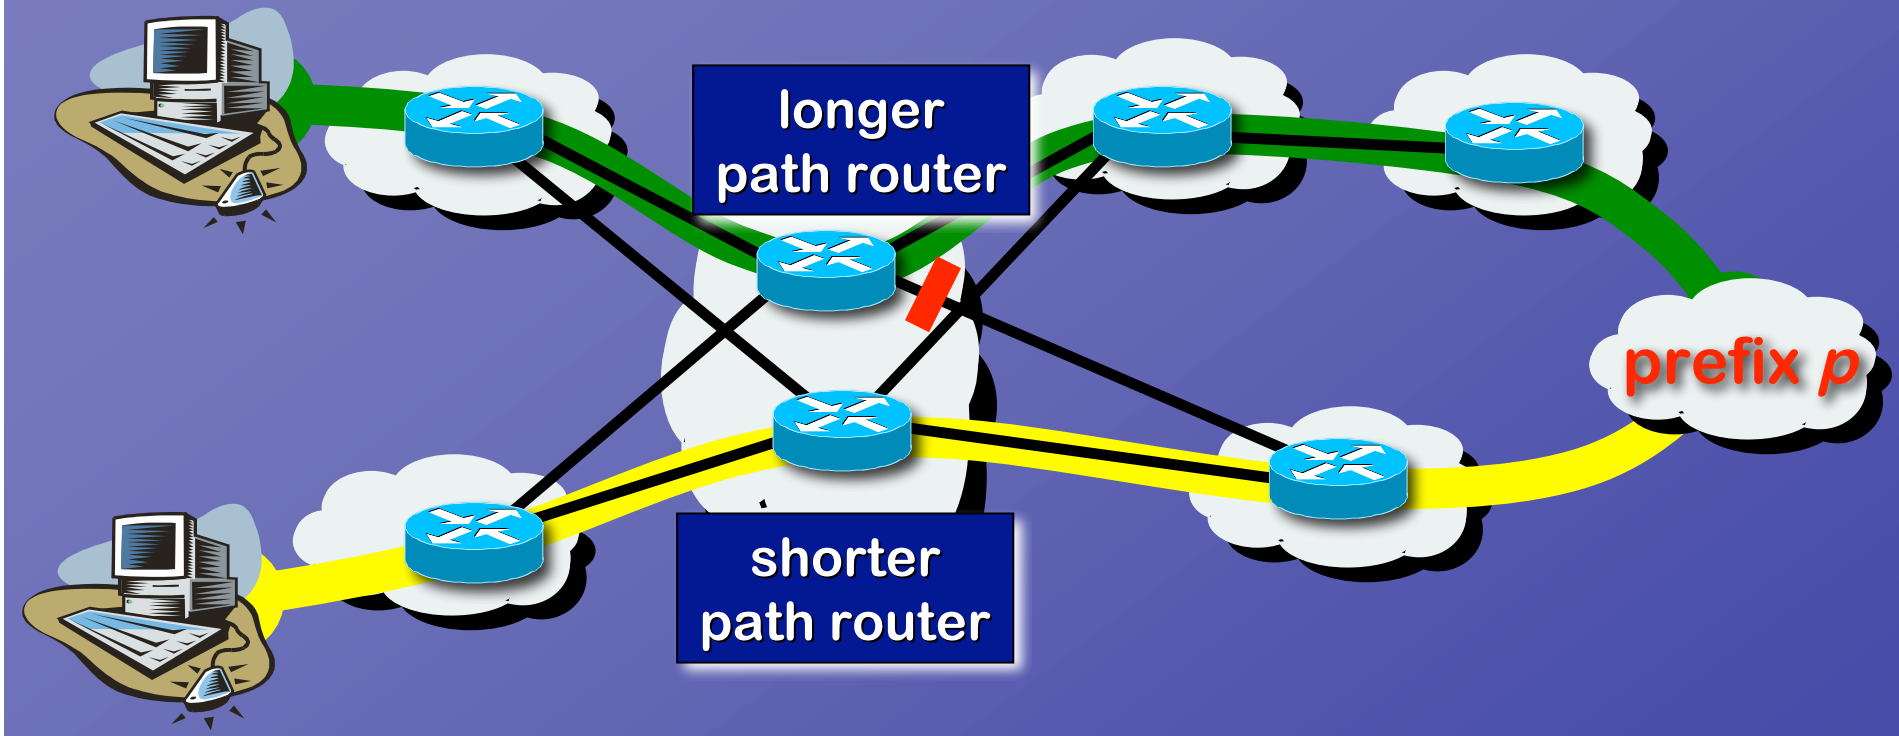

How to apply “our policies” ??

=> filter shorter path on “longer path router” (at ingress)

How to propagate longer paths?

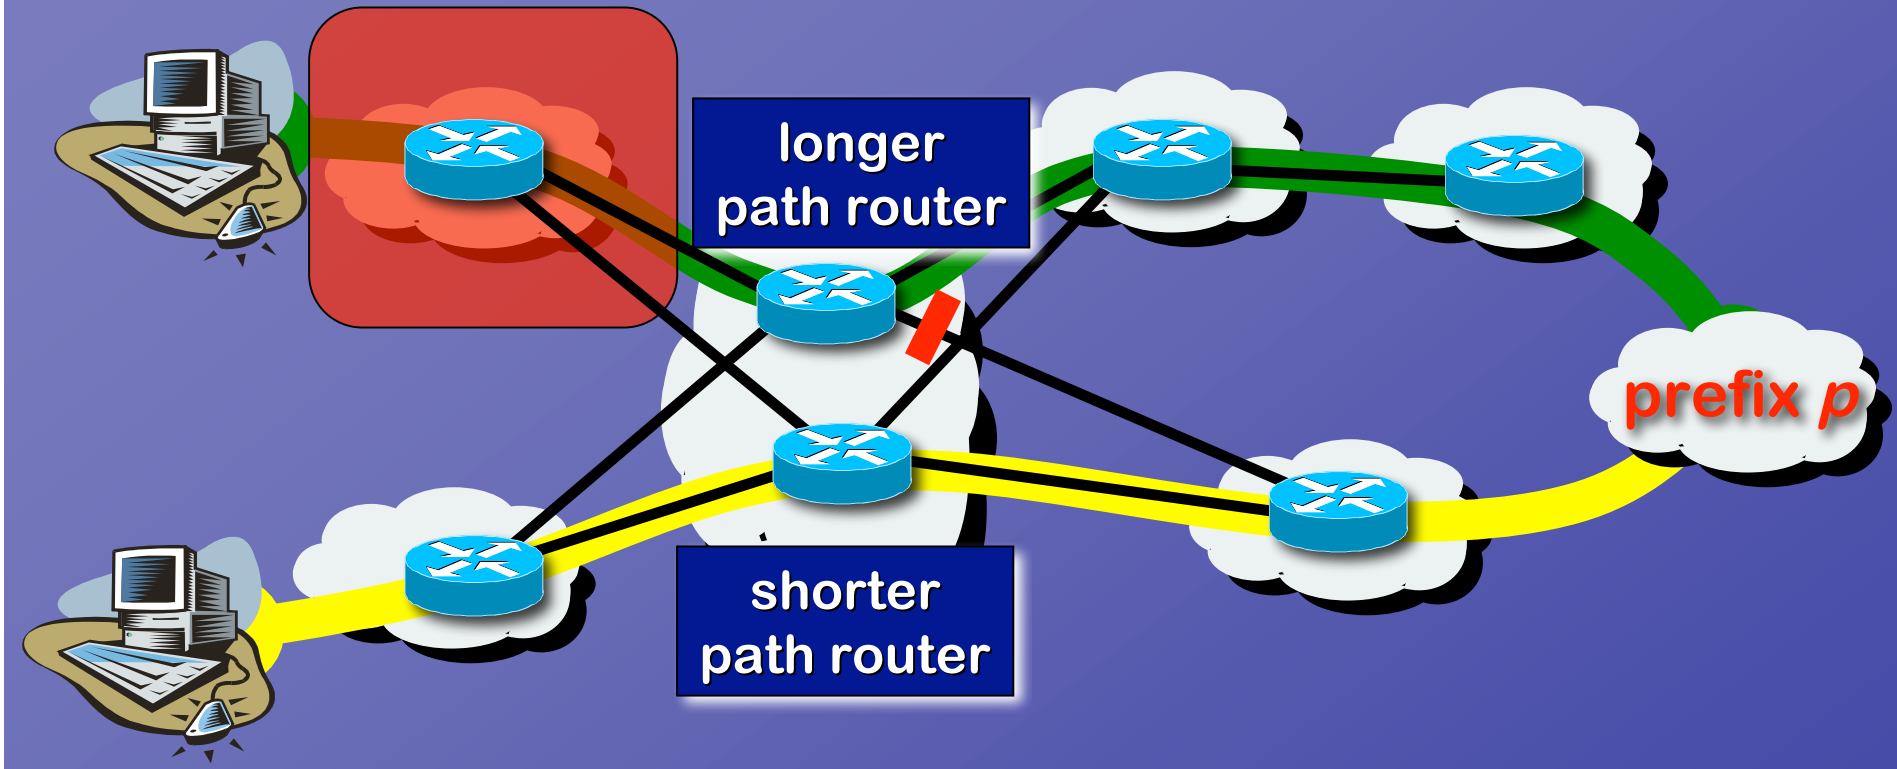

Propagate paths to appropriate neighbors!

How to propagate longer path?

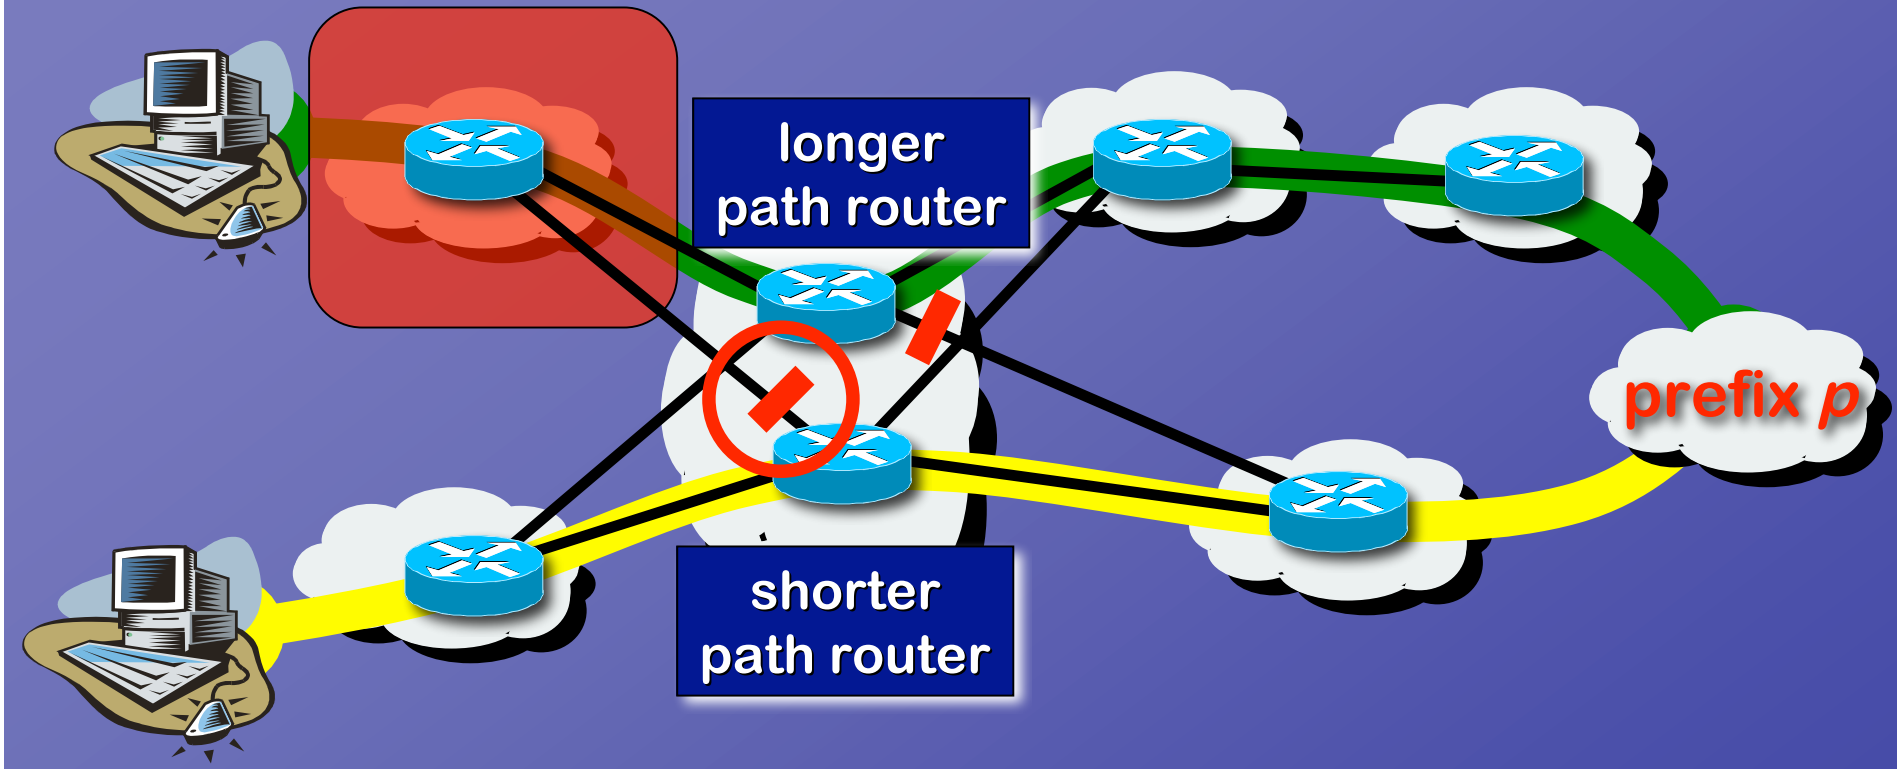

Filter also on “egress”-part of shorter path router.

Lowest Neighbor ID

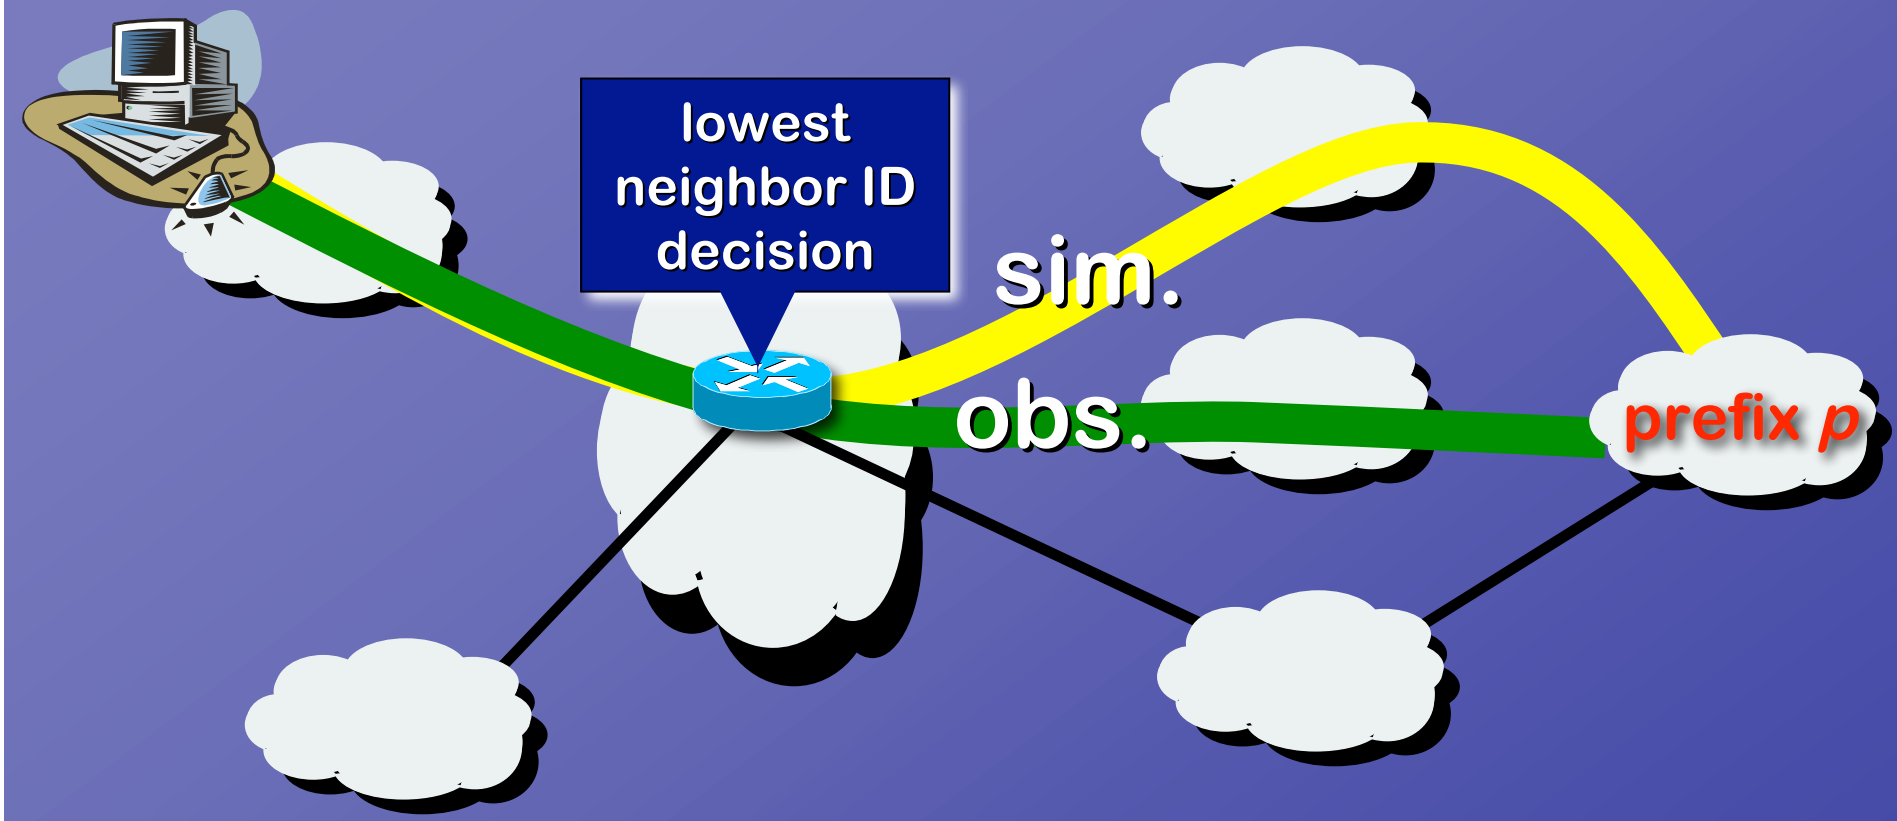

The decision was made by cBGP “randomly”

=> Fix random decision, if supporting information available

Outline

1. Why modelling inter-domain routing?
2. A model of the Internet
3. Some results
4. Conclusion

Terminology

- **Best match:** simulation selects a path that was observed in reality
- **RIB-IN:** simulation learns a path that was observed in reality, but did not pick that path as best
- **not found:** No router at the considered AS in the simulation learns about the path that was observed in reality

- **training data set:** paths that were used to build model (real observed paths)
- **validation data set:** paths that were *unknown* to the model, but that are real observed paths

Observation points

Training-Data-Set : used as input to model

Validation-Data-Set : NOT used as input

Source	training	validation	total
Akamai	600	302	902
RIPE	239	120	359
RouteViews	65	28	93
GEANT	16	6	22
Abilene	8	2	10
SpaceNet	4	1	5
total	932	459	1391

Initial Results

Training Data-Set	% unique paths
RIB-IN (learned)	99.99%
best path (selected)	99.98%
Validation Data-Set	
RIB-IN (learned)	93.84 %
best path (selected)	63 %

Lessons learned

- **Possible to model observable paths**
  - selecting the correct best path
  - propagate path among possible alternatives
- Requirement: **sufficient information**
  - enough observation points
  - diverse location of observation points

=> It is possible to construct a model that can answers some what-if questions

What next?

- 63% of the paths in the validation data-set were correctly predicted.
- Reasons:
  - We do not reverse engineer the Internet!
  - We do not know the real policies!
  - We "only" have policies which are consistent with our observations…

=> It is not easy to reverse engineer policies of any AS without sufficient observation points!

Outline

1. Why modelling inter-domain routing?
2. A model of the Internet
3. Some results
4. Conclusion

Conclusion

- C-BGP is a scalable simulator for large intra- and inter-domain topologies
- To answer what-if questions three information sources are required:
  - intra-domain topology (incl. router configs)
  - inter-domain topology (incl. local BGP views)
  - traffic information
- To make this practical, we need your help!

Thanks for your attention!