

Preventing the unnecessary propagation of BGP withdraws

V. Van den Schrieck, P. François, C. Pelsser, O. Bonaventure
<http://inl.info.ucl.ac.be>

Agenda

- Introduction
- Presentation and evaluation of iBGP-caused withdraws
- Preventing iBGP-caused withdraws
 - ✓ Advertising more than one path in iBGP (IETF)
 - ✓ Our proposal : Advertising the availability of alternate paths in iBGP


Agenda

- **Introduction**
- Presentation and evaluation of iBGP-caused withdraws
- Preventing iBGP-caused withdraws
 - ✓ Advertising more than one path in iBGP (IETF)
 - ✓ Our proposal : Advertising the availability of alternate paths in iBGP


BGP, the glue of the Internet

- Border Gateway Protocol - Interdomain routing protocol used in the Internet
- Exchange paths to destinations between Autonomous Systems (AS)
 - ✓ Each router advertises its best path to its BGP peers
 - ✓ Selection process : Several rules
- A path is described by the AS Path, i.e. the sequence of AS that leads to the path destination


BGP




BGP




BGP




BGP




BGP



BGP




BGP



iBGP vs eBGP

- An Autonomous System has several routers
- All routers of an AS need to learn the paths to all destination advertised to the AS
- ✓ BGP sessions between routers of the AS : iBGP
 - Specific iBGP rule : Do not advertise an iBGP-received path on an iBGP session
 - Exception : Route Reflectors
- ✓ BGP sessions between routers of different AS : eBGP

iBGP vs eBGP




Agenda

- Introduction
- **Presentation and evaluation of iBGP-caused withdraws**
- Preventing iBGP-caused withdraws
 - ✓ Advertising more than one path in iBGP (IETF)
 - ✓ Our proposal : Advertising the availability of alternate paths in iBGP

BGP Churn

- Churn is one of today's BGP issues
- Churn : High number of unnecessary BGP messages
 - ✓ Unnecessary : Do not impact the final routing state
- Consequences :
 - ✓ Load on router's CPU
 - ✓ Transient losses of reachability

BGP Churn




<http://bgpupdates.potaroo.net/instability/bgpupd.html>

- Average daily number of BGP updates : 400.000

Origin of BGP churn

- Path exploration during BGP convergence
- Flapping prefixes
- Bad iBGP propagation
 - ✓ Routers without path diversity could send unnecessary BGP withdraws in case of failure
 - ✓ **Even tough alternate paths exist in the AS**

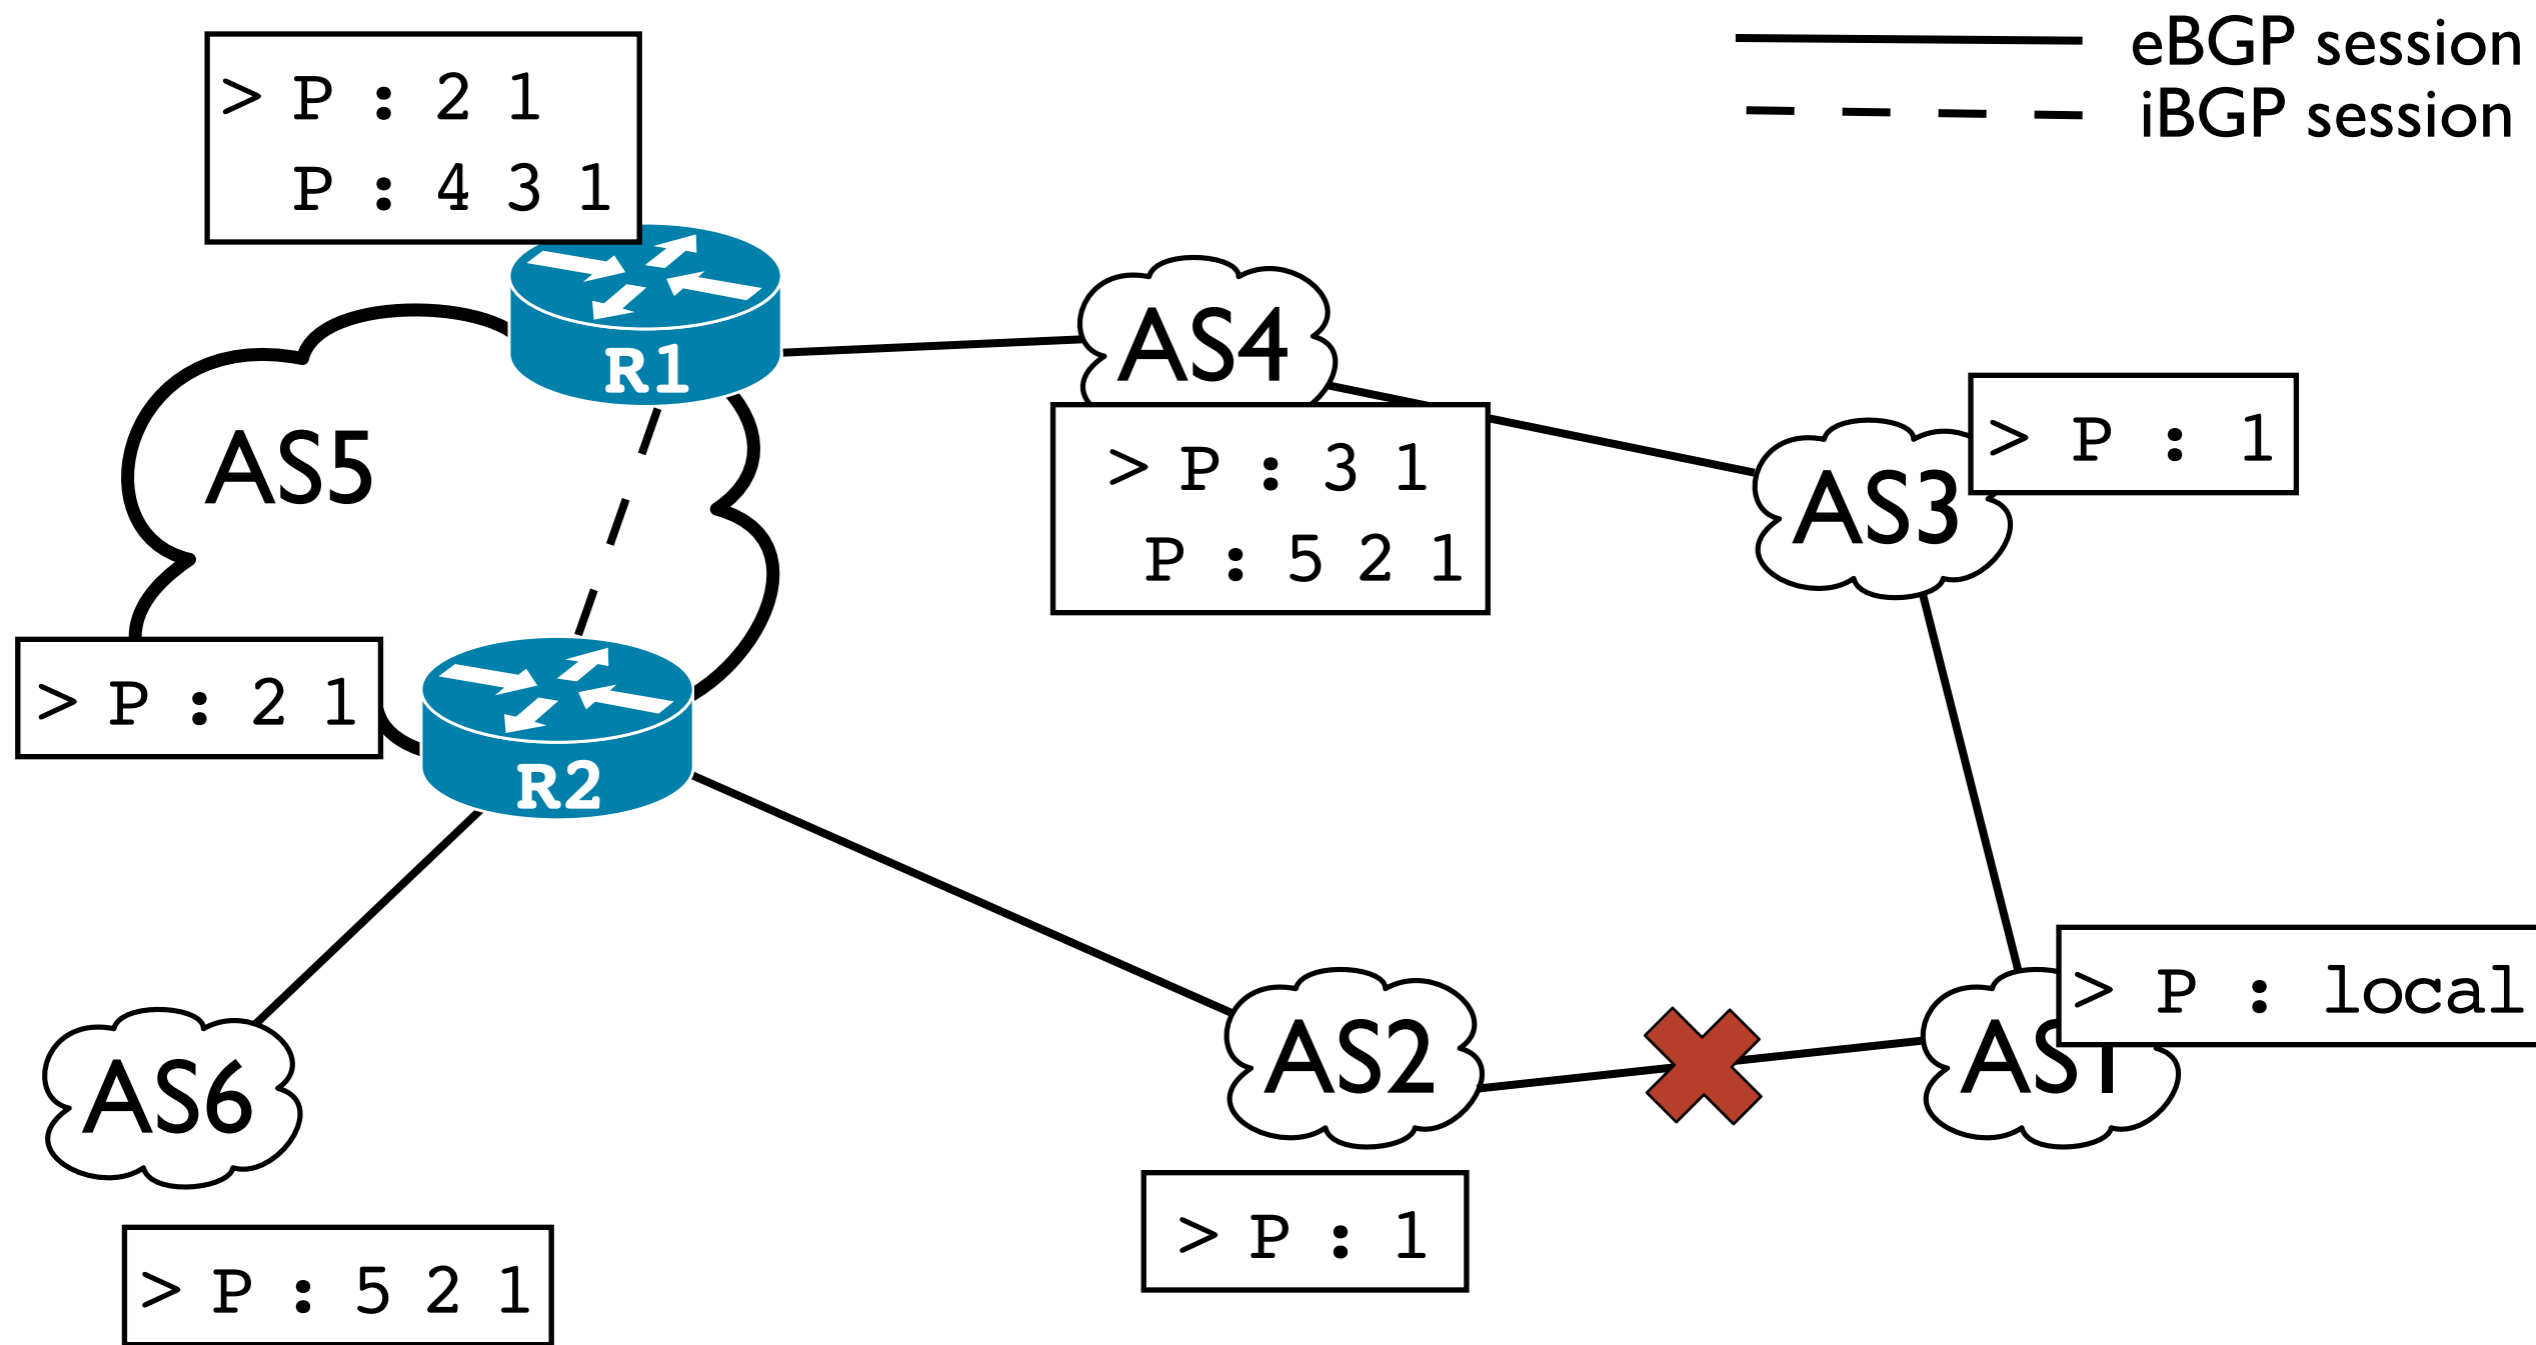
Bad iBGP propagation



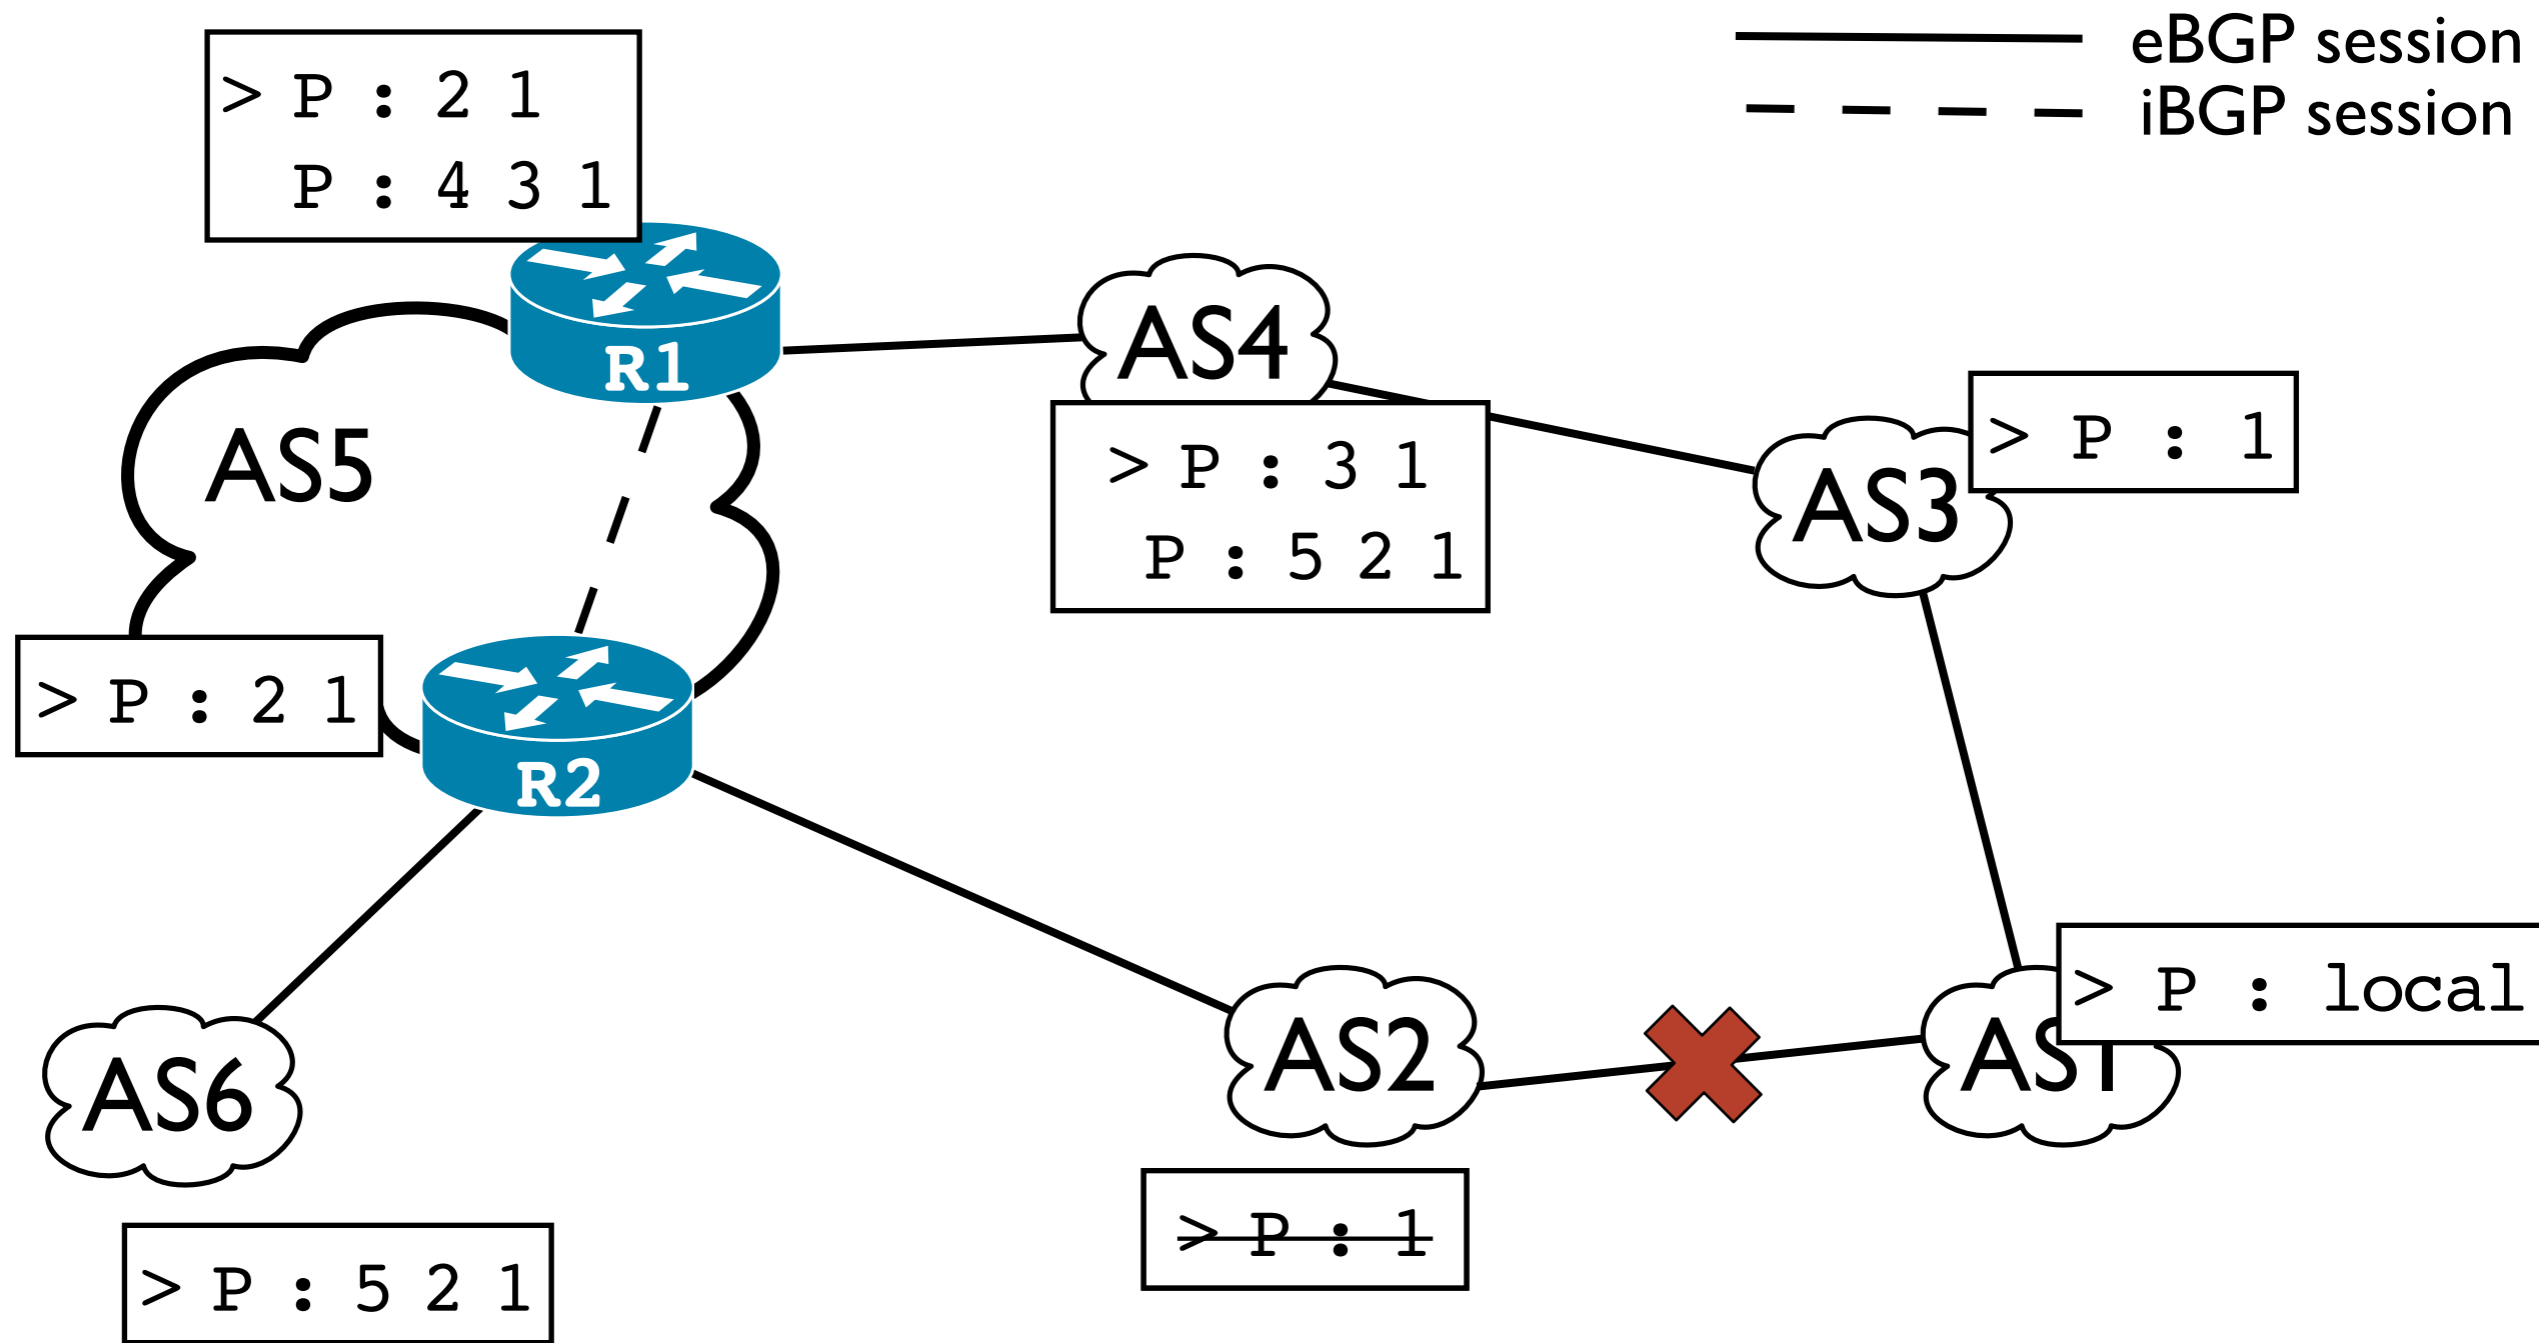
Origins of bad iBGP propagation

- Only best path announced on iBGP sessions
- One path preferred over all others => alternate paths are never propagated
 - ✓ Path quality (AS path length)
 - ✓ BGP policies (Local preference)
- Complex iBGP organizations
 - ✓ Paths hidden by Route Reflectors

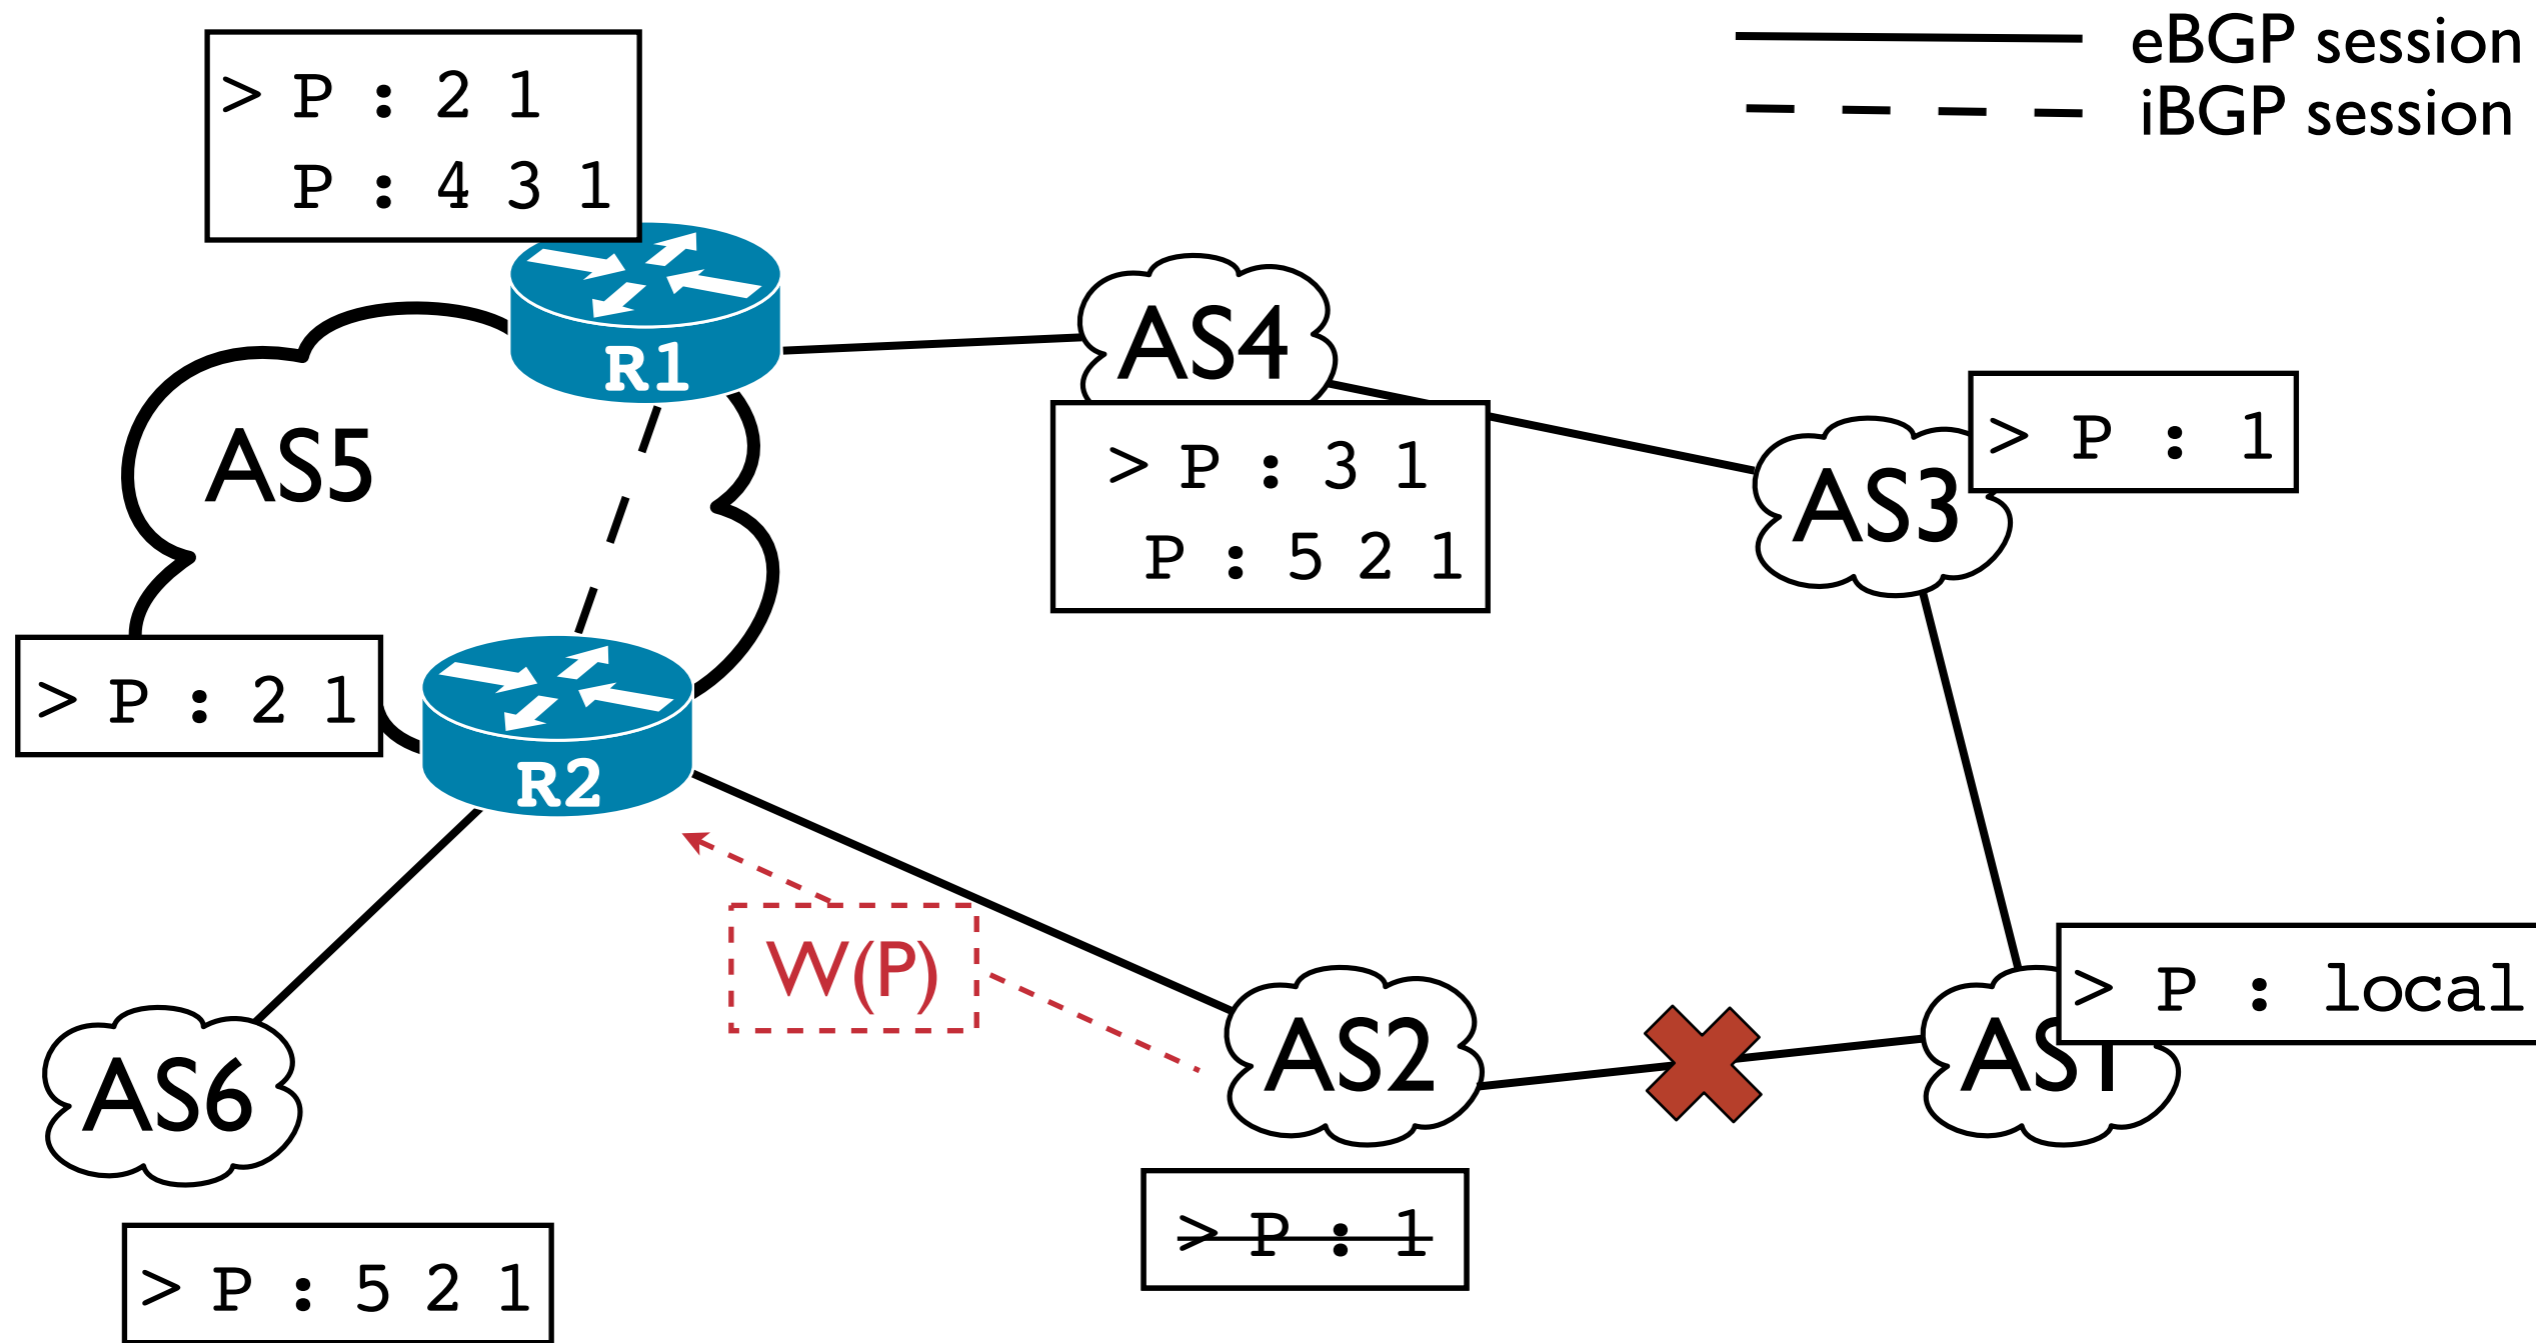
Withdraw Propagation when alternate path exists



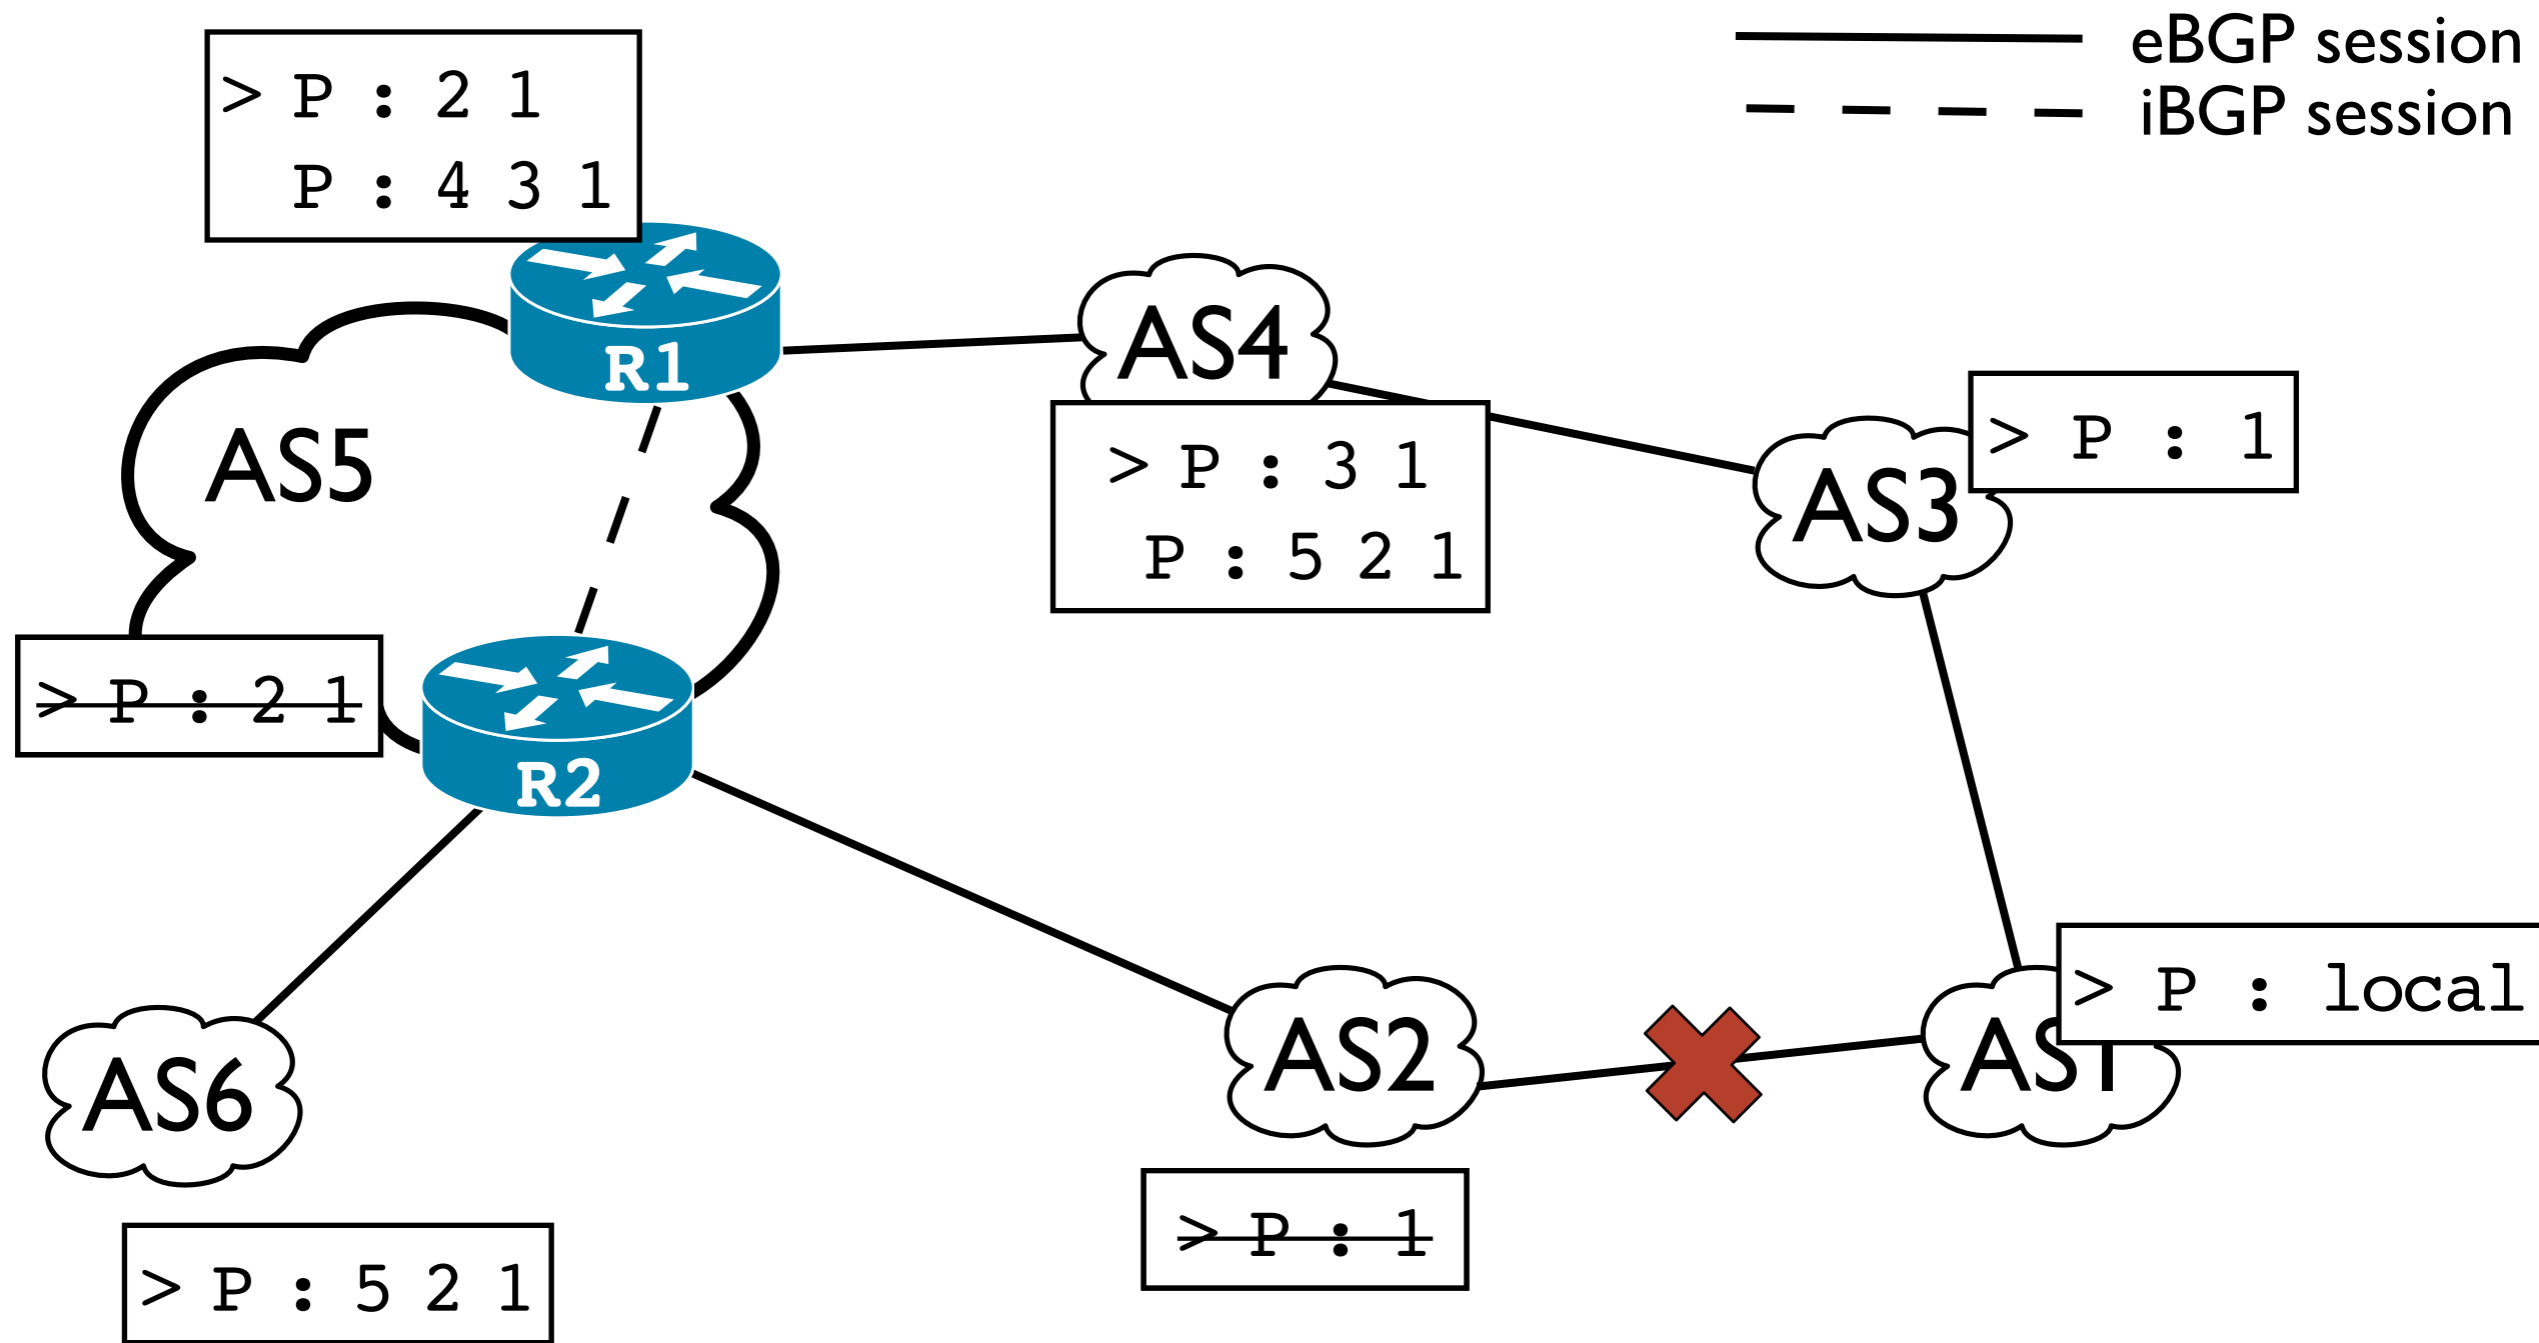
Withdraw Propagation when alternate path exists



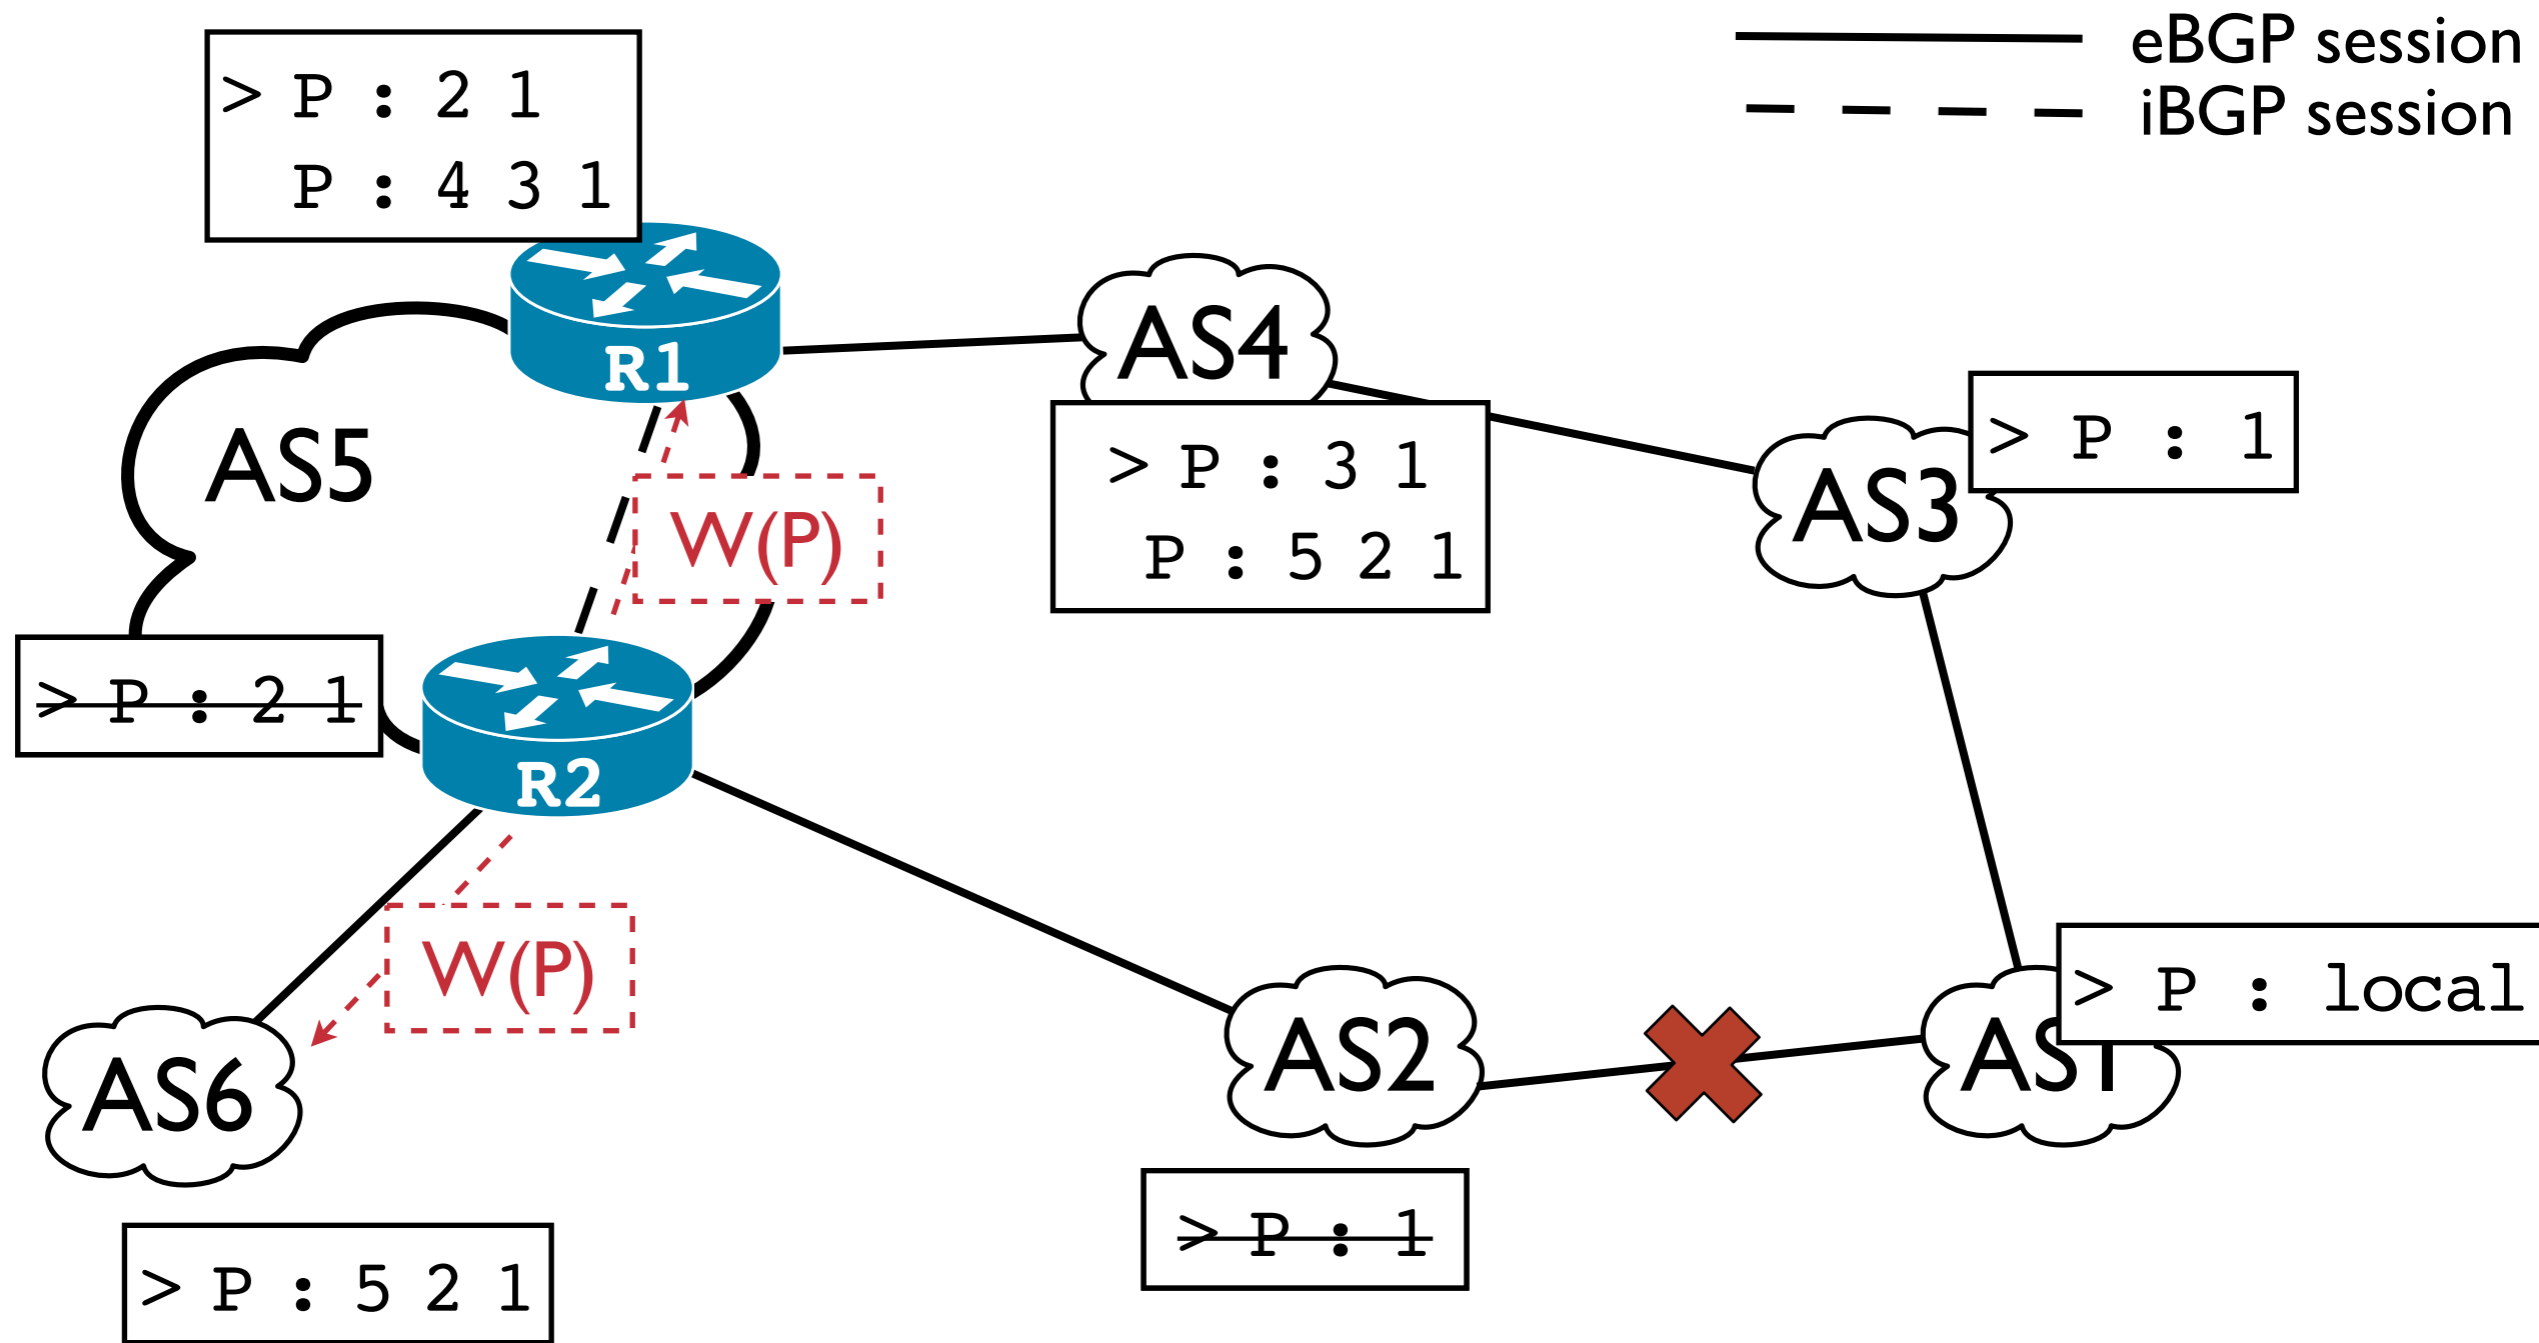
Withdraw Propagation when alternate path exists



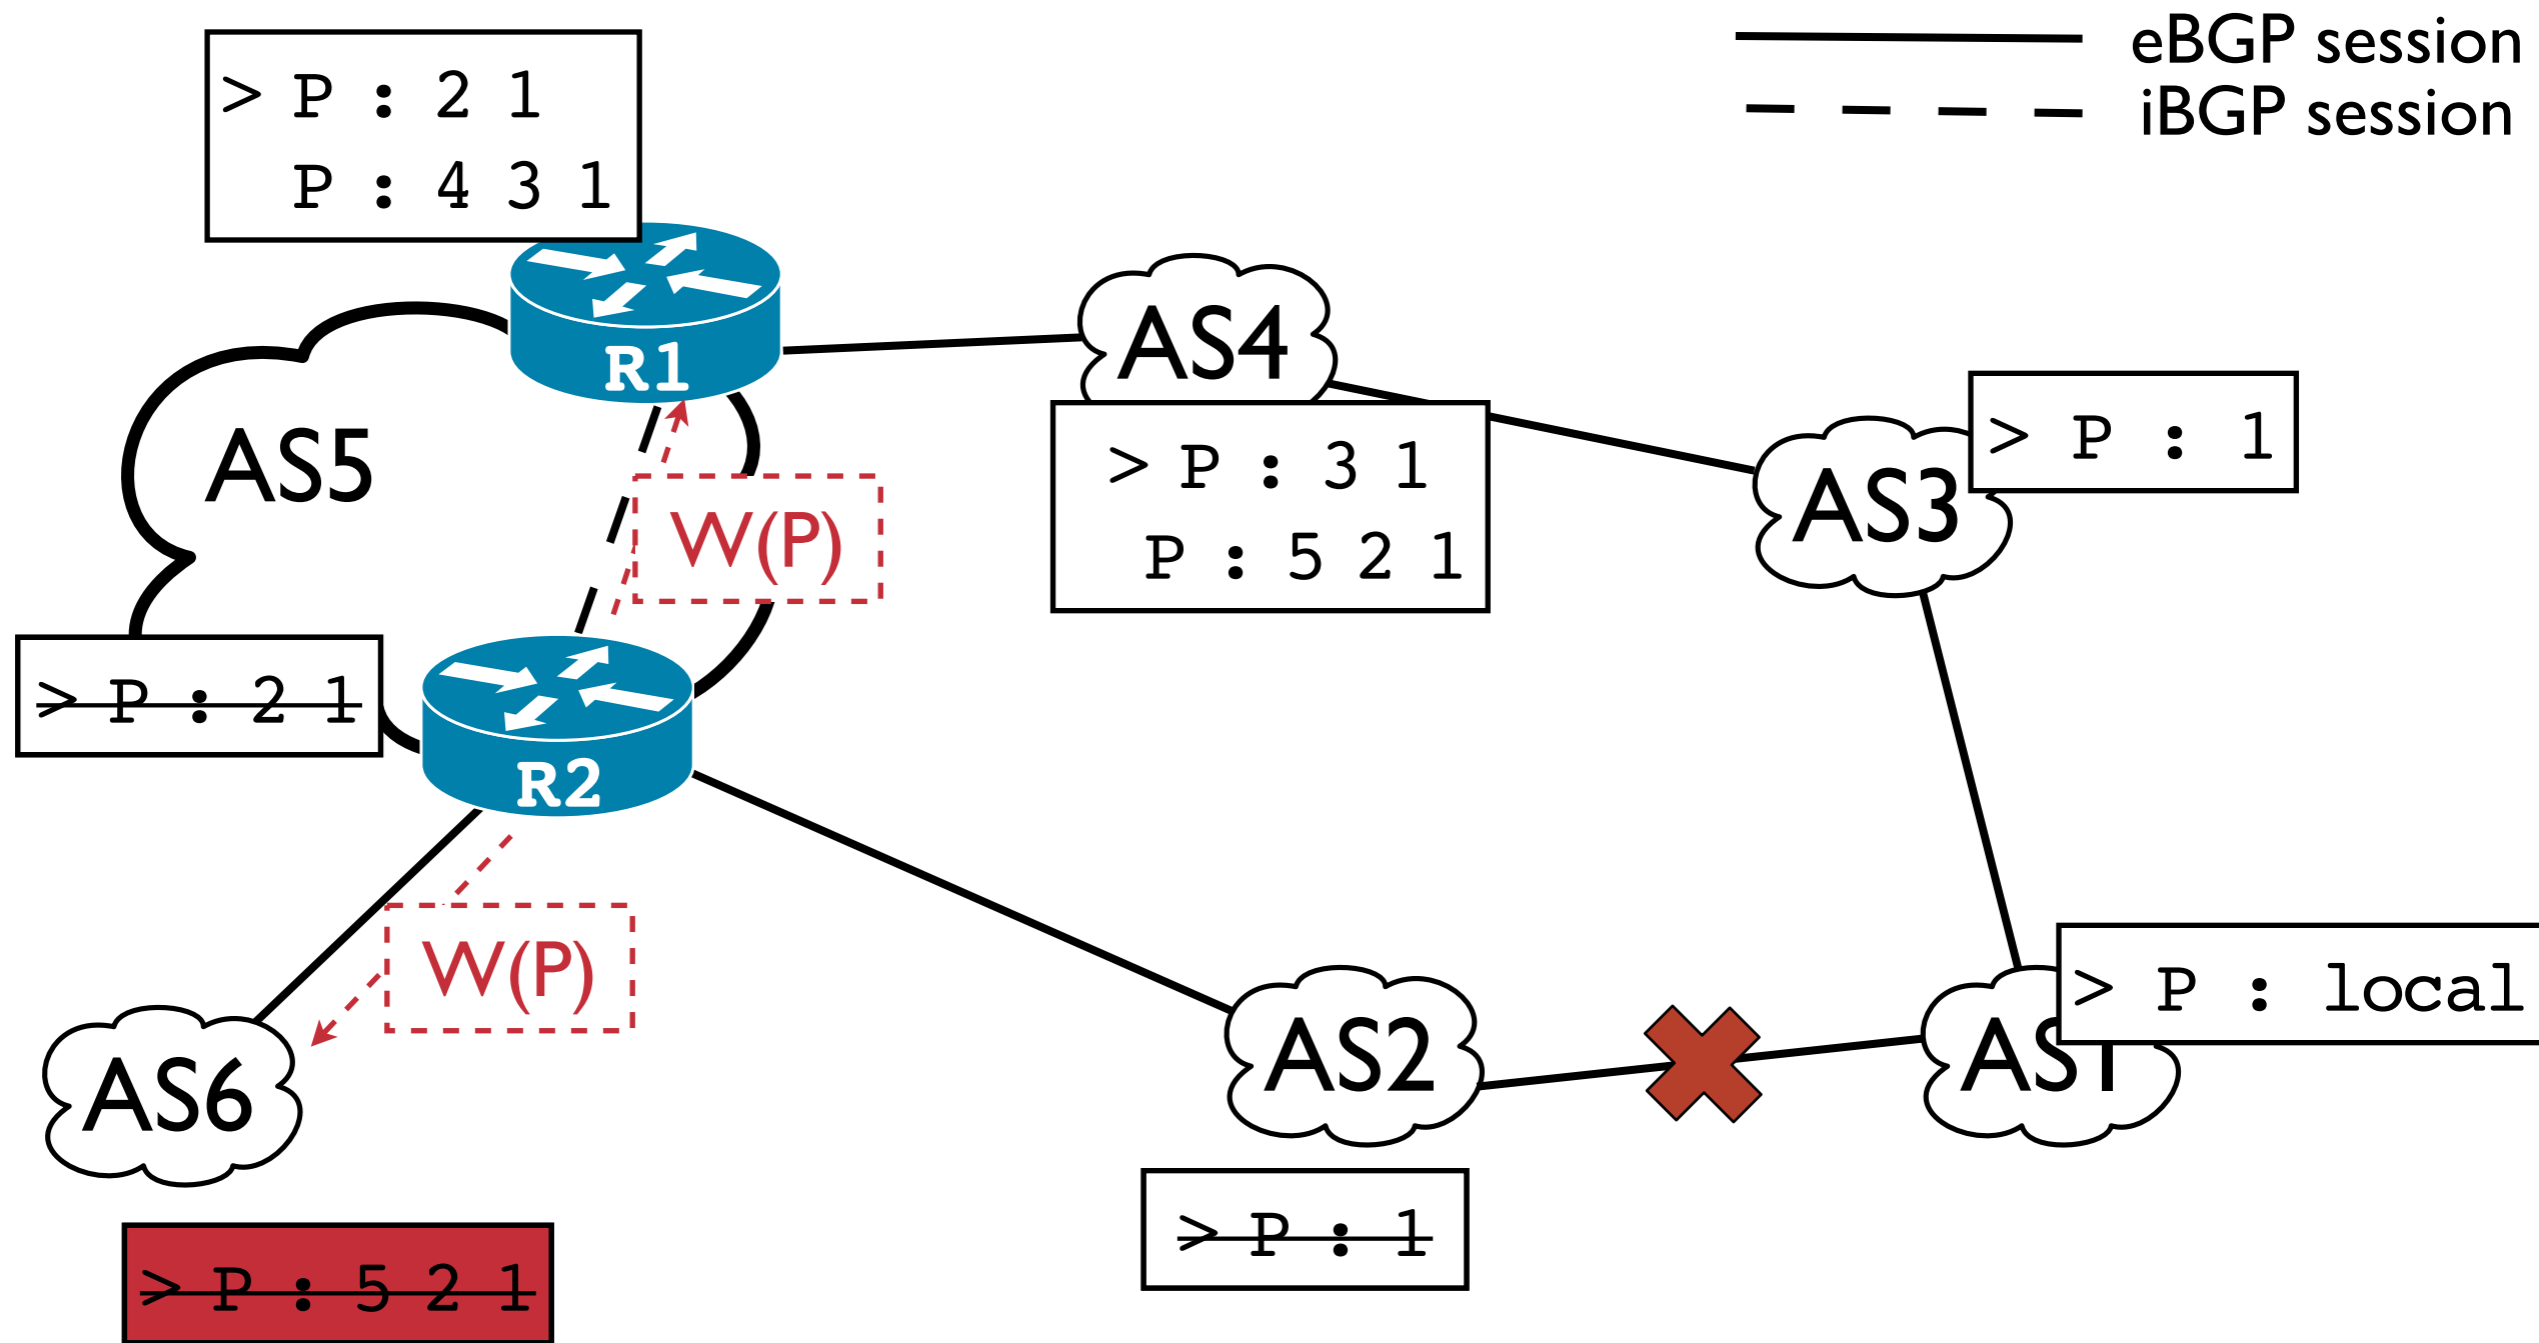
Withdraw Propagation when alternate path exists



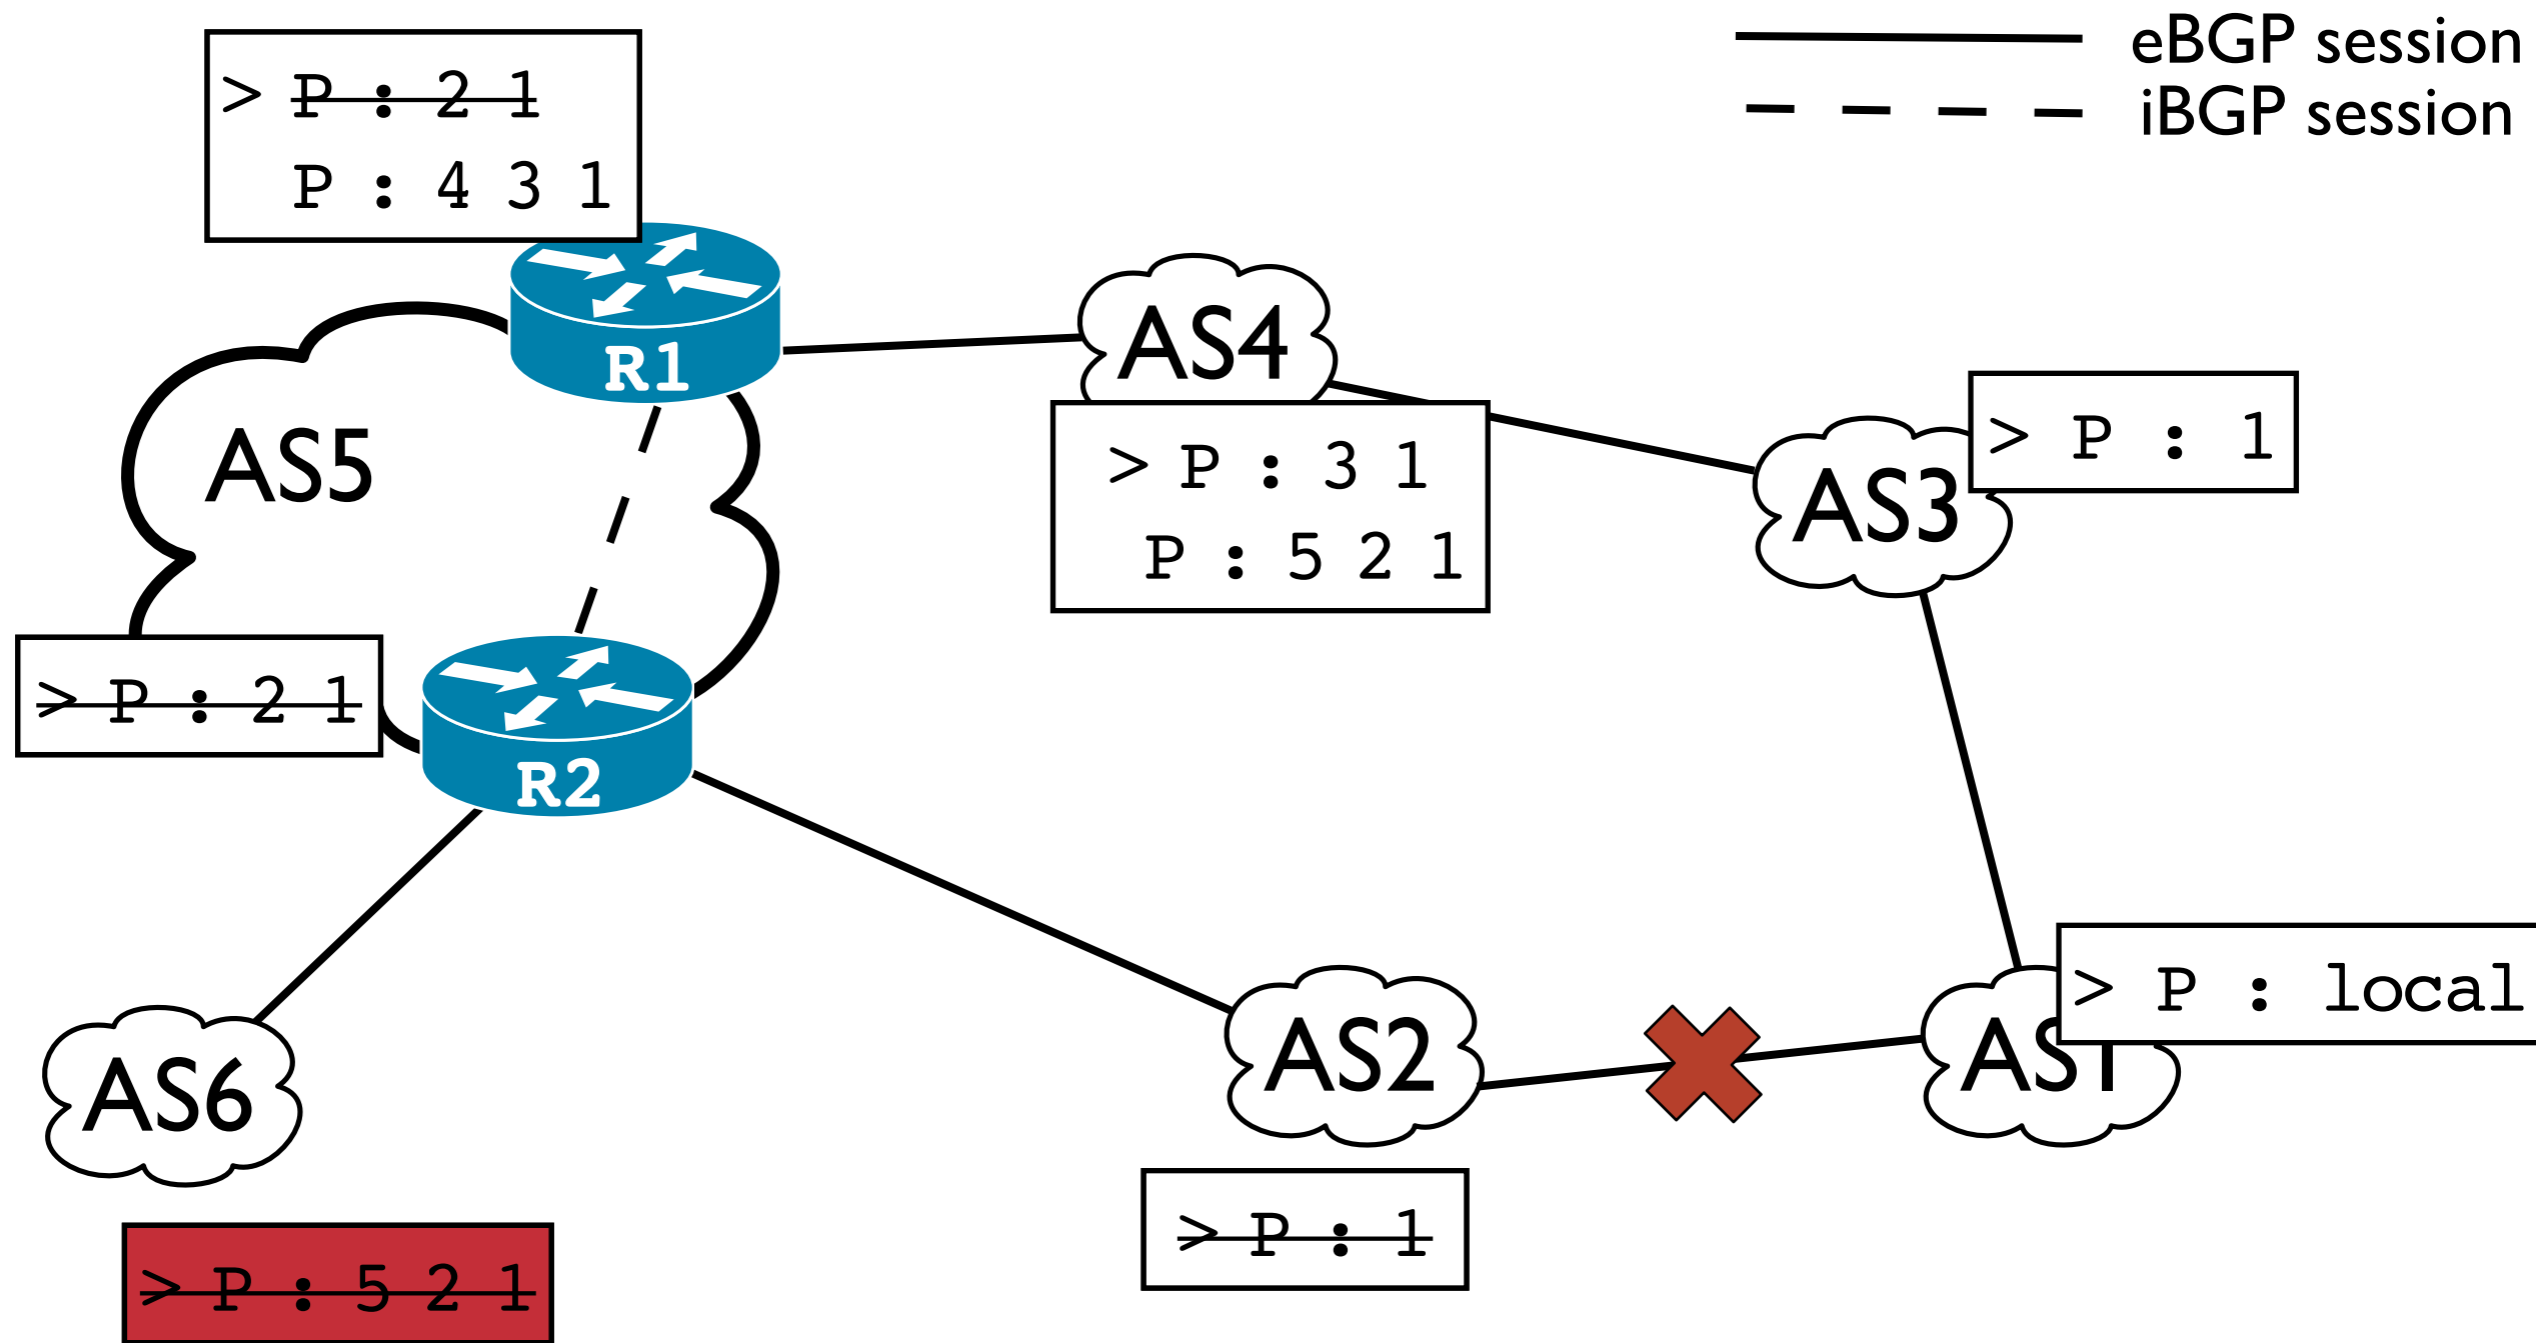
Withdraw Propagation when alternate path exists



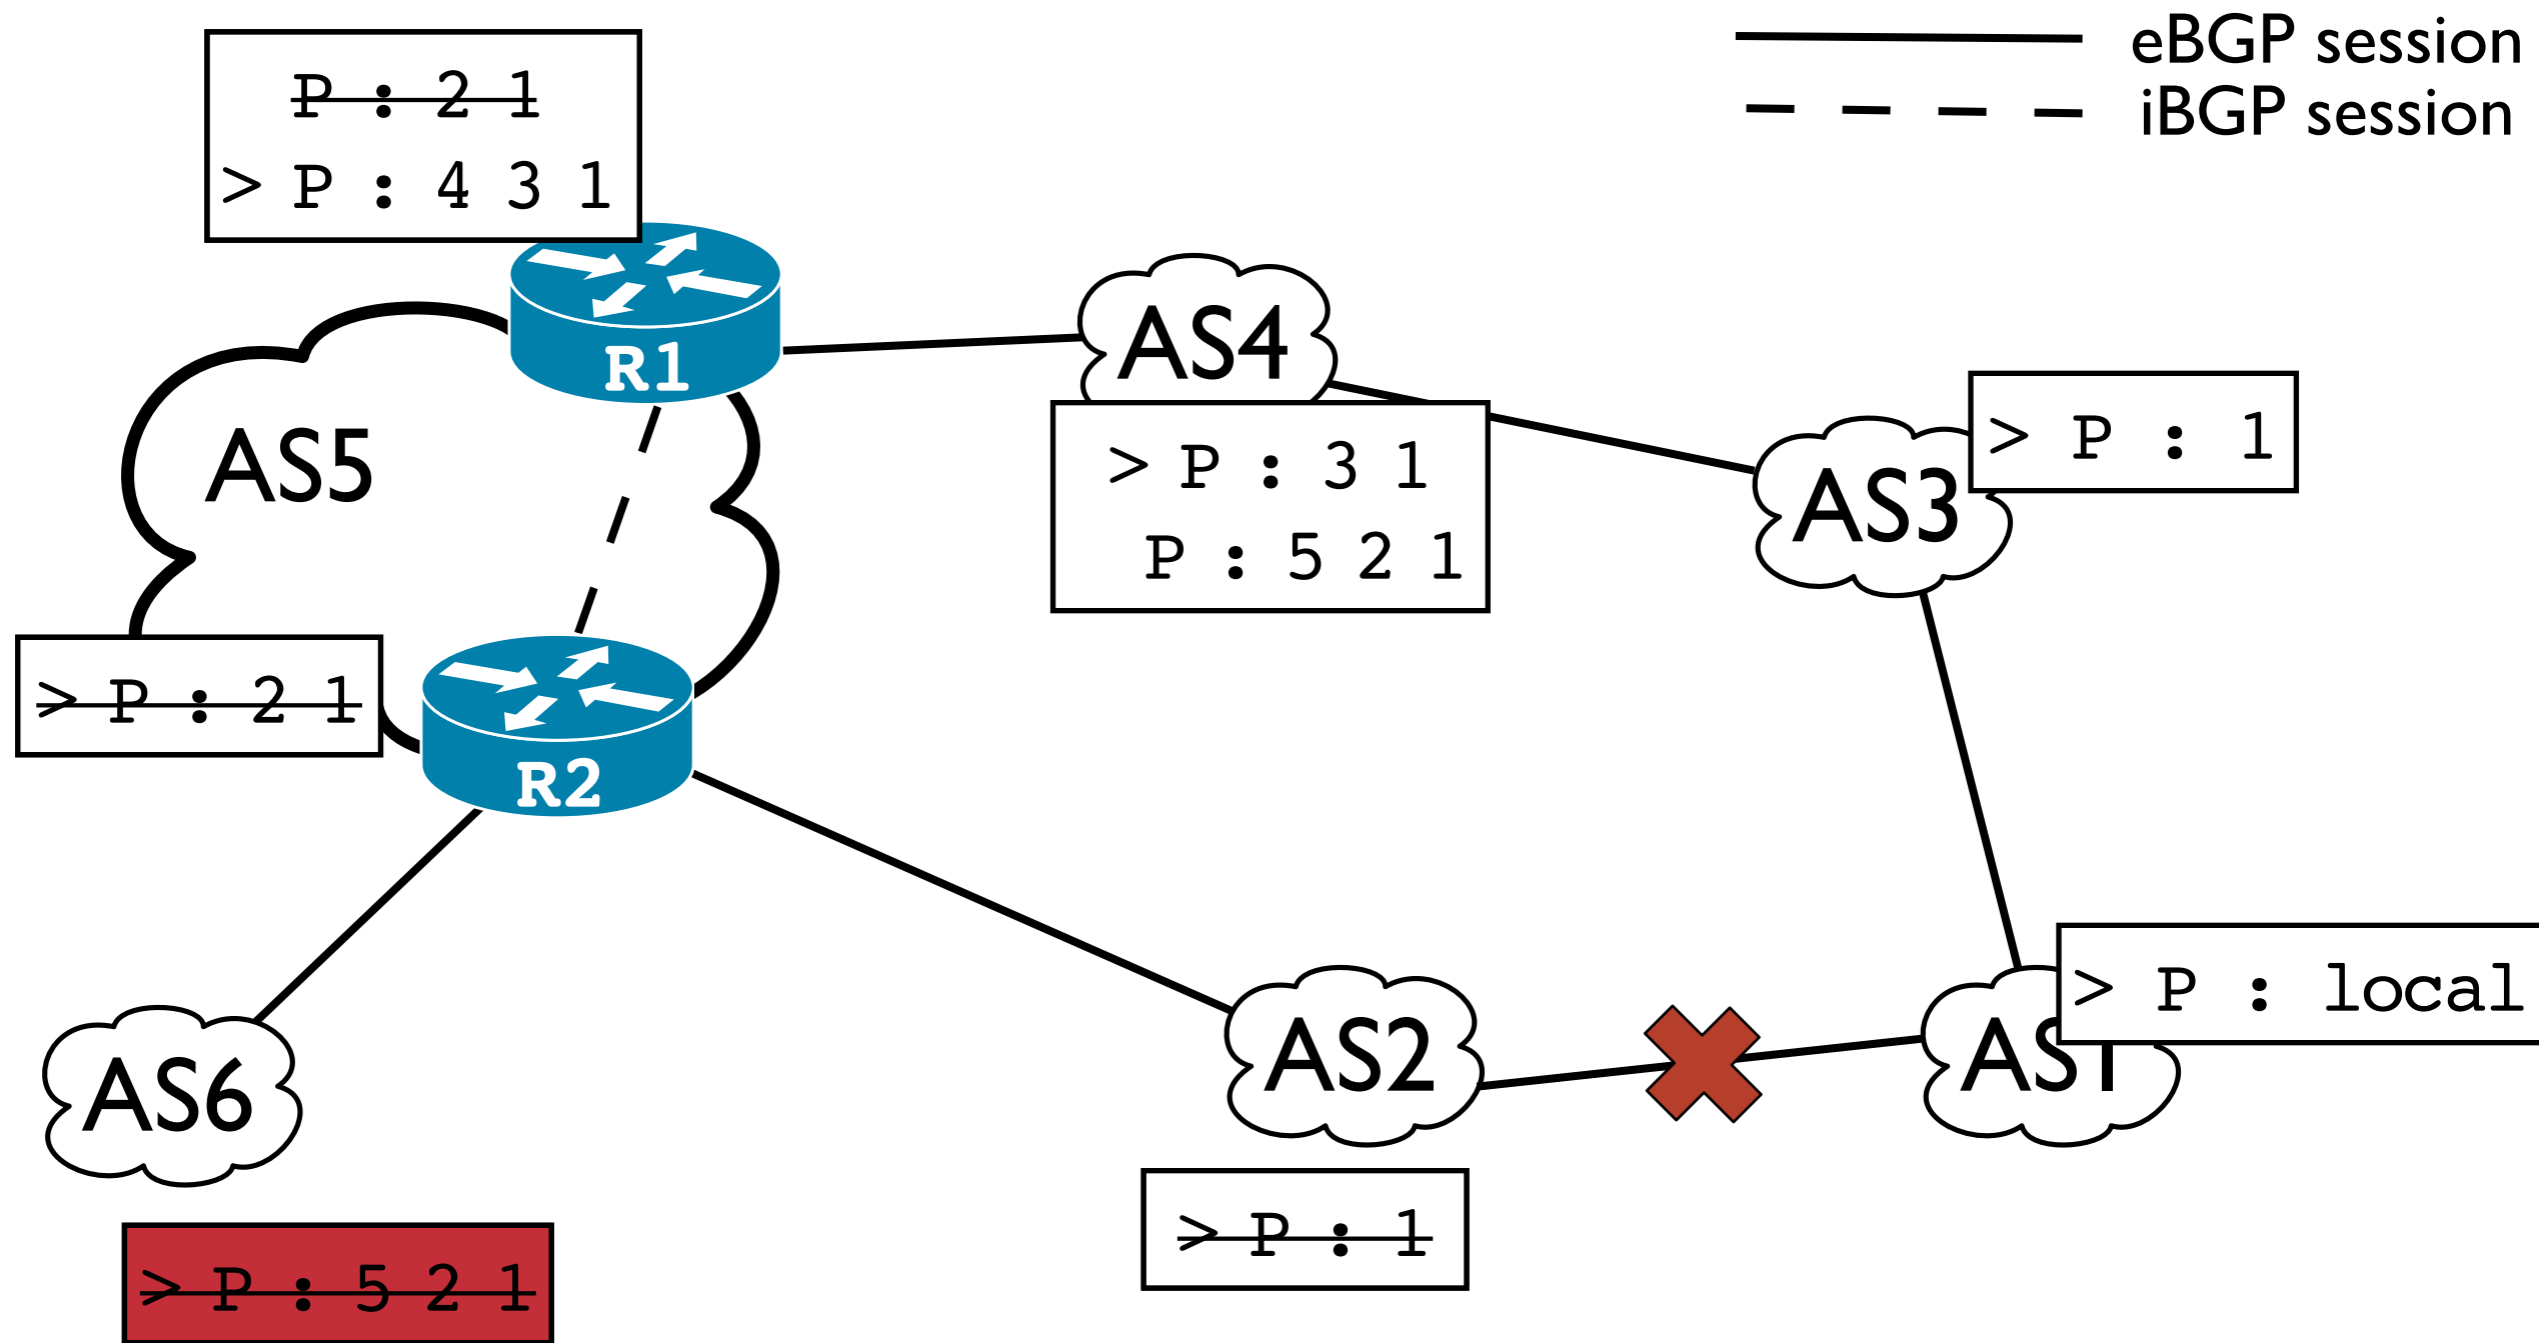
Withdraw Propagation when alternate path exists



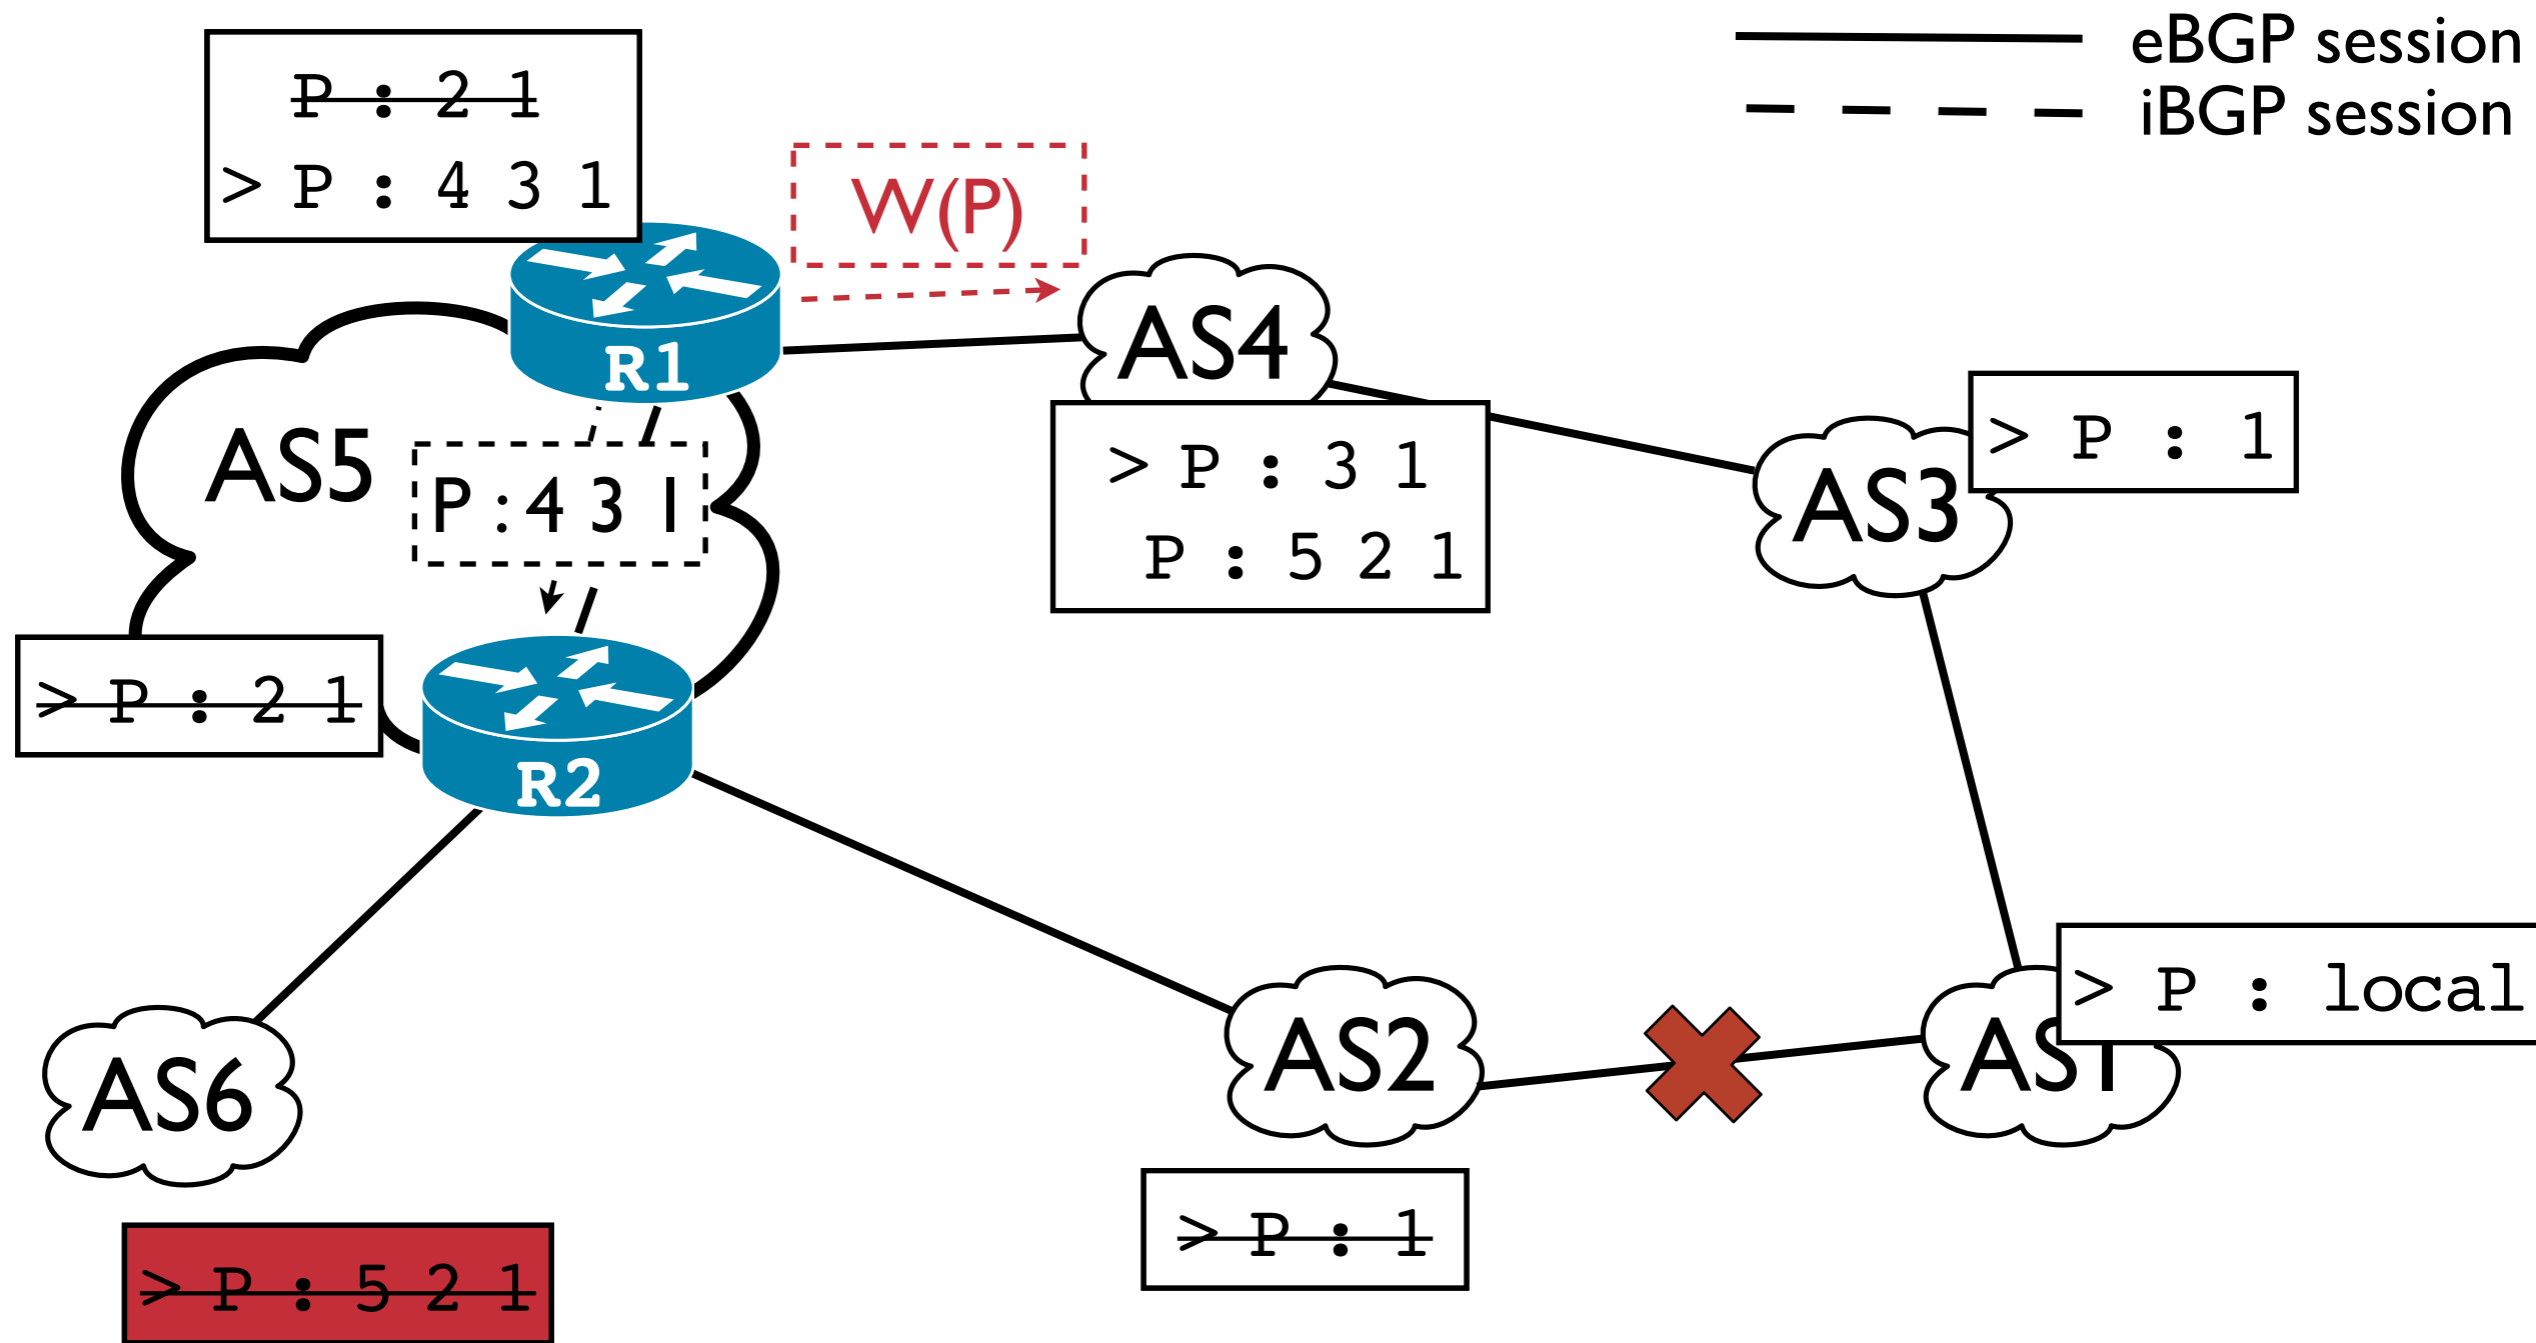
Withdraw Propagation when alternate path exists



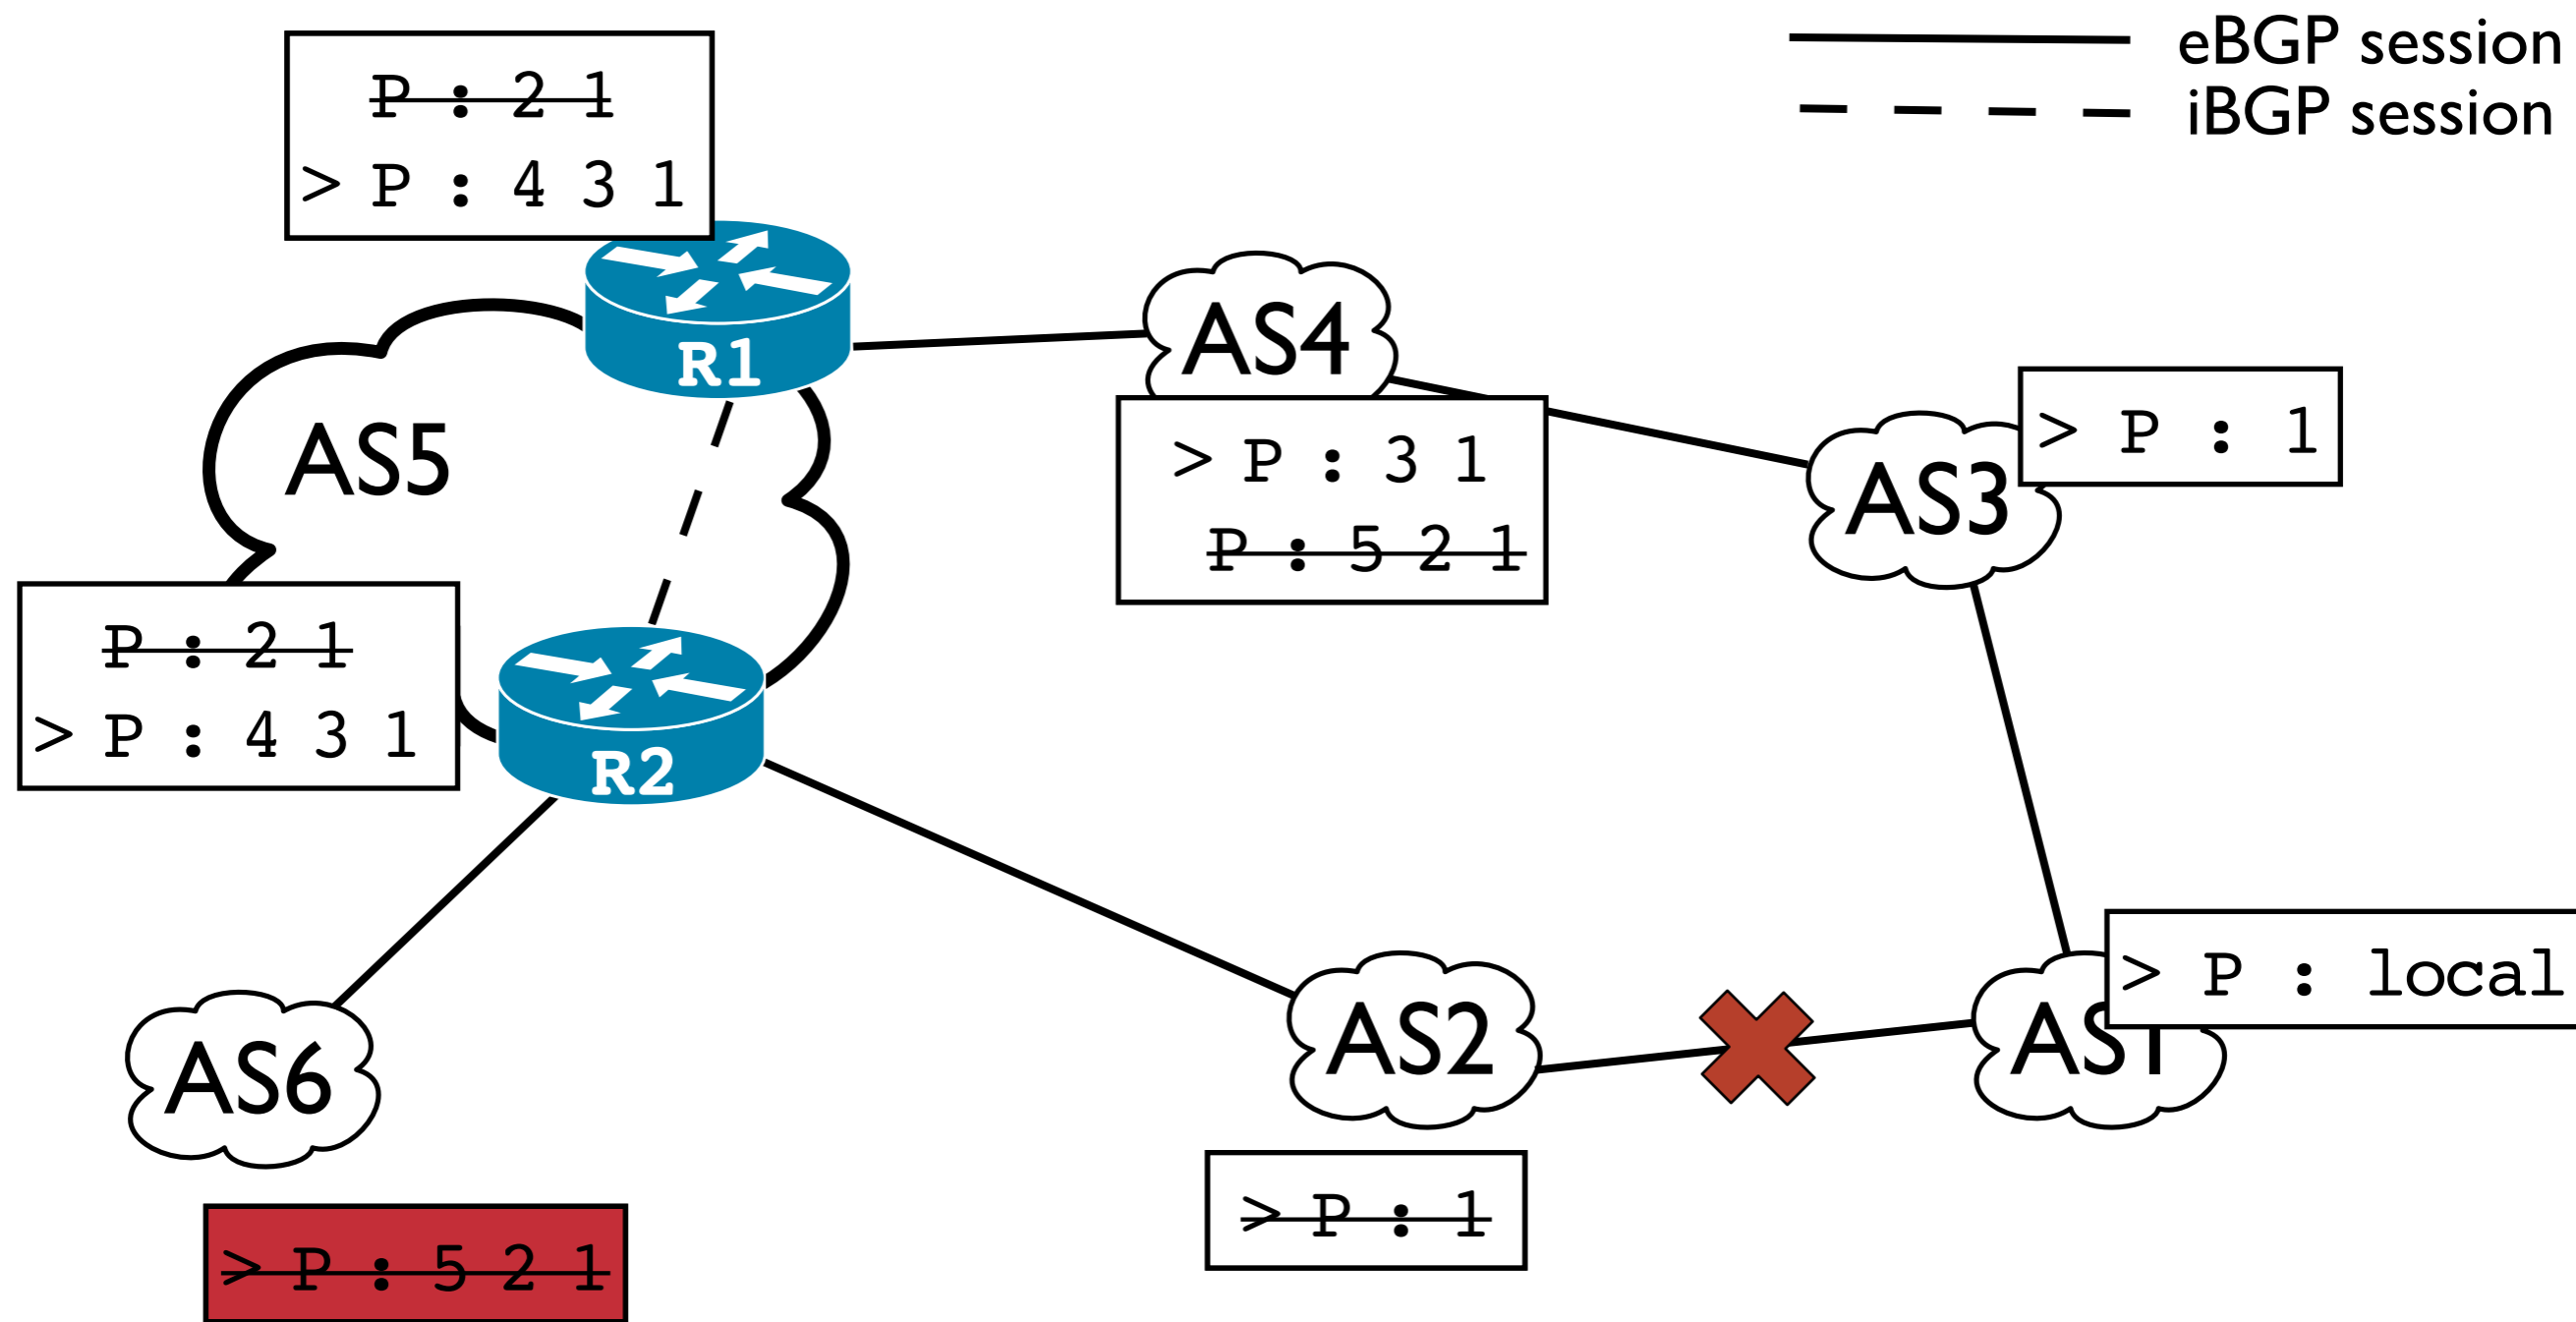
Withdraw Propagation when alternate path exists



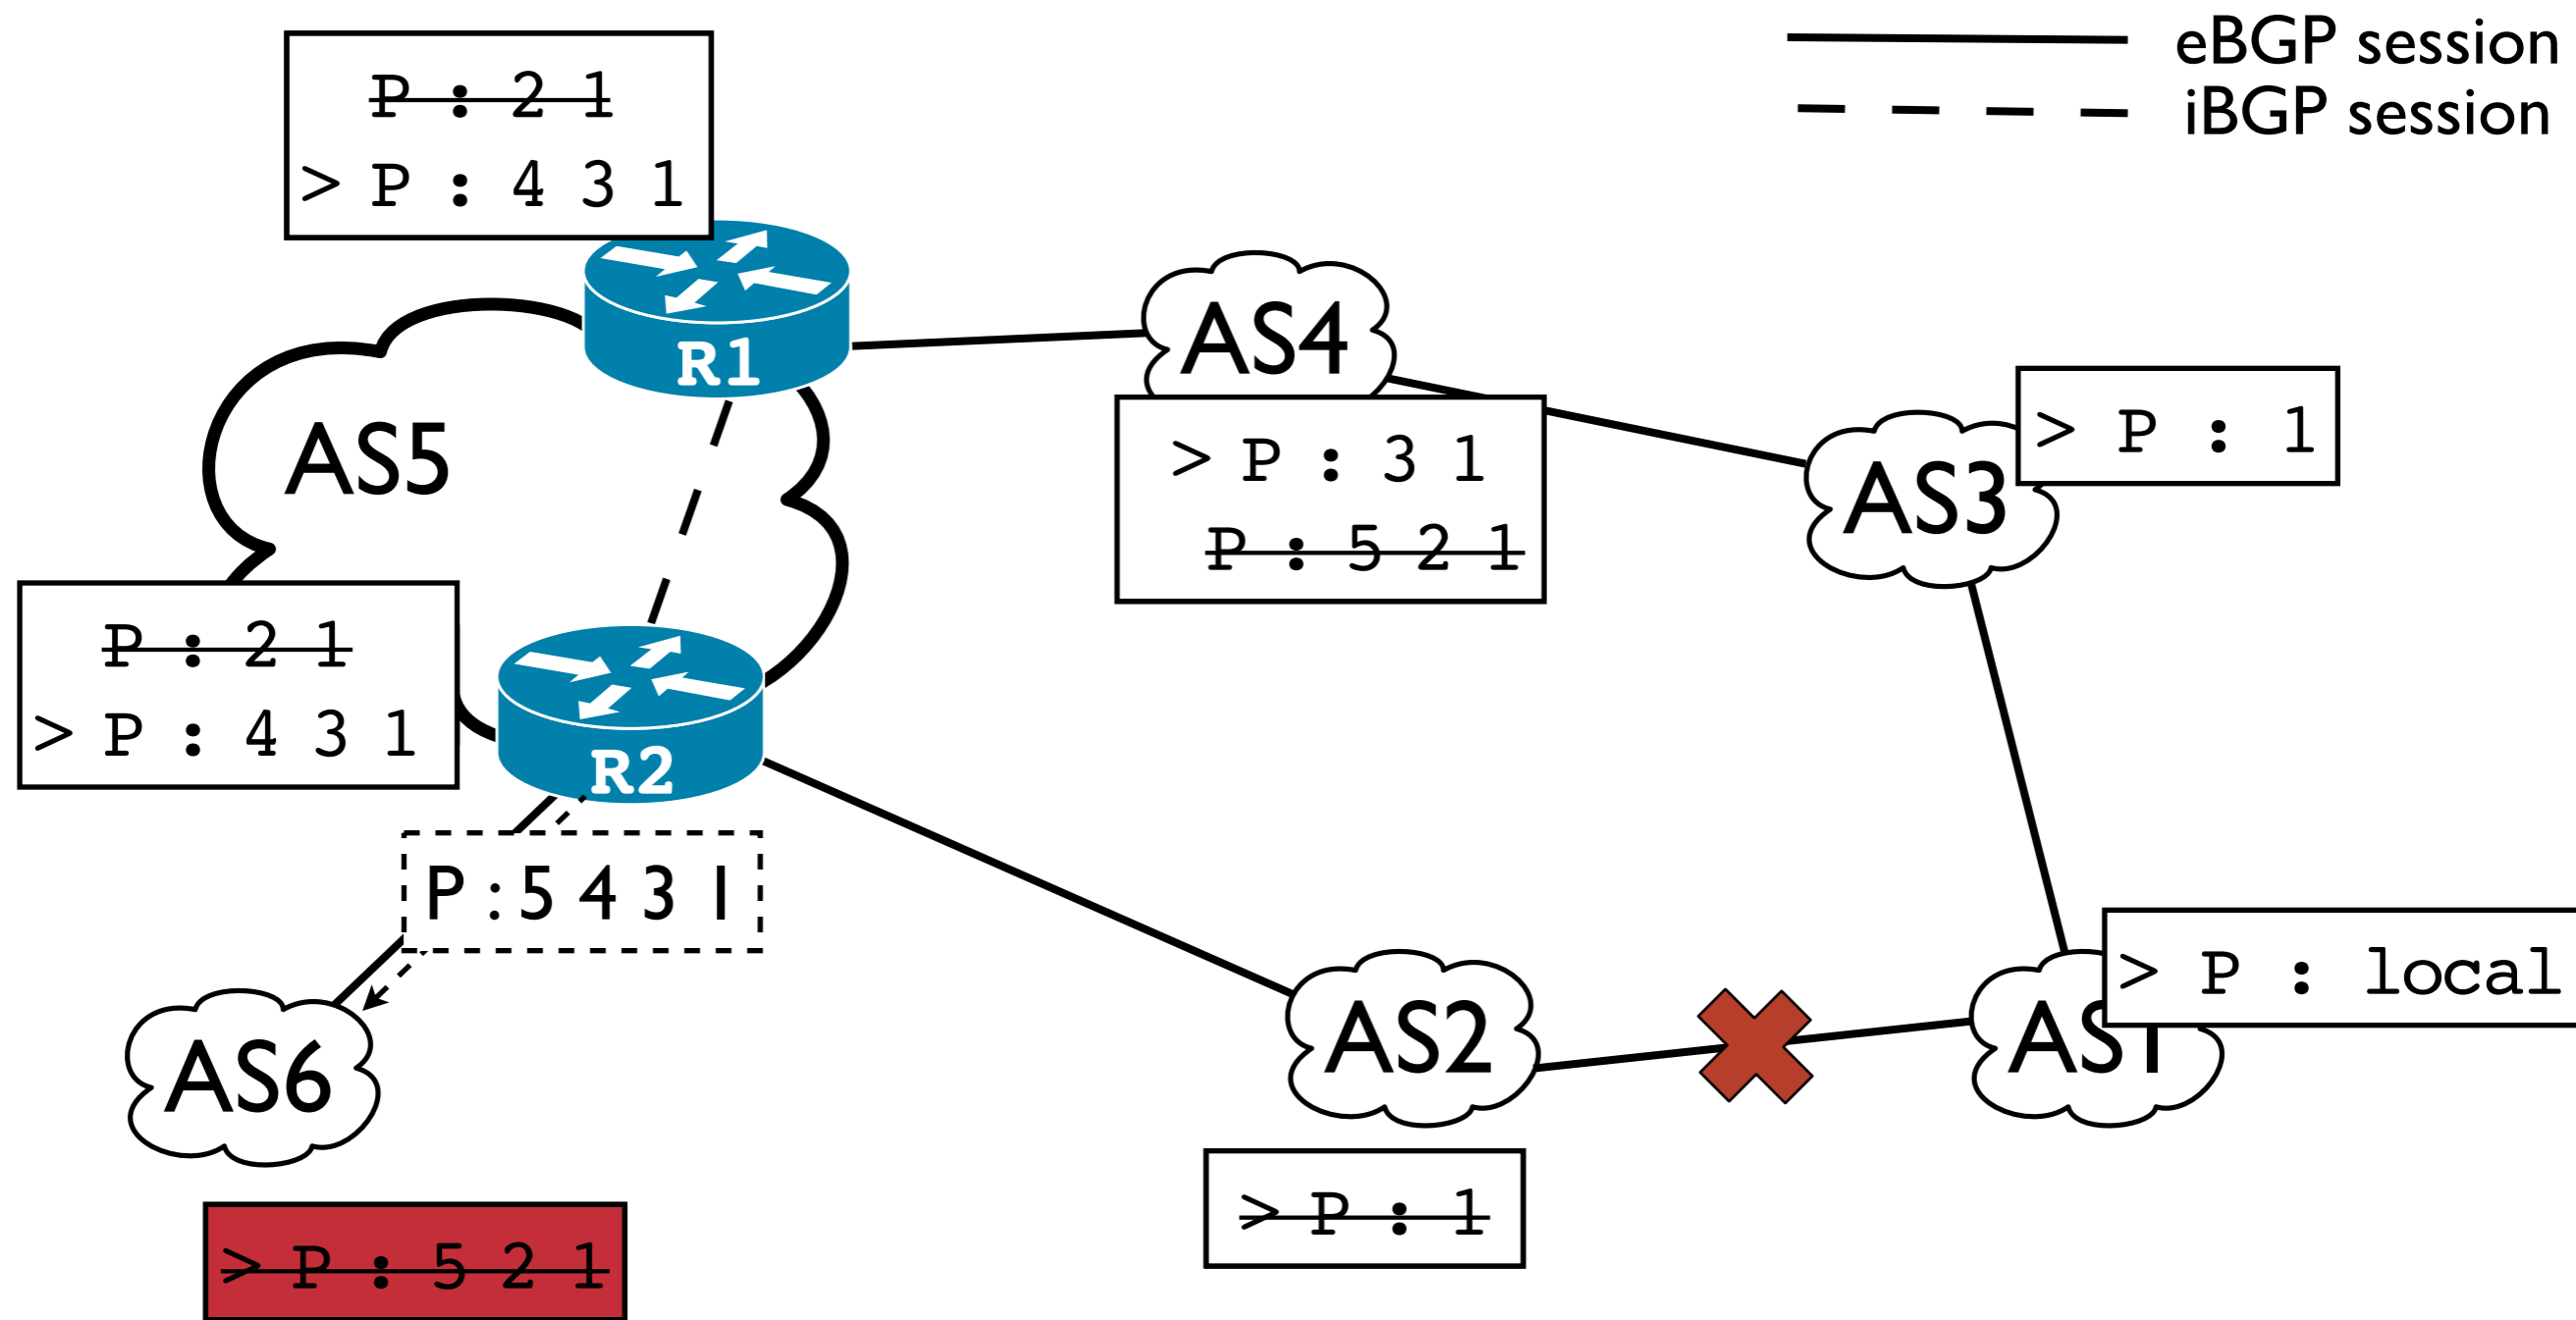
Withdraw Propagation when alternate path exists



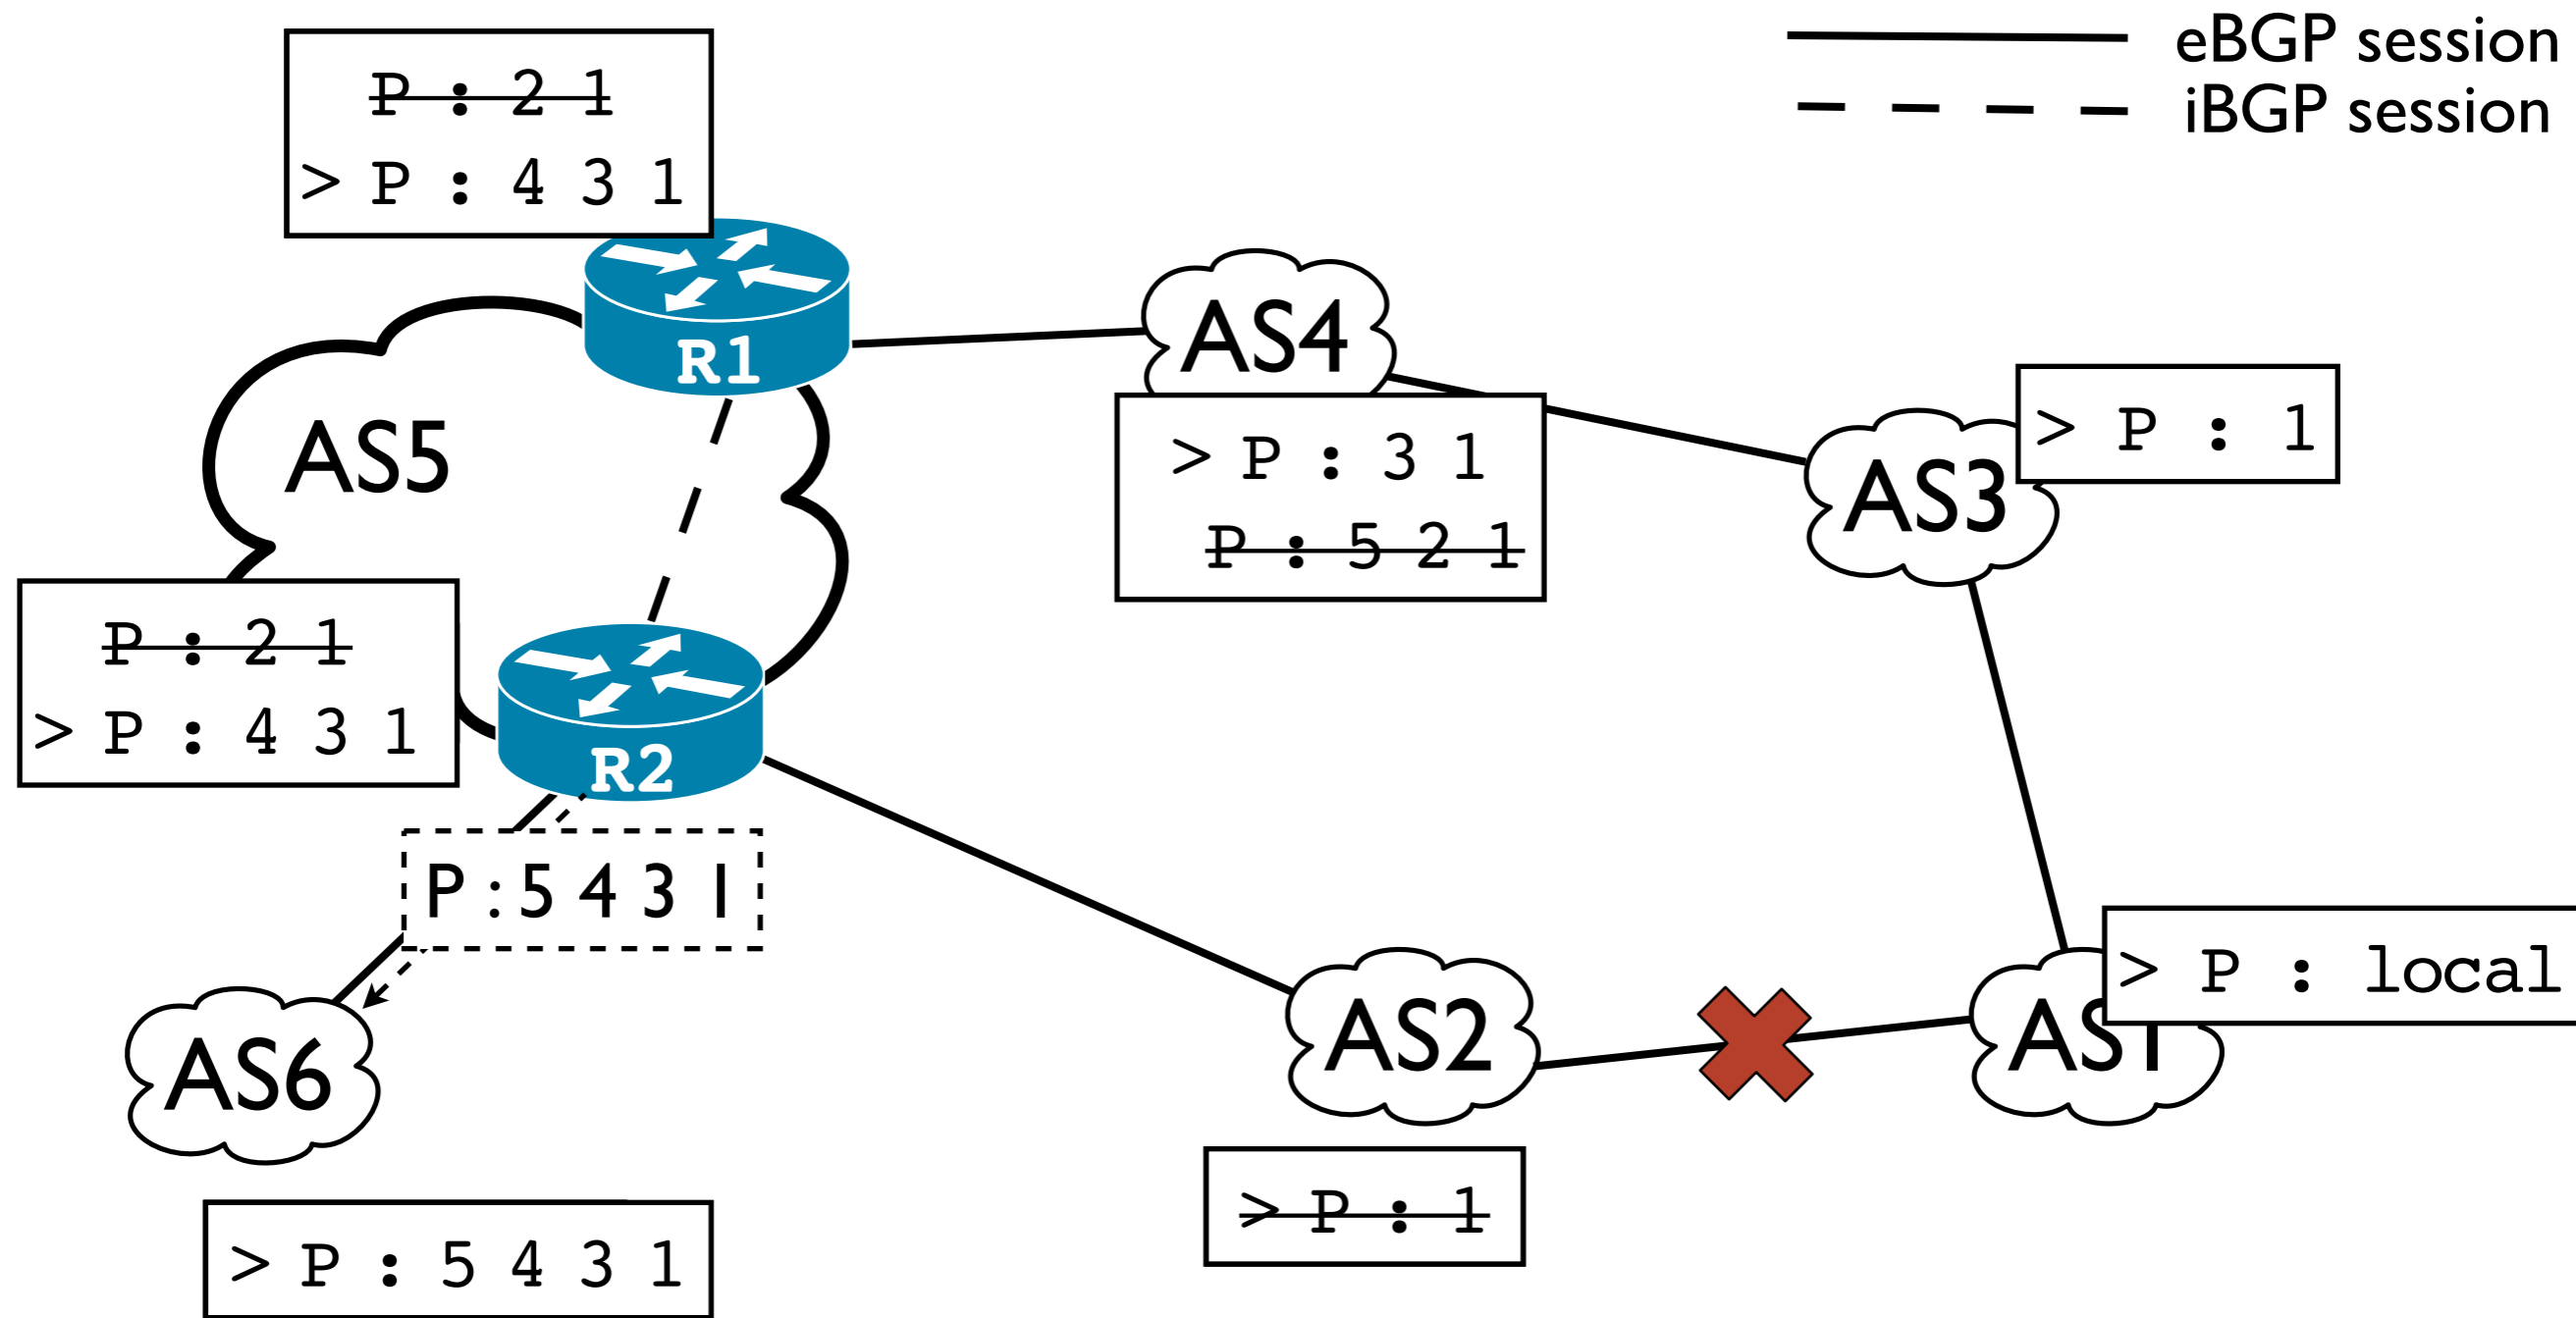
Withdraw Propagation when alternate path exists



Withdraw Propagation when alternate path exists

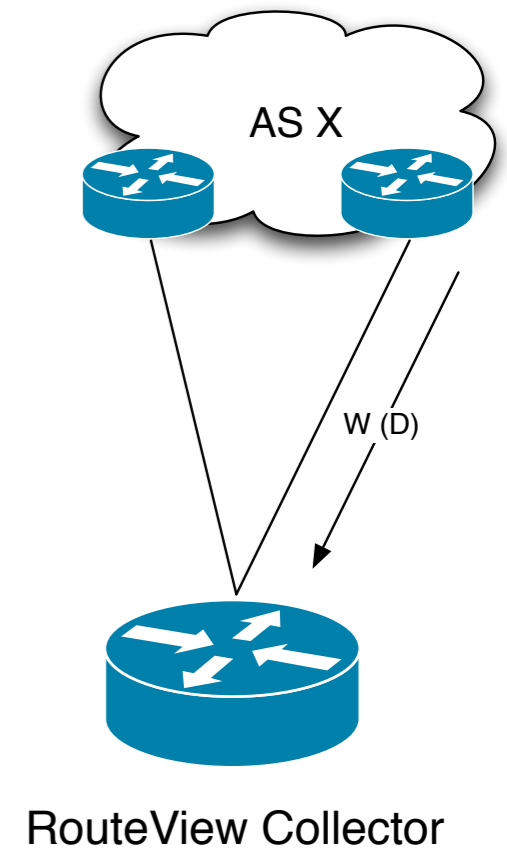


Withdraw Propagation when alternate path exists

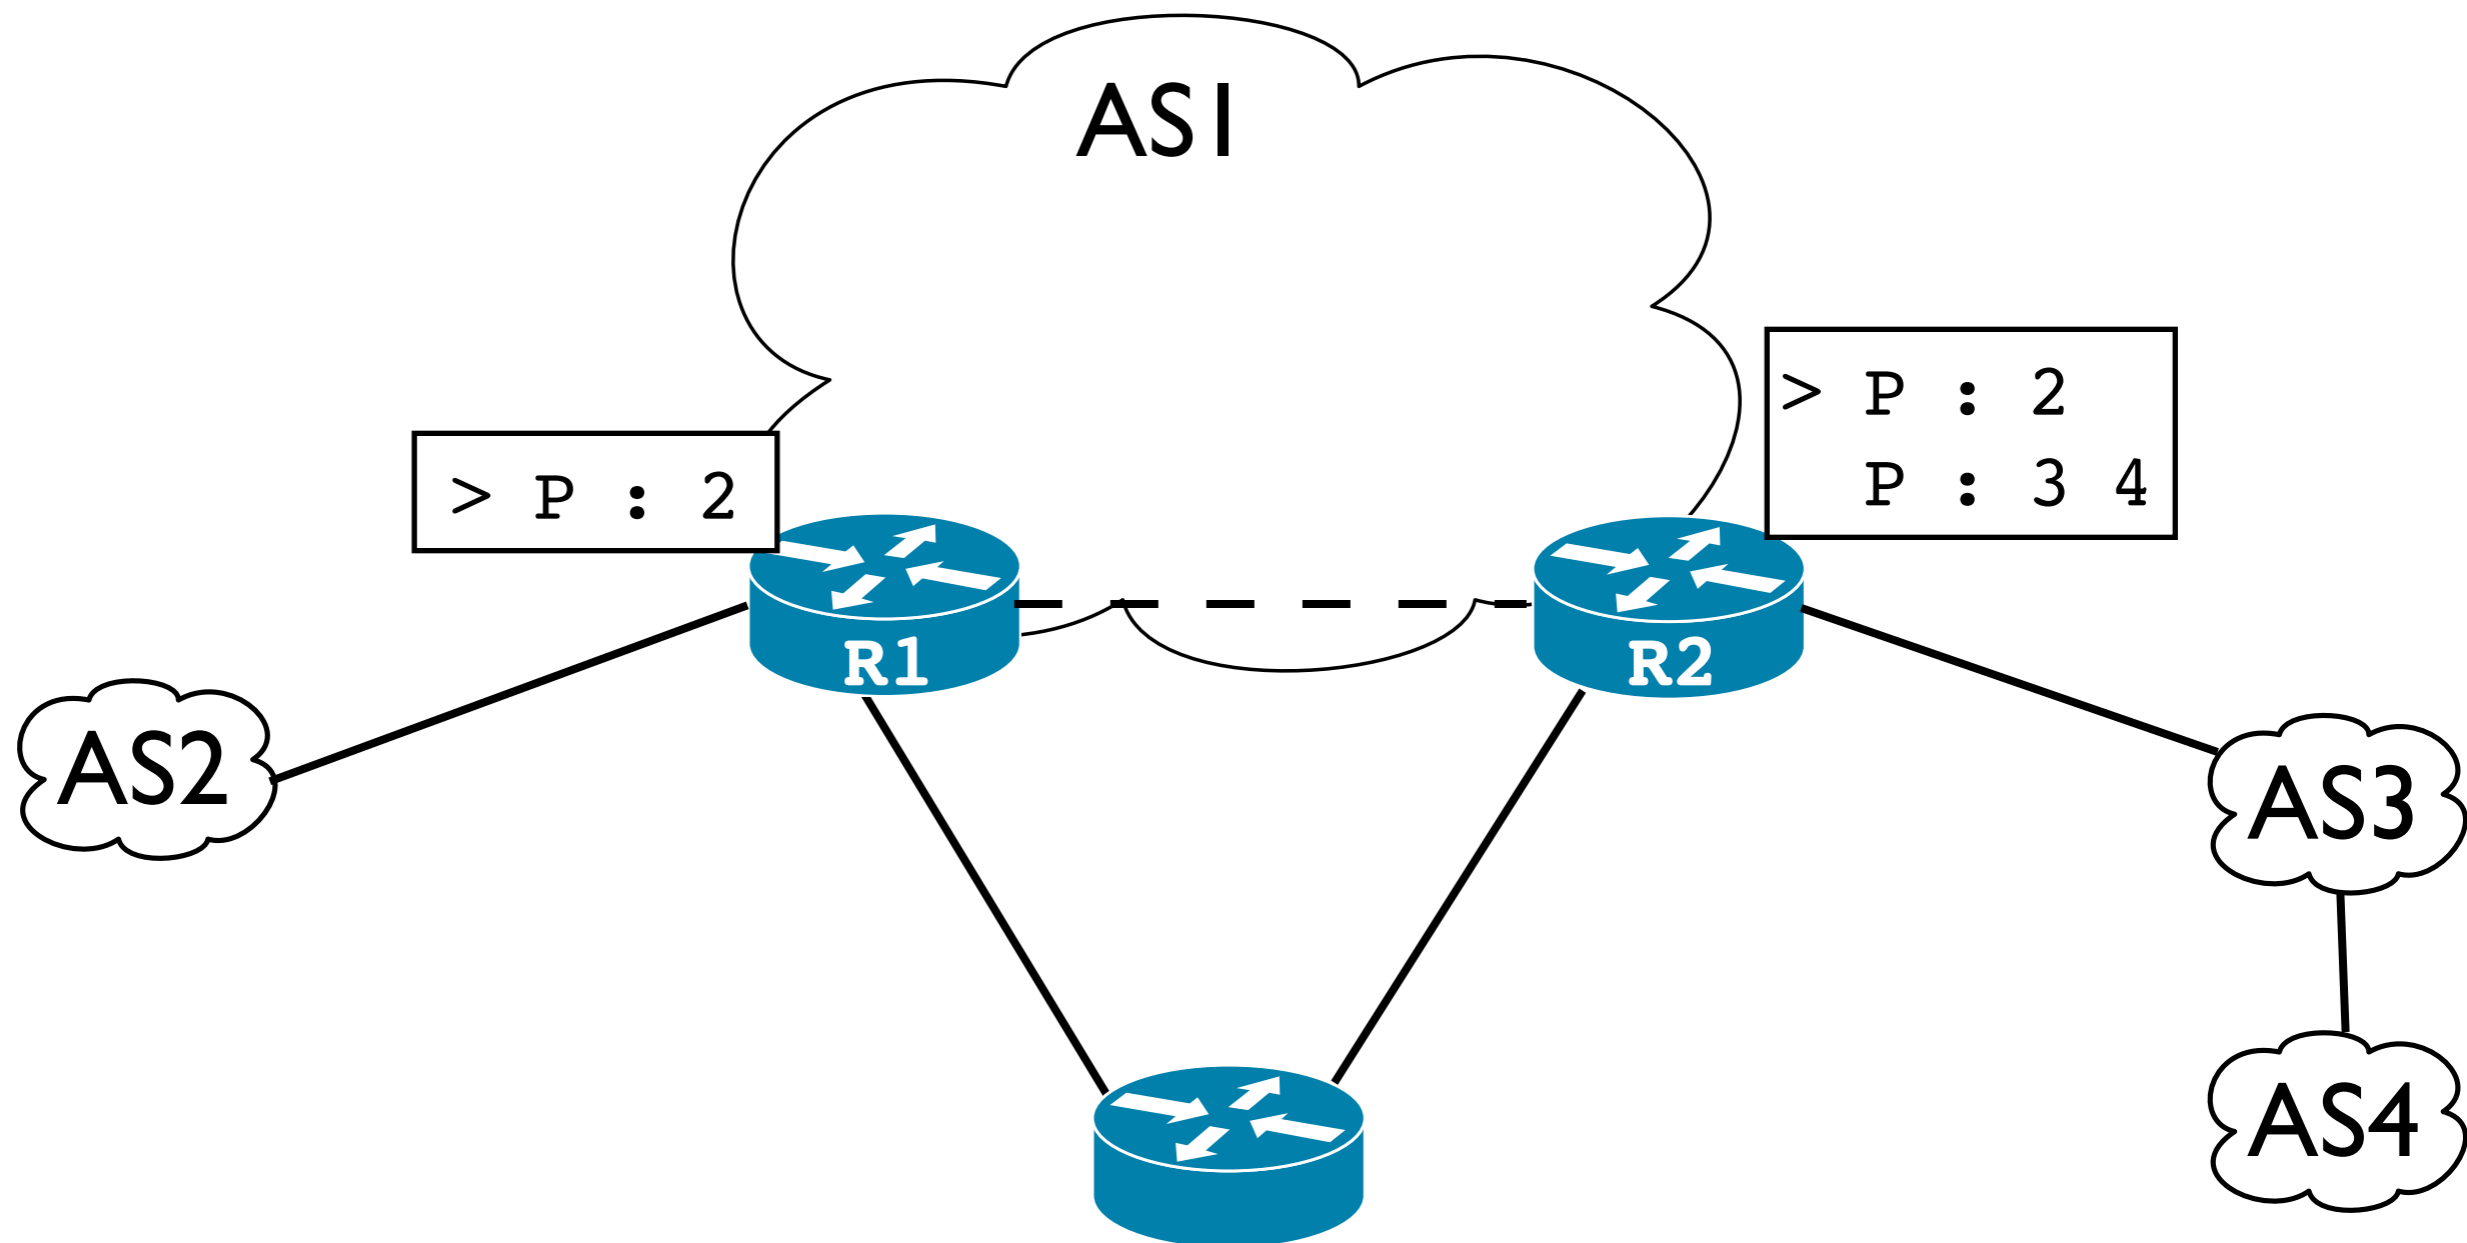


Quantification of unnecessary iBGP-caused Withdraws

- **RouteViews Dataset** : BGP feeds of parallel sessions (2 sessions with the same AS)
- Withdraw defined as iBGP-caused if the destination is withdrawn on one session, but stay stable on the other during 30 seconds before and after
- Withdraw is unnecessary because an alternate path exists in the AS




Detecting unnecessary iBGP-caused withdraws




RouteView Collector

Detecting unnecessary iBGP-caused withdraws




RouteView Collector

Detecting unnecessary iBGP-caused withdraws



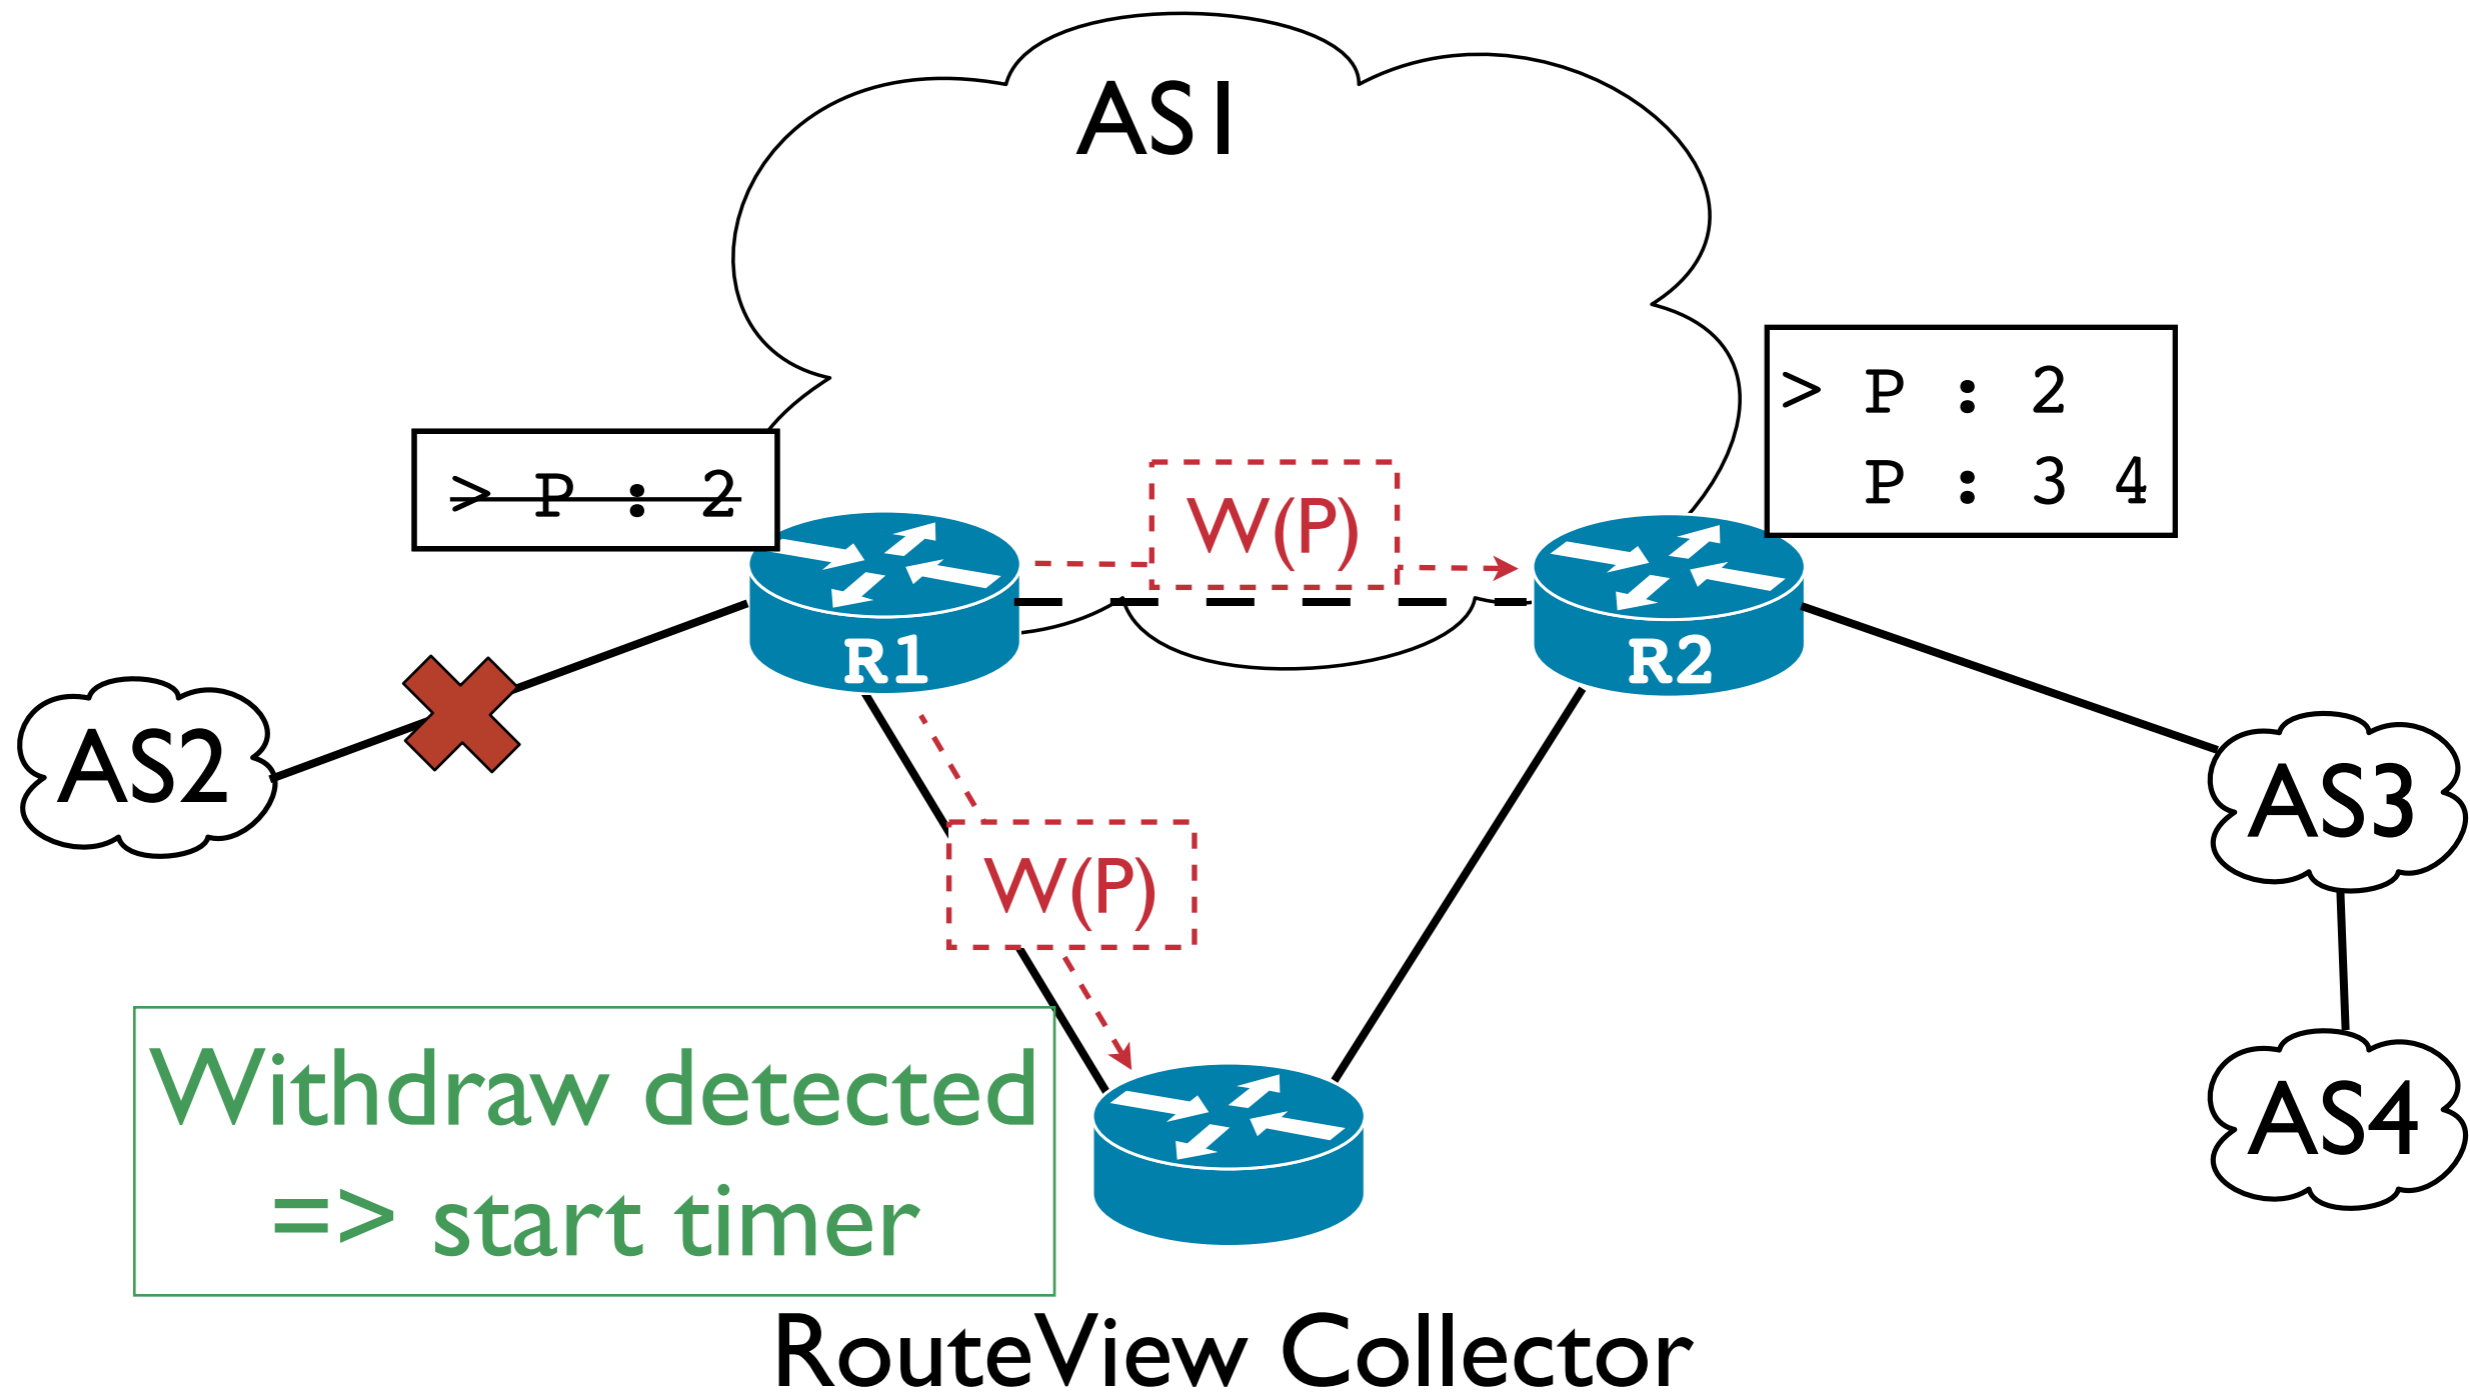
RouteView Collector

Detecting unnecessary iBGP-caused withdraws




RouteView Collector

Detecting unnecessary iBGP-caused withdraws

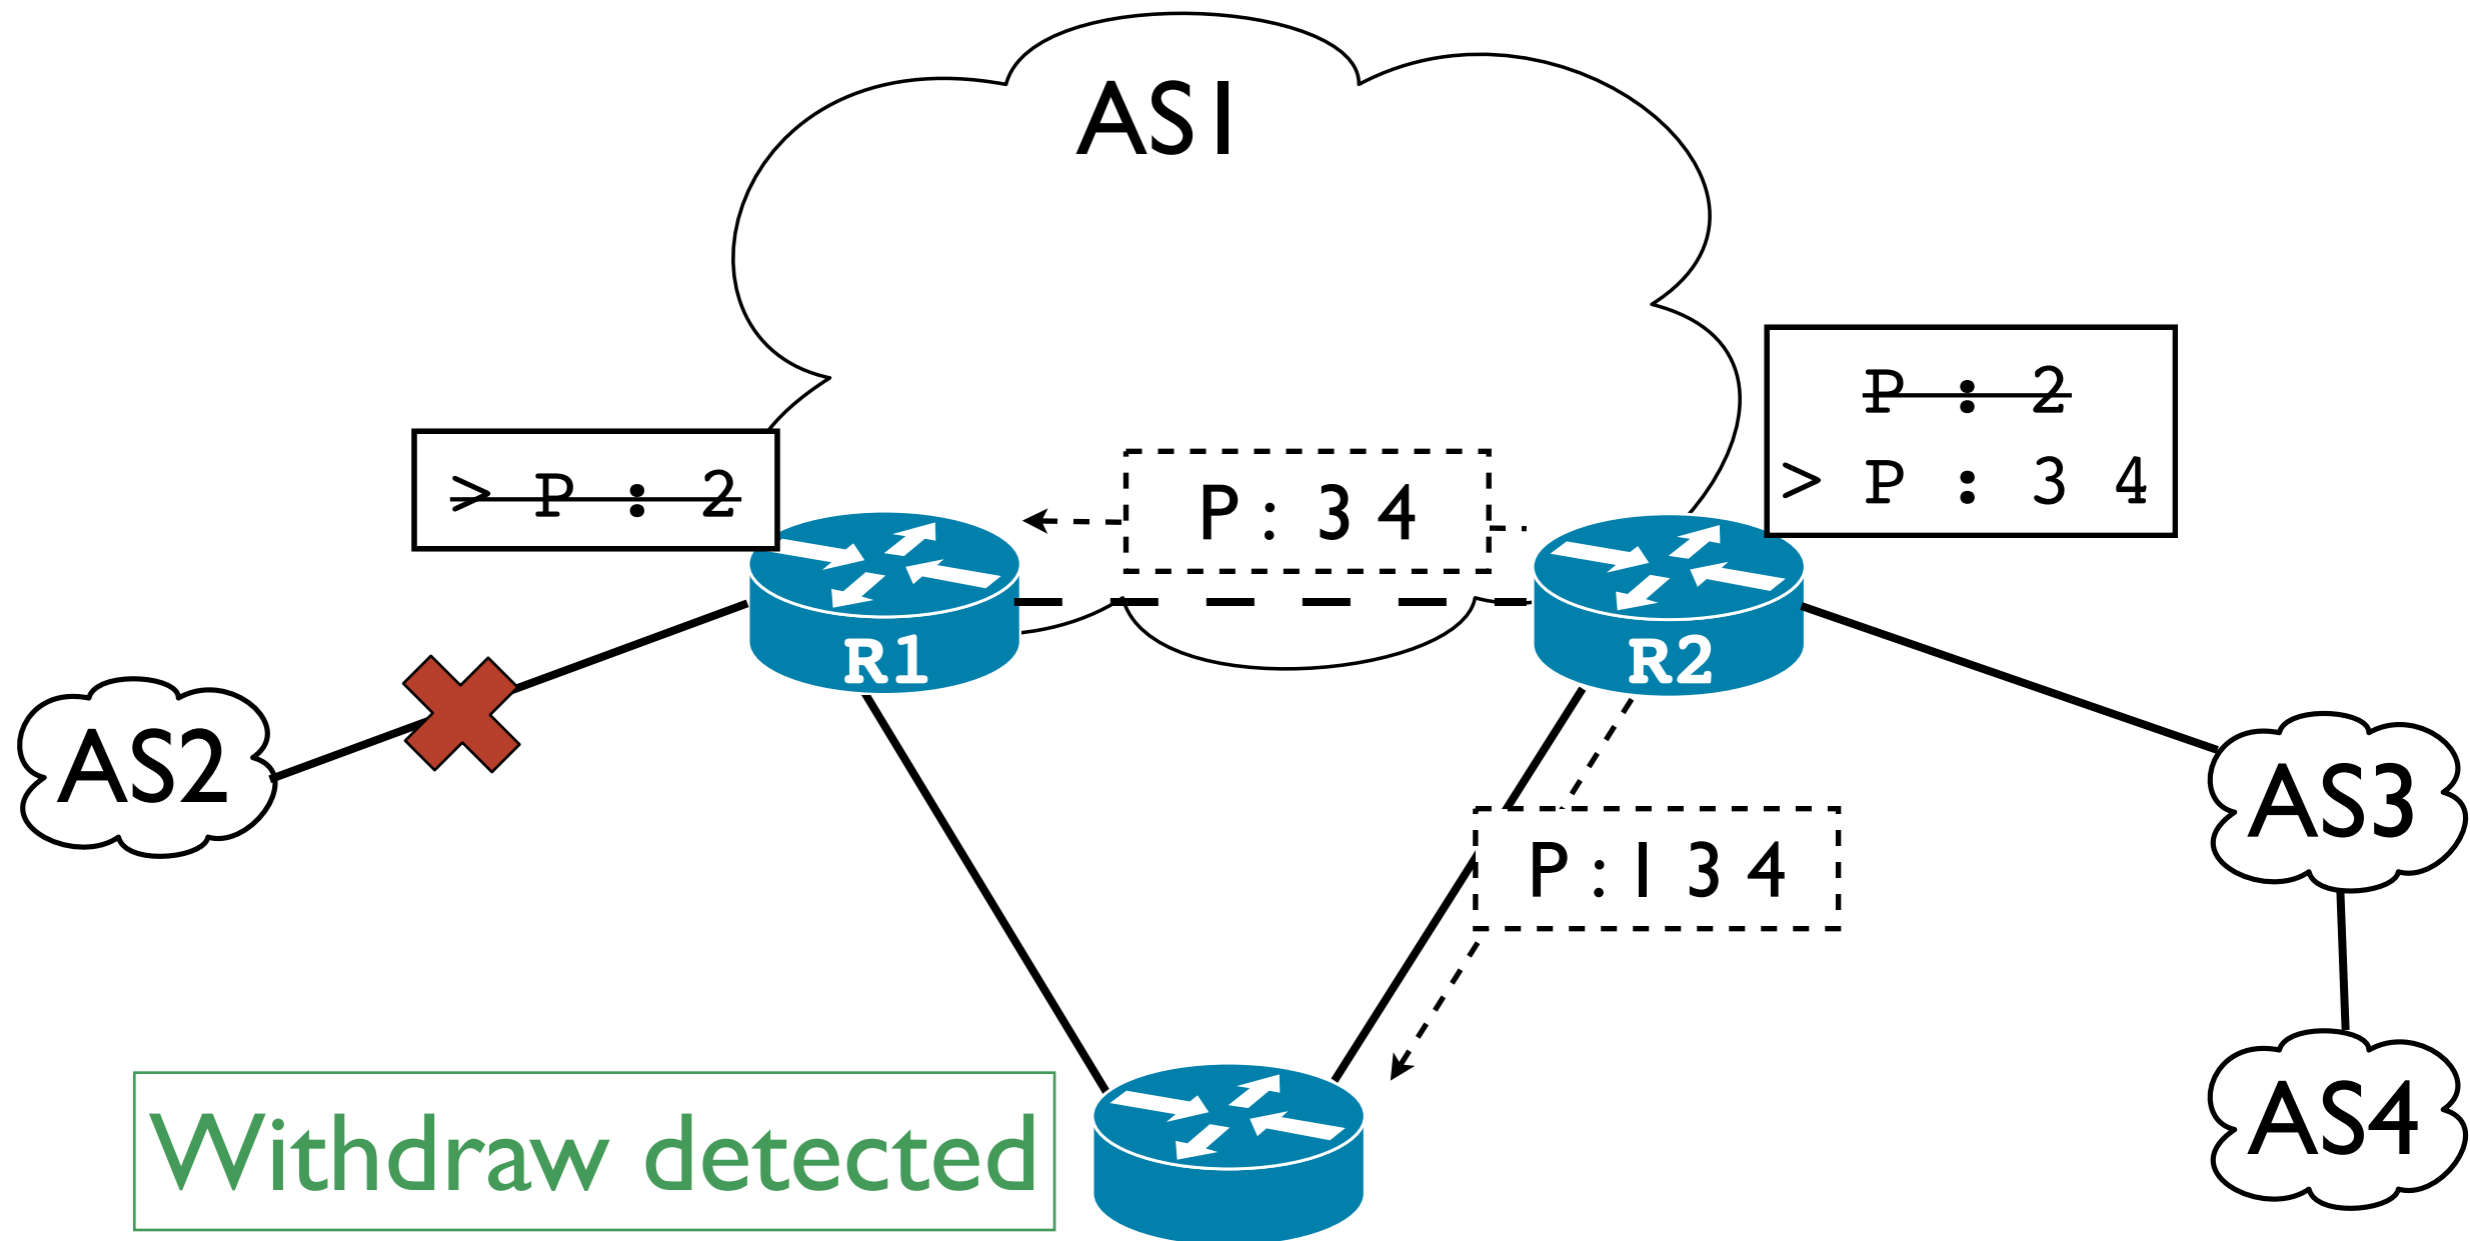


Detecting unnecessary iBGP-caused withdraws




RouteView Collector

Detecting unnecessary iBGP-caused withdraws




RouteView Collector

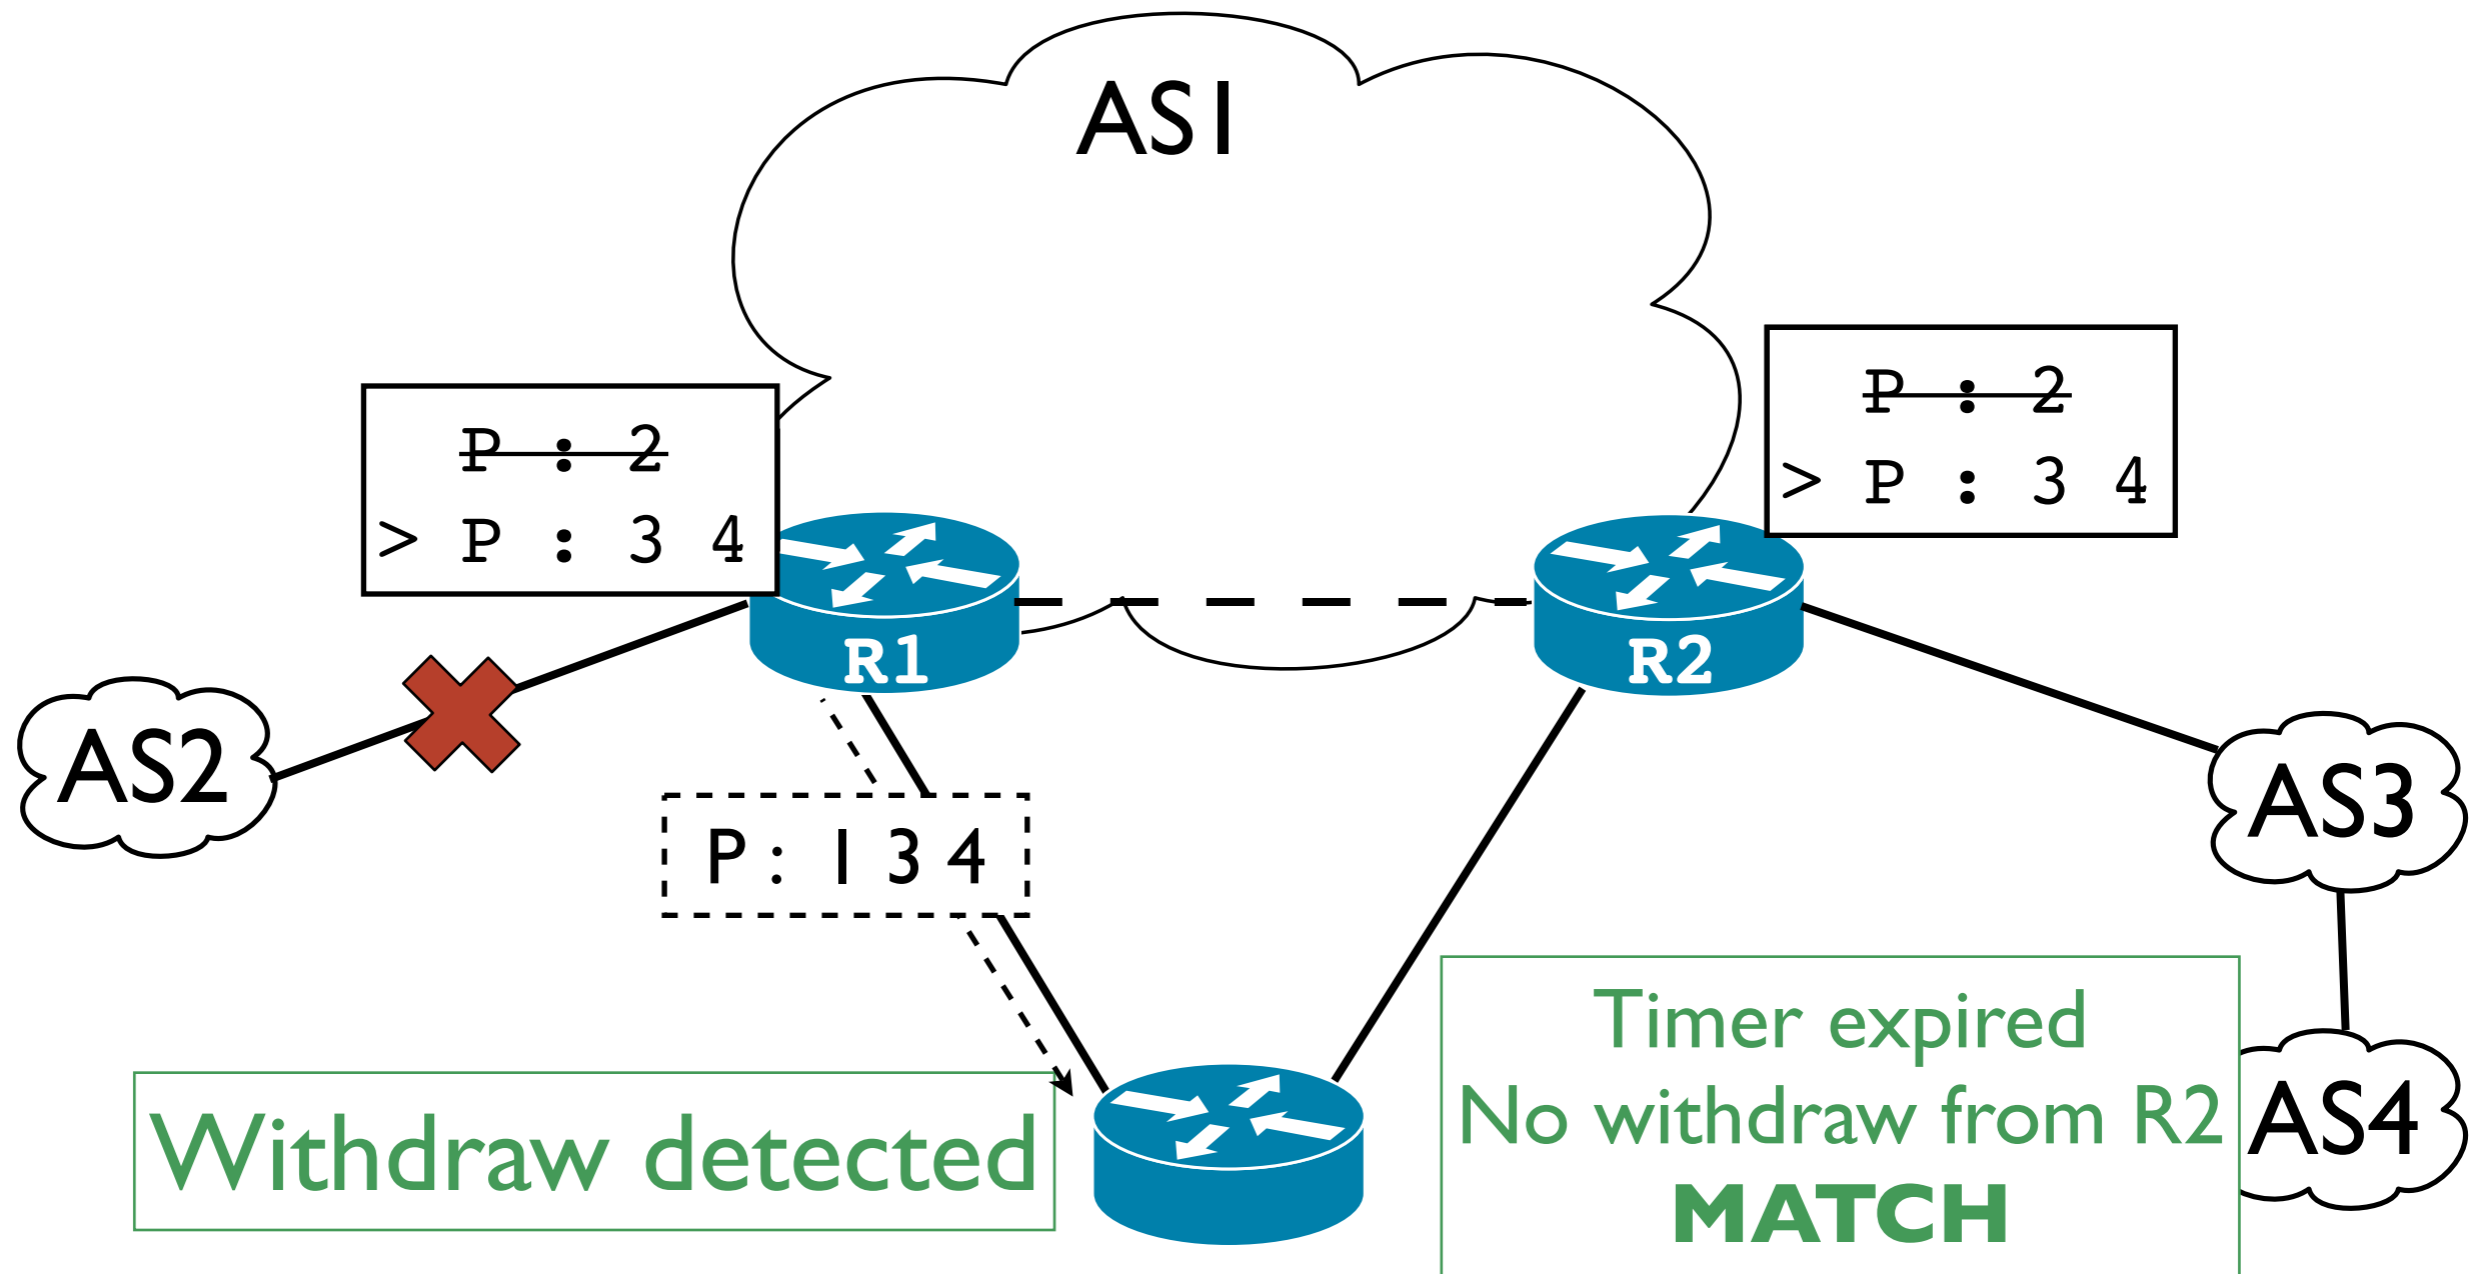
Detecting unnecessary iBGP-caused withdraws



Detecting unnecessary iBGP-caused withdraws

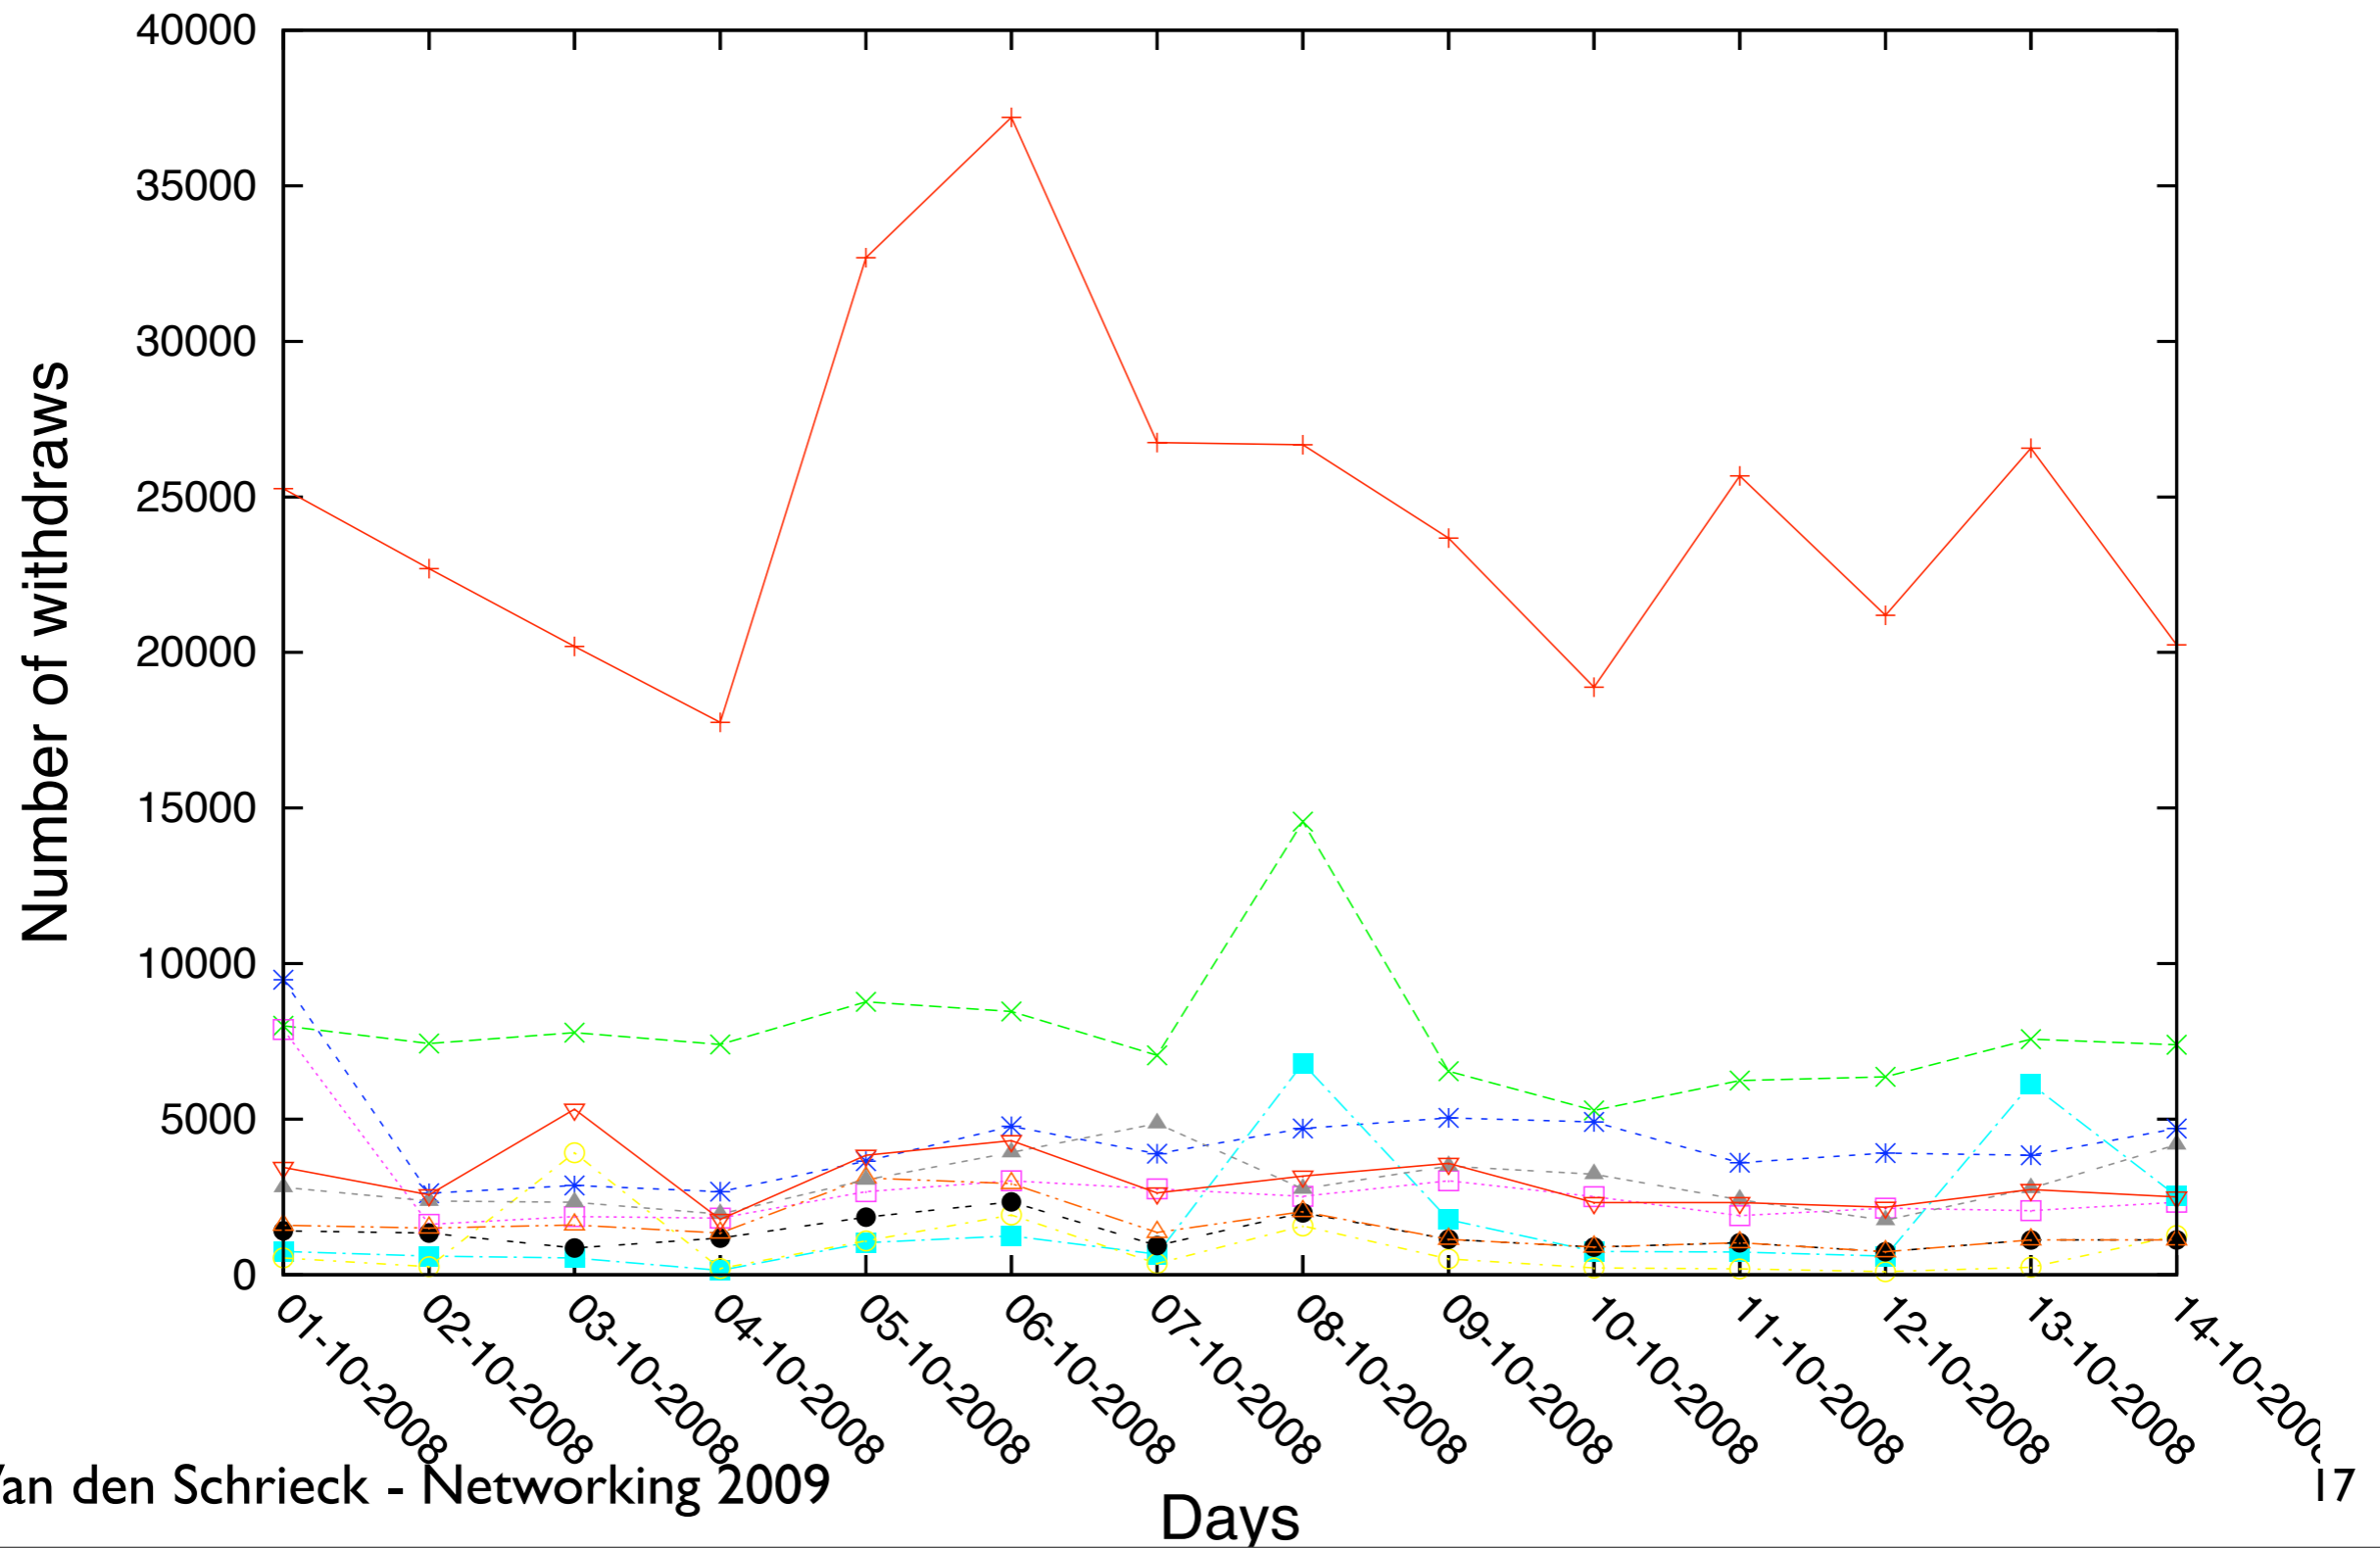


Detecting unnecessary iBGP-caused withdraws



RouteView Collector


Quantification of iBGP-caused withdraws on RouteViews sessions



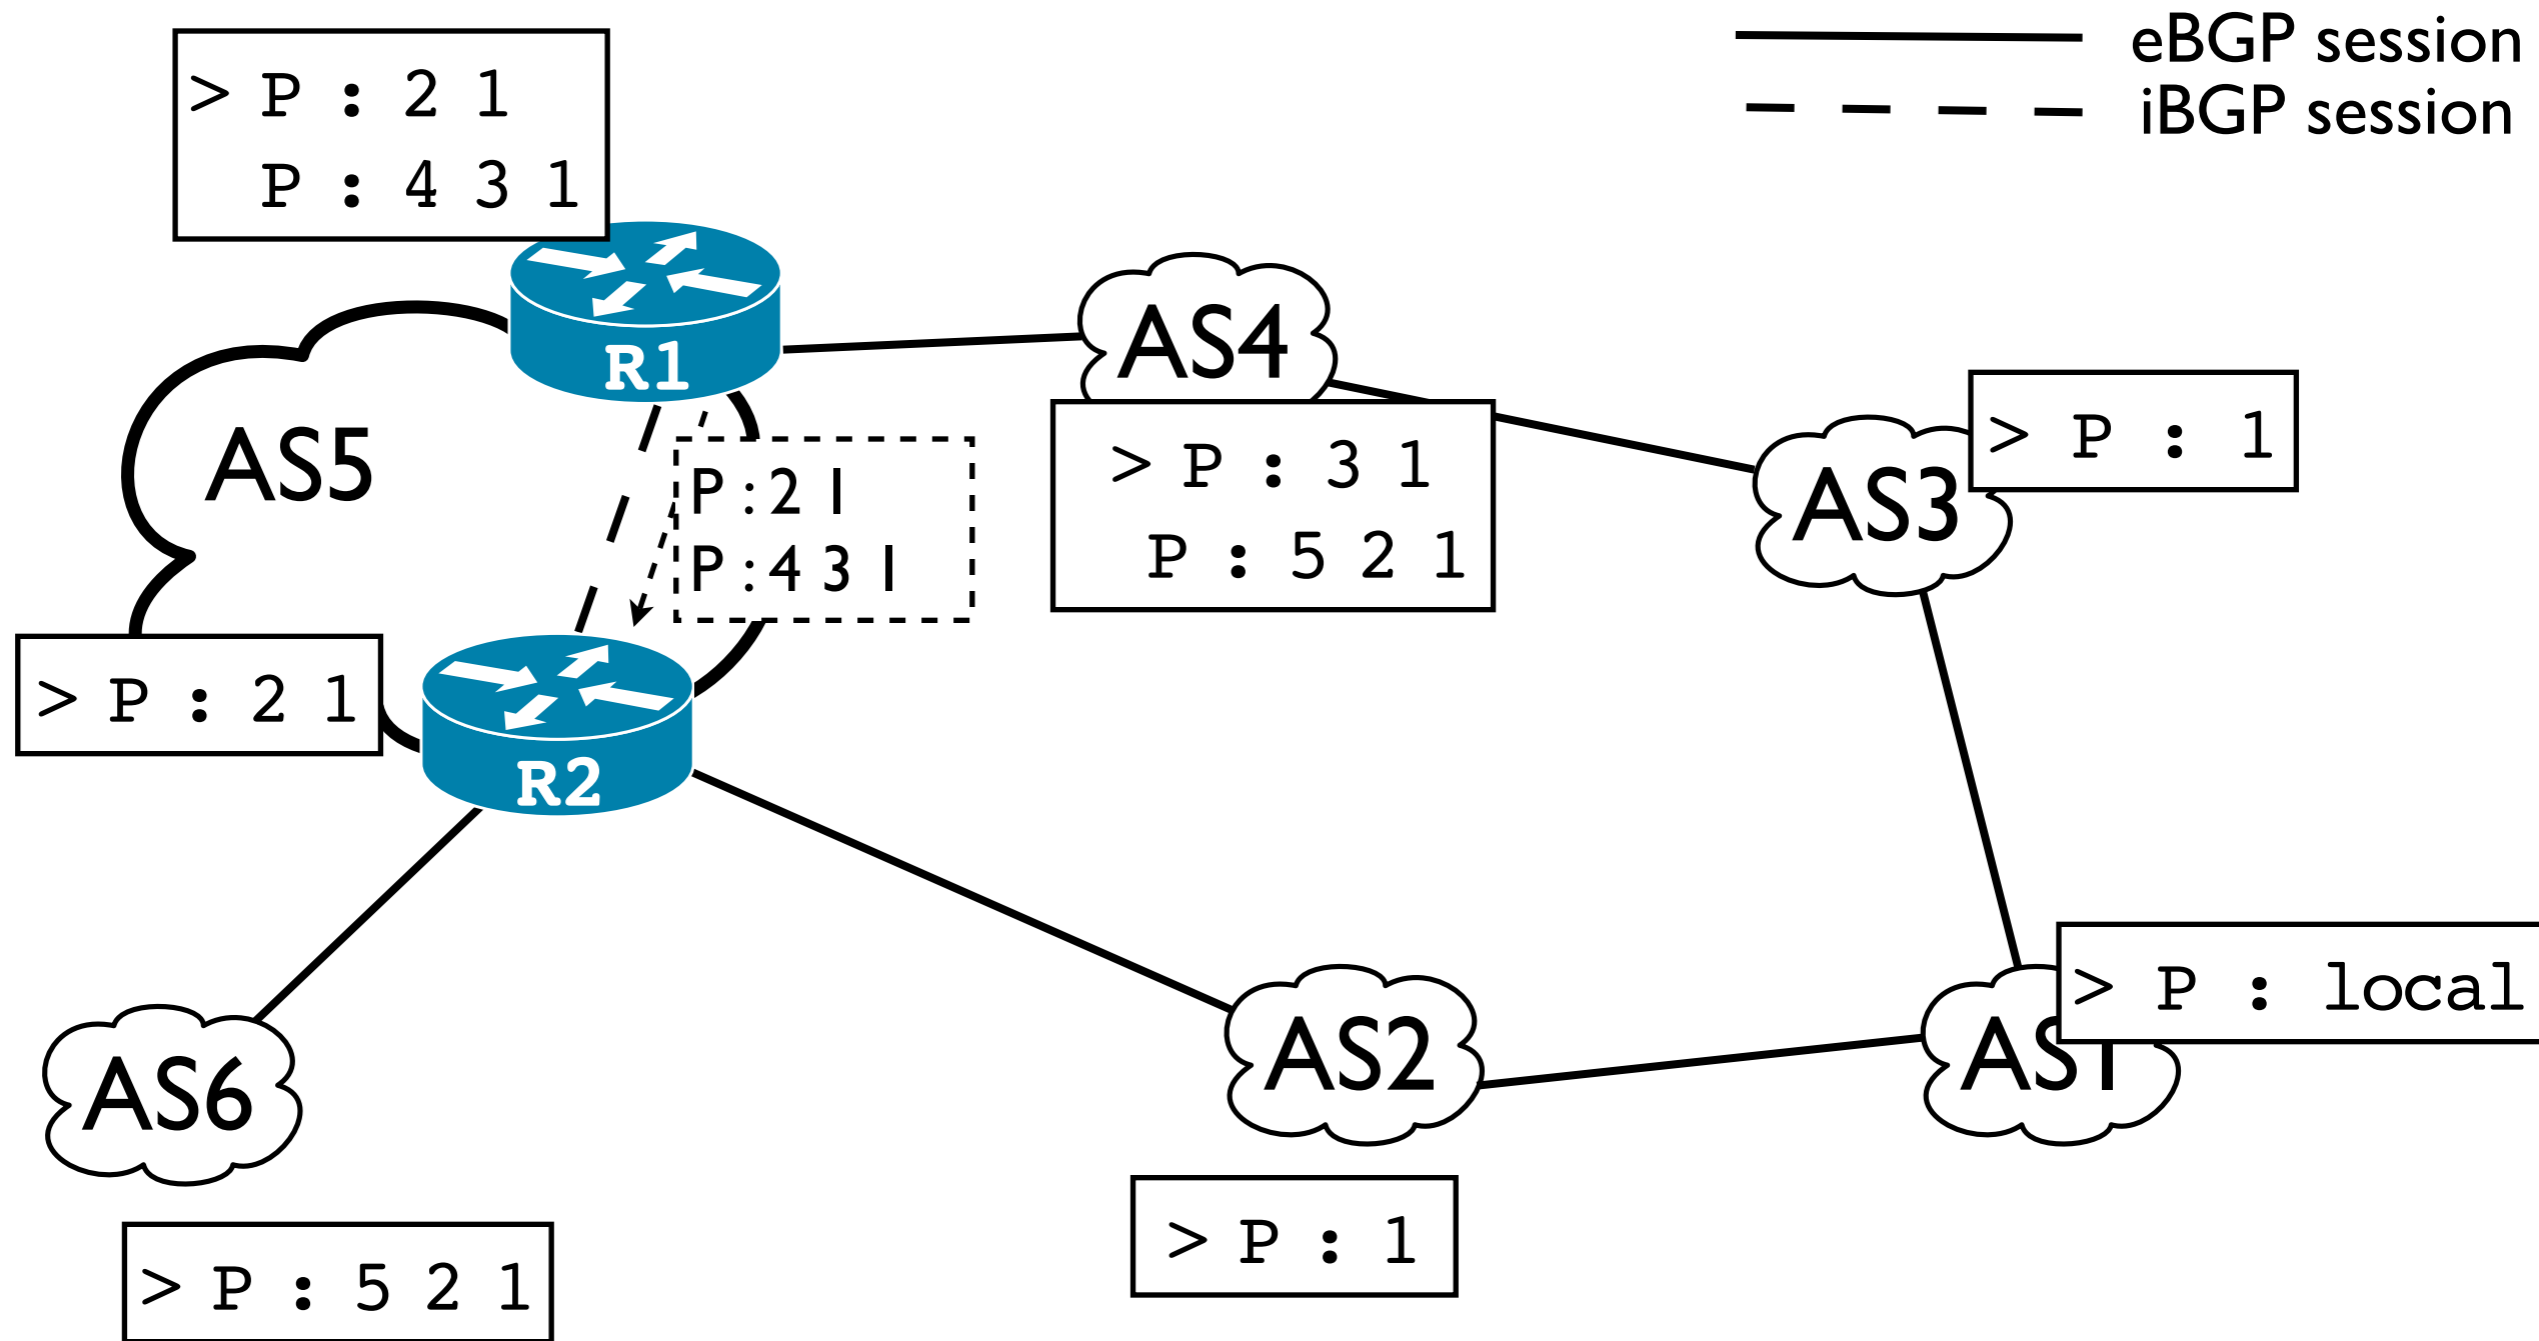
Agenda

- Introduction
- Presentation and evaluation of iBGP-caused withdraws
- **Preventing iBGP-caused withdraws**
 - ✓ **Advertising more than one path in iBGP (IETF)**
 - ✓ Our proposal : Advertising the availability of alternate paths in iBGP


Advertising multiple paths in iBGP



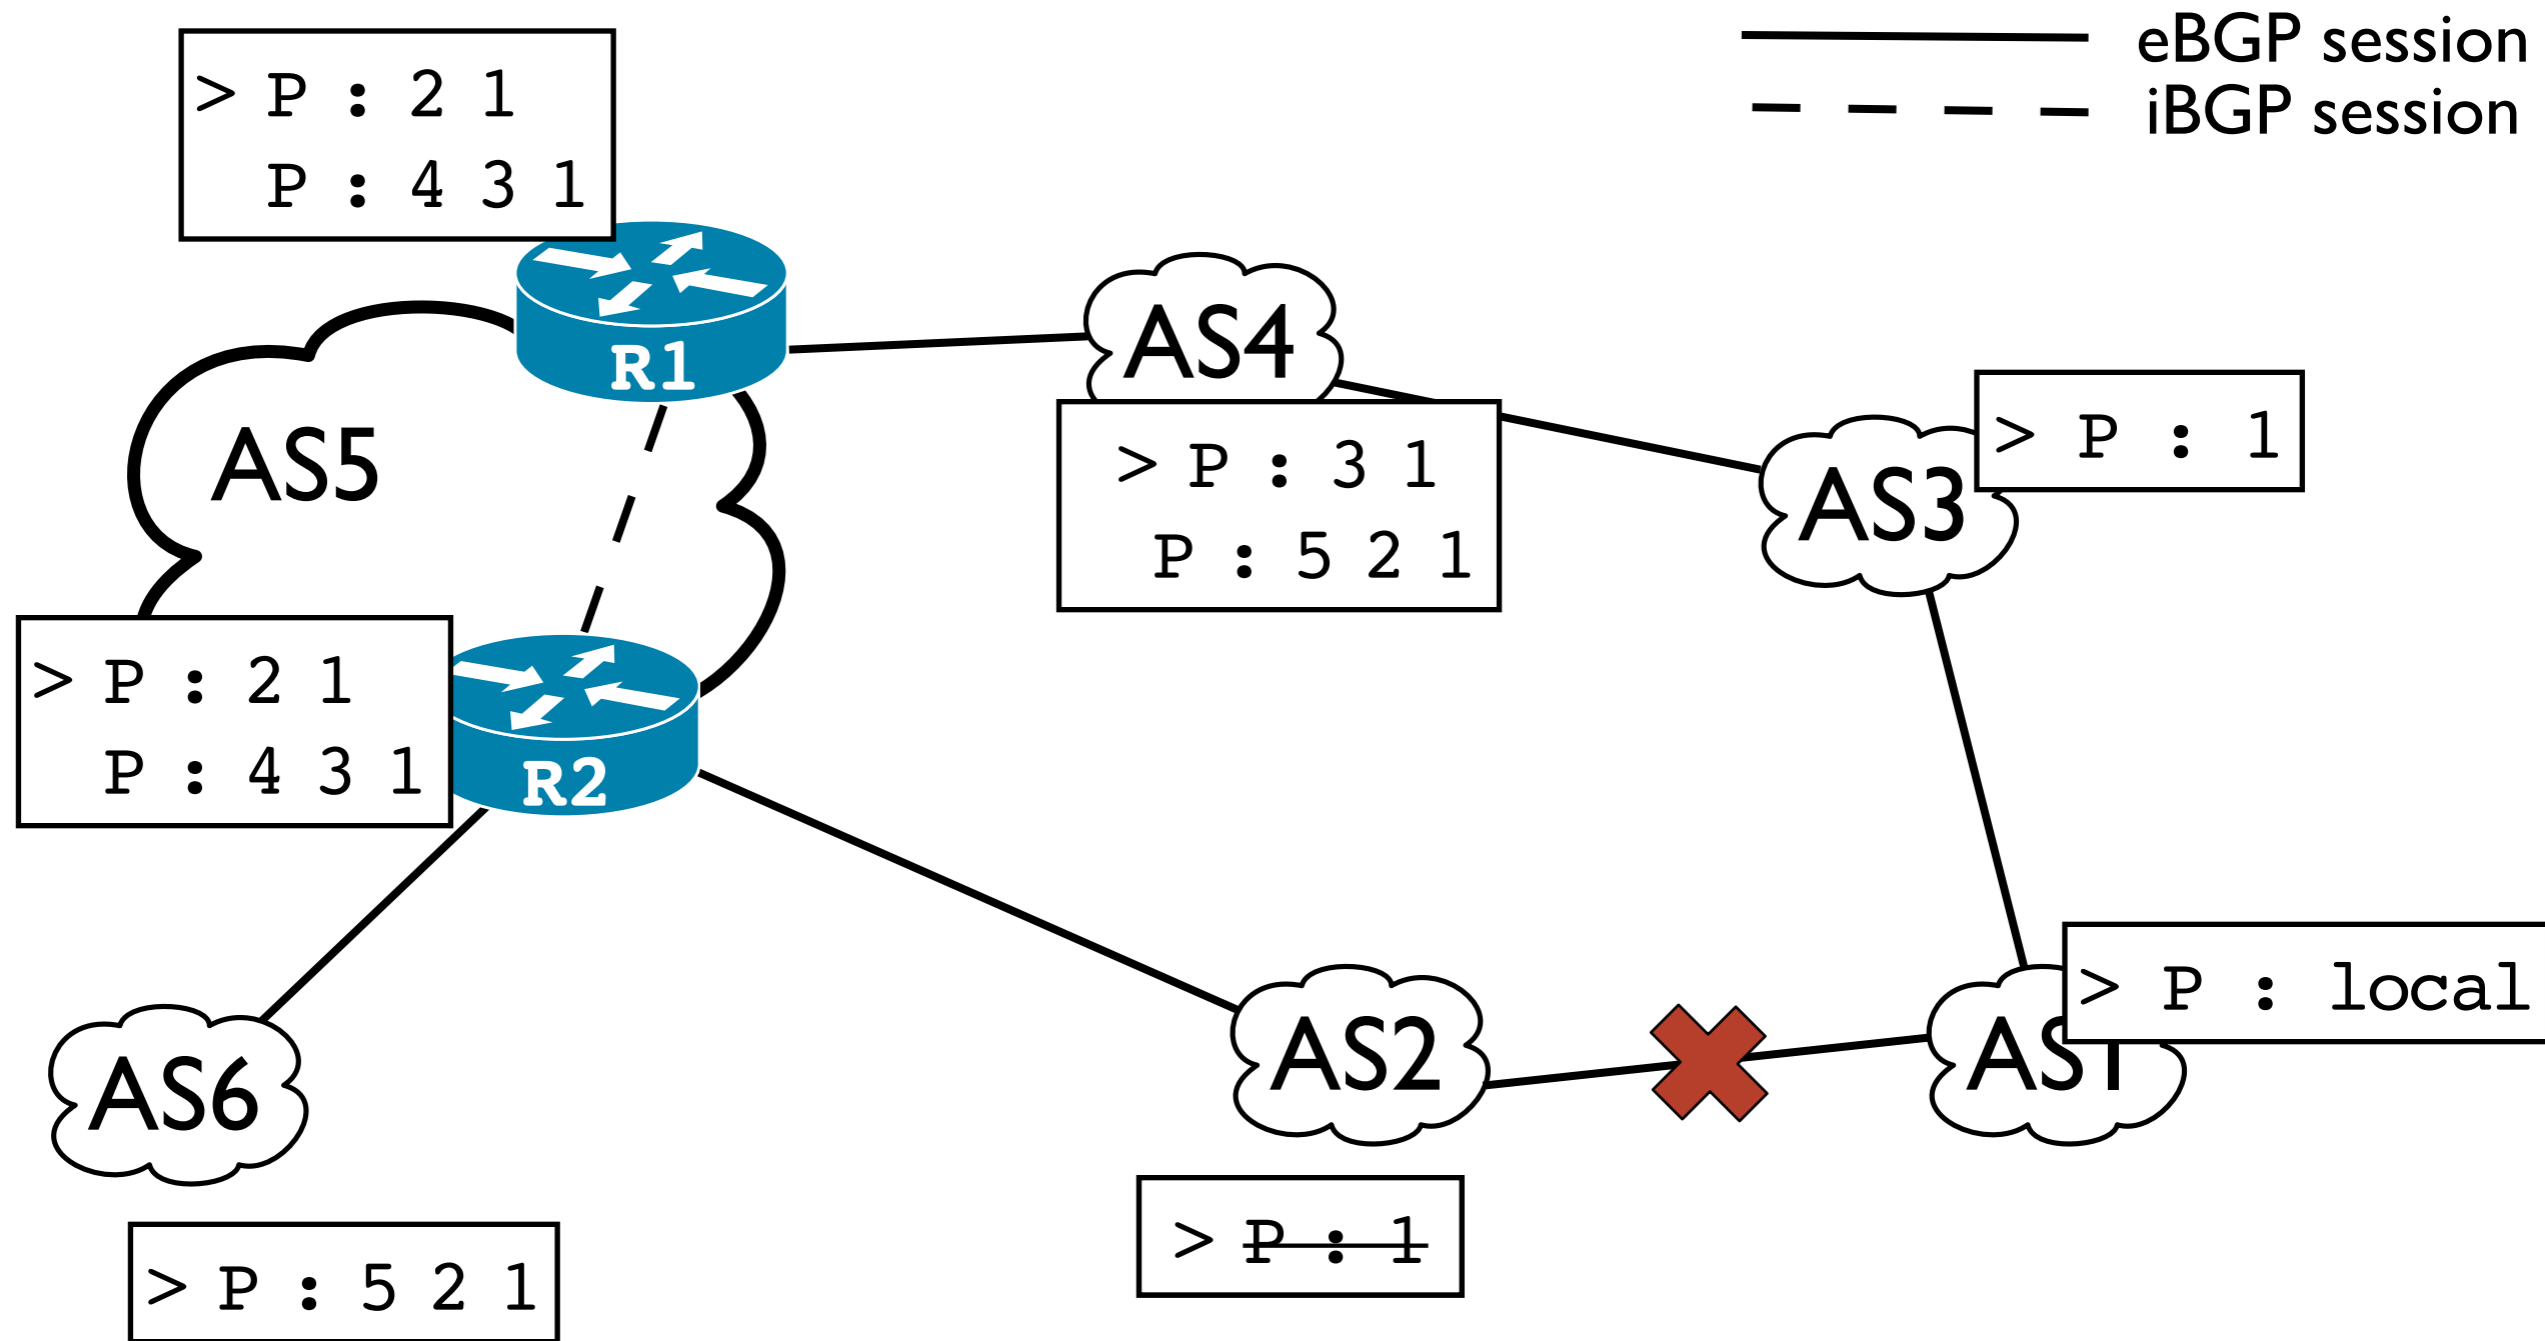
Advertising multiple paths in iBGP



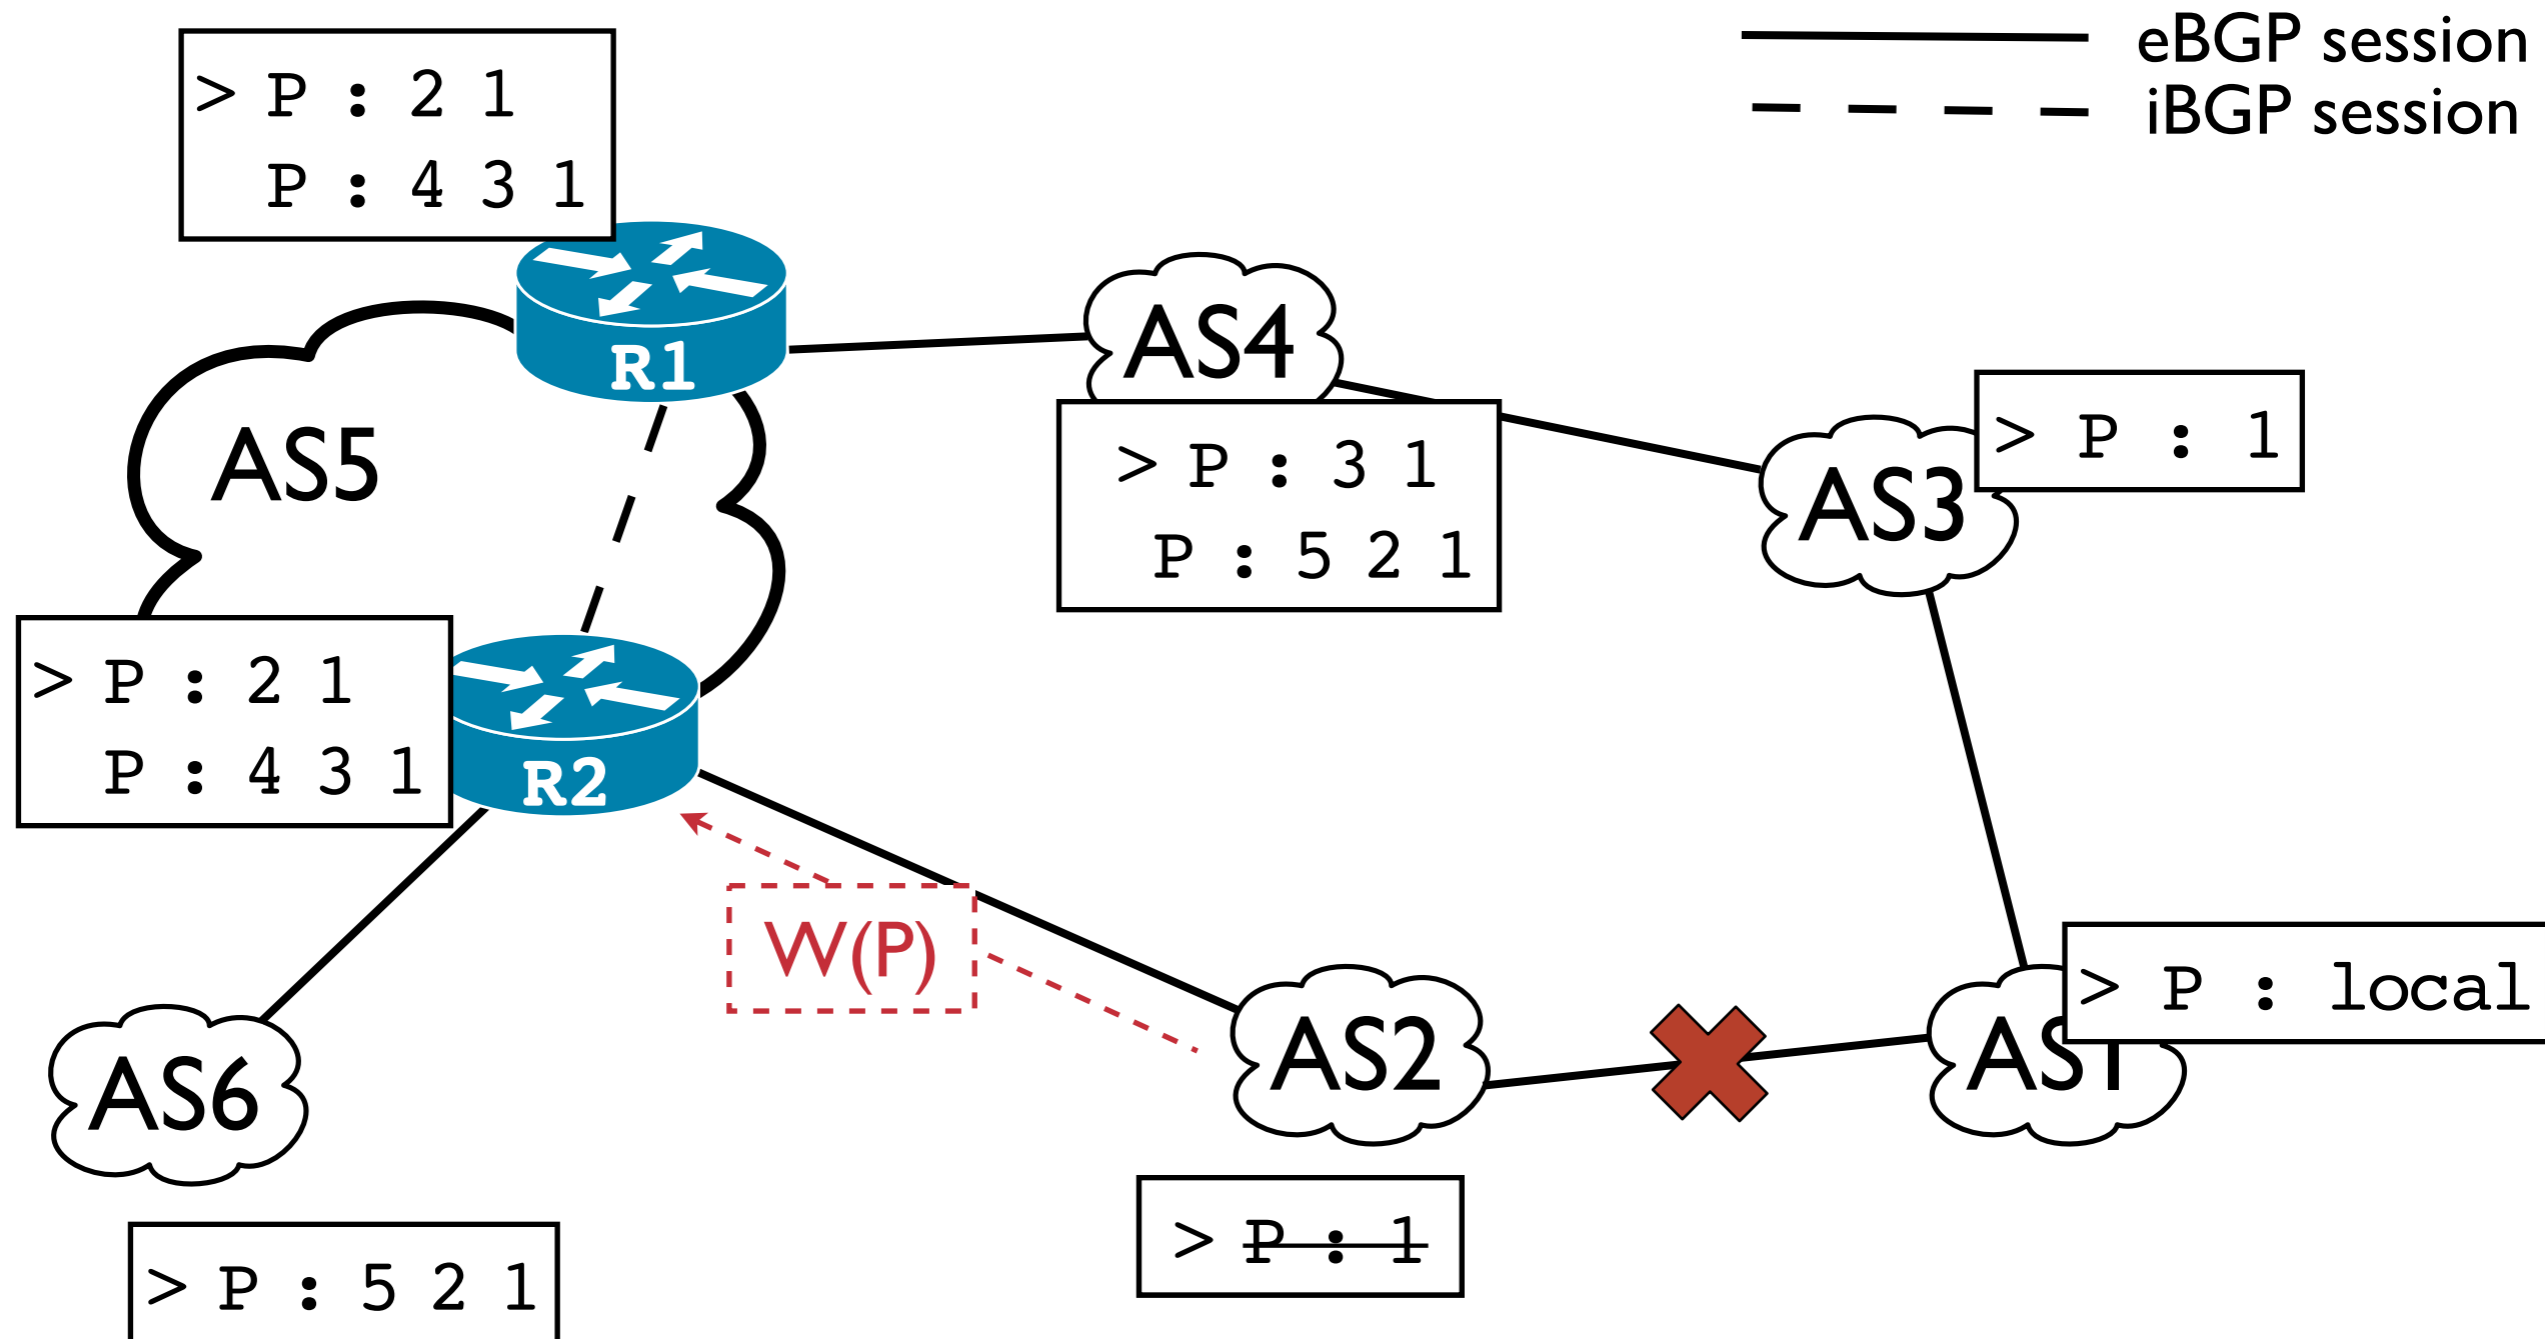
Advertising multiple paths in iBGP



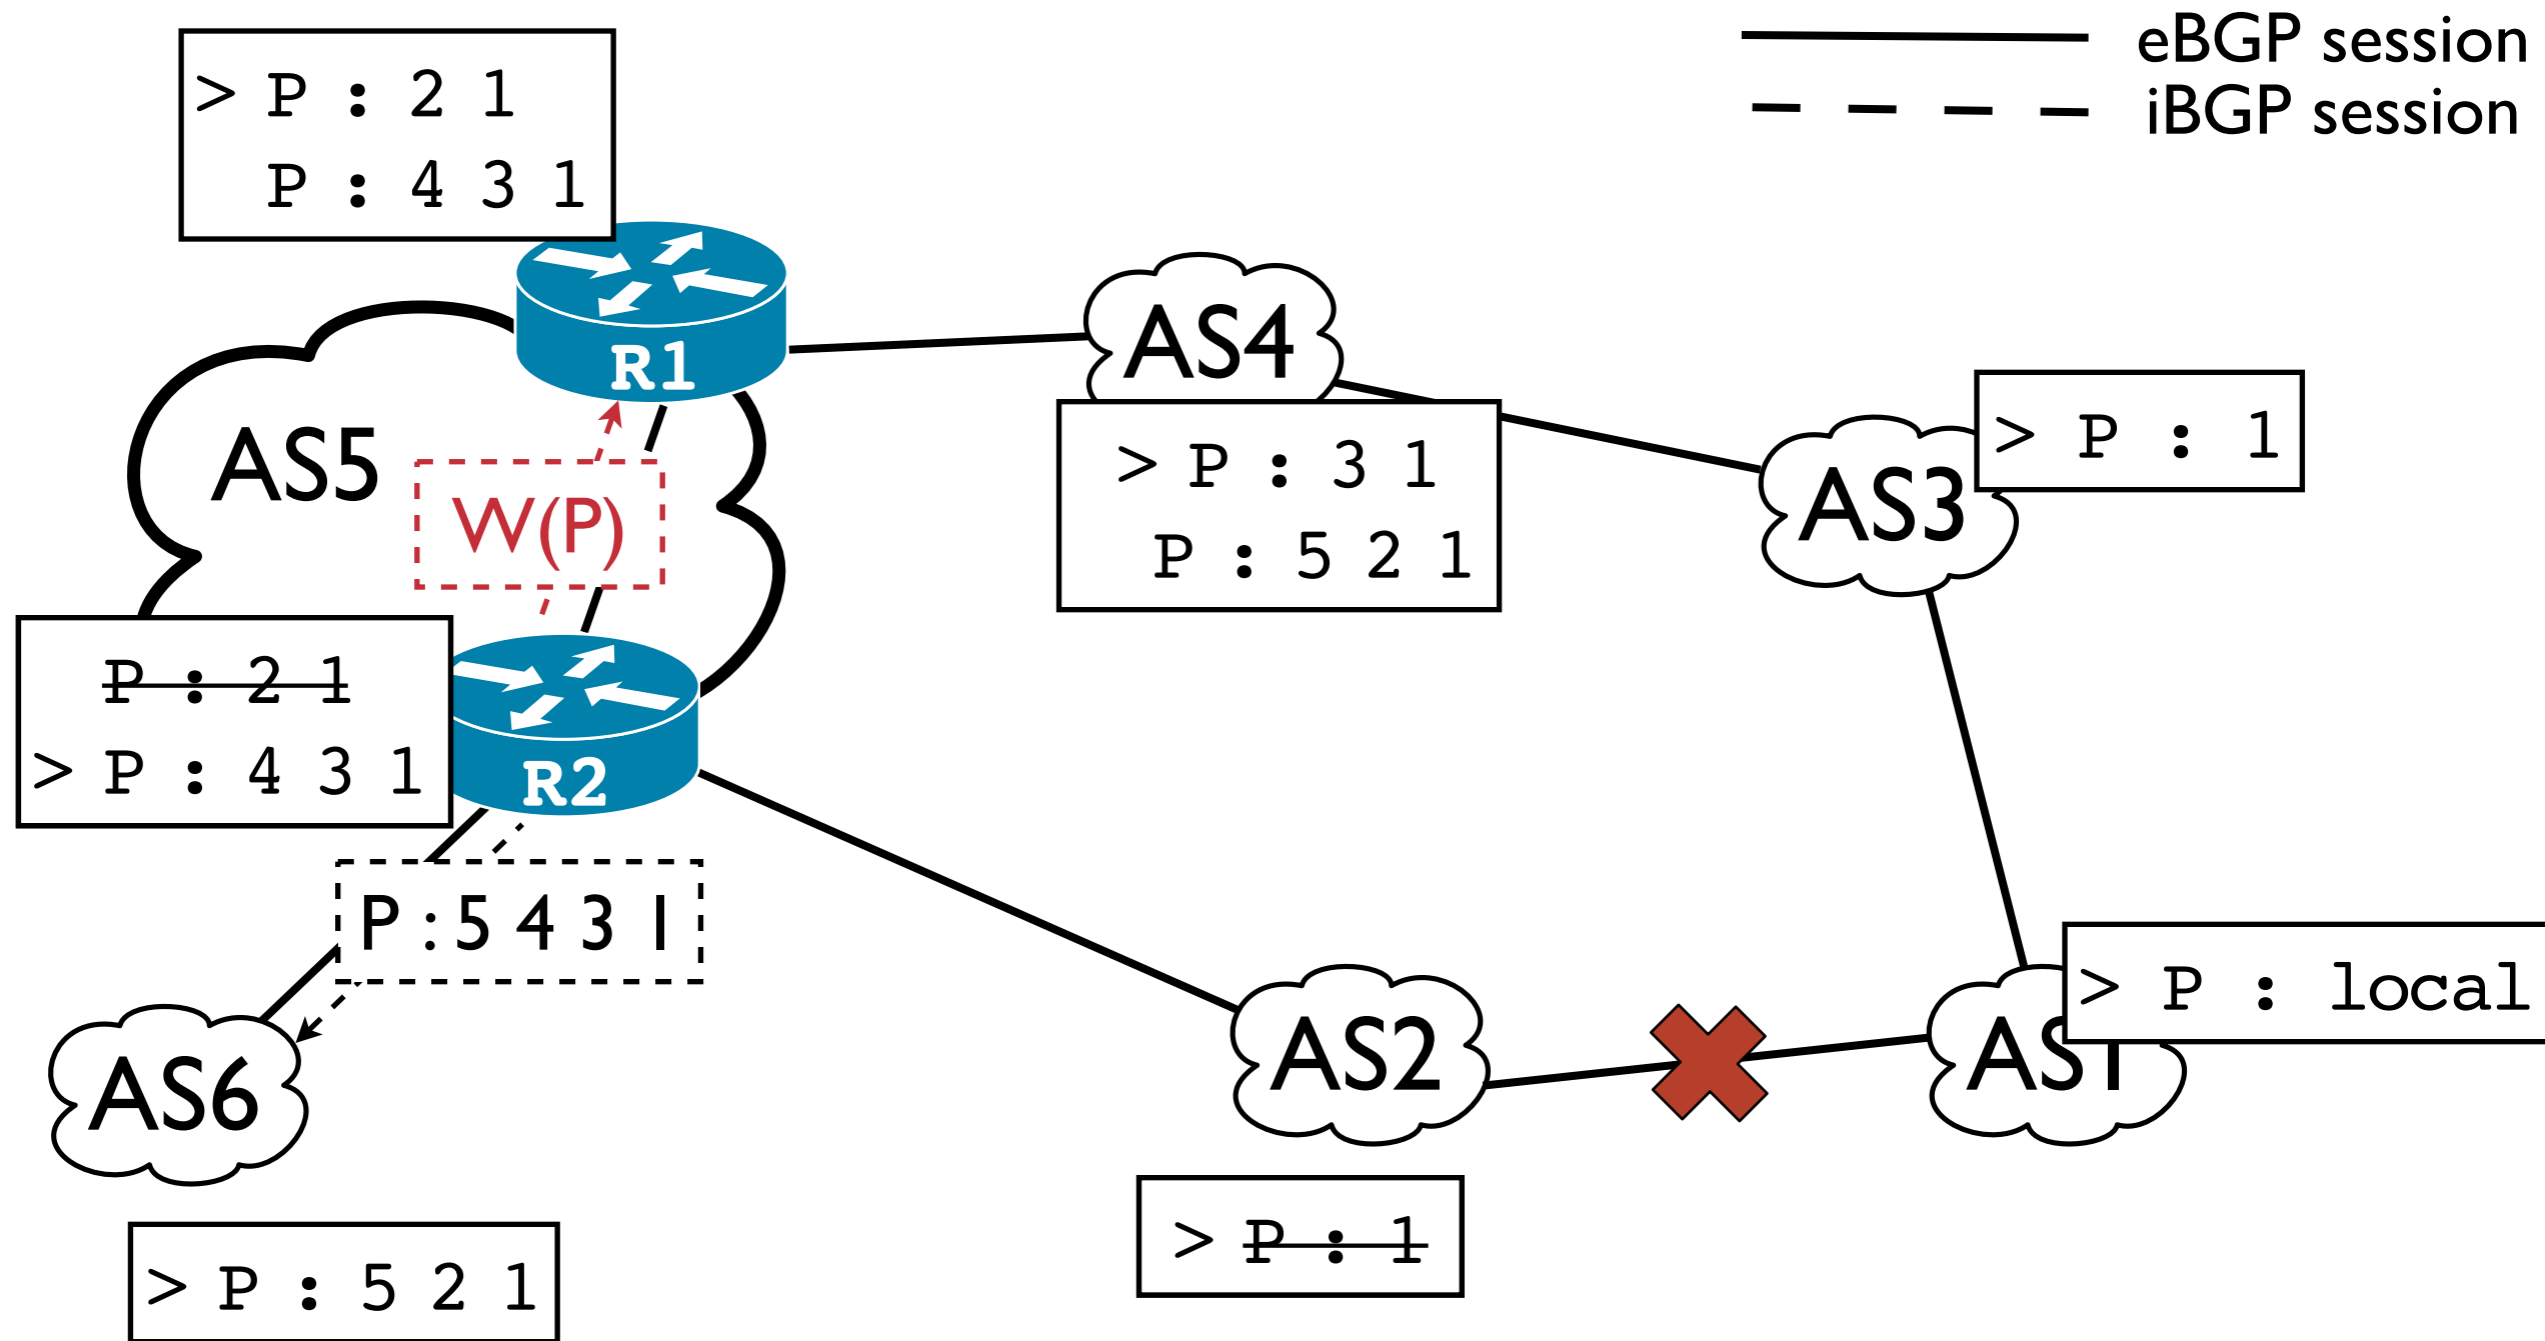
Advertising multiple paths in iBGP



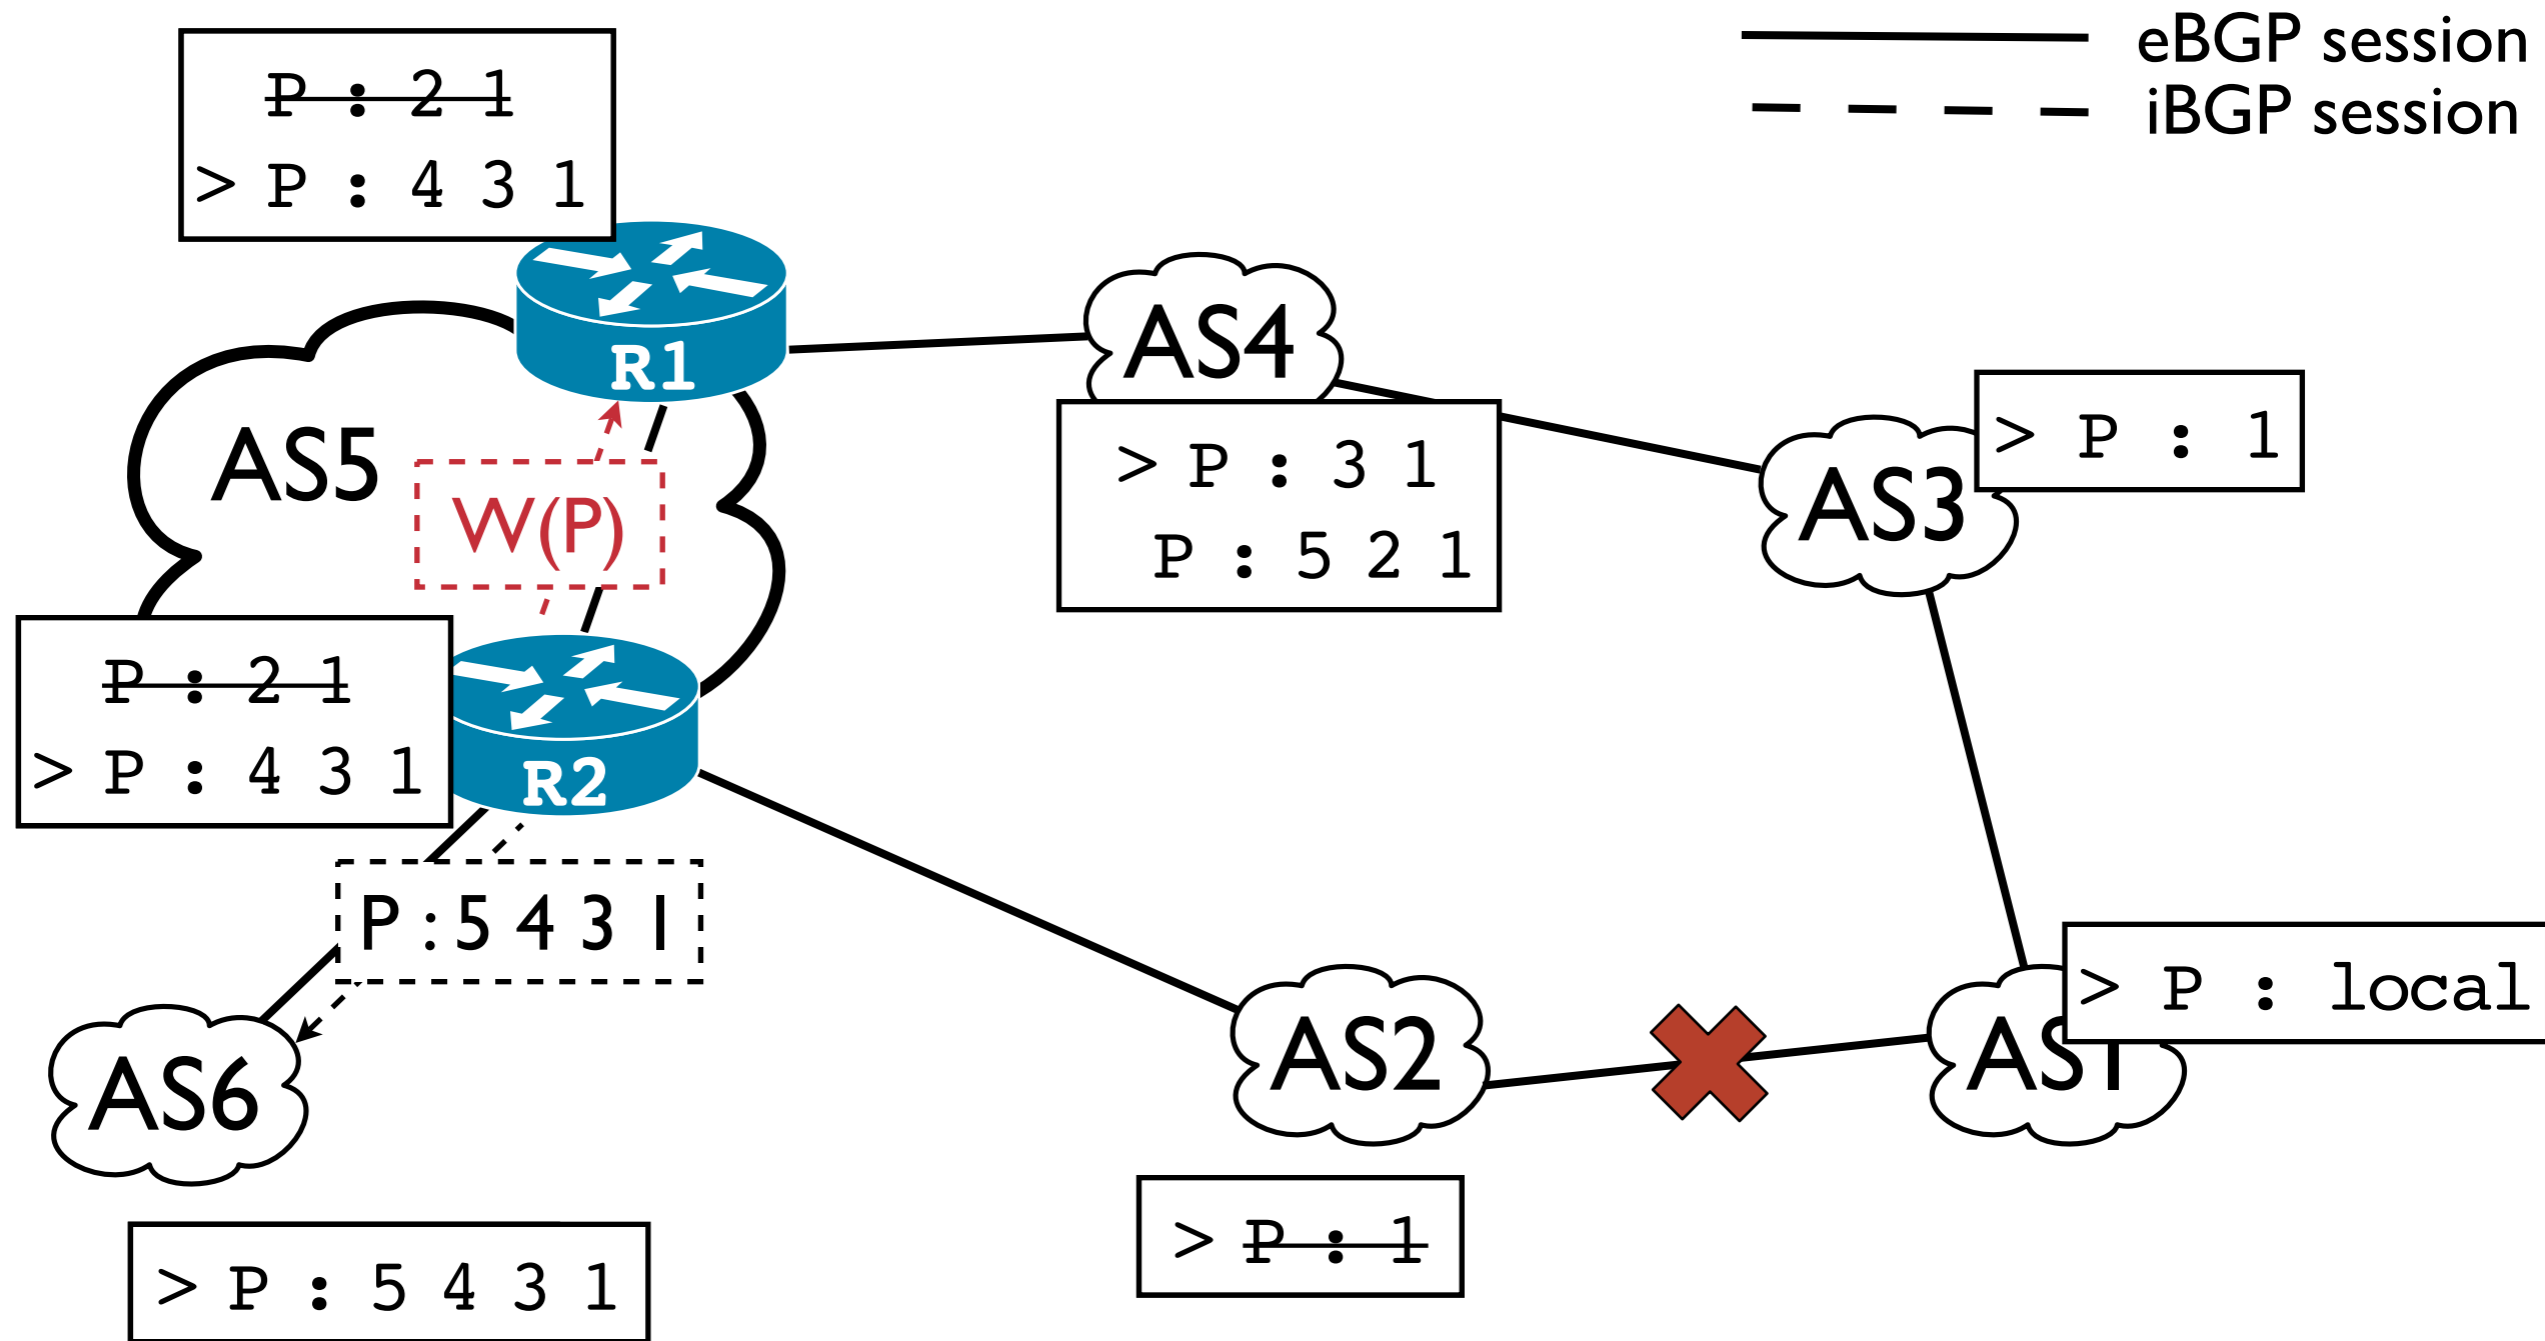
Advertising multiple paths in iBGP



Advertising multiple paths in iBGP



Advertising multiple paths in iBGP



Advertising multiple paths in iBGP

- Very promising solution, but :
 - ✓ Need to modify the BGP protocol to support advertisement of multiple paths for a prefix
 - ✓ All the routers of an AS need to be upgraded accordingly

Agenda

- Introduction
- Presentation and evaluation of iBGP-caused withdraws
- **Preventing iBGP-caused withdraws**
 - ✓ Advertising more than one path in iBGP (IETF)
 - ✓ **Our proposal : Advertising the availability of alternate paths in iBGP**

How to prevent iBGP-caused withdraws? (2)


- Instead of advertising alternate paths, advertise the **availability of alternate paths**
- BGP Community attribute used to tag this information to the path : DIVERSITY_AVAILABLE
- Lighter solution than the previous one, can be incrementally deployed
 - ✓ No modification to BGP protocol
 - ✓ DIVERSITY_AVAILABLE BGP community is simply not considered by legacy routers

Advertising the availability of alternate paths

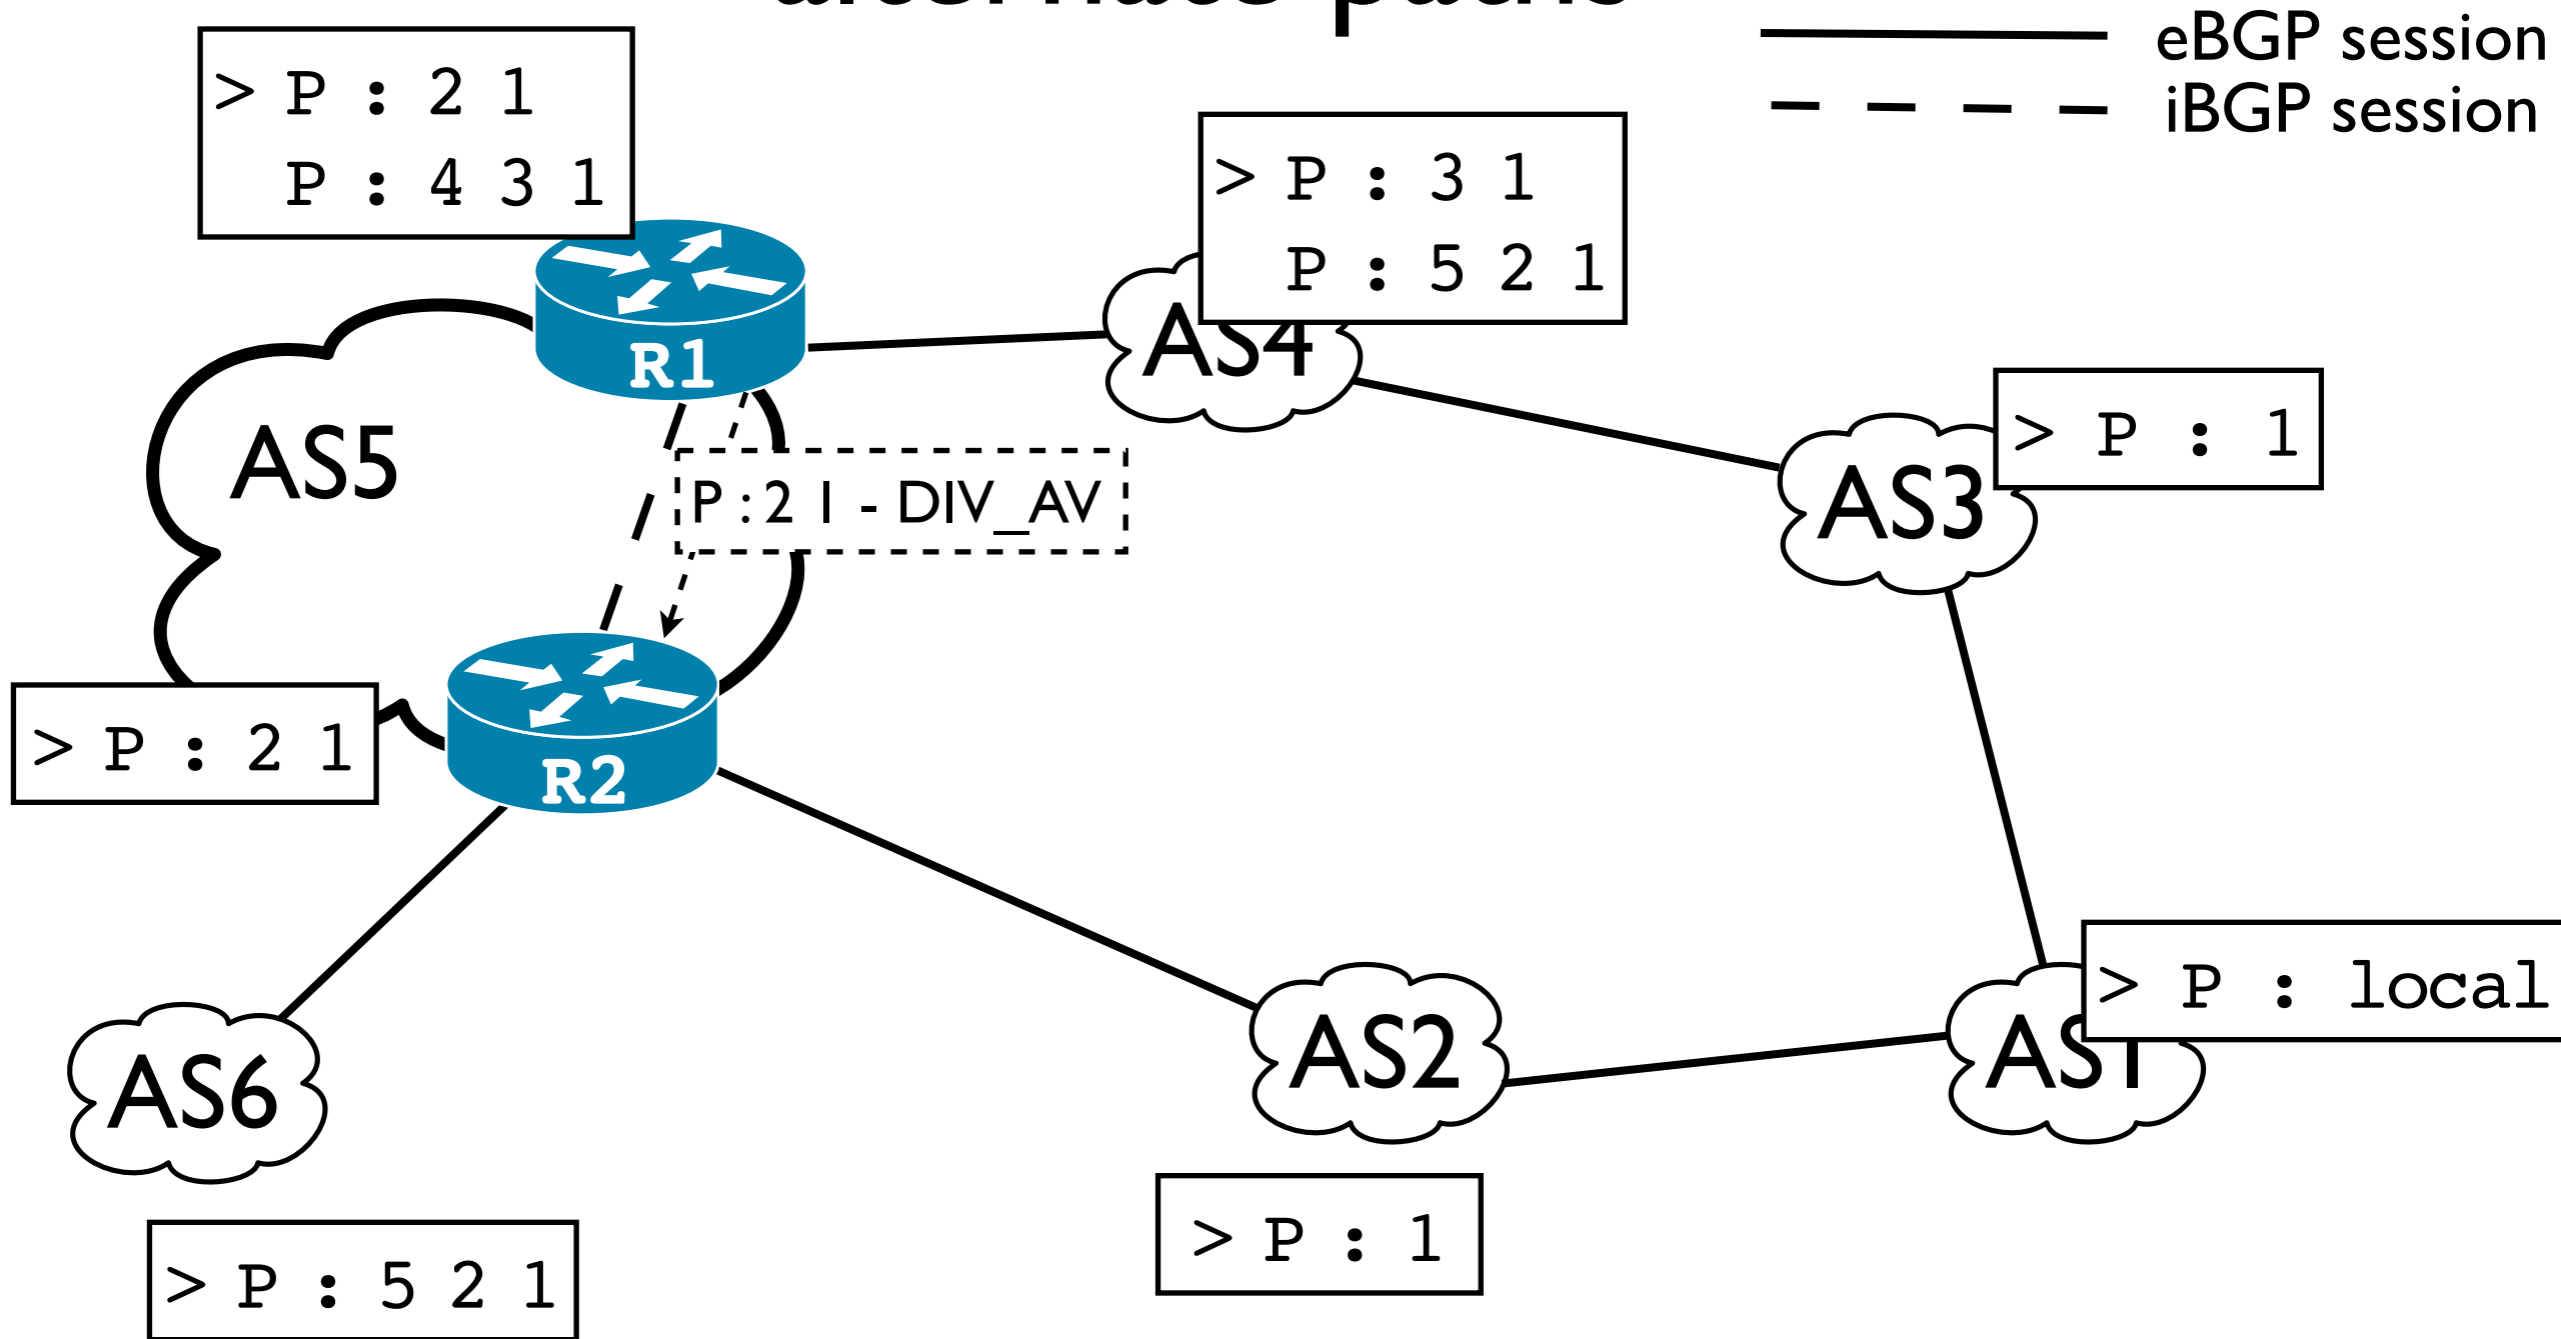
- **Principle :**

- ✓ A router knowing an alternate path
 - Tags the community `DIVERSITY_AVAILABLE` to the primary path
 - Then, it re-advertises the primary path to all its iBGP peers

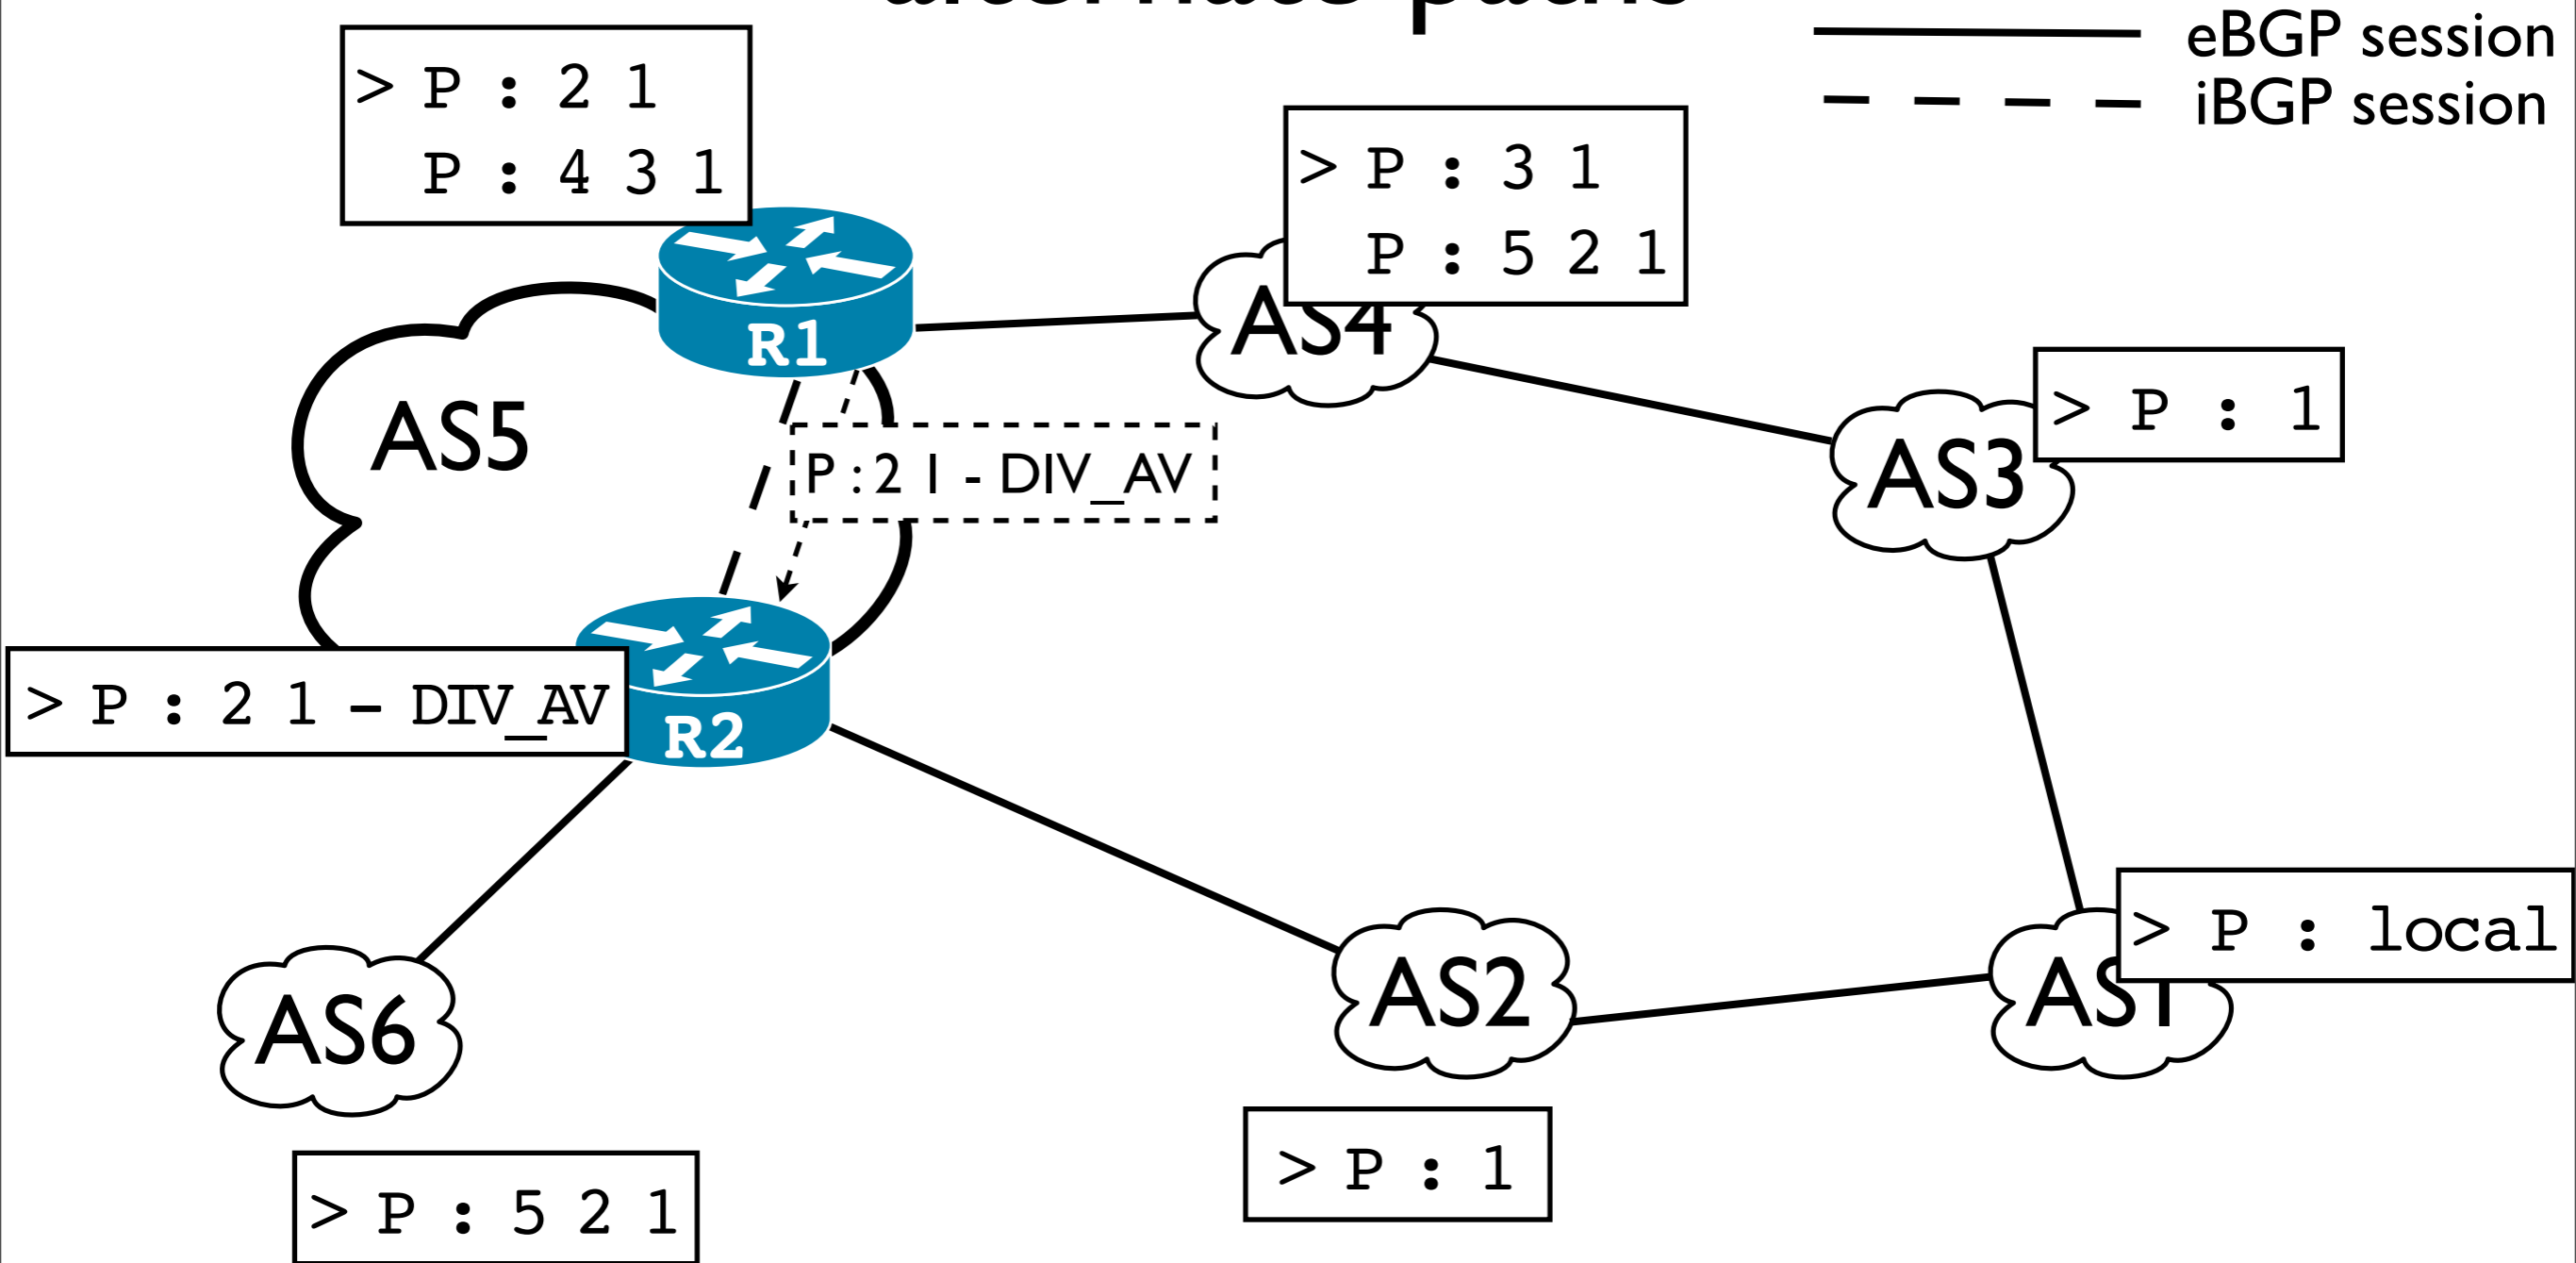
Advertising the availability of alternate paths



Advertising the availability of alternate paths



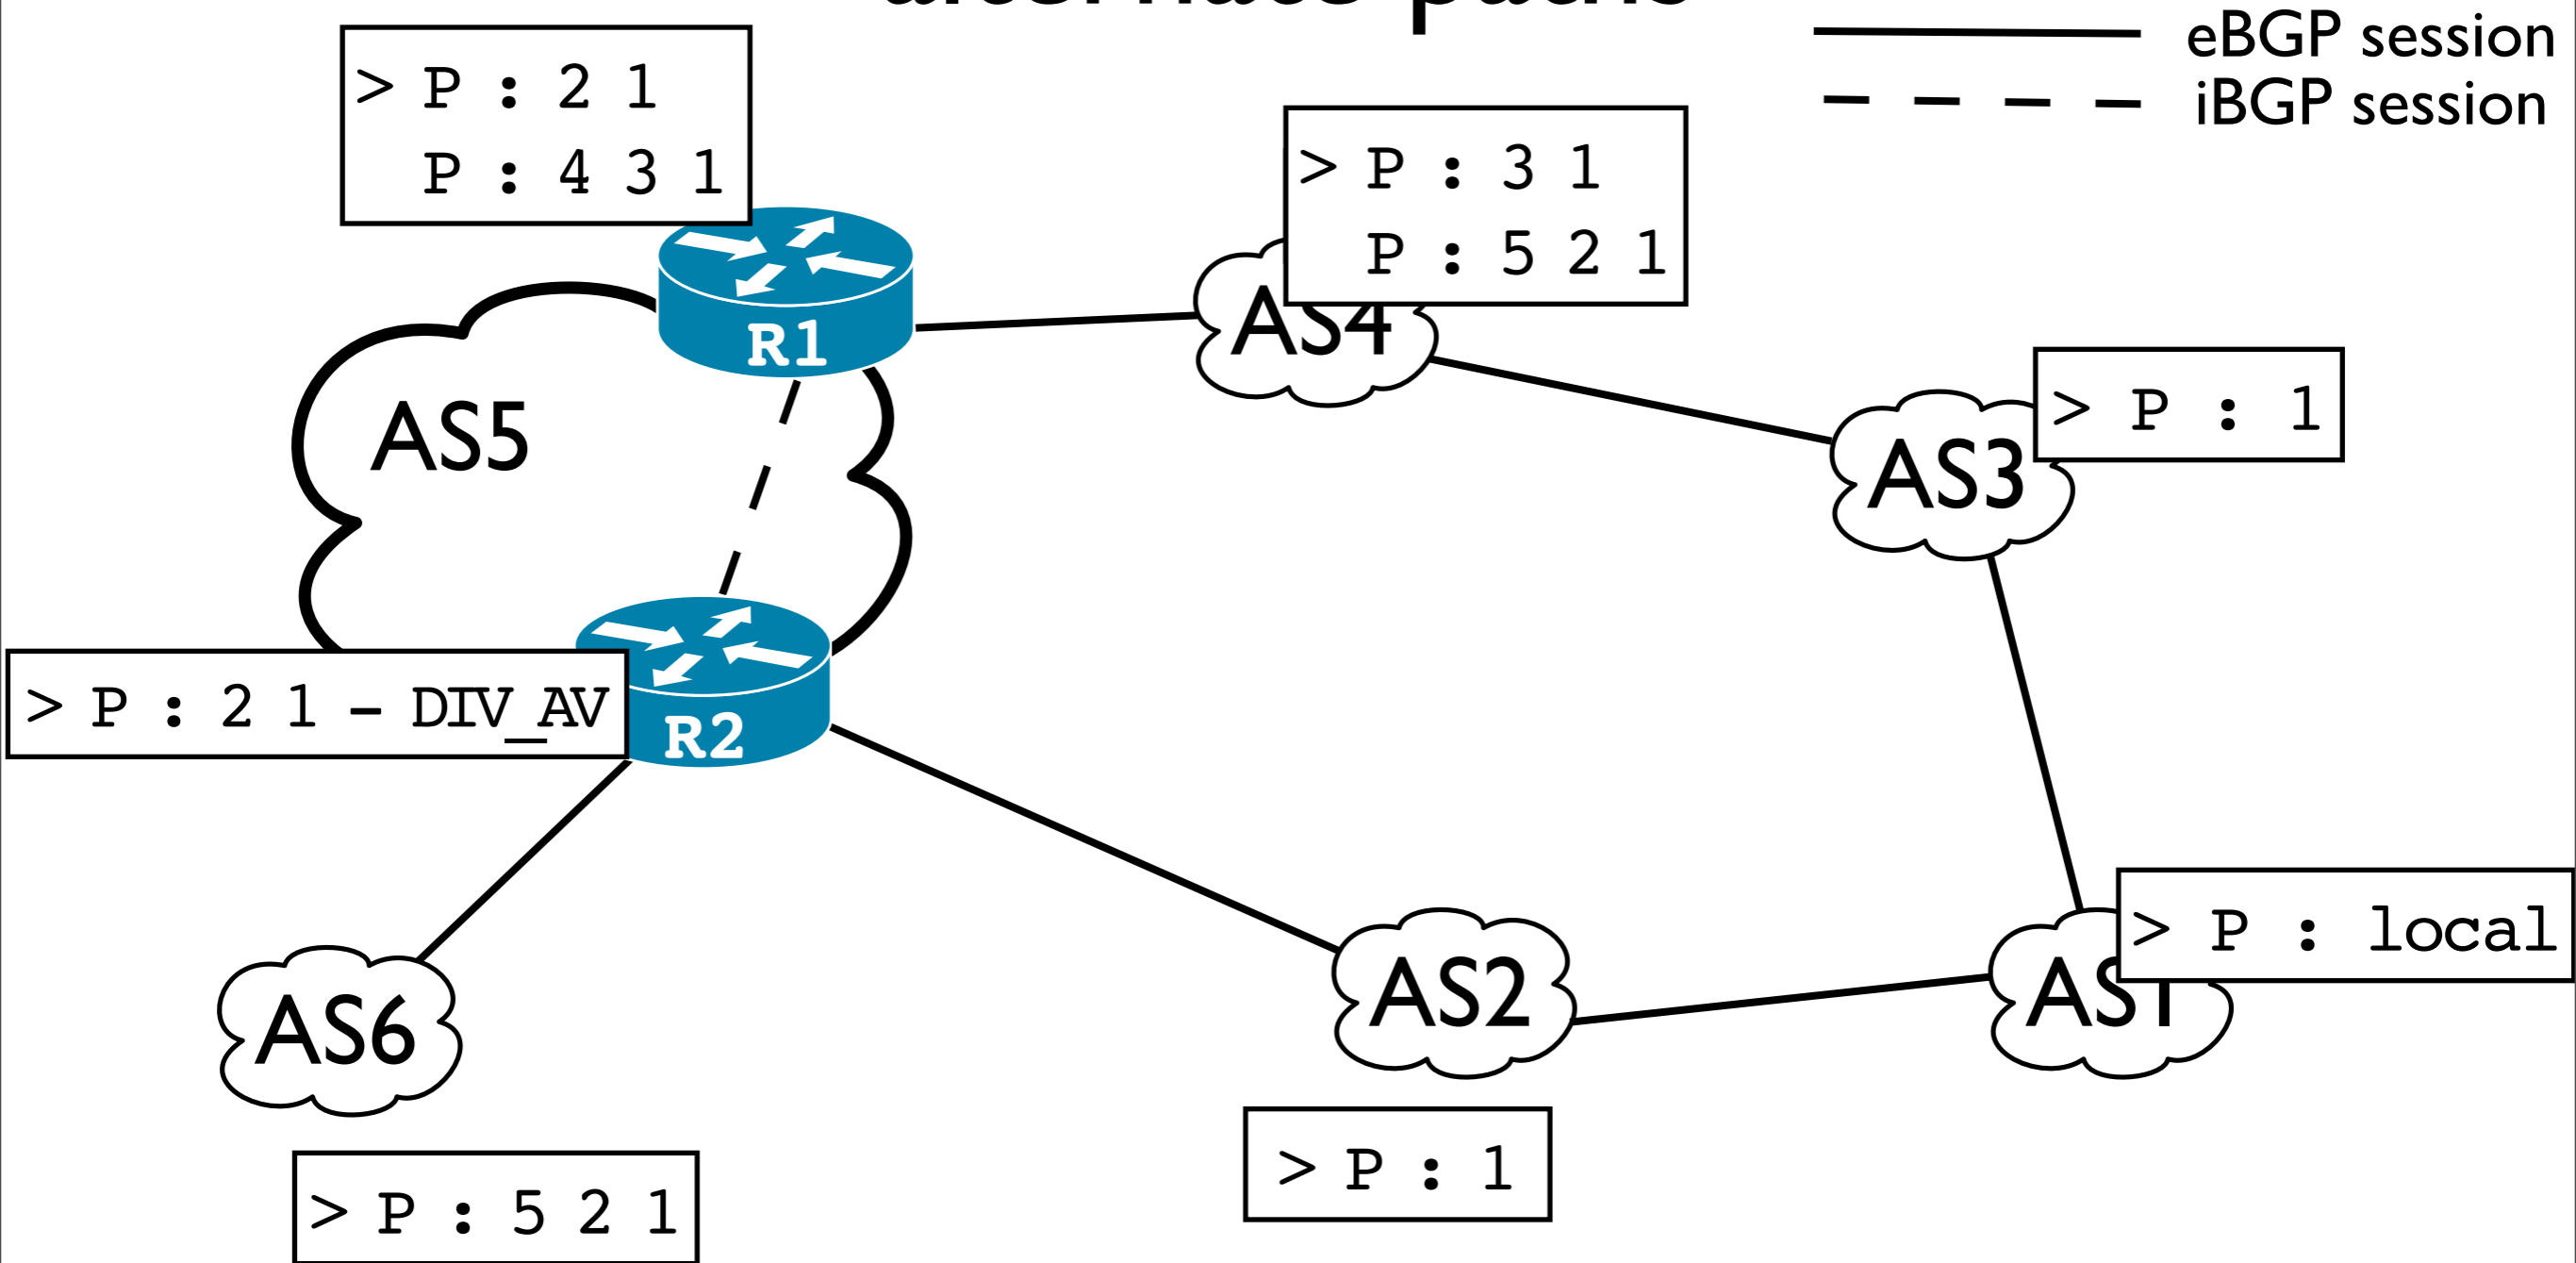
Advertising the availability of alternate paths



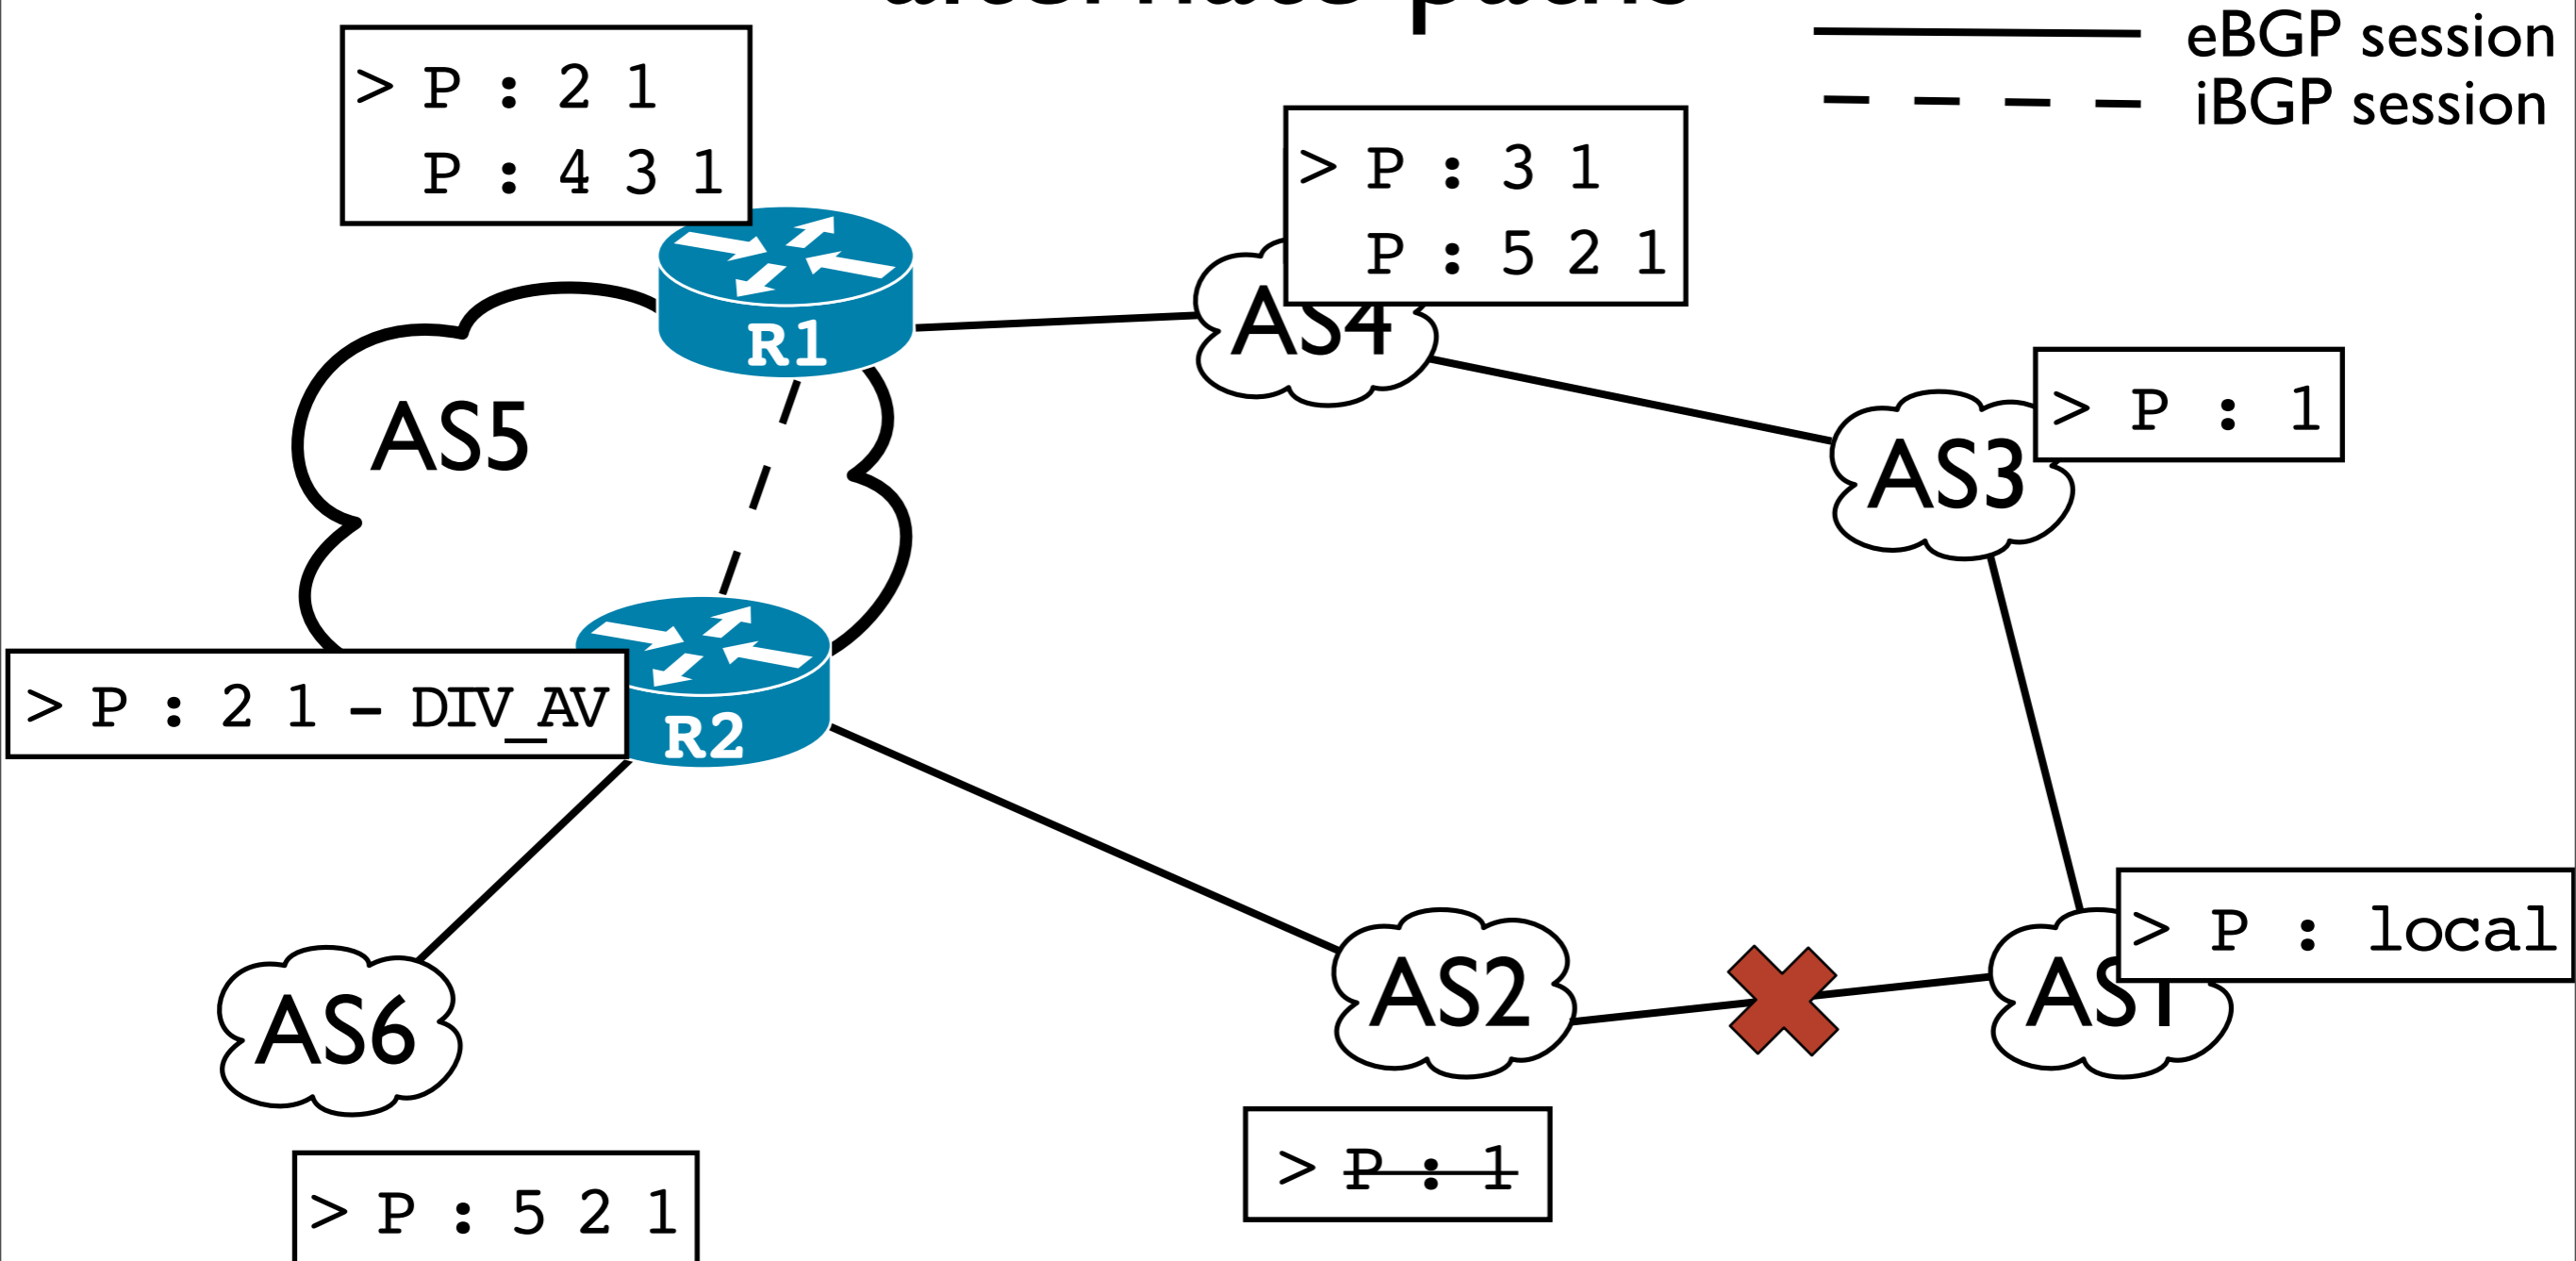
Advertising the availability of alternate paths

- **When the primary path fails :**
 - Routers without an alternate path check the community
 - ✓ If present : Wait for the alternate path before sending withdraws on eBGP
 - ✓ Otherwise : send the withdraw on eBGP sessions immediately

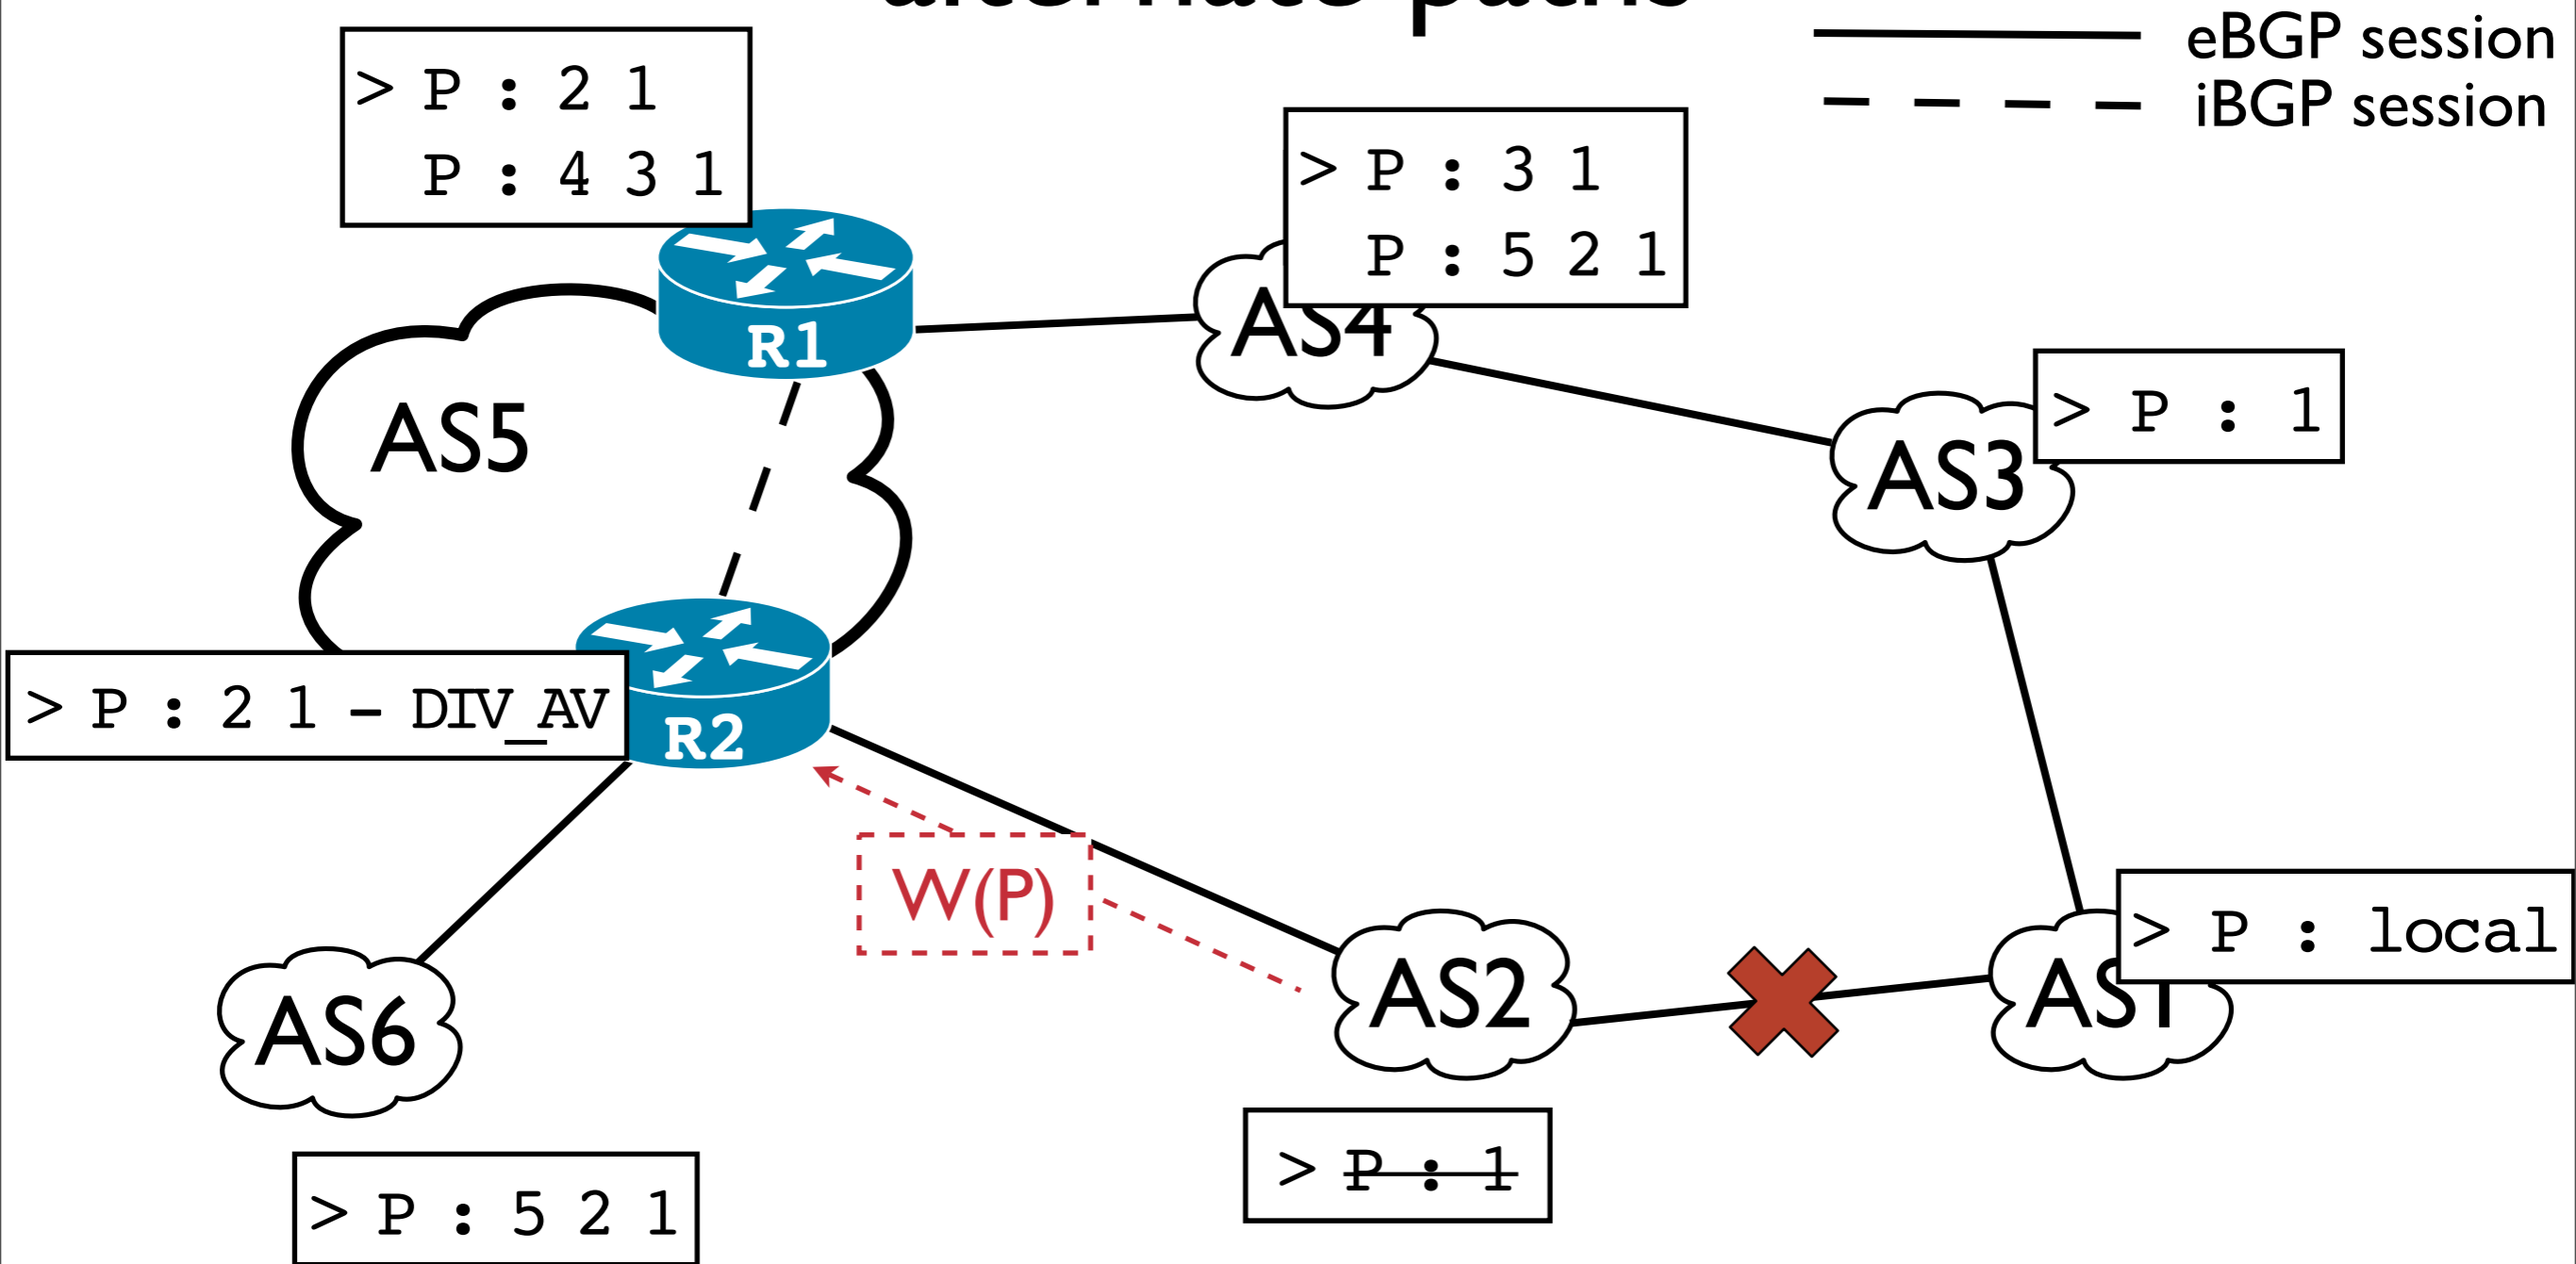
Advertising the availability of alternate paths



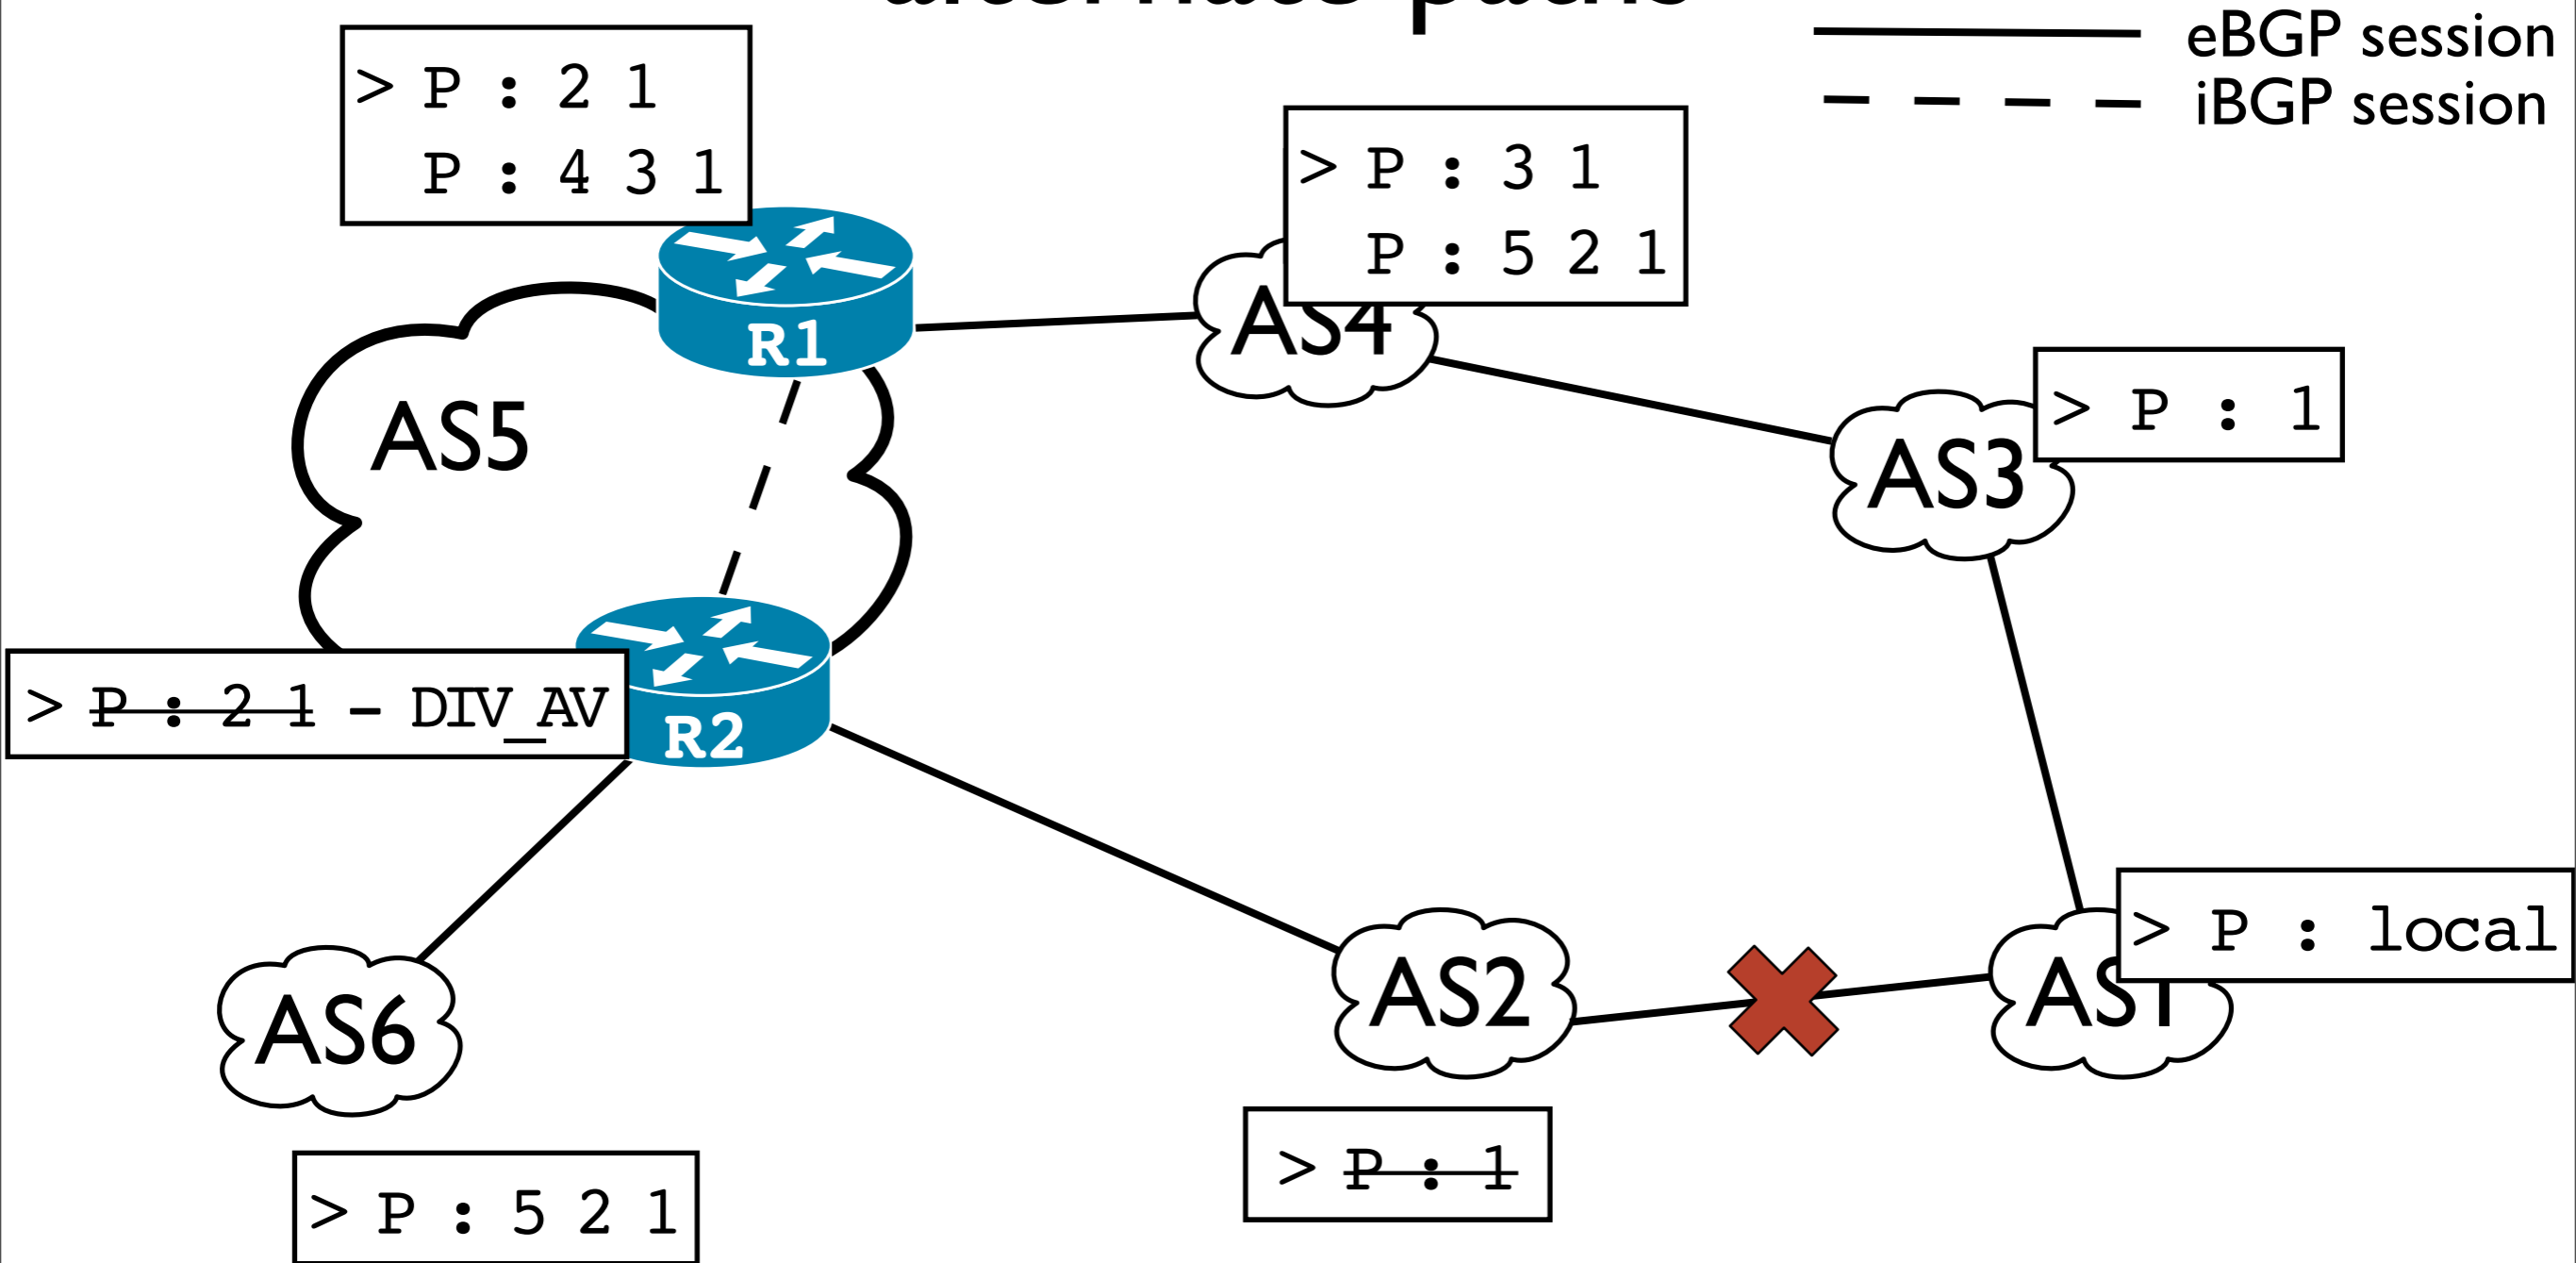
Advertising the availability of alternate paths



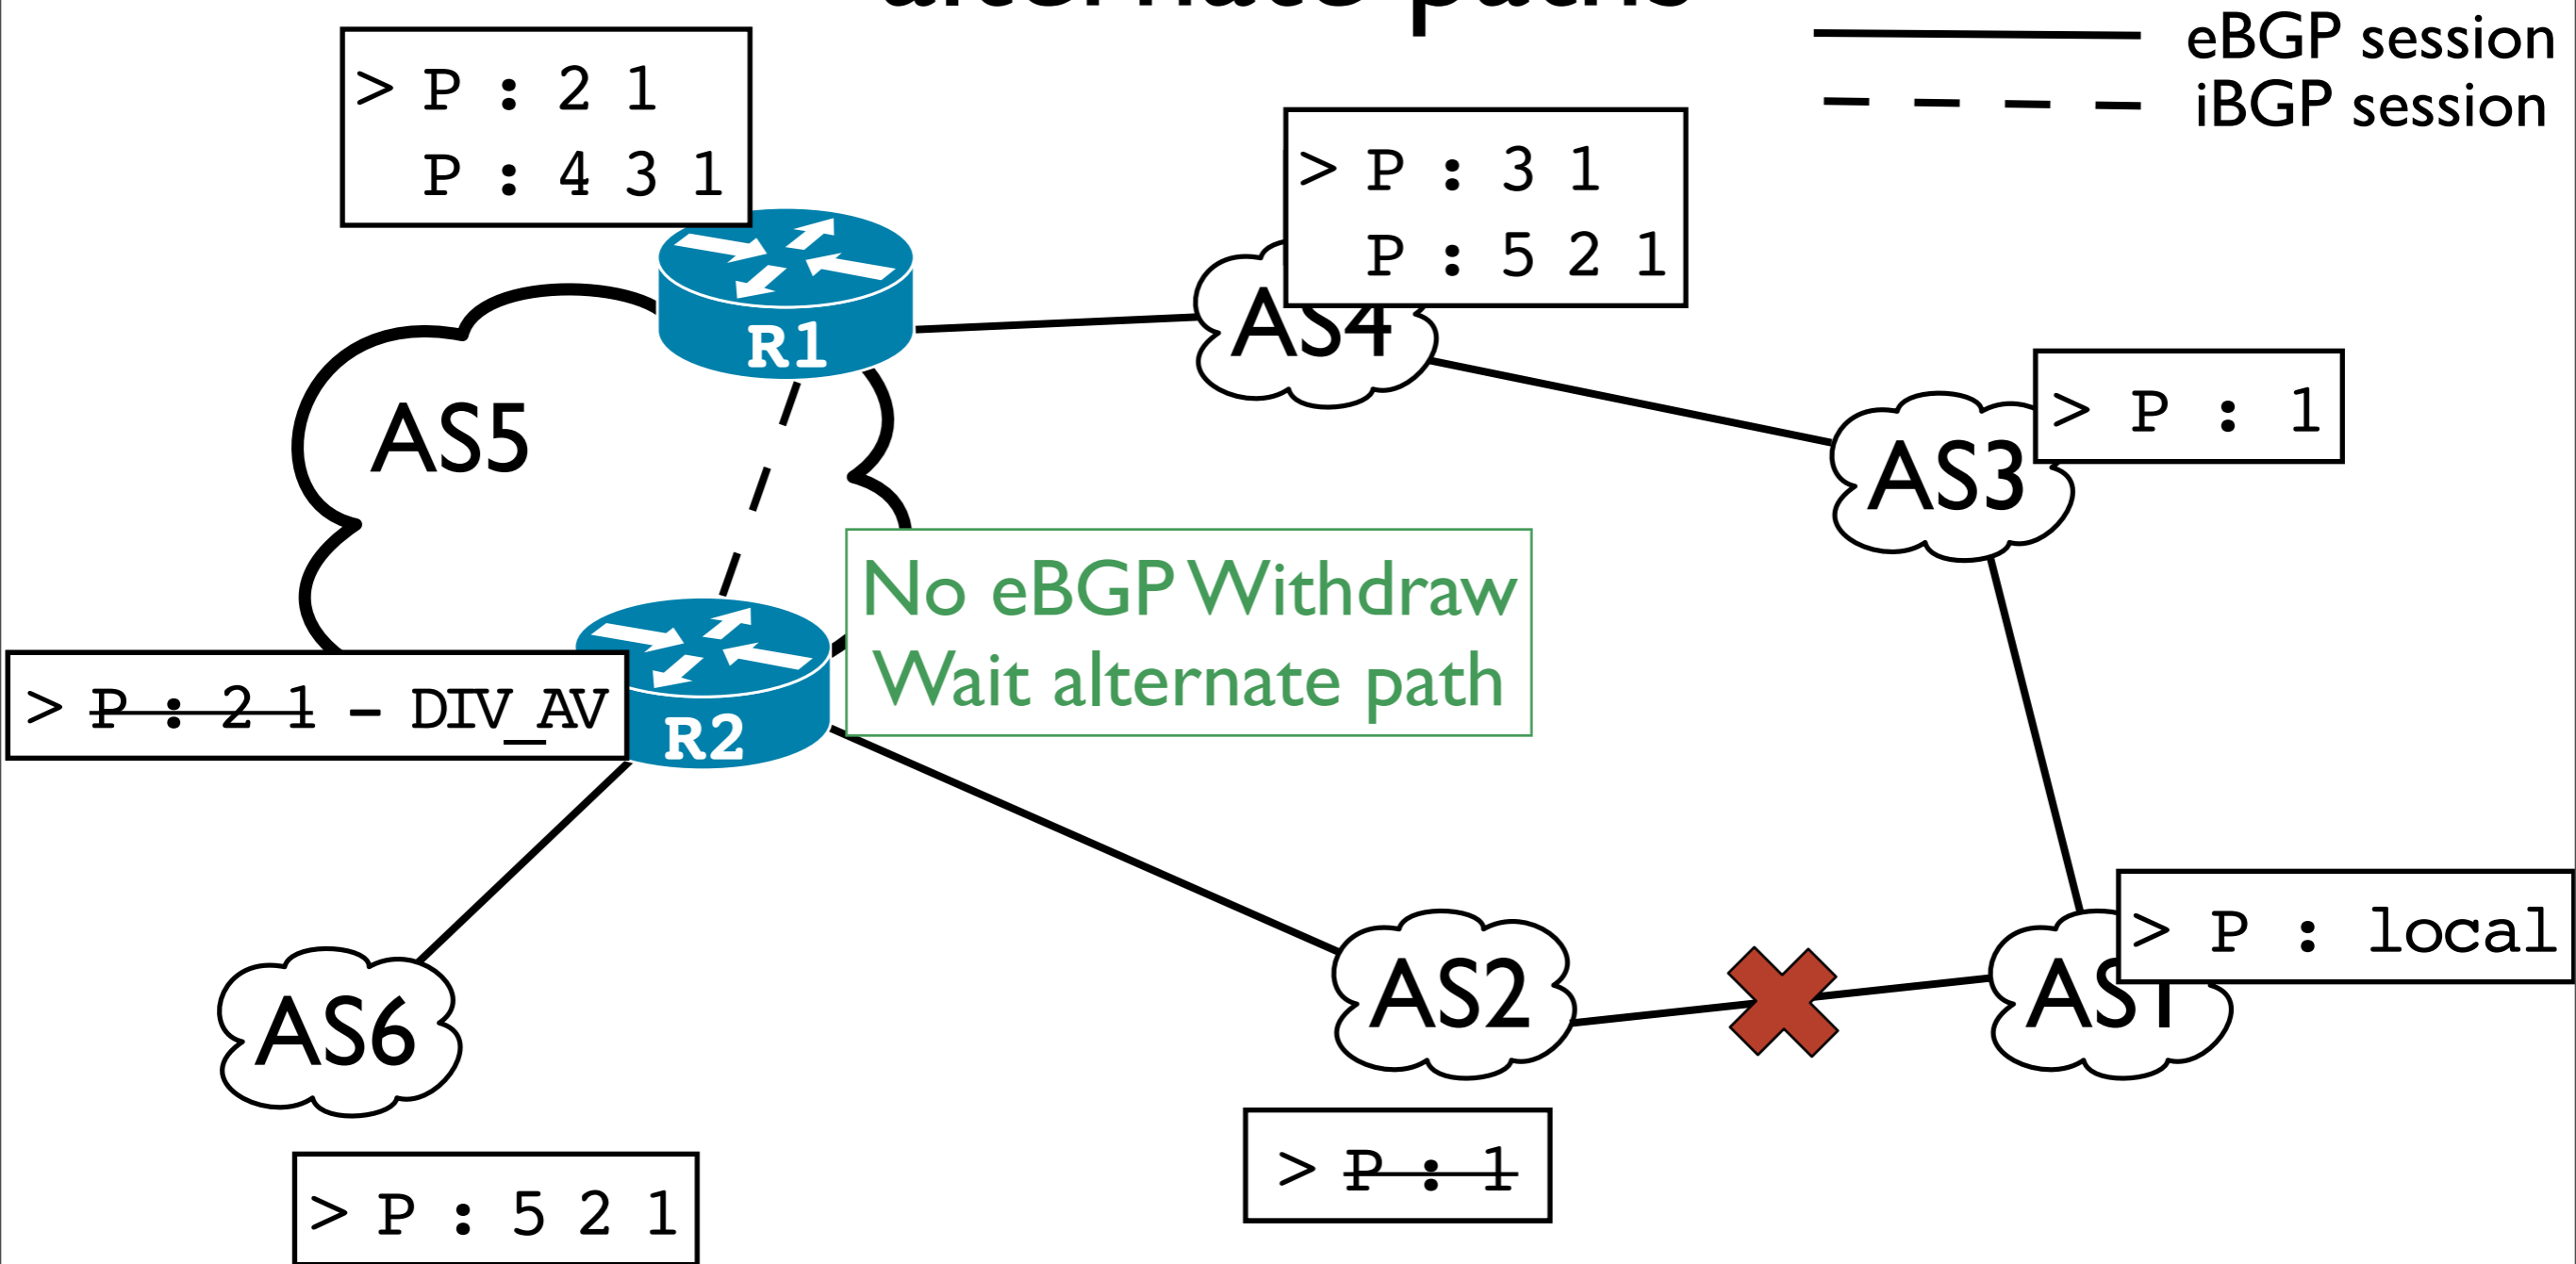
Advertising the availability of alternate paths



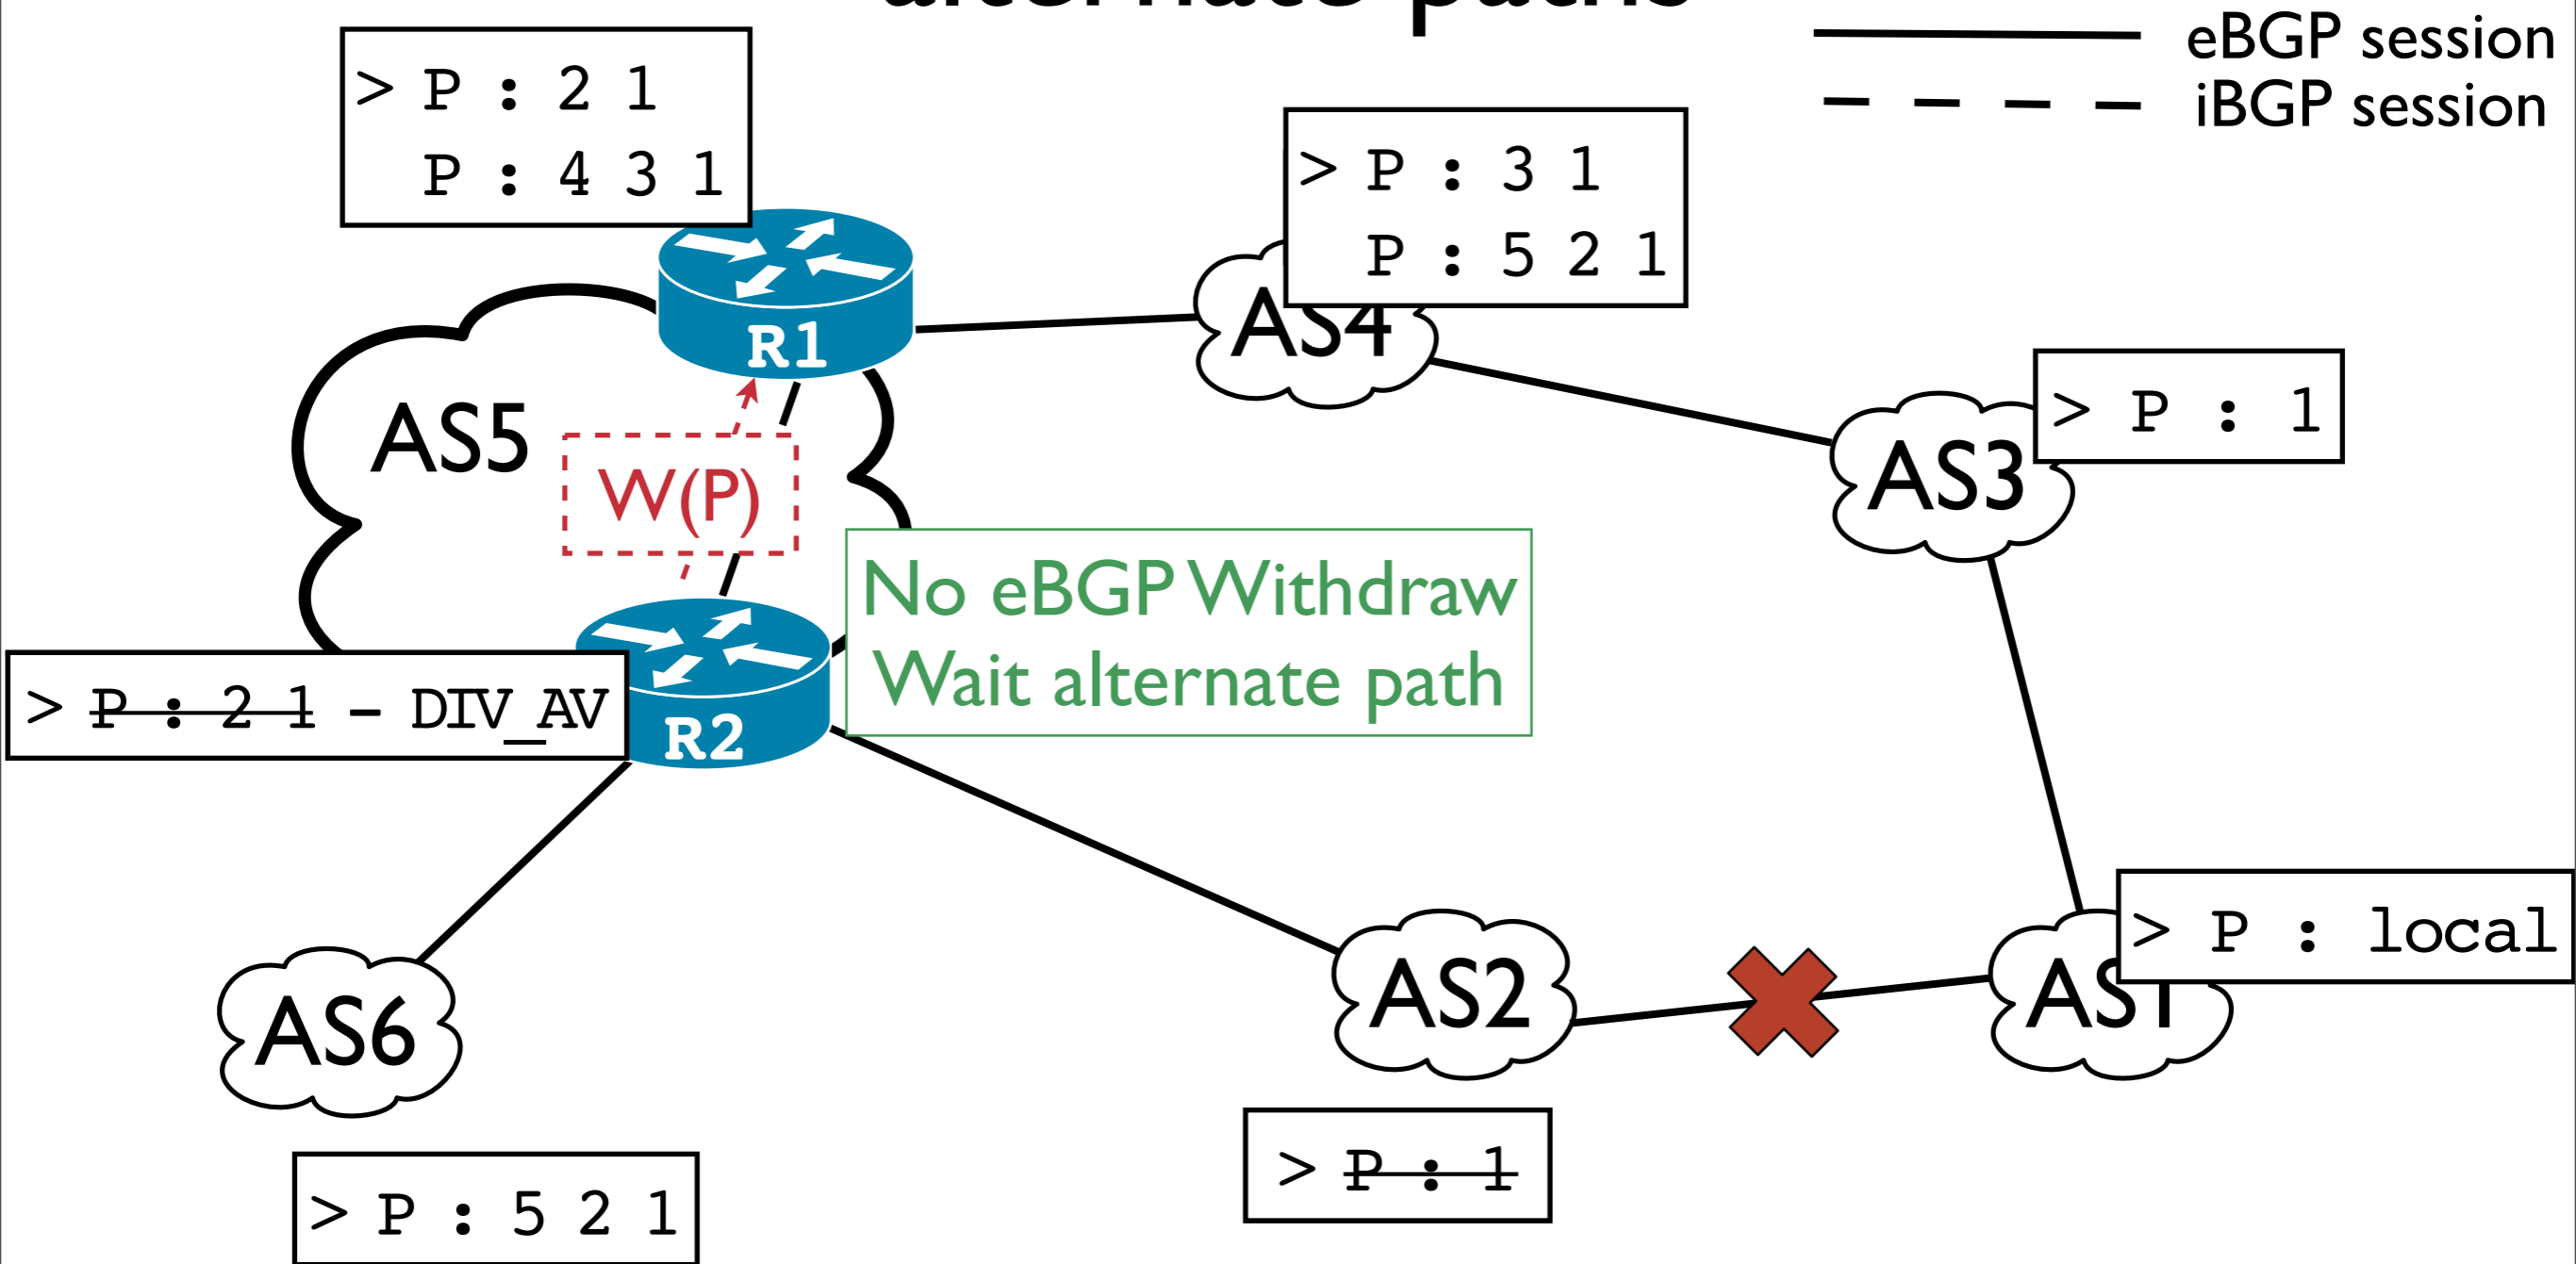
Advertising the availability of alternate paths



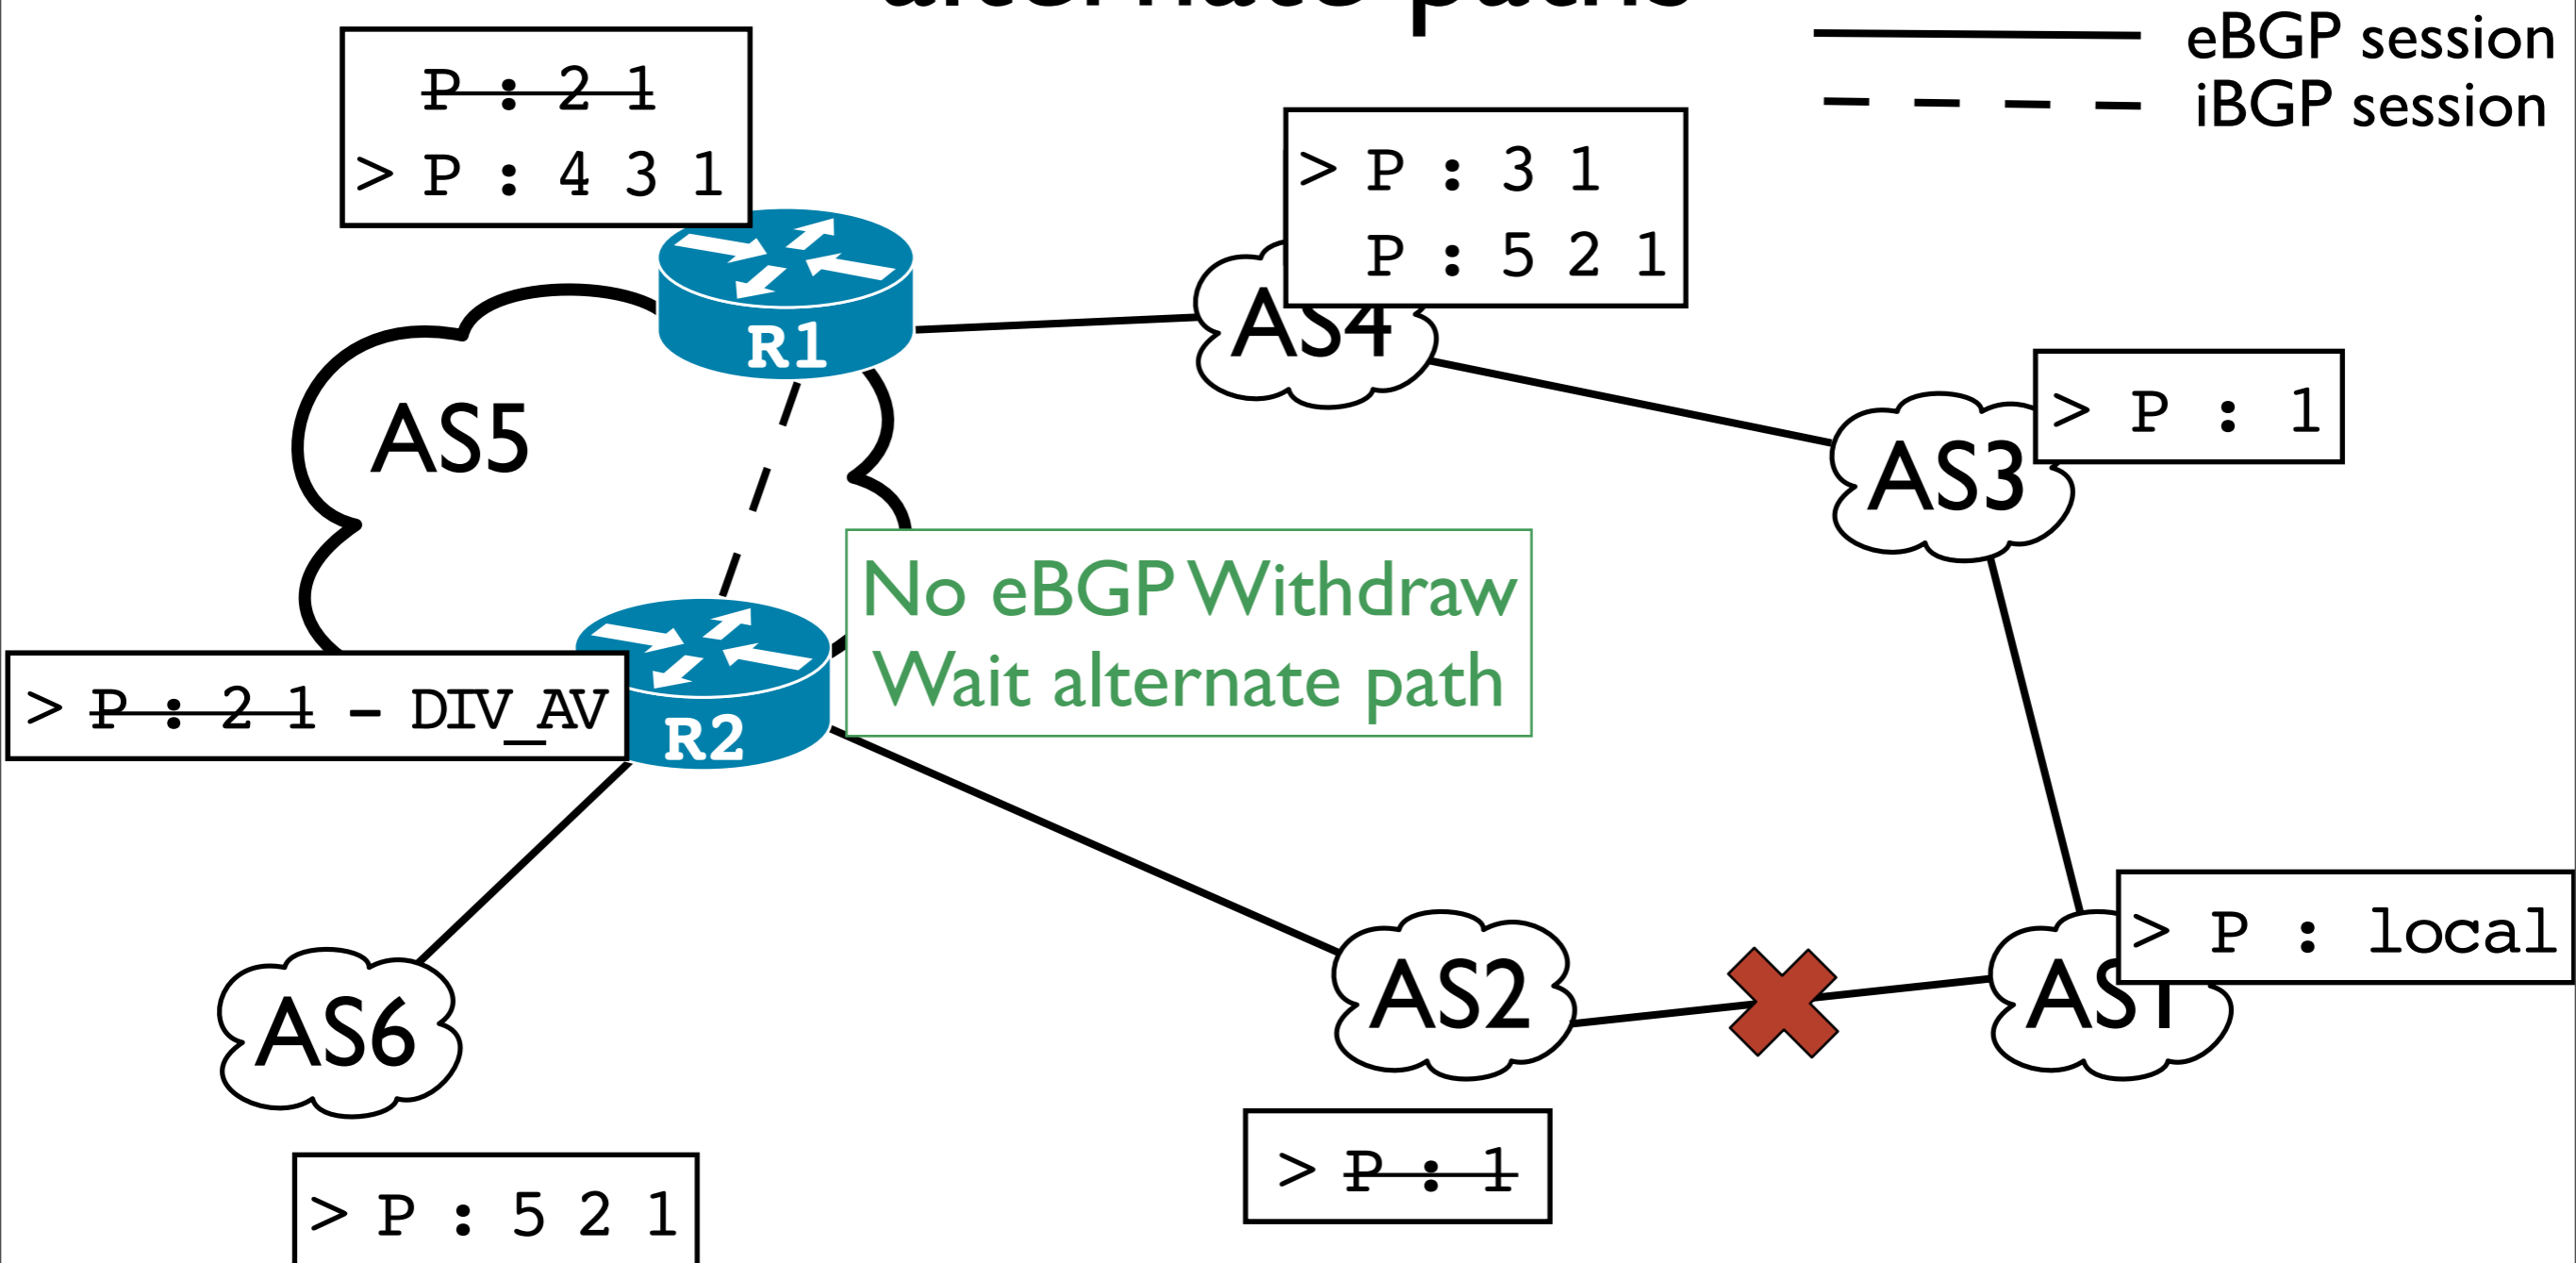
Advertising the availability of alternate paths



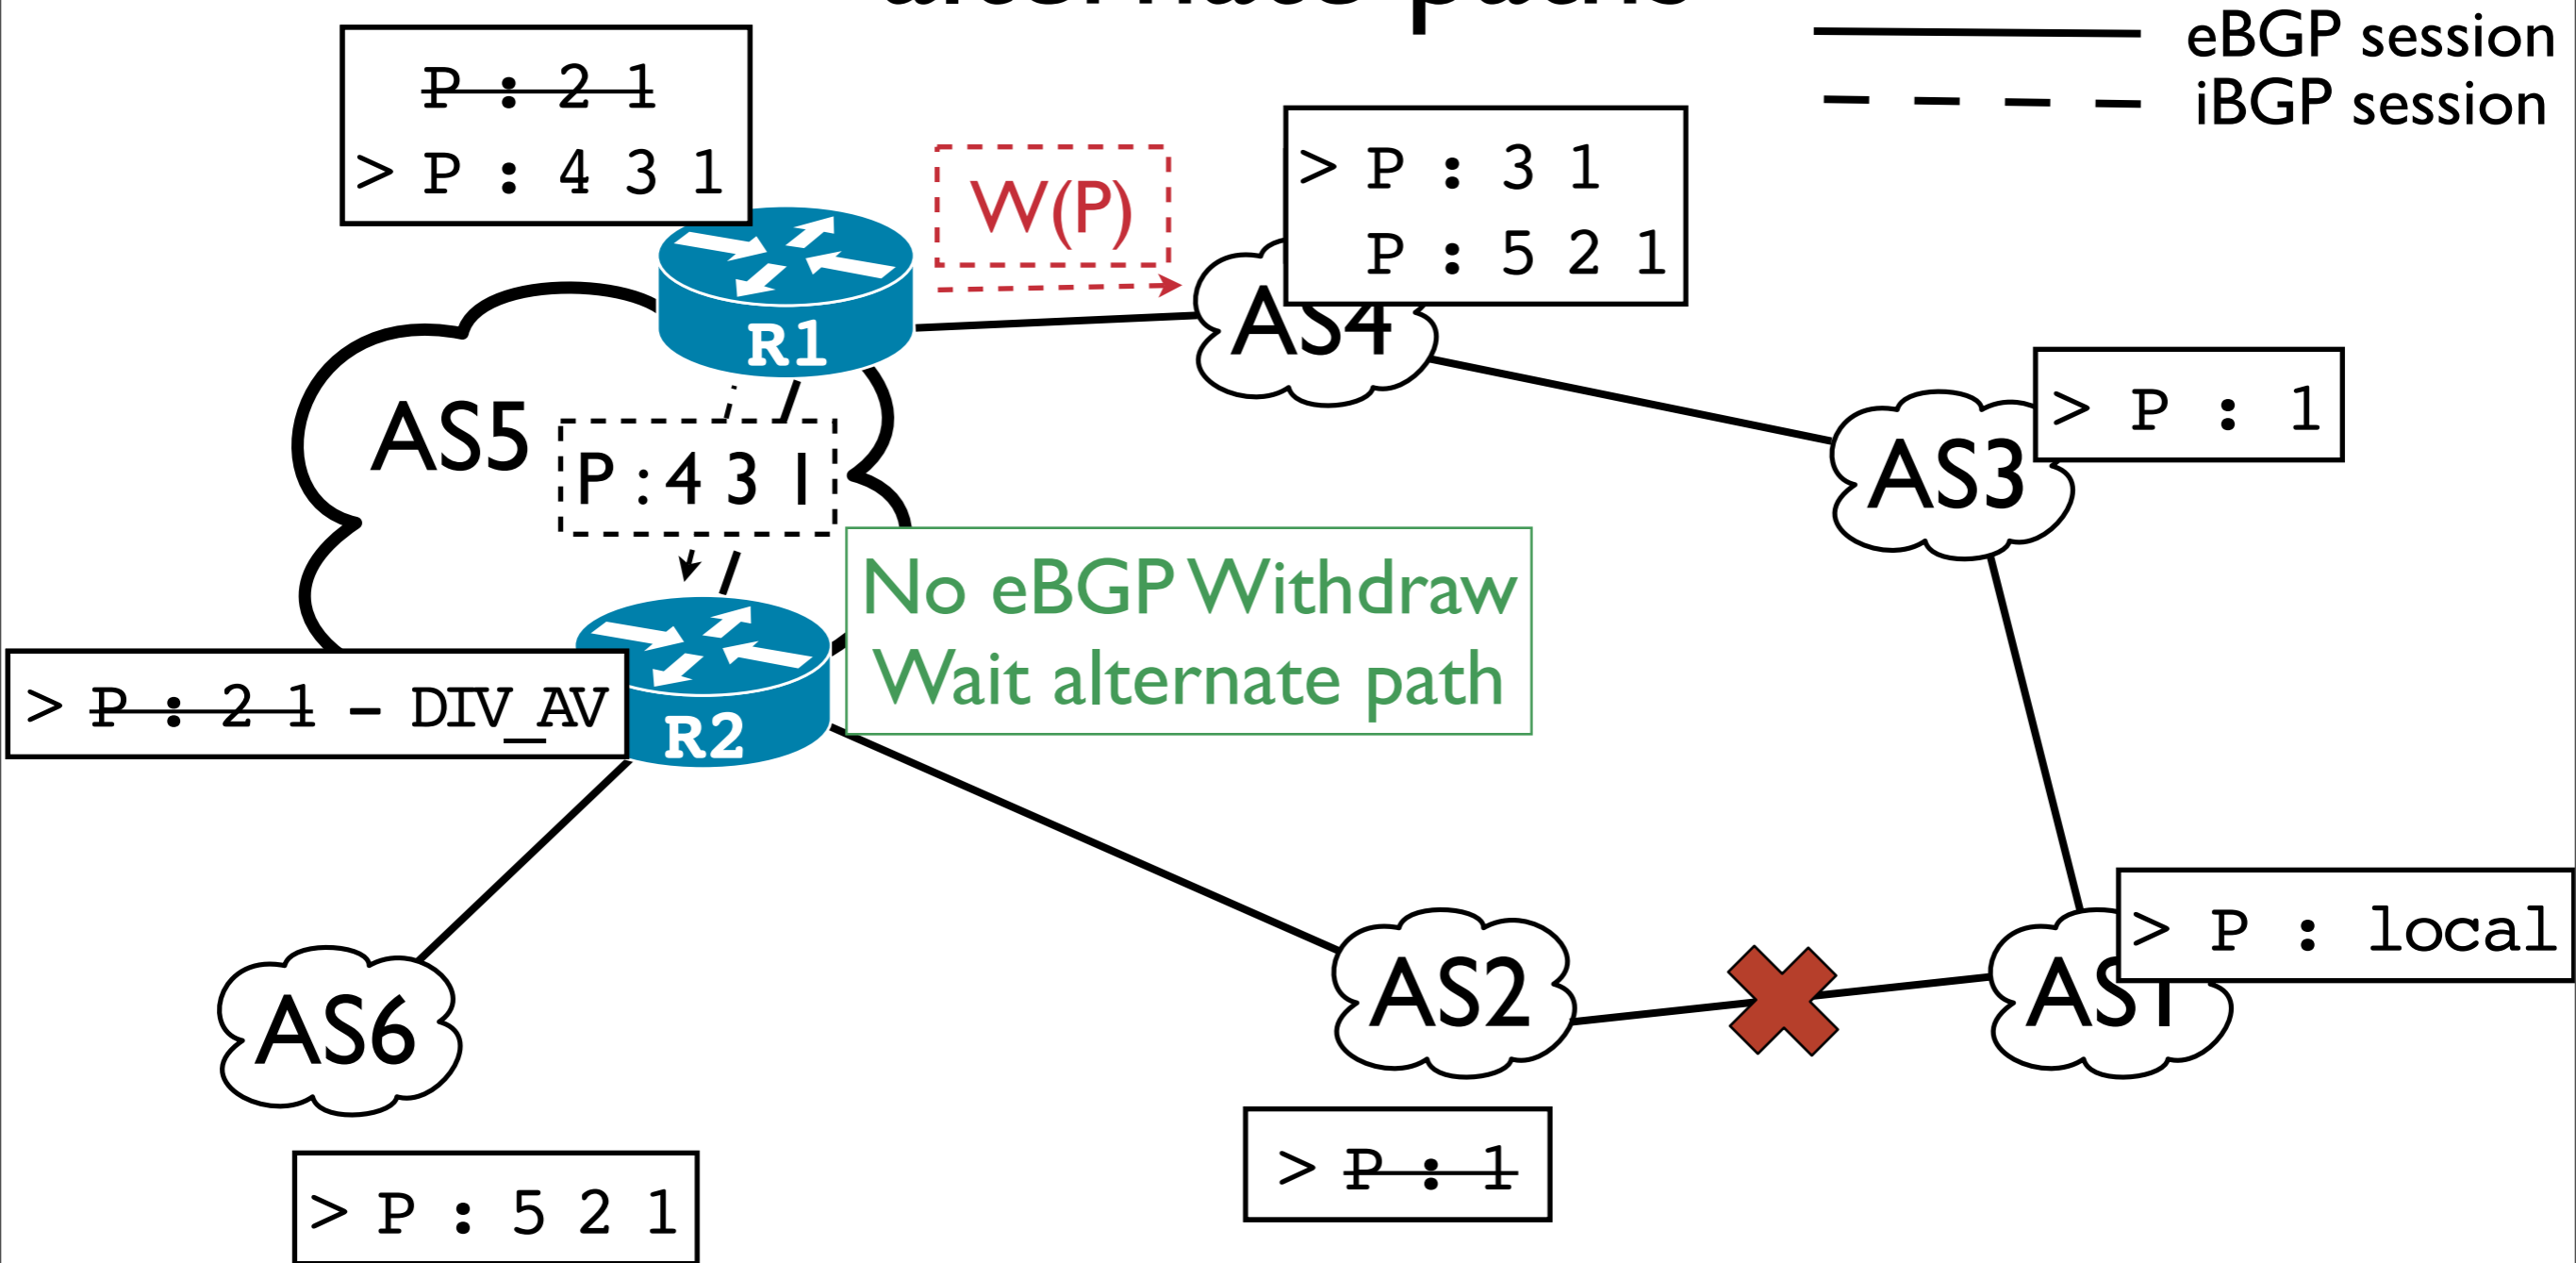
Advertising the availability of alternate paths



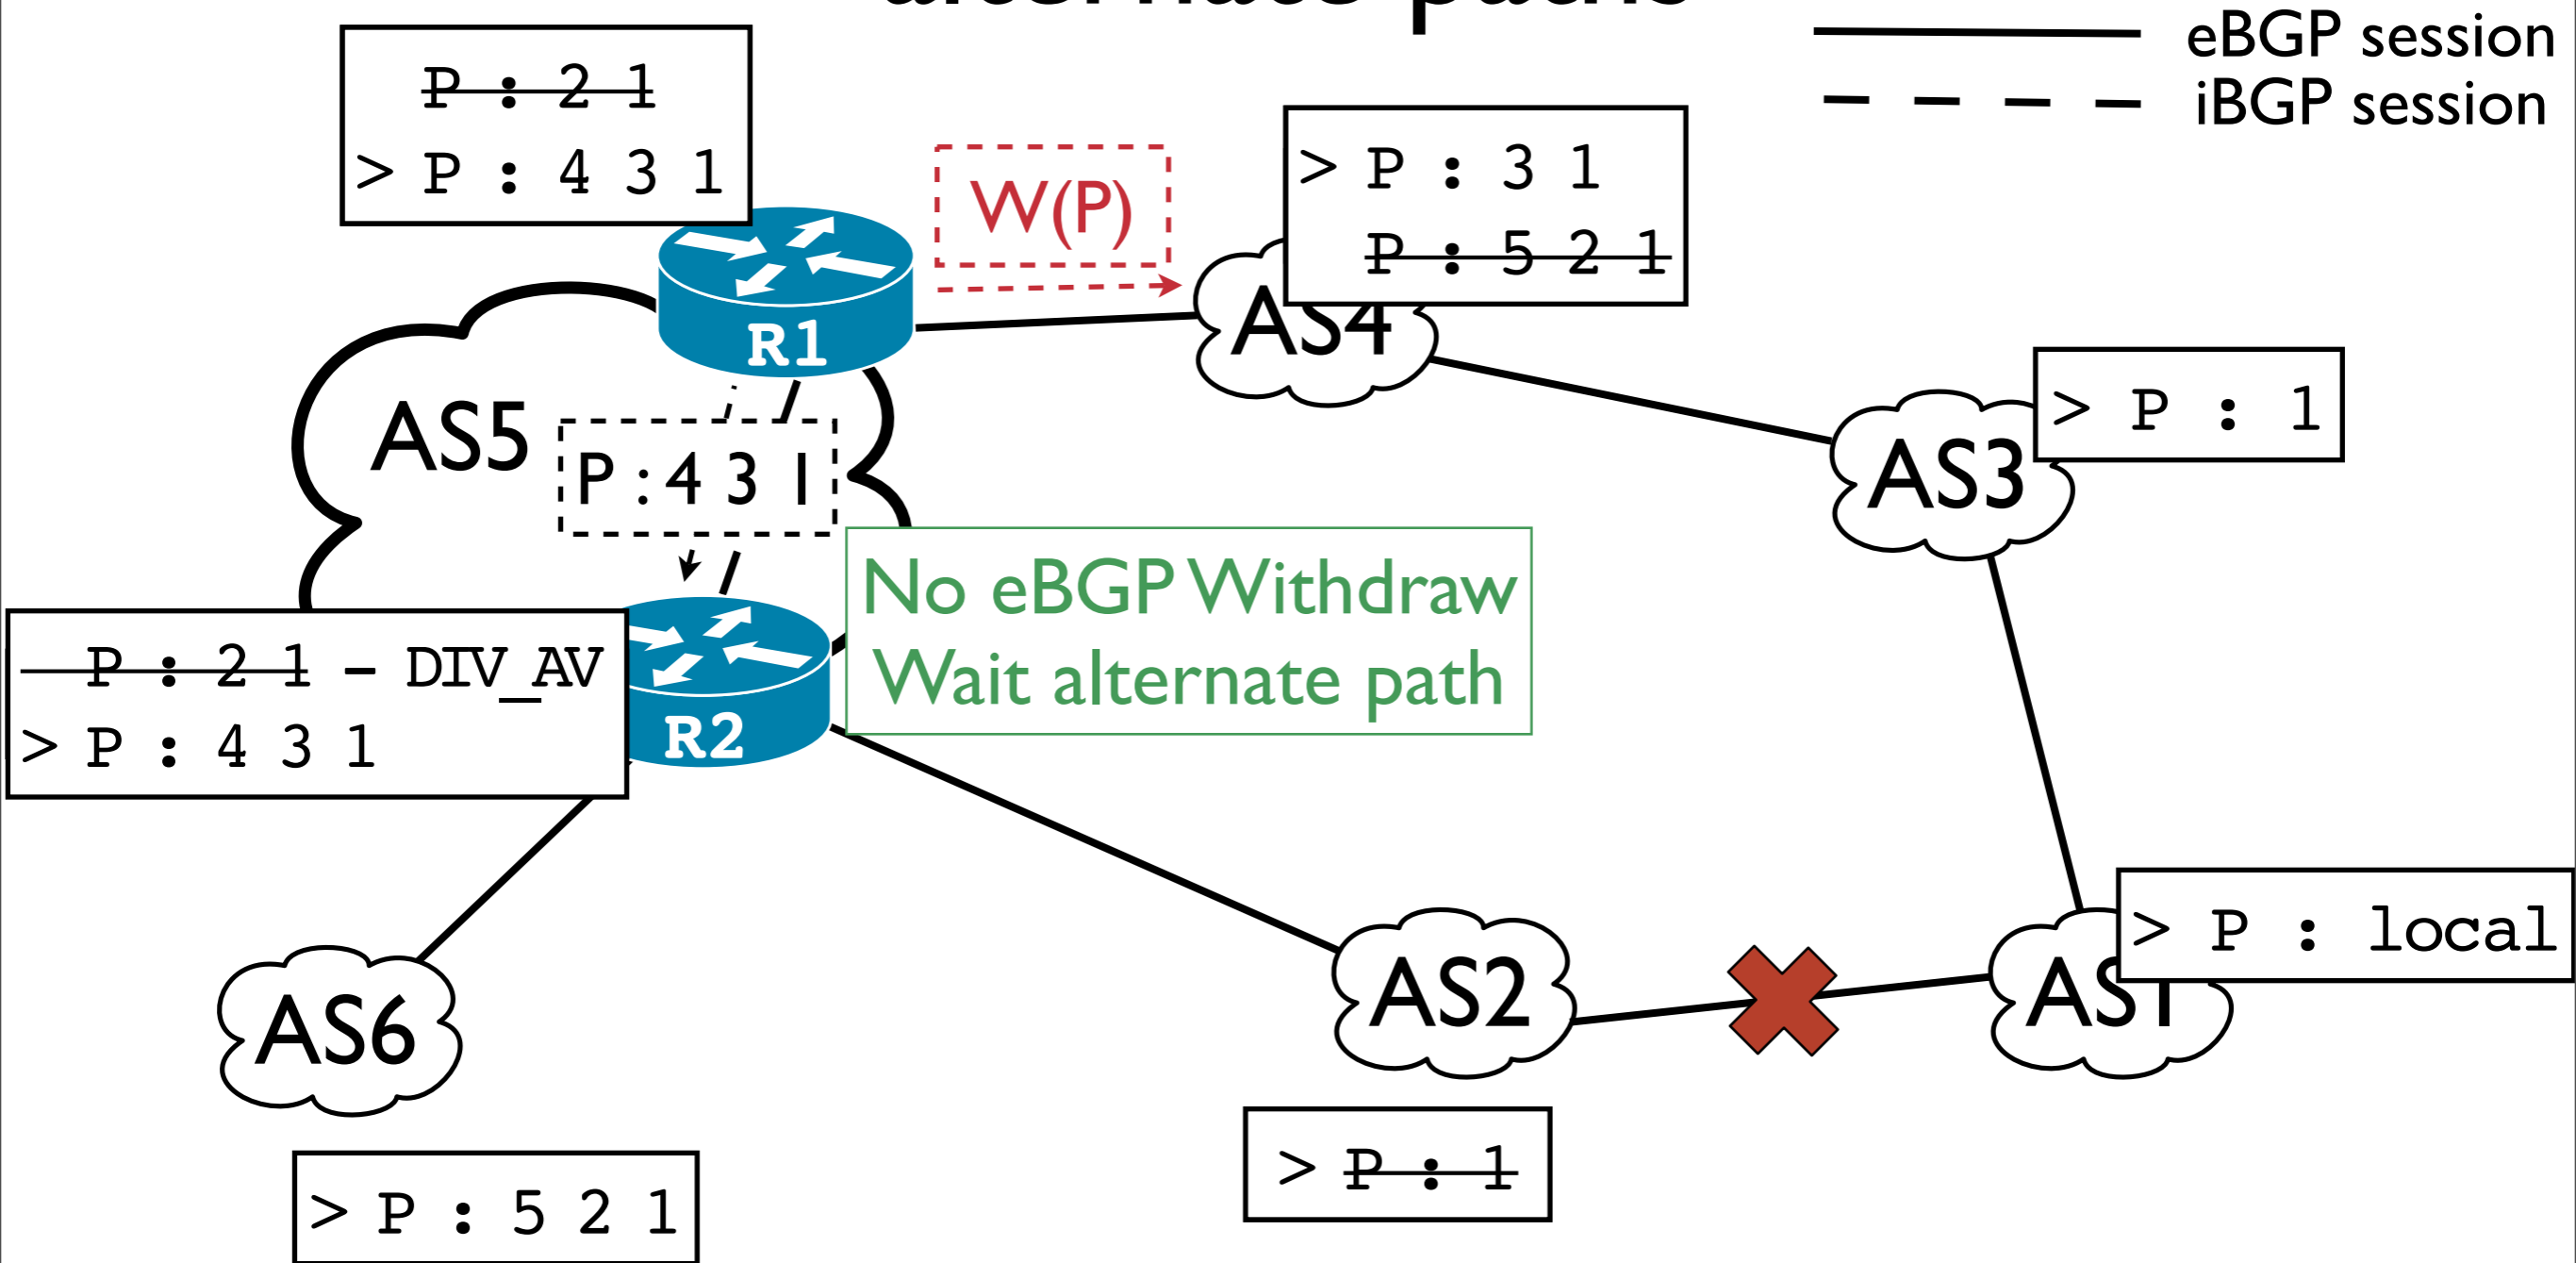
Advertising the availability of alternate paths



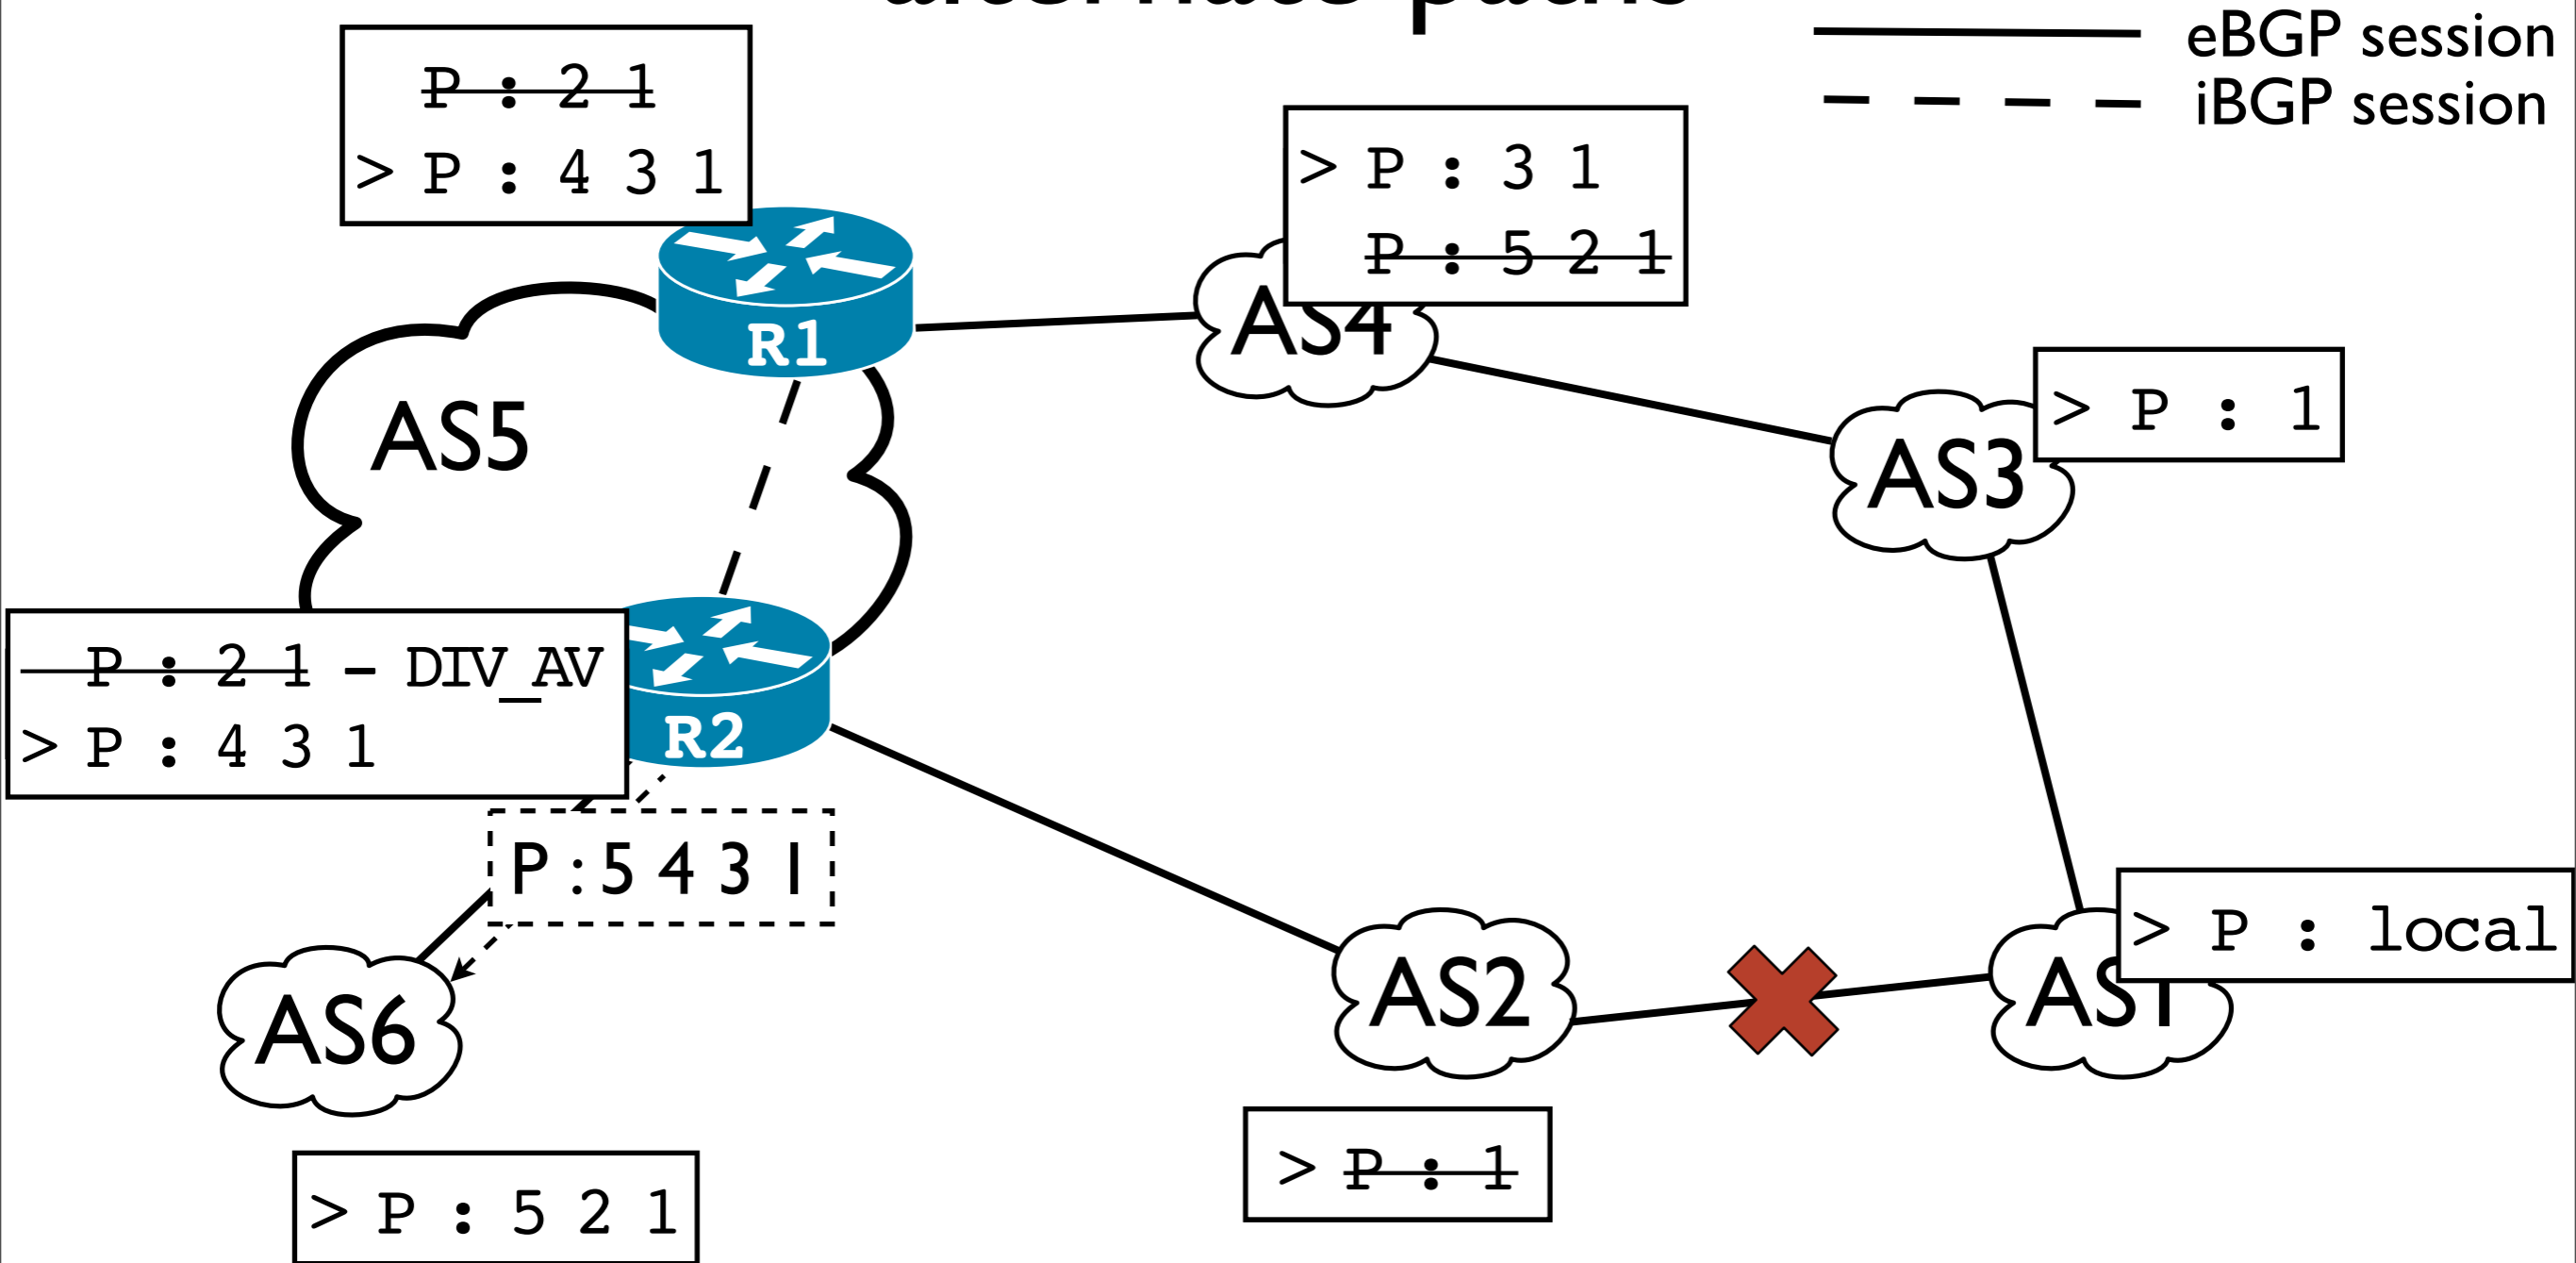
Advertising the availability of alternate paths



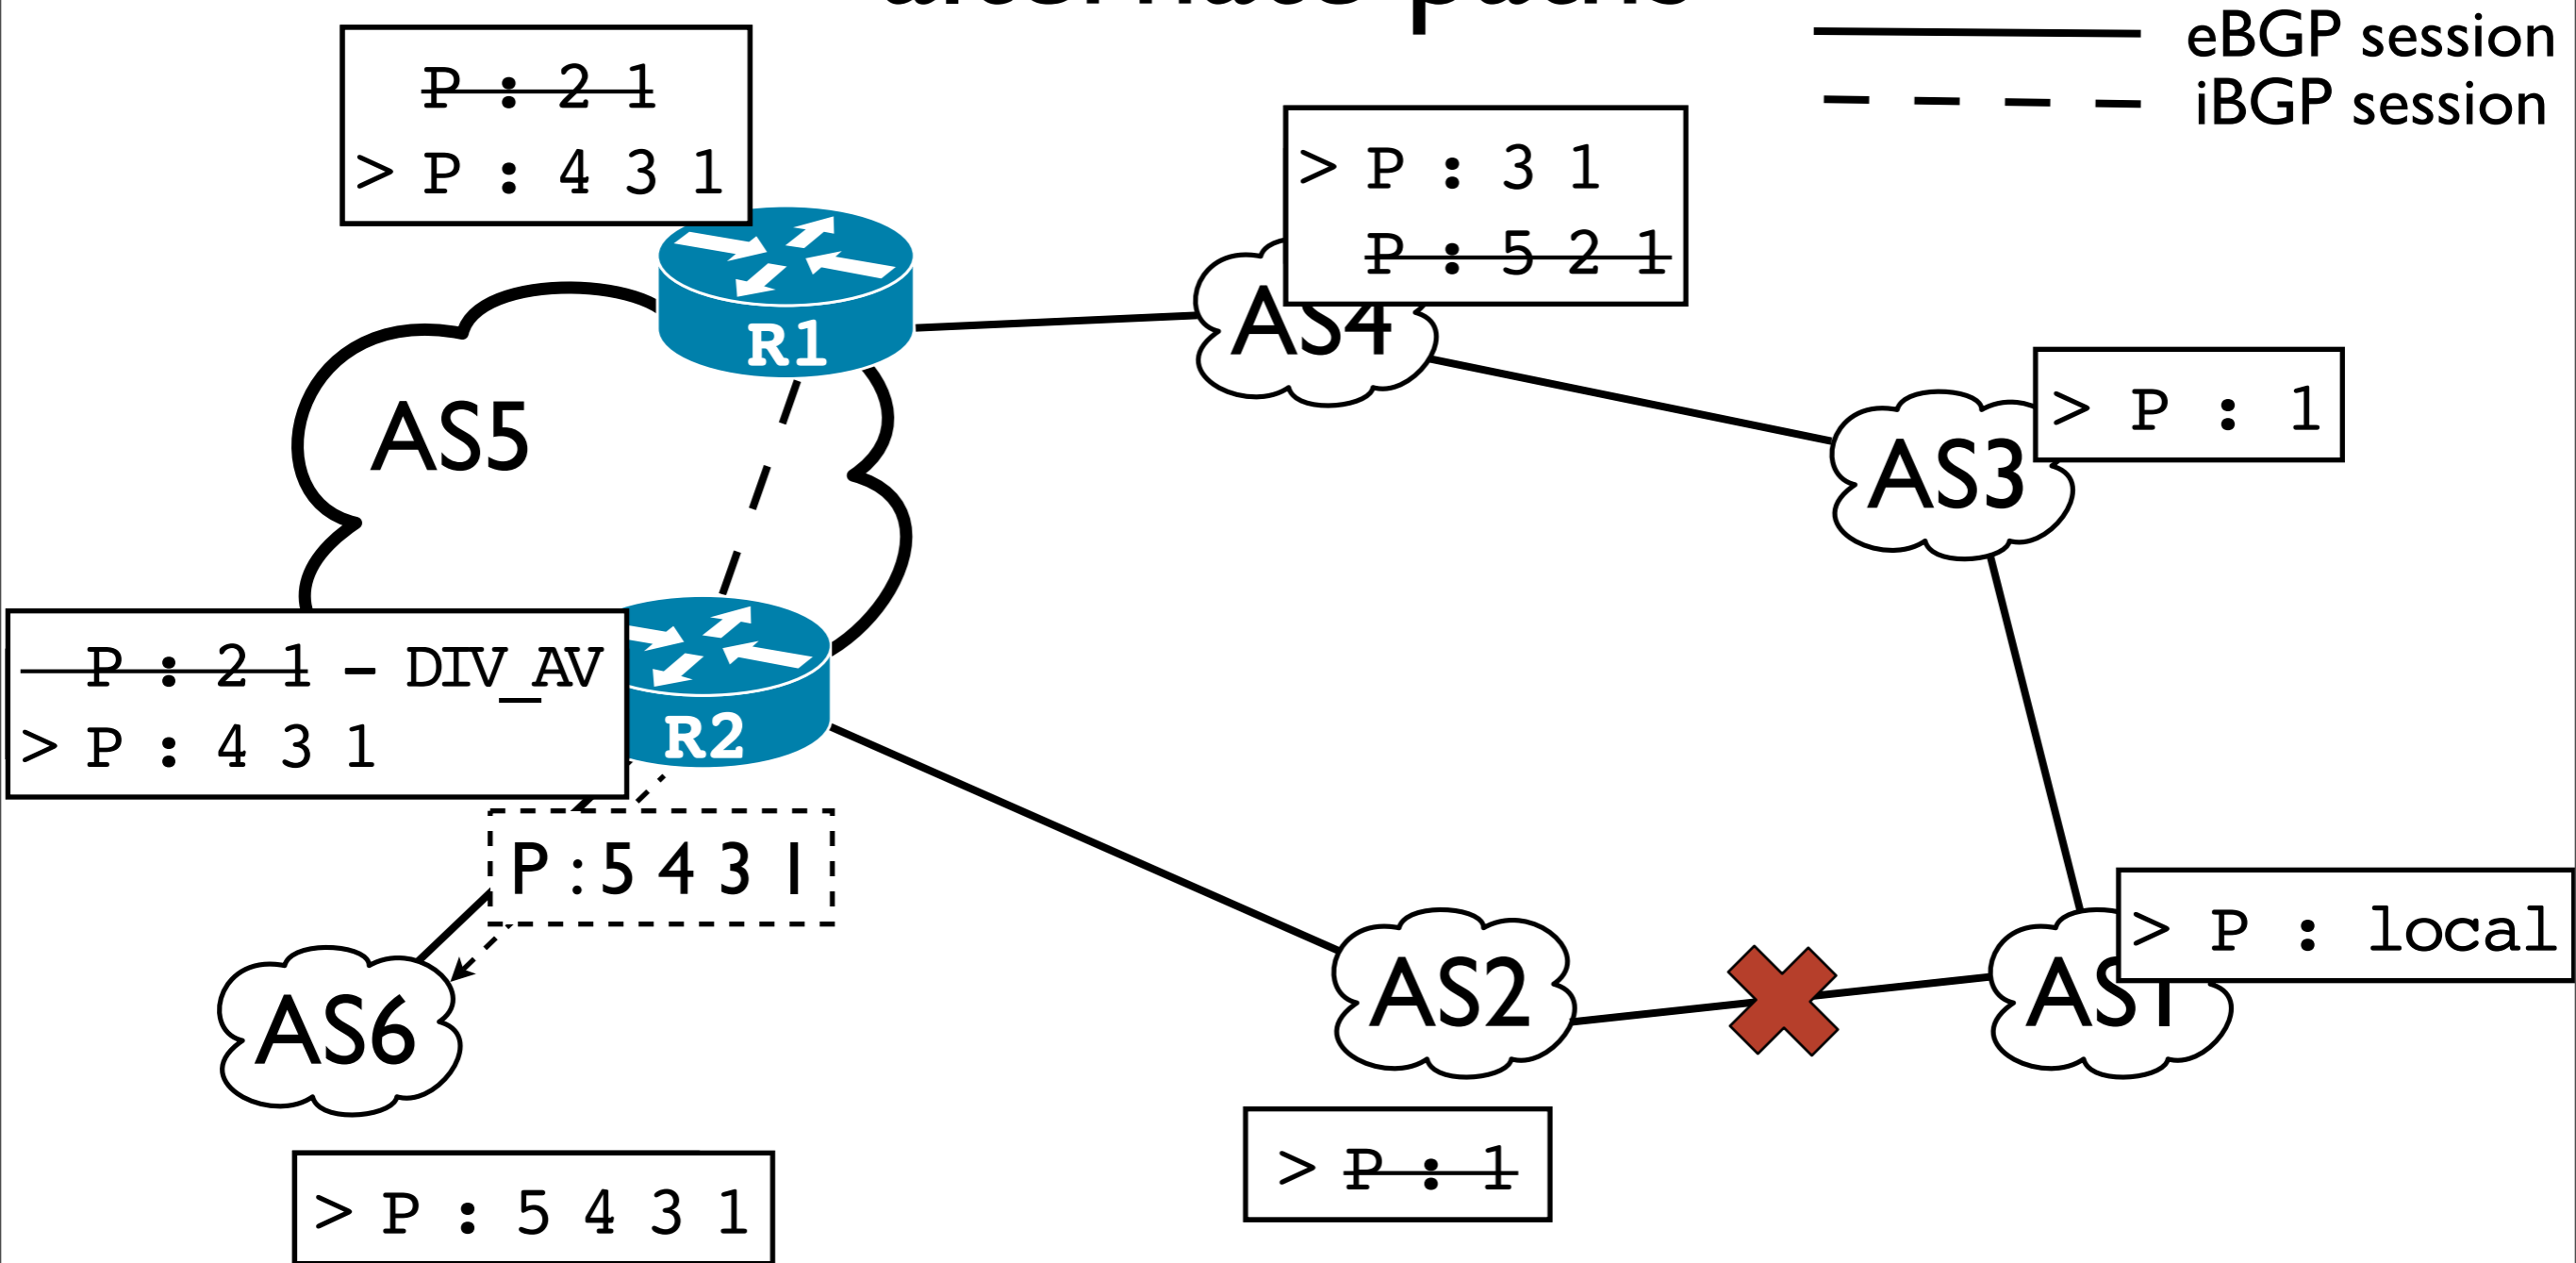
Advertising the availability of alternate paths



Advertising the availability of alternate paths



Advertising the availability of alternate paths



Policy-equivalent paths

- In practice : due to policies, alternate paths can still trigger eBGP withdraws
 - ✓ Use two communities instead of one

Conclusion

- We highlighted the contribution of iBGP-caused withdraws to BGP churn
- Two solutions to prevent this unnecessary withdraw propagation :
 - ✓ Add-Paths (IETF)
 - ✓ Our proposal : Diversity communities
- Work in progress : Evaluate the reduction of churn resulting from using our solution
 - ✓ Simulations on generated Internet topologies

Questions?

- More information on : <http://inl.info.ucl.ac.be>