BGP Safety with Spurious Updates

The Conditions of BGP Convergence

Martin Suchara

in collaboration with:
Alex Fabrikant and
Jennifer Rexford

January 12, 2011

Local Control vs. Global Properties

- The Internet is a "network of networks"
  - ~35,000 independently administered ASes
  - Competitive cooperation to find routes

Local Control

Interdomain policies,
Intradomain routing

Global Properties

Stability, reliability,
performance, etc.

Interdomain Routing

- Autonomous systems (AS) have different goals
  - Different views on which path is best
- Interdomain routing: agree on a set of paths

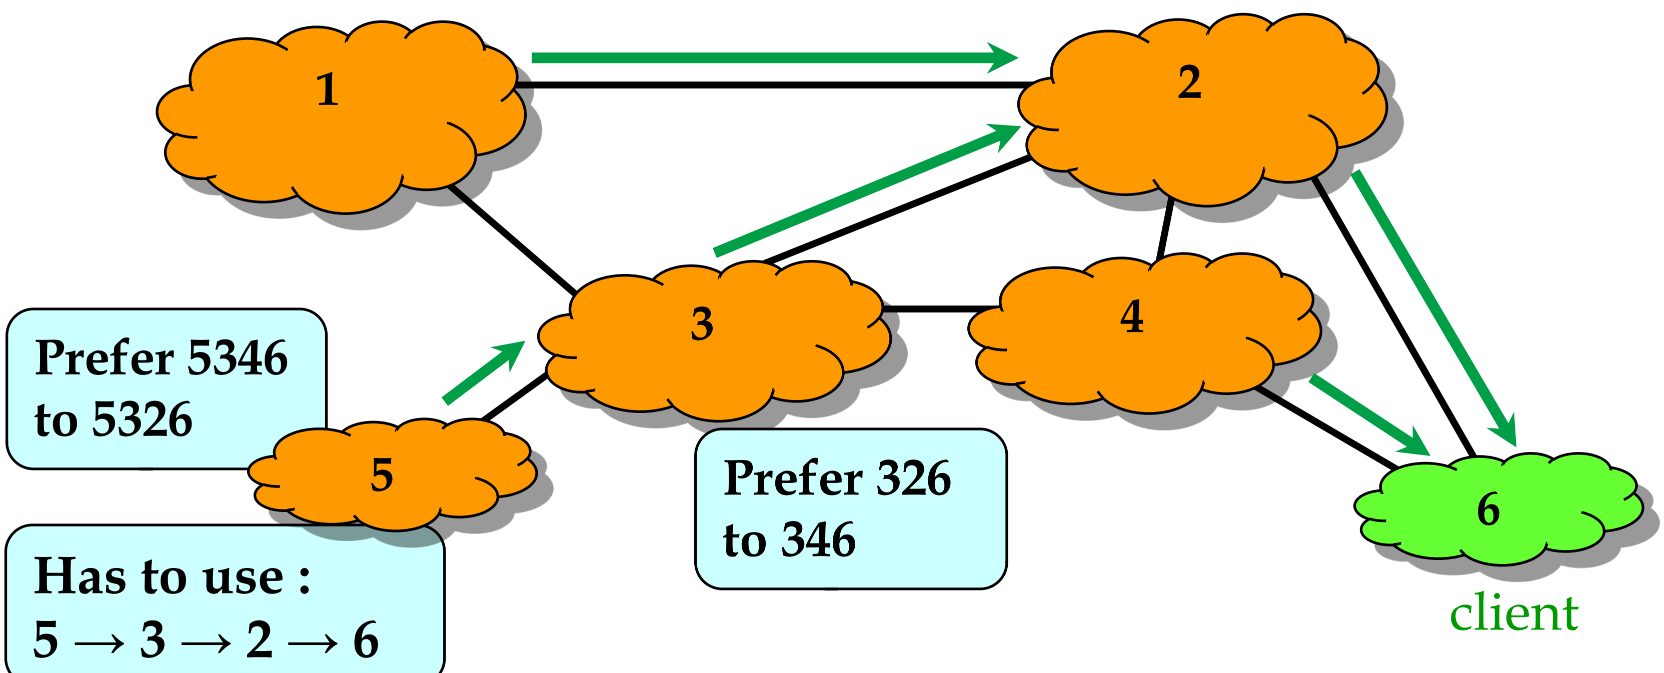

The Border Gateway Protocol (BGP)

- ASes exchange information about paths

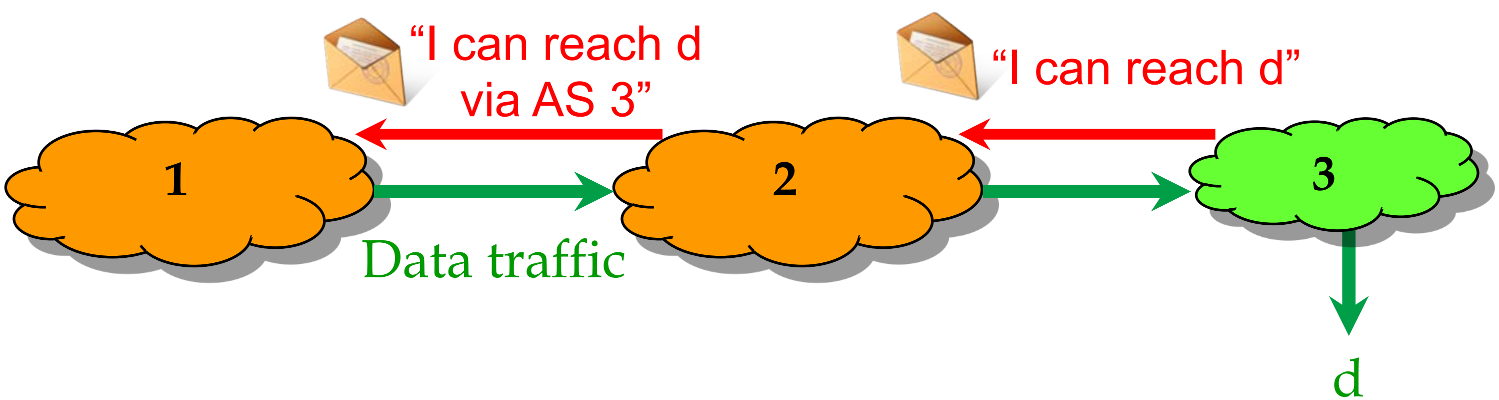

- Policy configurations provided by AS operators
  - Path selection: which path do I choose?
  - Path export: which neighbors do I tell?

Business Driven Policies of ASes

- Customer-Provider Relationship
  - Provider exports its customer's routes to everybody
  - Customer exports provider's routes only to downstream customers

- Peer-Peer Relationship
  - Export only customer routers to a peer
  - Export peer routes only to customers

BGP Safety Challenges

- 35,000 ASes and 300,000 address blocks
  - Flexible AS policies
- Routing convergence usually takes minutes
  - But the system does not always converge…

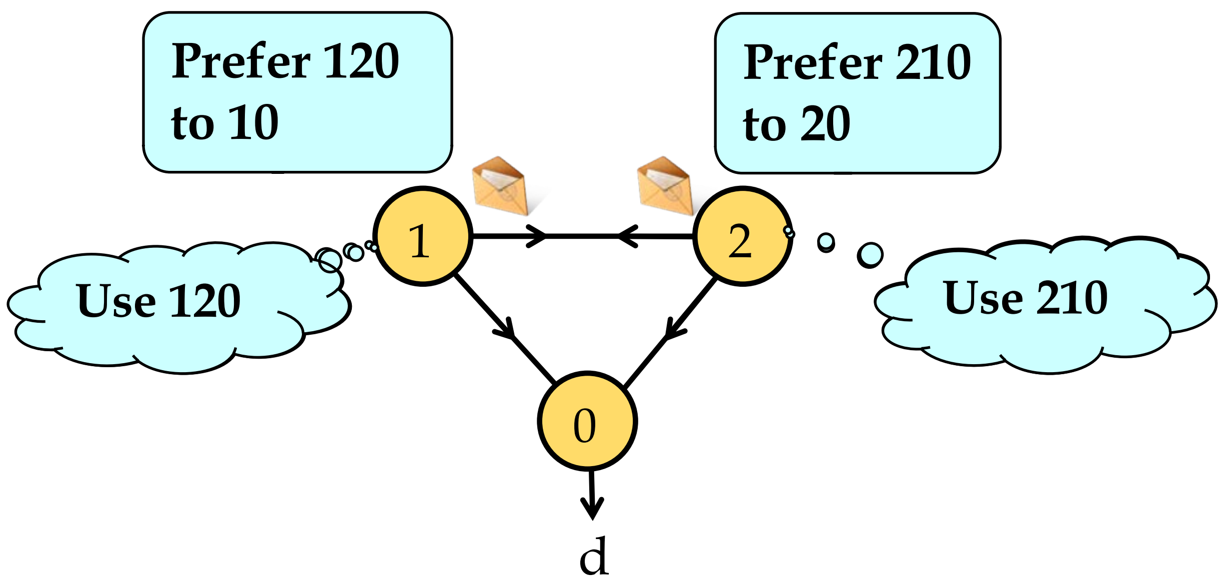

Results on BGP Safety

- Safety verification important to network operators
- Absence of a "dispute wheel" sufficient for safety *(Griffin, Shepherd, Wilfong, 2002)*
- Necessary or sufficient conditions of safety
  (Gao and Rexford, 2001), (Gao, Griffin and Rexford, 2001), (Griffin, Jaggard and Ramachandran, 2003), (Feamster, Johari and Balakrishnan, 2005), (Sobrinho, 2005), (Fabrikant and Papadimitriou, 2008), (Cittadini, Battista, Rimondini and Vissicchio, 2009), …

Models of BGP

- Existing models (variants of SPVP)
  - Widely used to analyze BGP properties
  - Simple but do not capture spurious behavior of BGP
- This work
  - A new model of BGP with spurious updates
  - Spurious updates have major consequences
  - More accurate model makes proofs easier!

Overview

I. Classical model of BGP: *the SPVP*

II. Spurious BGP updates: *what are they?*

III. The surprise: *networks believed to be safe oscillate!*

IV. The consequences: *applicability of earlier results*

V. Convergence conditions: *polynomial time verifiable*

VI. Conclusion

SPVP– Traditional Model of BGP

(Griffin and Wilfong, 2000)

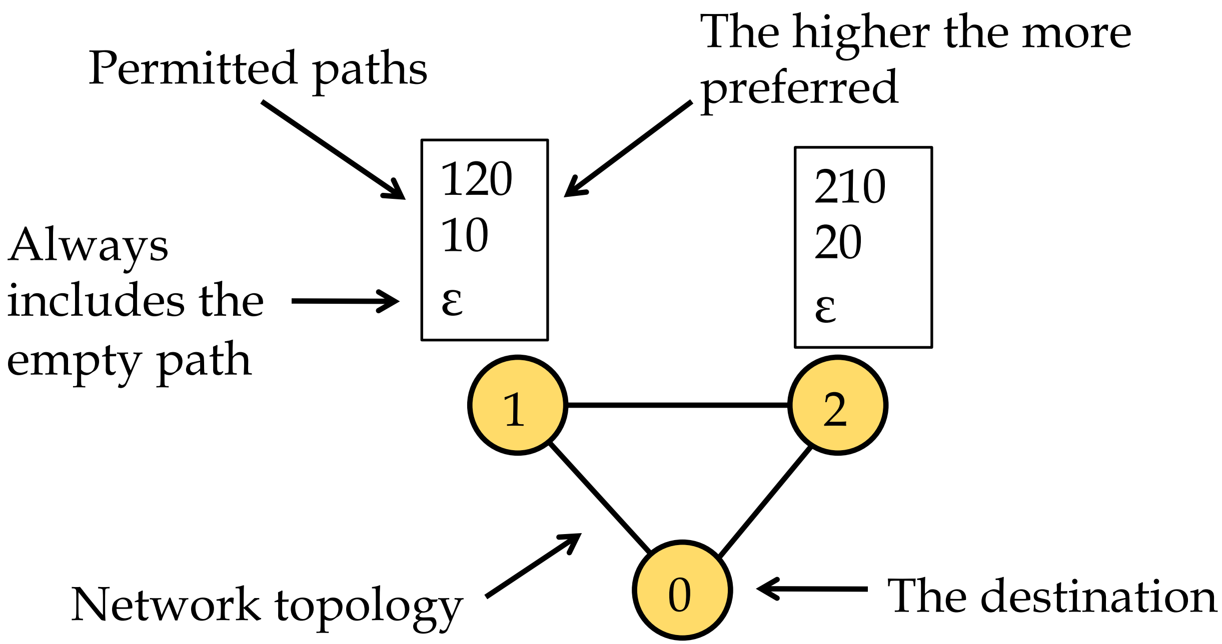

SPVP– Traditional Model of BGP

(Griffin and Wilfong, 2000)

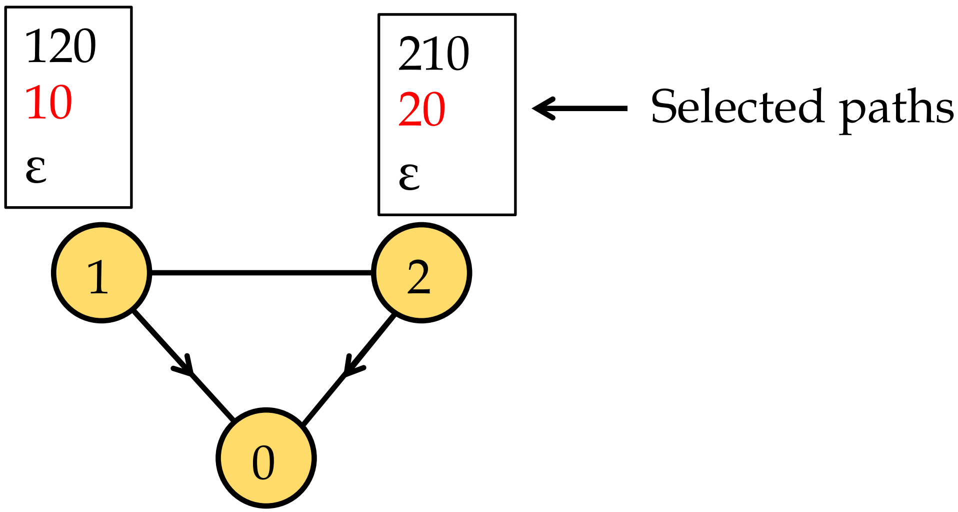

SPVP– Traditional Model of BGP

(Griffin and Wilfong, 2000)

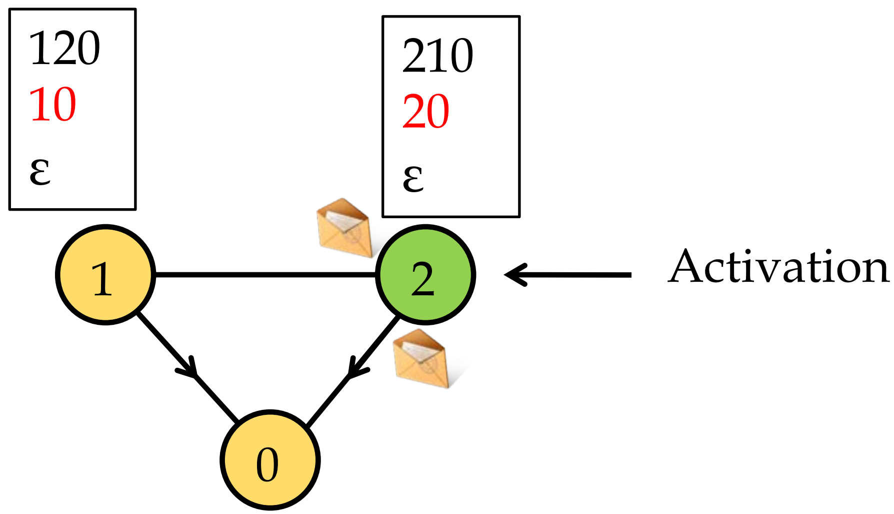

- Activation models the processing of BGP update messages sent by neighbors
- Vertex or edge activations

SPVP– Traditional Model of BGP

(Griffin and Wilfong, 2000)

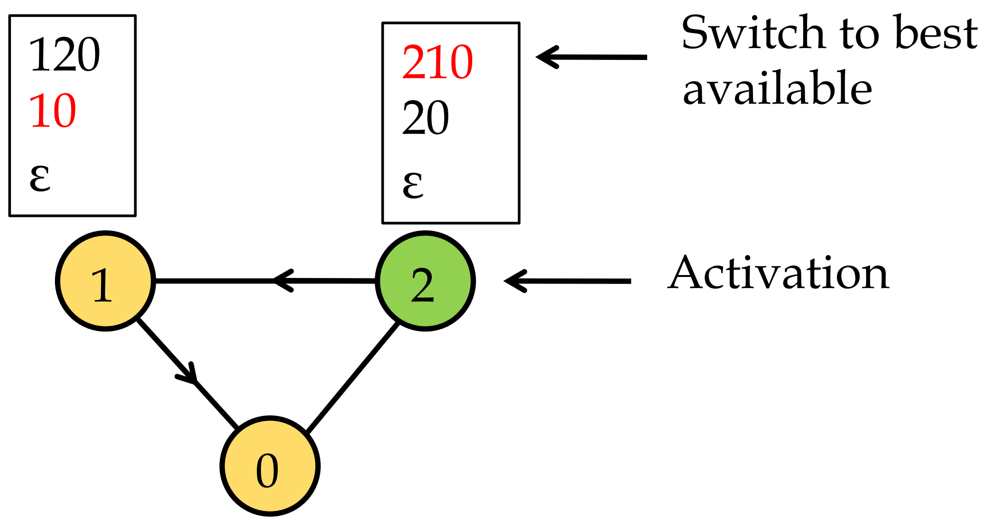

- Activation models the processing of BGP update messages sent by neighbors
- Vertex or edge activations

SPVP– Traditional Model of BGP

(Griffin and Wilfong, 2000)

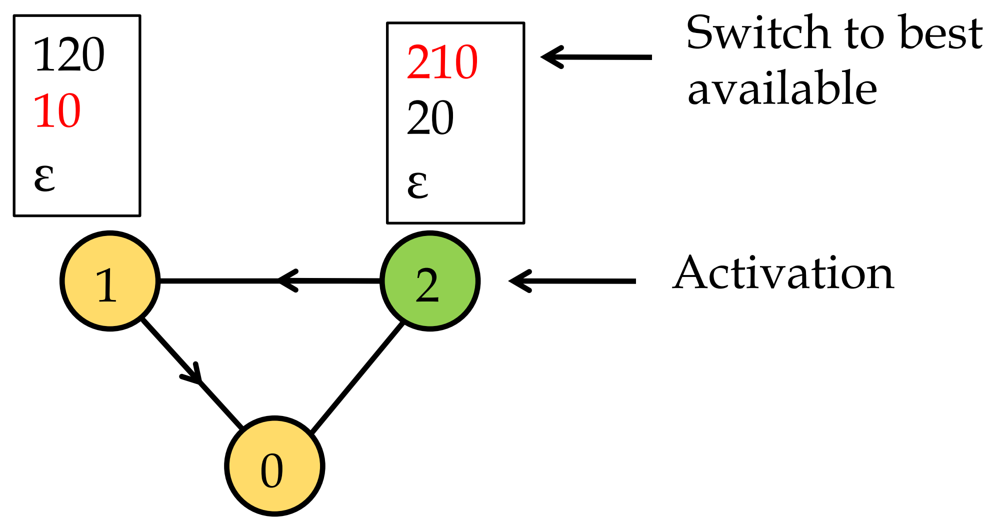

- System is safe if all “fair” activation sequences lead to a stable path assignment

Overview

I. Classical model of BGP: *the SPVP*

II. Spurious BGP updates: *what are they?*

III. The surprise: *networks believed to be safe oscillate!*

IV. The consequences: *applicability of earlier results*

V. Convergence conditions: *polynomial time verifiable*

VI. Conclusion

What are Spurious Updates?

- A phenomenon: router announces a route other than the highest ranked one

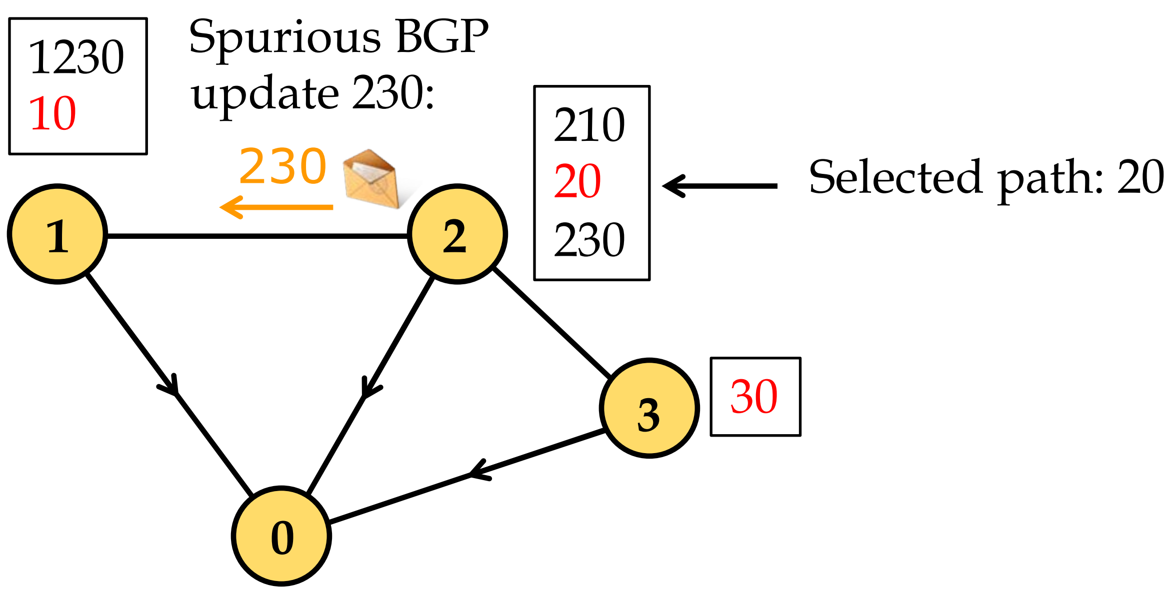

- Behavior not allowed in SPVP

What Causes Spurious Updates?

1. Limited visibility to improve scalability
   - Internal structure of ASes
   - Cluster-based router architectures
2. Timers and delays to prevent instabilities and reduce overhead
   - Route flap damping
   - MRAI timers
   - Grouping updates to priority classes
   - Finite size message queues in routers

Cause 1 – Limited Visibility

- The internal structure of ASes improves scalability while reducing visibility

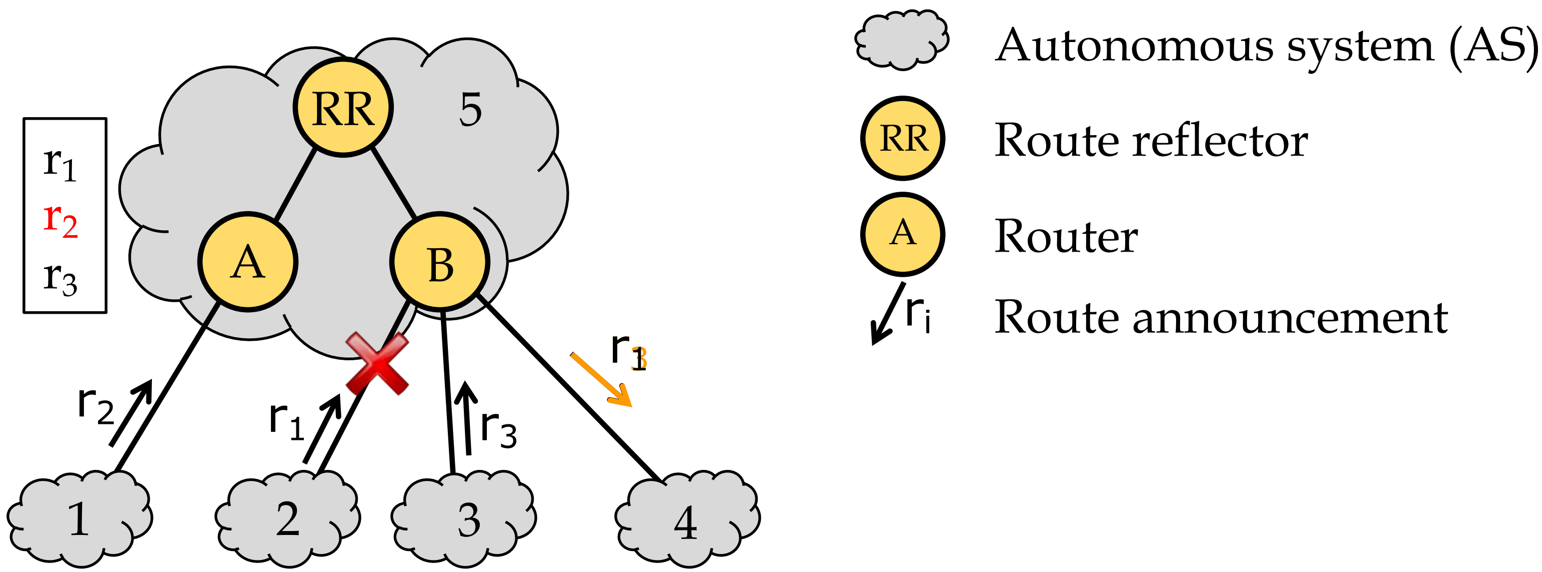

- After route r_1 is withdrawn, router B temporarily announces r_3

Cause 1 – Limited Visibility

- The internal structure of ASes improves scalability while reducing visibility

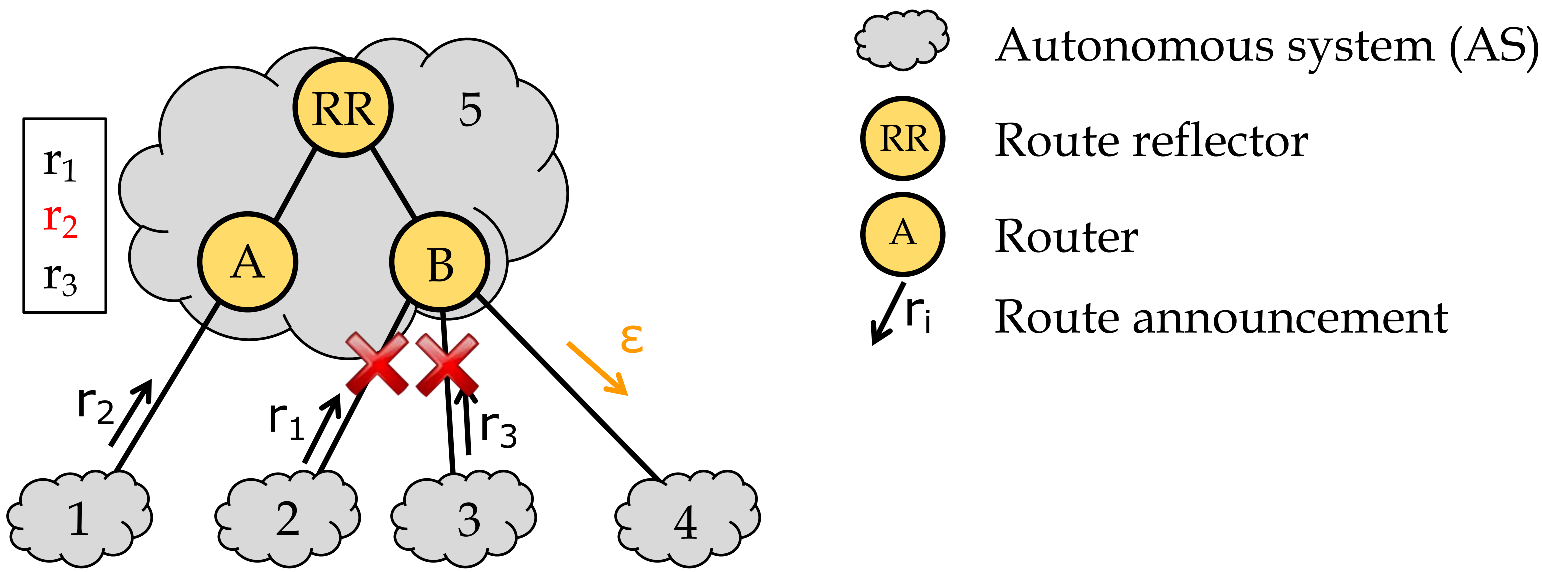

- After withdrawals of routes r_1 and r_3, router B temporarily withdraws the route

Cause 2 – Delays

- Route flap damping temporarily suppresses all routes learned from a neighbor

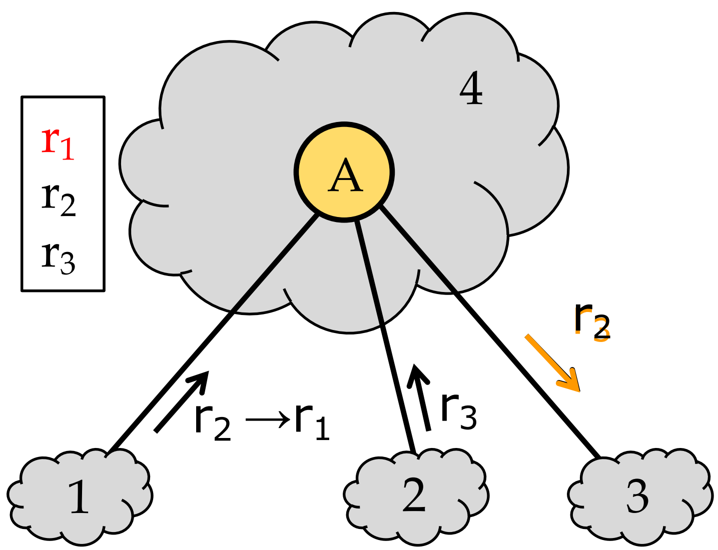

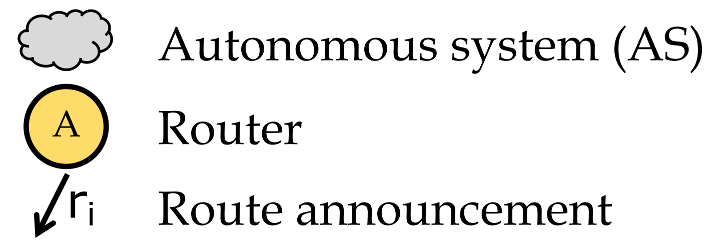

- After the update $r_2 \rightarrow r_1$ the less preferred route r_3 is temporarily selected

DPVP– A More General Model of BGP

- DPVP = Dynamic Path Vector Protocol
  - Generalizes the earlier model (SPVP)
  - Spurious update with a less preferred route that was recently available
- Spurious updates allowed in transient period τ after last route change
  - Safety is independent of numerical value τ

DPVP– A More General Model of BGP

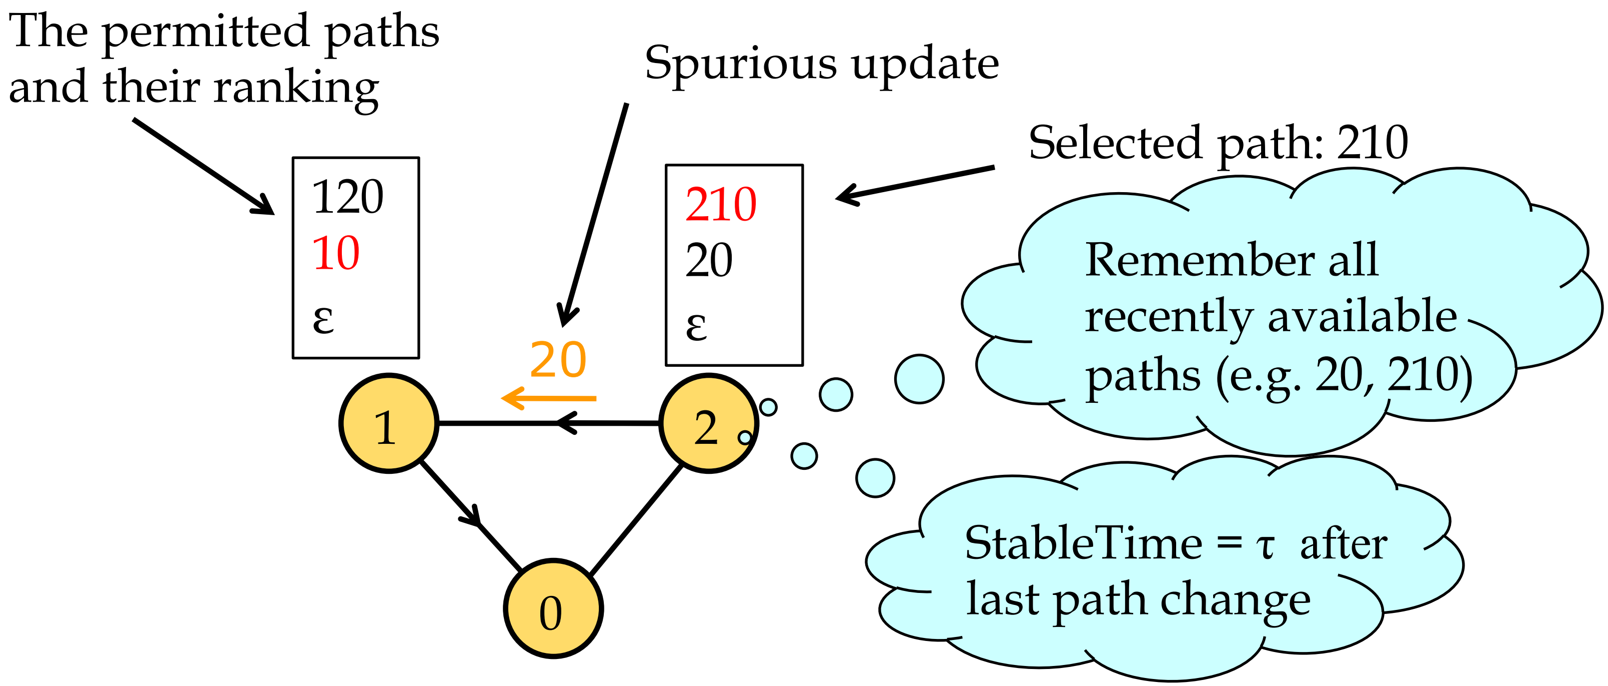

- Spurious updates are allowed only if current time < StableTime
- Spurious updates may include paths that were recently available or the empty path

DPVP– A More General Model of BGP

- Behavior captured irrespective of cause
  - Simple future-proof model independent of underlying network technologies
- For every allowed spurious behavior in DPVP we can find a possible cause
  - Details in our technical report:

    TR-881-10, Dept. of Comp. Sci., Princeton, July 2010

    www.cs.princeton.edu/~msuchara/publications.html

Overview

I. Classical model of BGP: *the SPVP*

II. Spurious BGP updates: *what are they?*

III. The surprise: *networks believed to be safe oscillate!*

IV. The consequences: *applicability of earlier results*

V. Convergence conditions: *polynomial time verifiable*

VI. Conclusion

Consequences of Spurious Updates

- Spurious behavior is temporary
- Tempting to conclude that it cannot have long term consequences
- The surprise: spurious behavior may trigger permanent oscillations!

The Surprise: Spurious Announcements Trigger Permanent Oscillations!

- Safe instance in all classical models of routing:

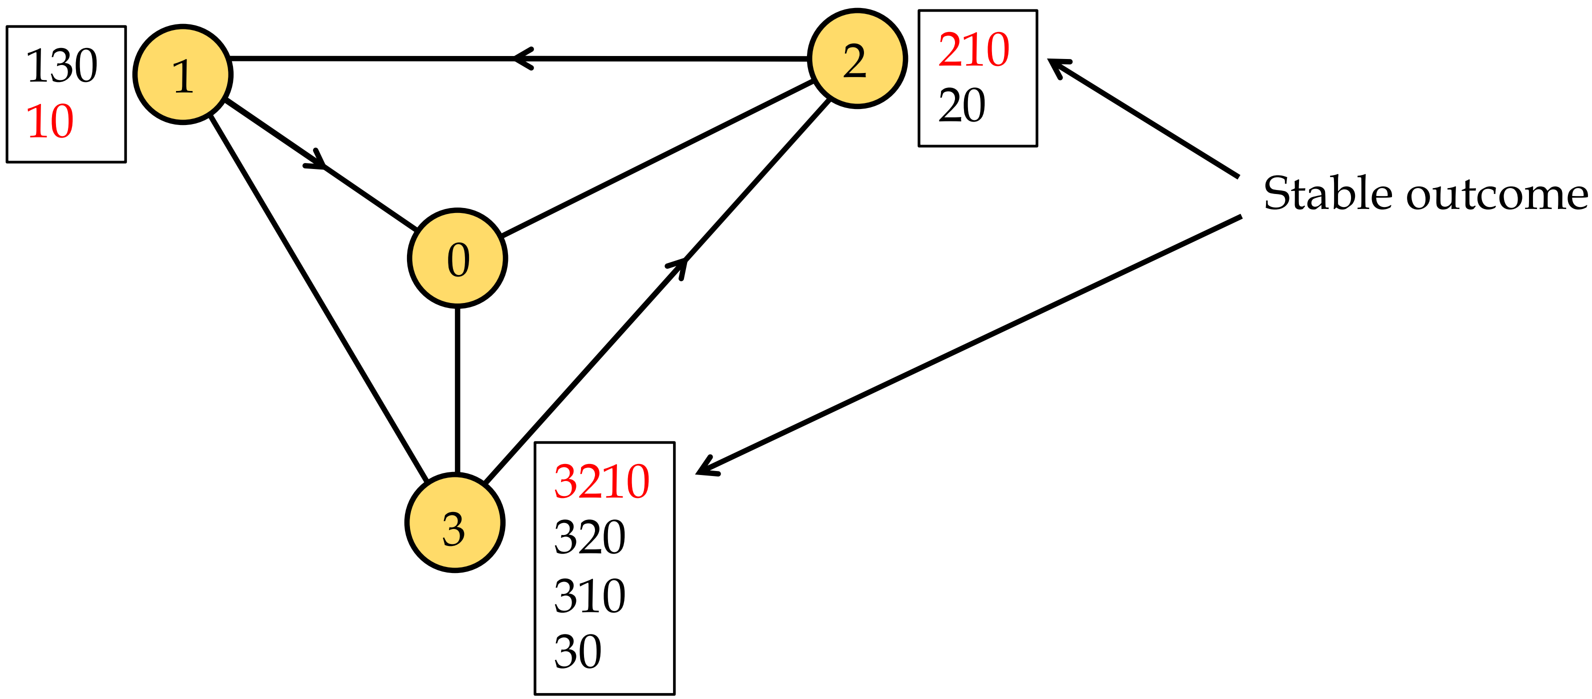

- Permanent oscillations with spurious behavior

Example of Oscillation

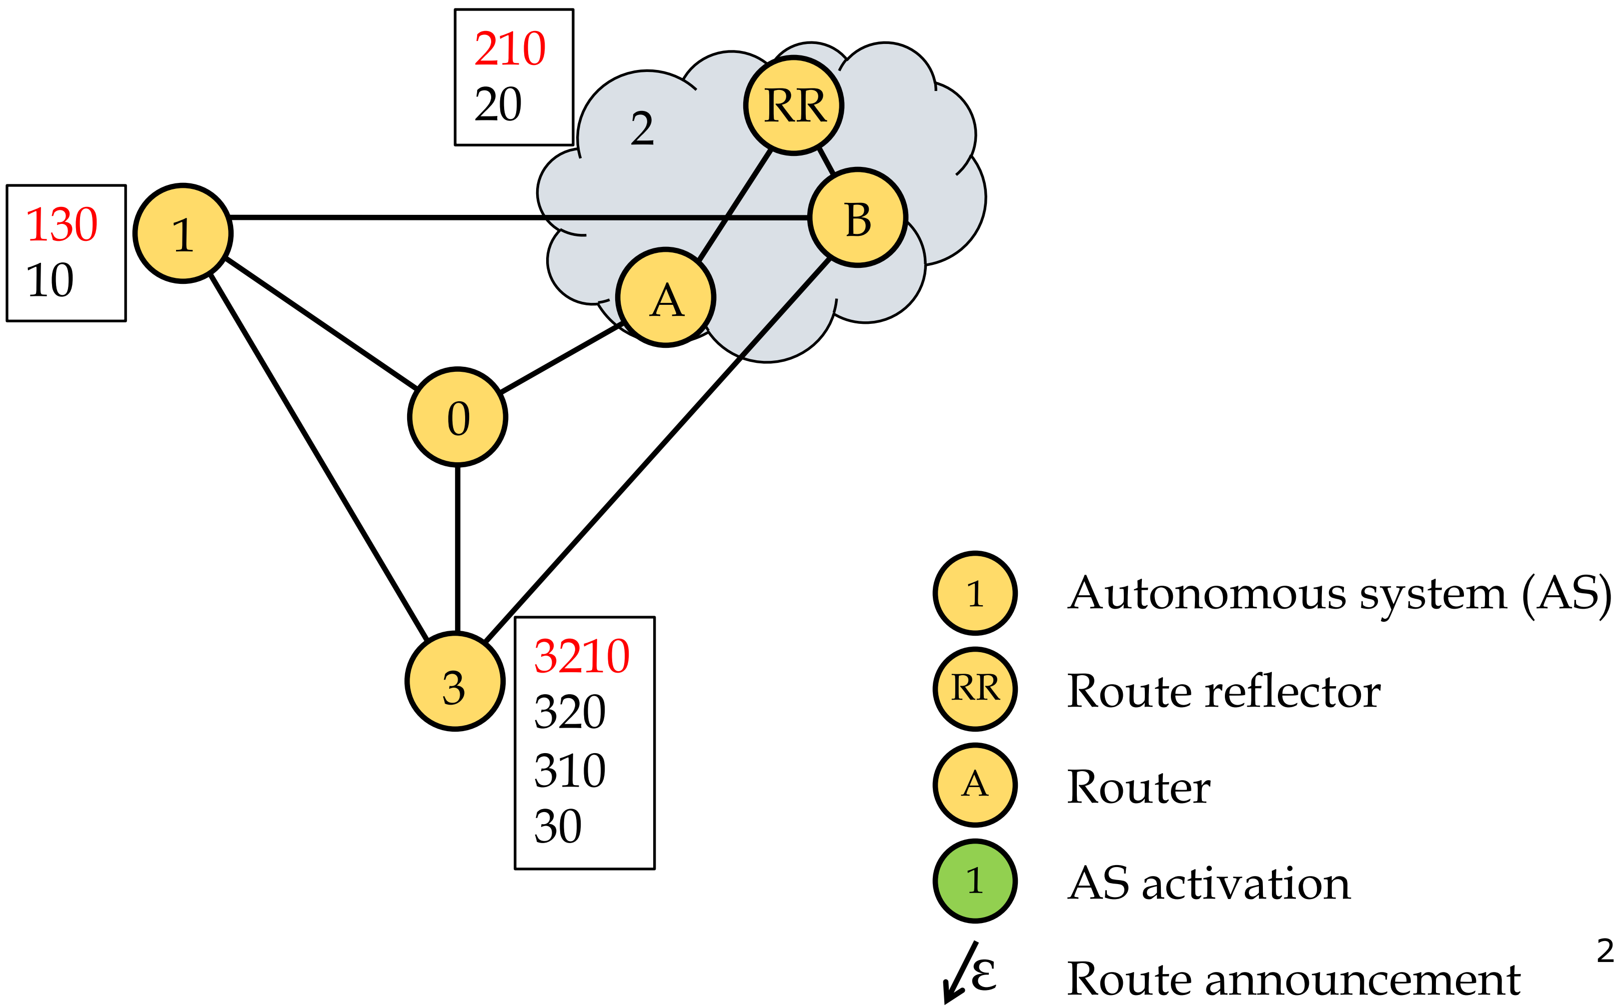

Example of Oscillation

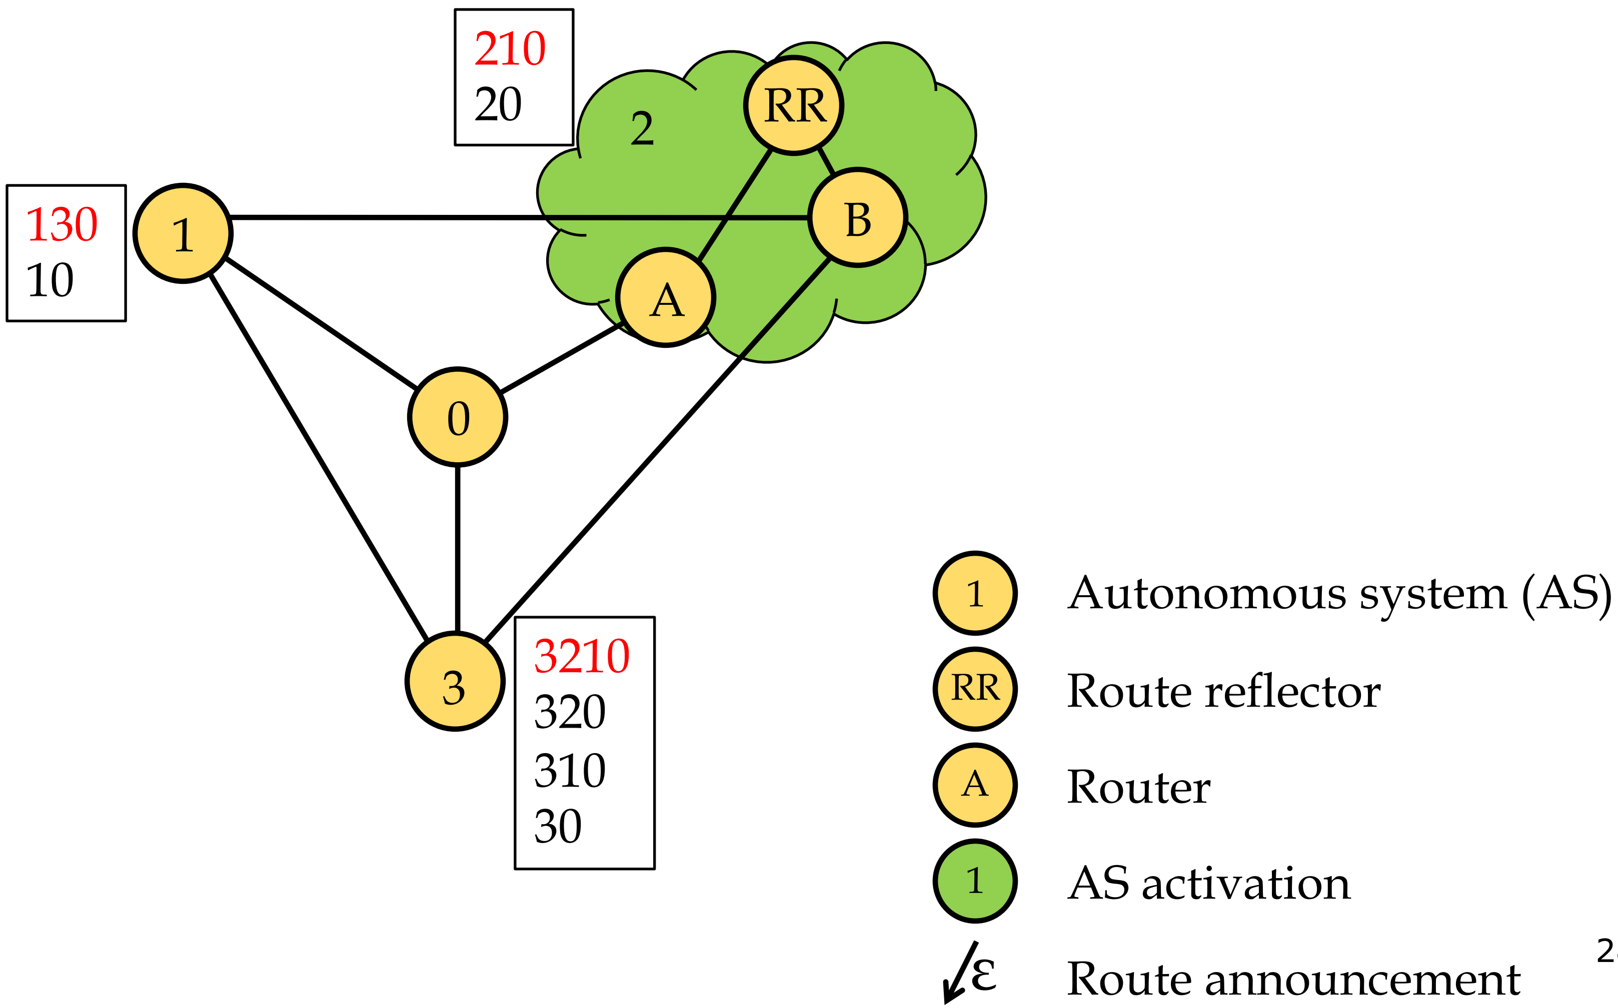

Example of Oscillation

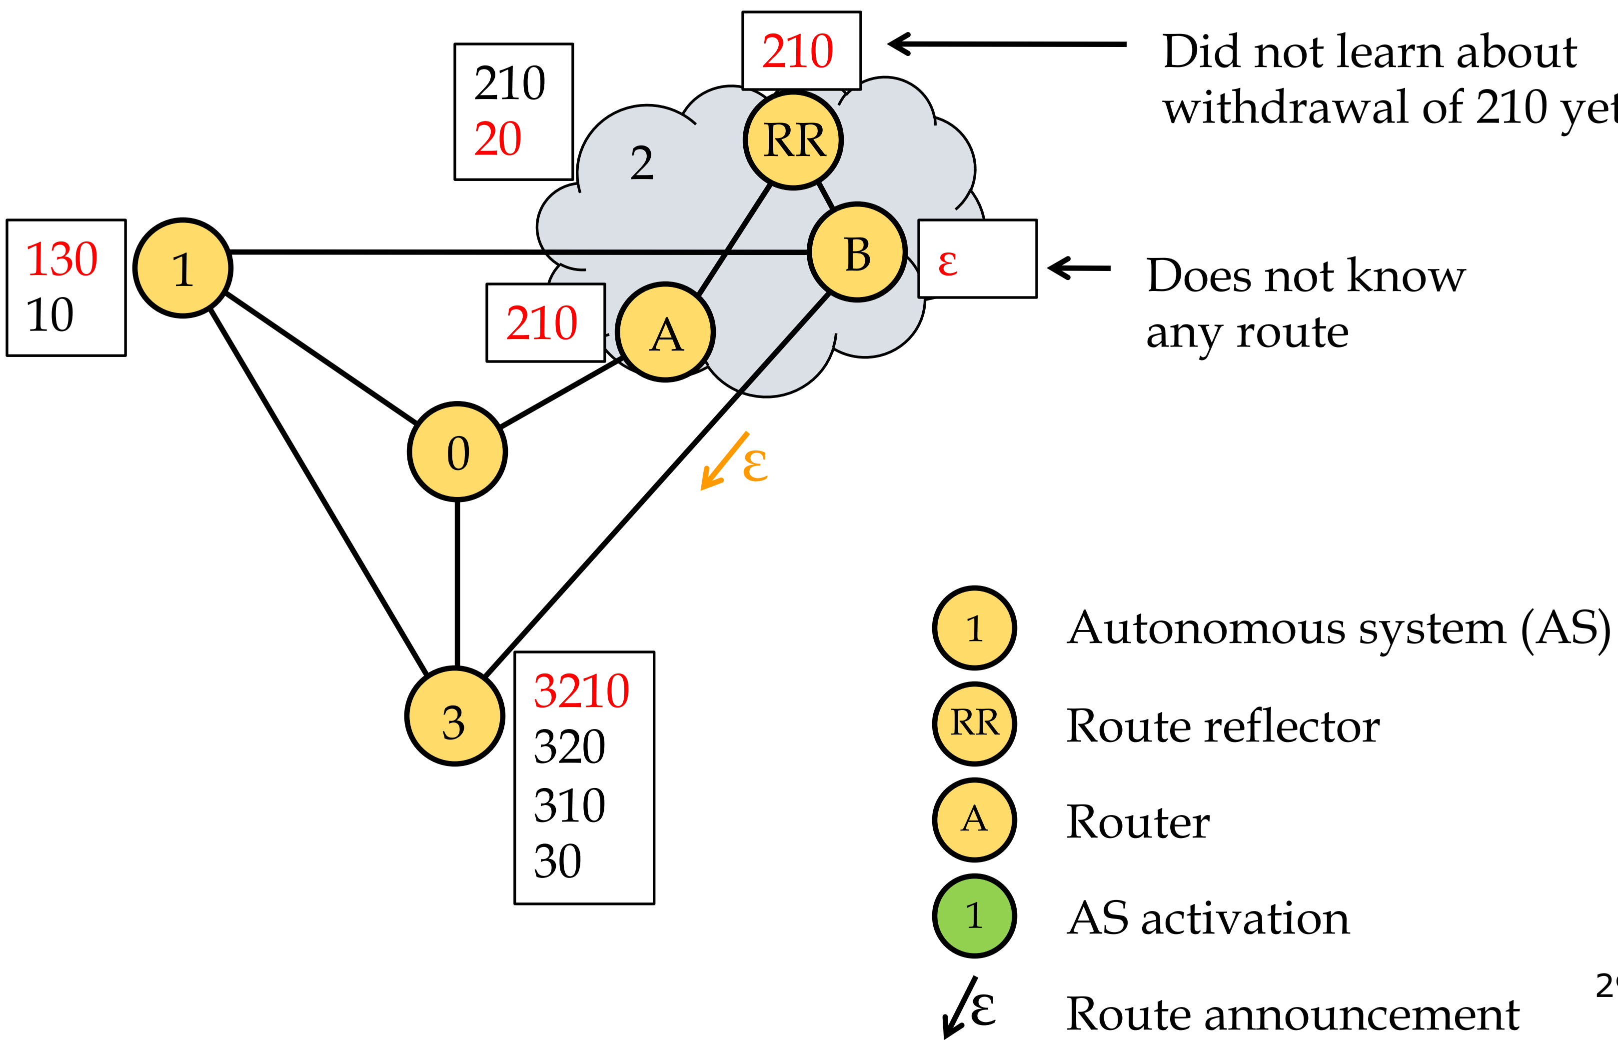

Example of Oscillation

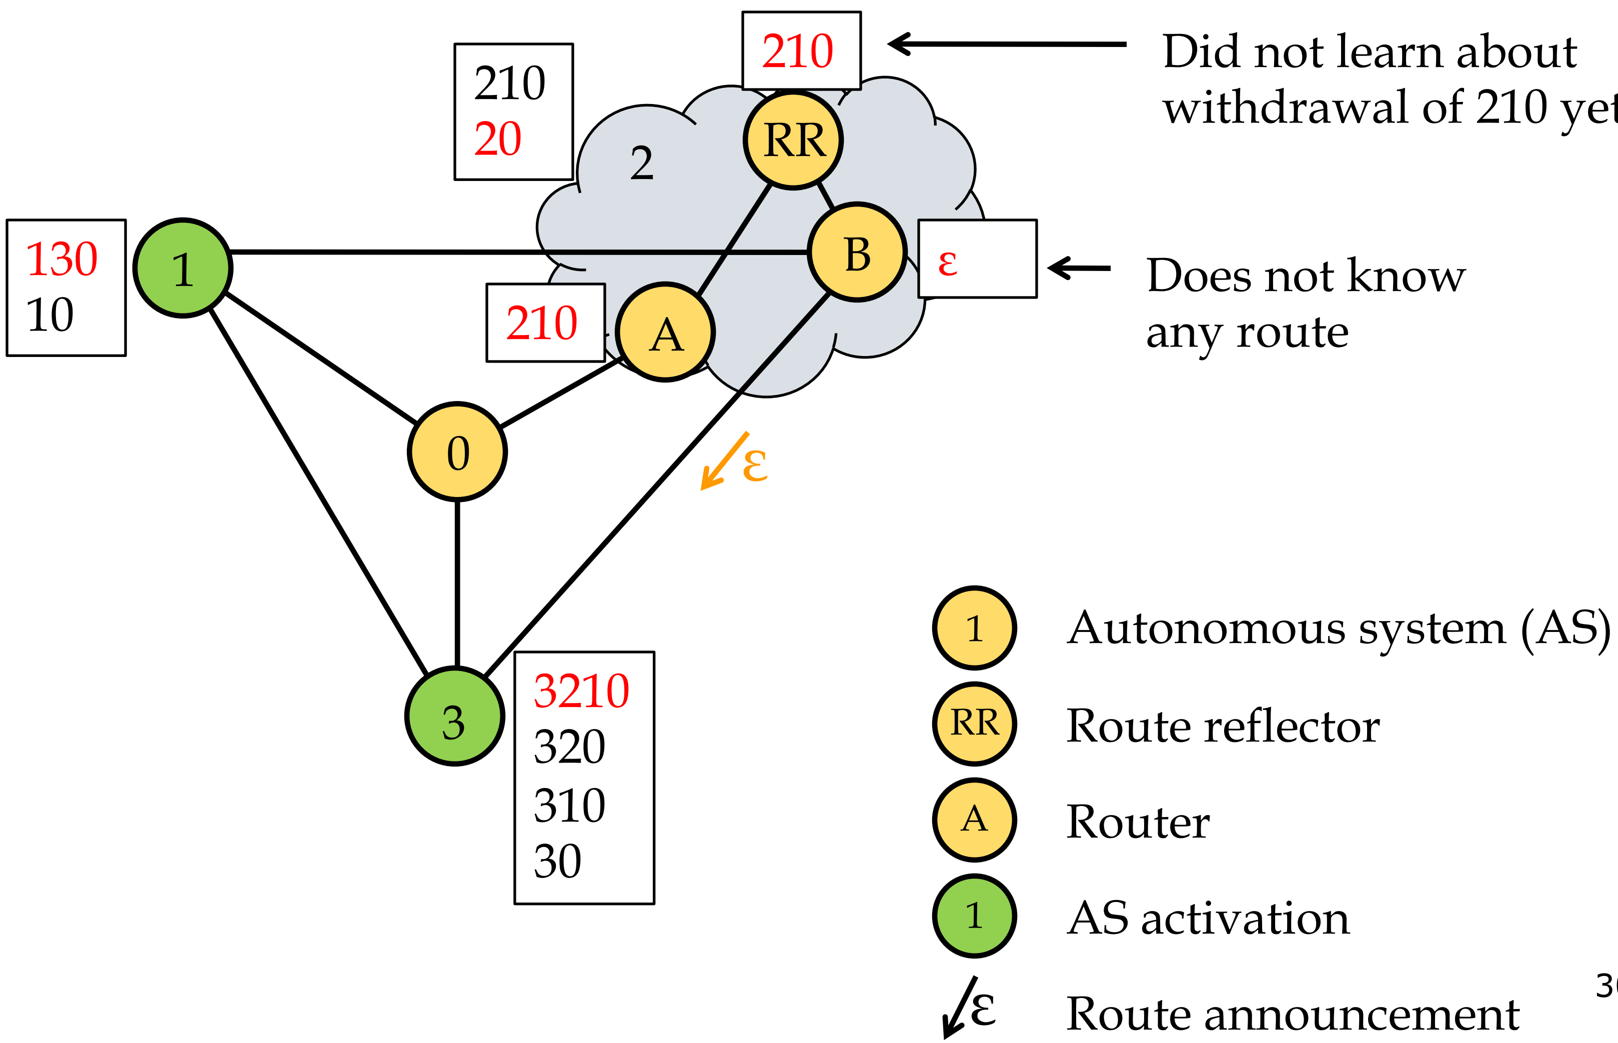

Example of Oscillation

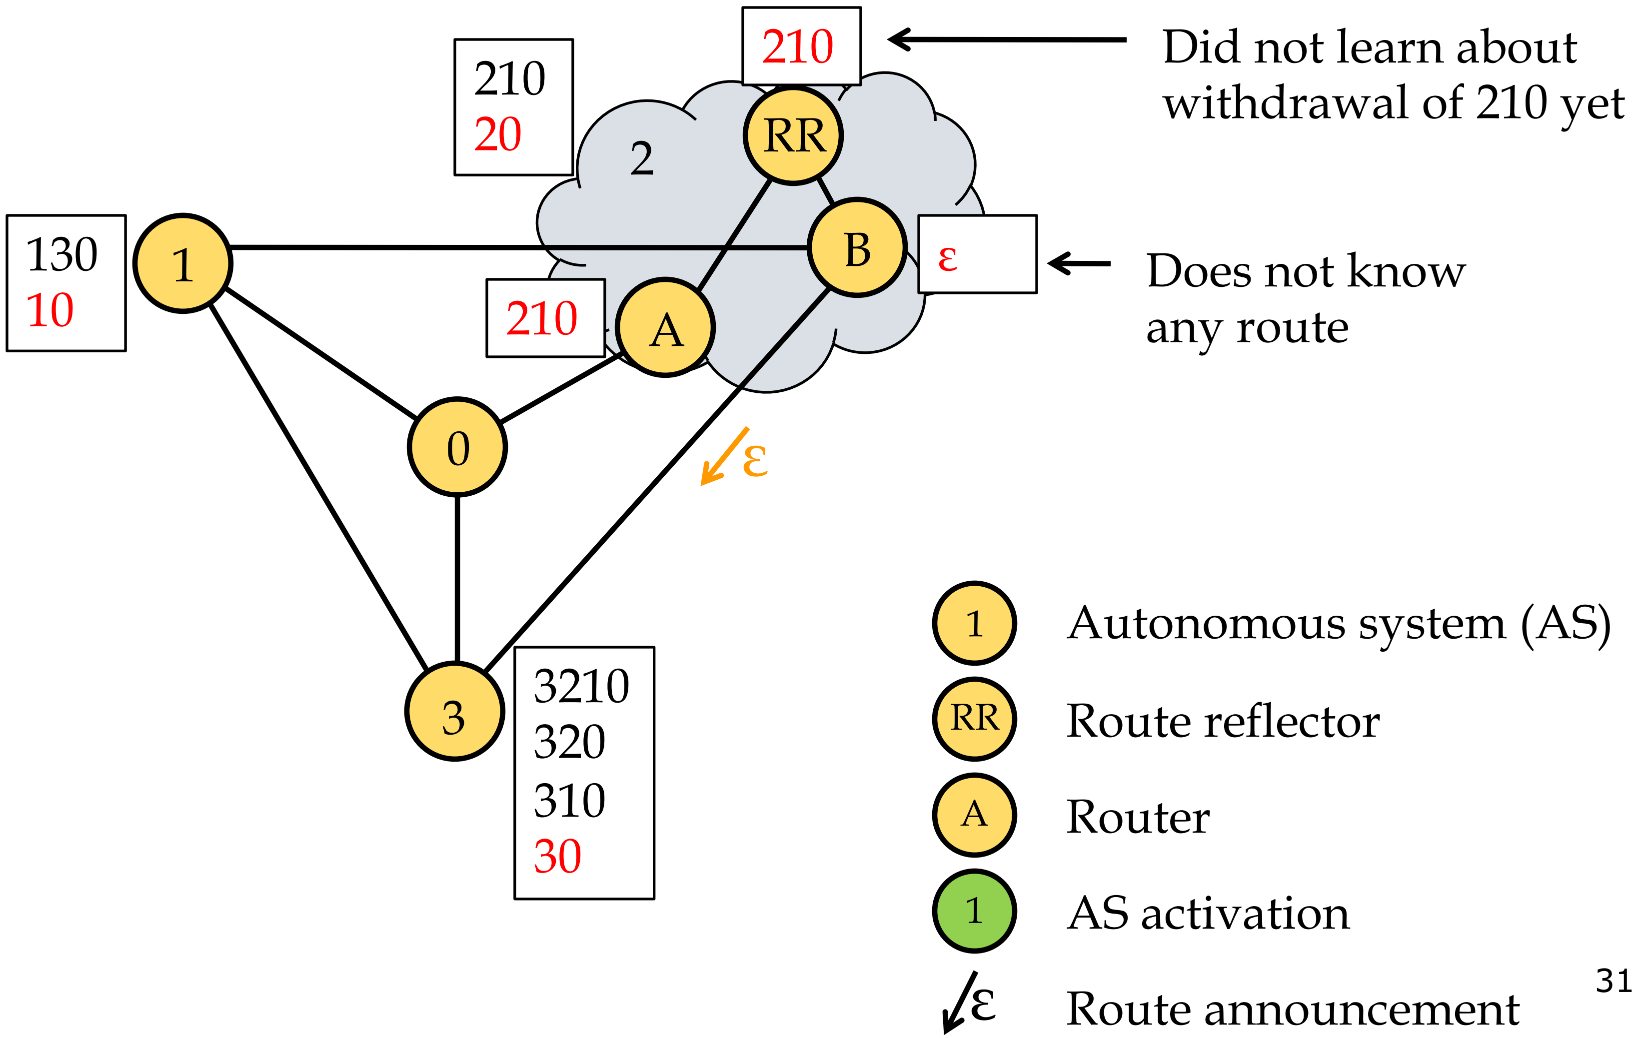

Example of Oscillation

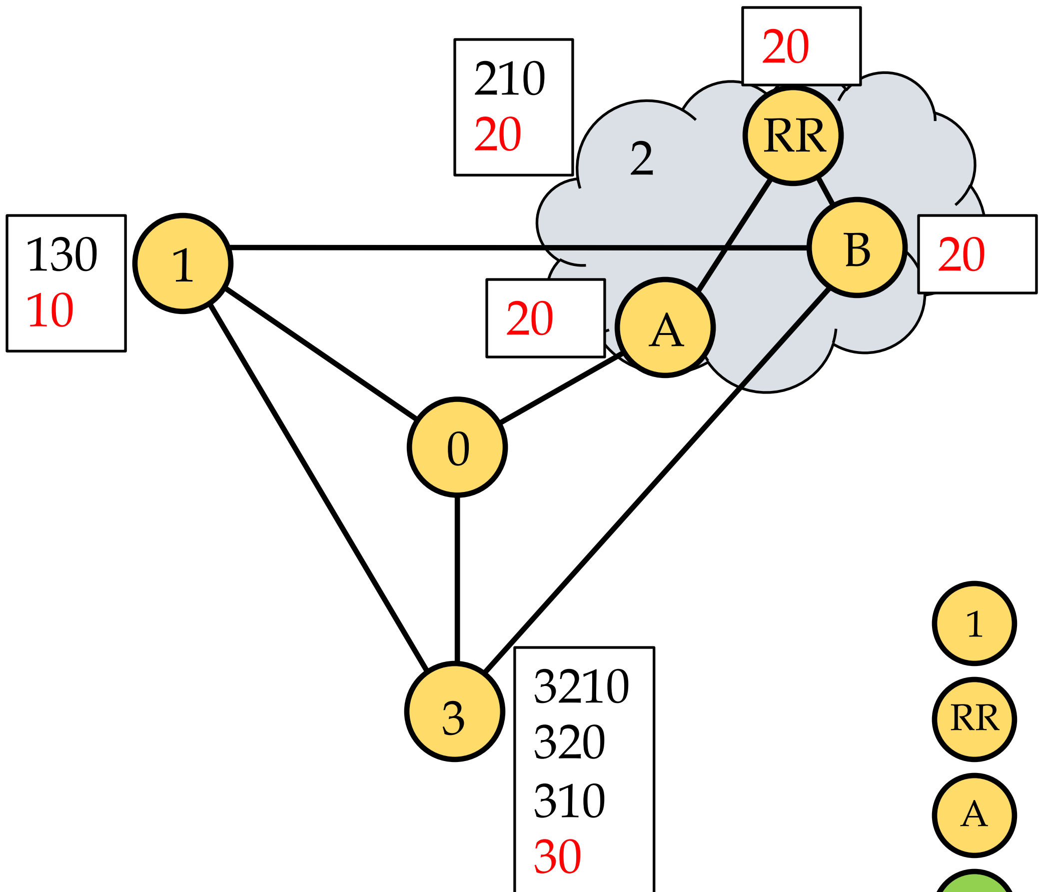

Autonomous system (AS)

Route reflector

Router

AS activation

Route announcement

Example of Oscillation

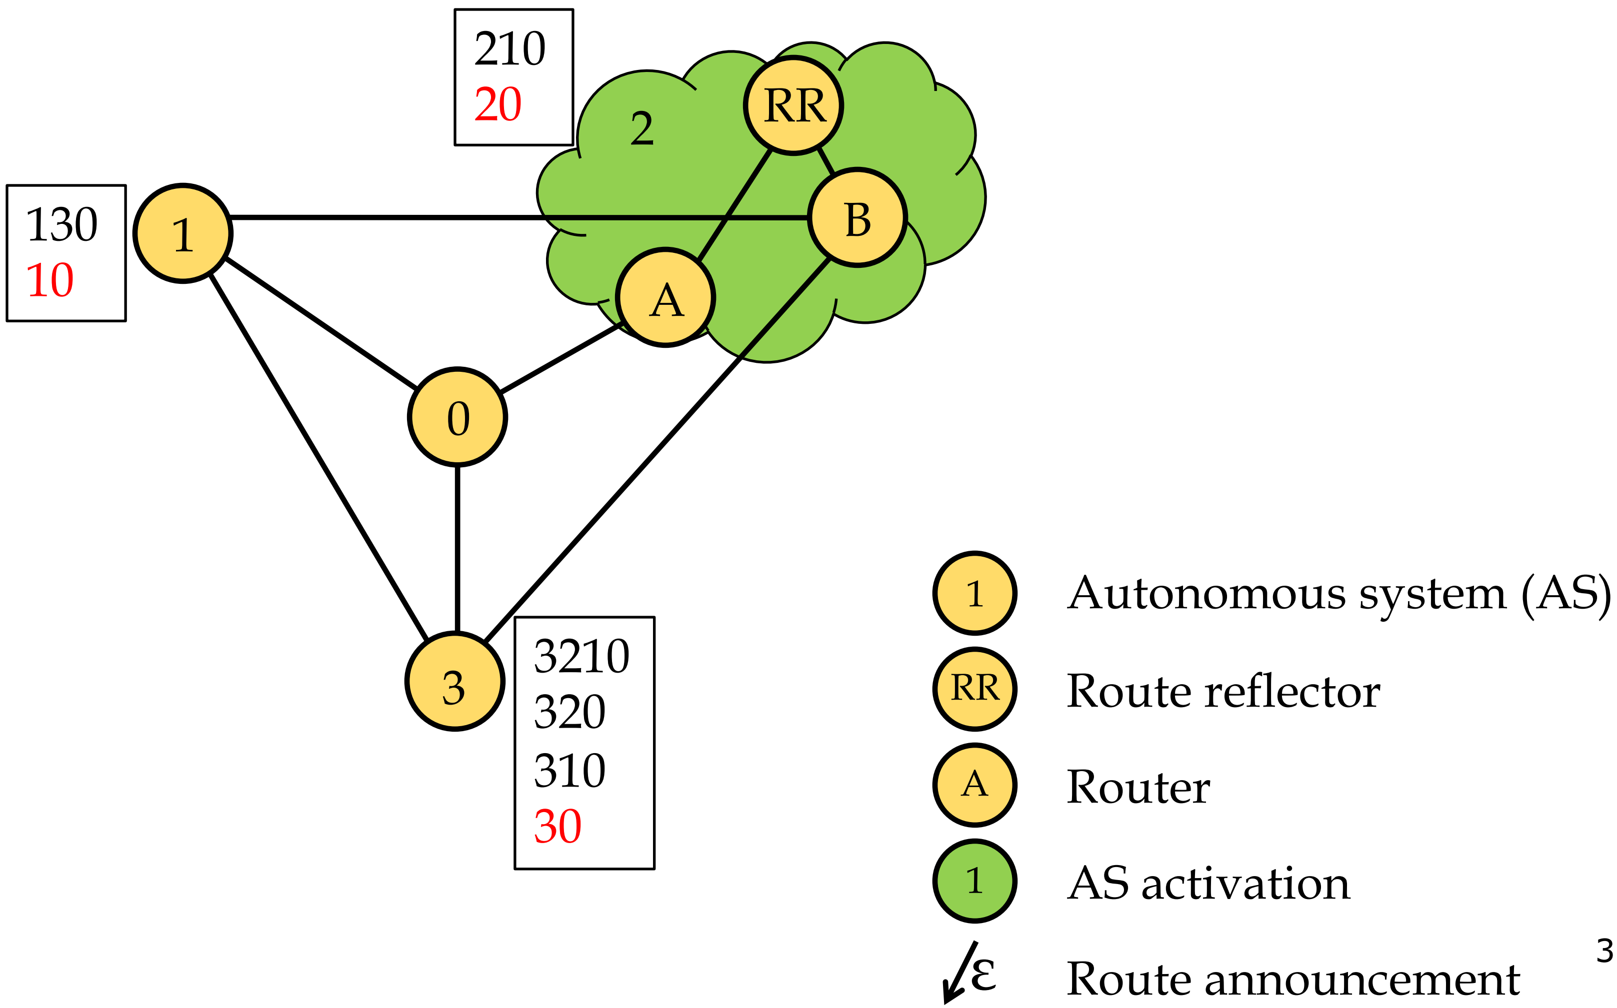

Example of Oscillation

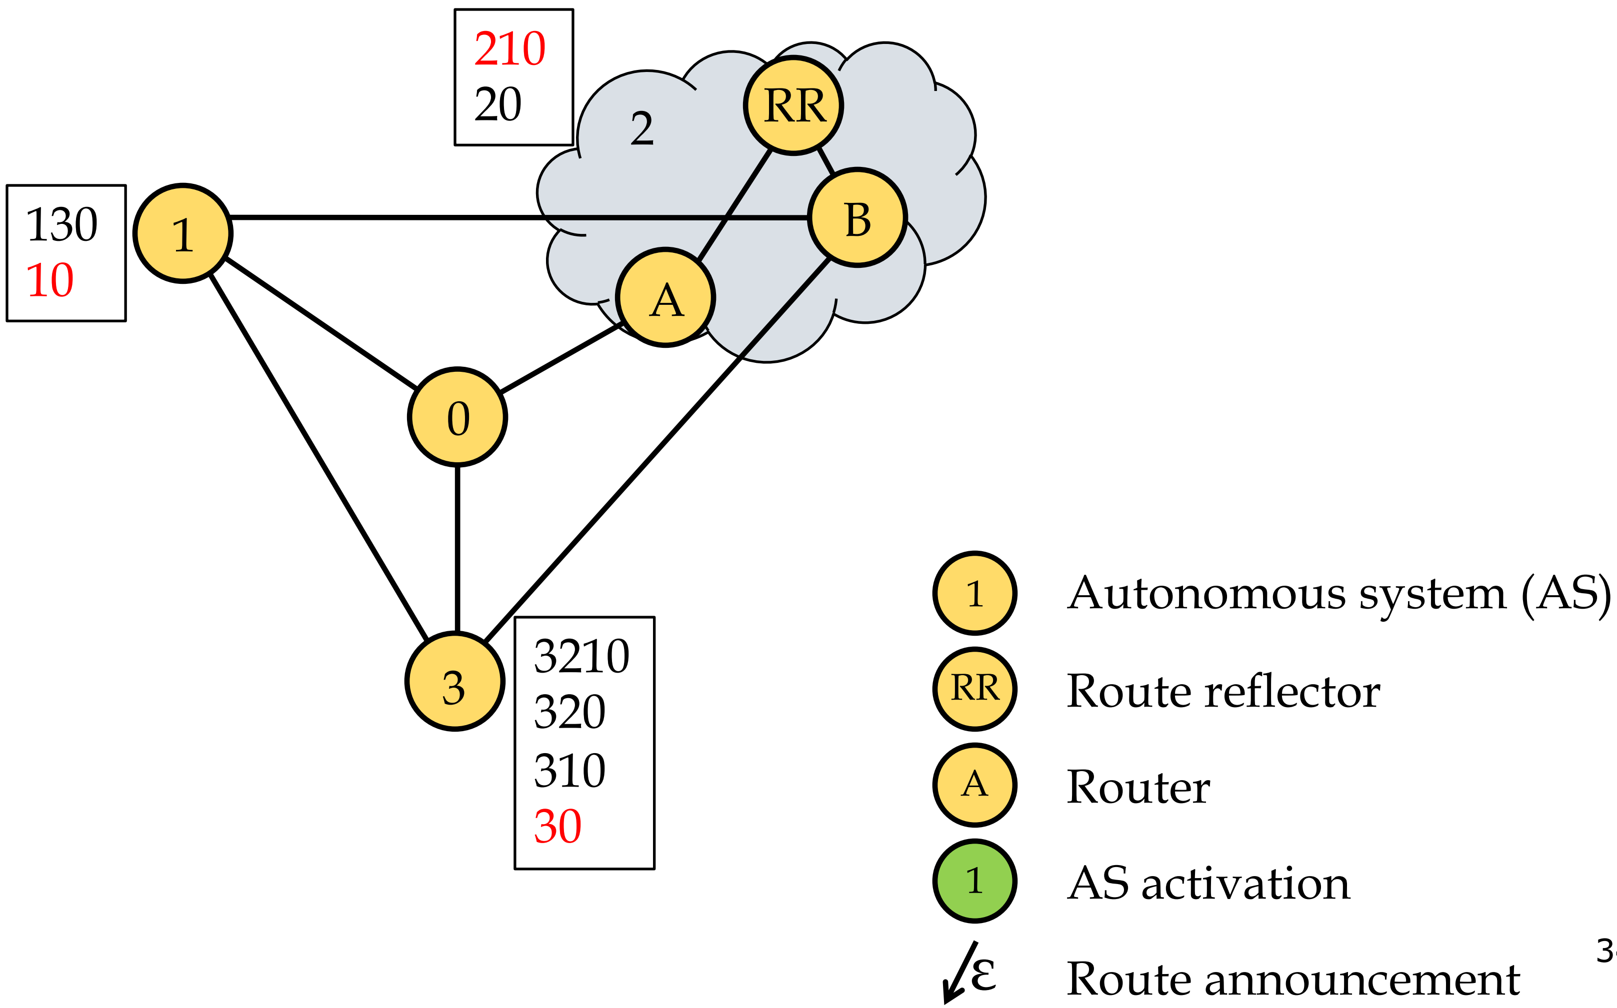

Example of Oscillation

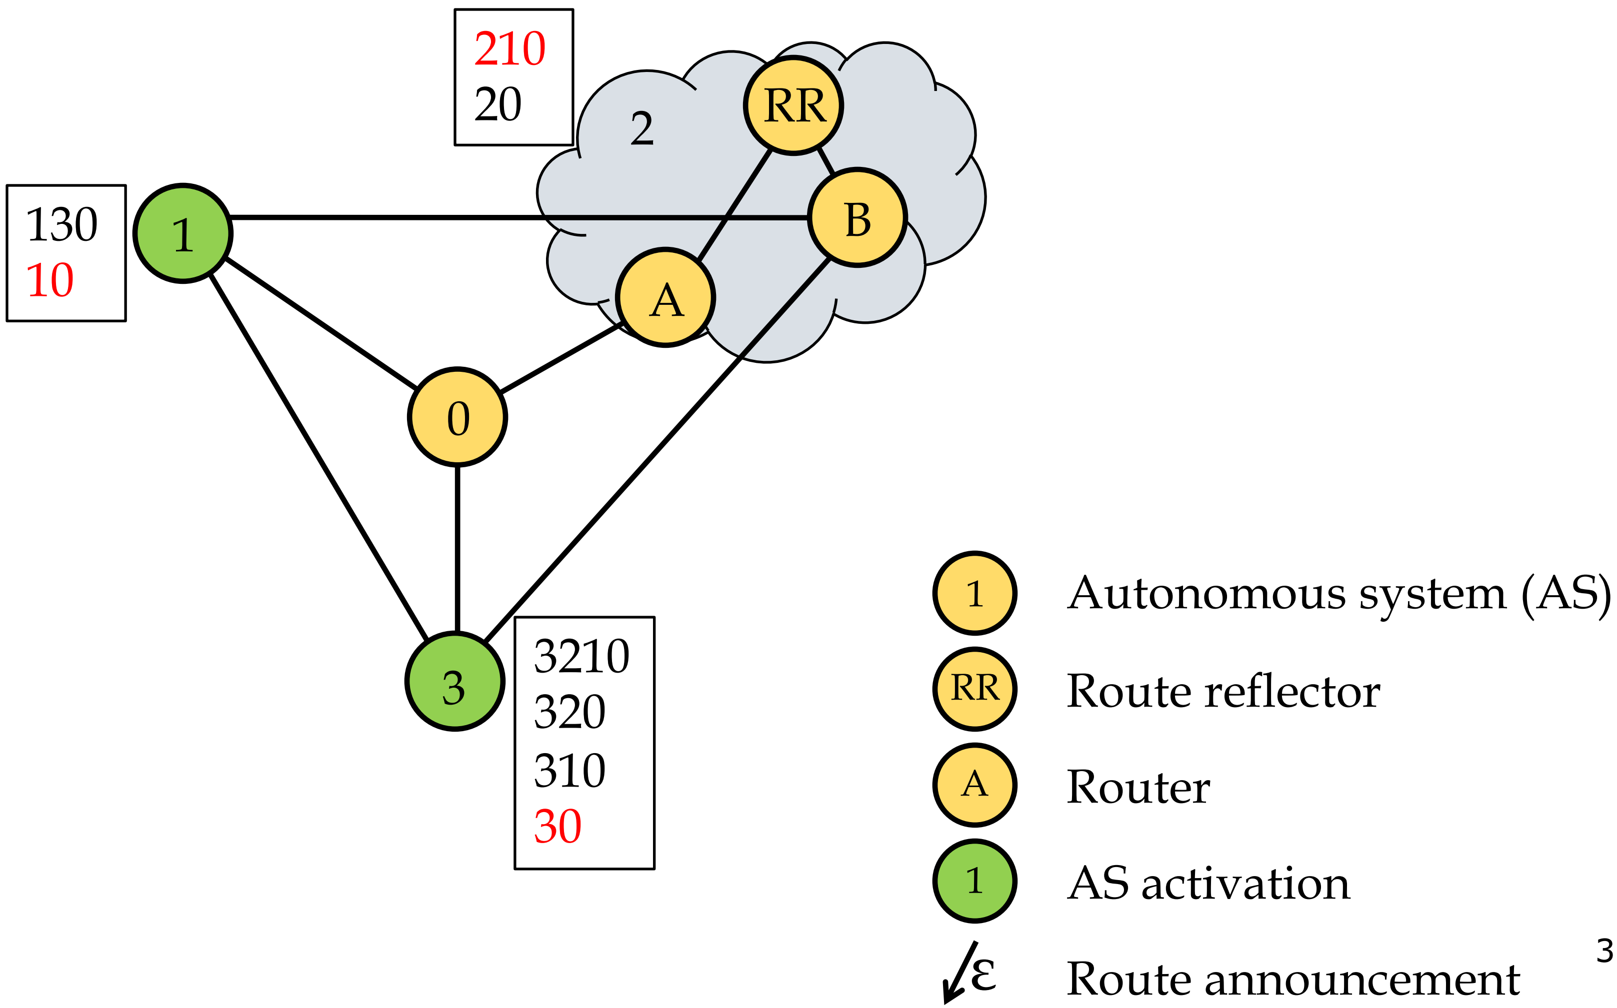

Example of Oscillation

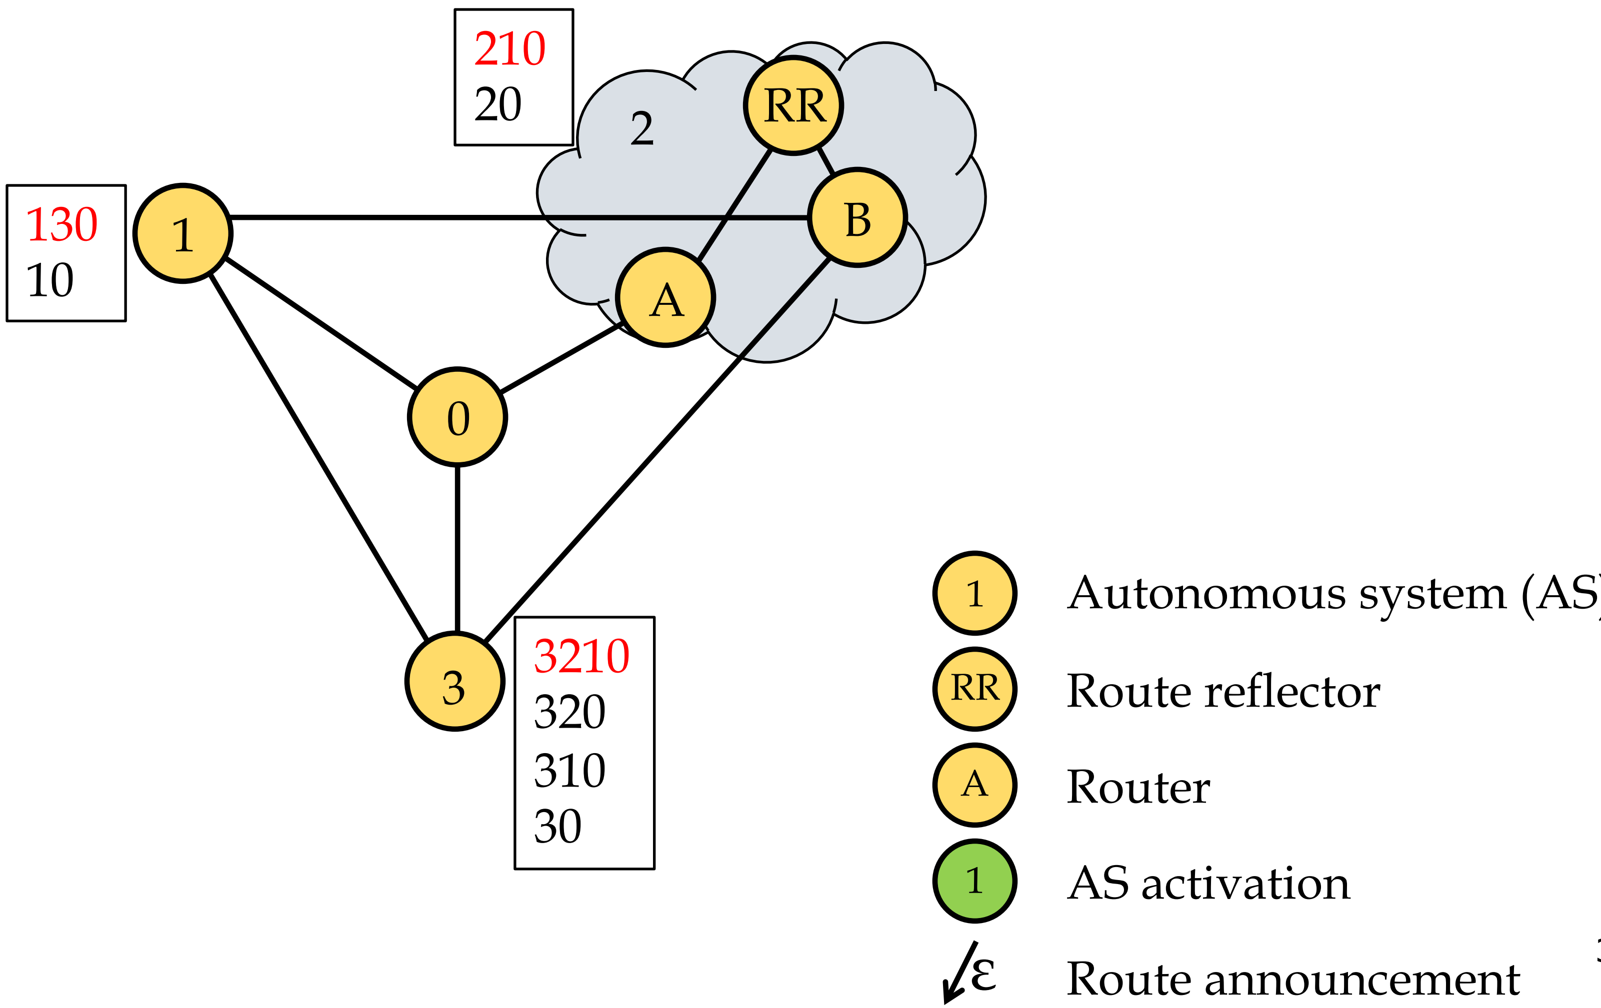

Consequences of Spurious Updates

- Temporary behavior may cause permanent oscillations
- The number of oscillating nodes and / or frequency of oscillations may increase
- Which results do not hold in the new model?

Overview

I. Classical model of BGP: *the SPVP*

II. Spurious BGP updates: *what are they?*

III. The surprise: *networks believed to be safe oscillate!*

IV. The consequences: *applicability of earlier results*

V. Convergence conditions: *polynomial time verifiable*

VI. Conclusion

Which SPVP Results Hold in DPVP?

- Most previous results in SPVP also hold for DPVP
  - Formal justification later in the talk

- Some results cannot be extended
  - Slightly different conditions of convergence
  - Exponentially slower convergence possible

Case Study I
Different Conditions of Convergence

- Safety under filtering: is instance safe under any filtering?

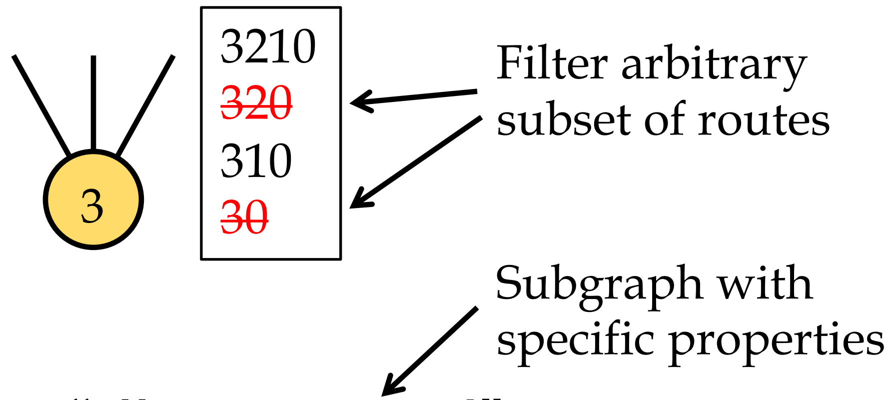

- Absence of a "dispute reel" necessary and sufficient for safety under filtering in SPVP *(Cittadini et al., 2009)*
- Our result: permanent oscillations in DPVP even without a reel

Case Study I
Different Conditions of Convergence

- Example of a "safe" topology, Cittadini et al.:

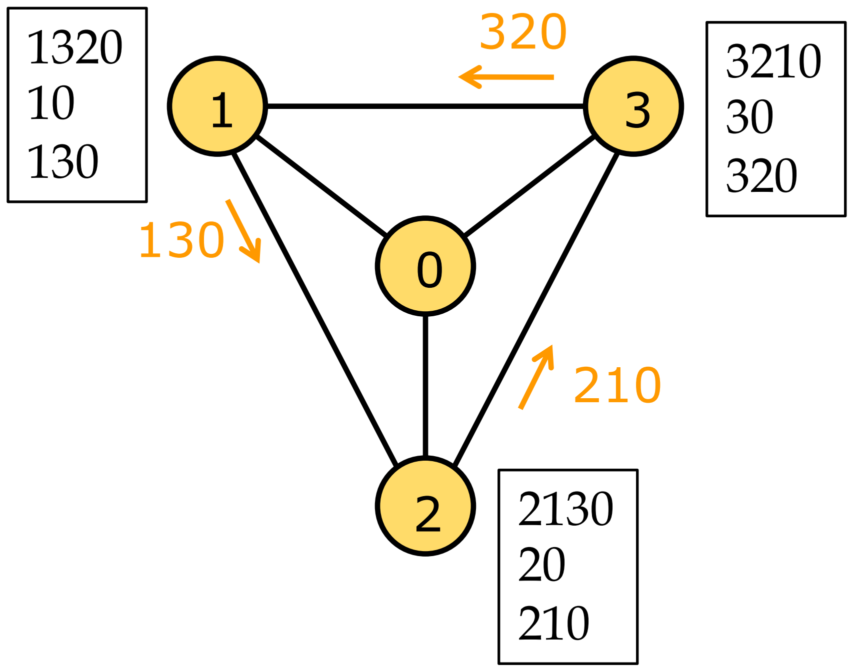

- Spurious updates cause oscillations

Case Study II
Exponential Slowdown of Convergence

- BGP converges in $2l + 2$ phases
  (Sami et al., 2009)
  - l: length of longest customer-provider chain
  - Phase: each node processes and sends updates
  - Assumes standard business relationships
- With spurious updates exponential slowdown to $(2k + 1)^{l-2}$ phases
  - k: max. # of spurious updates after route change

Overview

I. Classical model of BGP: *the SPVP*

II. Spurious BGP updates: *what are they?*

III. The surprise: *networks believed to be safe oscillate!*

IV. The consequences: *applicability of earlier results*

V. Convergence conditions: *polynomial time verifiable*

VI. Conclusion

Convergence Conditions

- Absence of a "dispute wheel" sufficient for safety in SPVP *(Griffin, Shepherd, Wilfong, 2002)*
  - One of the most cited results
- Absence of a "dispute wheel" is still sufficient for safety in DPVP
  - Most of the previous results of the past decade still hold under DPVP!
- Other stronger results in DPVP next

Why are Proofs Easier in DPVP?

- No need to prove that:
  - Announced route is the highest ranked one
  - Announced route is the last one learned from the downstream neighbor
- We changed the problem
  - PSPACE complete vs. NP complete
- Next the necessary and sufficient conditions

Necessary and Sufficient Conditions

- How can we prove a system may oscillate?
  - Classify each node as "stable" or "coy"
  - At least one "coy" node exists
  - Prove that "stable" nodes must be stable
  - Prove that "coy" nodes may oscillate

Easy in a model with spurious announcements

- Next: a formal definition of a construction that captures this intuition

Necessary and Sufficient Conditions

- Definition: CoyOTE is a triple (C, S, Π) satisfying several conditions

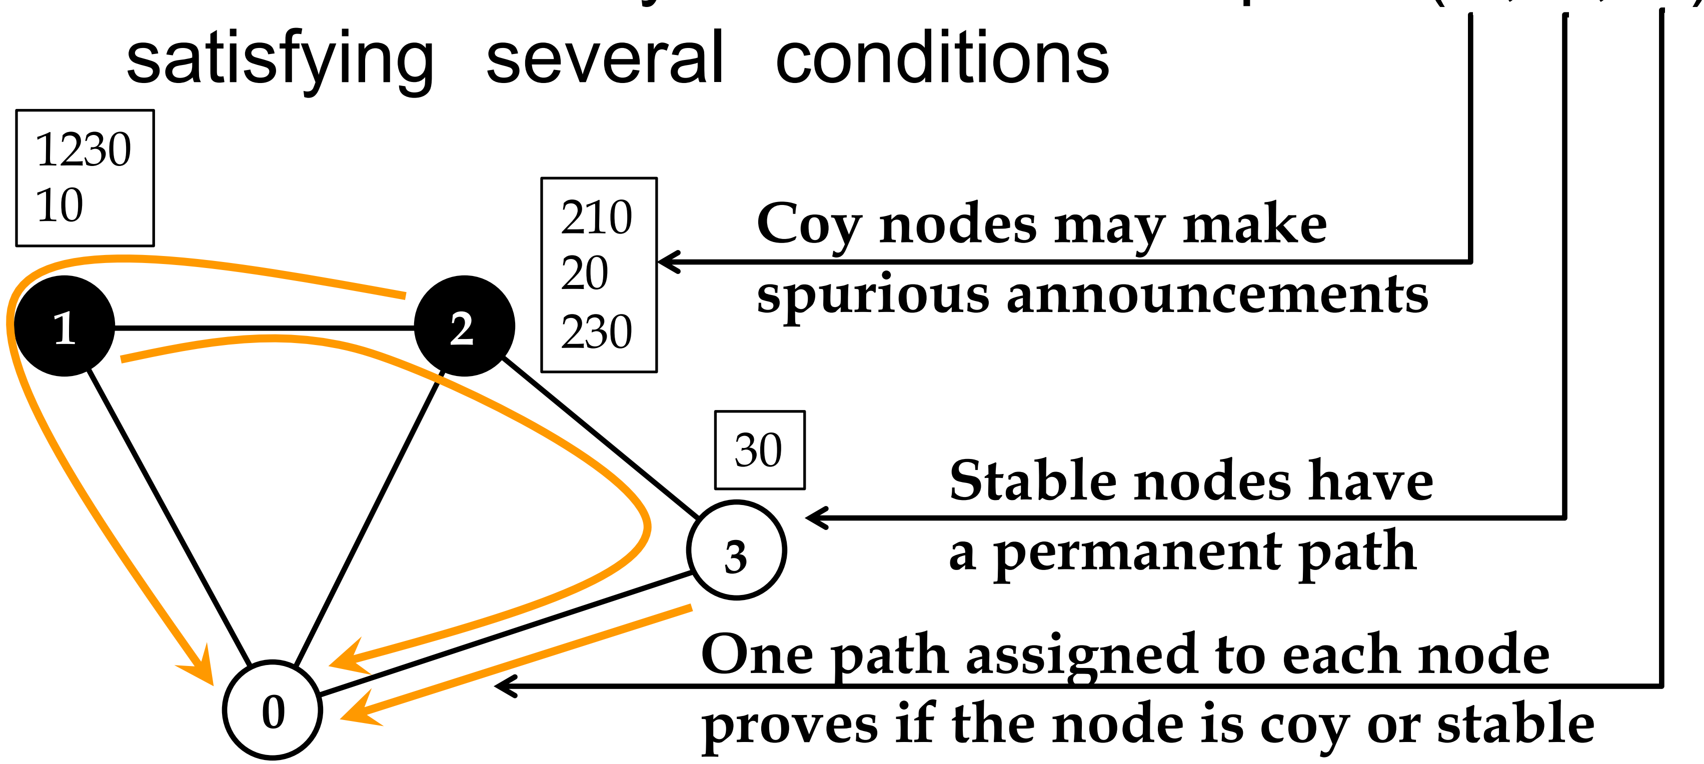

- Theorem: DPVP oscillates if and only if it has a CoyOTE

Necessary and Sufficient Conditions

Definition: CoyOTE satisfies these conditions:

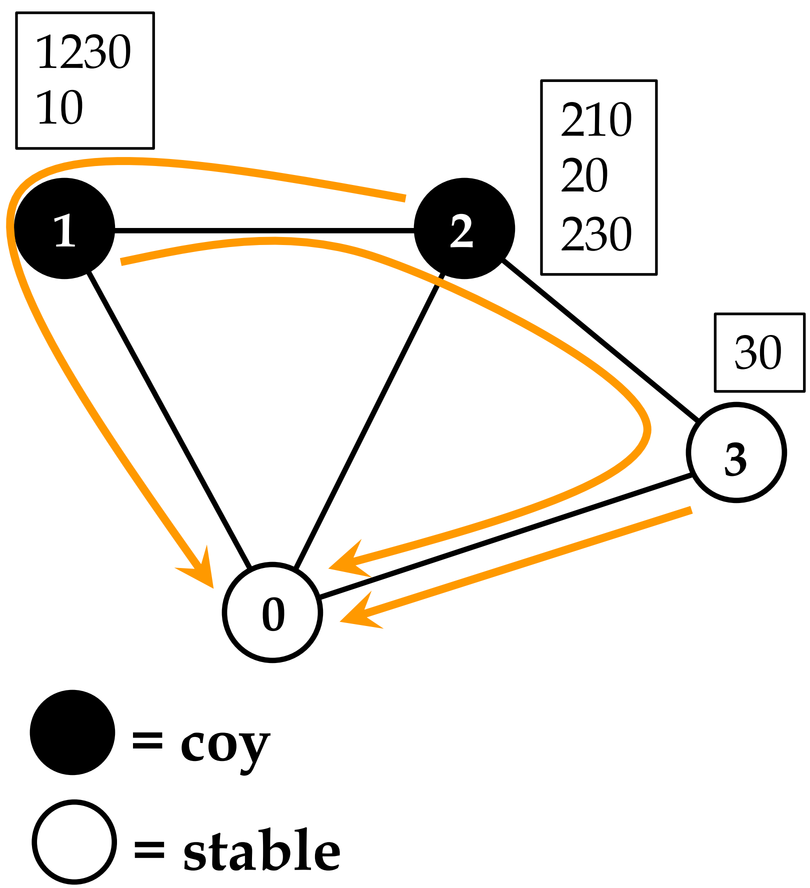

1) The best stable path assigned to each stable node

2) Coy node is assigned a coy path:
   - more preferred than the best stable path
   - consistent with the paths of stable nodes

3) Origin is stable

Verifying the Convergence Conditions = Finding a CoyOTE

- In general an NP-hard problem
  - Compact inputs with regular expressions
- Can be checked in polynomial time for most "reasonable" network configurations!

↙ **e.g.**

(i)	filter paths violating business relationships
(ii)	prefer paths not containing certain AS numbers
(iii)	prefer paths from certain groups of neighbors
(iv)	prefer shorter paths over longer ones
(v)	prefer paths from a lowest AS number neighbor

DeCoy – Safety Verification Algorithm

- Goal: verify safety in polynomial time
- Main idea: find the maximal stable set S by expanding it in a greedy fashion
  - If all nodes are stable system is safe
  - Otherwise system may oscillate

DeCoy – Safety Verification Algorithm

- Goal: verify safety in polynomial time

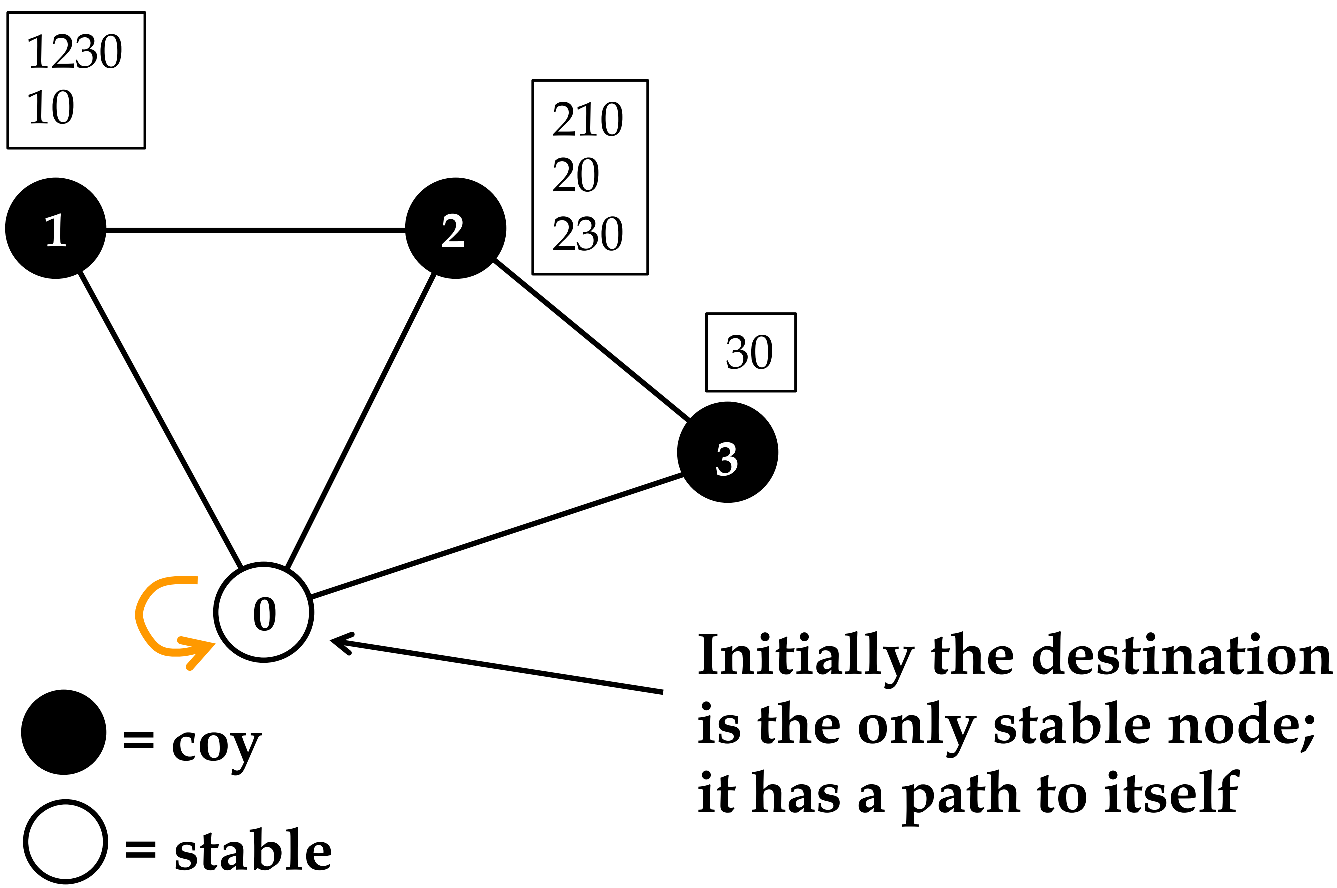

DeCoy – Safety Verification Algorithm

- Goal: verify safety in polynomial time

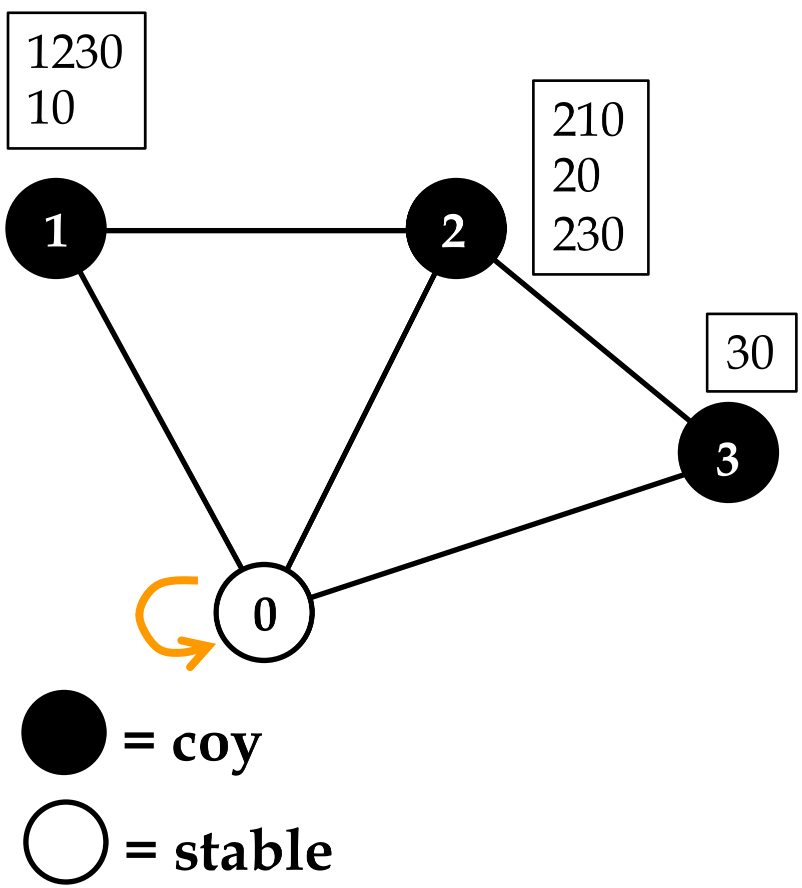

Find a coy node where its most preferred path consistent with the paths of all stable nodes is a stable path

DeCoy – Safety Verification Algorithm

- Goal: verify safety in polynomial time

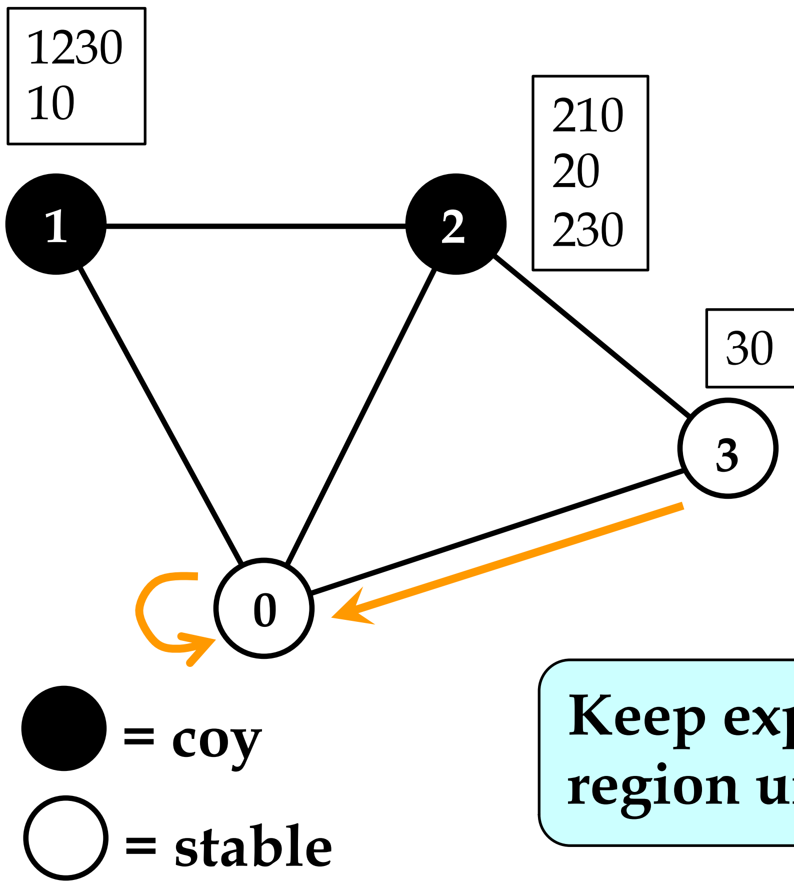

Find a coy node where its most preferred path consistent with the paths of all stable nodes is a stable path

Assign this path to the node and mark the node as stable

Keep expanding the stable region until we get stuck

DeCoy – Safety Verification Algorithm

- Goal: verify safety in polynomial time
- Theorem: if all nodes are added to the stable set the system is safe. Otherwise it is not safe.
- Open question: find a distributed version of the algorithm that preserves privacy of the nodes
  - Business relationships are secret

Overview

I. Classical model of BGP: *the SPVP*

II. Spurious BGP updates: *what are they?*

III. The surprise: *networks believed to be safe oscillate!*

IV. The consequences: *applicability of earlier results*

V. Convergence conditions: *polynomial time verifiable*

VI. Conclusion

Conclusion

- DPVP: best of both worlds
  - More accurate model of BGP
  - Model simplifies theoretical analysis
- Key results

(i)	Spurious announcements are real
(ii)	Safe instances in SPVP may oscillate in DPVP
(iii)	No dispute wheel $\rightarrow$ safety
(iv)	Necessary and sufficient conditions of convergence, can be found in polynomial time

Thank You!

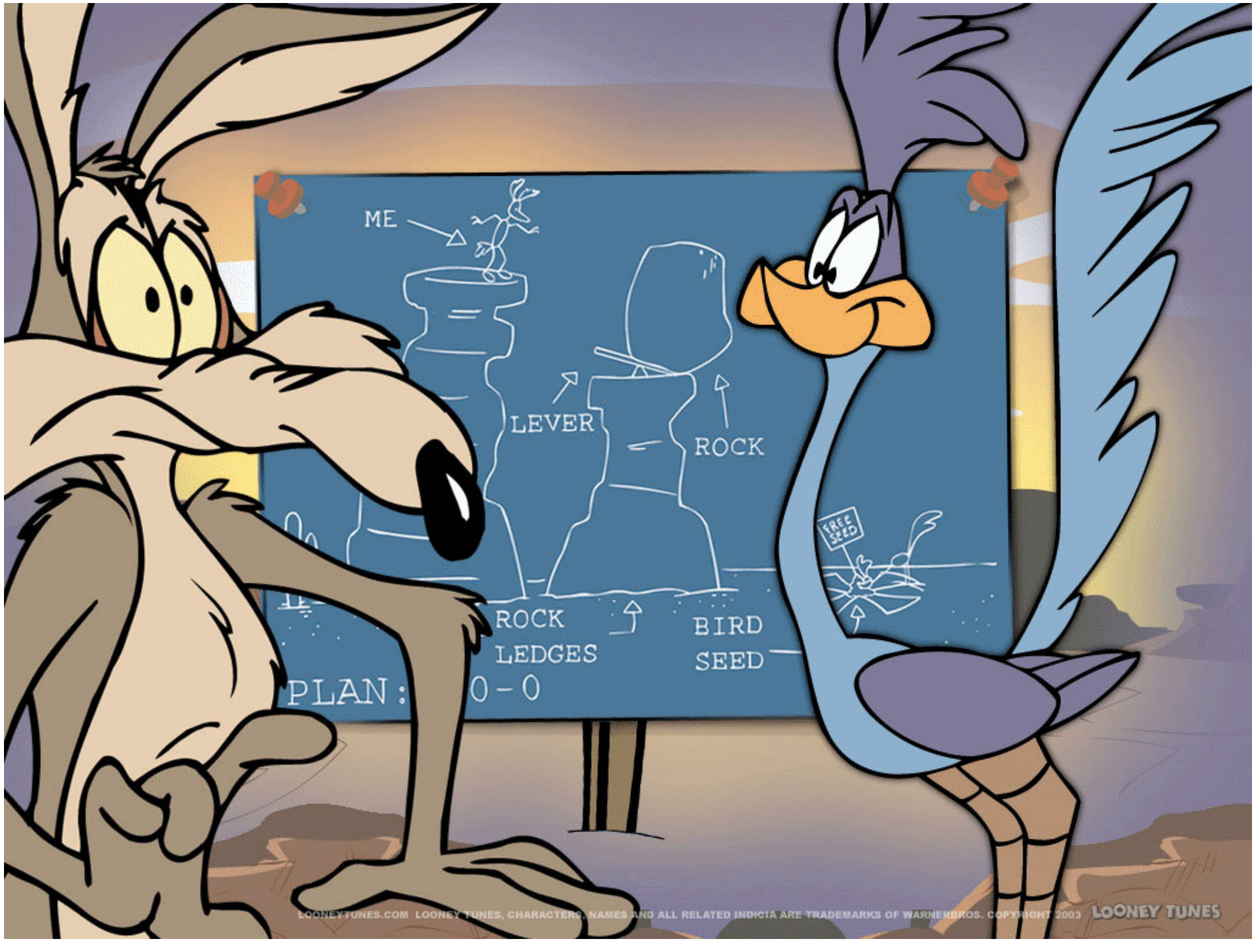

ME
LEVER
ROCK
FREE SEED
ROCK LEDGES
BIRD SEED
PLAN: 0-0
LOONEY TUNES