Reducing the Complexity of BGP Stability Analysis with Hybrid Combinatorial-Algebraic Models

Debbie Perouli, Stefano Vissicchio, Alexander Gurney,
Olaf Maennel, Timothy G. Griffin, Iain Phillips,
Sonia Fahmy, and Cristel Pelsser

WRiPE 2012

BGP Stability

Safety: convergence to a *unique*, *stable* routing solution

Shortest path theory is not adequate to model policy-based routing.

Instead:

- Graph Theoretic Models based on SPP[1]
  - Major drawback: require path enumeration
- Algebraic Models
  - can do without it at cases (e.g. iBGP)

Both approaches provide *sufficient* conditions for safety.

1: T. G. Griffin, F. B. Shepherd, and G. Wilfong. Policy Disputes in Path-Vector Protocols. ICNP 1999.

Internet Routing

BGP: interdomain routing protocol of the Internet

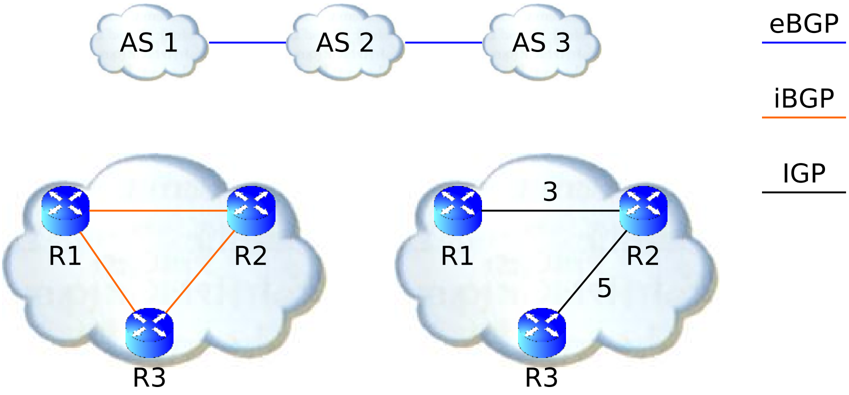

Decision Process Steps:

1. Highest Local Preference 2. Shortest AS path length ...

Policies and Complexities in iBGP

- Route reflection limits route visibility.

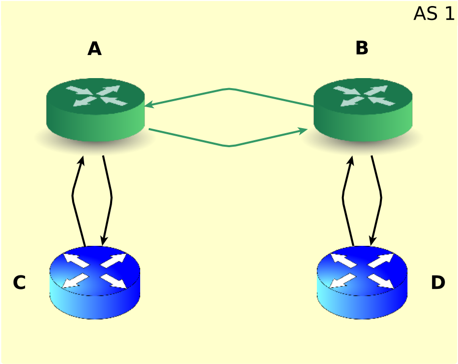

- Route reflection combined with the IGP weight distance metric can cause routing and forwarding anomalies.
- Policies are deployed (manipulating the local preference attribute) even in iBGP.

Goal and Contributions

How can an ISP apply the theory of algebraic frameworks, SSPP in particular, to verify if the iBGP configuration is guaranteed to be safe?

- Define a data structure to systematically check whether sufficient conditions for safety are met.
- Extend the SSPP model
  - increase the expressive power of the policies it can describe (through the use of communities)
  - model attributes important for iBGP such as IGP weight

More efficient checking of configuration correctness.

The Stratified Shortest Paths Problem (SSPP)

Algebraic model of policy-based routing based on the Semiring theory.

Path

Instead of *distance*, its attributes are (stratum, distance).

Arc

Instead of *weight*, it is characterized by (function, weight).

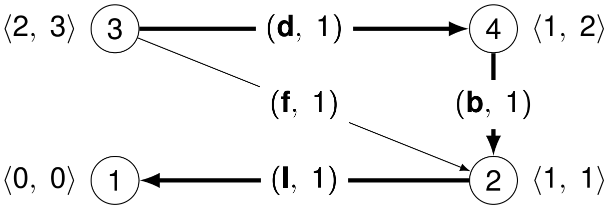

T.G. Griffin, The Stratified Shortest-Paths Problem, COMSNETS 2010

Sufficient Condition for Safety in SSPP

	0	1	2
b	0	1	∞
d	0	2	∞
f	0	∞	∞
l	1	∞	∞

For example,

(**l**, 1) $\langle 0, 0\rangle = \langle 1, 1\rangle$

(**d**, 1) $\langle 1, 2\rangle = \langle 2, 3\rangle$

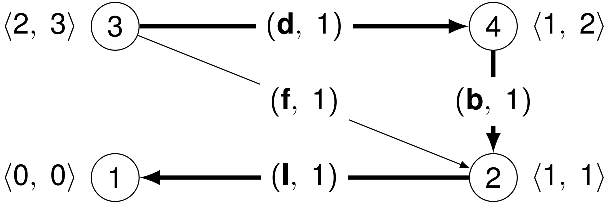

Safety

The sufficient condition for safety requires the strata function to be *inflationary*.

Example: Direct Application of SSPP

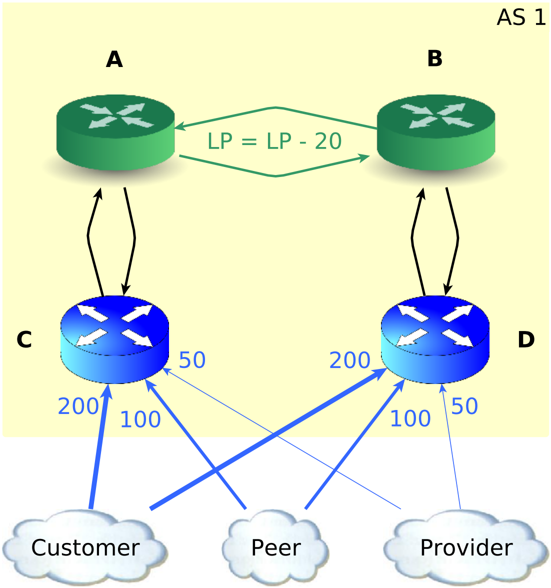

Example: Direct Application of SSPP

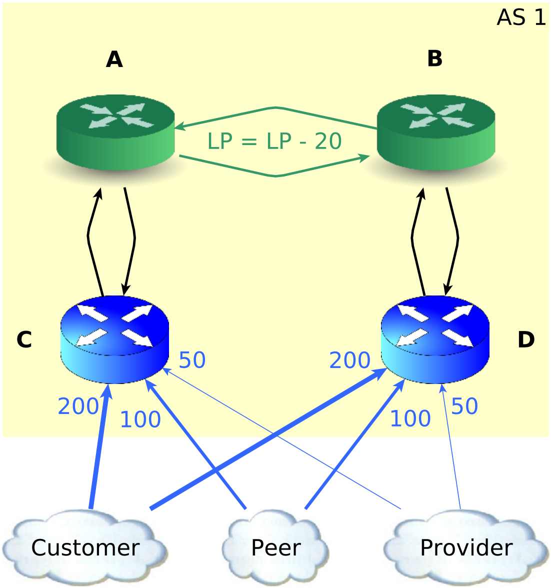

LP	Stratum
200	0
180	1
100	2
80	3
50	4
30	5

Example: Direct Application of SSPP

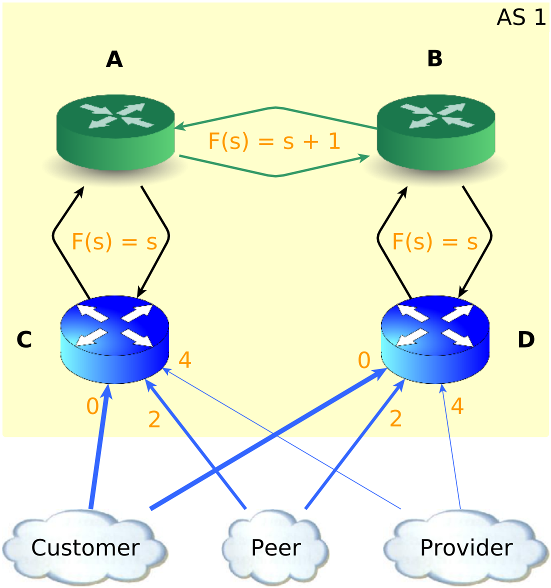

LP	Stratum
200	0
180	1
100	2
80	3
50	4
30	5

Example: Direct Application of SSPP

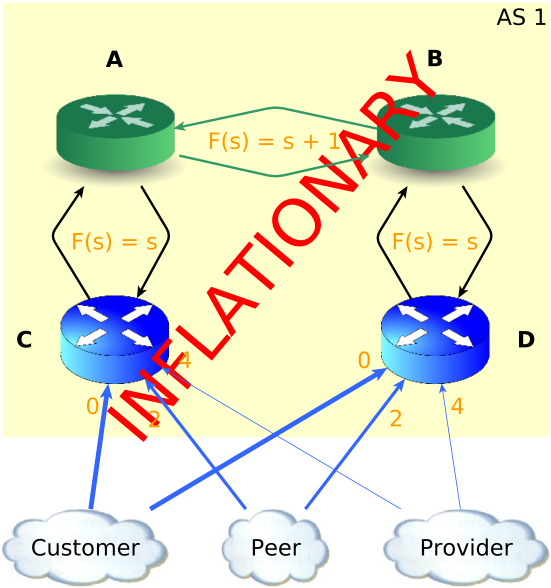

LP	Stratum
200	0
180	1
100	2
80	3
50	4
30	5

Same Policy, Different Configuration

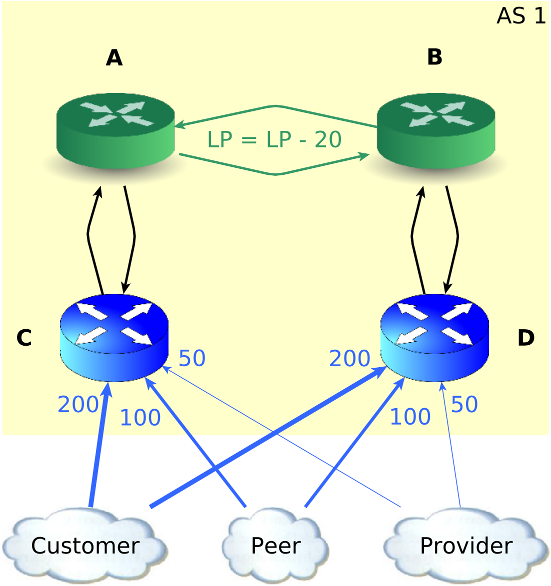

Same Policy, Different Configuration

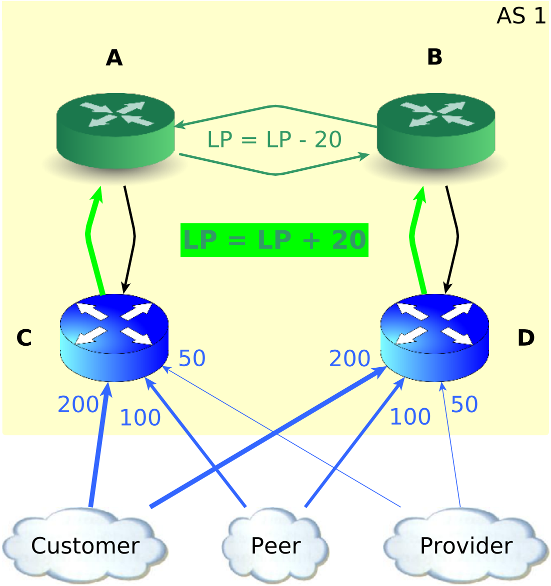

Same Policy, Different Configuration

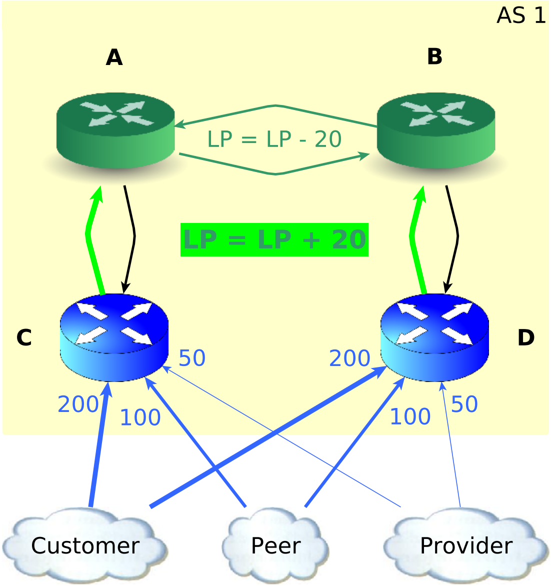

Previous LPs:
$\{200, 180, 100, 80, 50, 30\}$

New LPs:
$\{220, 200, 120, 100, 70, 50\}$

Same Policy, Different Configuration

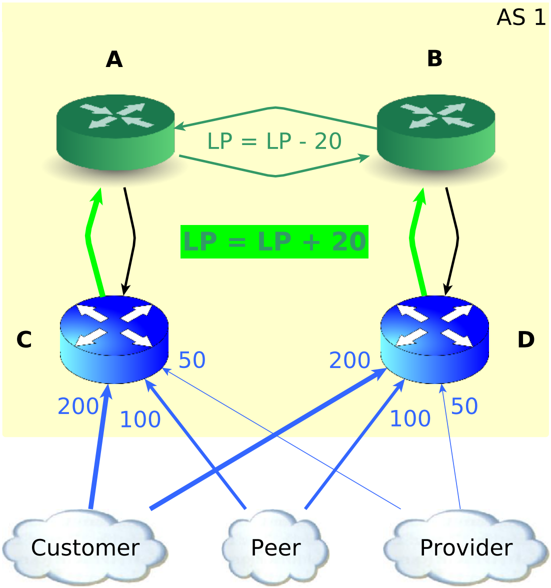

Previous LPs:
{200, 180, 100, 80, 50, 30}

New LPs:
{220, 200, 120, 100, 70, 50}

Try Again with Six Strata:
{0, 1, 2, 3, 4, 5}

Same Policy, Different Configuration

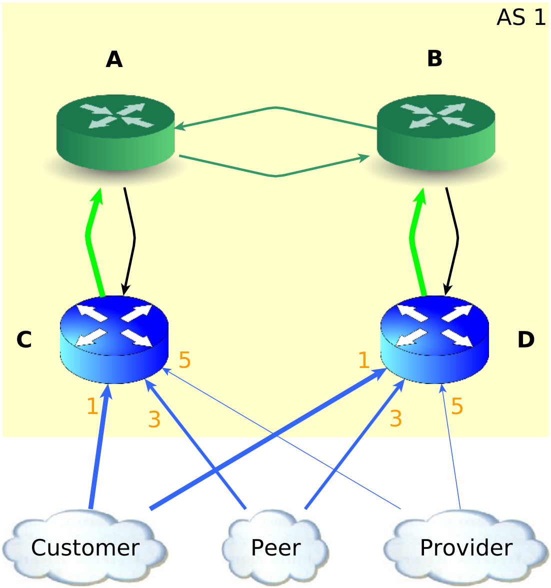

A → B, B → A:
F(0) = 1 | F(2) = 3 | F(4) = 5

A → C, B → D:
F(1) = 1 | F(3) = 3 | F(5) = 5

C → A, D → B:
F(1) = 0 | F(3) = 2 | F(5) = 4

Same Policy, Different Configuration

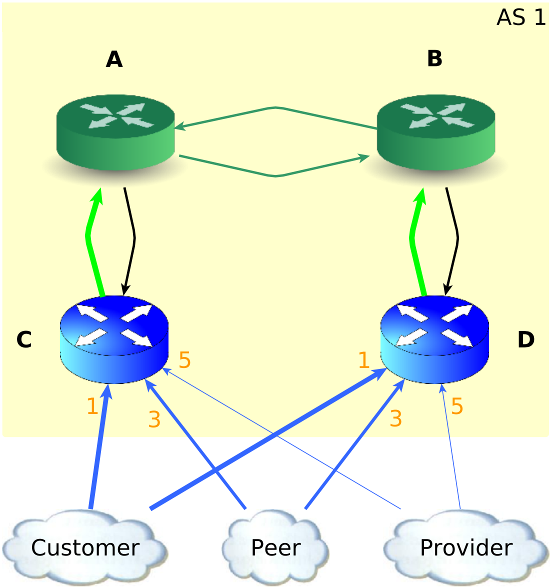

Fix C $\rightarrow$ A: F(1) = 0

Same Policy, Different Configuration

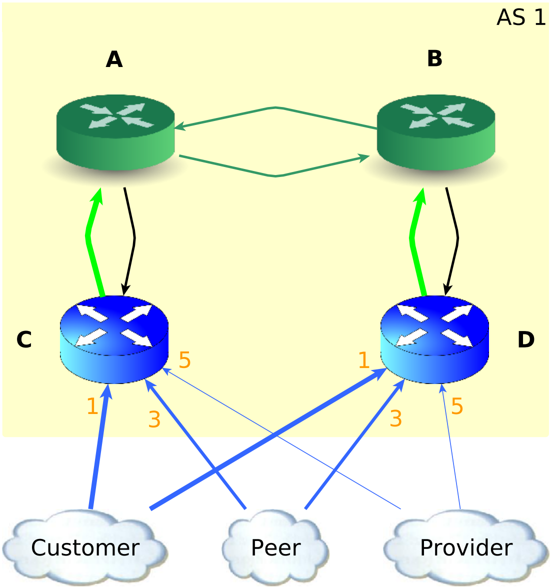

Fix C $\rightarrow$ A: F(1) = 0

- Replace with F(1) = 1

  Then, we change the semantics of A's preferences, because A $\rightarrow$ B: F(0) = 1

Same Policy, Different Configuration

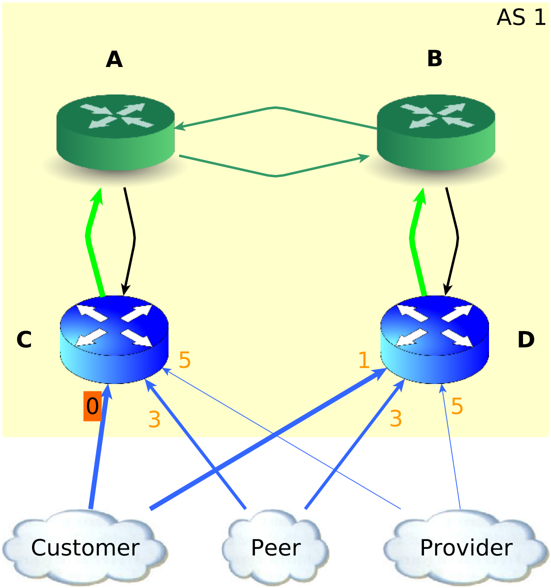

Fix C $\rightarrow$ A: F(1) = 0

- Replace with F(1) = 1

  Then, we change the semantics of A's preferences, because A $\rightarrow$ B: F(0) = 1

- Replace with F(0) = 0

  Then, the policy on the A $\rightarrow$ C link becomes non inflationary: F(1) = 0

A Data Structure to Capture Strata Dependencies

Two kinds of dependencies:

1. order of local preference values:
   if $l_{i1} > l_{i2} > l_{i3} > ...$ then $s_{i1} < s_{i2} < s_{i3} < ...$
2. inflationary property across BGP sessions:
   when a route with stratum s_a in router r_i is announced to r_i' and receives stratum s_b, then $s_b \geq s_a$

Strata Digraph

Nodes: strata values | Links: inequalities

If there is a cycle that involves strict inequalities, then there is no strata assignment to satisfy all dependencies.

Acyclic Strata Digraph

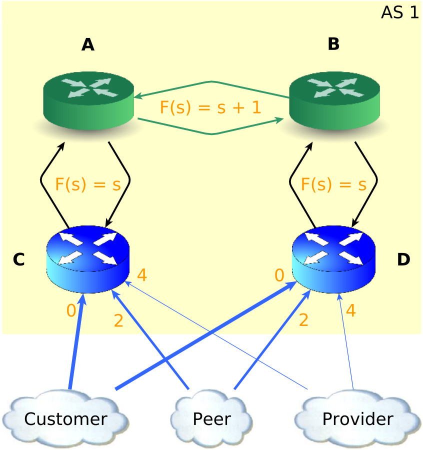

Acyclic Strata Digraph

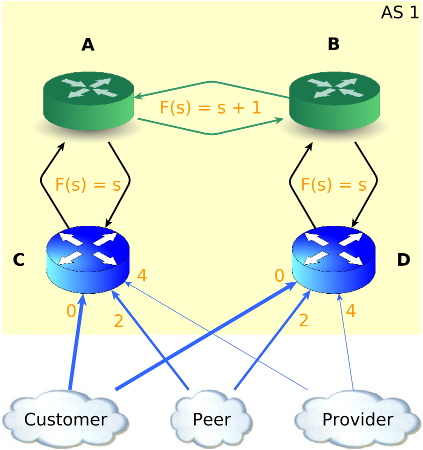

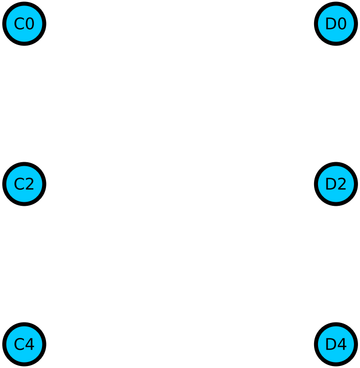

Acyclic Strata Digraph

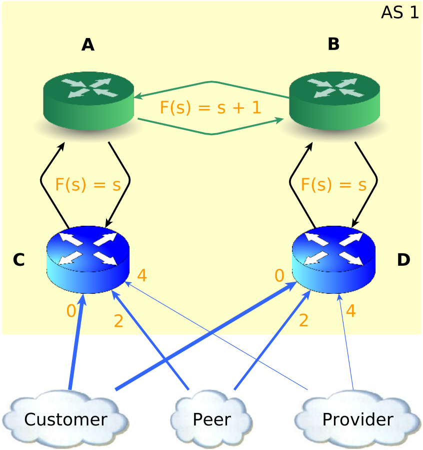

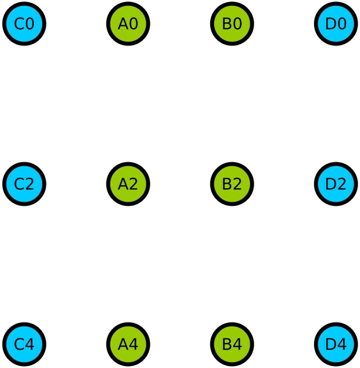

Acyclic Strata Digraph

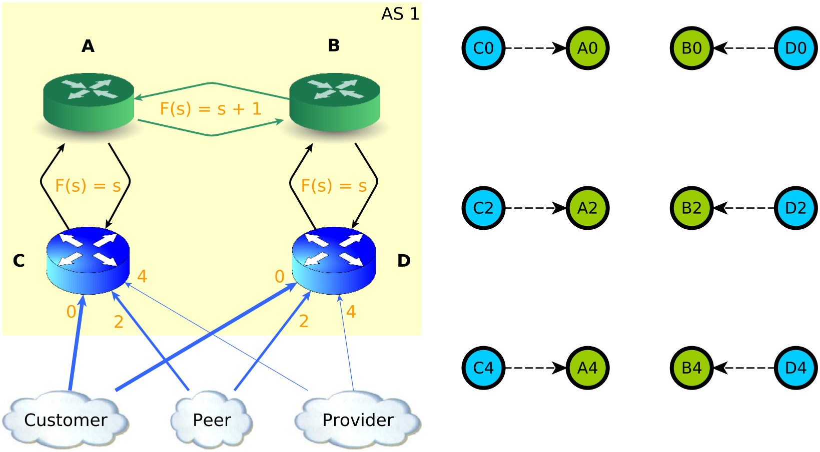

Acyclic Strata Digraph

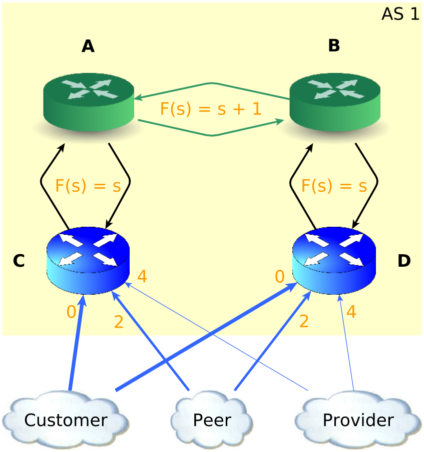

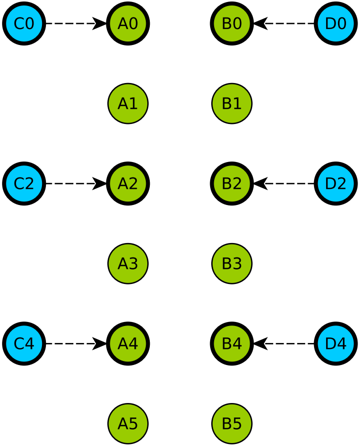

Acyclic Strata Digraph

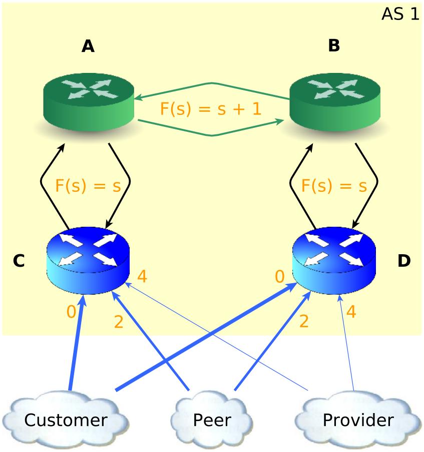

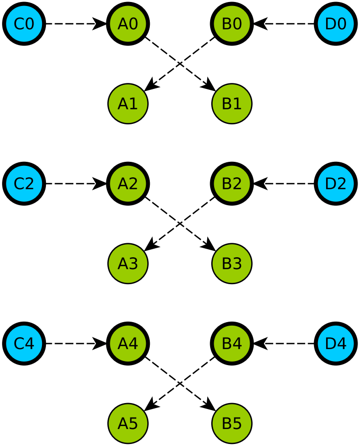

Acyclic Strata Digraph

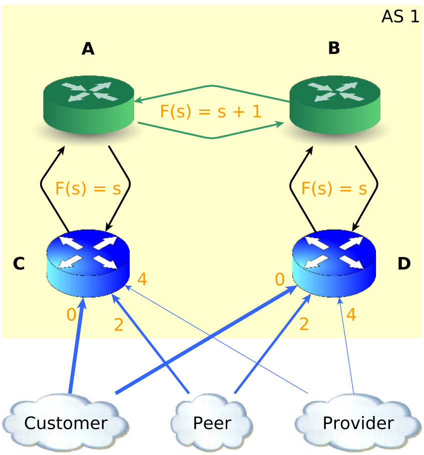

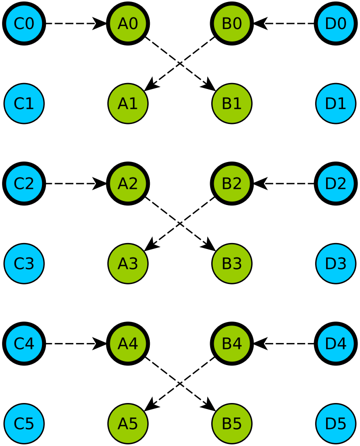

Acyclic Strata Digraph

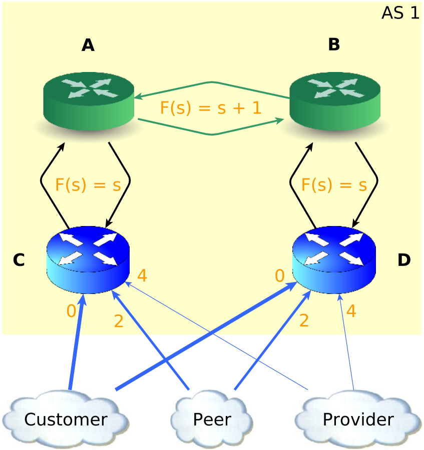

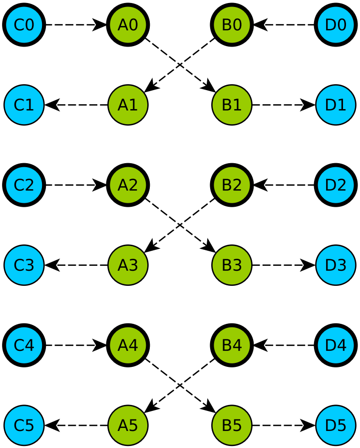

Acyclic Strata Digraph

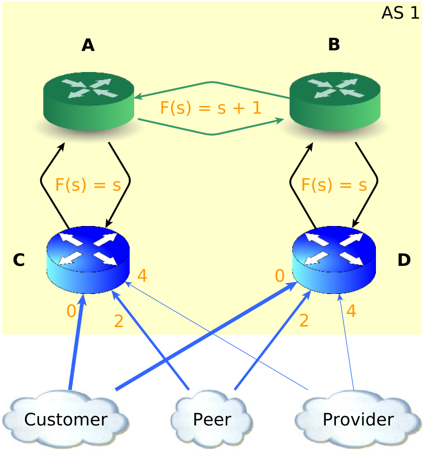

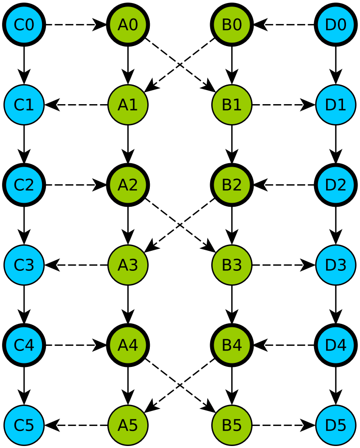

Strata Digraph With (Strict) Cycles

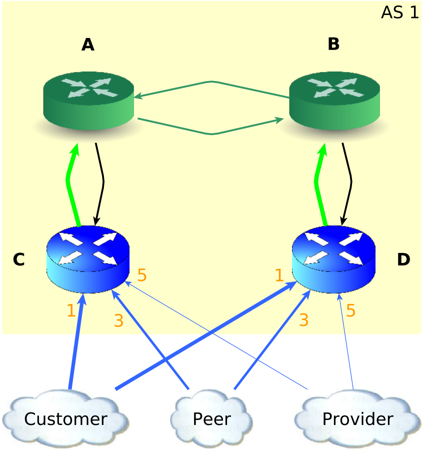

Strata Digraph With (Strict) Cycles

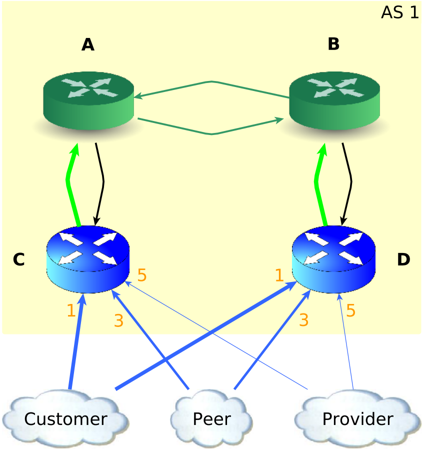

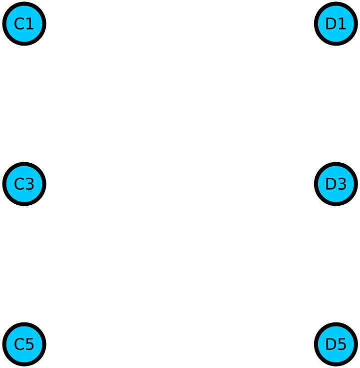

Strata Digraph With (Strict) Cycles

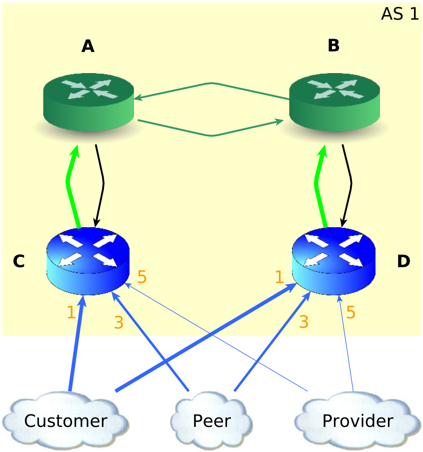

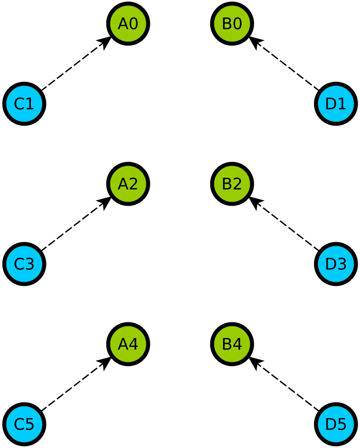

Strata Digraph With (Strict) Cycles

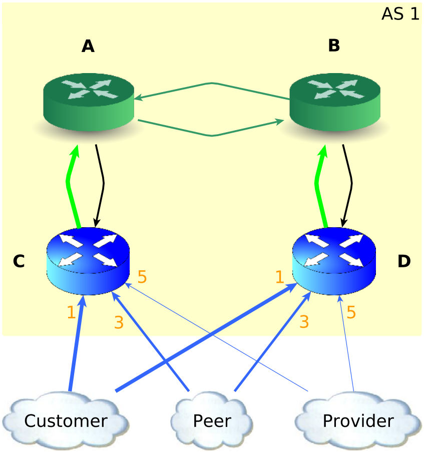

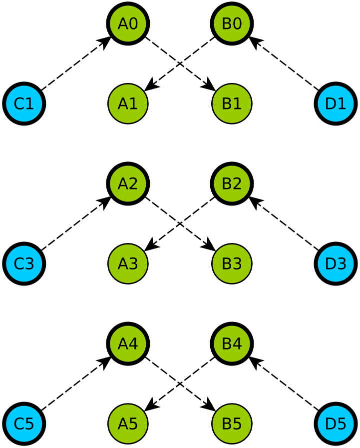

Strata Digraph With (Strict) Cycles

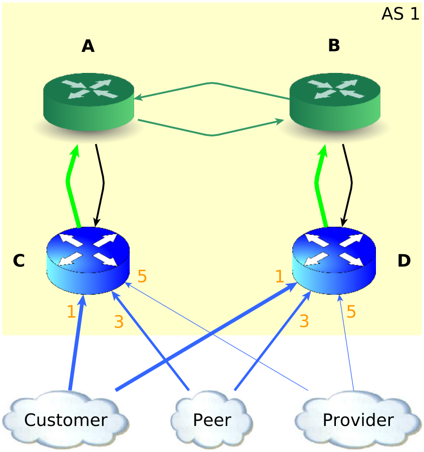

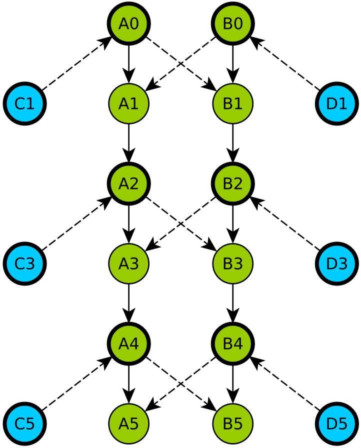

Strata Digraph With (Strict) Cycles

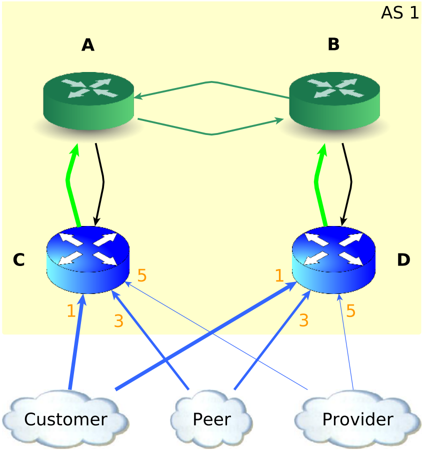

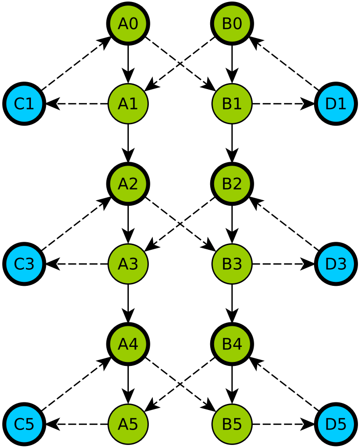

Strata Digraph With (Strict) Cycles

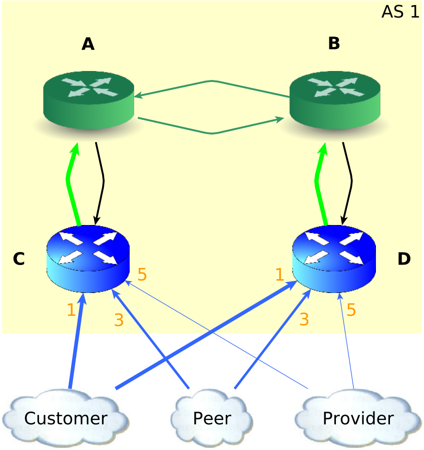

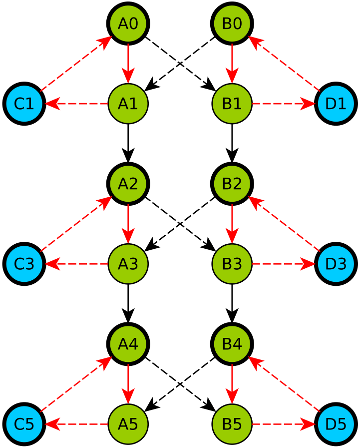

$$C5 \leq A4 < A5 \leq C5$$

Does AS 1 need to change the configuration to guarantee safety?

Not necessarily.

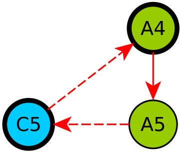

Step v can break the cycle.

We can model this additional step without adding elements to the model.

BGP Decision Process

(Weight)

i Highest Local Preference

ii Shortest AS path length

iii Lowest Origin type

iv Lowest MED

v eBGP-learned over iBGP-learned

vi ...

Does AS 1 need to change the configuration to guarantee safety?

Not necessarily.

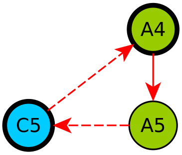

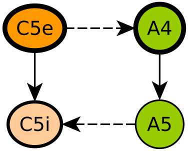

Step v can break the cycle.

We can model this additional step without adding elements to the model.

C5e: assigned to provider paths learned through *eBGP*

C5i: assigned to provider paths learned through *iBGP*

Cycles Gone

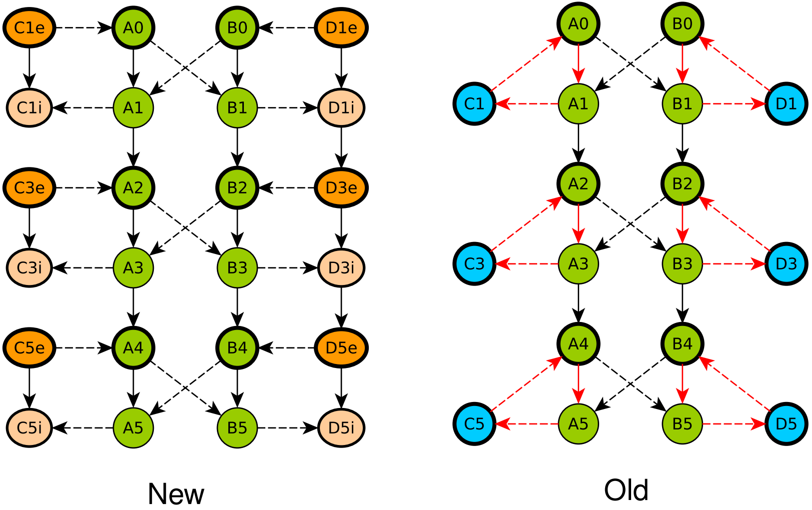
C1e
A0
B0
D1e
C1i
A1
B1
D1i
C3e
A2
B2
D3e
C3i
A3
B3
D3i
C5e
A4
B4
D5e
C5i
A5
B5
D5i
New
A0
B0
C1
A1
B1
D1
A2
B2
C3
A3
B3
D3
A4
B4
C5
A5
B5
D5
Old

Roadmap

So far:

- Introduced a systematic way to apply the SSPP model in iBGP configurations to check for safety.
- Used the same data structure to model BGP decision process steps *indirectly* without changing the model.

Next:

Increase expressive power of SSPP by adding elements to it.

Limitation 1

F(neighbor's stratum) = my stratum

What if I wish to apply different strata to routes that have the same stratum on my neighbor's side?

Adding Communities: iBGP

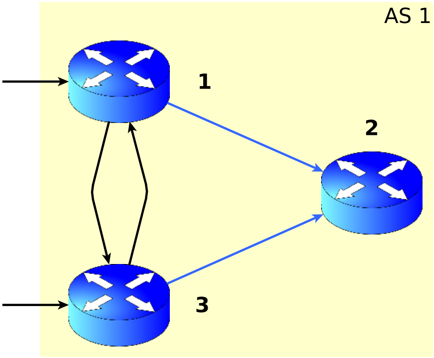

$1 \rightarrow 2$, $3 \rightarrow 2$:

$$F(s) = \begin{cases} \infty & \text{if comm-ibgp} \\ s & \text{otherwise} \end{cases}$$

Adding Communities: eBGP

For routes that receive the same stratum in one node, enable:

loop prevention and filtering

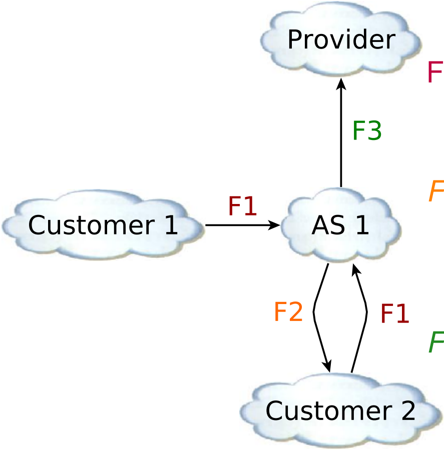

$F1(s) = s_1$

$$F2(s_1) = \begin{cases} \infty & \text{if comm-fromCust2} \\ s_2 & \text{otherwise} \end{cases}$$

$$F3(s_1) = \begin{cases} \infty & \text{if comm-fromCust1} \\ s_3 & \text{otherwise} \end{cases}$$

Modeling the IGP Weight

Limitation 2: need to model additional attributes in the BGP decision process for analysis of iBGP configurations

Solution: Each path is associated with a triple (s, d, w).

	Inter-AS Arcs	Within an AS Arcs
s	$f(s)$	$f(s)$
d	increase	no change
w	set to zero	increase by non-zero value

The w component must also be strictly inflationary. Require w:

1. to be strictly greater than zero, and
2. to increase when the iBGP path becomes longer.

Conclusion

- It is possible to check the safety of iBGP policies without path enumeration, using SSPP and a systematic methodology.
- We can model additional steps of the BGP decision process without adding features to the SSPP model under certain conditions.
- The extension of SSPP with communities allows it to model more iBGP and eBGP policies.
- When there is congruence between IGP and iBGP paths, the IGP weight step of the BGP decision process can also be added to the model.

Questions?

Thank you

Debbie Perouli
depe@purdue.edu

Distributivity: Locally *and* Globally Optimal Solution

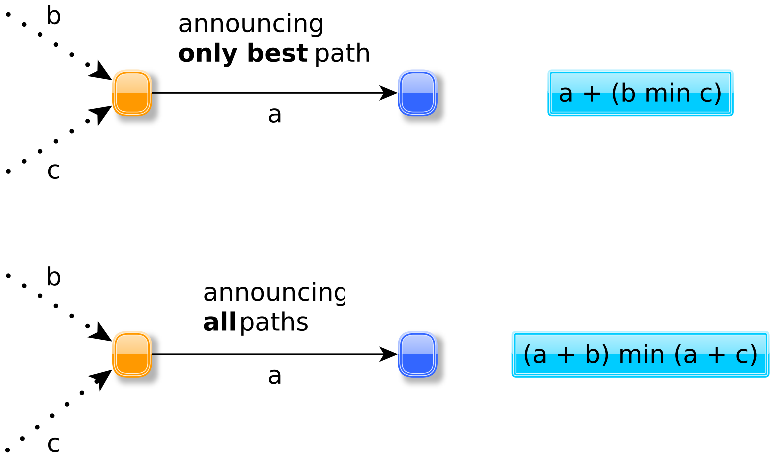

Shortest paths routing modeled through a $(\min, +)$ Semiring.

All Safe Policy Functions for Three Strata

	0	1	2		0	1	2
a	0	1	2	**m**	2	1	2
b	0	1	∞	**n**	2	1	∞
c	0	2	2	**o**	**2**	**2**	**2**
d	0	2	∞	**p**	2	2	∞
e	0	∞	2	**q**	2	∞	2
f	**0**	∞	∞	**r**	2	∞	∞
g	1	1	2	**s**	∞	1	2
h	1	1	∞	**t**	∞	1	∞
i	1	2	2	**u**	∞	2	2
j	1	2	∞	**v**	∞	2	∞
k	1	∞	2	**w**	∞	∞	2
l	**1**	∞	∞	**x**	∞	∞	∞

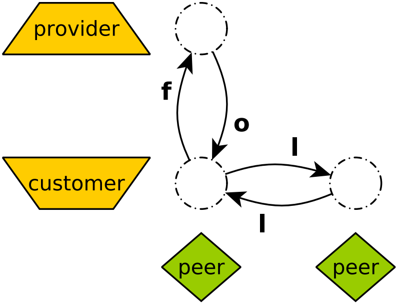

(direction of routing path)