A Distributed Method for Dynamic Resolution of BGP Oscillations

Ehoud Ahronovitz, Jean-Claude König, Clément Saad
University Montpellier 2 - LIRMM

International Parallel & Distributed Processing Symposium

25-29 april 2006

1. BGP presentation
2. Oscillation problem
   - SPP (Stable Paths Problem)
   - Dispute digraph
3. Our solution
   - Maintaining path local stateful information (PLSI)
   - Token principles
   - Coherence between routing policies
4. Conclusion

BGP presentation

Autonomous System (AS) is a set of machines managed by unique administration.

Each AS chooses its own internal routing (RIP,OSPF,...).

BGP used for **external routing** allows each AS to define its own routing policy.

Oscillations are due to **incoherences between policies**.

BGP presentation

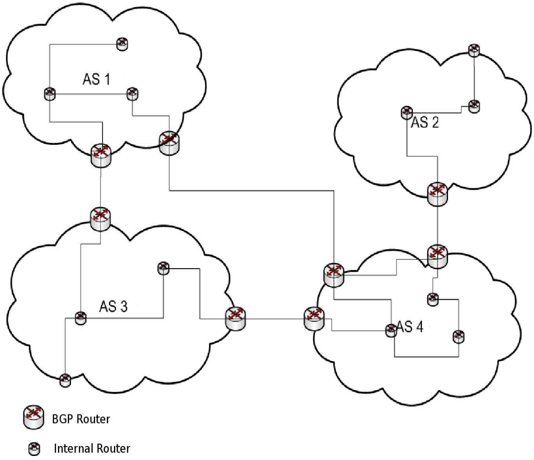

BGP presentation

Autonomous System (AS) is a set of machines managed by unique administration.

Each AS chooses its own internal routing (RIP,OSPF,...).

BGP used for **external routing** allows each AS to define its own routing policy.

Oscillations are due to **incoherences between policies**.

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

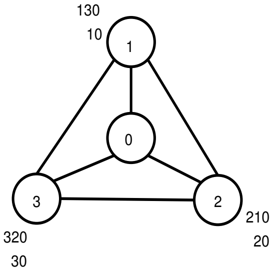

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

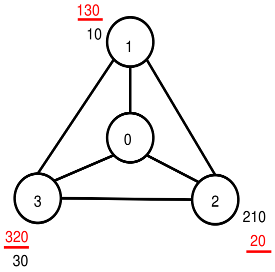

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

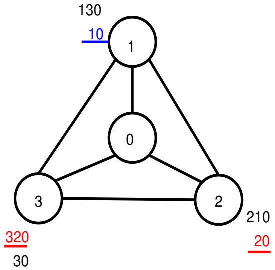

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

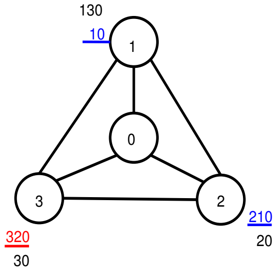

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

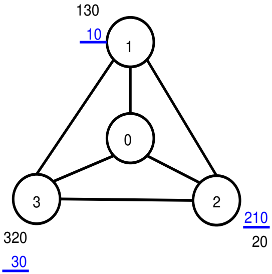

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

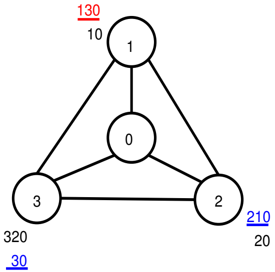

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

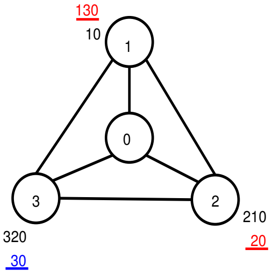

SPP (Stable Paths Problem)

SPP (Stable Paths Problem) (Griffin & Wilfong [1, 2])

- Each node represents an AS and each edge represents a BGP link.
- AS defines a list of paths ordered by preference related to its own policy.

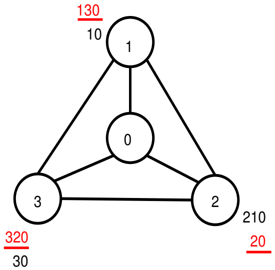

Dispute digraph

- Each node represents a path.
- Dotted lines represent transmission arcs.
- Full lines represent dispute arcs.

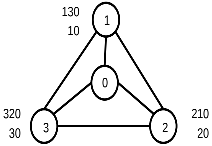

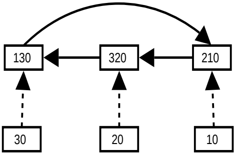

Dispute digraph

Theorem

If the dispute digraph related to an instance S of SPP is acyclic, then S contains a stable solution.

1. BGP presentation
2. Oscillation problem
   - SPP (Stable Paths Problem)
   - Dispute digraph
3. Our solution
   - Maintaining path local stateful information (PLSI)
   - Token principles
   - Coherence between routing policies
4. Conclusion

Our solution :

- allows to detect oscillation due to cycles in the dispute digraph.
- marks barred one path in the cycle, in order to break it.

Maintaining path local stateful information (PLSI)

We need only local information to detect an oscillation and thus we respect private policy choices as imposed by BGP.

- Each AS manages **locally** states (+ or -) of its paths.

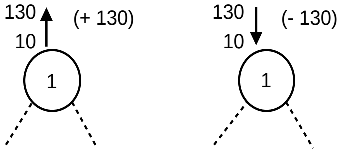

- If an AS detects a state change from state + to state - on one of its paths then it concludes that this path oscillates.

Maintaining path local stateful information (PLSI)

	AS1			AS2			AS3		
step	130	10	rib-in	210	20	rib-in	320	30	rib-in
1	*	*	10	*	*	20	*	*	320
2	*	*	10	+	*	210	*	*	320
3	*	*	10	+	*	210	-	*	30
4	+	*	130	+	*	210	-	*	30
5	+	*	130	-	*	20	-	*	30
6	+	*	130	-	*	20	+	*	320
7	-	*	10	-	*	20	+	*	320

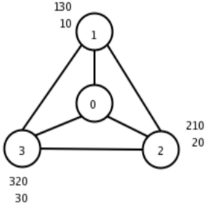

rib-in : current path

Maintaining path local stateful information (PLSI)

Two important questions :

- When a path oscillates, how can we know that it belongs to a cycle in the dispute digraph ?
- When an oscillation occurs, all paths belonging to the cycle will oscillate. Which path should we mark barred ?

Token principles

Detection of cycle :

1. When an AS detects an oscillation on path X, it generates a token jX and joins it with its new BGP announce.

Token principles

Detection of cycle :

1. When an AS detects an oscillation on path X, it generates a token jX and joins it with its new BGP announce.
2. After reception of this message, if an AS has to adopt a new path, it forwards the token jX with the announce of its new path.

Token principles

Detection of cycle :

1. When an AS detects an oscillation on path X, it generates a token jX and joins it with its new BGP announce.
2. After reception of this message, if an AS has to adopt a new path, it forwards the token jX with the announce of its new path.
3. If the generator of jX retreives its token and has to adopt path X, it concludes that X belongs to a cycle and marks X barred.

Token principles

Detection of cycle :

1. When an AS detects an oscillation on path X, it generates a token jX and joins it with its new BGP announce.
2. After reception of this message, if an AS has to adopt a new path, it forwards the token jX with the announce of its new path.
3. If the generator of jX retreives its token and has to adopt path X, it concludes that X belongs to a cycle and marks X barred.

Please note that the value jX can be assigned with a hashtable, in order to respect confidentiality.

Token principles

Only one path should be marked barred :

- All ASes having a path belonging to a cycle will retreive their tokens.
- It is required to define a total order relation on tokens in order to mark barred only one path.
- When an AS receives a token, it checks if this token has a higher priority than its own token. If yes, it forwards this token, otherwise it ignores it.

Token principles

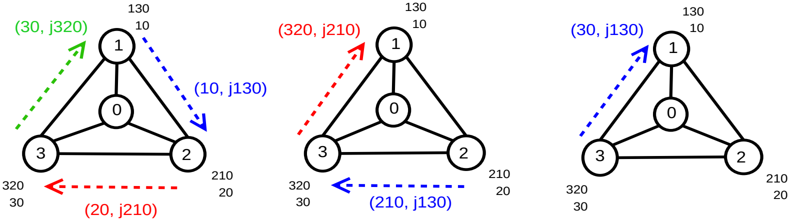

Coherence between routing policies

Coherence between routing policies

Definition of relation $<_\alpha$

locally : Let A be an AS ; $\forall$ P, Q paths of A, if P is preferred to Q then $P <_\alpha Q$.

globally : $\forall$ P, Q paths belonging to two differents Ases, if P is a sub-path of Q then $P <_\alpha Q$.

Coherence between routing policies

Coherence between routing policies

The local relation respects private policy of each AS,

the global relation maintains the coherence between policies of differents ASes.
In fact, whatever the policy, a sub-path of any path is obviously better than the whole path.

Coherence between routing policies

Theorem

If $<_\alpha$ is a strict order relation then the policies are coherent between themselves.

Coherence between routing policies

New dispute digraph

$$130 \underbrace{<_\alpha}_{local} 10 \underbrace{<_\alpha}_{global} 210 \underbrace{<_\alpha}_{local} 20 \underbrace{<_\alpha}_{global} 320 \underbrace{<_\alpha}_{local} 30 \underbrace{<_\alpha}_{global} 130$$

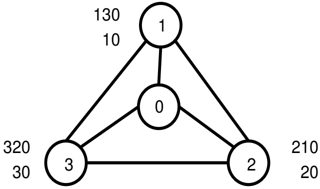

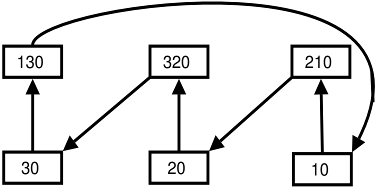

1. BGP presentation
2. Oscillation problem
   - SPP (Stable Paths Problem)
   - Dispute digraph
3. Our solution
   - Maintaining path local stateful information (PLSI)
   - Token principles
   - Coherence between routing policies
4. Conclusion

Conclusion

Summary :

1. Detection and resolution of oscillations
   - Dynamic and distributed method
   - Maintaining path local stateful information (PLSI)
   - Tokens are added to BGP messages
2. Coherence between routing policies

Conclusion

Perspectives :

- We adapted our solution in order to take into account failures or appearances of links
- We must check our solution on a simulator
- Manage Byzantin beahaviors
- Manage the connectivity problem in BGP

- Timothy G. Griffin and Gordon Wilfong, *A Safe Path Vector Protocol*, Proc. IEEE INFOCOM, 2000.
- Timothy G. Griffin and F. Bruce Sherpherd and Gordon Wilfong, *Policy Disputes in Path-Vector Protocols*, Proc. 7th Int. Conf. Network Protocols (ICNP'99).
- Lixin Gao and Jennifer Rexford, *Stable Internet Routing Without Global Coordination*, in Proc. ACM SIGMETRICS, 2000.
- Selma Yilmaz and Ibrahim Matta, *A Randomized Solution to BGP Divergence*, in Proc. of the 2nd IASTED Int. Conf. on Communication and Computer Networks (CCN'04).