

Resolving Inter-Domain Policy Disputes

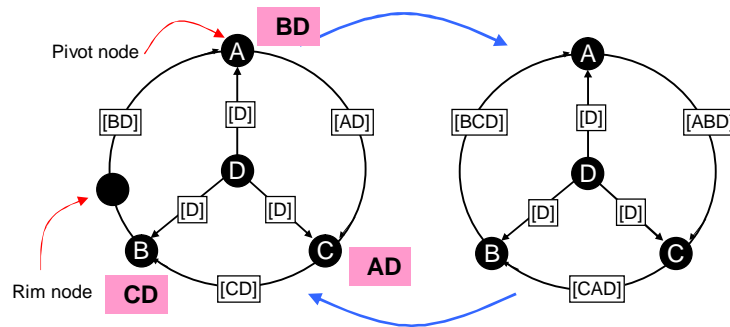
Cheng Tien Ee, **Byung-Gon Chun**, Vijay Ramachandran,
Kaushik Lakshminarayanan, Scott Shenker

Introduction

- Internet divided into different autonomous systems (ASes), controlled by competing Internet service providers (ISPs)
- ISPs administer their networks independently
 - No global view of all the ASes
 - No global control
 - ISPs' policies reflect business relations, thus not revealed
- Conflicts in policies can cause persistent route oscillations
 - Hard to detect and to resolve

Problem

- Assume a node represents an AS
- Conflicts between AS policies can result in dispute wheels (Griffin, Shepherd, Wilfong 2002)



- * Prefer anti-clockwise neighbor
- * Filter routes with length > 2

Sigcomm '07

3

Related Work: Good News

- Hierarchical business structure results in convergence (Gao, Rexford 2001)
 - Relationship either peer-to-peer, customer-provider, or backup
- But
 - Verification still requires exposing business relationship
 - Misconfigurations can still cause oscillations
 - Prefer to separate economics and network convergence

Sigcomm '07

4

Related Work: Bad News

- Search for local constraints
 - Able to check locally, yet sufficient for global convergence
- With complete autonomy and filtering expressiveness, essentially only shortest path routing can guarantee convergence (Feamster, Johari, and Balakrishnan, 2005)
 - ISPs cannot use shortest path since they are unlikely to relinquish policy autonomy, for economic reasons

Sigcomm '07

5

Our Approach


- A simple extension to BGP that constrains policy choices only when a persistent oscillation is detected
- Design Goals
 1. Do not reveal any ISP policies
 2. Distributed, online dispute detection and resolution
 3. If no dispute resulting in oscillations exist, have BGP decision as it is now
 4. Account for transient oscillations, don't unnecessarily dismiss routes

Sigcomm '07

6

High-level Idea

- Node notices that it selects alternatively routes that are more / less preferred than previous route
- Node locally thinks it *may* be involved in dispute-induced oscillations
- Precedence metric
- Indicates this involvement by incrementing a precedence value in route advertisement
- If in turn receives route from neighbor that also thinks it's involved in dispute, confirms presence of dispute wheel



7

Precedence Metric

- Maintain *history table* of routes encountered during oscillations
- Precedence metric consists of global and local values

Ranked based on
current BGP
decision process ↓


AS Path	Global Precedence (Incoming)	Local Precedence (Computed)
P0	2	0
P1	1	1
P2	0	1

- Incoming global value used in first step of route selection decision process
- Choose route (P2)
- Advertise route (P2) with route's (P2's) outgoing global value
= incoming global value + route's local value = 0+1 = 1

Sigcomm '07

8

Precedence Metric in Action



C's History Table


AS Path	Global Precedence	Local Precedence
[AD]	0 1	0
[D]	0	1

Sigcomm '07

9

No Dispute

- On the other hand, if no disputes exist...



- By induction, all nodes eventually pick their most preferred next-hops
- Global precedence values advertised are all 0

Sigcomm '07

10

Properties of Precedence Metric

- Uses the circular nature of dispute wheels to force shortest path, and only when disputes that cause oscillations exist
- If routes encountered during previous policy-induced oscillations are stored and we use the Precedence Metric, no further policy-induced oscillations can occur
- If, after convergence, non-zero global precedence values exist, \Rightarrow dispute(s) exist
- Only global value advertised, no other routes or policies revealed

Sigcomm '07

11

Challenges in Practice

1. Of the routes encountered during oscillation, which routes to store and for how long?
2. Distinguish between transient and persistent oscillations
3. Provide sufficient visibility to ease troubleshooting, but should not reveal any ISP policies

Sigcomm '07

12

Which Routes and How Long?

- Which routes to store for dispute detection:
 - More preferred, unavailable routes stay in memory to detect disputes
 - These routes are used to confirm disputes resulting in oscillations
- How long should routes be stored?
 - Short-term: time taken to detect disputes (bounded by propagation time around wheel)
 - Long-term: as long as disputes exist

Transient Oscillations

- More preferred but unavailable routes eventually expire
 - The advertised routes' precedence not incremented
- Long-term: If any pivot changes policy, clear states associated with previously resolved disputes

Visibility

- What level of visibility is provided for troubleshooting?
 - Route numbers (router identifier, AS #, seq #) issued by pivots
 - Thus, identities of pivot nodes advertised together with routes
- No ISP policies revealed
 - Route numbers refer, or point to, routes that were previously advertised

Evaluation

- Built an event-based, packet-level, asynchronous simulator
- Parameters:
 - Route updates batched, occurs every Minimum Route Advertisement Interval (MRAI) – 30 seconds
 - Processing delay jitter – randomly selected in $[0,1]$ second
 - Link propagation delay – 10 ms
- Two types of graphs:
 - Simple graph that consists of a destination and rim and pivot nodes
 - AS-level network topology from RouteViews

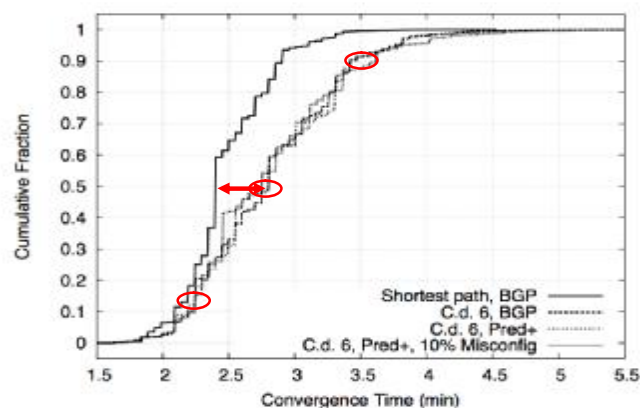
RouteView Graph

- Retrieved route dump on 3rd Jan '07
 - Consists of 24307 ASes, 56914 inter-AS links
- AS policies determine possible network paths, but precise policies hard to obtain
- Need to construct reasonable paths (but not always shortest)
 - Aim to obtain similar route inflation ratio reported by Tangmunarunkit et al. (2001)
- Use a depth + breadth-first search to determine next hop, set as most locally preferred

Sigcomm '07

17

Convergence Time (RouteView)



- Compared to Shortest path + BGP, half of nodes take 20% more time to converge
- Insignificant change to convergence time (with disputes)

Sigcomm '07

18

Conclusion

- Precedence solution enforces stability only when disputes causing oscillations exist
 - No additional ISP policies revealed
 - No need for global view of Internet
 - Routes due to transient oscillations are not dismissed
- Provide just enough visibility to pinpoint routers where policies conflict
 - Eases network troubleshooting
 - Allows ISPs to have precedence value 0 by changing their local policies
- Practical and minimal-impact way to solve BGP policy-induced stability problems



Sigcomm '07

19

Thanks for listening

Simple Graph

- Simple graphs




- Convergence time depends on number of rim nodes between two adjacent pivots, not total number of nodes

Sigcomm '07

21

Required Memory (RouteView)




- In average, about 50% more memory required

Sigcomm '07

22

Multiple Nodes in an AS

- Distinguish between ingress and egress global precedence values



- Ingress global precedence value used for all route decisions within an AS
- Output of those decisions stored in an egress global precedence value and sent to neighboring routers in other ASes

Sigcomm '07

23

MED Handling

- MED involves non-strict preferences; ranking of route A with respect to B depends on another route C
- First detect such a scenario
- Ignore an incoming route causing the previous most locally preferred route demoted in computing the global precedence value
- Remember the previous most preferred route which in this case may still be feasible, but no longer most preferred

Sigcomm '07

24

Handling Misbehaving ASes


- Misbehaving AS: incoming route with lowest global value not selected, global values not properly updated
- Requires a monitor that keeps track of hashed incoming and outgoing routes of the AS
- Operation of a monitor: beyond the scope of our work

SPVP vs. Precedence


SPVP	Precedence
protocol for solving abstract stable path problems sketch the basic ideas of a safe BGP under simplifying assumptions	fleshed out all the details to apply to BGP
path history (route change events) carried on all route advertisements	a single global value when there is no dispute Kicked in only when disputes causing oscillations exist; thus minimal-impact
	tell a dispute wheel exists easily tell pivot nodes for troubleshooting
message size prop. to the number of nodes in the oscillation	message size prop. to the number of pivot nodes in the oscillation

Adaptive Filtering

- Adaptive filtering
 - Filtering and choosing most preferred route may result in no feasible route




- Solution: monitor routes advertised, reuse of supposedly filtered routes \Rightarrow adaptive filtering



Misbehavior

- Misbehavior
 - Incoming route with lowest global value not selected



- Might be permanent or adaptive filtering
- Adaptively filter so AS gets to select its most preferred route