Understanding BGP Misconfiguration

Ratul Mahajan
David Wetherall
Tom Anderson

University of Washington

Motivation

- BGP instabilities have widespread impact
- Misconfigurations can be a leading cause of unreliability
  - BGP is complex to configure
  - Known major incidents
- Little is known about misconfiguration in BGP
  - Only anecdotal evidence
- Use our experience to avoid future mishaps

Understanding BGP misconfiguration

- A systematic study to understand the problem
  - How common are misconfigurations?
  - What is their impact on connectivity and routing load?
  - Why do they happen?
  - How can we stop them?

- Approach
  - Leverage global visibility of BGP actions to detect misconfigs
    - Data from 23 BGP speakers in the backbone
  - Obtain operator feedback through an email survey

Border Gateway Protocol (BGP)

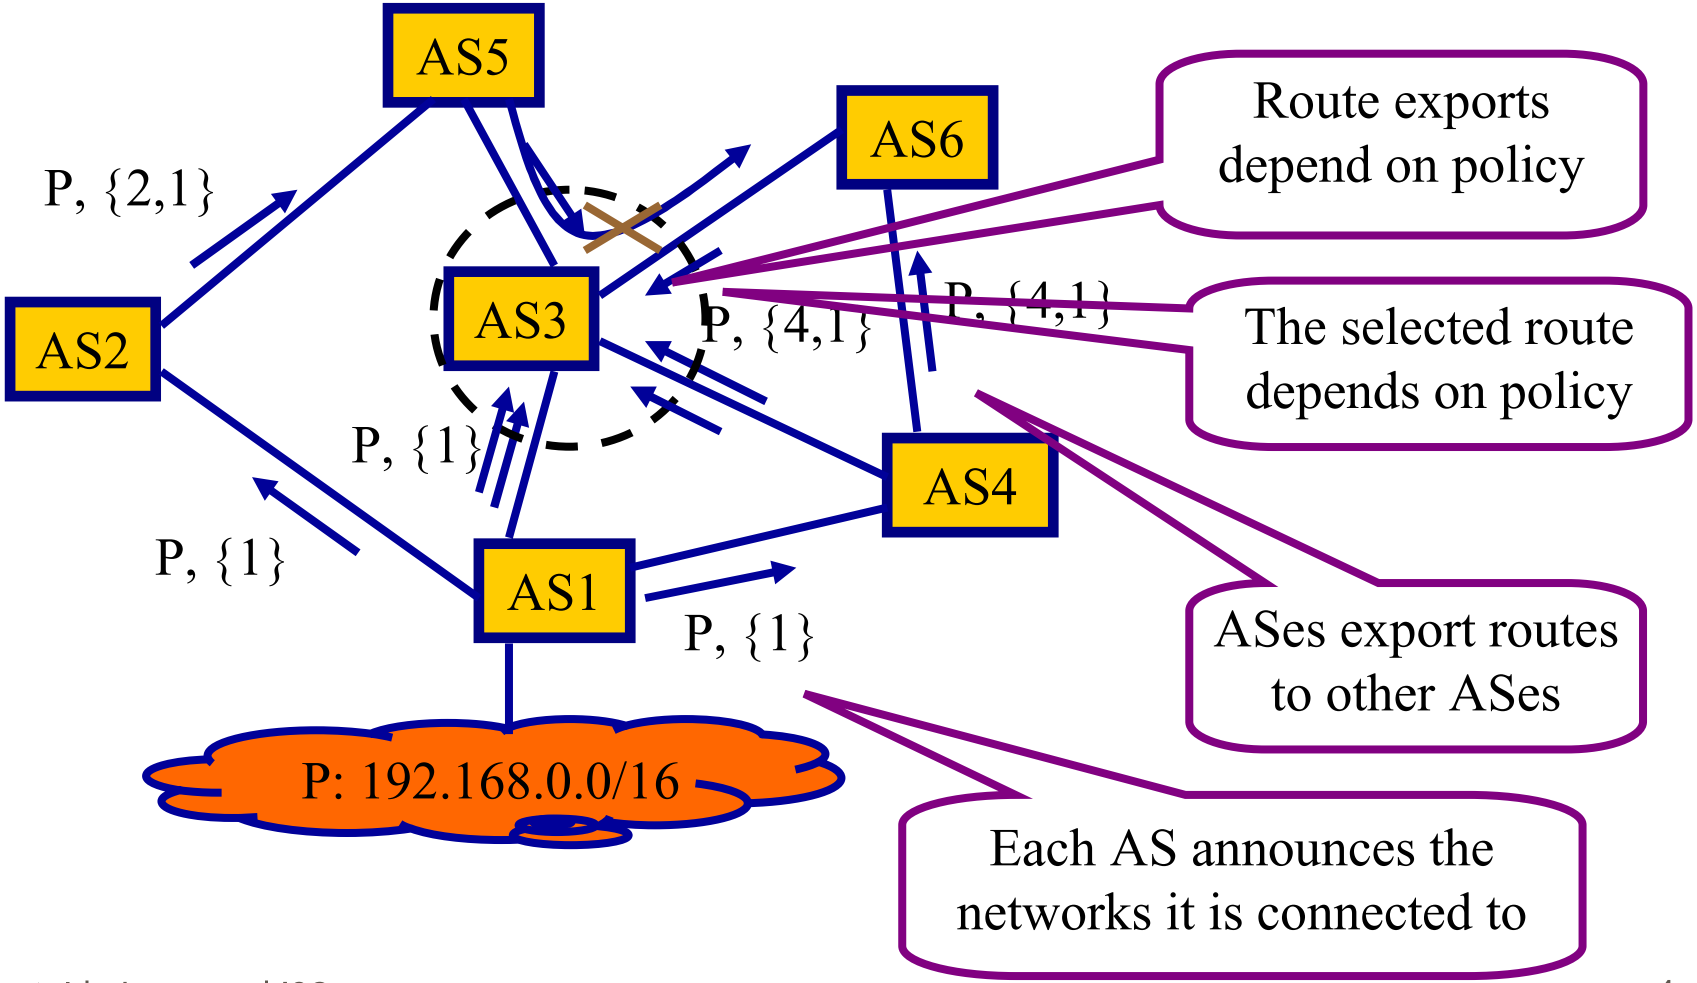

Route exports depend on policy

The selected route depends on policy

ASes export routes to other ASes

Each AS announces the networks it is connected to

BGP Misconfiguration

- No universally accepted list of “Dos & Don'ts”
- Defined as behavior unintended by the operator
  - Includes both *slips* (inadvertent errors) and *mistakes* (erroneous plan)

- We study two broad classes of globally visible faults
  - Origin misconfiguration
  - Export misconfiguration

BGP Misconfiguration (2)

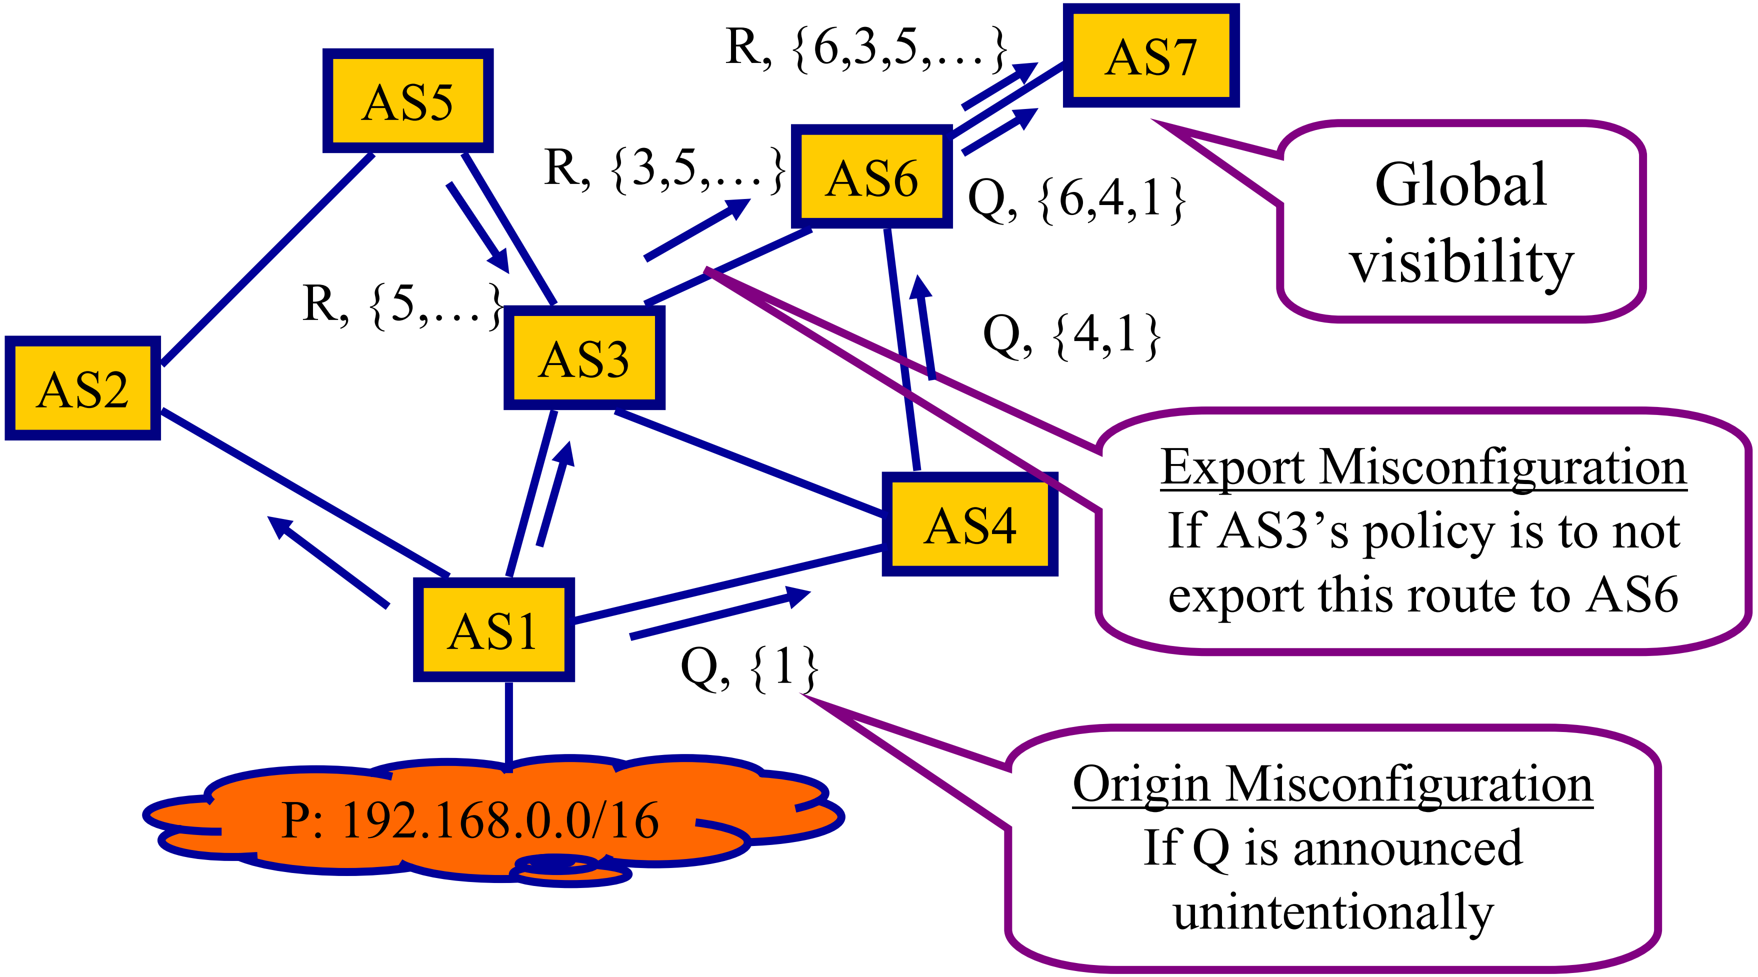

Global visibility

Export Misconfiguration
If AS3's policy is to not export this route to AS6

Origin Misconfiguration
If Q is announced unintentionally

Methodology

- Analyze updates from 23 BGP speakers for 21 days [route-views]
  - Rich view of backbone routing
  - Ability to observe even very short-lived events

- Identifying misconfiguration
  - IRRs are inaccurate or outdated
  - Instead use signature of misconfigs in the update stream
    - Policy changes have similar signature but bigger timescales

Methodology (2)

1. Identify short-lived (< 24hrs) changes as ***potential*** misconfigs
   - Origin misconfiguration
     - § Short-lived new route – new prefix or new origin for a prefix
   - Export misconfiguration
     - § Short-lived AS-path that violates policy
     - § Infer AS relationships using Gao’s heuristics

2. Email verification through operators
   - Was it a misconfig? Connectivity disrupted? What caused it?

3. Use email responses to discover underlying causes

u Test connectivity using public traceroute servers

- Coarse independent verification of email responses

Results: Origin misconfiguration

	Potential misconfigs per day	Email responses (% of potential)	Misconfigs (% of email)	Connectivity (% of misconfigs)
Prefixes	605	352 (58%)	339 (96%)	13 (4%)
Incidents	178	52 (29%)	45 (86%)	6 (13%)

- Misconfiguration detection accuracy is high
- Large number of misconfigurations
    - Extrapolated estimate is 580 (605 * 0.96) prefixes per day
    - 3 in 4 new routes seen in a day result from misconfigs
- Most misconfigurations don’t disrupt connectivity

Results: Export misconfiguration

	Potential misconfigs per day	Email responses (% of potential)	Misconfigs (% of email)	Connectivity
Paths	96	64 (66%)	61 (96%)	-
Incidents	35	12 (36%)	10 (86%)	-

- Misconfiguration detection accuracy is high
- Almost no impact on connectivity
  - But congestion experienced

Routing load

- Defined as fraction of updates due to misconfigs
  - = (bad updates) / (total updates)

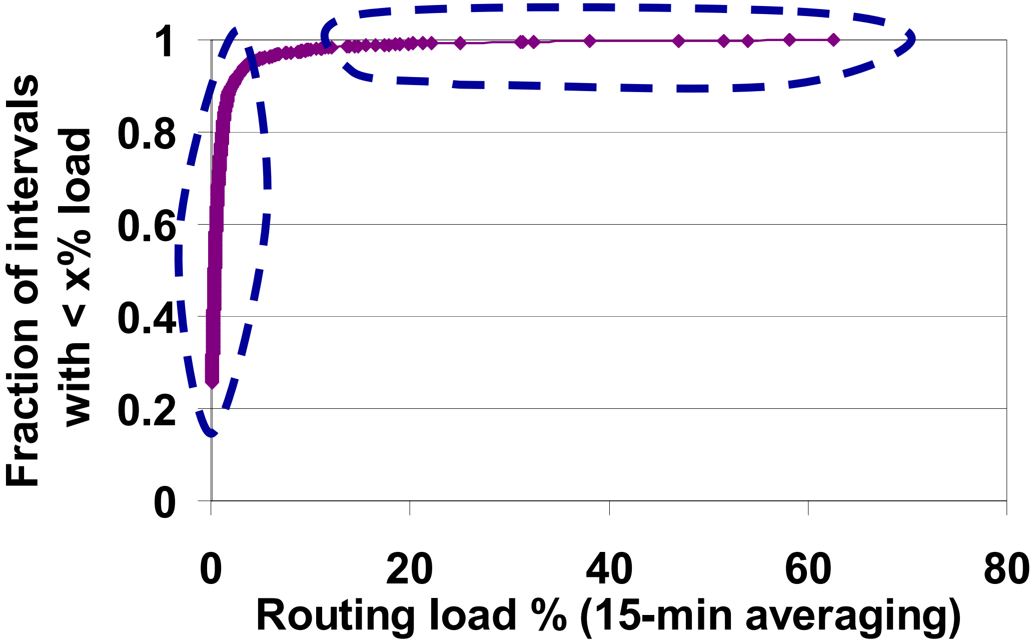

Some misconfigs cause extreme short-term routing load

Causes: Origin misconfiguration

- Faulty redistribution (32% prefixes/ 5% incidents)
  - Errors in propagating IGP routes into BGP
- Initialization bug (22% / 5%)
  - Leaking routes temporarily during boot-up or maintenance
- Reliance on upstream filtering (14% / 46%)
  - Announcing routes assuming upstream would filter them
- Hijacks (1% / 6%)
  - Announcing somebody else’s address space
- Old configuration (1% / 4%)
  - Reactivation of stale configuration

Prefix based (mis)configuration

- Prefix based configuration was responsible for 22% of the export misconfig incidents

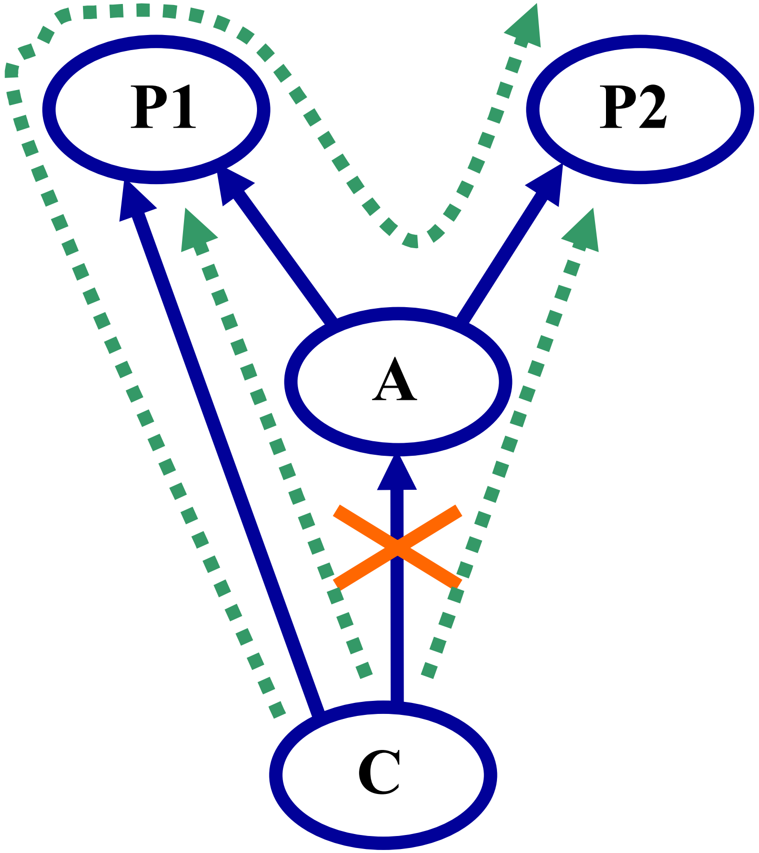

Intended policy at A: Provide transit of C through link A-C

Configuration: Export all prefixes originated by C to P1 and P2

The misconfiguration is exposed when the link A-C fails

Fixes (largely speculative)

- User interfaces
  - Basic principles need to be followed
  - High-level configuration tools built into the routers
- Configuration checker
- Automated verification

- Expose errors
- Appropriate configuration semantics
- Consistent databases and updated registries

Conclusions

- Misconfigurations are commonplace
- Connectivity is surprisingly robust to most misconfigs but routing load can be significant
- The causes of misconfigurations are diverse
  - Much needs to be done to improve the operational reliability of the Internet

On email surveys

- Don’t worry. That was a configuration error of our upstream ISP.
- Yes, we know this is not a recommended way of doing things; but the packet monster of the internet must be fed.
- I am writing to thank you for your letter and say that I am glad that someone apart from me is interested in our BGP announcements.
- Hope you enjoy living in Seattle; it’s a beautiful city.