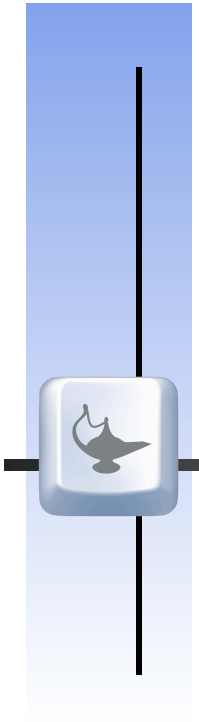

NetReview: Detecting when interdomain routing goes wrong

Andreas Haeberlen
MPI-SWS / Rice

Ioannis Avramopoulos
Deutsche Telekom Laboratories

Jennifer Rexford
Princeton

Peter Druschel
MPI-SWS

Motivation

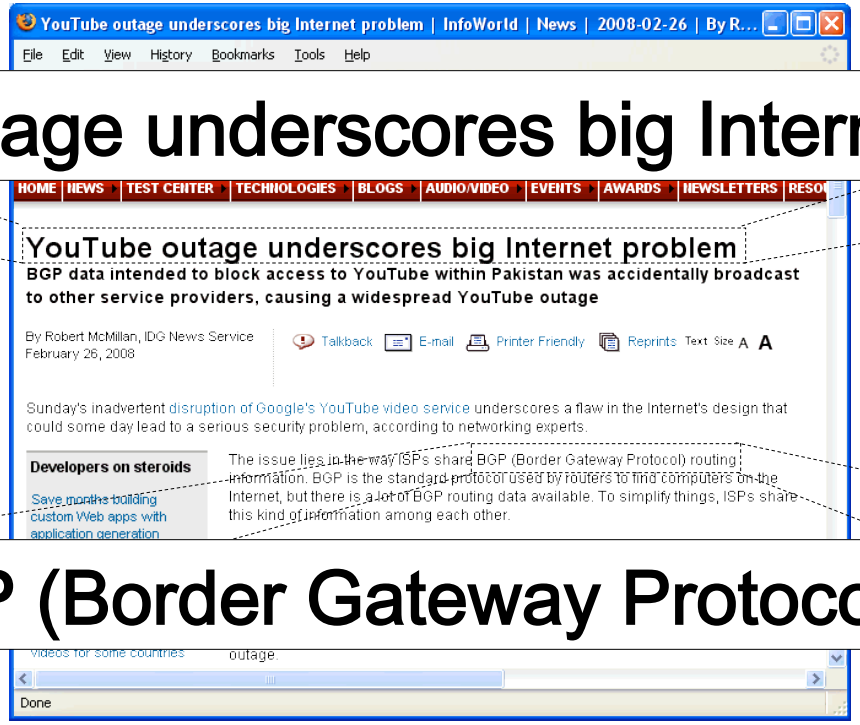

YouTube outage underscores big Internet problem

BGP (Border Gateway Protocol) routing

- This is just the tip of the iceberg
- A considerable fraction of Internet prefixes is affected by routing problems every day

BGP and its problems

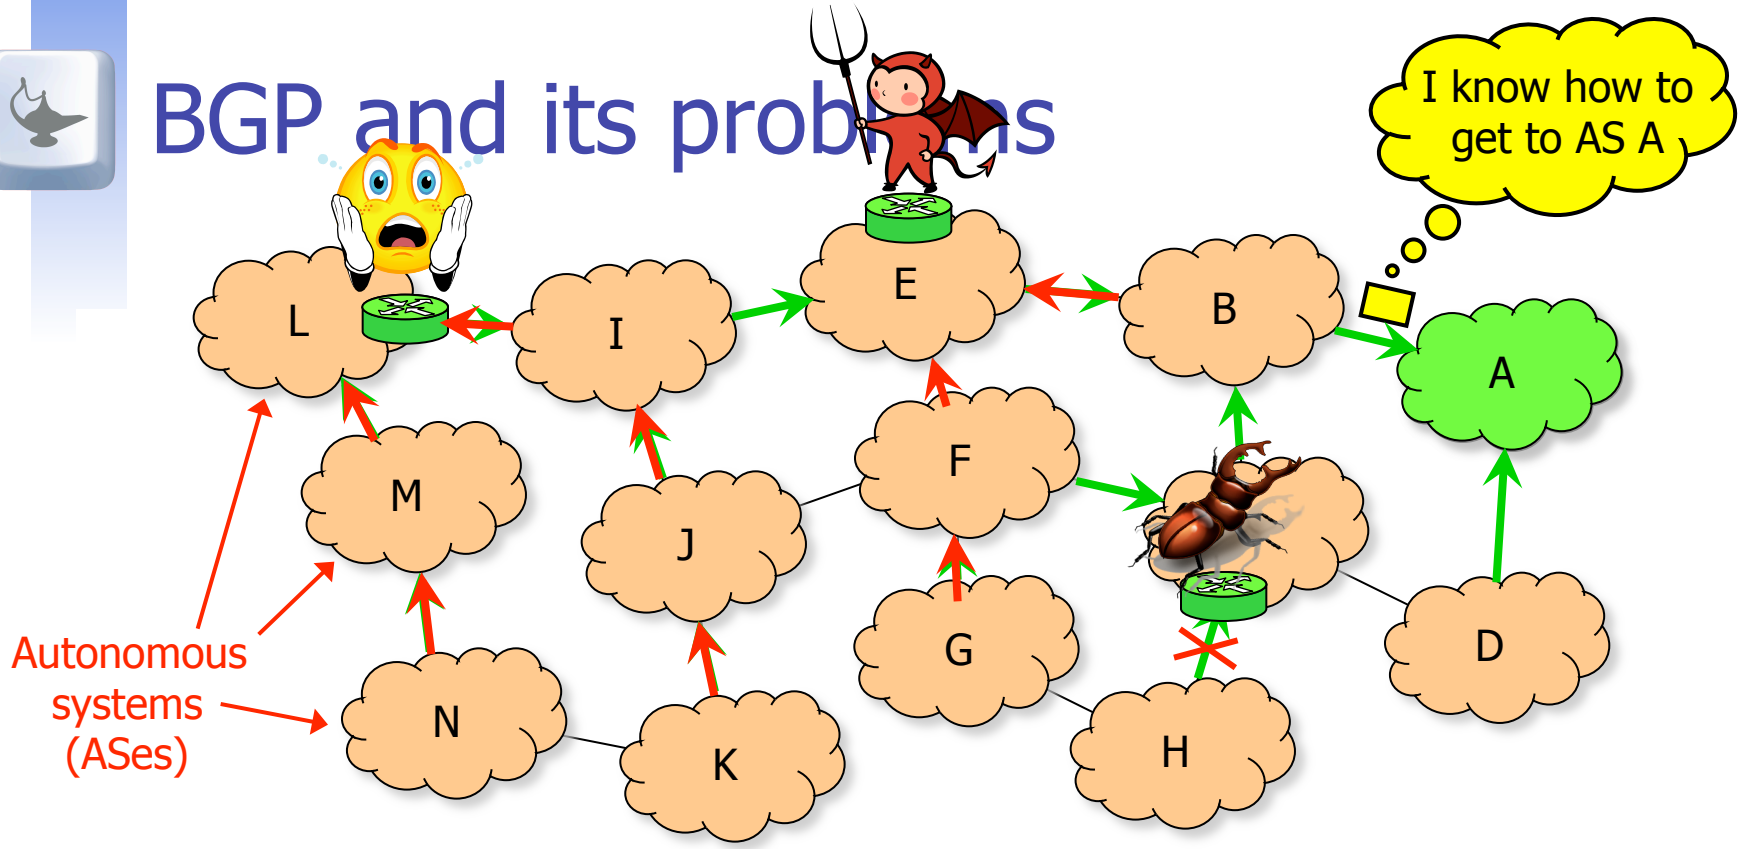

- ASes exchange routing information via BGP
- BGP routing is plagued with many problems:
  - Misconfigurations, bugs, attacks by spammers, instabilities, hijacks, oscillation, equipment failures, policy conflicts, ...

Approach: Fault detection

- Goals:
  1. Reliably detect each routing problem, and
  2. link it to the AS that caused it

- Benefits:
  - ASes can respond to problems quickly
  - No need to diagnose faults manually
  - Works for a very broad class of problems
  - Provides an incentive for reliable routing
  - Easy to deploy incrementally

Challenges in BGP fault detection

- Idea: Upload all router logs to a central entity, who inspects them for problems
  - Sufficient to find almost any routing problem

- Why wouldn't this work in practice?
  - Privacy: Logs contain sensitive information
  - Reliability: Logs may be inaccurate (bugs, hackers)
  - Automation: Can't manually inspect that much data
  - Deployability: Can't assume global deployment
  - Decentralization: ASes wouldn't accept a single detector entity

NetReview from 10,000 feet

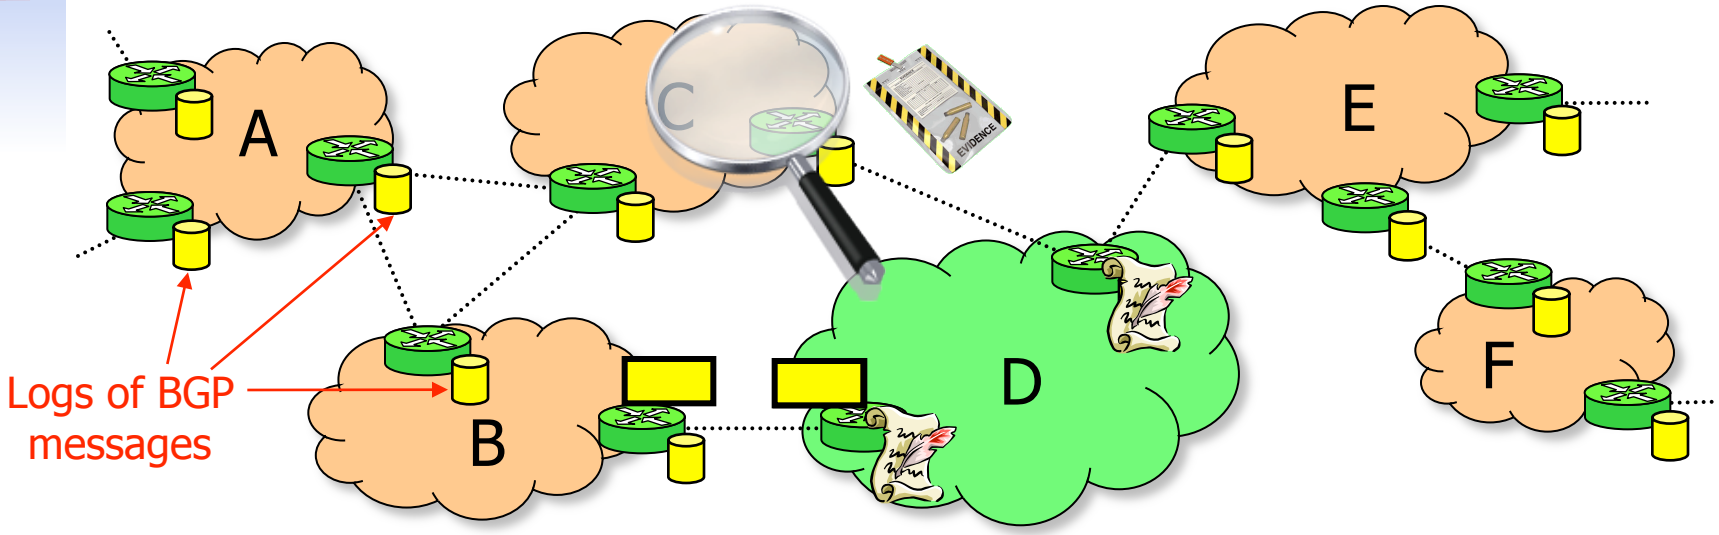

- Border routers maintain logs of all BGP messages
  - Logs are tamper-evident → can reliably detect & obtain proof if faulty routers omit, forge, or modify log entries
- Neighbors periodically audit each other's logs and check them for routing problems
  - If a problem is found, auditor can prove its existence to a third party

Outline

- Introduction
  - Motivation: Internet routing problems
  - Approach: Fault detection
- What is a BGP fault?
- The NetReview system
- Practical challenges
- Evaluation
- Summary

BGP routing policies

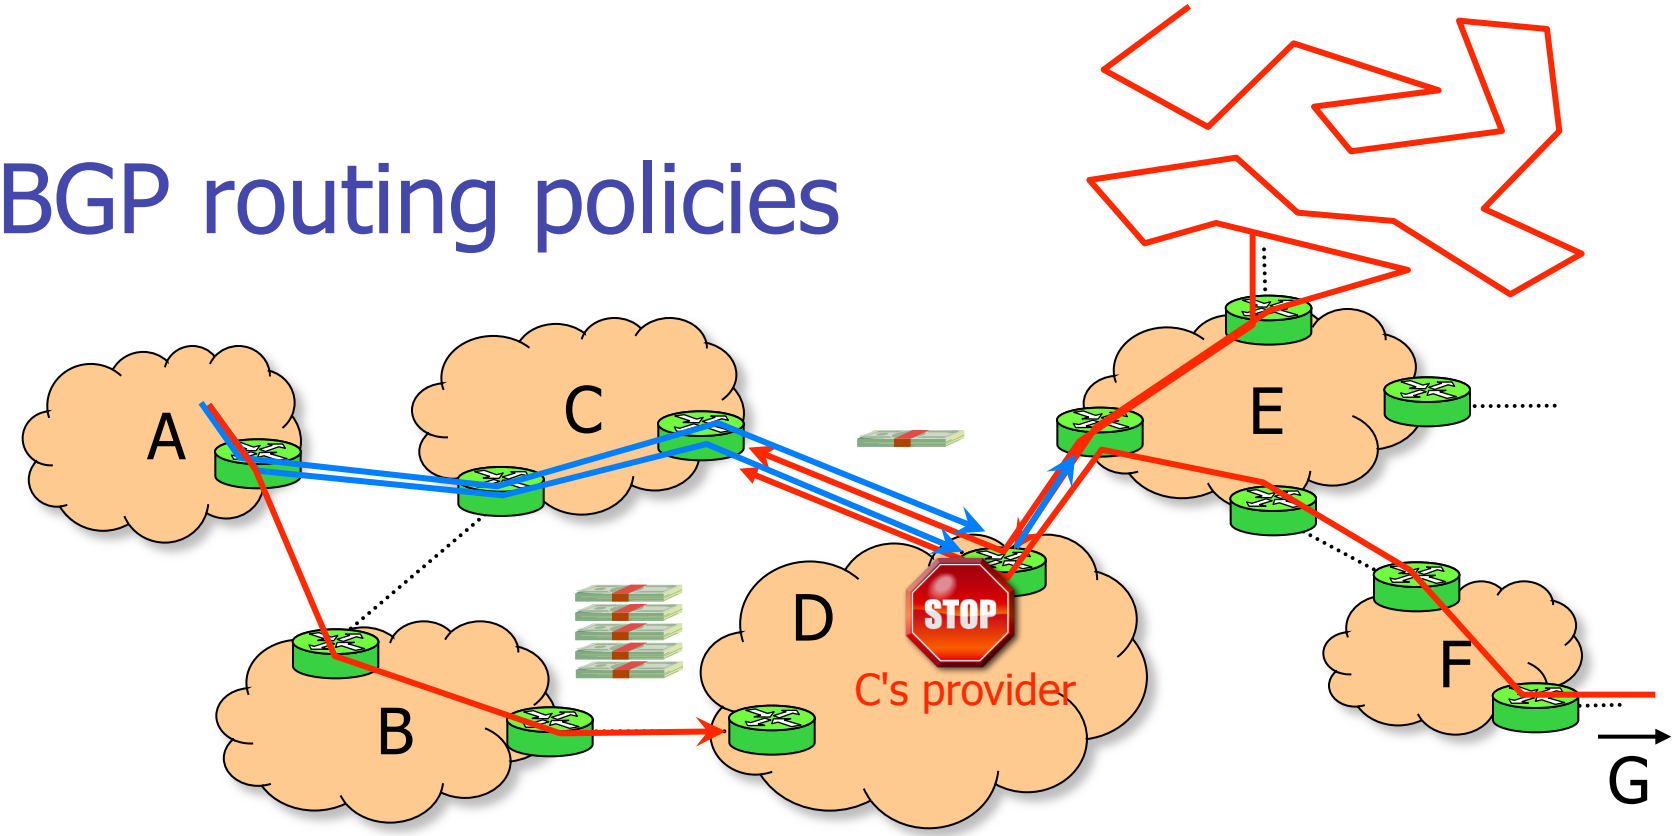

- How do ASes decide what to announce via BGP?
- Each AS has a routing policy, which is based on:
  - Peering agreements: Customer/provider, ...
  - Best practices: Limited path length, ...
  - Internal goals: Choose the shortest/cheapest path, ...
  - Address assignments: IP address prefixes, ...

What is a BGP fault?

- Expected behavior of the AS := Combination of its peering agreements, best practices, internal goals, ...
- BGP fault := The BGP messages sent by the AS do not conform to its expected behavior

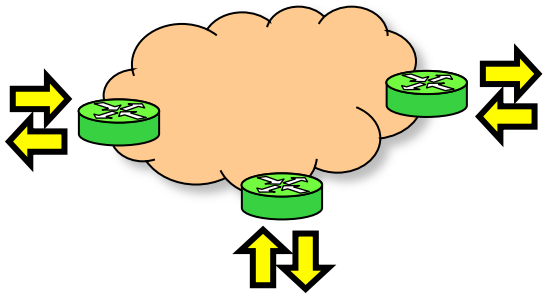

- How do we know what BGP messages the AS sent?
  - Need a complete+accurate message trace even if some routers are faulty in arbitrary, unknown ways
  - Requires a robust+secure tracing mechanism
- How do we know what its expected behavior is?
  - Different for every AS → need a specification

BGP rules

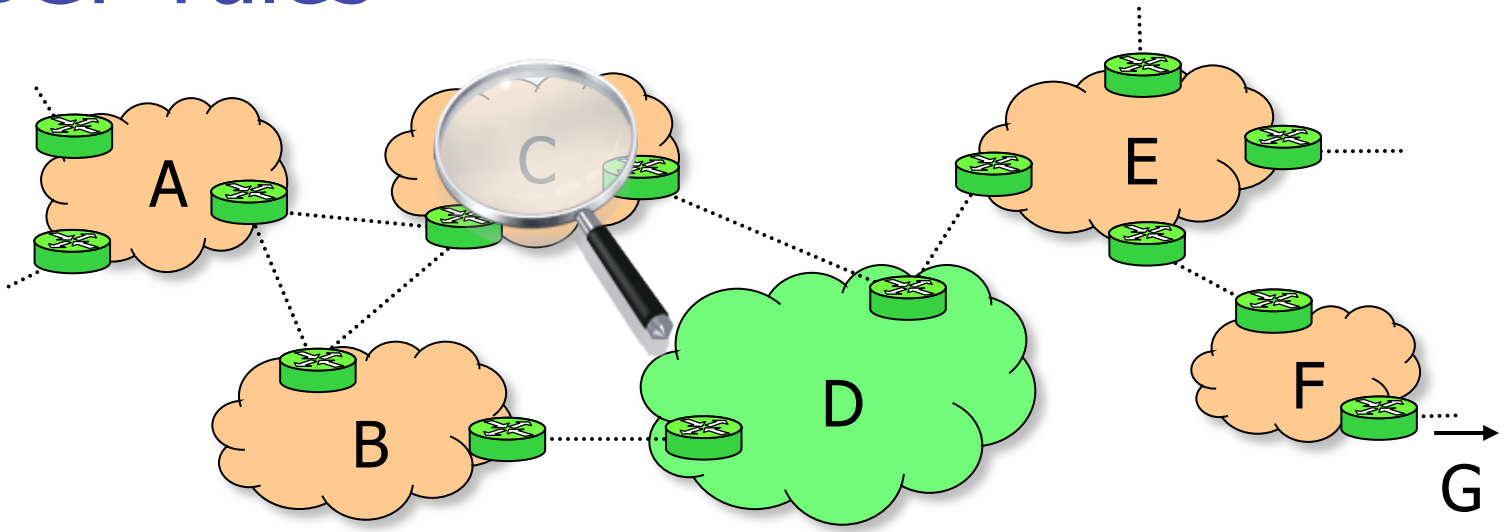

- For example, D might specify the following:
  - "I will filter out routes with excessive paths" (best practice)
  - "I will act as C's provider" (peering agreement)
  - "I will prefer routes through B, if available" (internal)

"Rules"

- Some rules may be confidential, but the AS need not reveal all of them to each auditor

Outline

- Introduction
- What is a BGP fault?
- The NetReview system
- Practical challenges
- Evaluation
- Summary

The tamper-evident log

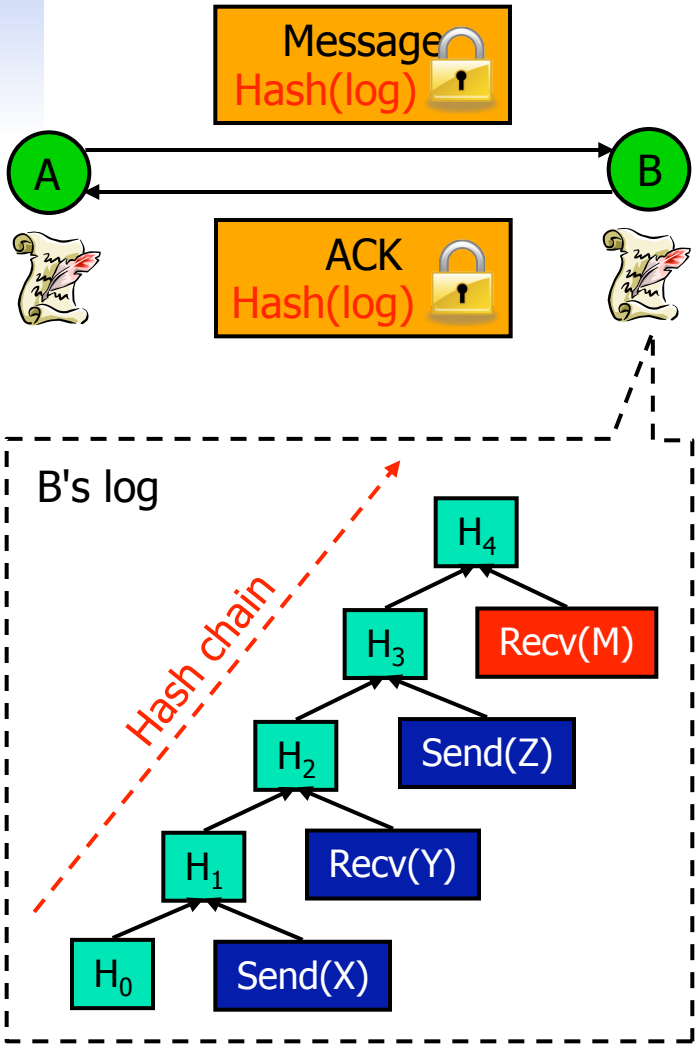

- Based on the tamper-evident log in PeerReview [SOSP'07]
  - If router omits, modifies, or forges entries, neighbors can detect this and obtain evidence
- Log entries form a hash chain
  - Messages include signed hash
  - Tampering breaks the hash chain and is thus detectable
- Messages are acknowledged
  - Detects if message is ignored
- Neighbors gossip about the hash values they've seen

Writing rules

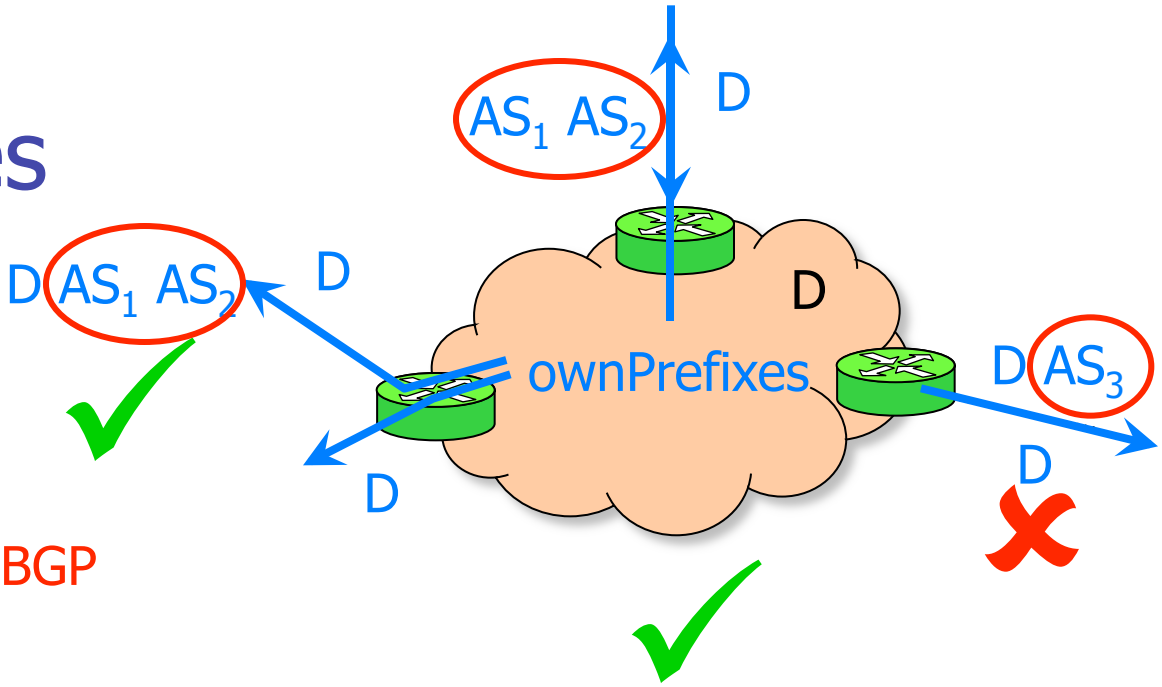

Describes everything that S-BGP can check, and more!

$$\forall a \forall p \forall r \in \text{outRIB}(a, p, \cap[t-40, t]) : (|\text{as\_path}(r)| = 1 \wedge \text{prefix}(r) \in \text{ownPrefixes}) \vee (\exists a' \exists p' \exists r' \in \text{inRIB}(a', p', \cup[t-40, t+5]) : \text{prefix}(r) = \text{prefix}(r') \wedge \text{startsWith}(r, r') \wedge (\forall n \in r - r' : n \in \text{ownPrefixes}))$$

- Rules are predicates on the AS's routing state
  - Declarative; easy to get correct
  - Even simple rules can be very powerful

Auditing and rule evaluation

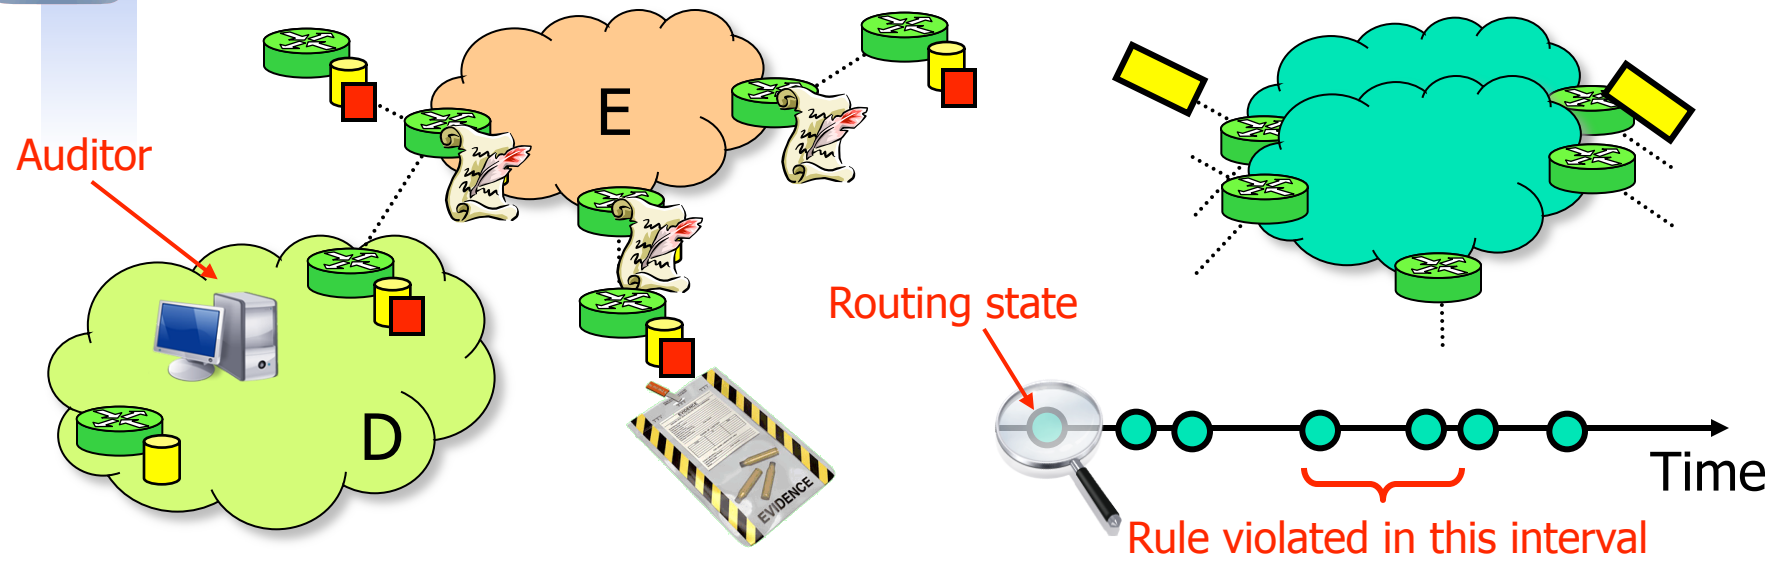

- To audit a neighboring AS:
  1. Auditor requests the logs from each border router
  2. Auditor checks logs for inconsistencies and tampering
  3. Auditor locally replays the logs → series of routing states
  4. Auditor evaluates the rules over each routing state
  5. If a rule is violated during some time interval, auditor extracts verifiable evidence from the logs

Outline

- Introduction
- What is a BGP fault?
- The NetReview system
- Practical challenges
  - Incentives for incremental deployment
  - Partial deployment
  - Working without a certificate authority
  - Using existing routers
- Evaluation
- Summary

Incremental deployment

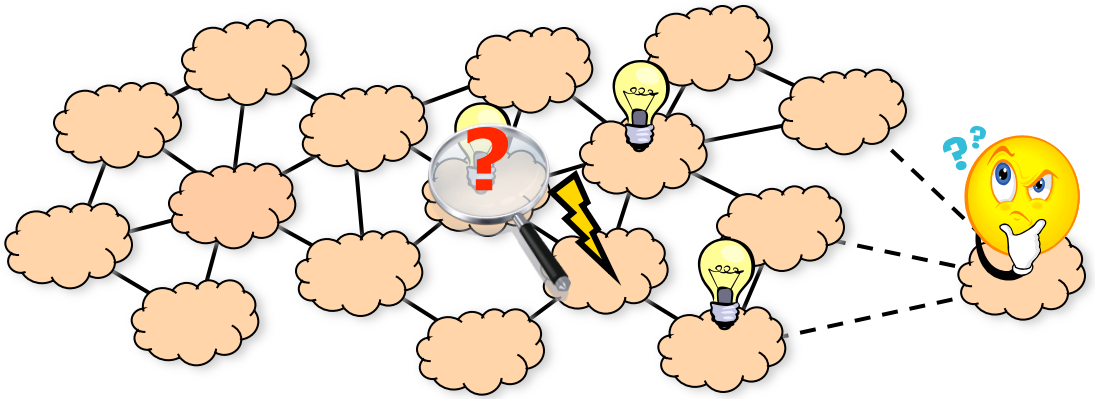

- What is the smallest useful deployment?
  - One AS can find bugs, misconfigurations, ...
  - Two adjacent ASes can check peering agreements, ...
- What are the incentives for deployment?
  - Reliable ASes can attract more customers
  - Logs can be used for root-cause analysis

Outline

- Introduction
- What is a BGP fault?
- The NetReview system
- Practical challenges
- Evaluation
- Summary

Experimental setup

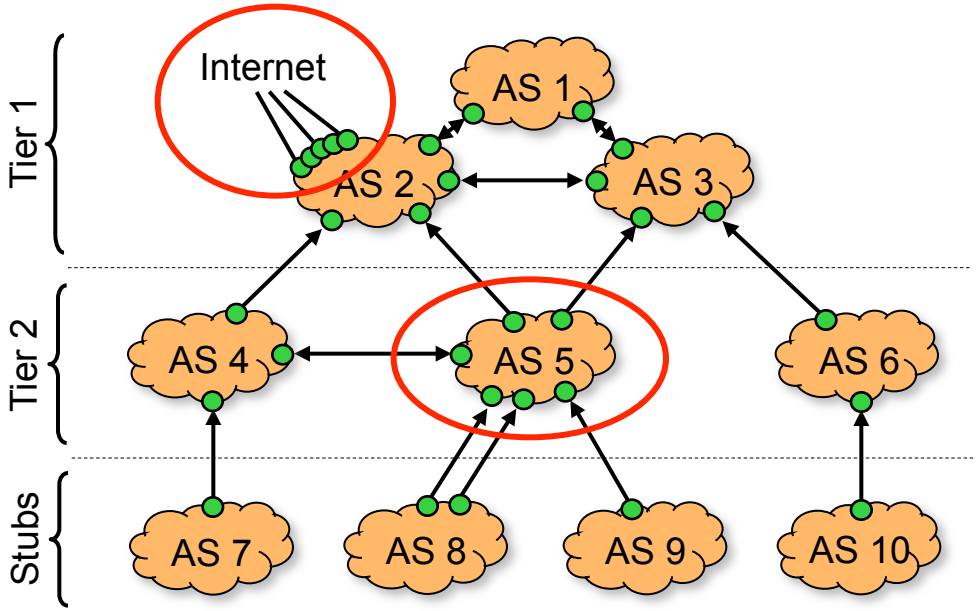

- Synthetic network of 35 Zebra BGP daemons
- Default routing policies (Gao-Rexford)
- Injected real BGP trace (Equinix) to get scale
- Results in this talk are from AS 5 (92% of Internet ASes have degree five or less)

Evaluation: Functionality check

- Fault injection experiment with five rules based on common routing problems:
  - No origin misconfiguration
  - Export customer routes
  - Honor NO_ADVERTISE community
  - Consistent path length
  - Backup link
- NetReview detected all the injected faults
  - Also produced diagnostic information, such as time when the fault occurred, and prefixes that were affected

Evaluation: Overhead

- Processing power: 15-minute log segment can be checked in 41.5s on a P4
  - A single commodity PC is sufficient for small networks
- Storage space: 710kB/minute, ≈356 GB/year
  - Fits comfortably on a single hard disk
- Bandwidth: 420kbps, including BGP updates
  - Insignificant compared to typical traffic volume

Related Work

- Fault prevention
  - Secure routing protocols: S-BGP, soBGP, SPV, ...
  - Trusted monitors: N-BGP
- Heuristic fault detection
  - Anomaly detection
  - Root-cause analysis
- Accountability
  - PeerReview, AIP, AudIt, ...

Summary

- NetReview: A fault detection system for interdomain routing
  - Automatically detects a wide variety of routing problems
  - Links each problem to the responsible AS
  - Not a heuristic - produces proof of each fault
- NetReview is practical
  - Easy to deploy incrementally
  - No PKI required
  - Reasonable overhead

Thank you!