REIN: Reliability as an Interdomain Service

Jia Wang
with
Hao Wang, Yang Richard Yang, Paul H. Liu,
Alexandre Gerber, Albert Greenberg
Yale University
AT&T Labs - Research
Microsoft Research

ACM SIGCOMM 2007

"Any future Internet should attain the highest possible level of availability, so that it can be used for mission-critical activities, and it can serve the nation in times of crisis."

- GENI, 2006

"The 3 elements which carriers are most concerned about when deploying communication services are:

- Network reliability
- Network usability
- Network fault processing capabilities"

-Telemark, 2006

The top 3 all belong to reliability!

Failures in IP Networks

- Part of everyday life of IP networks
  - e.g., 675,000 excavation accidents in 2004 [Common Ground Alliance]
  - Network cable cuts every few days …
- However, major failures can lead to substantial disruption
  - E.g., Jan. 9, 2006, two link failures in a major US ISP led to disconnection of millions of wireless users, partition of many corporate networks

To Handle Failures, We Need

- Network redundancy
  - Redundant resources to make up for the failure
    - Diversity of physical connectivity
    - Over-provision of bandwidth
  - Challenge: significant investments
    - Extra equipment for over-provisioning
    - Expense & difficulty to obtain rights of way for connectivity
- Efficient utilization of network resources
  - IP layer techniques: restoration and protection
  - Challenge: good traffic engineering for reliability

Our Approach: REIN

REliability as an INterdomain Service

- Objective
  - Focuses on intradoman failures
  - Increase the redundancy available to an IP network at low cost
- Basic Idea
  - Observation: IP networks overlap, yet they differ
  - IP networks provide redundancy for each other through interdomain bypass paths
  - Analogy: insurance, airline alliance
  - Effects: Sharing improves reliability and reduces costs

Example: Jan. 9, 2006 of a Major US ISP

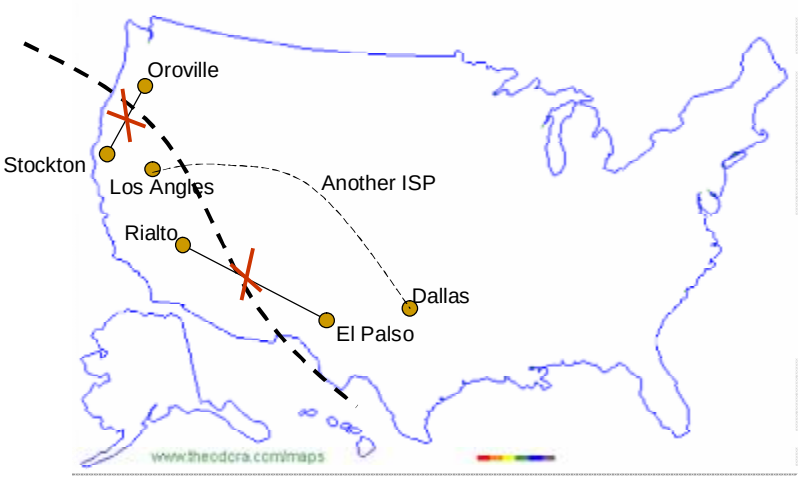

How to Make REIN Work: the Details

1. Why would IP networks share interdomain bypass paths?
2. What is the signaling protocol to share these paths?
3. How can an interdomain bypass path be used in the intradomain forwarding path?
4. After an IP network imports a set of such paths, how does it effectively utilize them in improving reliability?
5. How to minimize the number of such paths?

REIN Business Model: Three Possibilities

- Peering
  - Mutual backup w/o financial settlement
  - Incentive: improve reliability of both at low cost
  - Symmetry in backup paths provisioning & usage
- Cost-free
  - One-sided, volunteer and/or public service
- Customer-Provider
  - Fixed or usage-based pricing
  - Pricing should limit abuse

Interdomain Bypass Path Signaling

- Many possibilities, e.g.,
  - Manual configuration
  - A new protocol
  - Utilize BGP communities

BGP Bypass Path Signaling

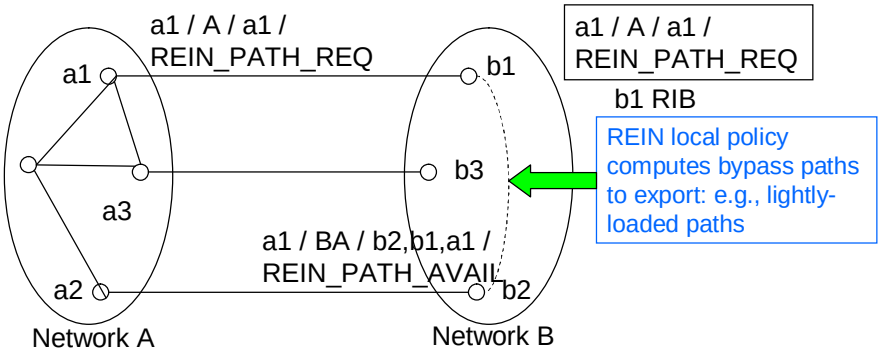

B provides interdomain bypass paths to A.
Task of A: discover a path to a1 through B

a1 / BA / b2,b1,a1 /-
a2 RIB

BGP announcement: Dest. / AS path / Bypass path / Tag
Additional attr.: desired starting point (e.g. a2), bw, etc.

REIN Data Forwarding

- Main capability needed: Allow traffic to leave and re-enter a network
  - Not supported under hierarchical routing of the current Internet because of potential loops
- REIN forwarding mechanism
  - Interdomain GMPLS
  - IP tunneling
  - Either way, only need agreement b/w neighboring networks
    - Incrementally deployable

Traffic Engineering for Reliability (TE-R)

- Objectives
  - Efficient utilization of all redundant resources
  - Scalable and implementable in current Internet
  - Protection: fast ReRouting for high-priority failure scenarios
  - Restoration: routing convergence for other failure scenarios
  - QoS guarantee for important traffic (e.g., VPN), if possible

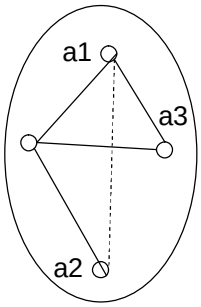

Network topology for TE-R

Intradomain link

REIN virtual link

Our TE-R Algorithm: Features

- Robust normal-case routing f*
  - Based on COPE [Wang et al. '06]
  - Guarantee bandwidth provisioning for hose-model VPN under f*
- Robust fast rerouting under failures on top of f*
  - Important traffic purely intradomain if possible
- Novel coverage-based techniques for computational feasibility and implementability
  - Use flow-based routing to compute optimal solution
  - Coverage to generate implementation with performance guarantee
- For details, please see paper.

Further Optimization: Minimize Interdomain Bypass Paths

- Motivation
  - REIN may provide many alternatives
  - Only a few may be necessary
    - Reduce configuration overhead & budget constraints
- Step 1: Connectivity objective
  - Preset connectivity requirement
  - Cost assoc. w/ interdomain paths
  - Meet connectivity requirement + minimizing total cost
  - Formulated as a Mixed Integer Programming (MIP)
- Step 2: TE-R objective
  - Sort interdomain paths according to a scoring function
  - Greedy selection until TE-R has desired performance

Evaluation Methodology

- Dataset
  - US-ISP
    - Hourly PoP-level TMs for a tier-1 ISP (1 month in 2007)
  - Abilene
    - 5-min router-level TMs on Abilene (6 months: Mar – Sep. 2004)
  - RocketFuel PoP-level topologies
- TE algorithms
  - TE-R (*robust*)
  - Oblivious routing/bypassing (*oblivious*)
  - COPE + Constrained Shortest Path First rerouting (*CSPF*)
  - Flow-based optimal routing (*optimal*)

Why Need a TE-R (Abilene 1-link failure)

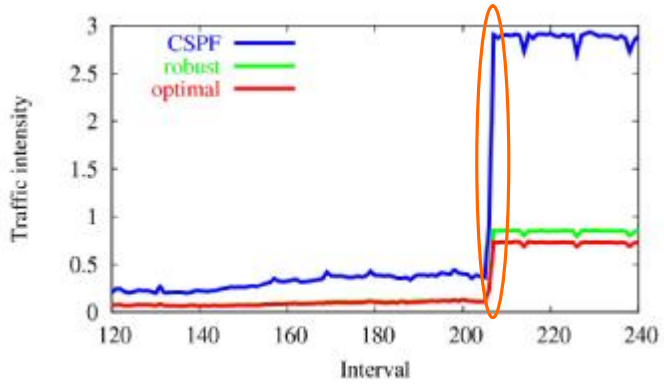

Abilene bottleneck link traffic intensity: 1-link failures, Tuesday August 31, 2004

CSPF overloads bottleneck link by ~300%
vs.
robust TE-R successfully reroutes all traffic

Why REIN: Connectivity Improvements

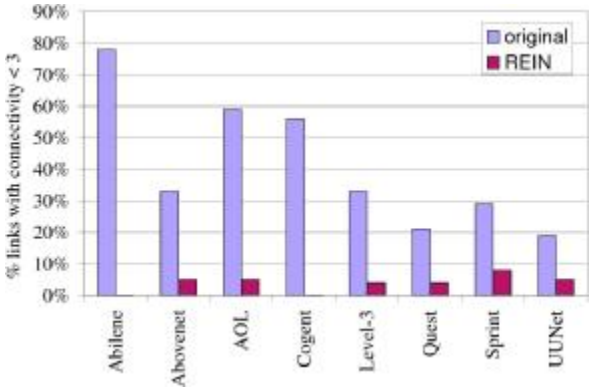

- Actual topology for Abilene, RocketFuel inferred for all others and may underestimate connectivity
- Links with conn. < 3 ==> possible partition under 2 fiber cuts
- As high as 60% of links w/ conn. < 3 in some smaller networks
- A few (<= 7) backup routes from neighboring networks help a lot

Why REIN: Overload Prevention (Abilene 2-link)

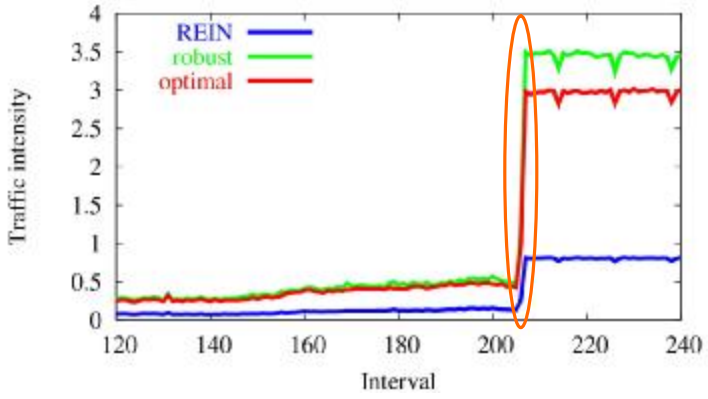

Abilene bottleneck link traffic intensity: 2-link failures, Tuesday, August 31, 2004

Without REIN, even optimal routing overload bottleneck links by ~300%.
With 10 interdomain bypass path of 2Gbps each, REIN reduces MLU to ~80%

Why REIN: Overload Prevention (US-ISP failure log)

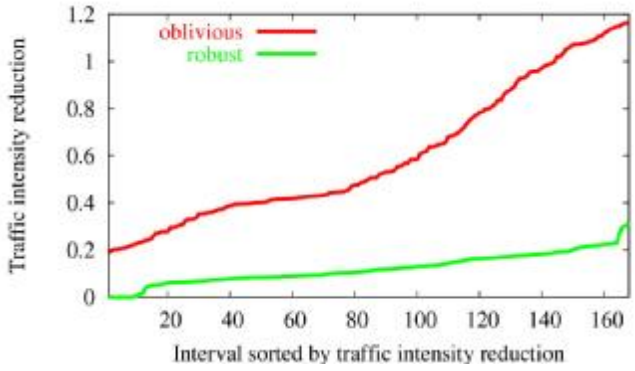

Improvement of traffic intensity by REIN for a week in January 2007 for US-ISP

REIN can reduce normalized traffic intensity by 118% and 35%, depending on the TE algorithms used.

Conclusions & Future Work

- REIN
  - An interdomain service to improve the redundancy of IP networks at low cost
  - Significantly improves network reliability, esp. when used with our TE-R to utilize network resources under failures
- Ongoing & future work
  - A thorough study of the effects of cross-provider shared-risk link group data
  - Further Improve TE-R performance

Thank you!