On the Impact of Filters on Analyzing Prefix Reachability in the Internet

Ravish Khosla, Sonia Fahmy, Y. Charlie Hu
Purdue University
ICCCN 2009

Background

- Border Gateway Protocol (BGP)
  - Inter-domain policy based Routing Protocol
  - Advertises IP prefixes belonging to Autonomous Systems (ASes)

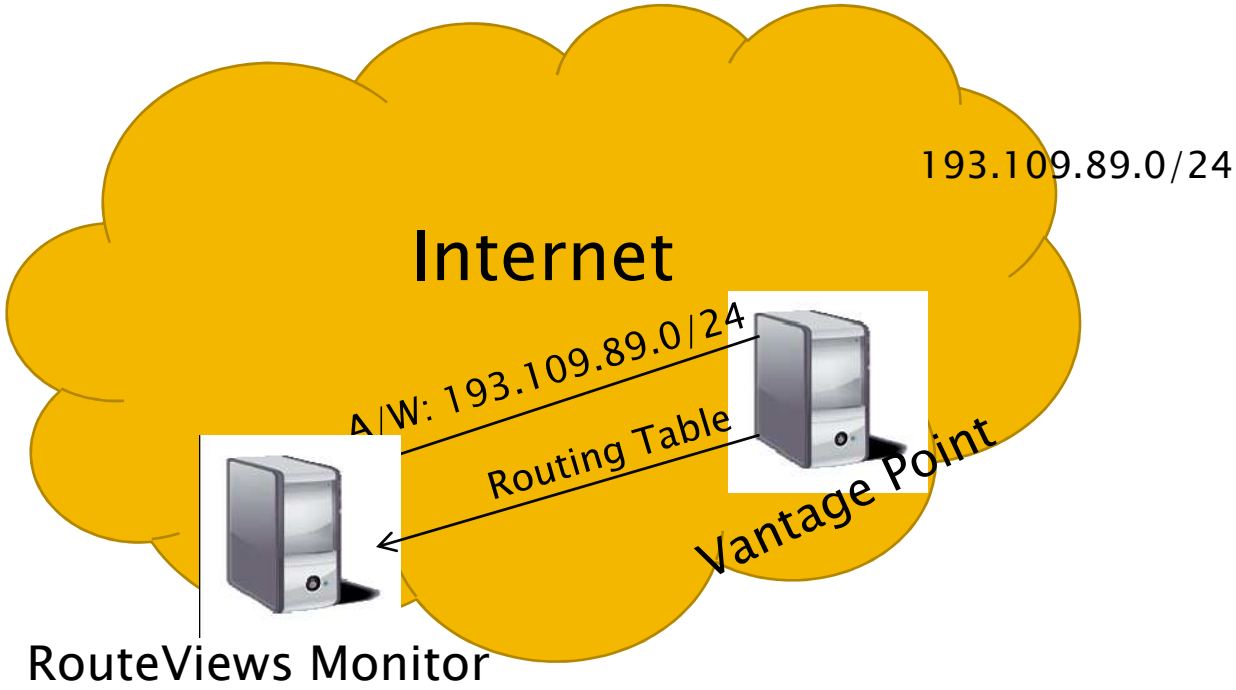

Goal

- Study prefix reachability
  - Existence of announced paths from vantage points
  - Impacts prefix availability
- For those paths:
  - Mean Time to Failure (MTTF)
  - Mean Time to Recovery (MTTR)
- Dataset: RouteViews
  - Routing tables sampled every 2 hours
  - Updates collected in 15 minute durations
    - Announcements
    - Withdrawals

Contributions

- How to process data to compute average prefix reachability in the Internet?
  - Which *prefixes* to consider?
  - Do we have *sufficient* data for these prefixes?
  - Which *updates* to consider?
  - *Filter:* Remove biased data points
- Tuning filter strength (parameters)
  - Impact of filter strength on reachability results
- Internet reachability results
  - Normal operation
  - With stress events
    - Undersea cable cut in the Middle East (2008)

Preprocessing Datasets

- Removing routing table transfers
  - Updates include those caused by peering session resets between a peer and monitor
- Used *Minimum Collection Time (MCT)* algorithm[1] to identify transfers
- Peer: Any vantage point that exists in one routing table entry and at least one update
  - 45–47 peers in our dataset
- Execute scripts w.r.t. every peer
  - Filter updates in identified transfer

[1] B. Zhang, V. Kambhampati, M. Lad, D. Massey, and L. Zhang, "Identifying BGP routing table transfers," *ACM SIGCOMM MineNet workshop, 2005*

Computing MTTF and MTTR

- Unit of *combination*: (peer, prefix) tuple
- Combination state: Up (U)/Down (D):
  - Currently advertised path exists/doesn't exist to the prefix by the peer

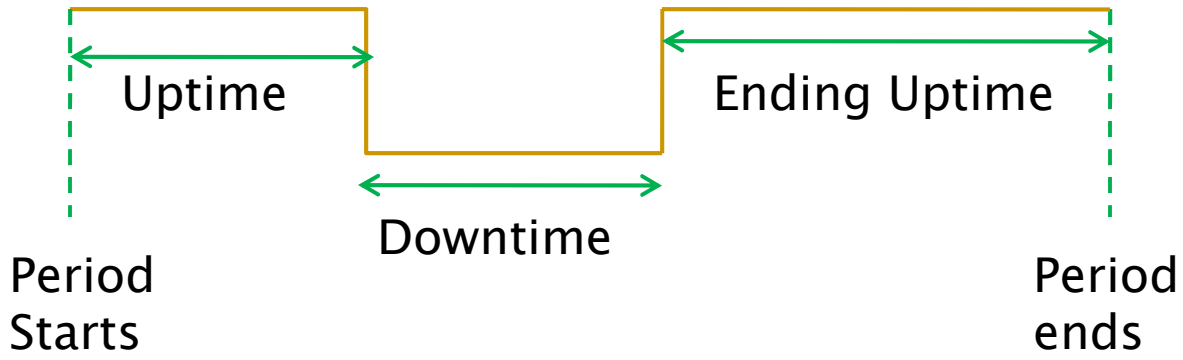

- Boundary effects
  - Initial state and ending state

Filter Design

- Observed Combinations filter
  - Eliminates boundary effects
  - Look at combinations in the first routing file or first α % of update files
  - Results with Observed Combinations filter called *"Initial Filtering"* results
- Stable prefixes filter
  - Prefixes existing in more than β % of routing tables
- Route Convergence filter
  - Count multiple failures within γ seconds as a single failure
  - Remove updates while routes are still converging

Effect of filters on MTTF,MTTR

Filter	Effect on MTTF	Effect on MTTR
Observed Combinations	$\alpha \uparrow \Rightarrow$ MTTF $\downarrow$	$\alpha \uparrow \Rightarrow$ MTTR $\uparrow$
Stable Prefixes	$\beta \uparrow \Rightarrow$ MTTF $\uparrow$	$\beta \uparrow \Rightarrow$ MTTR $\downarrow$
Route Convergence	$\gamma \uparrow \Rightarrow$ MTTF $\uparrow$	$\gamma \uparrow \Rightarrow$ MTTR $\uparrow$

Choosing Filter Parameters

- Observed Combinations Filter α
  - Enough data for the combinations
    - Eliminates combinations first advertised at end of dataset
  - Try α = 10,25 & 50 %
  - Studied 20 cases of α with/without the other two filters and used all typical values of β & γ
    - Which α gives largest increase in (MTTF-MTTR) w.r.t. Initial Filtering case on average?
    - MTTF-MTTR monotonic in α, so choose α=10%
    - Retain 91% (9 million) of combinations for Jan. 07 data

Filter Parameters (Contd.)

- Stable Prefixes Parameter β
  - Implemented observed combinations filter with α = 10% with this filter and different β
  - β = 0% different than initial filtering
    - Transient "hidden" prefixes
  - β = 100% typically doesn't yield any prefixes
  - Greatest increase in MTTF-MTTR vs. Initial filtering was for high β
  - ∴Choose β as high as possible

Filter Parameters (Contd.)

- # of prefixes visible vs β:

Jan. 07

β (%)	# of output prefixes	% of total prefixes
0	233,537	100
30	225,685	96.64
60	221,620	94.89
90	217,607	93.18
98	212,883	91.16

Jan. 05

β (%)	# of output prefixes	% of total prefixes
0	180,229	100
30	171,289	95.04
60	166,447	92.35
90	162,619	90.23
98	145,633	80.8

Filter Parameters (Contd.)

- Route convergence filter parameter γ
  - Both MTTF and MTTR increase on increasing γ
  - MTTF-MTTR ↓ as γ ↑
  - Choose smallest γ indicative of route convergence
  - Typical values: 200-300 sec [1,2]

[1] S. Burkle, "BGP convergence analysis," Ph.D. dissertation, 2003.
[2] C. Labovitz, A. Ahuja, A. Abose, and F. Jahanian, "The Problem with BGP," http://www.nanog.org/mtg-0002/converge2.html, 2000

Results for a 9 month period

- Mar.–Nov. 07
- Longer duration
  - # of combinations exceed combinations in one month
  - % prefixes available reduction more with β

β (%)	# of output prefixes	% of total prefixes
0	341,122	100
60	227,291	66.63
99.5	180,489	52.91

Mar. – Nov. 07 results

- Initial Filtering results
  - MTTF : 9.5 days
  - MTTR : 1.5 days
  - Median uptime : 16 minutes
  - Median downtime : 1 minute
- Stable prefixes filter –53% of the more stable prefixes have
  - MTTF : 2 weeks
  - MTTR :10 hours
  - Median uptime : 1.4 hours
  - Median downtime : 53 seconds

Mar. – Nov. 07 results (Contd.)

- Stable prefixes + route convergence filter
  - MTTF goes up with β and γ
    - Min MTTF is 25.5 days
    - Max MTTF is 60 days
  - MTTR ↑ than I.F. for γ > 300 s except β=0%
    - Min MTTR: 27 hours (–25%)
      - β = 0 % ; γ = 200 s
    - Max MTTR: 6.6 days (340%)
      - β = 0 % ; γ = 900 s
  - Highest MTTF–MTTR
    - β = 99.5 % ; γ = 900 s
    - MTTF is 60 days ⇒ 540% ↑ than Initial Filtering (I.F.)
    - MTTR is 2.1 days

Mar. – Nov. 07 results (Contd.)

- Comparison with Labovitz's results in [1]
  - MTTF :12 days , MTTR :15 minutes
  - β=99.5 % ; γ=900 s ⇒ MTTF 410% ↑, MTTR 19,300% ↑
  - MTTF ↑ ⇒ Internet is getting healthier
  - Backbone paths studied in [1] explains MTTR ↑
  - Our median downtime is 17 minutes to 1 hour

[1] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental Study of Internet Stability and Backbone Failures," FTCS, 1999

Effect on availability

- MTTF ↑ & MTTR ↑
  - How does it impact availability?

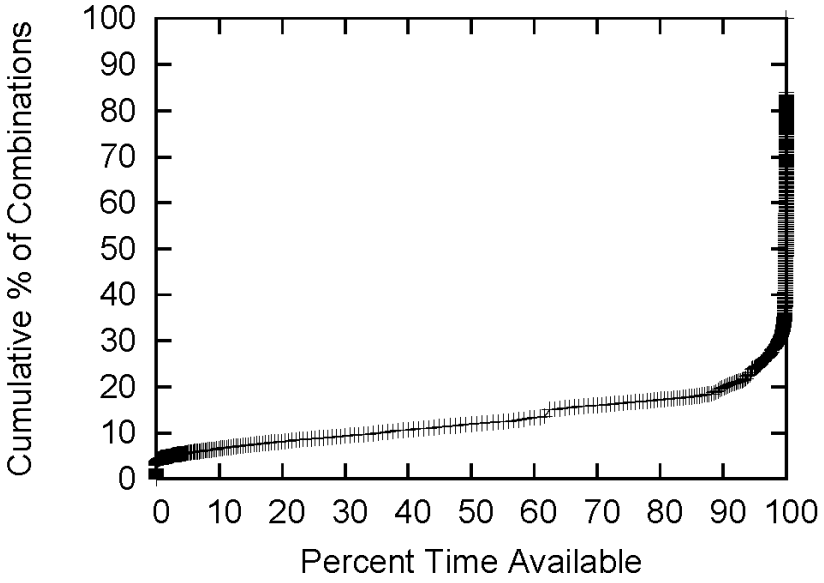

Initial Filtering case

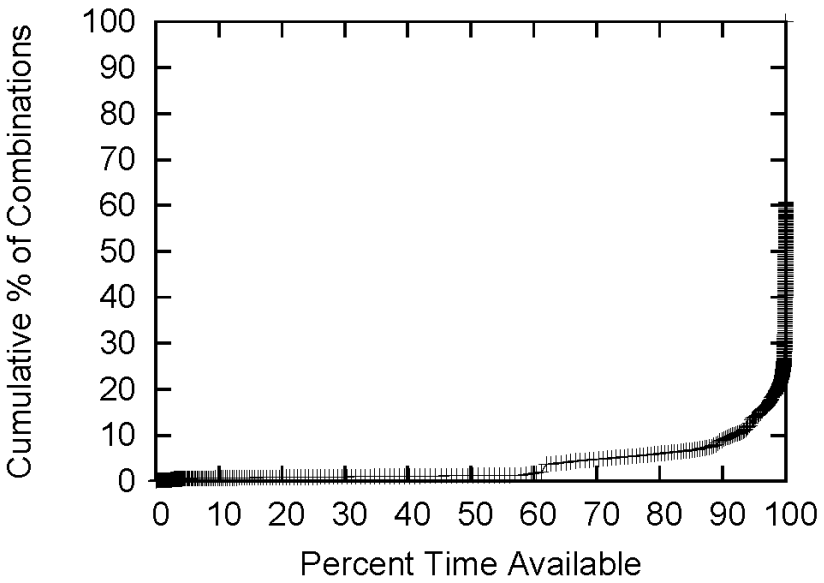

$\beta = 60\%$; $\gamma = 900$ sec

Effect on availability (Contd.)

- Availability histograms

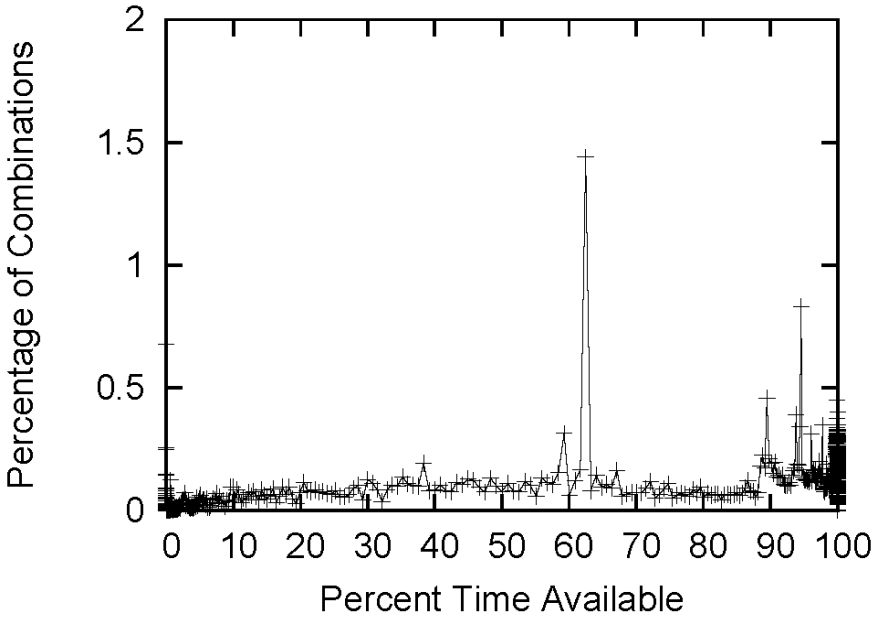

Initial Filtering case

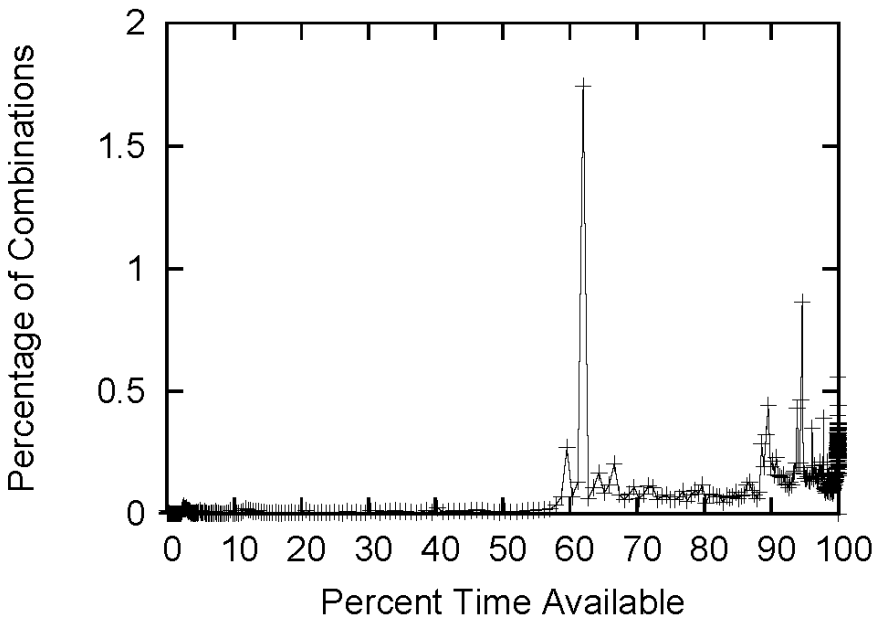

β = 60% ; γ = 900 sec

- Spike around 60% : peer 206.24.210.99
  - AS 3561 (SAVVIS Corporation)

Impact of cable cuts

- Jan. and Feb. 2008 data
- 1 week sliding window
  - 60 days & 54 windows

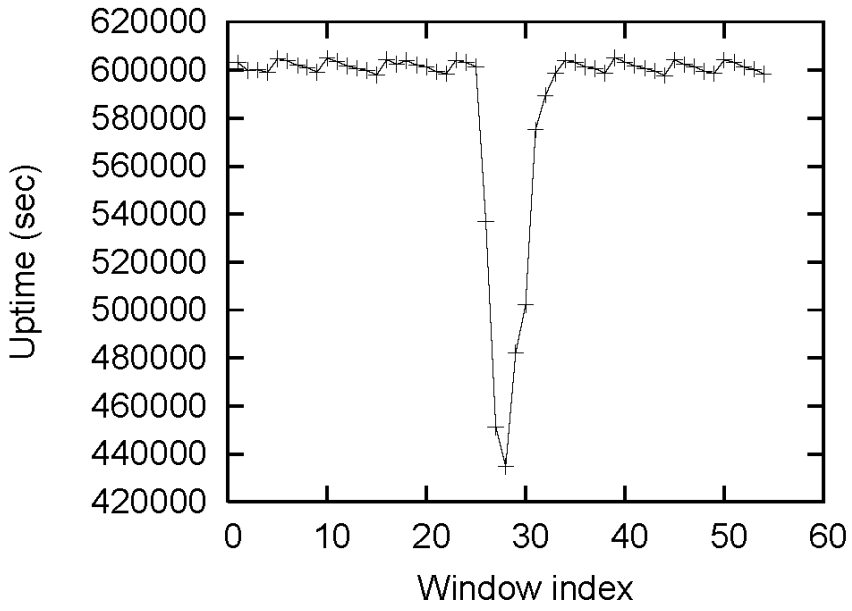

Initial Filtering results

- Lowest uptime : 28th window and is about 28% lower
- Cable cut happened around 26th window

Filtering for cable cuts data

- Used stable prefixes and route convergence filters
  - β=60%, γ=900 s (same as in Labovitz's work)

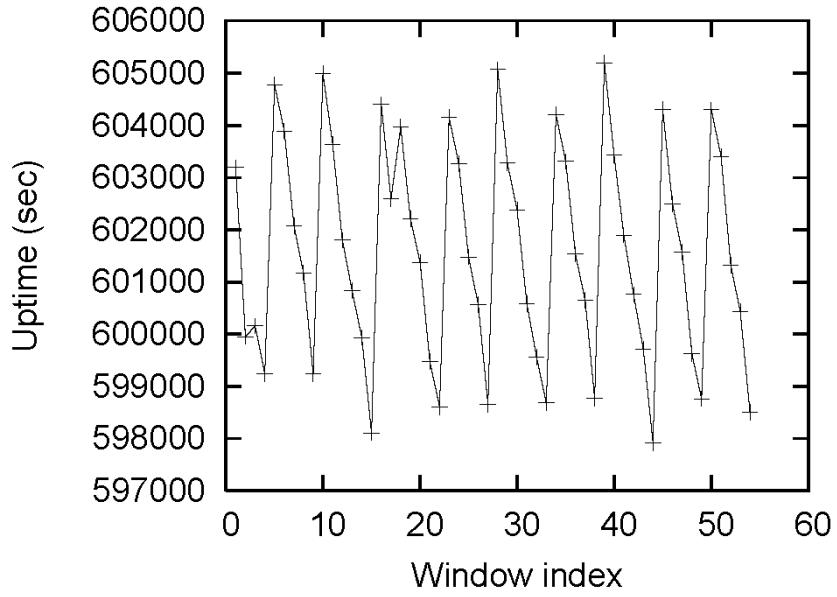

➢Uptime drop filtered out

Implications

- Filter parameters have to be chosen according to the goals of the study
- Goal: Observe stress events
  - Use Initial filtering case
- Goal: Study general health of the Internet
  - Use filters with carefully chosen parameters
  - Filter parameters can be chosen to make Internet look healthy or unhealthy
    - Results can vary by orders of magnitude

Filter parameters revisited

Filter	Effect on MTTF	Effect on MTTR	Values in [1]	"Healthy" values
Observed Combinations	$\alpha \uparrow \Rightarrow$ MTTF $\downarrow$	$\alpha \uparrow \Rightarrow$ MTTR $\uparrow$	Not used	10%
Stable Prefixes	$\beta \uparrow \Rightarrow$ MTTF $\uparrow$	$\beta \uparrow \Rightarrow$ MTTR $\downarrow$	60 %	~99.5%
Route Convergence	$\gamma \uparrow \Rightarrow$ MTTF $\uparrow$	$\gamma \uparrow \Rightarrow$ MTTR $\uparrow$	900 s	200s

[1] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental Study of Internet Stability and Backbone Failures," FTCS, 1999

Conclusions

- Filters offer significant power in eliminating pathological updates and unstable prefixes
  - However, they should be used judiciously
- Internet is now “healthier”
- 53% of the most “healthy” prefixes
  - MTTF : two weeks
  - MTTR :10 hours
  - Median uptime : 1.4 hours
  - Median downtime : 53 seconds
- Future work: Tying prefix reachability to Internet resilience

Questions

Backup Slide 1: Data Sets

- RouteViews[1]
  - Routing tables sampled every 2 hours
  - Updates collected in 15 minute durations
  - About 25 GB per month of zipped data
- Months studied
  - Mar.–Nov. 07: 9 months to match Labovitz's work
  - Jan and Feb. 08 for cable cuts
- Preprocessing
  - Convert to text format
  - Remove unused fields
    - Keep timestamp, peer, prefix, update type
  - 13–15 GB of gzipped, processed data/month

[1] University of Oregon Route Views Project, www.routeviews.org

Backup Slide 2: Future work

- Study prefix lifetime distribution and routing update arrivals for better filter parameter selection
- Studying prefix aggregation
- Tying prefix reachability to Internet resilience