

Internet Routing Instability

Craig Labovitz · G. Robert Malan · Farnam Jahanian

University of Michigan

Department of Electrical Engineering and Computer Science

`{labovit, rmalan, farnam}@eecs.umich.edu`

Outline

- ◆ **Routing Instability**
- ◆ **Methodology**
- ◆ **Pathological Routing Information**
- ◆ **Analysis of Instability**
- ◆ **Conclusion**

Routing Instability(1/2)

◆ 『Route flaps』

-- rapid change of network reachability and topology information

◆ Primary effect

- (1) increased packet loss
- (2) delays in the time for network convergence
- (3) additional resource overhead (memory, CPU, etc.)

Routing Instability(2/2)

◆ Route Caching Architecture

Cache miss result : (1) increased load on the CPU
(2) increased switching latency
(3) loss of packets

◆ **Severe CPU load** : delay routing Keep-Alive packets

◆ **route flap storm**

Solution

- ◆ **ASPATH attribute** : detect and prevent forwarding loops
- ◆ **Route dampening algorithms** : not a panacea
- ◆ **Increased use of route aggregation** :
Combines a number of smaller IP prefixes into a single ,
less specific route announcement ◦

Methodology

- ◆ based on data collected at Internet routing exchange points
- ◆ logged BGP routing messages exchanged
- ◆ concentrating on the logs of the largest exchange :
"Mae-East"
- ◆ analyze the BGP data
- ◆ analyze sample BGP backbone logs

Map of major U.S. Internet exchange points

AADS 、 Mae-East 、 Mae-West 、 PacBell 、 Sprint



Classifications of routing information

- ◆ Pathologic (or redundant) updates
- ◆ Policy fluctuation
- ◆ Forwarding Instability

BGP updates events

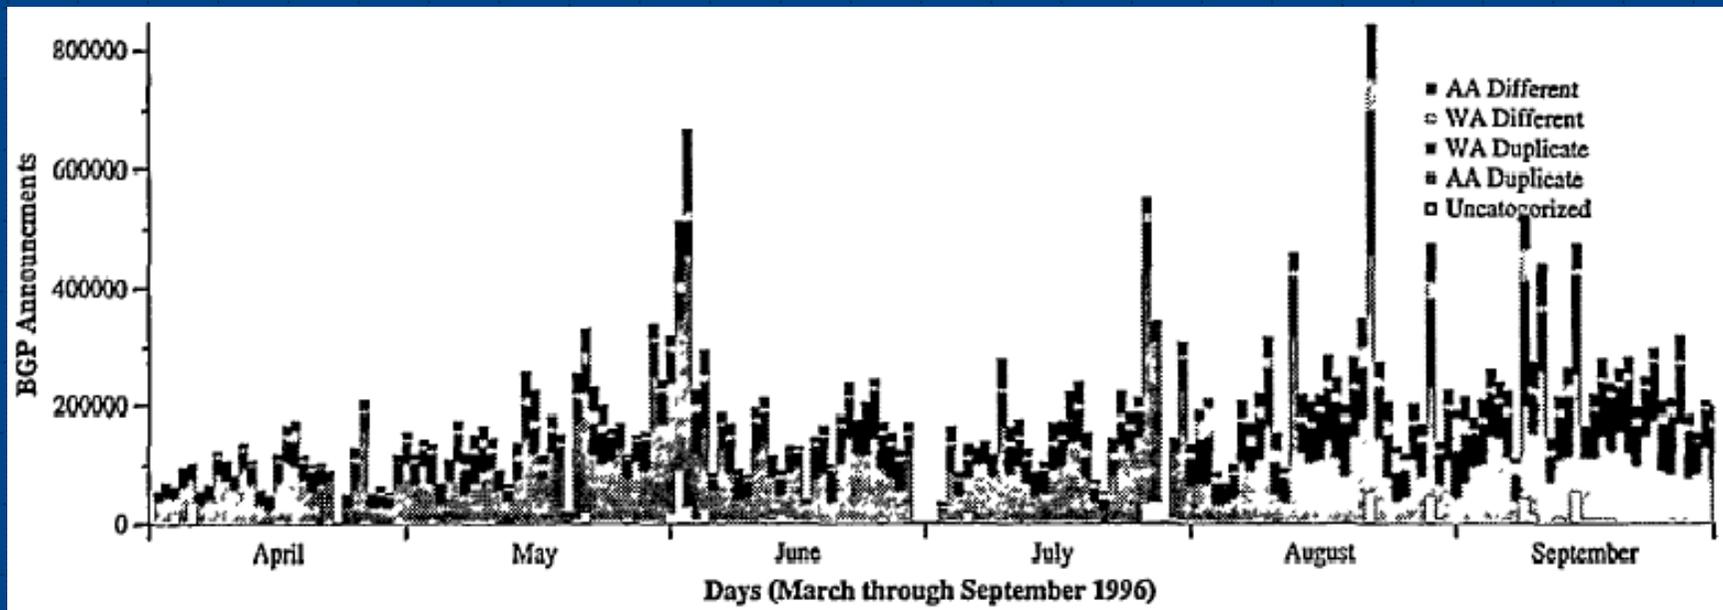
- ◆ **WADiff** : Forwarding instability
- ◆ **AADiff** : Forwarding instability
- ◆ **WADup** : oscillating reachability announcements
- ◆ **AADup** : duplicate path announcements
- ◆ **WWDup** : duplicate withdrawal announcements

Classless inter-domain routing

- ◆ 『Supernet』
- ◆ multi-homing
- ◆ Internet address space are not well-aggregated

Breakdown routing updates(1/2)

- ◆ AADup and WADup classifications consistently dominate other categories of routing instability
- ◆ AADup category is dominated by policy changes that do not directly affect forwarding instability



Breakdown routing updates(2/2)

- o The majority of BGP updates consist entirely of WWDup
- o All pathological routing incidents were caused by small service providers

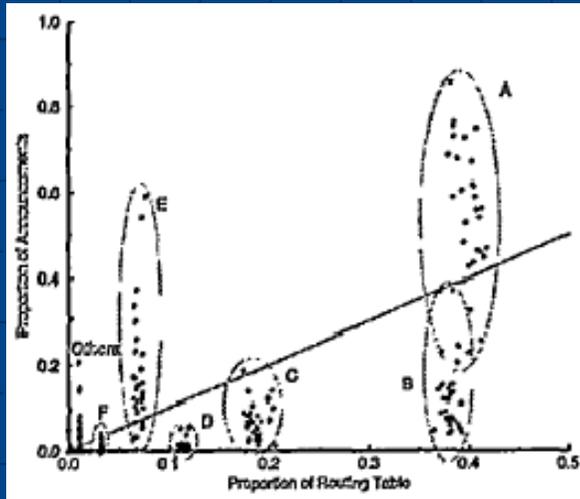
Network	Announce	Withdraw	Unique
Provider A	1127	23276	4344
Provider B	0	36776	8424
Provider C	32	10	12
Provider D	63	171	28
Provider E	1350	1351	6
Provider F	11	86417	12435
Provider G	2	61780	10659
Provider H	21107	77931	14030
Provider I	250	2479023	14112
Provider J	2335	1363	853

Possible Origins of Routing Pathologies

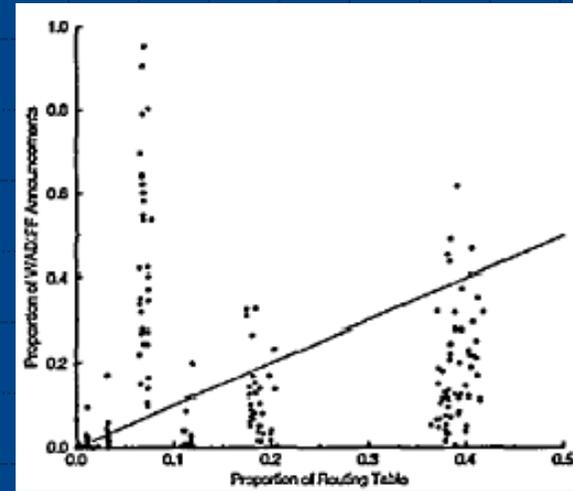
- ◆ **stateless BGP (IETF BGP standard)**
 - contribute an additional $O(N * U)$ updates
 - Router vendors : Partial state on BGP advertisements
 - limits distribution of WWDup updates
- ◆ **Channel Service Units (CSU) timer problems**
- ◆ **Self synchronization**

Fine-grained Instability Statistics

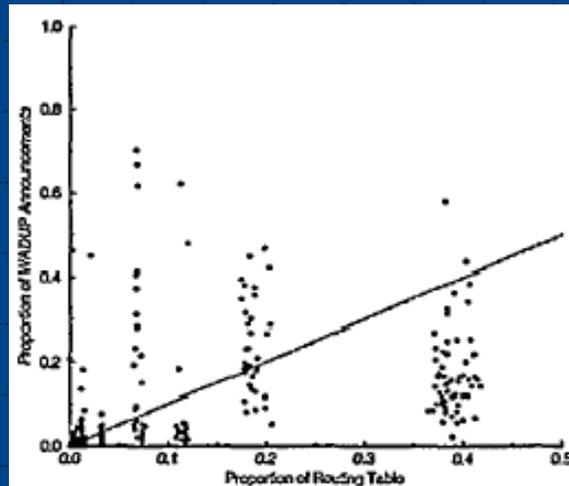
- ◆ No single autonomous system consistently dominates the instability statistics.
- ◆ There is not a correlation between the size of an AS (measured at the public exchange point as the number of routes which it announces to non-customer and non-transit peers) and its proportion of the instability statistics.
- ◆ A small set of paths or prefixes do not dominate the instability statistics; instability is evenly distributed across routes.



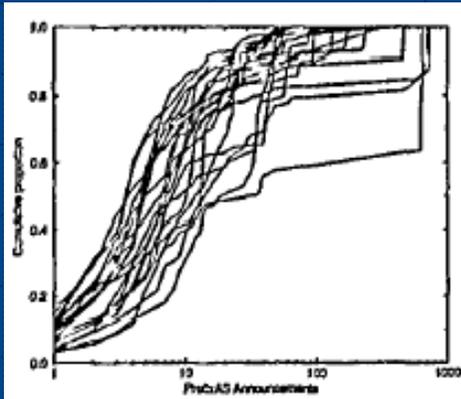
AADiff



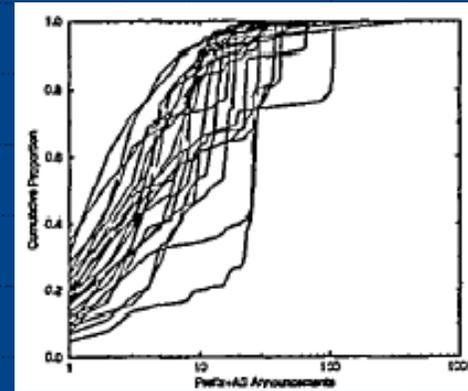
WADiff



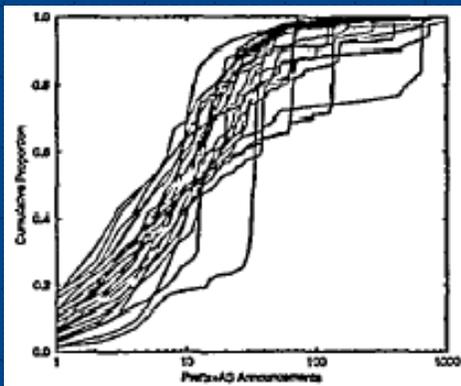
WADup



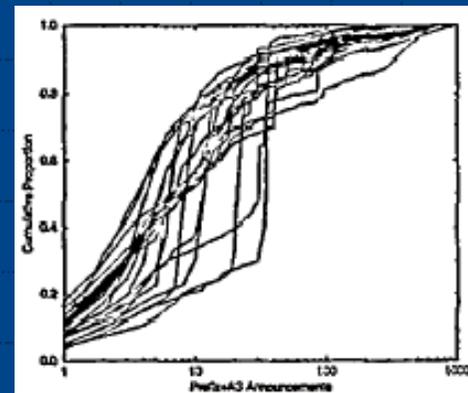
AADiff



WADiff

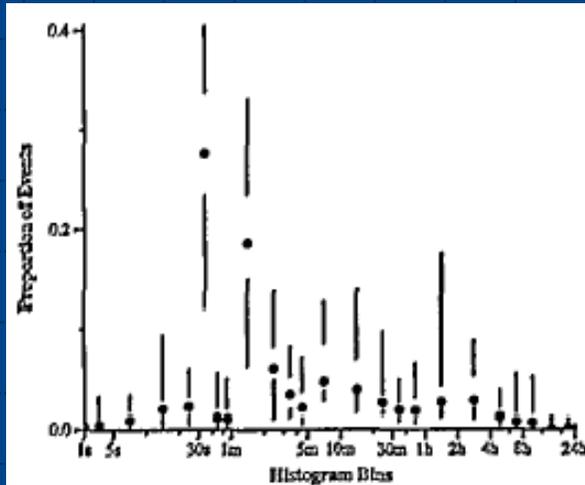


AADup

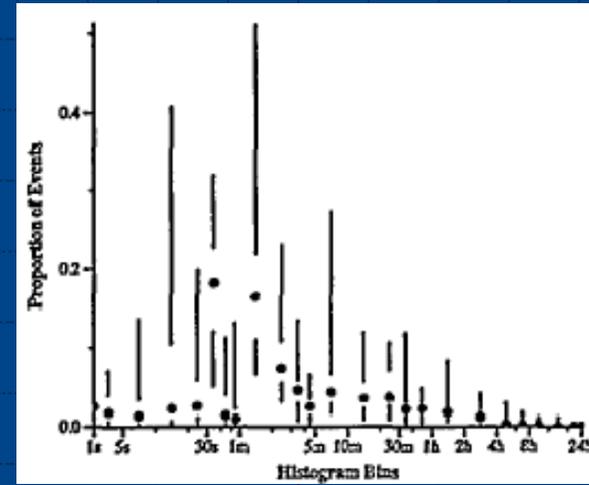


WADup

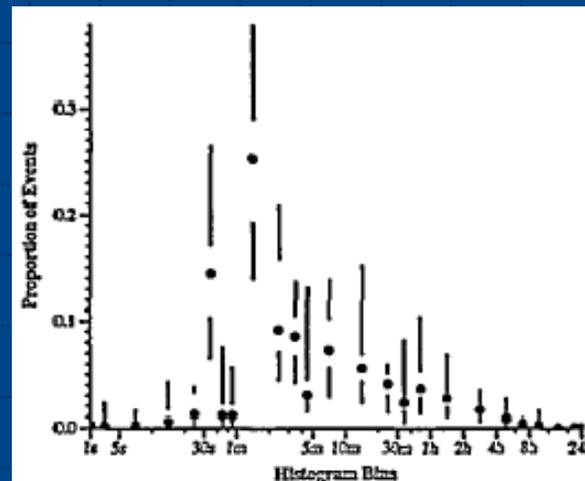
Temporal Properties of Instability Statistics



AADiff



WADiff



WADup

Conclusion

- ◆ Forwarding instability can have a significant deleterious impact on the Internet infrastructure
- ◆ majority (99 percent) of routing information is pathological and may not reflect real network topological changes
- ◆ sufficiently high rates of pathological updates (e.g. 300 updates per second) are enough to crash a widely deployed, high-end model of commercial Internet router