

# Internet Routing Protocols

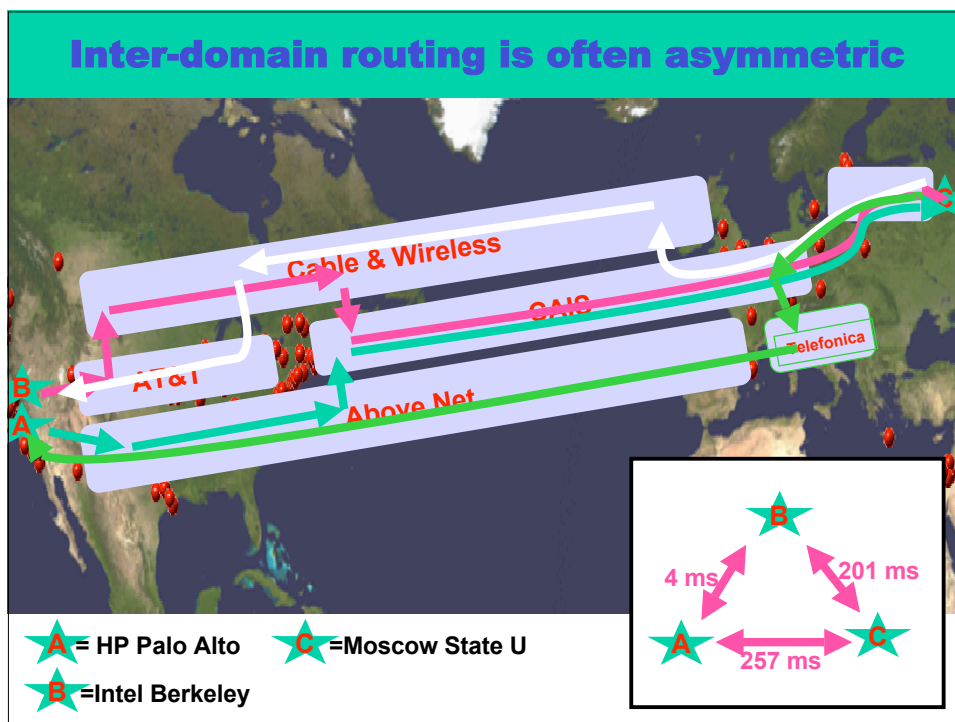
## Lecture 05

### BGP Dynamics

## Advanced Systems Topics

Lent Term, 2008

Timothy G. Griffin  
Computer Lab  
Cambridge UK




## BGP Dynamics

- How many updates are flying around the Internet?
- How long Does it take Routes to Change?

**The goals of**  
**(1) fast convergence**  
**(2) minimal updates**  
**(3) path redundancy**  
**are at odds.**

**Pick any two!!**


### Hourly Average of Per-Second Updated and Withdrawn Prefix Rate



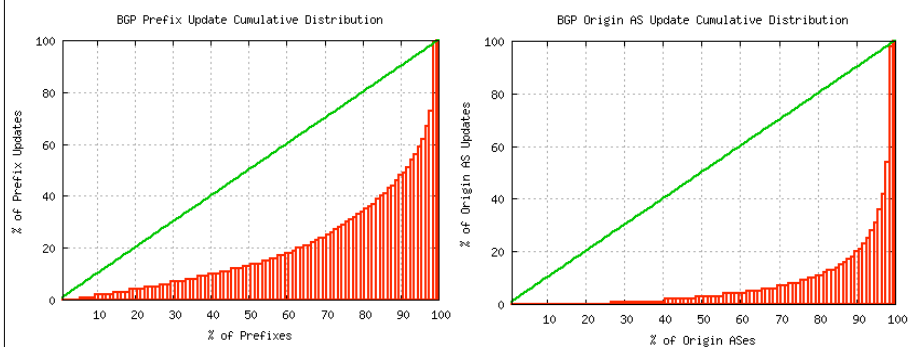
<http://bgpupdates.potaroo.net>

**FEB 6  
2008**

## Hourly Peak of Per-Second Updated and Withdrawn Prefix Rate



## Cumulative Distribution of per-Prefix and per-AS Updates




## Q: Why All the Updates?

- The Internet is large, so isn't there always something going on somewhere? (That is, BGP is just doing a good job of keeping things connected!)
- Is BGP exploring many alternate paths during convergence?
- Are IGP instabilities being exported to the interdomain world?
- Have bad tradeoffs been made in router software implementation?
- Are BGP sessions being reset due to congestion?
- Weird policy interactions like MED oscillation?
- Gnomes, sprites, and fairies
- ....


**A: NO ONE REALLY KNOWS ...**  
**BGP does a very good job hiding information!**

## Routing Change: Path Exploration

- Initial situation
  - Destination 0 is alive
  - All ASes use direct path
- When destination dies
  - All ASes lose direct path
  - All switch to longer paths
  - Eventually withdrawn
- E.g., AS 2
  - $(2,0) \rightarrow (2,1,0)$
  - $(2,1,0) \rightarrow (2,3,0)$
  - $(2,3,0) \rightarrow (2,1,3,0)$
  - $(2,1,3,0) \rightarrow \text{null}$




## IGP Tie Breaking Can Export Internal Instability to the Whole Wide World



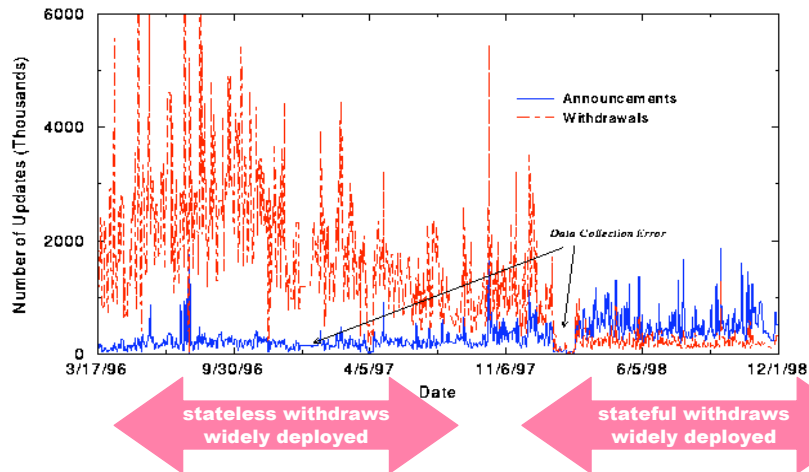
9

## MEDs Can Export Internal Instability




10

## Implementation Does Matter!



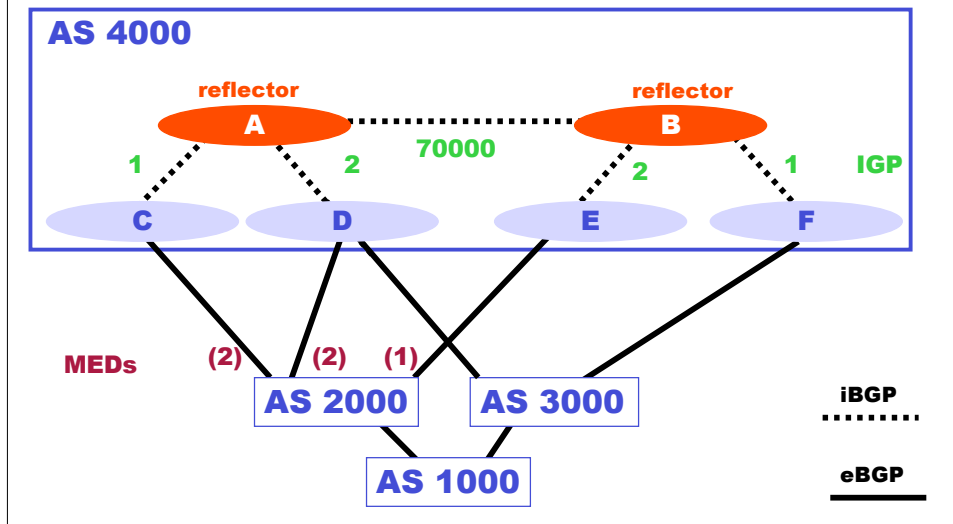
Thanks to Abha Ahuja and Craig Labovitz for this plot.

## Conjestion can take down BGP sessions! The SQL Slammer worm

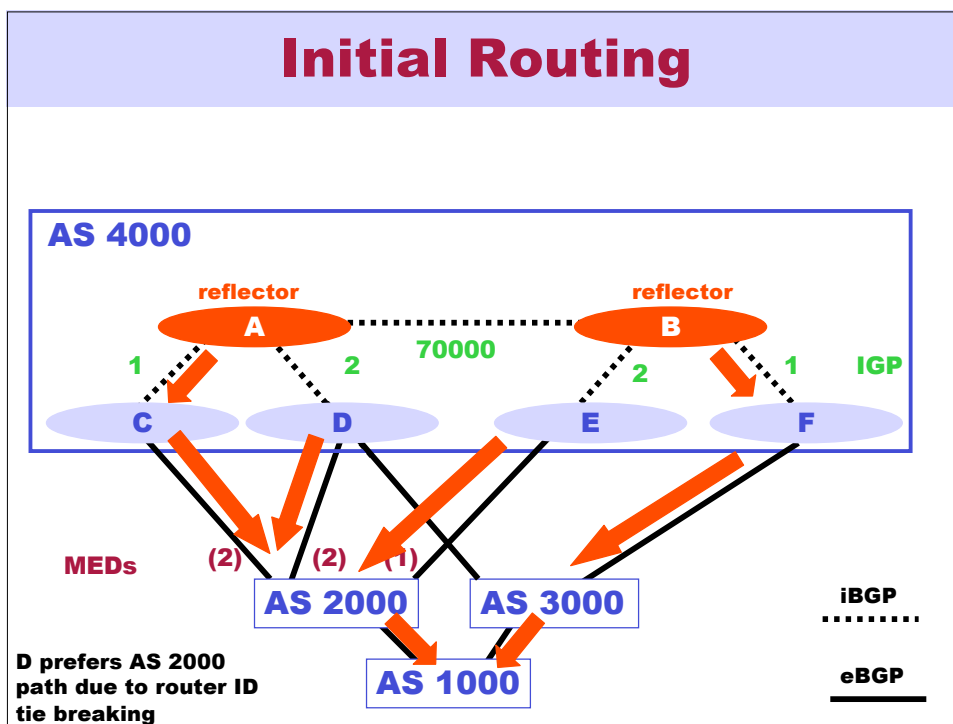


# BEWARE of MED OSCILLATION!

Only AS 2000 sends MEDs to AS 4000



## Initial Routing








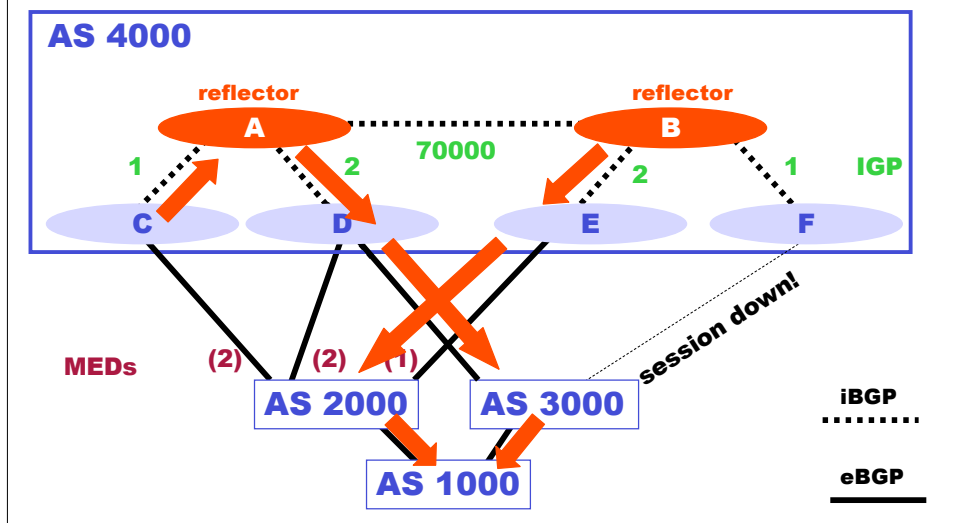
## C & D change routes

The MED 1 route from A knocks both MED 2 routes out of the picture for C & D ...



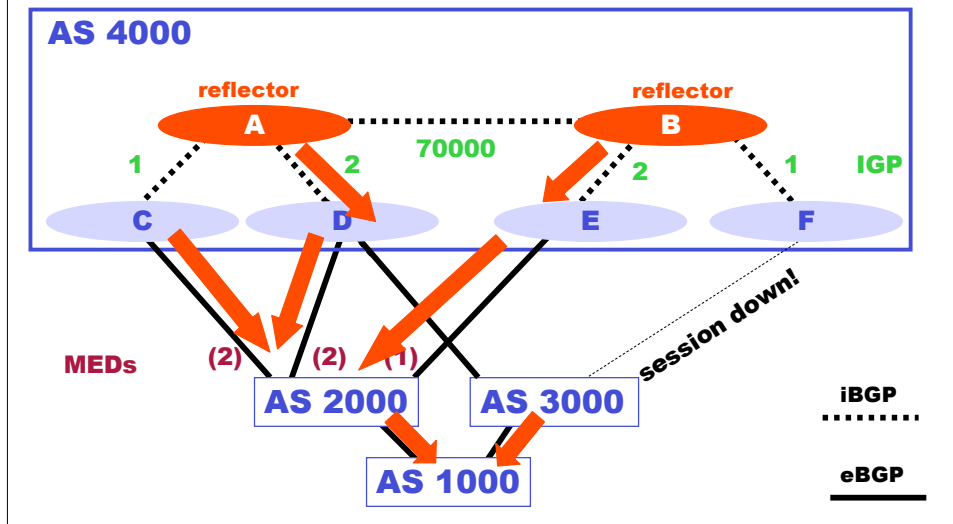
## A changes route again

A now sees the route from D through AS 3000, and it is closer IGP-wise than the route from B...



## C & D go back to initial routes

C & D no longer see MED 1 route from A, so they return to the eBGP routes with MED 2...



## Back to state zero!

A switches back to egress C, which is closer.

