BGP Routing Convergence Times [Labovitz00a]

CSci551: Computer Networks SP2006 Thursday Section John Heidemann

4d_Labovitz00a: CSci551 SP2006 © John Heidemann

Context • BGP widely deployed in the Internet • but poorly understood



1



















- BGP has a *minimum advertisement interval timer* – designed to limit updates
 - and to encourage aggregation
- How does it affect convergence?
 - by delaying announcements, routers figure out the pain sooner
 - see section 5.2
- result: waiting speeds convergence by allowing routers to make decisions on more complete information

4d_Labovitz00a: CSci551 SP2006 © John Heidemann

Does this explain measurements?

- Tup/Tshort converge quickly because they shorten path length and therefore are quickly accepted
- Tdown/Tlong converge slowly because BGP tries hard to find all alternatives
 - Tlong actually *sometimes* goes quicker if it's "not long enough" and can preempt some of the thrashing

38

4d_Labovitz00a: CSci551 SP2006 © John Heidemann

Other Observations

- Could do loop detection at *sender* side and not just receiver side
- how hard were the experiments?
 - need distributed hosts
 - need a fair amount of time to observe net
 - but doable (and more doable today)

4d_Labovitz00a: CSci551 SP2006 © John Heidemann

40

37