Large-Scale Topology Discovery

Benoit Donnet, Timur Friedman
LIP6-CNRS laboratory, UPMC, Paris

ISMA Workshop on Internet Topology (WIT)
CAIDA, UC San Diego - 10 May 2006

Context

- Network measurement
- Internet topology discovery
  - distributed traceroute monitors
  - IP interface level
- CONMI Workshop
  - community-oriented network measurements

Existing Tools

- skitter (CAIDA)
- scamper (WAND)
- TTM (Ripe NCC)
- AMP (NLANR)
- Rocketfuel (University of Washington)
- Scriptroute (University of Washington)
- DIMES (Tel Aviv University)

Problem: *Intra-monitor* Redundancy

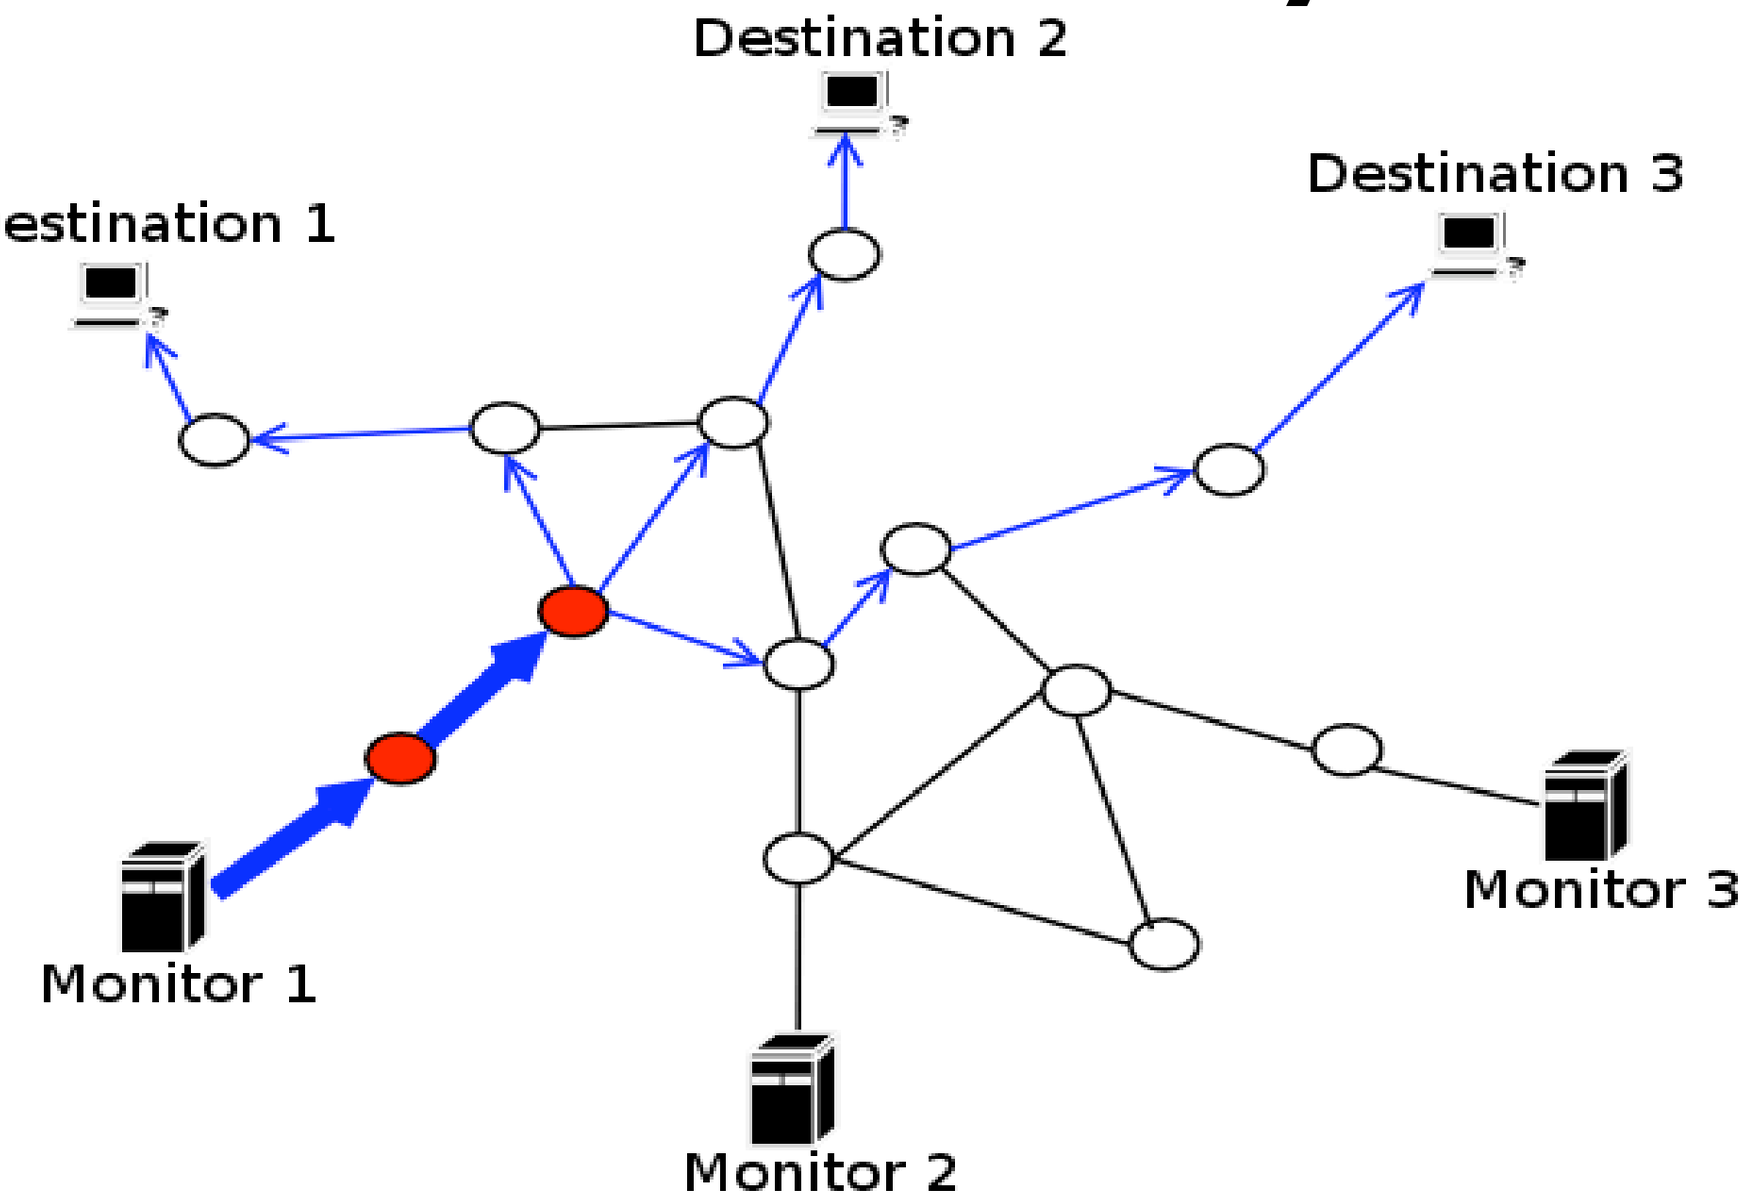

Problem: *Inter-Monitor* Redundancy

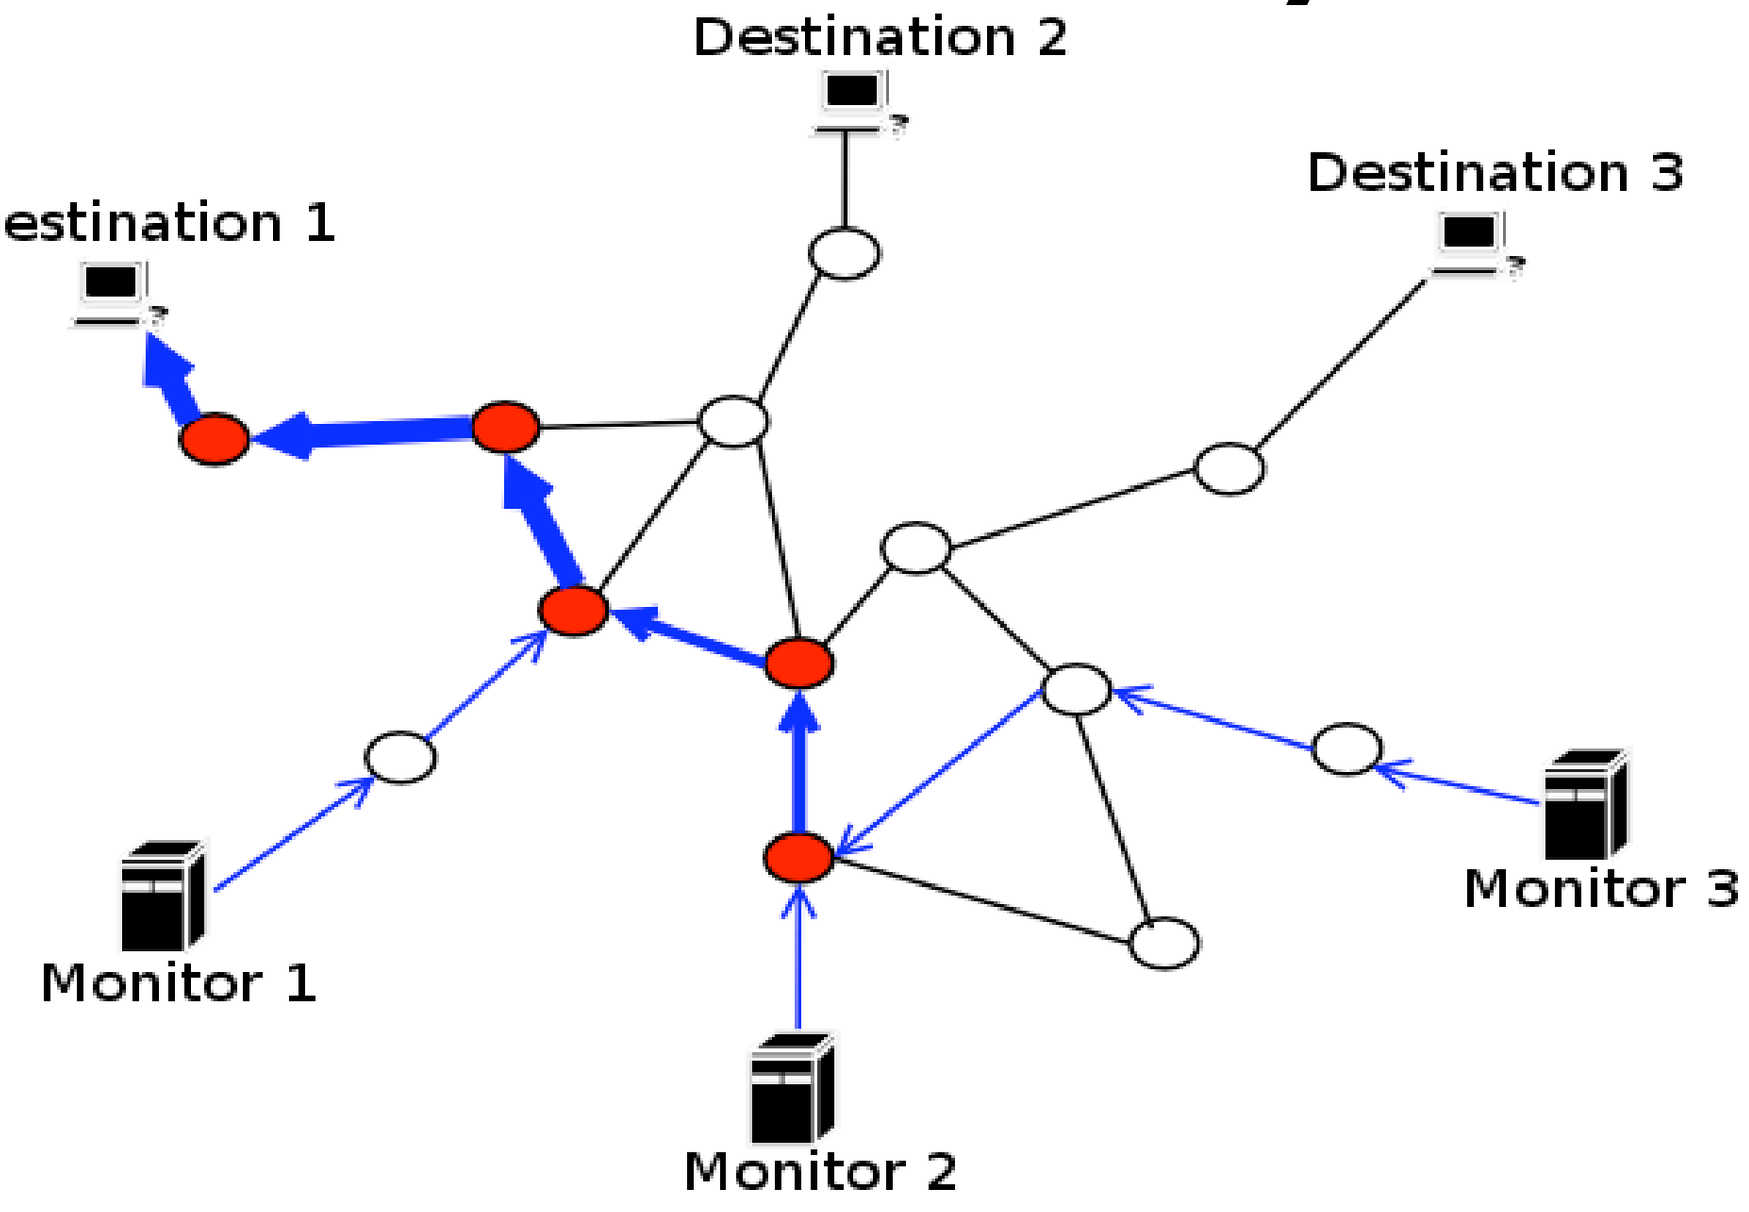

Statement

- Classical approaches are inefficient
- Wide scale deployment of traceroute tools
  - DIMES
- Scale up the number of monitors
  - little work on efficiency
- Be careful regarding
  - network resources
  - load on destinations

Doubletree: Basics

- Cooperative algorithm
- Goal: avoid paths already explored
- Exploits tree-like structure of routes in the internet:
  - from a monitor to a set of destinations
    - ➡ *monitor-rooted tree*
  - from a set of monitors to a destination
    - ➡ *destination-rooted tree*

Doubletree: *Monitor-rooted* Tree

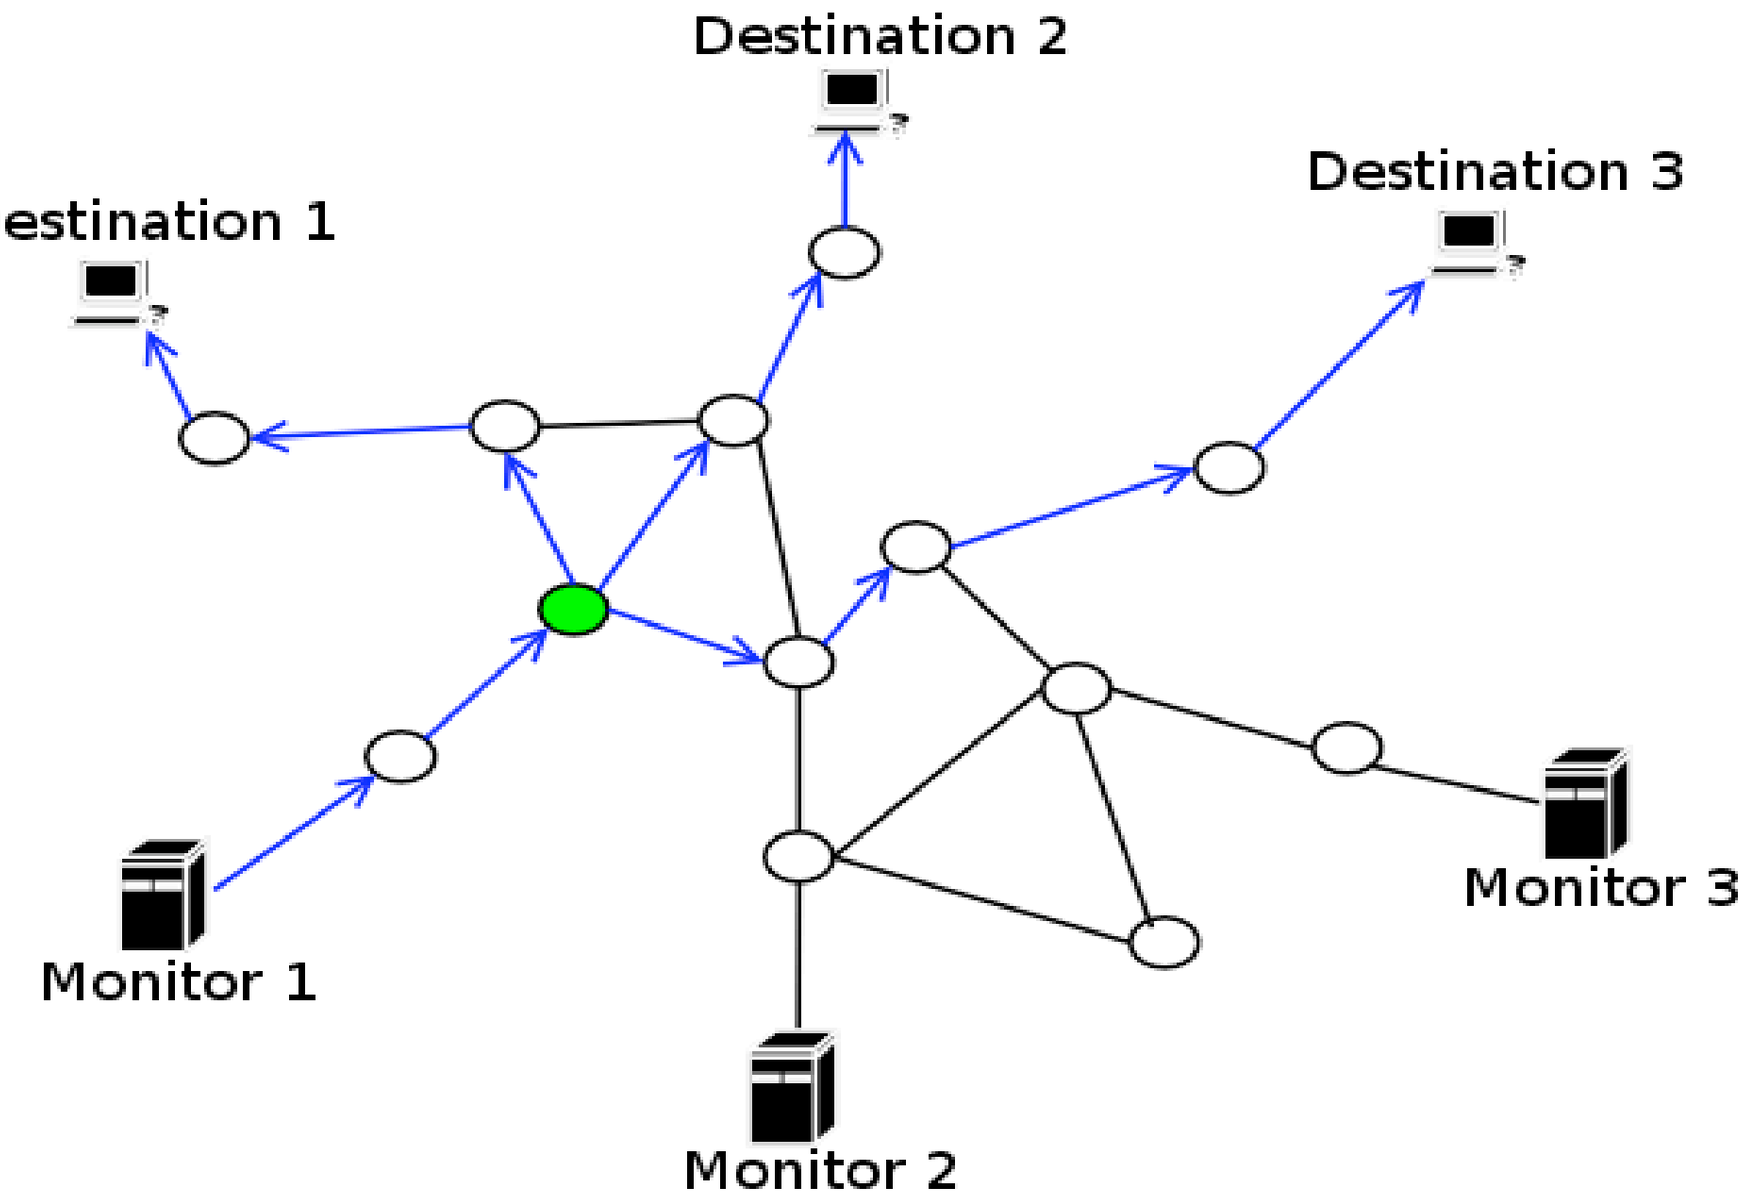

Doubletree: *Destination-rooted* Tree

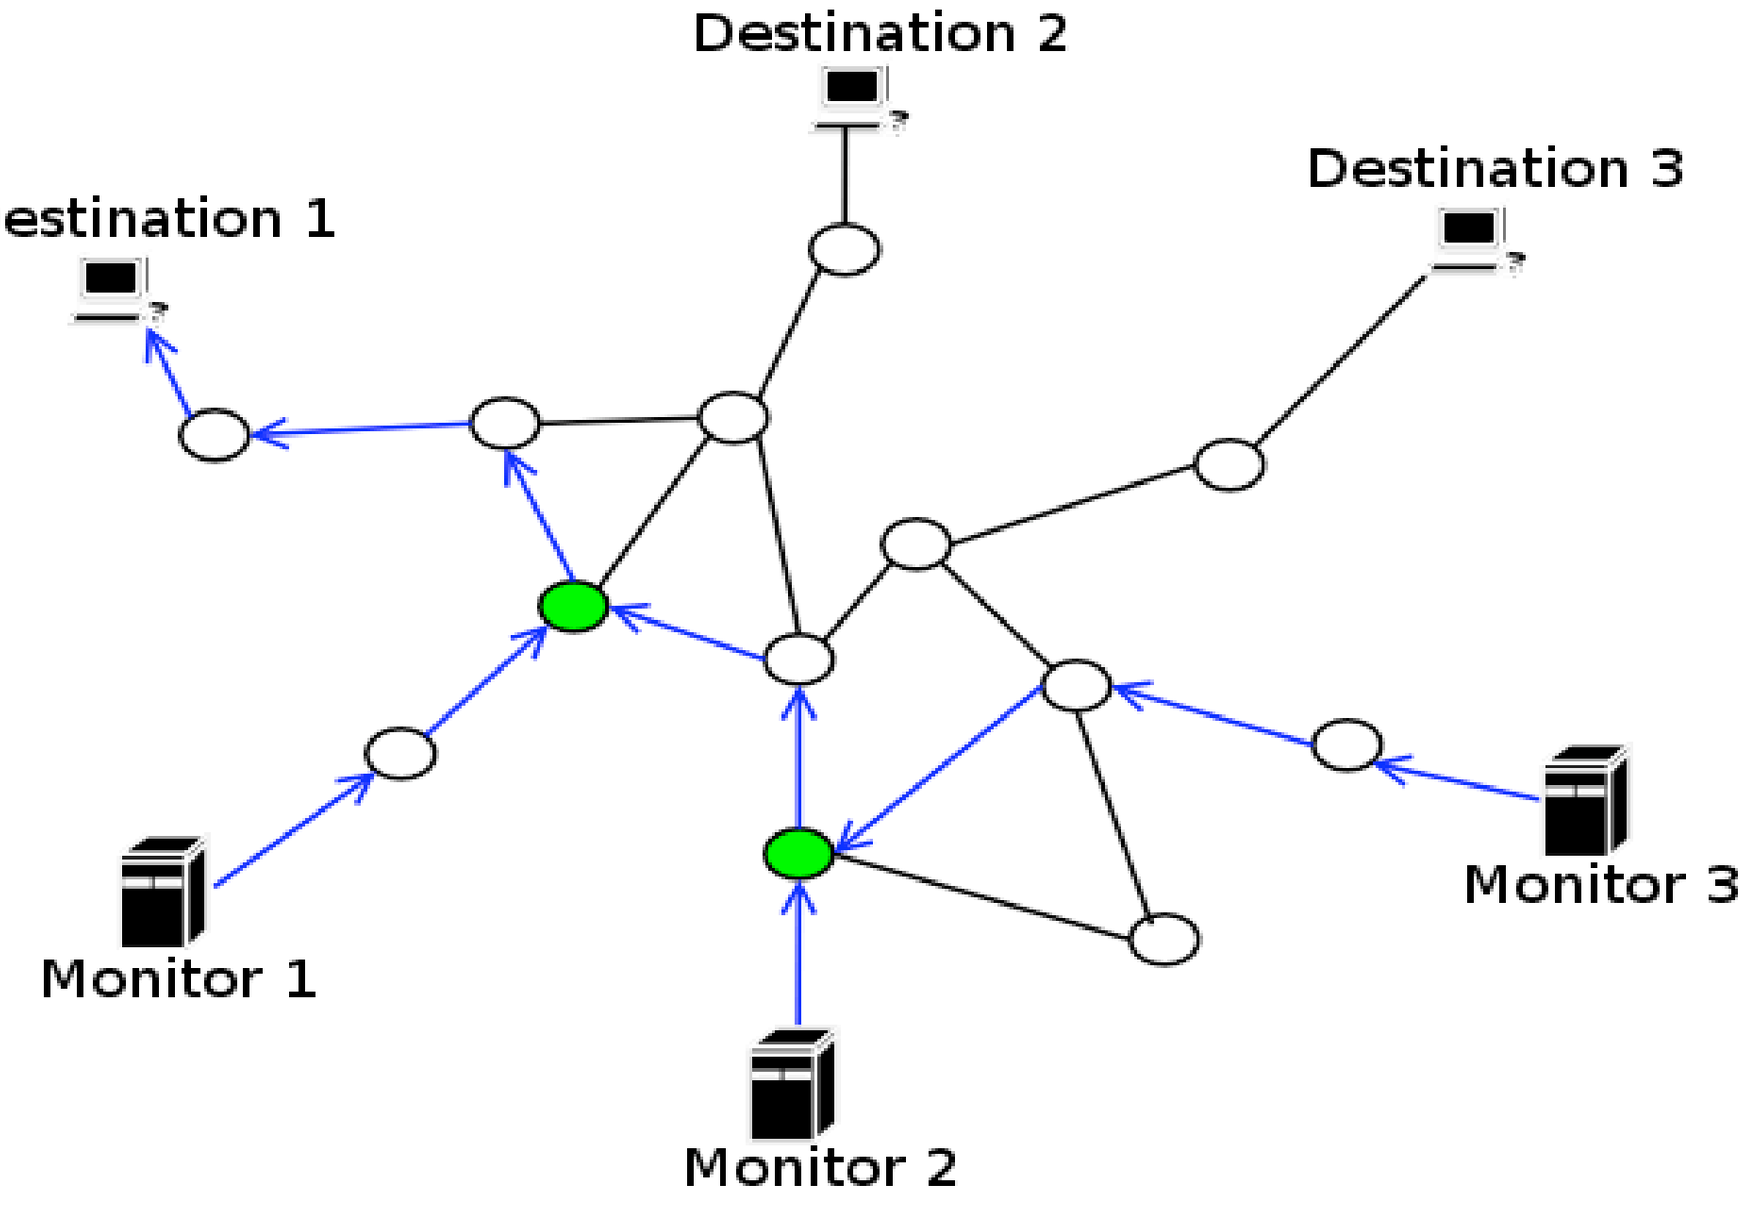

Doubletree: Reconciliation

- Backward and forward probing are opposite schemes
- How can we reconcile them?
  - Start probing at some hop *h*
    - First, probe forward from *h*
    - Second, probe backward from *h-1*

Doubletree: Reconciliation (2)

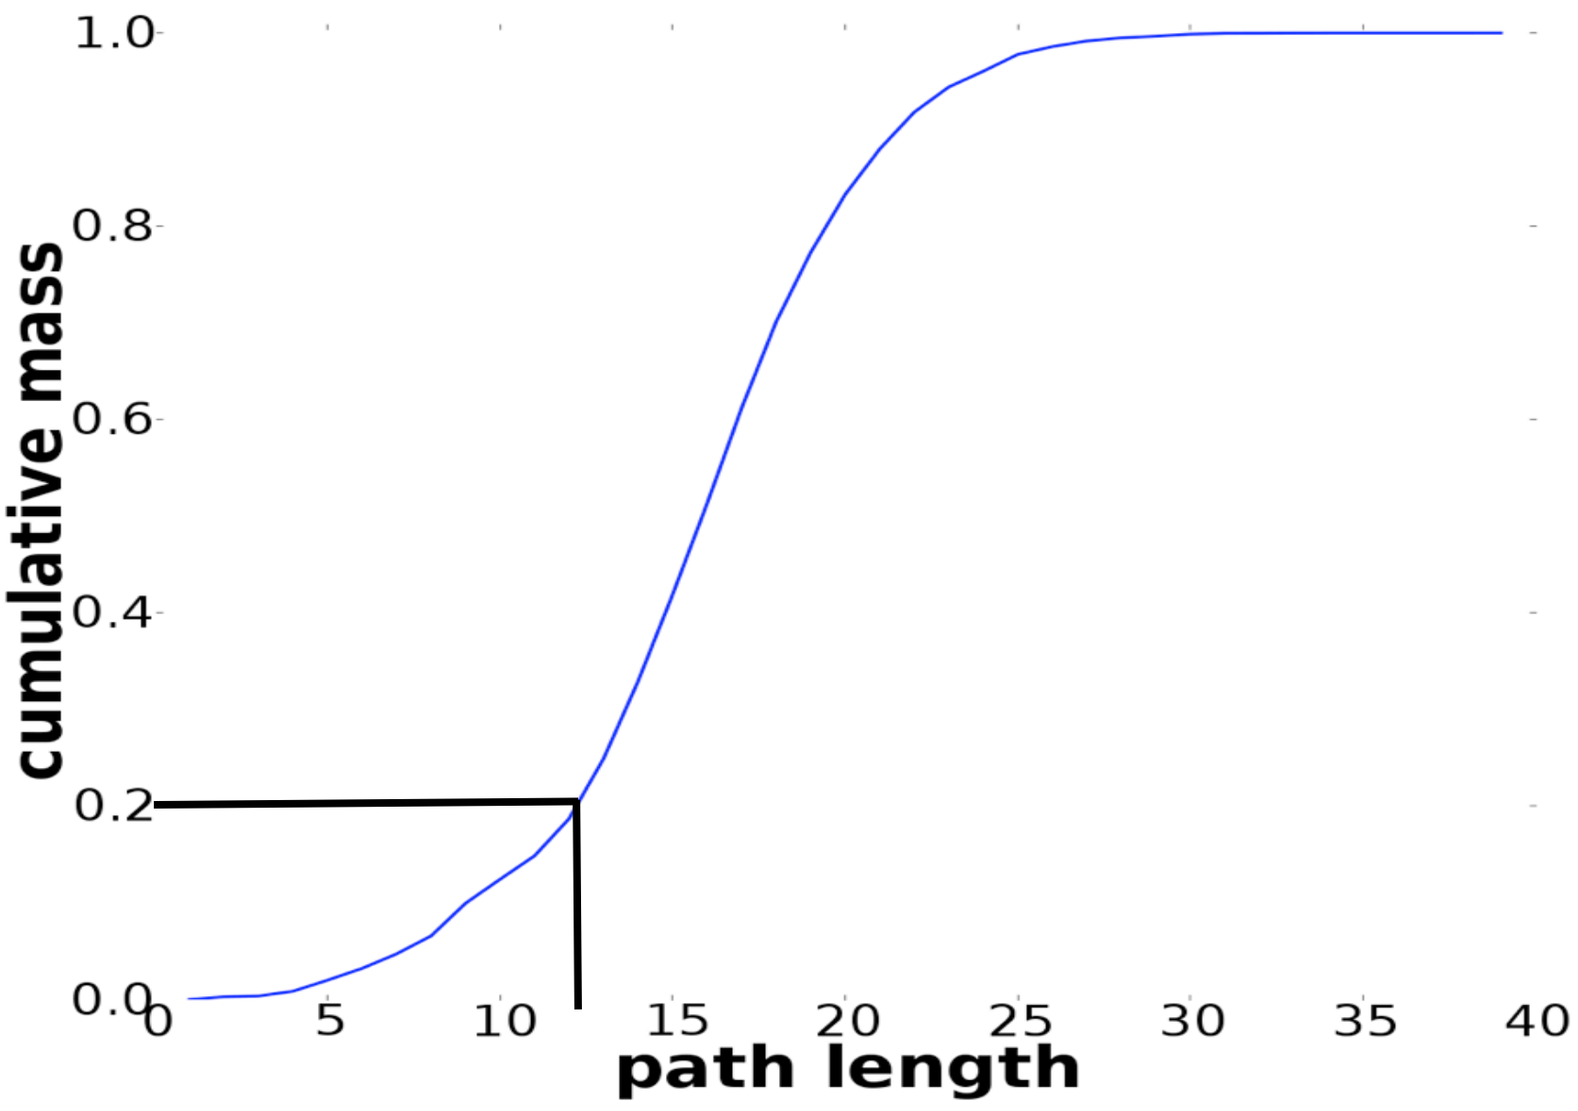

Doubletree: Stop Sets

- It is not necessary to maintain the whole tree structure
- Each monitor uses *stop sets*: {(interface, root)}
  - *Local Stop Set* B: {interface}
    - ➡ Backward probing
  - *Global Stop Set* F: {(interface, destination)}
    - ➡ Forward probing
    - ➡ Shared between monitors

Doubletree: Coverage

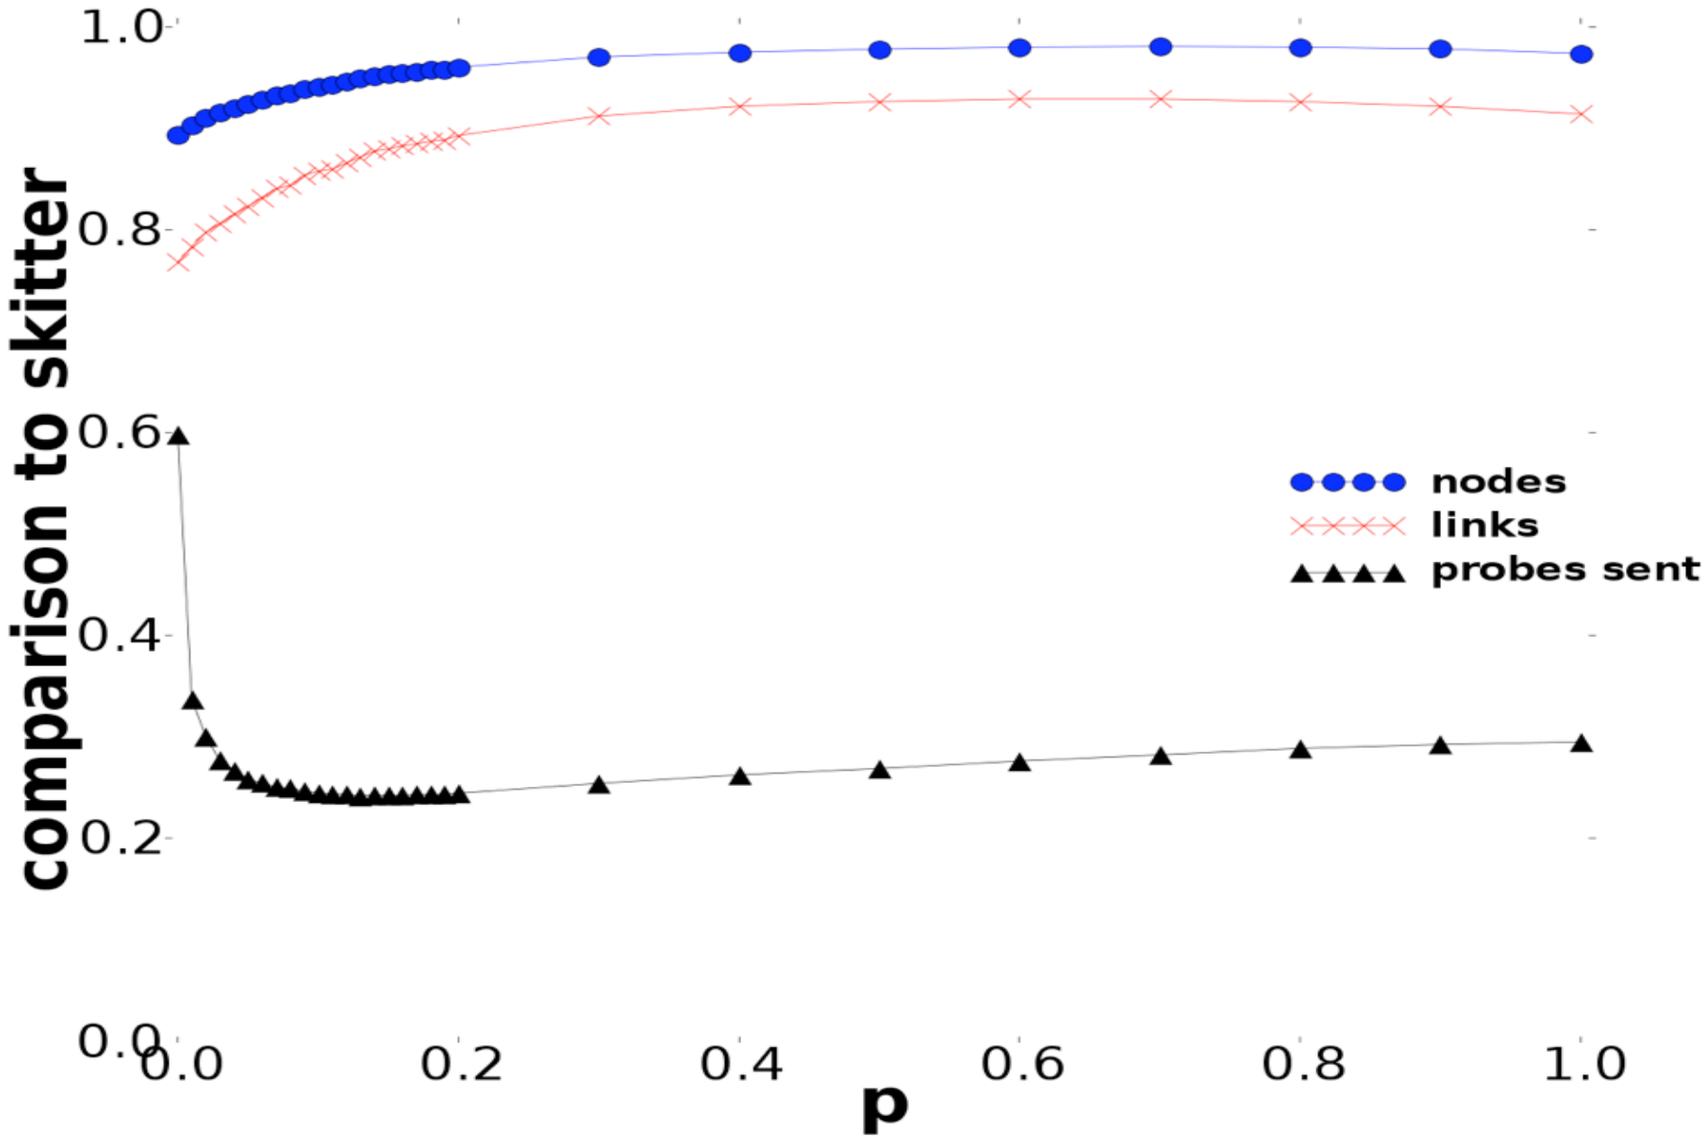

Doubletree: Redundancy

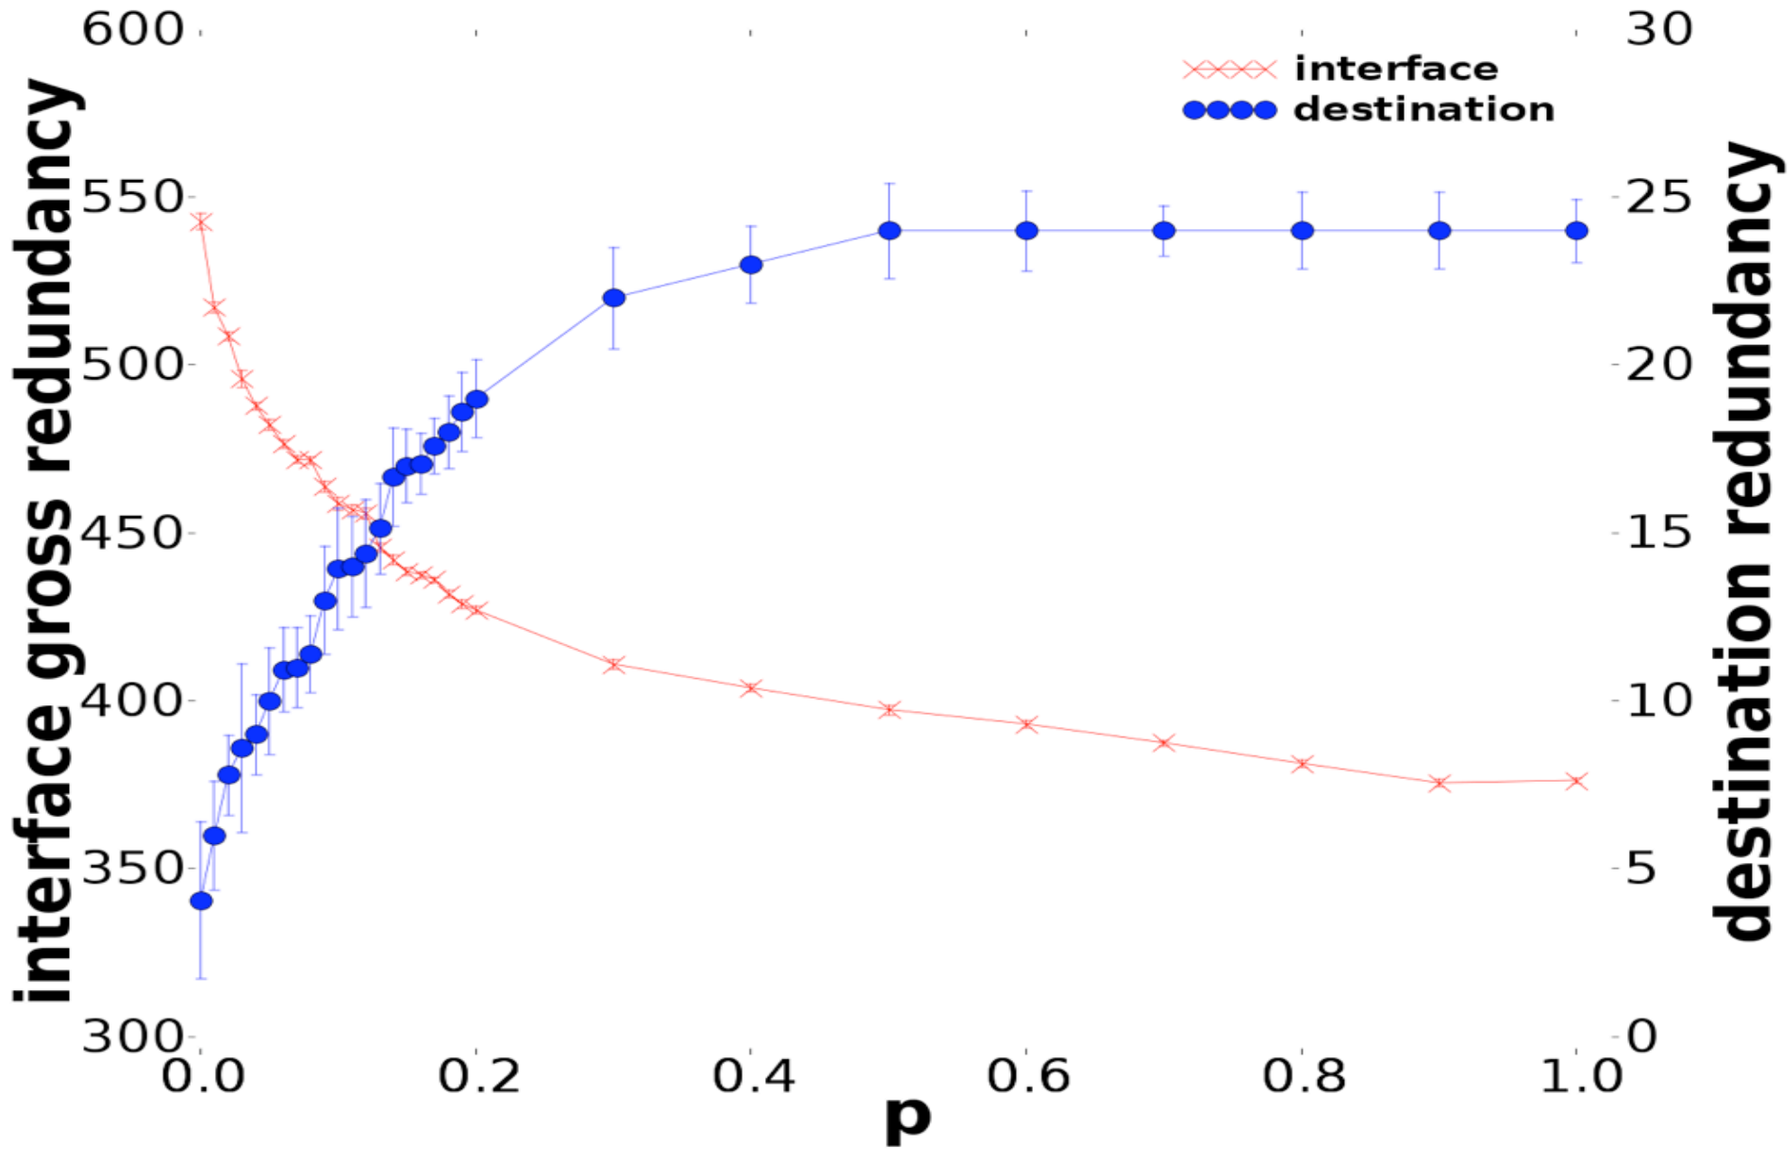

Doubletree: Deployment

- *traceroute@home*
  - Java program
  - Open source, BSD-like license
  - Freely available: http://trhome.sourceforge.net
- Deployed on the PlanetLab testbed
  - Doubletree behavior confirmed
  - Simulation results confirmed

Infrastructures: OneLab

- European project
  - Due to start in September 2006
  - Ten partners (university & industry)
- PlanetLab extensions
  - European administration
  - Wireless environments
  - Monitoring capabilities
    - ➡ Doubletree/traceroute@home integration

Future Work

- Improvement to Doubletree
  - Diamonds
- Capturing the network dynamics
- BGP-guided probing
- Overlay/P2P structure for traceroute@home
- Measurement tools standardization

Summary

- Standard approaches for tracerouting the network are inefficient
- Doubletree is an efficient and cooperative topology discovery algorithm
  - Measurement load reduced up to 76%
  - Coverage above 90%
- traceroute@home to be deployed
  - OneLab project

References

B. Donnet, P. Raoult, T. Friedman, M. Crovella.
Deployment of an Algorithm for Large-Scale Topology Discovery.
In **IEEE JSAC** Special Issue on Internet Sampling. Oct. 2006. to appear.

B. Donnet, T. Friedman.
Topology Discovery Using an Address Prefix Stopping Rule.
In **IFIP**, International Federation for Information Processing. vol. 196. Mar. 2006.

B. Donnet, T. Friedman.
Topology Discovery Using an Address Prefix Based Stopping Rule.
In Proc. **EUNICE** Worshop. Jul. 2005.

B. Donnet, P. Raoult, T. Friedman, M. Crovella.
Efficient Algorithms for Large-Scale Topology Discovery.
In Proc. **ACM SIGMETRICS**. Jun. 2005.

B. Donnet, T. Friedman.
A CIDR Stopping Rule for Topology Discovery.
In Proc. **Algotel**. May 2005.

B. Donnet, T. Friedman, M. Crovella.
Improved Algorithms for Network Topology Discovery.
In Proc. Passive and Active Measurements (**PAM**) Workshop. Mar. 2005.