On the sensitivity of transit ASs to internal failures

Steve Uhlig

suh@info.ucl.ac.be

http://www.info.ucl.ac.be/~suh

thanks to:

Bruno Quoitin
DANTE
Intel research Cambridge

Computing Science and engineering dept.
Université catholique de Louvain, Belgium

Agenda

- Network robustness of transit ASs
- Sensitivity to internal failures
- Sensitivity of the GEANT network
- Route-reflection
- Further work

Designing robust transit ASs

- IGP topology: spreading the load of the traffic over the available shortest paths
- iBGP sessions: distributing the available BGP routes inside the AS
- eBGP routes: rarely designed but imposed by economical constraints

Toy example transit AS with iBGP full-mesh

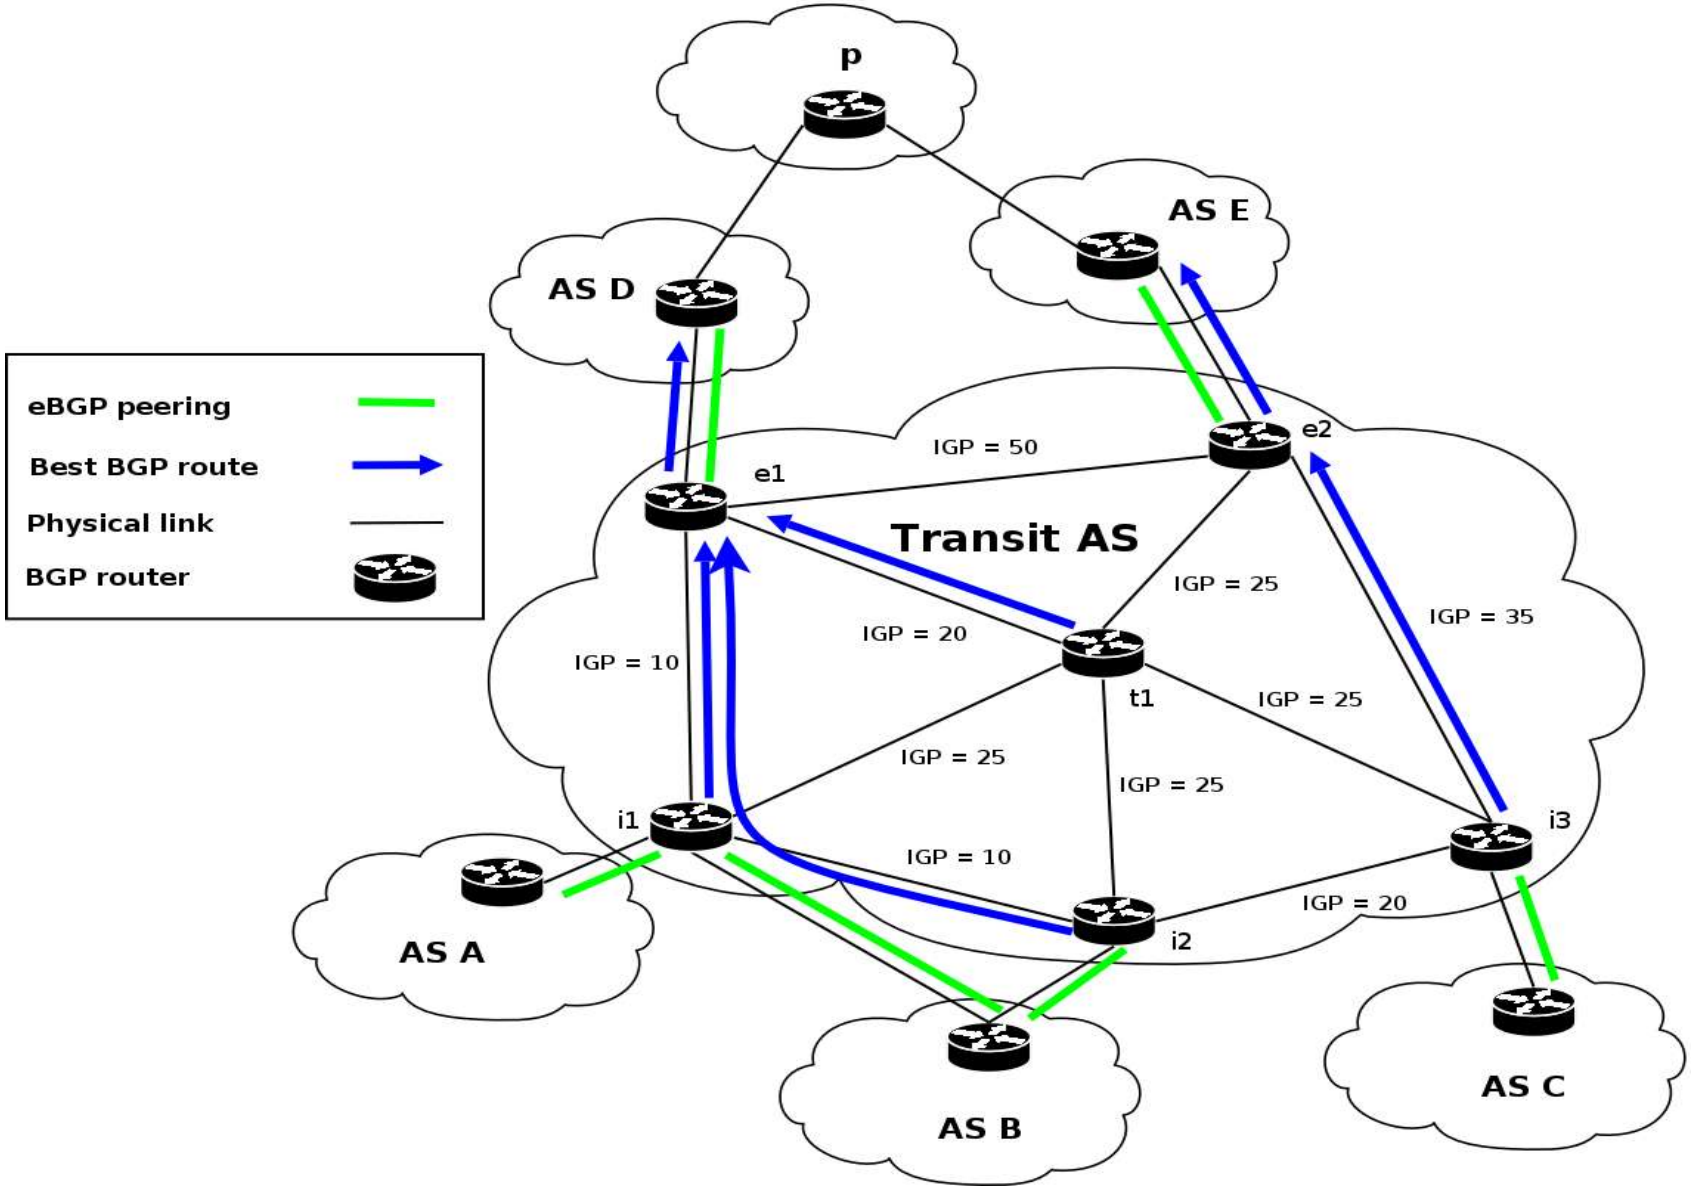

Previous work

- Potential sensitivity of a large tier-1 due to hot-potato disruptions [Teixeira et al., SIGMETRICS'04]
- Metrics to capture sensitivity to internal failures [Teixeira et al., SIGCOMM'04]
- Routing changes can have a significant impact on the traffic matrix [Teixeira et al., PAM'05]

Reproducing the routing of a transit AS

- Announce eBGP routes: important to have the actual diversity of the routes inside an AS
- Propagate routes inside iBGP structure
- Compute state of the Adj-RIB-ins of each BGP router

B. Quoitin and S. Uhlig. *Modeling the routing of an Autonomous System with C-BGP.* IEEE Network Magazine, November 2005.

Reproducing the routing of an AS with C-BGP

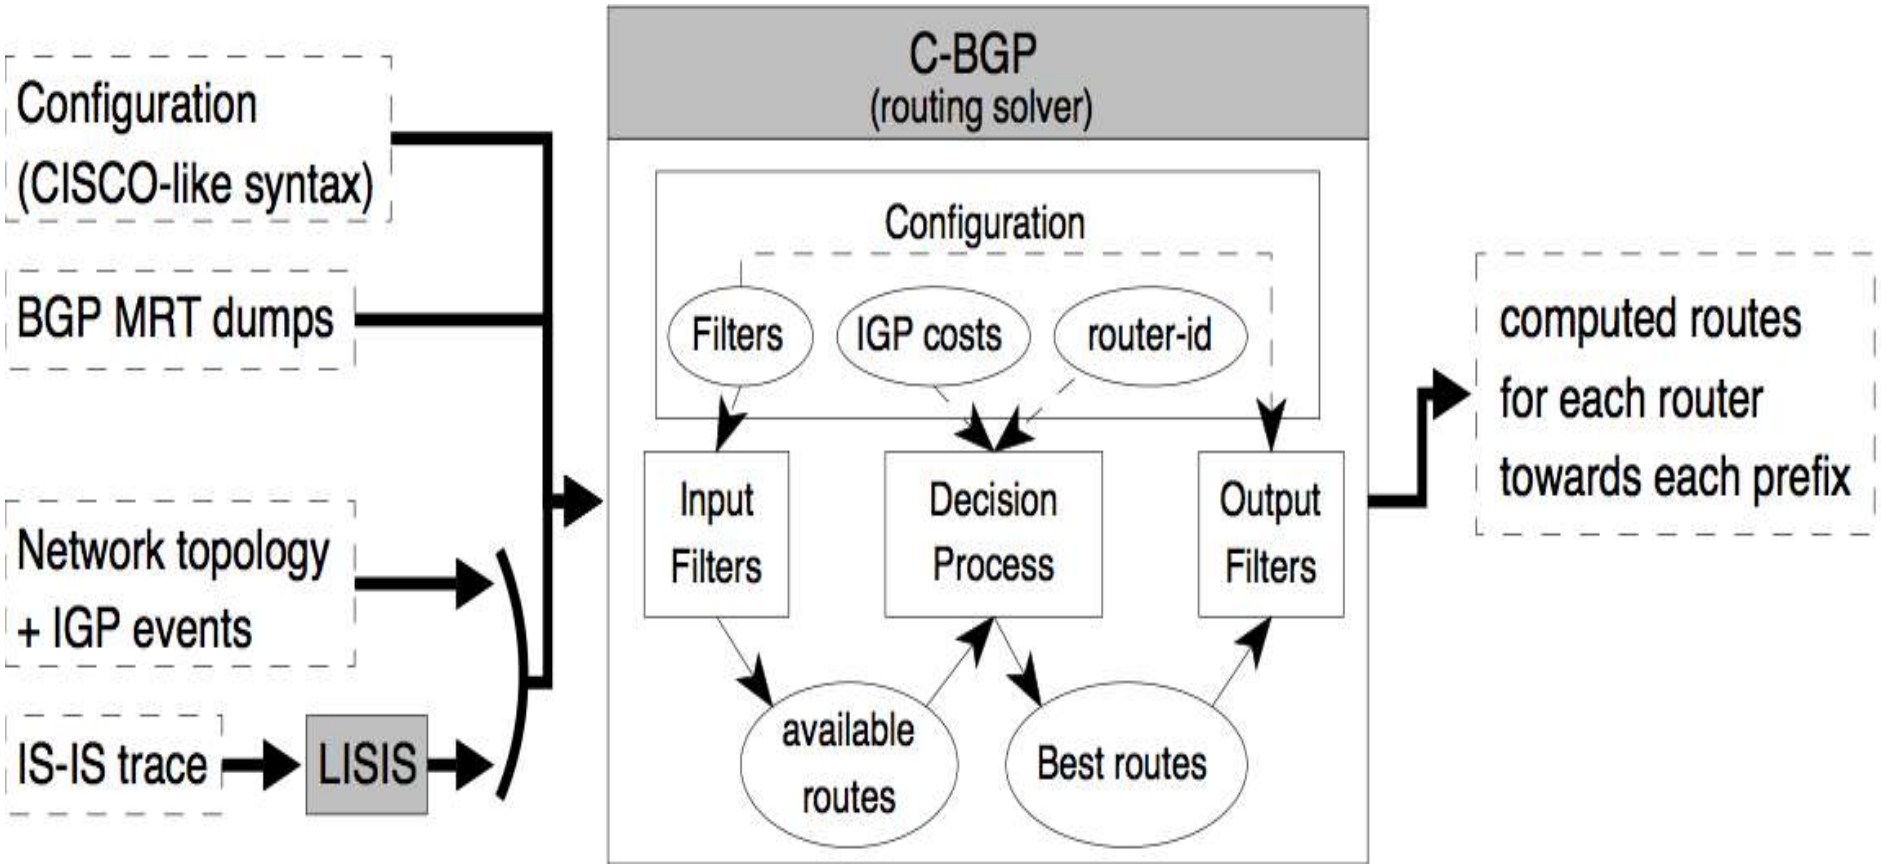

Sensitivity metrics to internal failures

- Consider a class of graph transformations on the graph G (denoted by $\triangle G$)
- Compute whether BGP changes its best route to reach prefix *p* after each graph transformation ∂G
- Metrics measure how graph transformations affect the graph (*impact*) and how each router is impacted (*sensitivity*)

Data from the GEANT network

- The GEANT network:
  - 23 POPs, 76 links, tens of eBGP peerings
  - iBGP full mesh
- 1 month of routing and traffic data in 2004:
  - ISIS and BGP capture
  - Netflow capture at all ingress interfaces

Routing impact of node failures

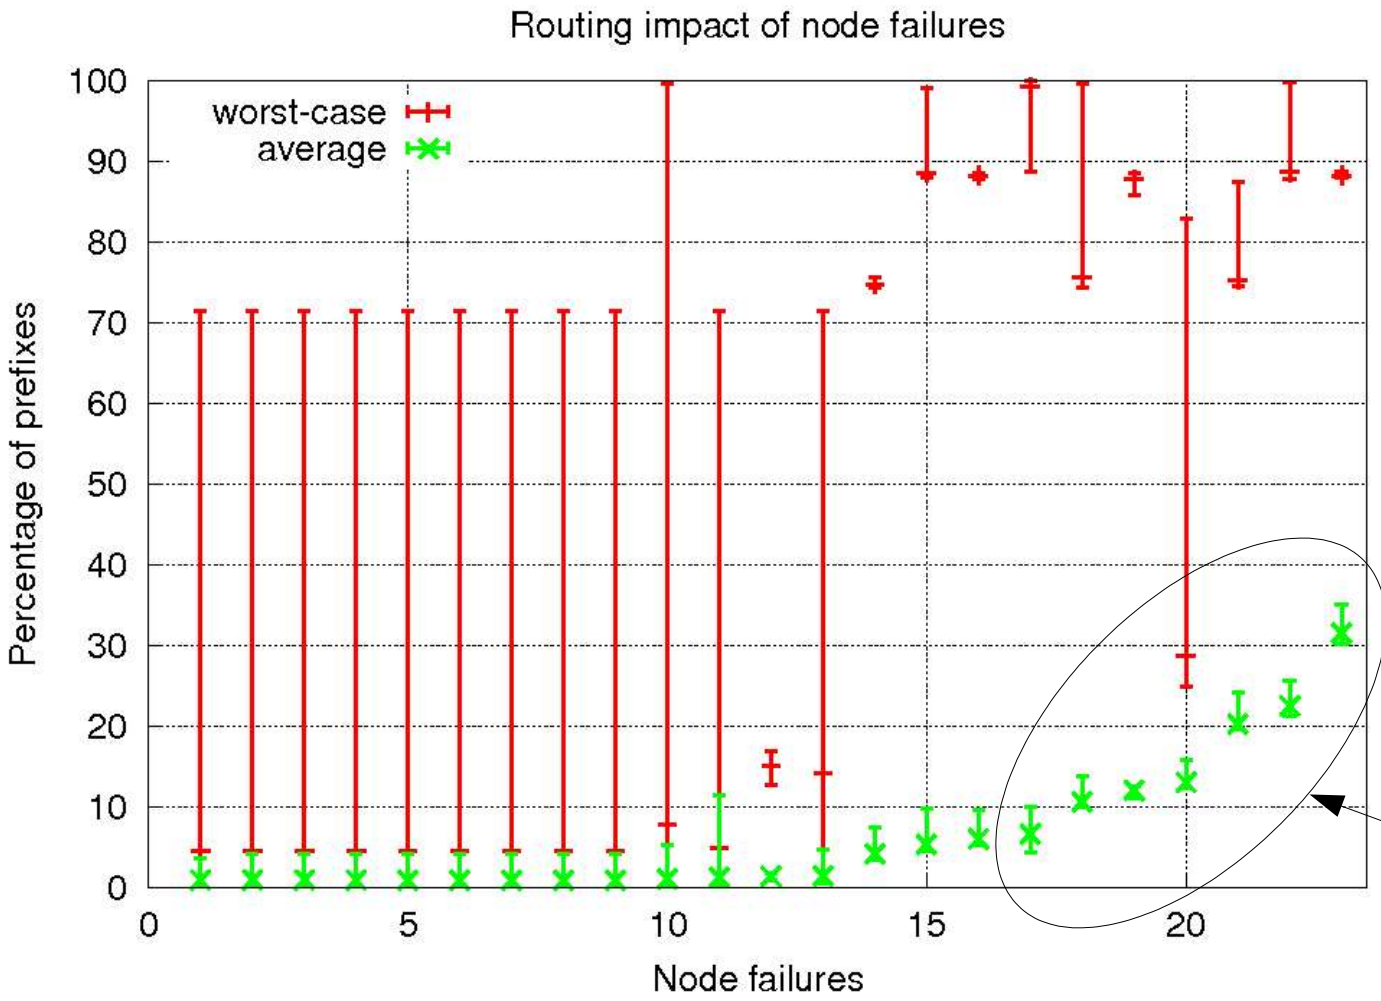

Routing impact of link failures

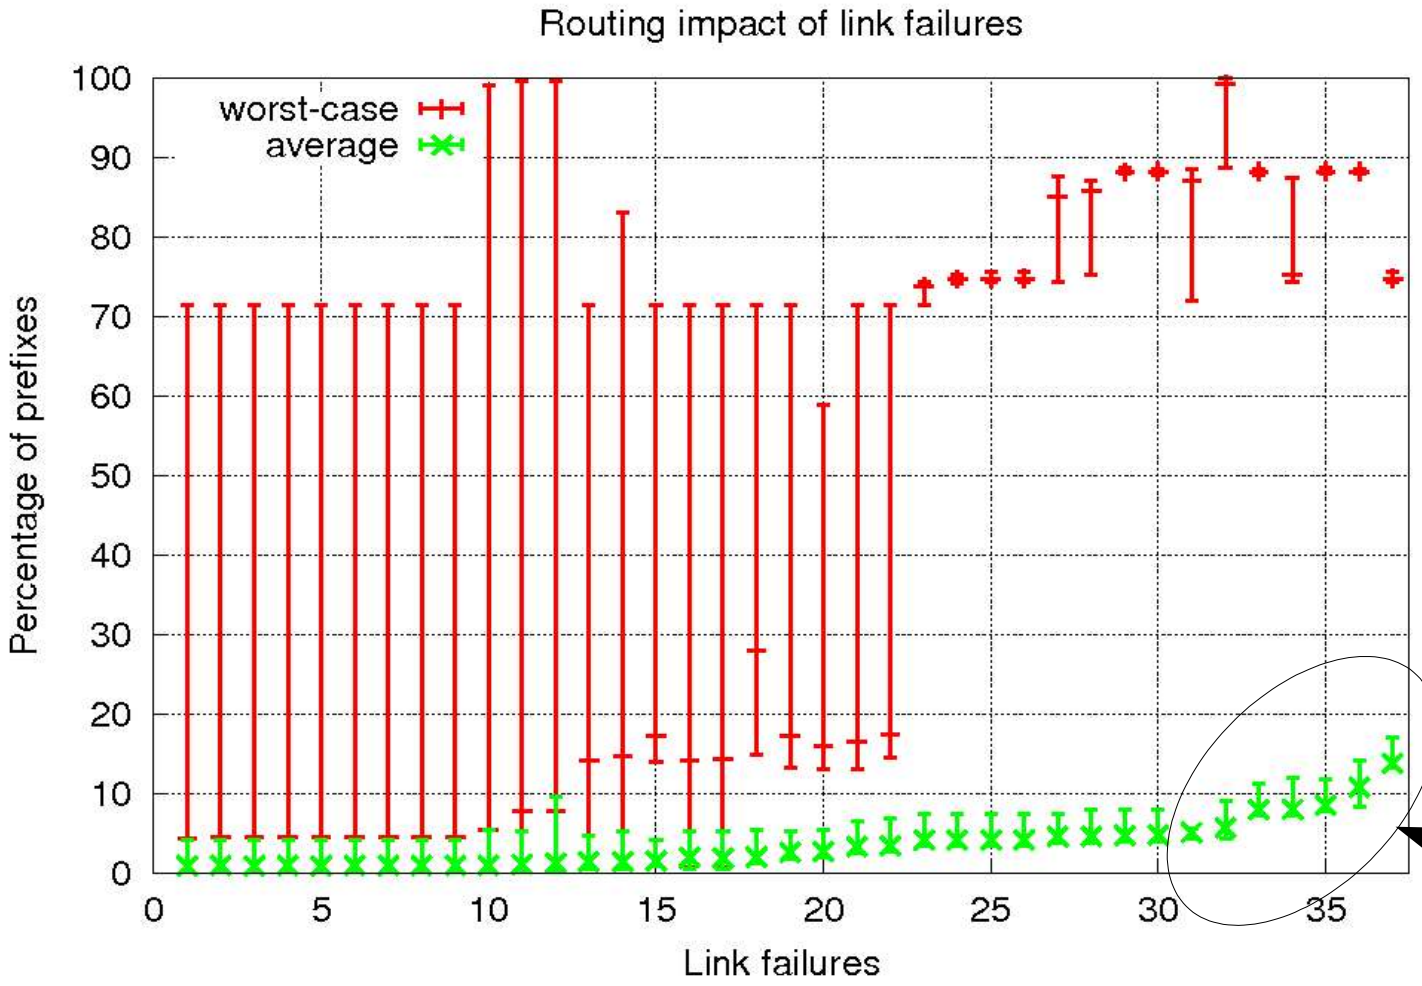

Router sensitivity to node failures

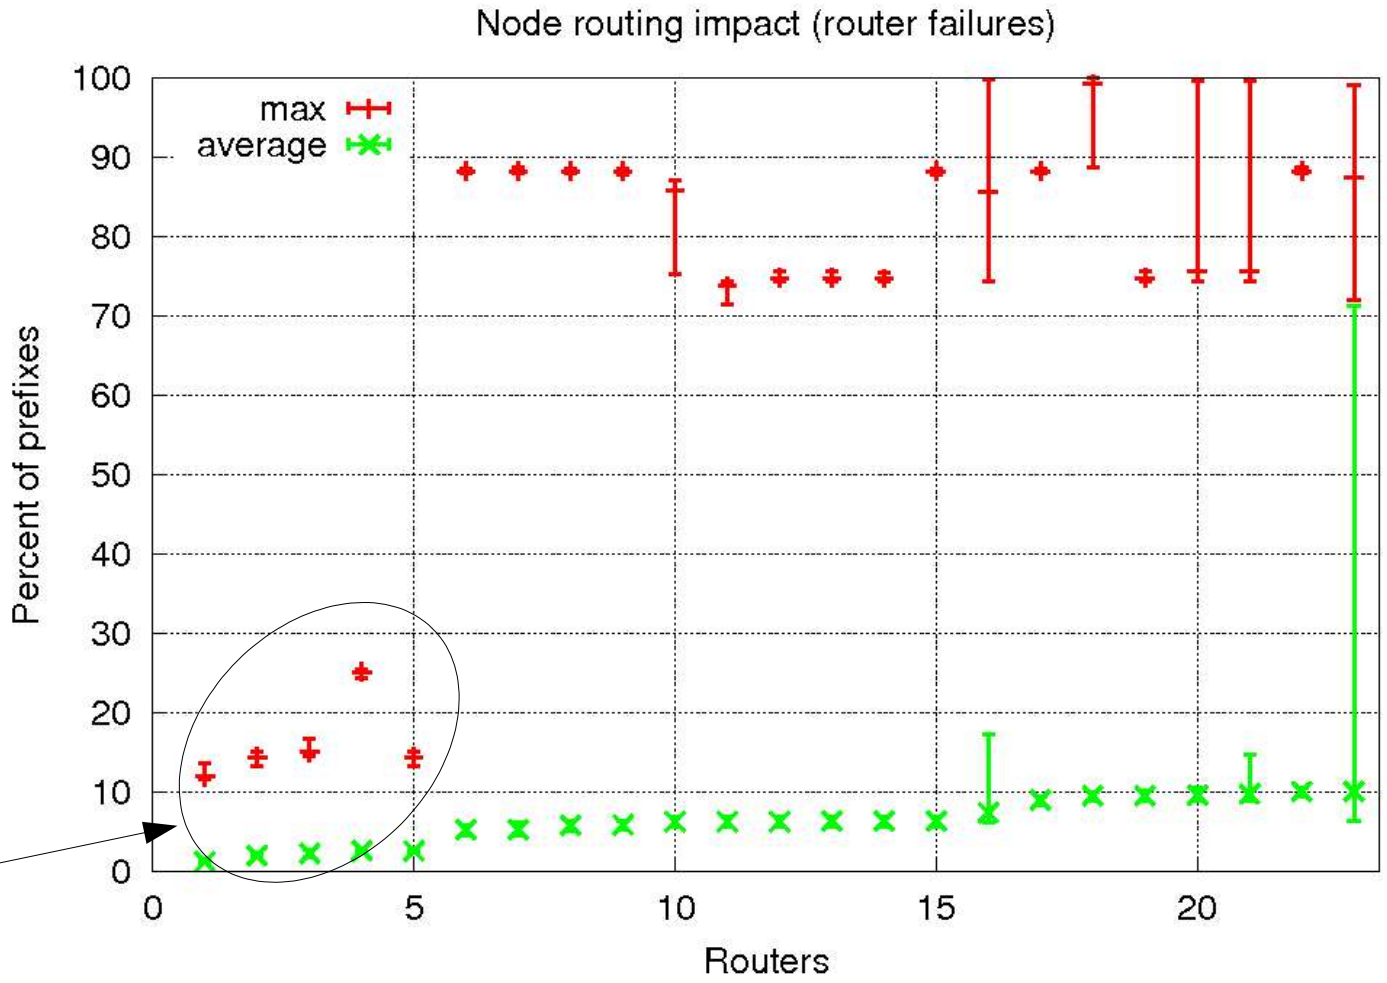

insensitive routers

Router sensitivity to link failures

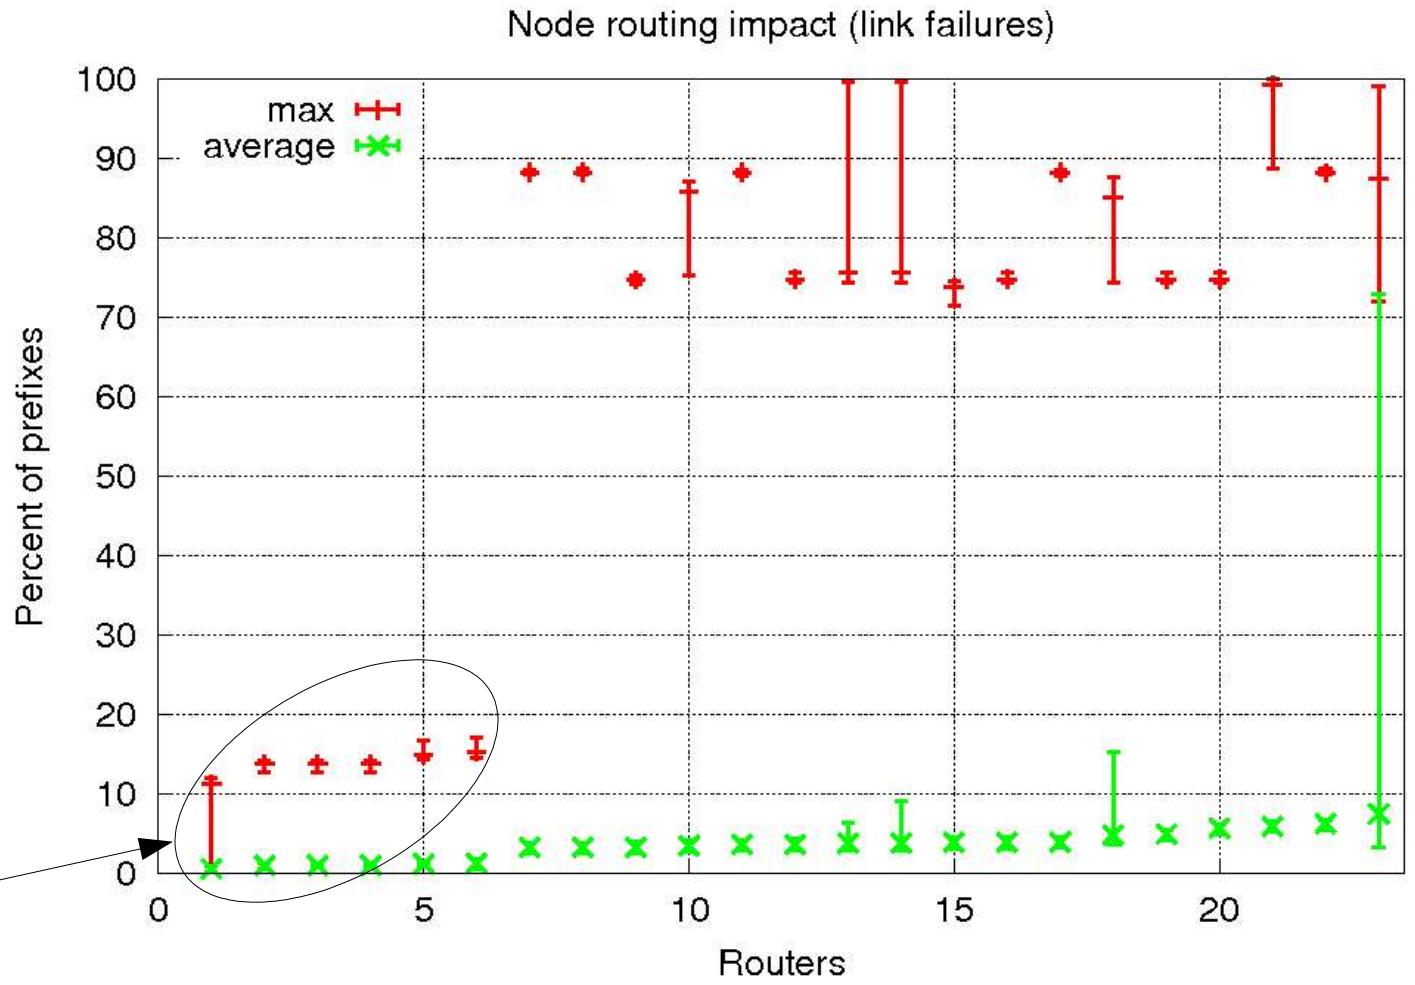

insensitive routers

Insight from sensitivity analysis

- Identifying critical links and routers
- Highly related to IGP weights setting (concentration of shortest IGP paths)
- Good enough for simple iBGP structures (full-mesh), helps find out whether to change network topopology
- What about complex transit ASs with route-reflection ?

Toy example transit AS with route-reflection

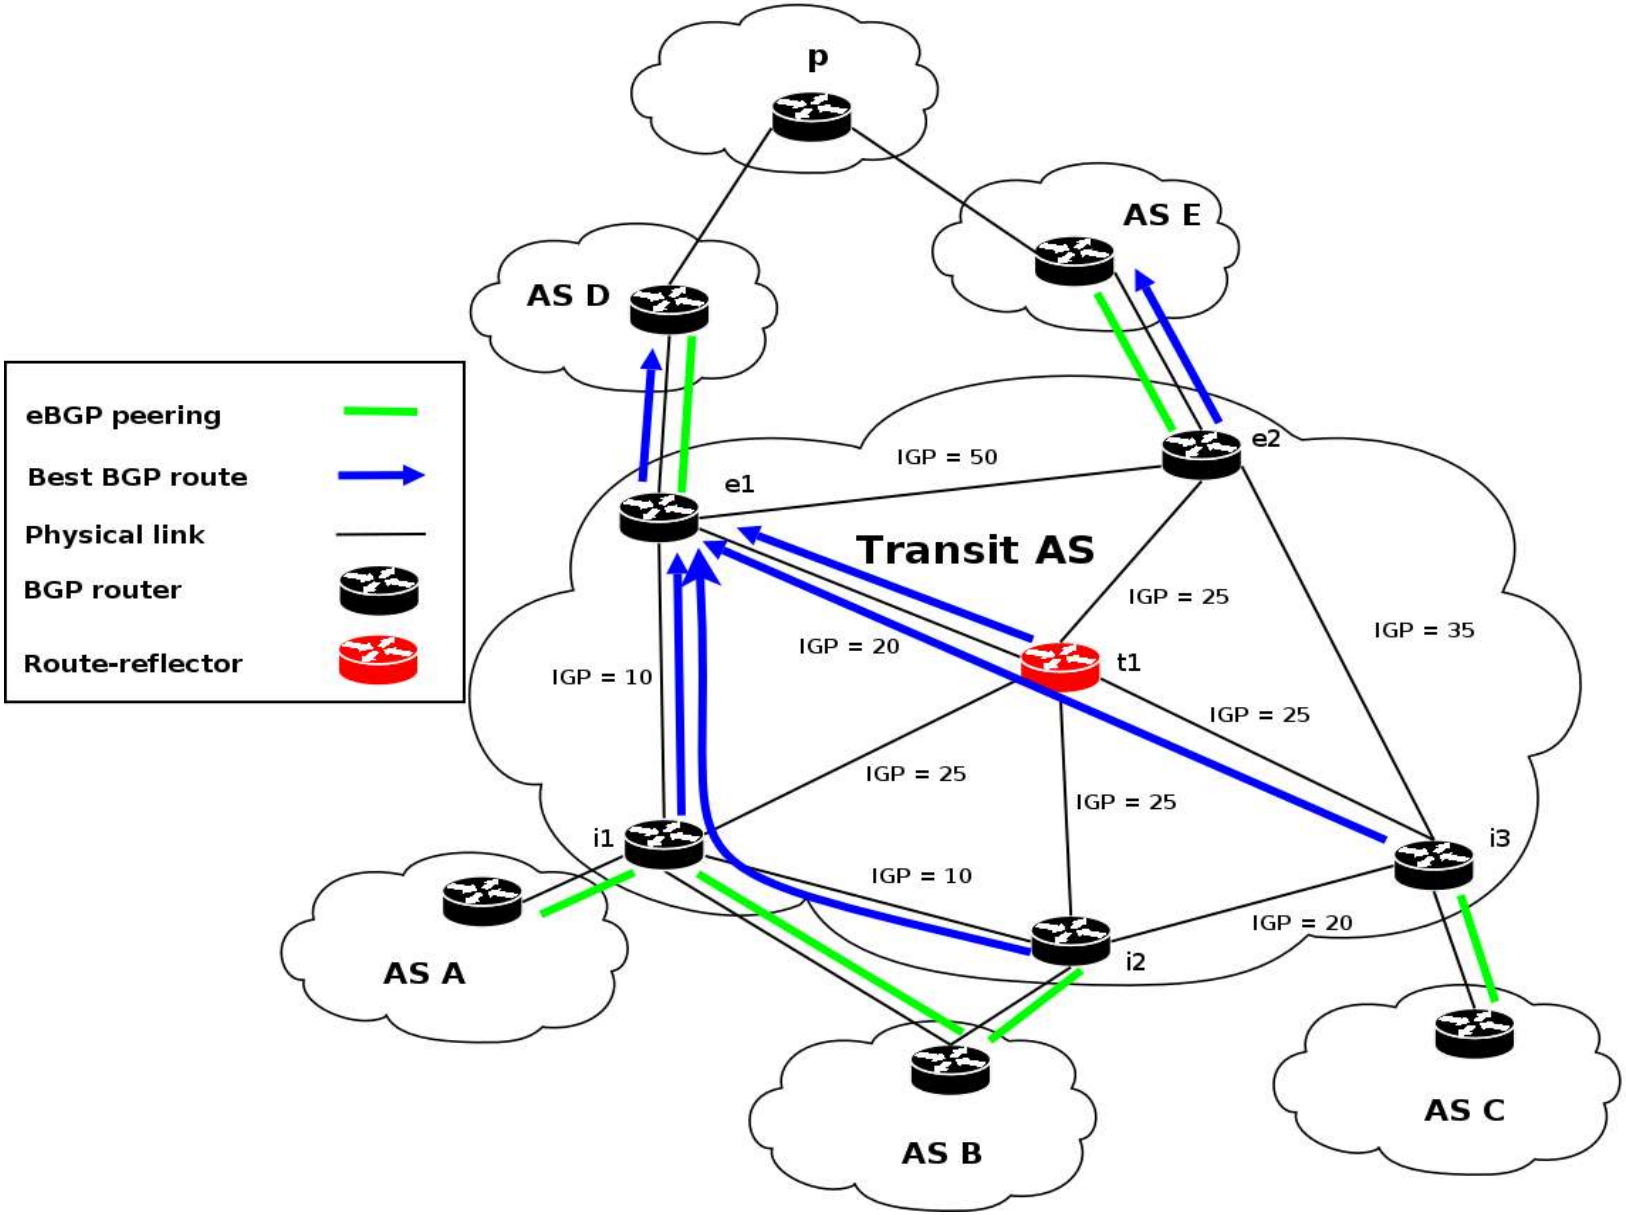

Dealing with complex ASs

- iBGP structure has non-trivial impact on best BGP routes choice inside AS
- RR initially introduced for scalability (limit number of iBGP sessions)
- RR performs tricky choice route selection and propagation inside large ASs
- No understanding today of how to design iBGP

Current and further work

- Need to understand what iBGP actually does
- Route-reflection is a necessity in large ASs
- Impact of RR unclear today, but it's not just about reducing number of iBGP sessions:
  - route diversity
  - sub-optimality of best routes
  - convergence in case of failures