

A critical view of the sensitivity of transit ASes to internal failures

Steve Uhlig

and

Sébastien Tandel



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

Agenda

- **Network robustness of transit ASes**
- **Sensitivity to internal failures**
- **Sensitivity of the GEANT network**
- **Route-reflection**

Designing robust transit ASes

- **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**

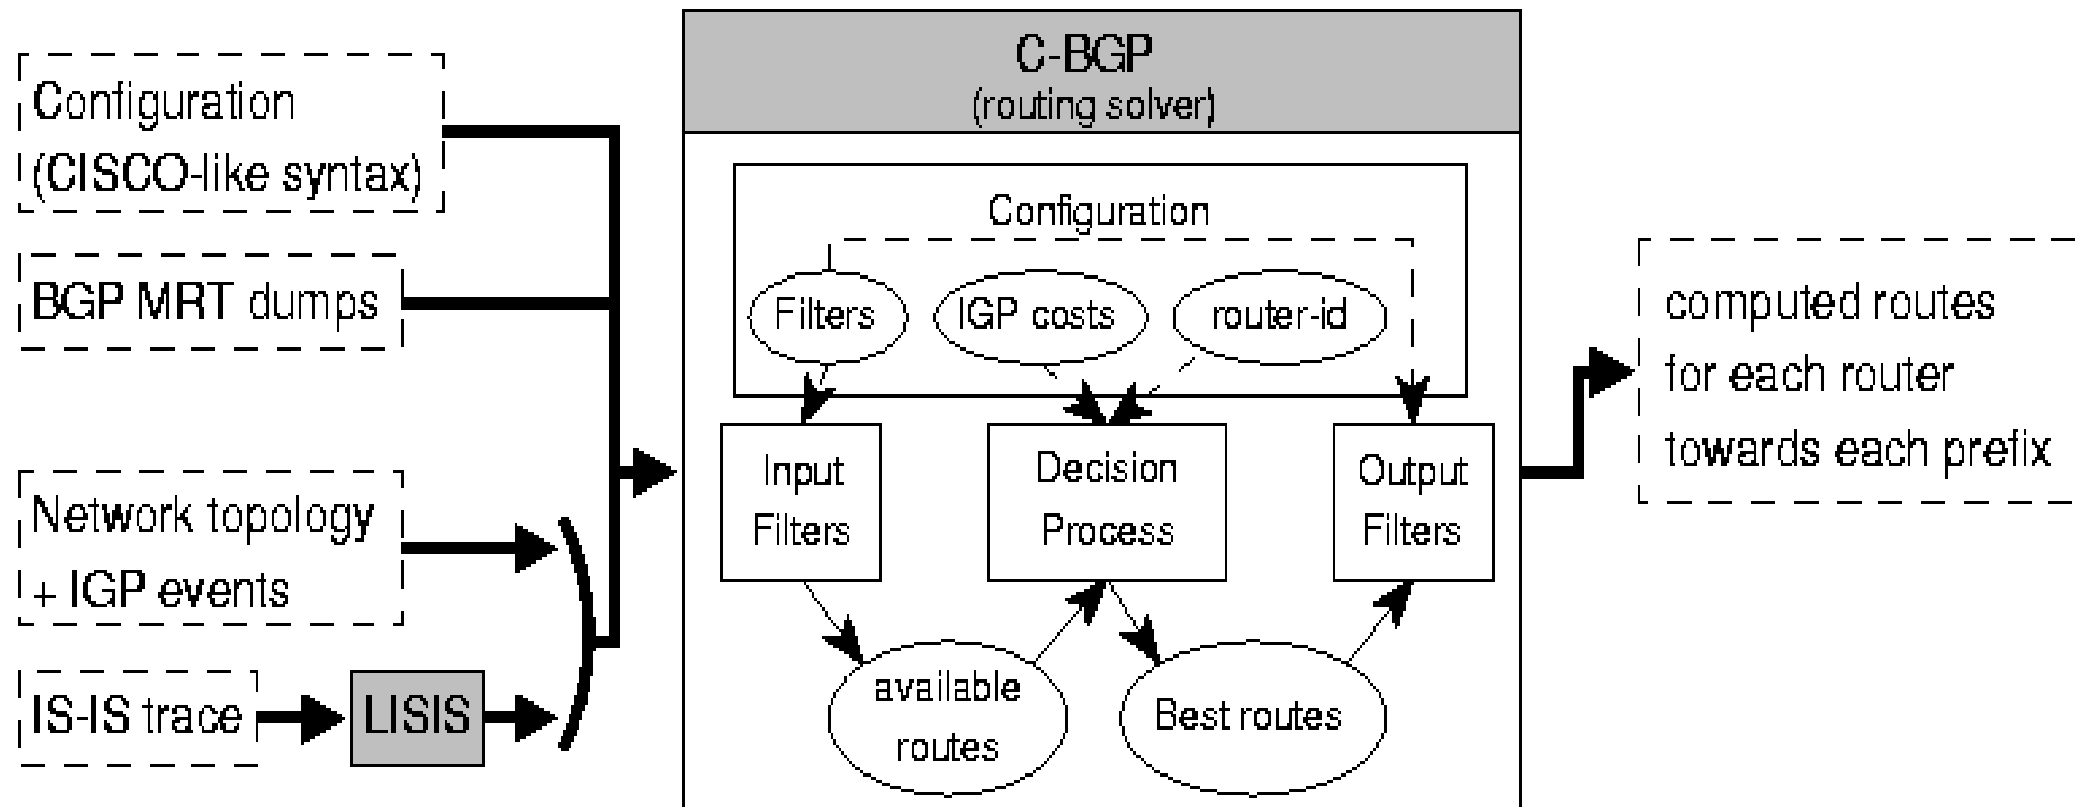
Sensitivity to internal failures

- **Sensitivity of a large tier-1 due to hot-potato disruptions [Teixeira et al., SIGMETRICS'04]**
- **Routing changes can have a significant impact on the traffic matrix [Teixeira et al., PAM2005]**
- **Metrics to capture sensitivity to internal failures [Teixeira et al., SIGCOMM'04]**

Reproducing the routing of transit ASes

- **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**

Reproducing the routing of transit ASes with CBGP



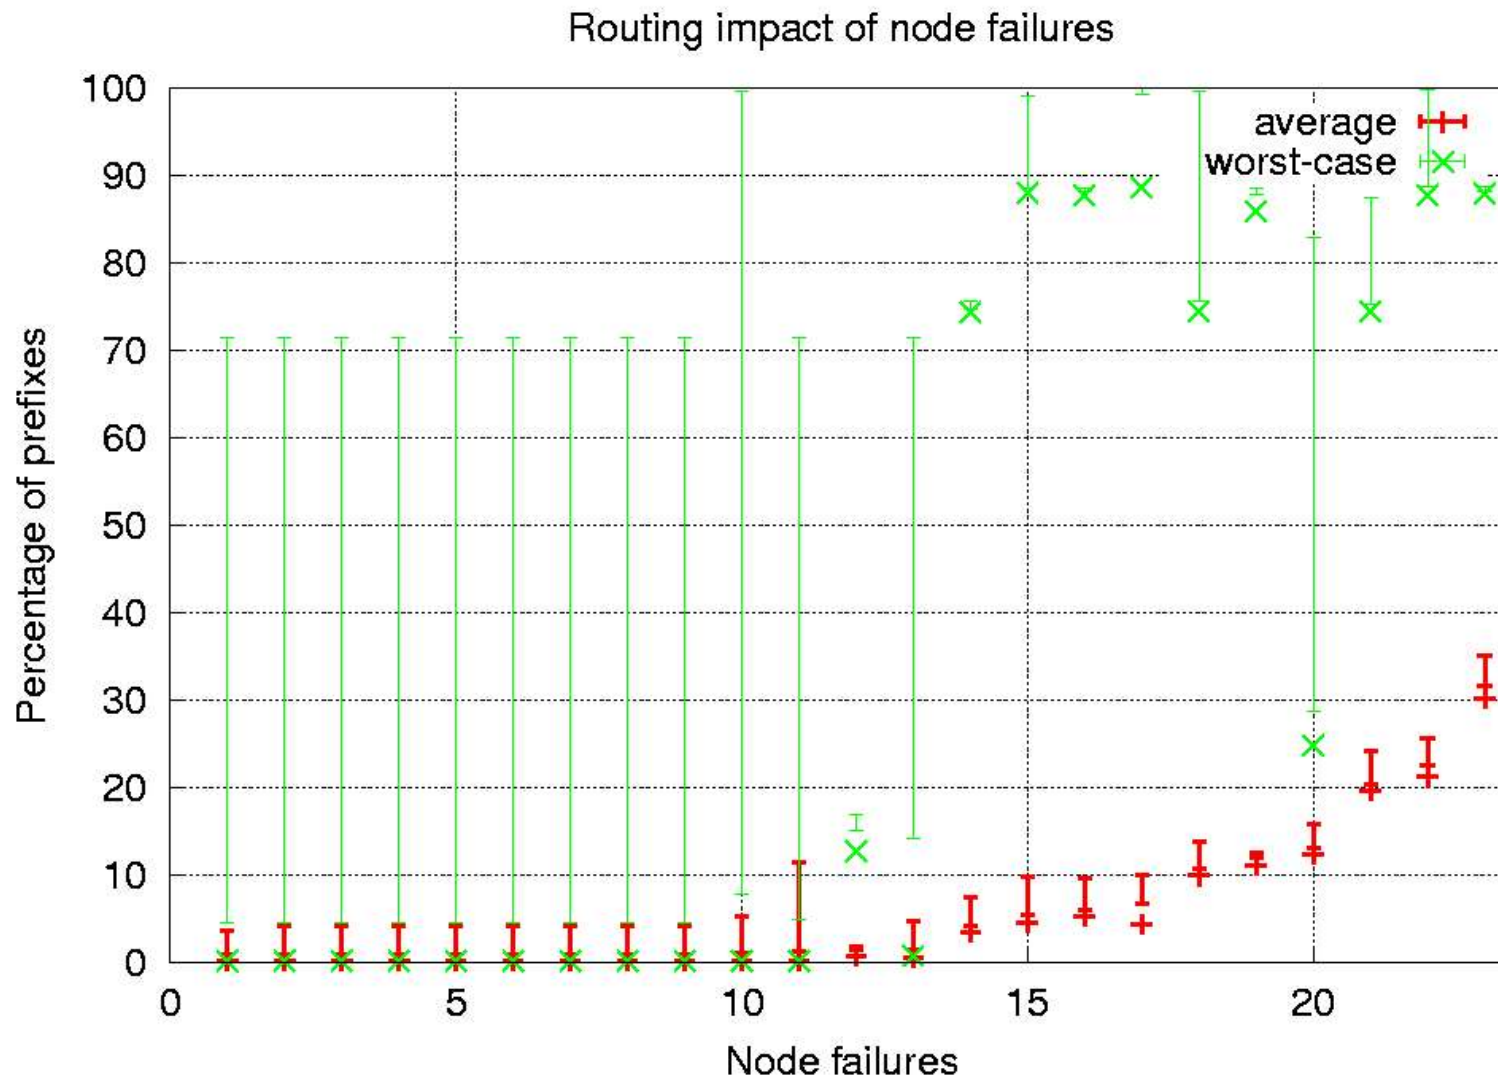
Sensitivity metrics to internal failures

- Consider a set of graph transformations on the graph G (denoted by ΔG).**
- 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).**

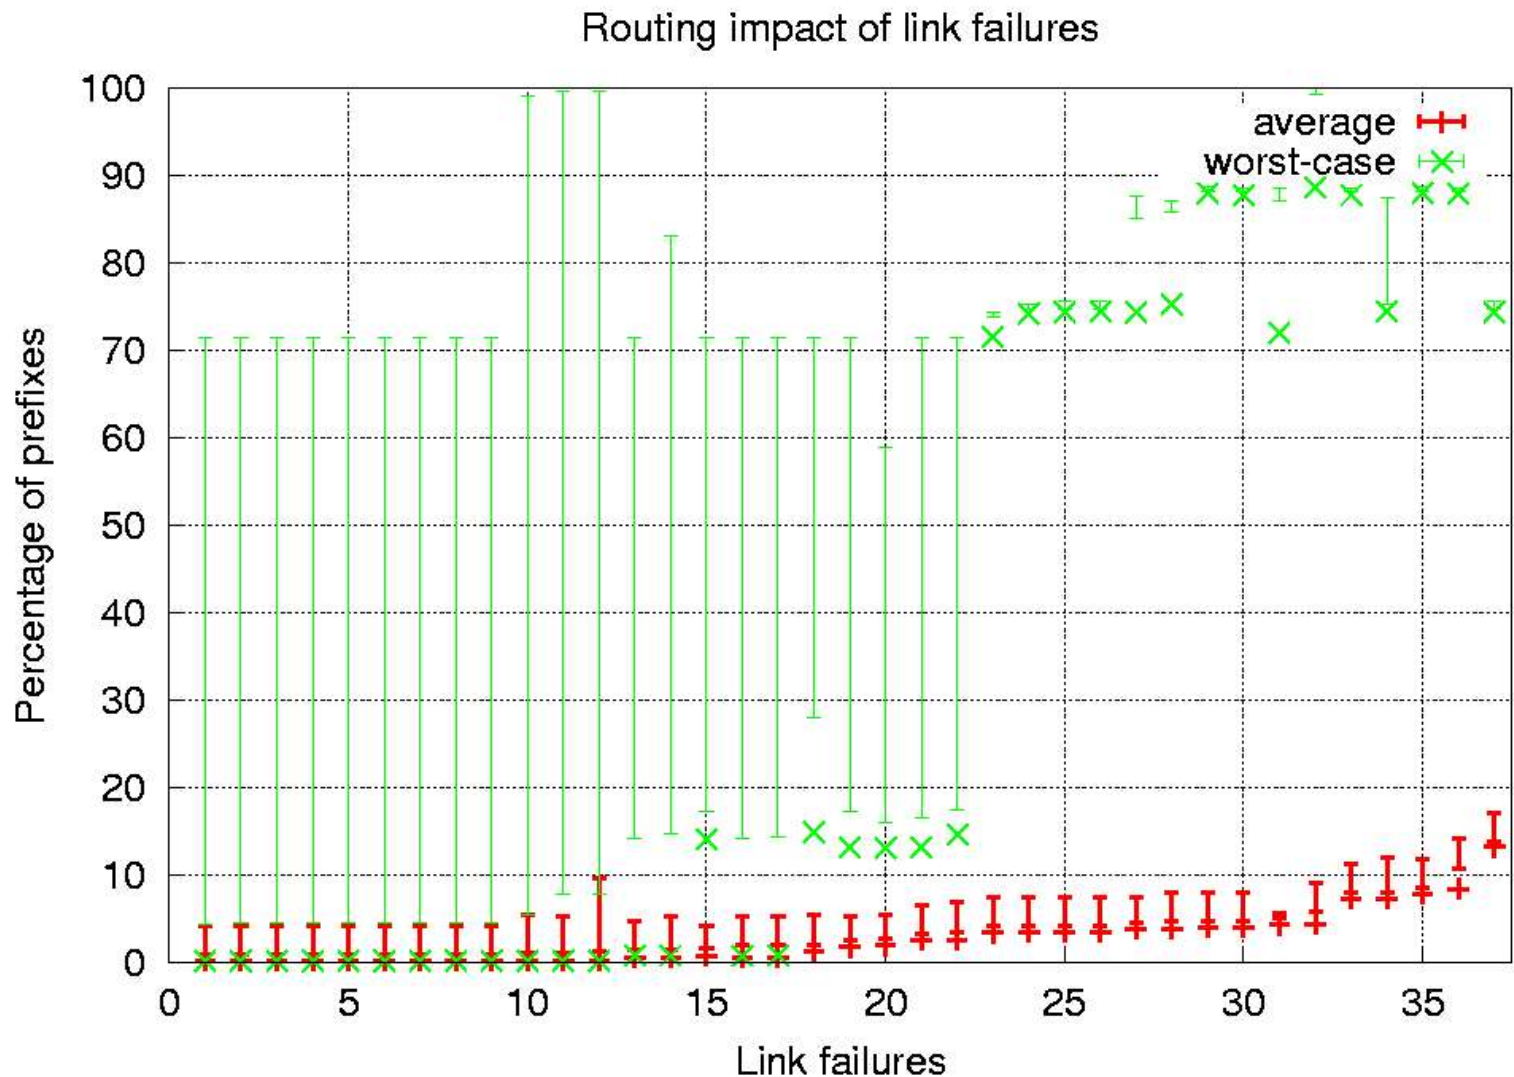
Data from the GEANT network

- 1 month period of routing and netflow data in 2004.
- The GEANT network:
 - 23 POPs, 76 links, tens of eBGP peerings
 - iBGP full mesh
- ISIS and BGP collected in Geneva


Sensitivity of the GEANT network



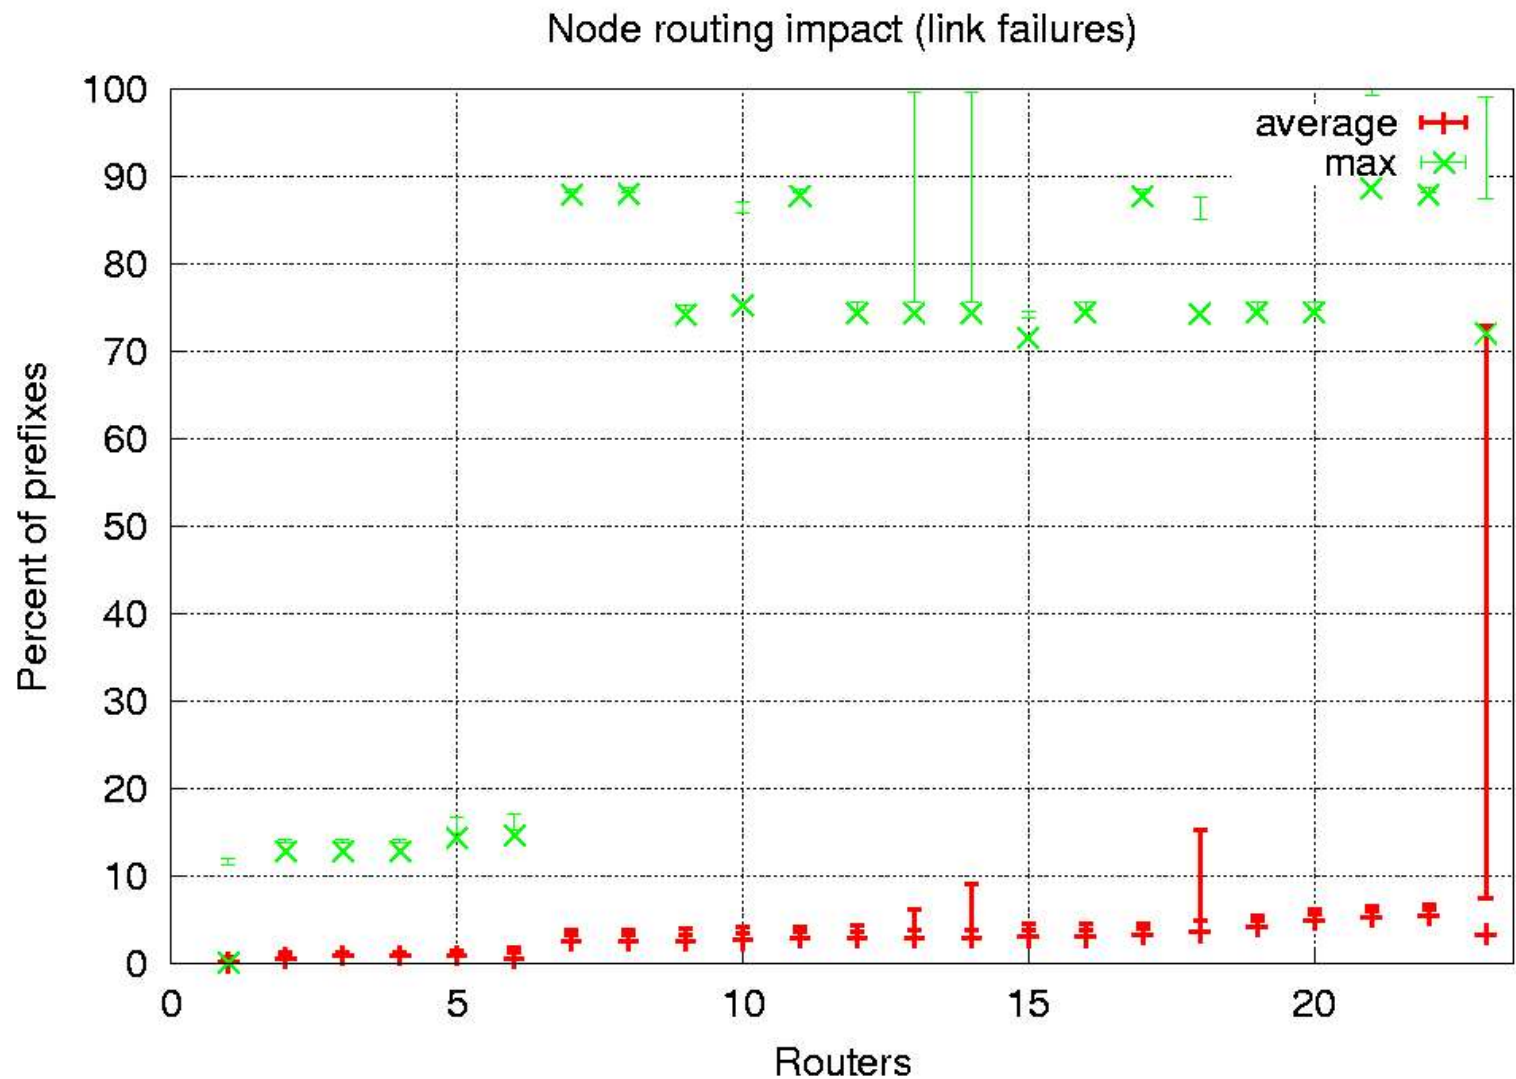
Sensitivity of the GEANT network



Sensitivity of the GEANT network




Sensitivity of the GEANT network



Insight from sensitivity analysis

- **Good insight of critical links and routers**
- **Highly connected to IGP cost design
(concentration of shortest IGP paths)**
- **Good enough for simple iBGP structures,
helps find out whether change network
topology**
- **What for complex transit ASes with
route-reflection ?**

Opacity of route-reflection



Dealing with complex ASes

- **iBGP structure has non-trivial impact on best BGP routes inside AS**
- **RRs was formerly introduced for scalability (limit number of iBGP sessions)**
- **Route-reflection performs tricky choice route selection and propagation inside large ASes**
- **No understanding today of link between iBGP design and impact on the network**

Conclusions

- **Need to understand impact of complex iBGP structures to understand network robustness**
- **Route-reflection is heavily used today in large networks, its impact goes however beyond reducing number of iBGP sessions**
- **Connections between IGP topology and BGP one need to be understood to design more robust networks**