



## **Impact of IS-IS on BGP** **case study on the Geant network**

**Bruno Quoitin**

([bqu@info.ucl.ac.be](mailto:bqu@info.ucl.ac.be))

Université Catholique de Louvain  
Computer Science and Engineering Dept.

with

**Olivier Bonaventure (UCL)**


**Sebastien Tandel (UCL)**

# Agenda

- ◆ **Motivations**
- ◆ **Model**
- ◆ **Case study: Geant**
  - ◆ Impact of Link/Router failures
  - ◆ Replay IS-IS trace
- ◆ **Conclusion and further work**

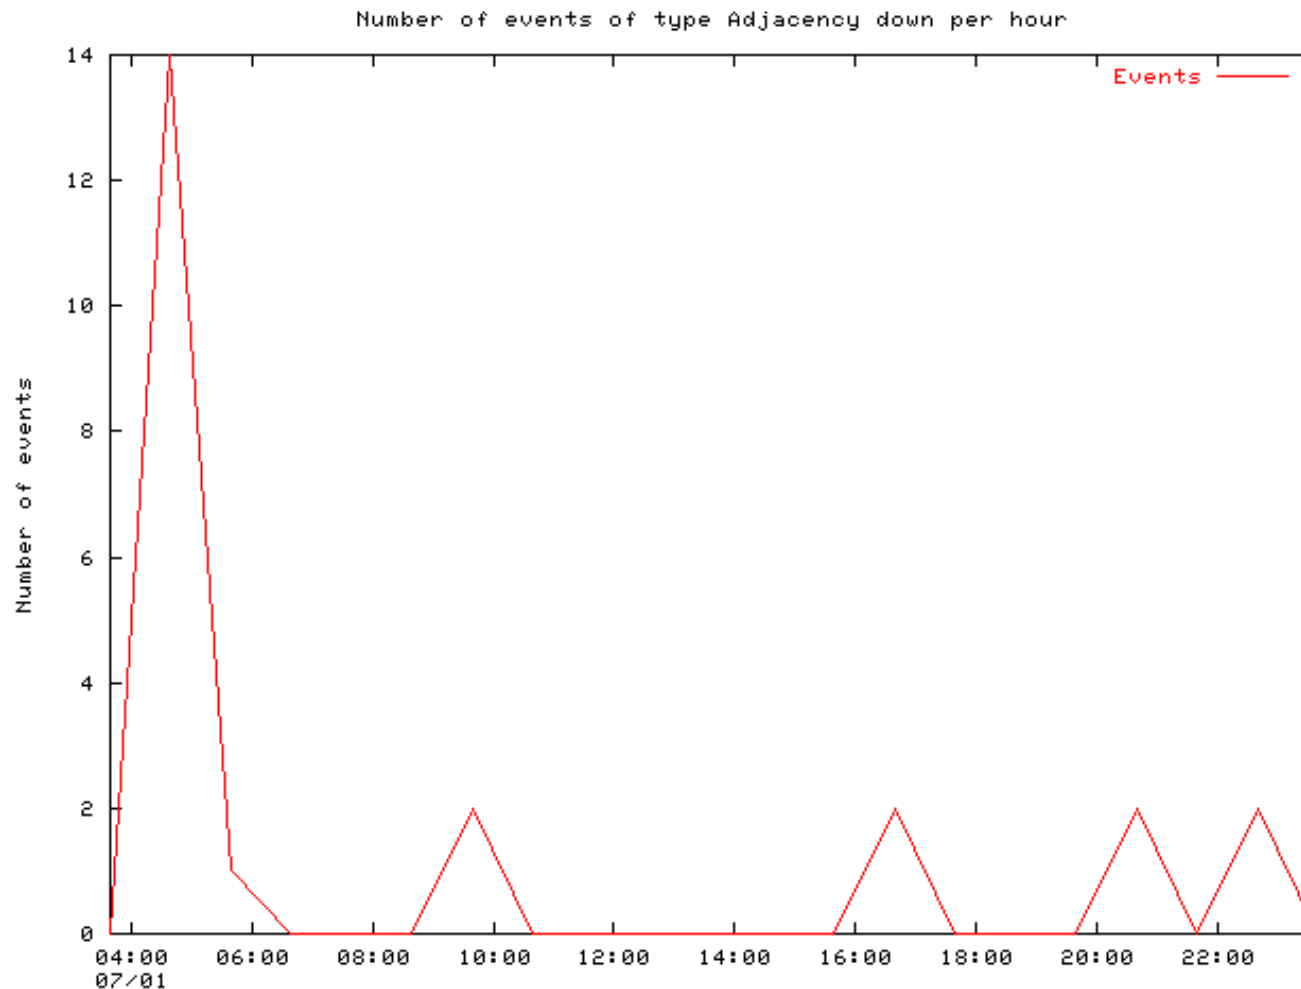
# Motivations (1)

## ◆ A closer look at GEANT's IS-IS events



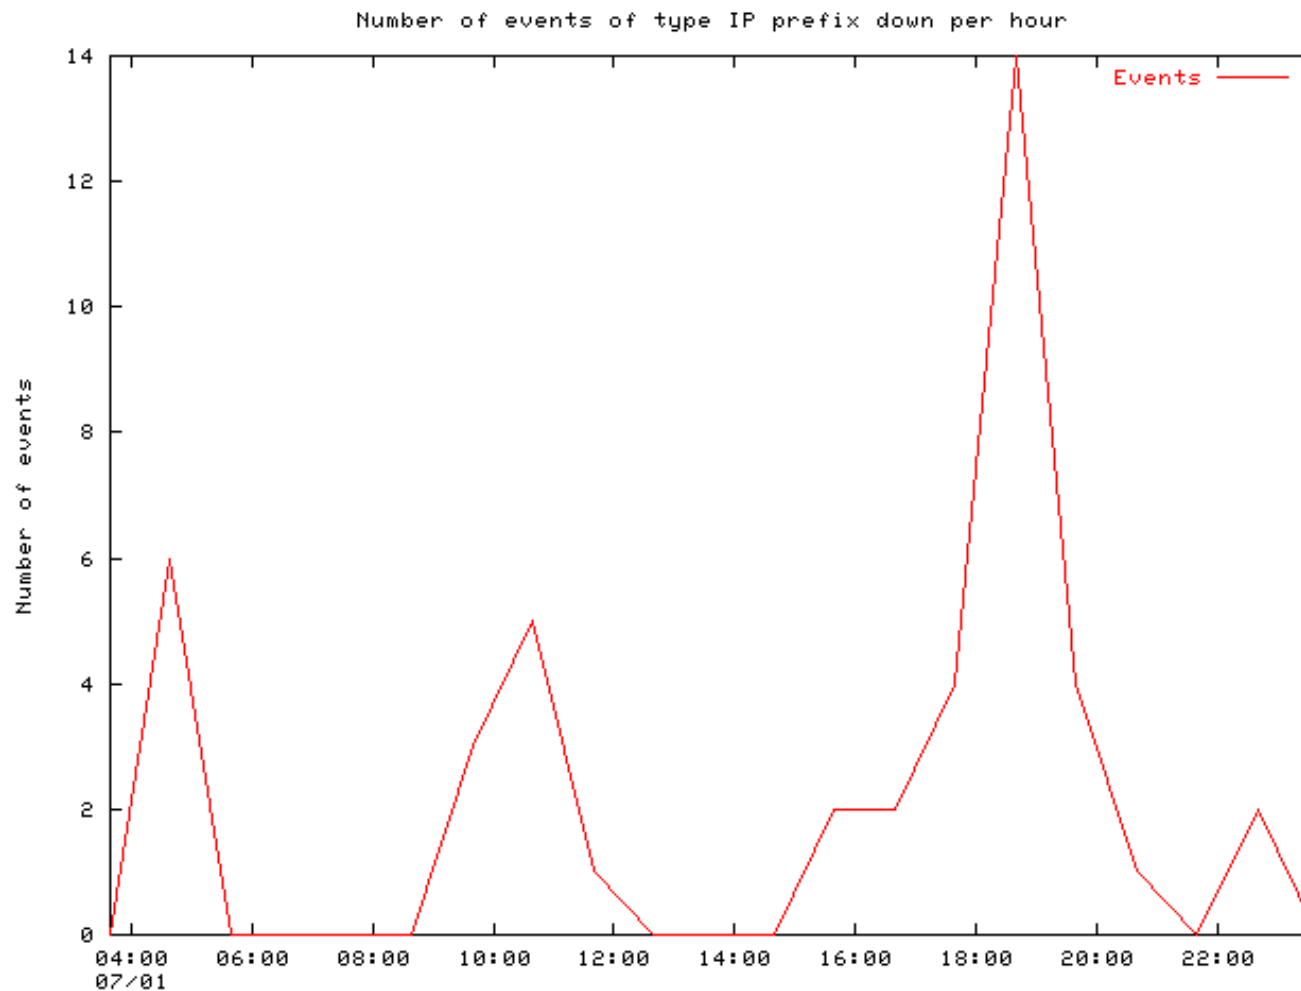
# Motivations (2)

## ◆ Adjacency down LSPs



# Motivations (3)

## ◆ IP Prefix Down LSPs



# Motivations (4)

- ◆ **Impact on BGP and Internet traffic**
  - ◆ A single IGP event
    - ◆ can trigger **BGP reconvergence**
    - ◆ can cause **BGP advertisements** outside the AS
    - ◆ can have a tremendous impact on the **distribution of interdomain traffic** inside the AS
  - ◆ Analytical study of the sensitivity of BGP to IGP events
    - ◆ simplified model: “*Network Sensitivity to Hot-potato Disruptions*”, R. Teixeira et al, SIGCOMM'04
    - ◆ **difficult** due to complex BGP policies, route-reflectors, and so on

# Another approach...



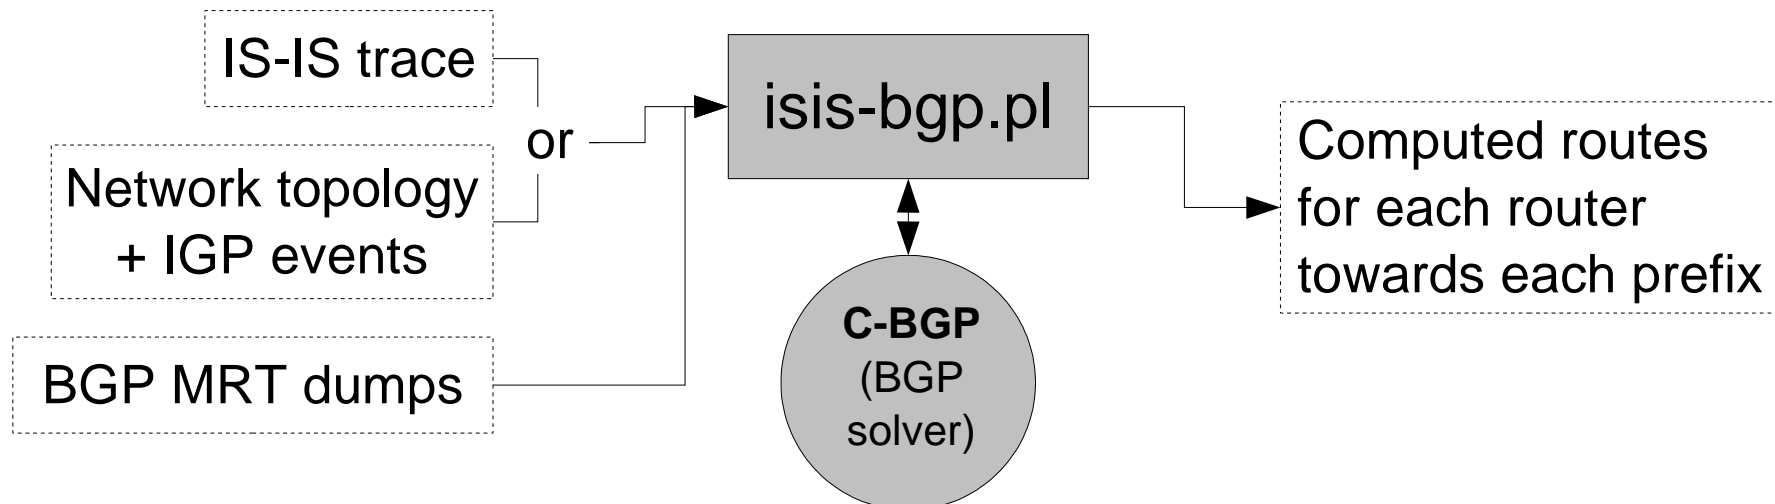
## ◆ Model

### ◆ Use BGP solver: C-BGP

- ◆ Supports complete BGP decision process, versatile policies, route-reflectors

### ◆ Requirements (inputs):

- ◆ Network topology (static or from IGP trace)
- ◆ External BGP routes



# BGP solver (1)



- ◆ **C-BGP: new BGP simulator**
  - ◆ Steady state
    - ◆ no protocol dynamics; only routing decisions
  - ◆ Efficiency
    - ◆ Not hindered by packet forwarding and simulation of TCP connections (like in SSFNet or J-Sim)
  - ◆ Large scale topologies
    - ◆ ~18k ASes, ~200k routers
  - ◆ Supports most of the real BGP features
    - ◆ all attributes, route-reflectors, versatile filters, ...
  - ◆ Public, open-source, tool (<http://cbgp.info.ucl.ac.be/>)

# BGP solver (2)



- ◆ **C-BGP: Interaction with real world**
  - ◆ Able to read BGP MRT dumps (RIB/messages)
  - ◆ Configuration syntax similar to CISCO
  - ◆ Dumps RIBs in CISCO/MRT format
- ◆ Possibility to convert CISCO/JunOS configs into C-BGP (work in progress)

# BGP solver (3)



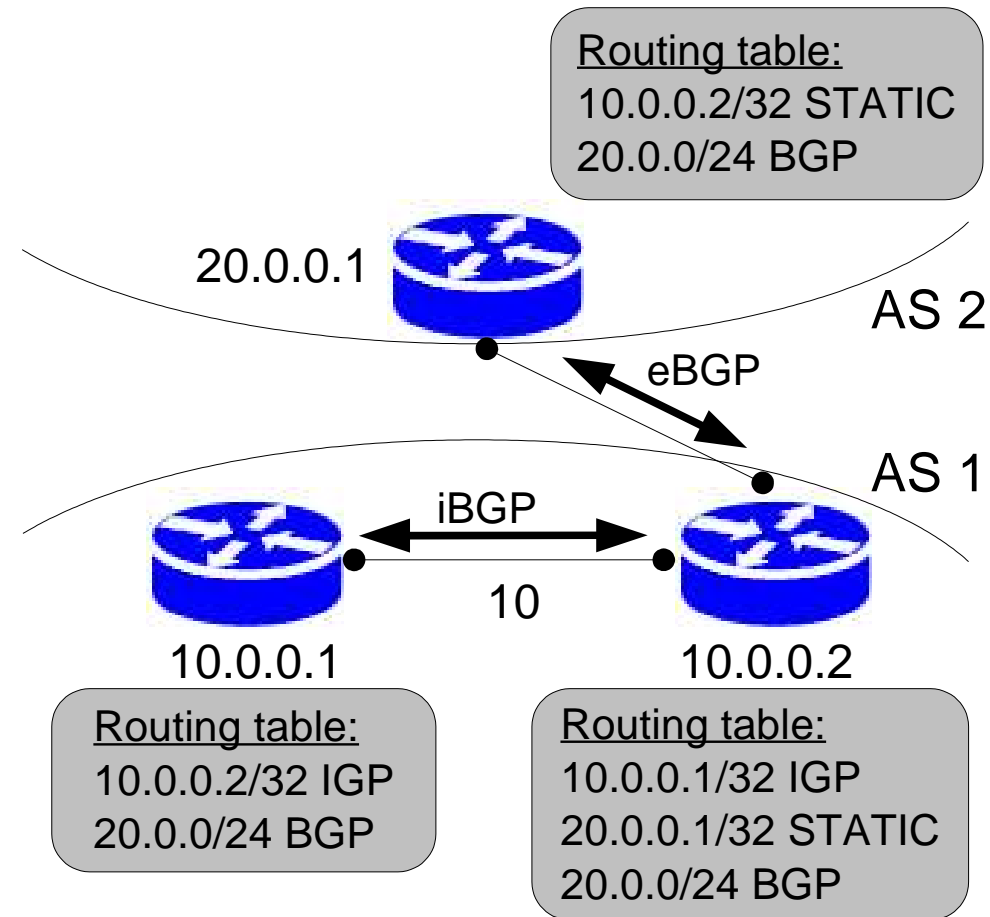
## Sample configuration

### Console

```
bgp router 10.0.0.2
  add network 10.0.0/24
  add peer 2 20.0.0.1
  peer 20.0.0.1 next-hop-self
  peer 20.0.0.1 up
```

```
bgp router 20.0.0.1
  add network 20.0.0/24
  add peer 1 10.0.0.2
  peer 10.0.0.2 up
```

sim run



# Methodology (1)

## ◆ Methodology

- ◆ Build accurate representation of the network in C-BGP
  - ◆ compute IGP routes
  - ◆ let BGP converge
- ◆ For each link/router failure or metric change
  - ◆ update IGP information in C-BGP
  - ◆ recompute IGP routes
  - ◆ let BGP reconverge
- ◆ Measure impact on chosen routes

# Methodology (2)

## ◆ Metric

- ◆ Count the number of route changes
- ◆ For each router and prefix, classify as
  - ◆ Prefix has become up/down
  - ◆ Next-hop AS has changed
  - ◆ Egress router has changed
  - ◆ Path towards egress router has changed
  - ◆ IGP cost of path towards egress router has changed
  - ◆ No change

# Case study (1)

## ◆ **GEANT network**

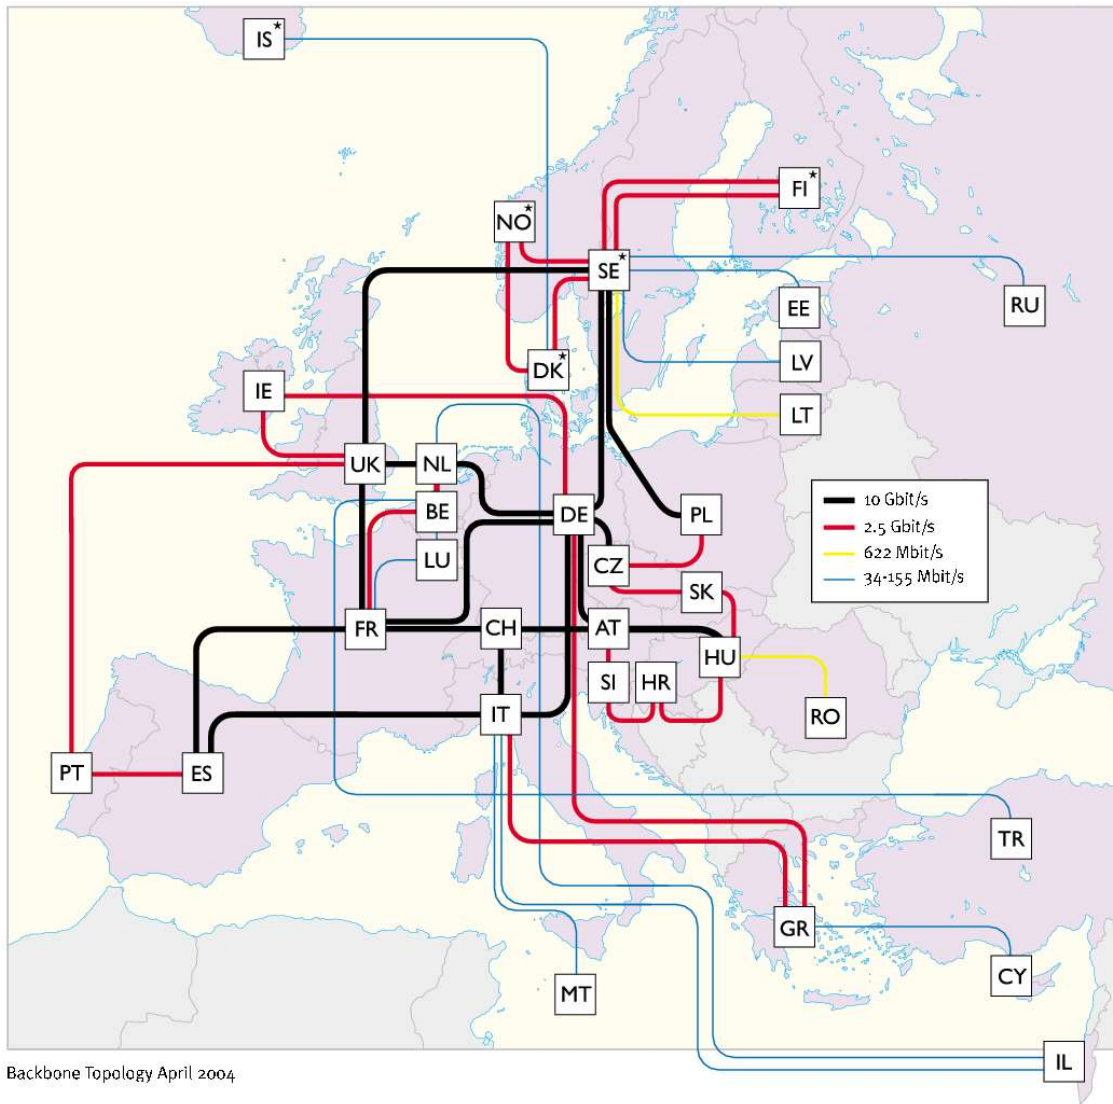
- ◆ European Research Transit network
- ◆ Peers with Abilene, NRENs and two commercial providers: GBLX and Telia

## ◆ **Data (collected on Geneva router)**

- ◆ IS-IS trace (in PyRT format)
- ◆ External BGP routes
- ◆ Period: July 2004
- ◆ Provided by T. Griffin & R. Gass (Intel Research)
- ◆ Policies are missing :-)

# Case study (2)


**GEANT network**  
(simplified)



# Case study (3)

- ◆ **What-if scenario: single link failures**
  - ◆ Impact on all routers


All routers



# Case study (4)

- ◆ **What-if scenario: single router failures**
  - ◆ Impact on all routers

All routers





# Replay IS-IS trace (1)

## ◆ Methodology

- ◆ 1). parse using LISIS (<http://totem.info.ucl.ac.be/>)
- ◆ 2). update the current state of the topology,
- ◆ 3). each 10 seconds (trace time), build a snapshot
- ◆ 4). update the IGP configuration of C-BGP
- ◆ 5). recompute IGP routes
- ◆ 6). recompute BGP routes
- ◆ 7). count BGP events between previous snapshot and current snapshot

# Replay IS-IS trace (2)

Replay IS-IS trace  
July 1<sup>st</sup> 04



# Further work

## ◆ Further work

### ◆ Impact on data plane

- ◆ We have all the paths, let's weight them with the traffic volume
- ◆ Geant Netflow statistics and BGP configurations should be released soon :-)
- ◆ All information already available for Abilene (<http://abilene.internet2.edu/observatory>)

# Conclusion

## ◆ Conclusion

- ◆ New approach to study the network sensitivity to hot-potato
- ◆ Supports BGP policies and iBGP hierarchy
- ◆ Open-source tools
- ◆ Can read MRT dumps, IS-IS traces and soon existing router configurations (CISCO/JunOS)

- ◆ If you are interested in testing these tools on your network, have a look at

<http://totem.info.ucl.ac.be/tools.html>

