Evaluating path diversity in the Internet: from an AS-level to a PoP-level granularity

Evaluation de la diversité de chemins sur Internet: d'une granularité au niveau des AS à une vision au niveau des points de présence

A.Fressancourt, C. Pelsser, M. Gagnaire

Journées non thématiques RESCOM

12-01-2016

Context

The Kumori architecture

- Architecture aiming at **enhancing the resiliency** of inter-DC communication
- Use of overlay of nodes located at **IXPs**
- Resiliency by routing around failure through IXP nodes

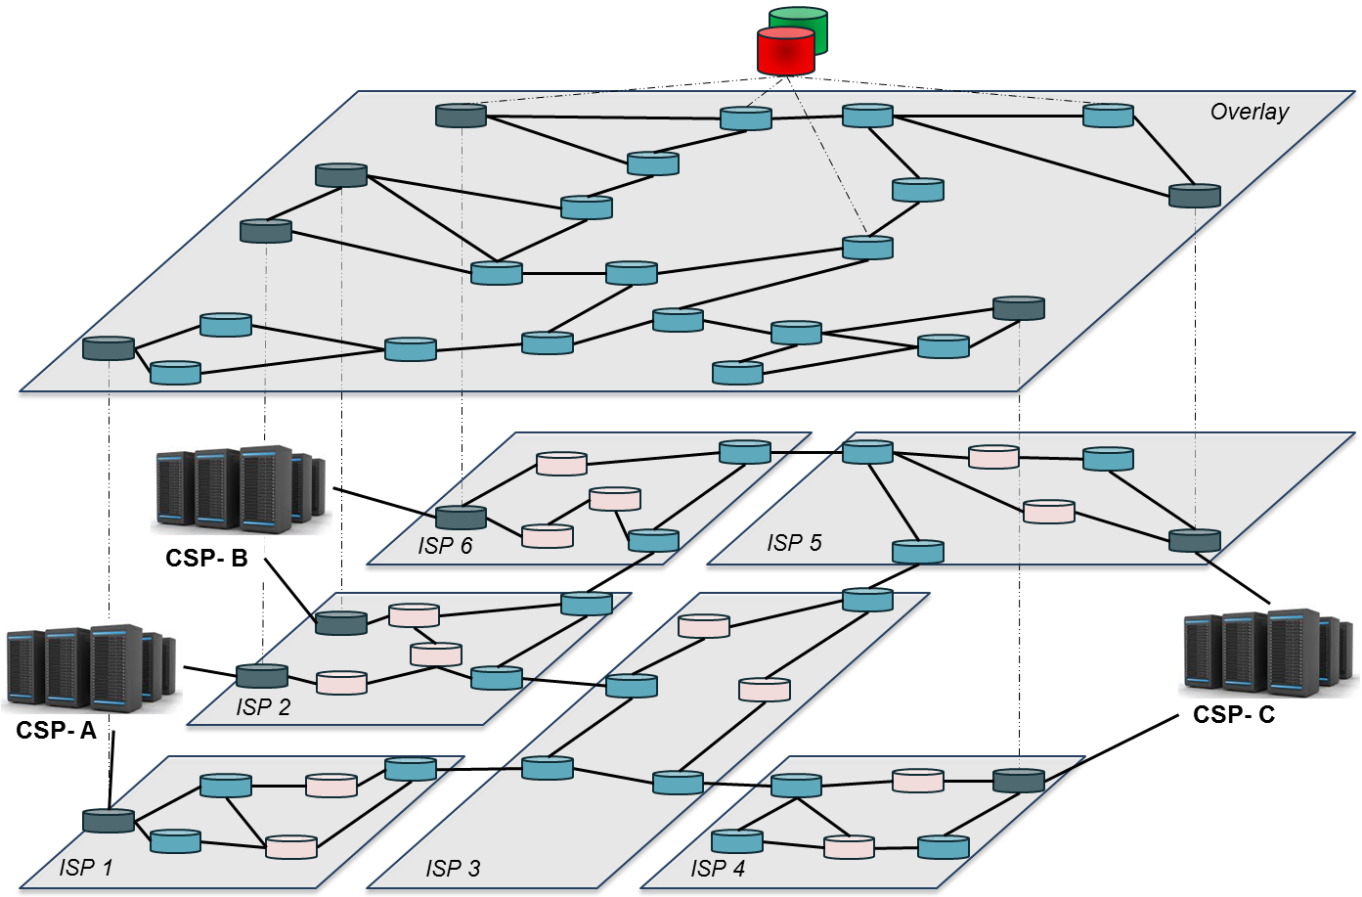

IIJ INNOVATION INSTITUTE TELECOM ParisTech worldline e-payment services

Initial evaluation of the Kumori architecture

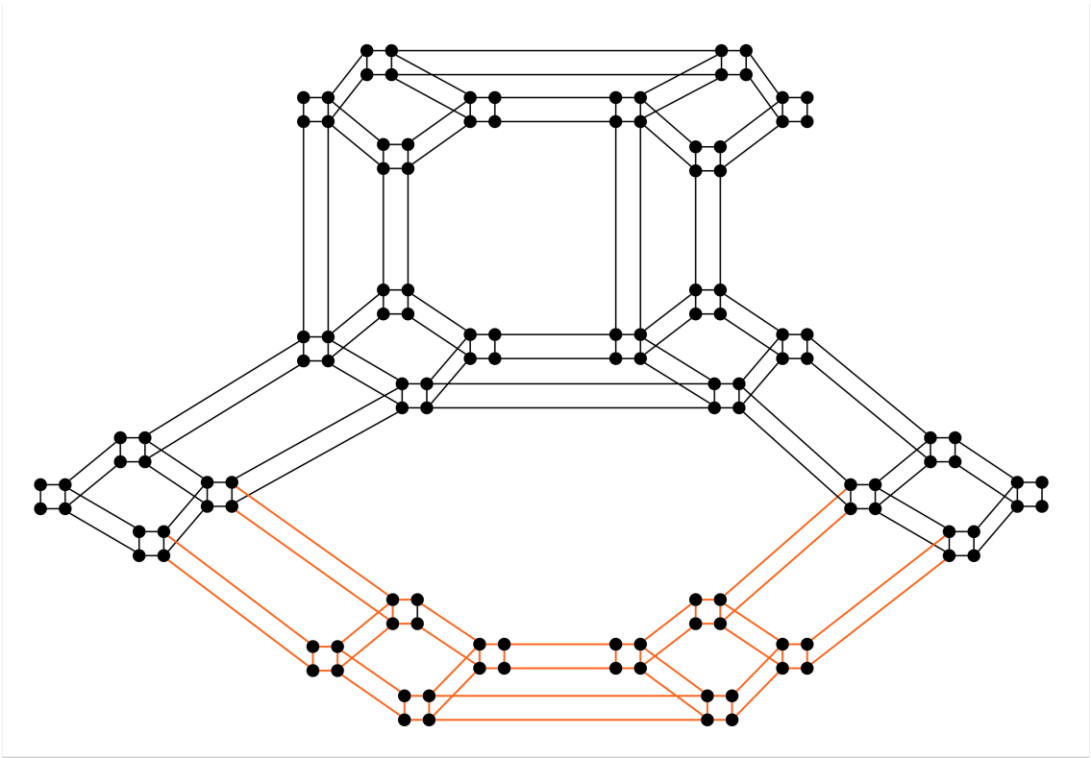

- Goal:
  - Comparison with RON (Andersen et al., *Resilient Overlay Networks*, SOSP 2001)
  - Evaluation of path diversity and path length
- iPlane measurement dataset
  - ➔ Router-level measurements of latency and packet loss on paths
  - ➔ Traceroutes between multiple vantage points in the Internet
- Undirected graph built
  - ➔ But the Internet is a **directed graph**

Madhyastha, Isdal et al., *iPlane: An Information Plane for Distributed Services, OSDI 2006*, November 2006

IIJ INNOVATION INSTITUTE TELECOM ParisTech worldline e-payment services

Building a PoP-level model of the Internet

Goal

- Build a **realistic graph** model of the Internet at PoP level

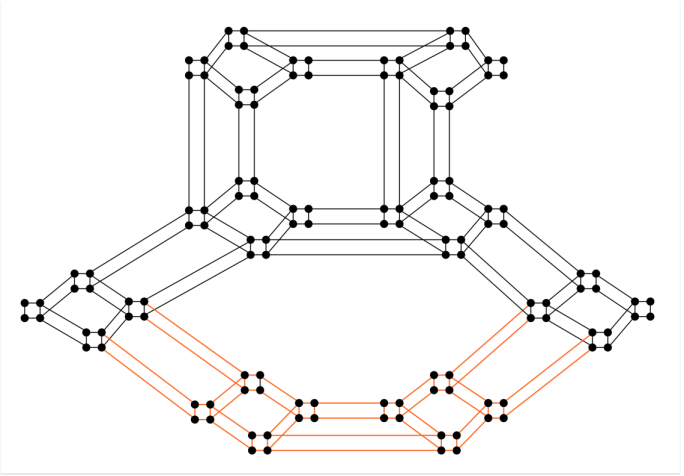

- Build a **directed graph** to render the client-provider relationships in the Internet

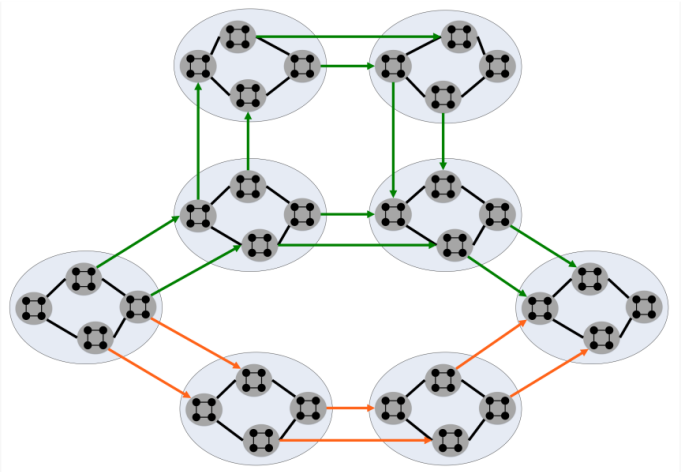

- Locate IXPs in the topology
- Share the topology with the research community

Data sources used

- **iPlane dataset:**
  - Madhyastha, Isdal *et al.*, *iPlane: An Information Plane for Distributed Services*, OSDI 2006, November 2006
  - Latency and packet loss measurements between routers
  - ➔ *Inter router, inter-PoP, inter AS links measurements, IP addresses of routers*

- **CAIDA topology variant:**
  - Sobrinho, Vanbever *et al., Distributed Route Aggregation on the Global Network*, ACM CoNEXT 2014, December 2014
  - Client – Provider hierarchy between AS
  - ➔ *Inter-AS link characterization*

- **PeeringDB:**
  - http://www.peeringdb.com
  - Provides information on IXP membership
  - ➔ *IXP IP addresses range, IXP membership*

IIJ INNOVATION INSTITUTE TELECOM ParisTech worldline e-payment services

The iPlane dataset

- Dataset characteristics
  - Summary of **traceroutes** done between the 13th of June and the 14th of July
  - We identified:
    - **417 638 routers**
    - **7 687 300 directed edges** between those routers

- Some **errors**
  - The **« AS 3303 » issue**
    - AS 3303 ➔ Swisscom
    - Roughly 8% of the routers in iPlane ➔ mistake
    - AS / router association from accurate data (iPlane's corrected association + Hurricane Electric)
  - Some **AS-prefix associations were unclear** ➔ Use of Hurricane Electric's data on those specific prefixes

Building our topology, step by step

***Step 1**: Extracting relevant data from our sources*

- **Goal:**
  - Build easily exploitable data structures for the next steps

- **Work done:**
  - Fix IP prefixes / AS matching relationship using Hurricane Electric's data
  - Build the router / AS association dictionary
  - Derive edges between routers
  - Using data provided by PeeringDB, associate routers with IXPs, and retrieve AS membership from router's presence at IXPs

➔ Data mixing iPlane and PeeringDB about relationship between routers

➔ Incomplete information on router's presence at each IXP.

IIJ INNOVATION INSTITUTE

TELECOM ParisTech

worldline e-payment services

Building our topology, step by step

***Step 2**: Inferring router IXP membership*

- **Goal:**
  - Associate routers with IXPs
  - Build missing inter-AS links at IXP according to *Anatomy of a Large European IXP* (Ager *et al.*)
- **Work done:**
  - First knowledge of router presence at IXP provided by PeeringDB
  - We first look at routers which are linked with a router at an IXP and that belong to a missing member AS
  - Then, for each IXP, we look at the missing Ases and we compare those AS' routers location with IXP locations using MaxMind's GeoIP database.

➔ Enhanced router / IXP association dictionary

➔ AS policy dictionary

Building our topology, step by step

***Step 3**: Building router clusters*

- **Goal:**
  - Reduce the size of the directed graph built from iPlane's view of the Internet
- **Work done:**
  - Clustering based on edge betweenness ➔ Very lengthy computations
  - We compared several methods to cluster routers together
    - Comparison with 2 ground truths: GEANT (30 PoPs) and Amazon (53 PoPs)
    - 2 algorithms give results that are close to the ground truth:
      - **Infomap**: Rosvall *et al.*, *Maps of information flow reveal community structure in complex networks*, 2008
      - **Walktrap**: Pons *et al.*, *Computing communities in large networks using random walks*, 2005

  ➔ Result:
    - 417 638 routers ➔ **148823 PoPs**
    - 7 687 300 directed edges ➔ **1 040 740 directed edges**

Building our topology, step by step

***Step 4**: Building a cluster-level directed graph*

- **Goal:**
  - Building a graph allowing us to evaluate path diversity between PoPs in the Internet
- **Work done:**
  - We retrieved AS relationships from the DRAGON topology
  - From the router links data file, we build inter-cluster directed links, and associate them with mean latency.

➔ Cluster-level directed graph

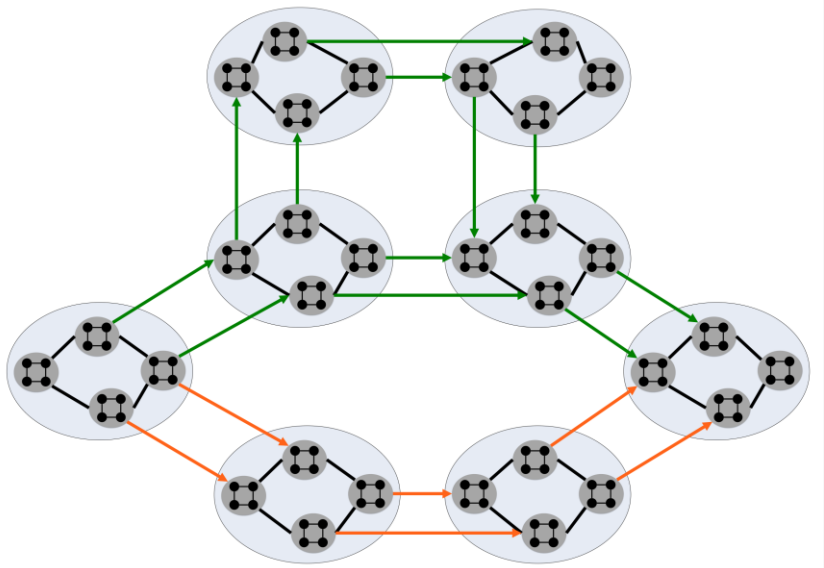

Evaluating path diversity at PoP-level

Customer / provider routing policy

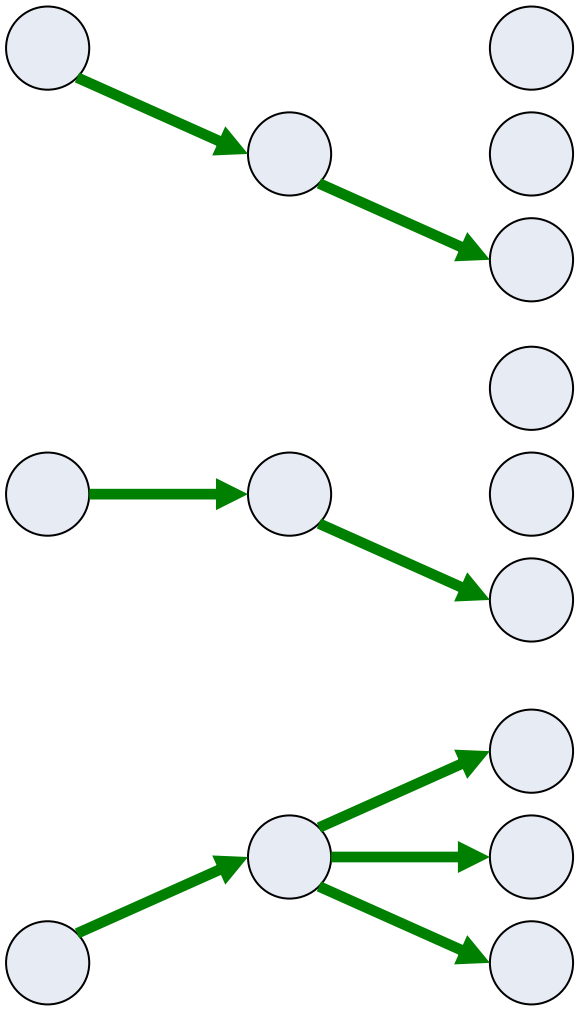

- For a packet to go from one Autonomous System (AS) to another, the later AS has to announce his prefixes to the first AS

- Gao and Rexford ➔ Basic prefix advertisment policy *(sometimes inaccurate)*
  - Customers announce their prefixes, and their customers' prefixes to their providers and peers.
  - Peers announce their prefixes and their customer's prefixes to one another.
  - Providers announce all the prefixes they know to their customer

Evaluating path diversity at PoP-level

***Step 1**: Evaluating AS-level path diversity for CSP cluster*

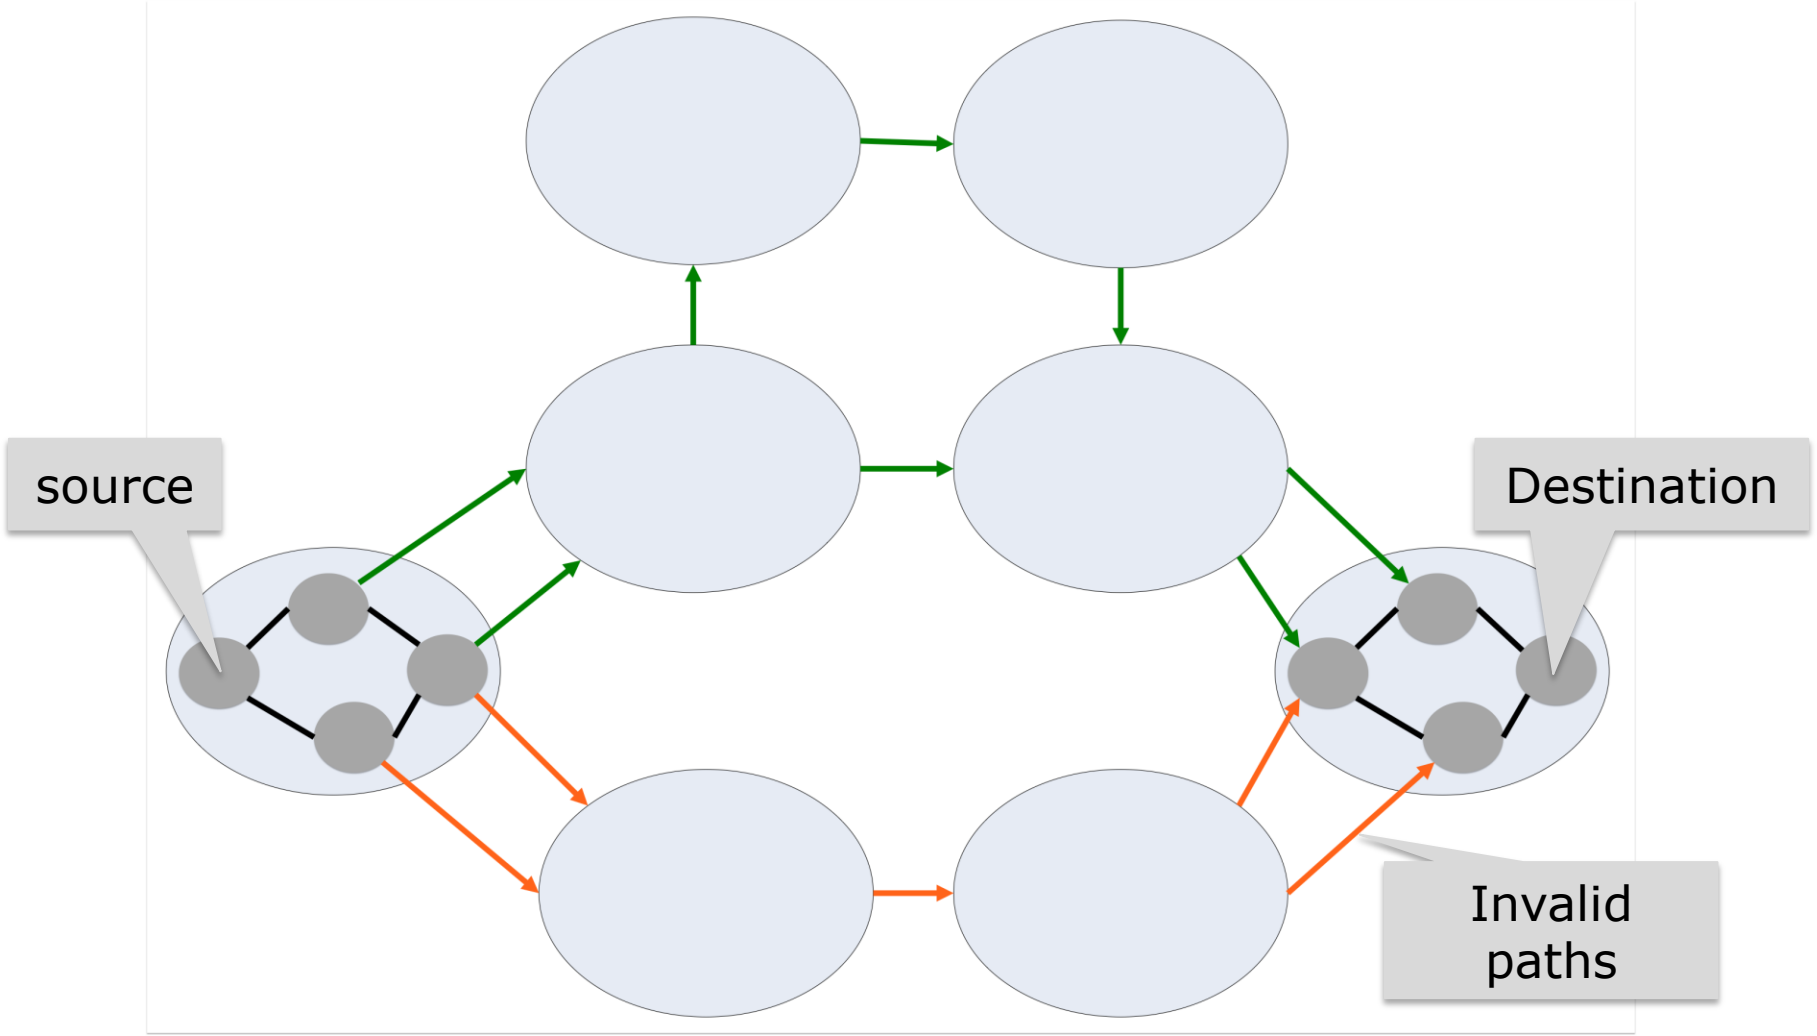

IIJ INNOVATION INSTITUTE TELECOM ParisTech worldline e-payment services

Evaluating path diversity at PoP-level

***Step 1**: Evaluating AS-level path diversity for CSP PoPs*

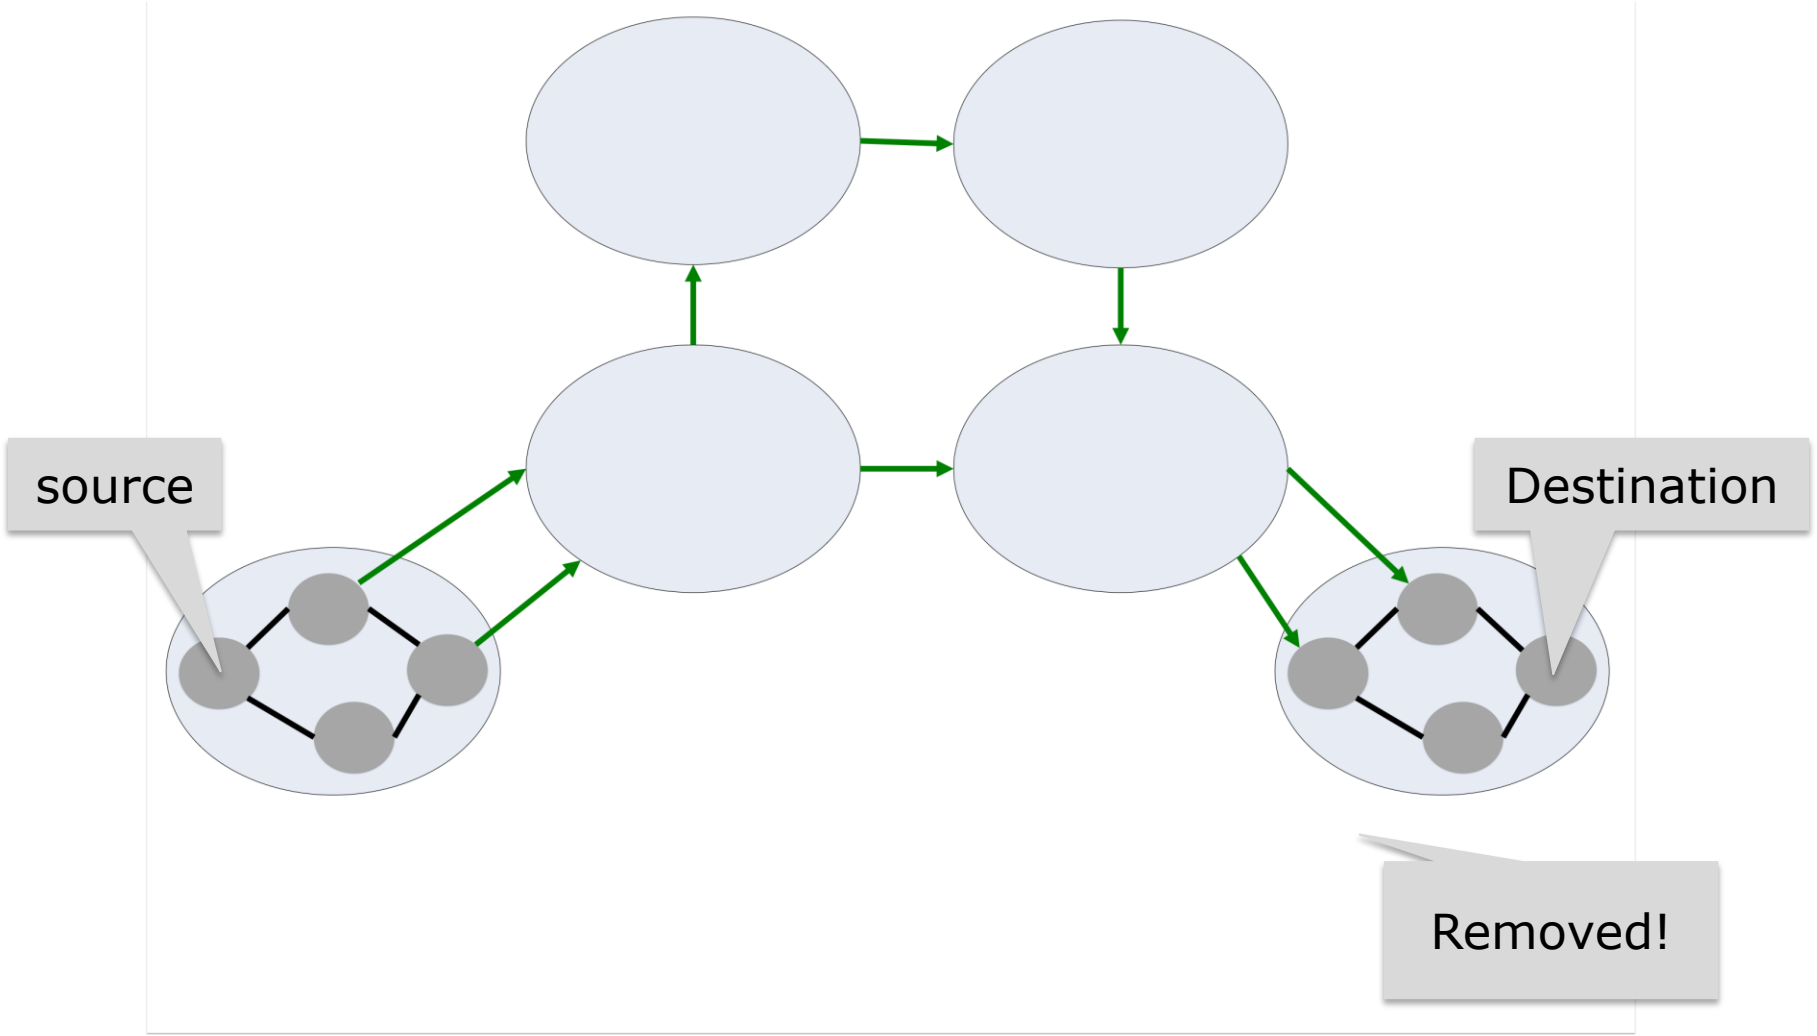

IIJ INNOVATION INSTITUTE
TELECOM ParisTech
worldline e-payment services

Evaluating path diversity at PoP-level

***Step 2**: Building PoP-level directed graph for valid ASes*

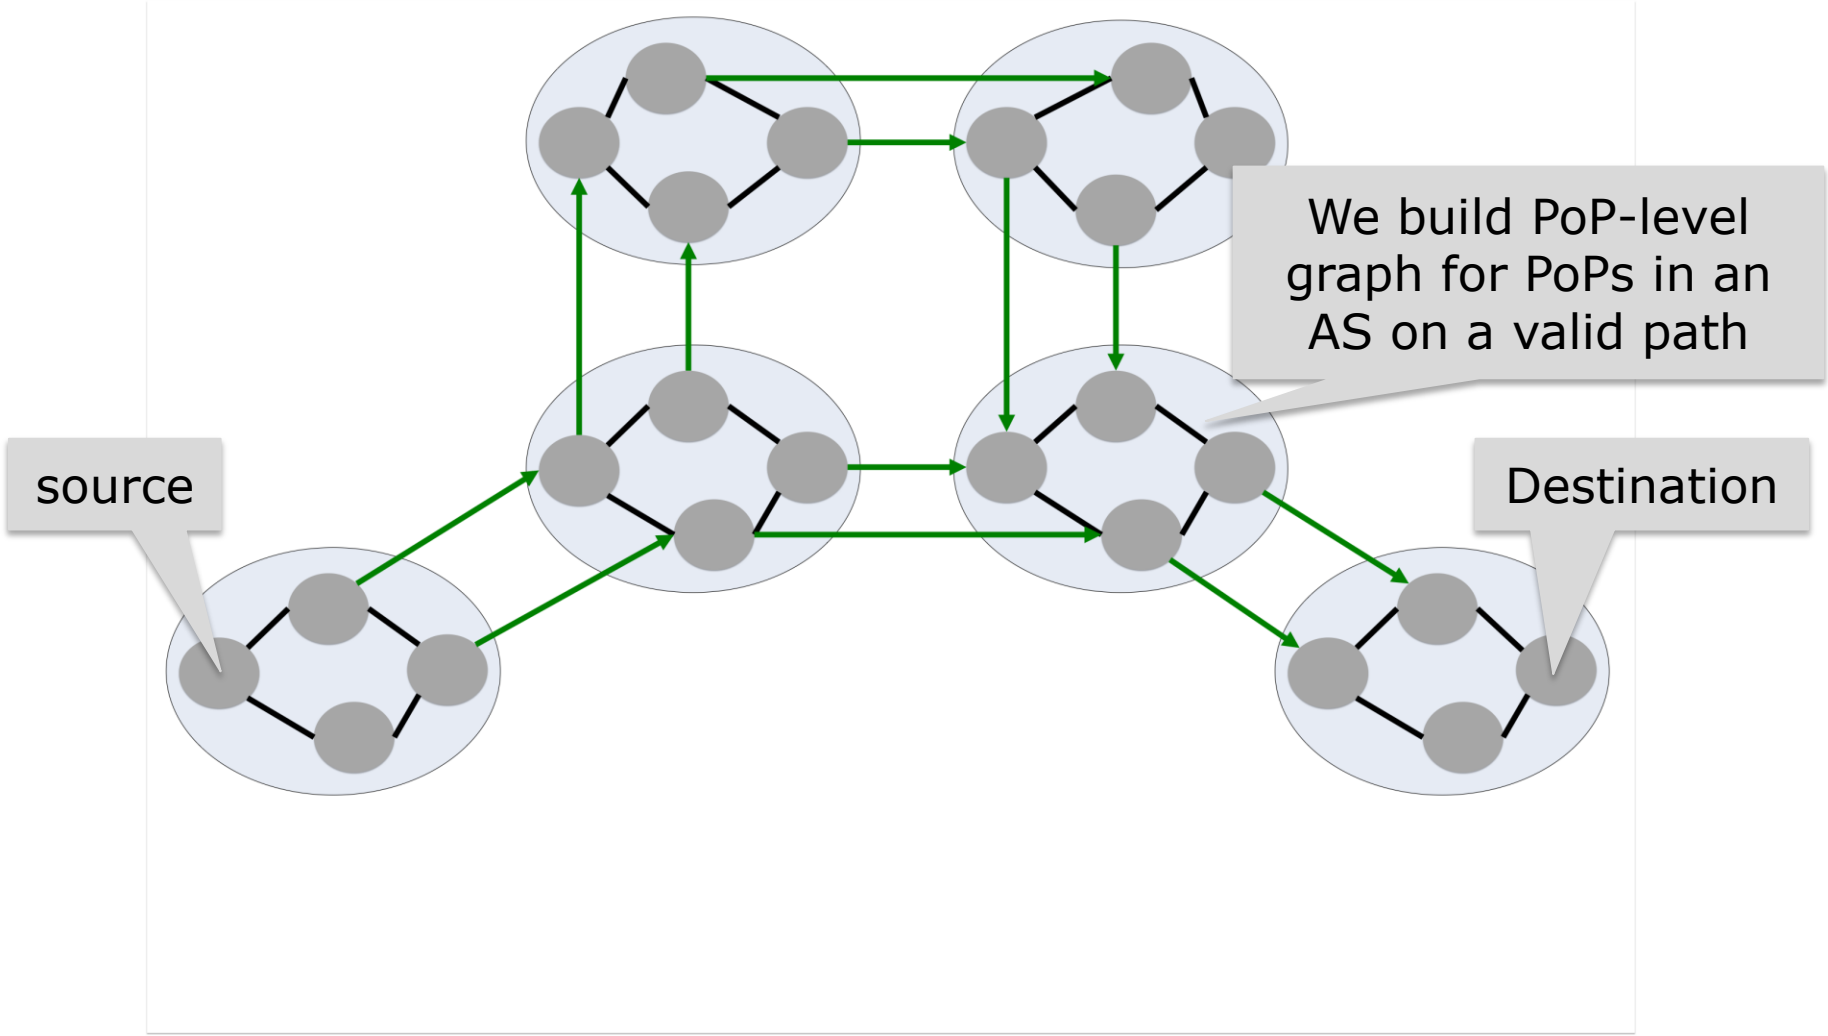

IIJ INNOVATION INSTITUTE

TELECOM ParisTech

worldline e-payment services

Future works and studies

Future work and studies

iPlane dataset

- What have we done yet?
  - We have built a PoP-level topology representing the Internet based on three datasets:
    - iPlane
    - DRAGON Topology
    - PeeringDB
  - We have gathered information about router presence at IXPs to help us evaluate the Intercloud architecture
- We will follow the evaluation of the Kumori architecture

Thank you

Antoine Fressancourt

antoine.fressancourt@worldline.com
antoine.fressancourt@telecom-paristech.fr

12-01-2016