On the Incompleteness of the AS-level graph: a Novel Methodology for BGP Route Collector Placement

Alessandro Improta

Enrico Gregori, Luciano Lenzini, Lorenzo Rossi, Luca Sani

UNIVERSITÀ DI PISA

Istituto di Informatica e Telematica del CNR

Institutions Markets Technologies IMT INSTITUTE FOR ADVANCED STUDIES LUCCA

Internet Measurement Conference - Boston - November 15th, 2012

Outline

1. BGP data collection overview
   - BGP route collector analysis
   - Feeder characteristics
   - Incompleteness and biases
2. Towards an ideal BGP measurement infrastructure
   - A new metric: p2c-distance
   - Tailored set covering problem
3. Quantifying the efforts required
   - Real world analysis

Outline

1. BGP data collection overview
   - BGP route collector analysis
   - Feeder characteristics
   - Incompleteness and biases
2. Towards an ideal BGP measurement infrastructure
   - A new metric: p2c-distance
   - Tailored set covering problem
3. Quantifying the efforts required
   - Real world analysis

> *"It is a capital mistake to theorize before you have all the evidence. It biases the judgment"*
>
> *(sir A.C. Doyle)*

BGP Route Collectors

> A Route Collector (RC) is a device which collects BGP routing data from co-operating ASes.

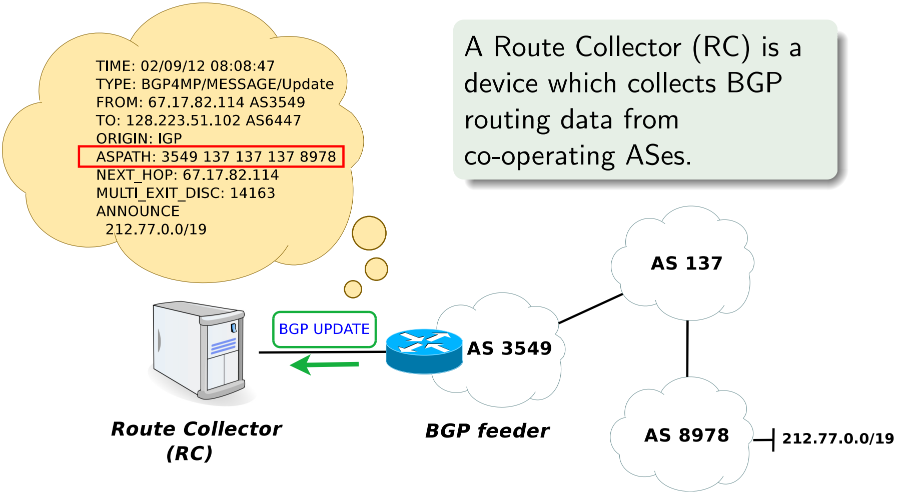

BGP Route Collector Status (Feb 2012)

	RouteViews	RIS	PCH
N. of RC	10	13	51
N. of feeders	313	299	1,842

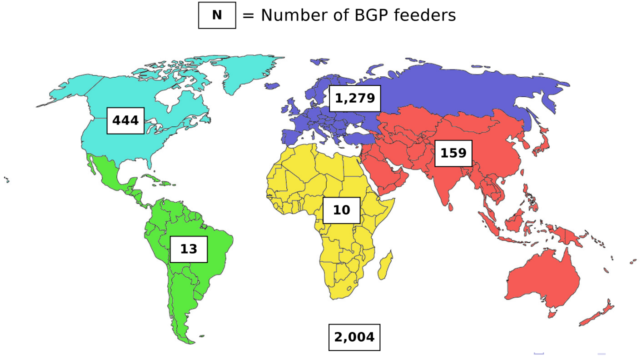

Feeder Contribution

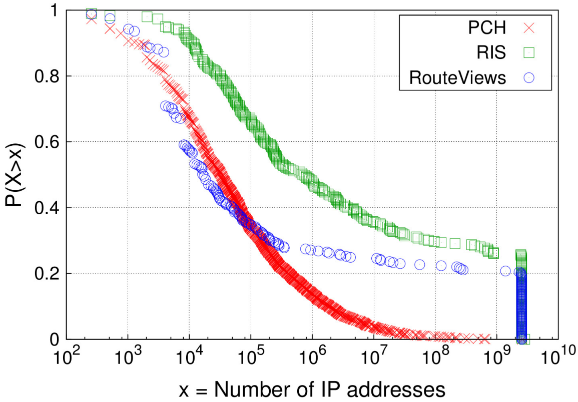

Only 120 feeders announce to the RCs their full routing table

Feeder Contribution

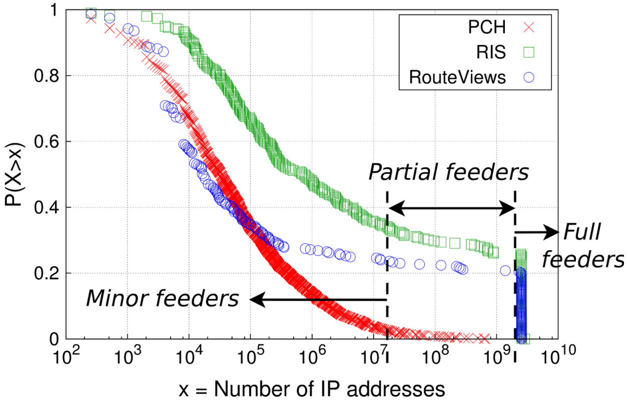

Minor feeders	Partial feeders	Full feeders
$\mathcal{S}_{IPv4} < 2^{24}$	$2^{24} < \mathcal{S}_{IPv4} < 2 * 10^9$	$\mathcal{S}_{IPv4} > 2 * 10^9$

Export Policies

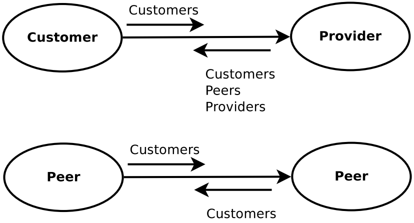

- PCH establishes only p2p connections
- RouteViews and RIS RCs are placed on IXPs

RCs need to be considered as *customers* by their feeders to gather a full routing table

Full feeder geographical distribution

Data collected represent mostly the Internet as viewed from Europe and North America than the real Internet

N (%) = Number of Full Feeders (%)

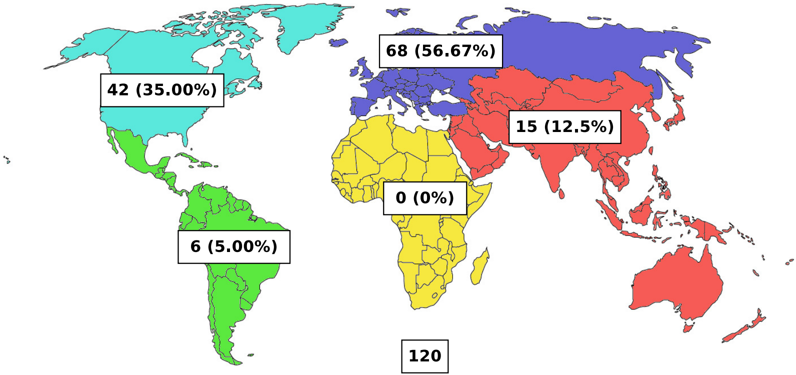

Feeder characterization

RouteViews

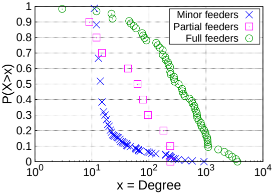

RIS

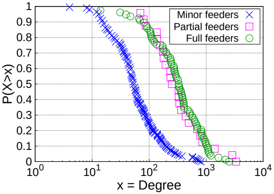

PCH

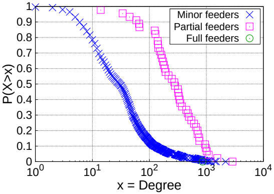

About 80% of full feeders have a degree higher than 100

A view from the top

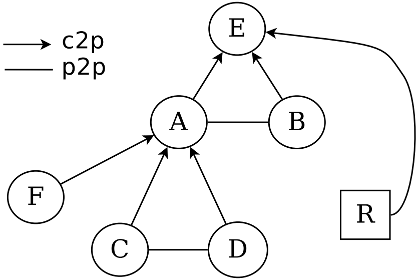

$$\frac{\textit{Connections that can be discovered}}{(A, C)\ (A, D)\ (A, E)\ (A, F)\ (B, E)}$$

RCs connected to large ISPs will fail to retrieve a large amount of p2p-connectivity

A view from the bottom

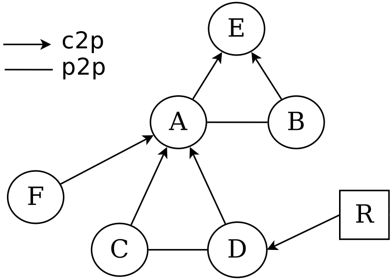

Connections that can be discovered

(A, B) (A, C) (A, D) (A, E) (A, F) (B, E) (C, D)

RCs need to be connected to ASes part of the lowest part of the Internet hierarchy to discover the missing p2p connectivity

Outline

"If you cannot measure it,
you cannot improve it"
(sir W. Thomson)

p2c distance of AS X from AS Y:

Minimum number of consecutive p2c links that connect X to Y

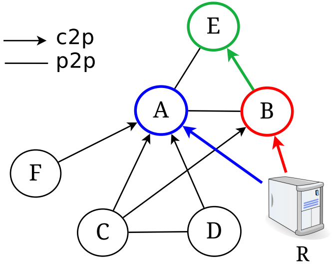

AS	p2c-distance from R
A	1
B	1
C	-
D	-
E	2
F	-

Farther an AS is from a RC, the greater are the chances to lose AS-level connectivity due to BGP decision processes

Focusing the target

Thoughts

- Every AS becomes feeder: unfeasible and unuseful
- The vast majority of missing links are p2p
- Stub ASes are not likely to establish many p2p connections (only 7% are members of at least an IXP)

Goals

- Discover the connectivity of **non-stub** ASes ...
- ... without connecting to all of them
- Note: Stub ASes may be still exploited as feeders to achieve this objective

	AF	*AP*	*EU*	*LA*	*NA*	*W*
ASes	770	6,576	17,657	2,490	16,032	41,127
Non-stub ASes	229	1,589	3,697	659	2,531	7,282

Tailored set covering problem

Goal rephrased

Select new BGP feeders such that each non-stub AS has a **finite and bounded** p2c distance from the route collector infrastructure

Set Covering

$$\text{Minimize} \quad \left(\sum_{AS_i \in \mathcal{U}} x_{AS_i} \right) \tag{1}$$

subject to

$$\sum_{AS_i : n \in S_{AS_i}^{(d)}} x_{AS_i} \geq 1 \qquad \forall n \in \mathcal{N} \tag{2}$$

$$x_{AS_i} \in \{0, 1\}, \qquad \forall AS_i \in \mathcal{U} \tag{3}$$

Output

- One optimal solution $\mathcal{P}$
- Set of candidates interchangeable with ASes in $\mathcal{P}$

Outline

1. BGP data collection overview
   - BGP route collector analysis
   - Feeder characteristics
   - Incompleteness and biases
2. Towards an ideal BGP measurement infrastructure
   - A new metric: p2c-distance
   - Tailored set covering problem
3. Quantifying the efforts required
   - Real world analysis

> *"First ask yourself: What is the worst that can happen? Then prepare to accept it. Then proceed to improve on the worst."*
>
> *(Dale Carnegie)*

Real World Analysis

Distance parameter

- $d_{p2c} = 2$: to obtain the best quality result without the need to establish a connection with every non-stub ASes

Economic topologies

- Global [1]
- Continental [2]

Scenarios

- Empty scenario:
  - current full feeders are ignored
- Full feeders scenario:
  - current full feeders are part of the solution set

[1] *"BGP and Inter-AS Economic Relationships"*, IFIP Networking 2011, pp. 54-67
[2] *"Inferring Geography from BGP Raw Data"*, IEEE INFOCOM NetSciCom, INFOCOM 2012, pp. 208-213

Empty Scenario

M / N = Cardinality of solution / N. of candidates

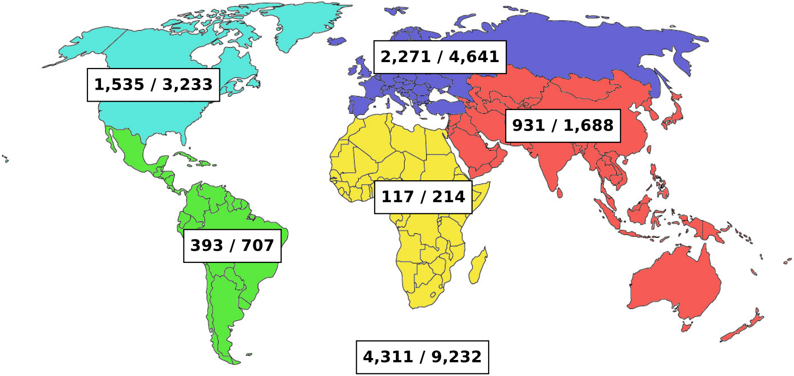

The number of feeders required heavily outnumbers the current number of (full) feeders

Candidate feeder details

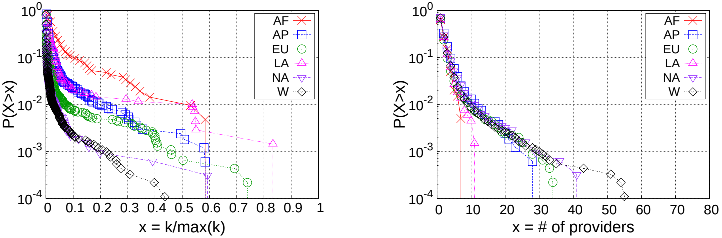

<table>
<tr><th rowspan="2">Region</th><th colspan="2">|Candidates|</th></tr>
<tr><th>On IXPs</th><th>Stubs</th></tr>
<tr><td>AF</td><td>27 (12.79 %)</td><td>114 (54.03 %)</td></tr>
<tr><td>AP</td><td>472 (28.04 %)</td><td>942 (55.97 %)</td></tr>
<tr><td>EU</td><td>1,931 (41.60 %)</td><td>2,250 (48.48 %)</td></tr>
<tr><td>LA</td><td>204 (29.14 %)</td><td>394 (56.29 %)</td></tr>
<tr><td>NA</td><td>406 (12.55 %)</td><td>1,509 (46.67 %)</td></tr>
<tr><td>W</td><td>2,944 (31.88 %)</td><td>4,221 (45.72 %)</td></tr>
</table>

These are exactly the ASes that rarely feed the RCs

Full Feeder Scenario

N = Number of additional full feeders required

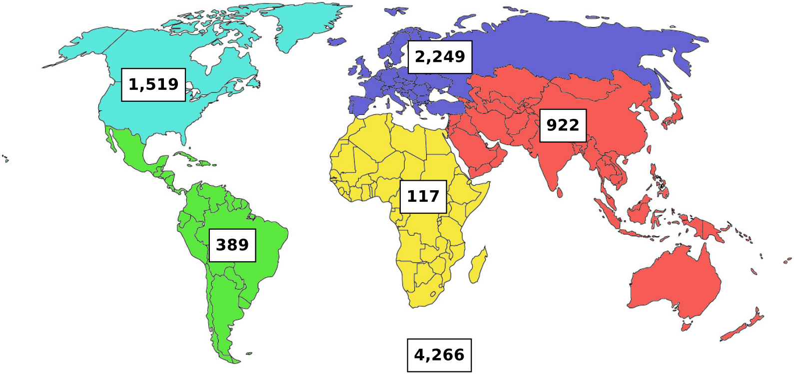

The introduction of full feeders in the solution set do not improve much the situation

Conclusions and Future works

Conclusions

- Several p2p-connectivity is hidden from RC sight
- Several Internet regions are basically uncovered
- The typical profile of an ideal feeder is a multi-homed stub AS

Future directions

- Analyze the impact of **traceroute** datasets to the results
- Analyze the IPv6 contribution of feeders
- Set up a **route collector project** in order to gather BGP data from identified ASes
  - do-ut-des rather than volunteer participation

Thank you for your attention

Isolario

Data presented in this paper and many others can be found at
www.isolario.it

Any question?

alessandro.improta@iet.unipi.it

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Phase a)

Identify covering sets ...

AS	Not stubs $\in \mathbf{S}^{(1)}_{\mathbf{AS_i}}$
A	{B}
B	{B,D}
C	{C}
D	{D}
E	{D,E,G,H}
F	{E}
G	{G}
H	{H}
I	{B}

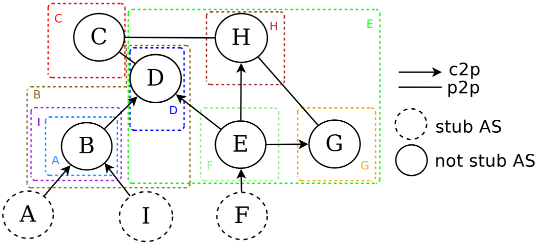

$$\mathcal{P} = \{\emptyset\}, \mathcal{D} = \{\emptyset\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Phase a)

... and ASes that uniquely cover a non-stub AS

AS	Not stubs $\in \mathbf{S}^{(1)}_{\mathbf{AS_i}}$
A	{B}
B	{B,D}
C	{C}
D	{D}
E	{D,E,G,H}
F	{E}
G	{G}
H	{H}
I	{B}

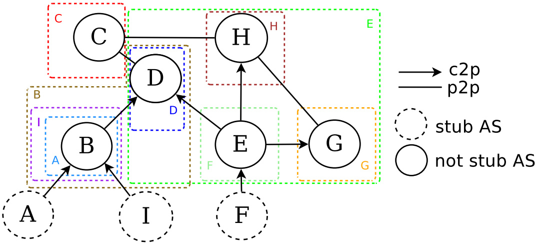

$$\mathcal{P} = \{\emptyset\}, \mathcal{D} = \{\emptyset\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Phase b)

Identify dominated covering sets ...

AS	**Not stubs** $\in \mathbf{S^{(1)}_{AS_i}}$
A	{**B**}
B	{**B**,**D**}
~~C~~	~~{C}~~
D	{**D**}
E	{**D**,**E**,**G**,**H**}
F	{**E**}
G	{**G**}
H	{**H**}
I	{**B**}

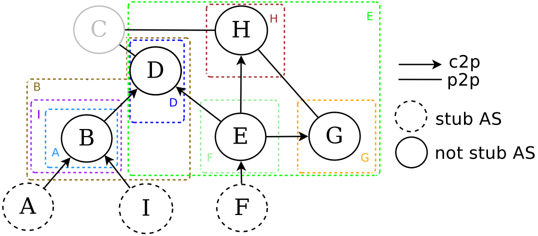

$$\mathcal{P} = \{C\}, \mathcal{D} = \{\emptyset\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Phase b)

... record and put them aside

AS	**Not stubs** $\in \mathbf{S}^{(1)}_{\mathbf{AS_i}}$
~~A~~	{~~B~~}
B	{B,D}
~~C~~	~~{C}~~
~~D~~	{~~D~~}
E	{D,E,G,H}
~~F~~	{~~E~~}
~~G~~	{~~G~~}
~~H~~	{~~H~~}
~~I~~	{~~B~~}

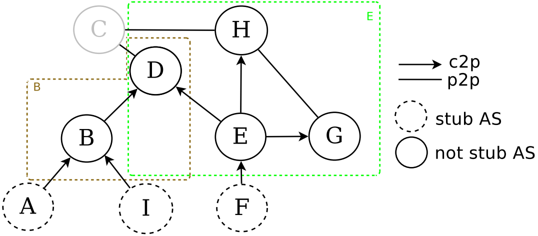

$$\mathcal{P} = \{C\}, \mathcal{D} = \{A, C, D, F, G, H, I\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Repeat previous steps until a solution is found or apply brute force approach (if needed)

AS	**Not stubs** $\in \mathbf{S}^{(1)}_{\mathbf{AS_i}}$
~~A~~	{~~B~~}
B	{**B**,D}
~~C~~	~~{C}~~
~~D~~	~~{D}~~
E	{D,**E**,**G**,**H**}
~~F~~	{~~E~~}
~~G~~	{~~G~~}
~~H~~	{~~H~~}
~~I~~	{~~B~~}

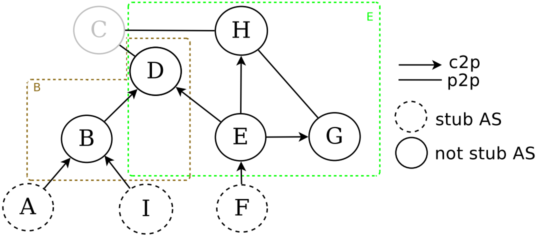

$$\mathcal{P} = \{C\}, \mathcal{D} = \{A, C, D, F, G, H, I\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Repeat previous steps until a solution is found or apply brute force approach (if needed)

AS	**Not stubs** $\in \mathbf{S}^{(1)}_{\mathbf{AS_i}}$
~~A~~	{~~B~~}
B	{B,D}
C	{C}
~~D~~	~~{D}~~
E	{D,E,G,H}
~~F~~	{~~E~~}
~~G~~	{~~G~~}
~~H~~	{~~H~~}
~~I~~	{~~B~~}

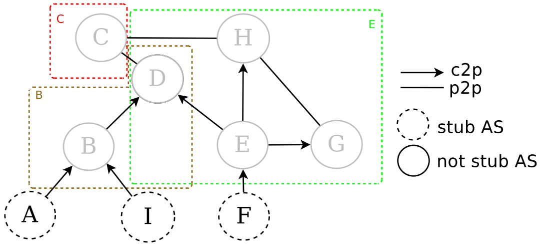

$$\mathcal{P} = \{B, C, E\}, \mathcal{D} = \{A, C, D, F, G, H, I\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Phase c)

Check if dominated covering sets can appear in a solution

AS	Not stubs uniquely covered $\in S^{(1)}_{AS_i}$
B	{**B**}
C	{C}
E	{E,G,H}

AS in $\mathcal{D}$	Not stubs $\in S^{(1)}_{AS_i}$
A	{**B**}
D	{D}
F	{E}
G	{G}
H	{H}
I	{**B**}

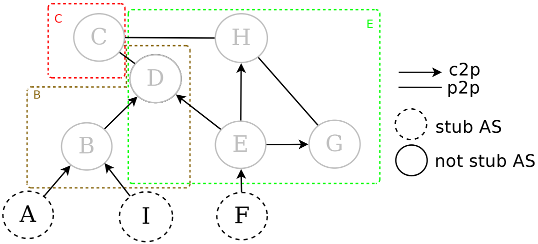

$$\mathcal{D} = \{A, C, D, F, G, H, I\}, \mathcal{C} = \{B, C, E\}$$

So, for example ...

Select the min number of feeders to have each **not stub AS** with $d_{p2c} = 2$ from the RCs (i.e. $d_{p2c} = 1$ from the feeders)

Phase c)

Check if dominated covering sets can appear in a solution

AS	Not stubs uniquely covered $\in S^{(1)}_{AS_i}$
B	{**B**}
C	{C}
E	{E,G,H}

AS in $\mathcal{D}$	Not stubs $\in S^{(1)}_{AS_i}$
A	{**B**}
D	{D}
F	{E}
G	{G}
H	{H}
I	{**B**}

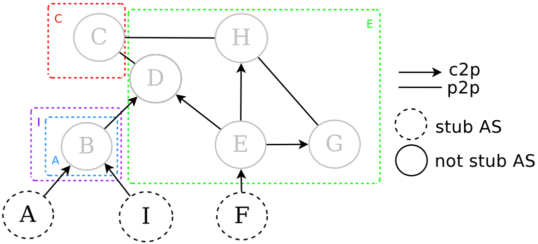

$$\mathcal{D} = \{C, D, F, G, H\}, \mathcal{C} = \{A, B, C, E, I\}$$