INTRIG
INFORMATION & NETWORKING TECHNOLOGIES RESEARCH & INNOVATION GROUP

SBRC 2015
33º SIMPÓSIO BRASILEIRO DE REDES DE COMPUTADORES E SISTEMAS DISTRIBUÍDOS

Anatomy of Internet eXchange Points (IXP) Ecosystem in Brazil

Samuel Henrique Bucke Brito
Mateus Augusto Silva Santos
Ramon dos Reis Fontes
Danny Alex Lachos Perez
Christian Esteve Rothenberg

Department of Computer Engineering and Industrial Automation (DCA)
Faculty of Electrical and Computer Engineering (FEEC)
University of Campinas (UNICAMP)

FEEC
Faculdade de Engenharia Elétrica e de Computação
UNICAMP

May 21, 2015

Agenda

- Introduction and Background
- Methodology
- Analysis and Results
- Conclusion
- Future Work

What Are IXPs?

- A neutral-carrier infrastructure
- Host the interconnection of many ASes
- Provides a layer-2 switch fabric for its members
- Allow bilateral agreements through direct connections
- Allow multilateral agreements through route-servers
- Improve the Internet performance in its location

What Are IXPs?

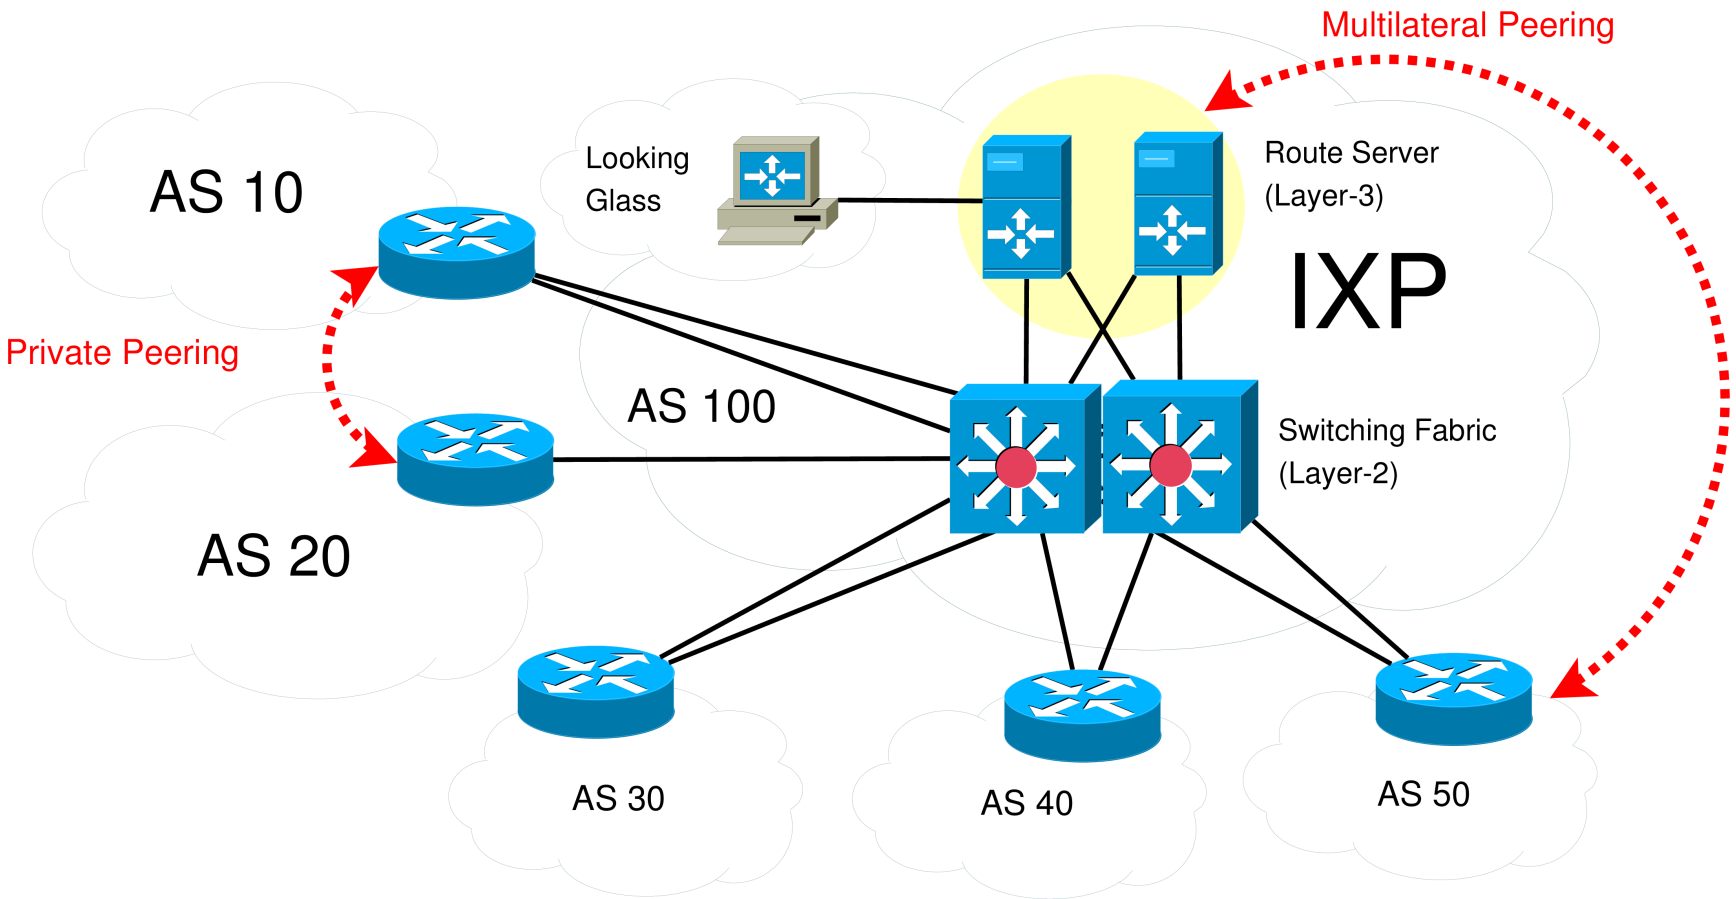

Figure 1a. General Architecture of an IXP

Why IXPs?

Because of the high cost of individual links between many ASes it is infeasible to achieve a full mesh topology. An IXP provides a layer-2 switch fabric to directly connect ASes in a star topology.

directly between ASes (without an IXP)

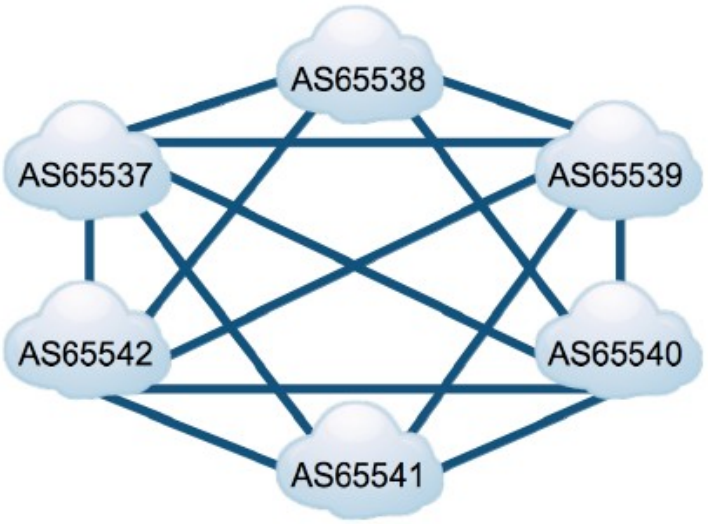

X

through an IXP

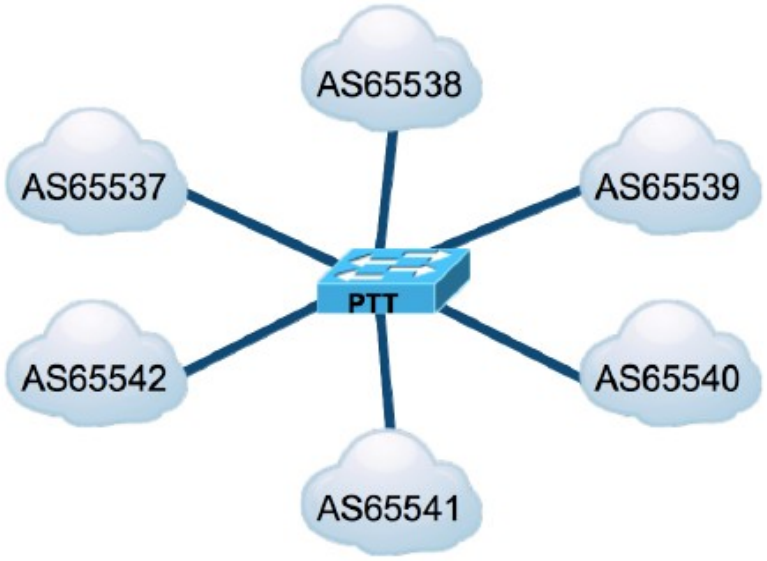

figures extracted from NIC.br Course for Autonomous Systems

Why IXPs?

The connection of many ASes through IXPs improve the Internet performance by minimizing the number of networks to traverse to reach a destination (diameter).

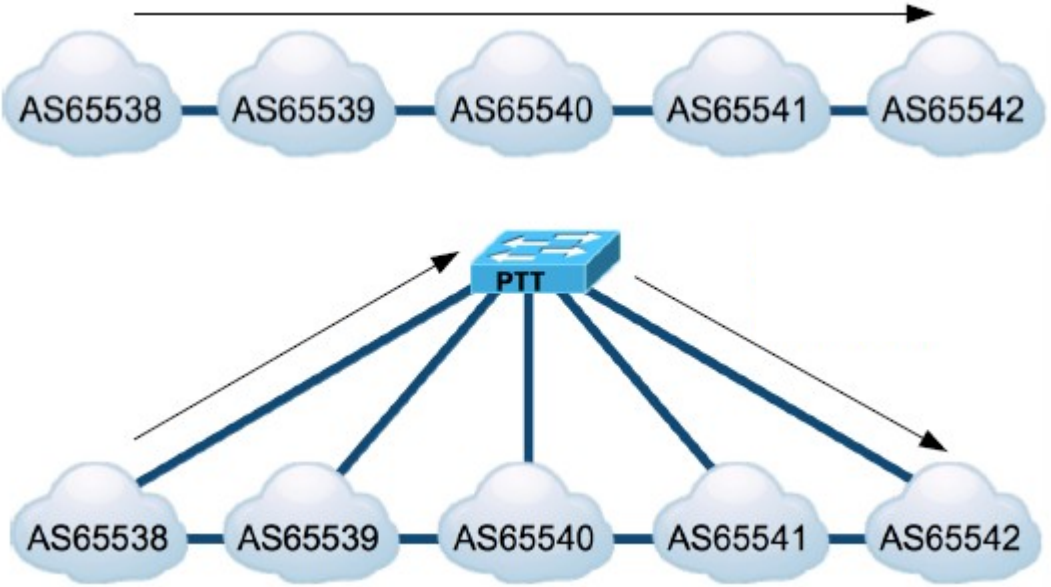

figures extracted from NIC.br Course for Autonomous Systems

PTTMetro in Brazil (PTT.br)

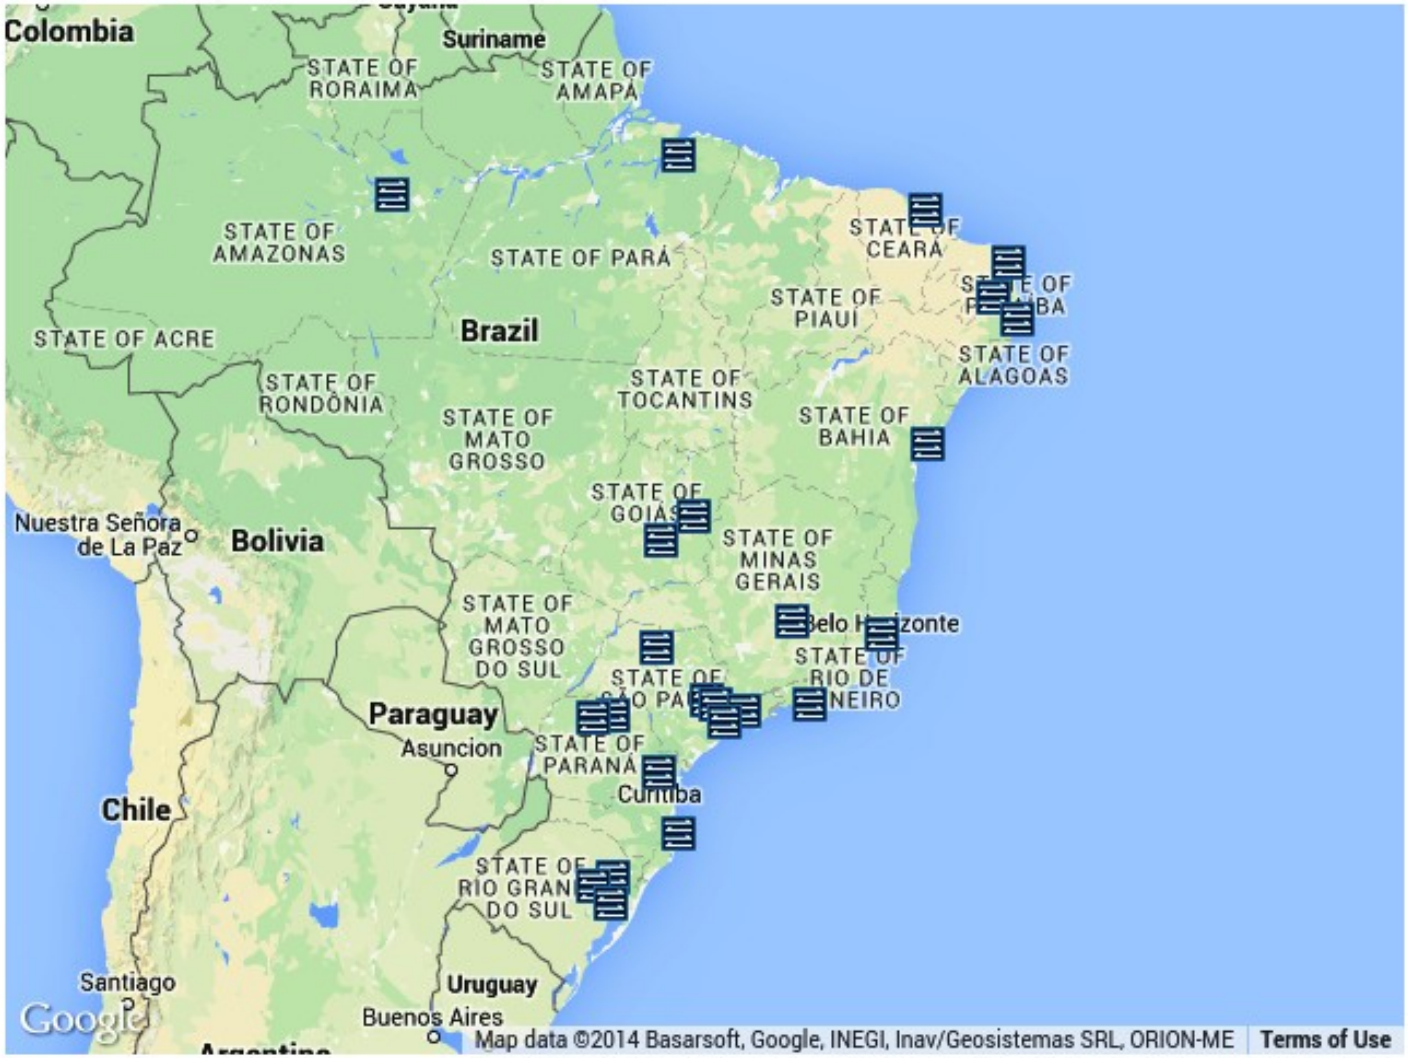

- Americana
- Belém
- Belo Horizonte
- Brasília
- Campina Grande
- Campinas
- Cuiabá
- Caxias do Sul
- Curitiba
- Florianópolis
- Fortaleza
- Goiânia
- Lajeado
- Londrina
- Manaus
- Maringa
- Natal
- Porto Alegre
- Recife
- Rio de Janeiro
- Salvador
- Paulista Central (São Carlos)
- São José dos Campos
- São José do Rio Preto
- São Paulo
- Vitória

Location of Each Brazilian IXP (figure extracted from http://www.ptt.br)

PTTMetro x World

Table 1: Comparison of traffic between some of the world's largest public IXPs (March 25, 2015).

IXP	Country	Members	Maximum Throughput (Gbps)			Average Throughput (Gbps)		
			Daily	Monthly	Yearly	Daily	Monthly	Yearly
(01) DE-CIX	Germany	600+	3,603.10	3,854.80	3,875.10	2,375.90	2,299.20	1,964.90
(02) AMS-IX	Netherlands	708	3,429.38	-	3,604.48	2,120.70	-	1,893.38
(03) LINX	United Kingdom	628	2,352.13	2,558.98	2,573.31	1,419.26	1,571.84	1,420.29
(04) MSK-IX	Russia	374	1,332.63	1,457.81	1,479.12	751.24	806.76	727.27
(05) NL-ix	Netherlands	476	801.14	-	-	-	456.48	-
(06) PTTMetro	**Brazil**	**1,142**	**678.50**	**685.67**	**467.46**	**393.32**	**432.21**	**360.62**
(07) HKIX	Hong Kong	220	416.96	421.12	441.03	288.18	285.36	220.06
(08) SIX	USA, Canada	193	347.84	366.11	366.11	250.94	257.39	202.37
(09) JPIX	Japan	138	303.78	-	-	186.89	-	-
(10) JINX	South Africa	24	11.80	17.90	9.70	6.40	6.20	5.60

(01) http://www.de-cix.net/about/statistics/
(02) https://ams-ix.net/technical/statistics
(03) https://www.linx.net/pubtools/trafficstats.html
(04) http://www.msk-ix.ru/network/traffic.html
(05) https://www.nl-ix.net/network/traffic/
(06) http://ptt.br/cgi-bin/all
(07) http://www.hkix.net/hkix/stat/aggt/hkix-aggregate.html
(08) http://www.seattleix.net/agg.htm
(09) http://www.jpix.ad.jp/en/technical/traffic.html
(10) http://stats.jinx.net.za/showtotal.php

Agenda

- Introduction and Background
- Methodology
- Analysis and Results
- Conclusion
- Future Work

Methodology: Data Sources

- 1st source of information:
  - PTT.br Official Data at www.ptt.br
  - PeeringDB (we found it was unreliable)

- 2nd (and most important) source of information:
  - Telnet Access to IXP's Looking Glasses
    = BGP Table, Paths Summary, Communities List
    (*) Currently we also have the IPv6 BGP Table

Methodology: AS-level Graphs

We used the following tools to build the connectivity graphs of each Brazilian IXP, providing the adjacency matrix as input:

1) NetworkX ::: https://networkx.github.io

2) Neo4j ::: http://neo4j.com

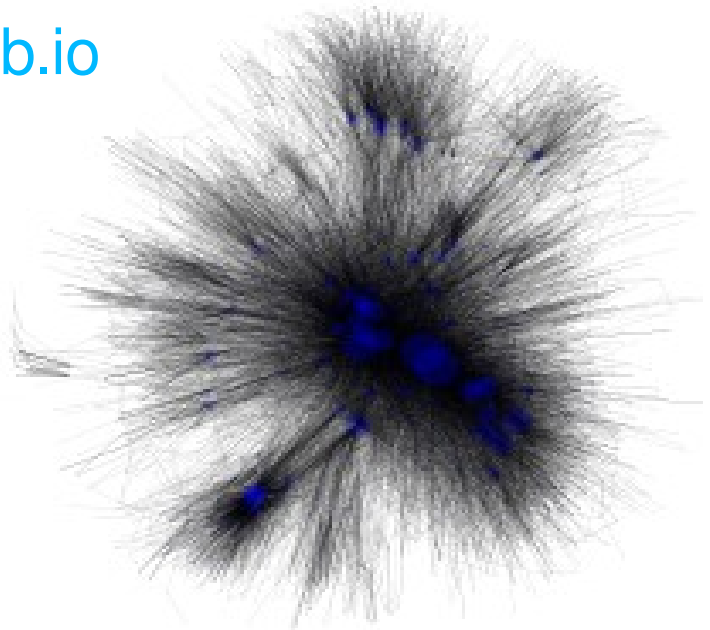

Figure 1b. Example Graph of PTT-VIX (Vitória, ES)

Methodology: AS-level Graphs

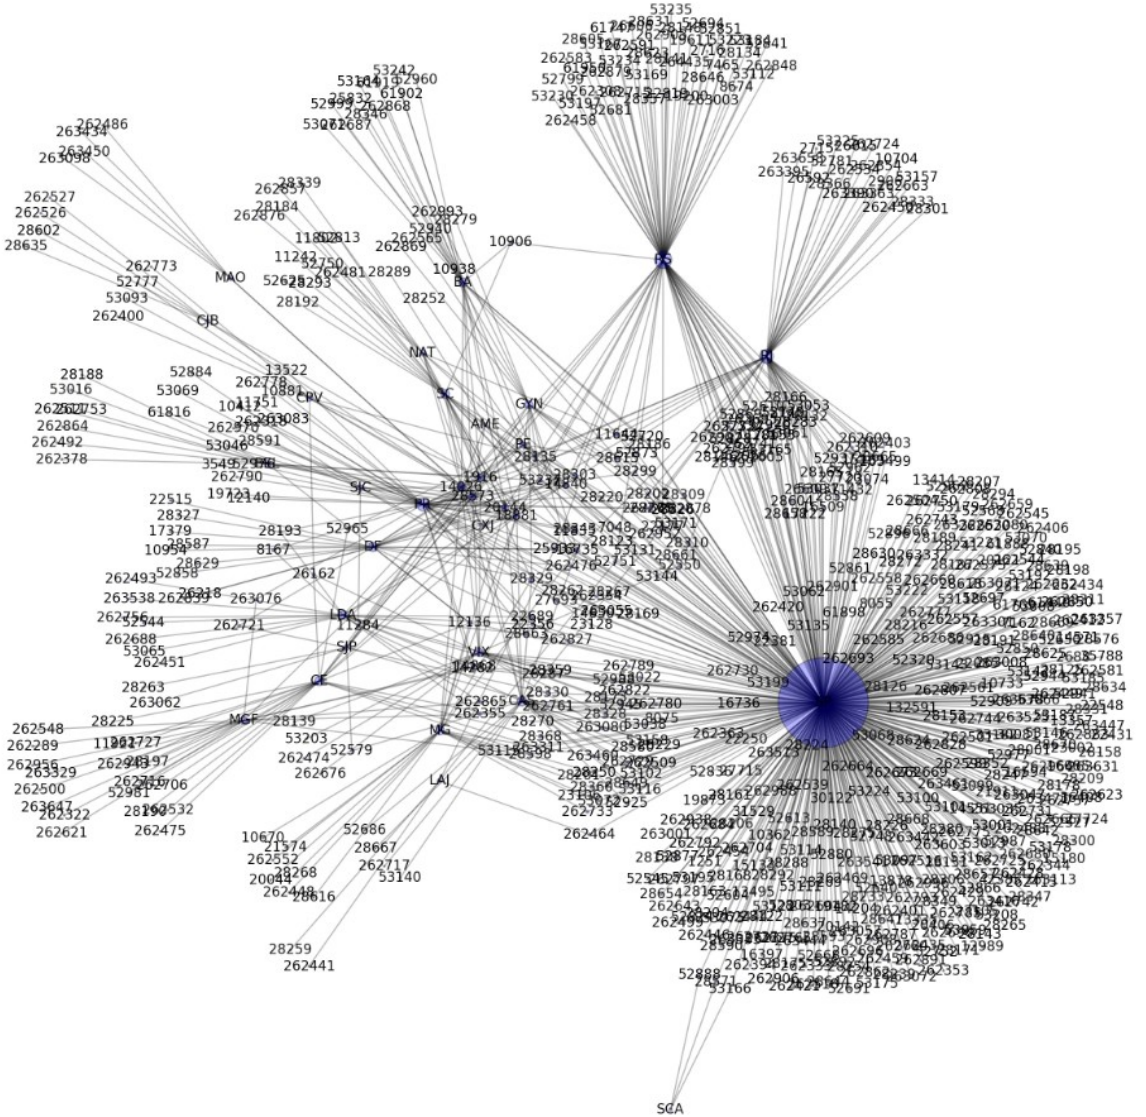

Methodology: Workflow

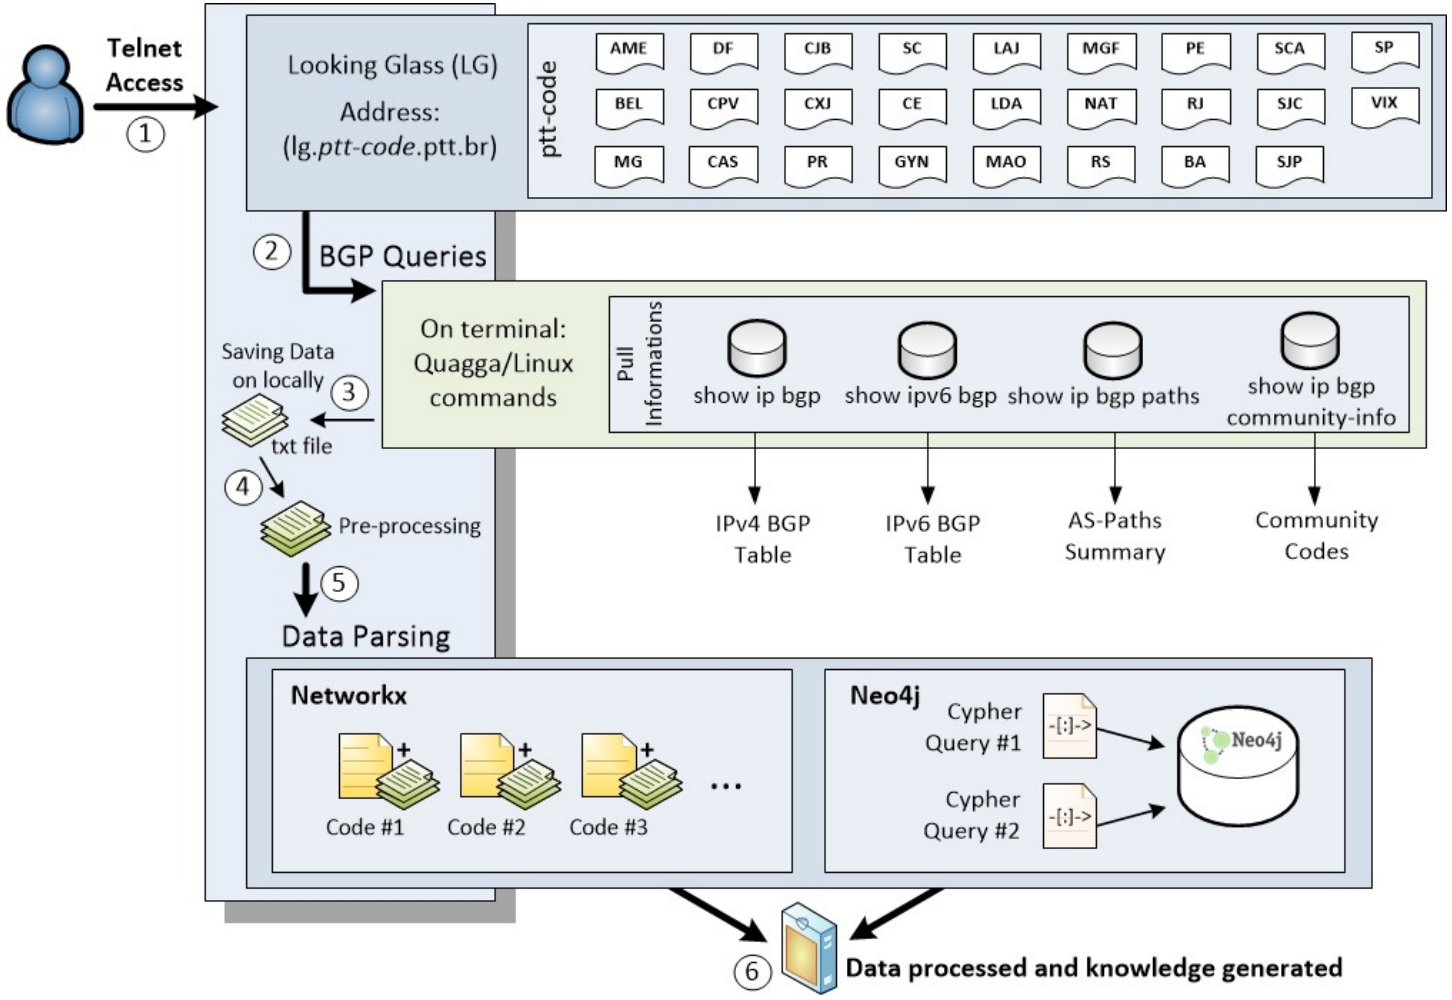

Agenda

- Introduction and Background
- Methodology
- Analysis and Results
- Conclusion
- Future Work

Results: Profile of IXP's Members

We first classified all the members of each Brazilian IXP according to the following sample table:

Classification	Brazil (*)		DF	MG	RJ	RS	SP	VIX
1. Internet Provider	743	(65.1% ± 20%)	37.5%	55.9%	51.9%	68.0%	73.1%	75.0%
1.1 Transit Provider	98	(8.6% ± 09%)	20.8%	14.7%	19.2%	5.0%	5.6%	10.0%
1.2 Access Provider	645	(56.5% ± 21%)	16.7%	41.2%	32.7%	63.0%	67.5%	65.0%
2. Services Provider	115	(10.1% ± 07%)	8.3%	8.8%	17.3%	5.0%	12.5%	5.0%
2.1 Content Provider	37	(3.2% ± 06%)	0.0%	2.9%	5.8%	3.0%	4.7%	0.0%
2.2 Hosting Provider	78	(6.8% ± 05%)	8.3%	5.9%	11.5%	2.0%	7.8%	5.0%
3. Public Organization	140	(12.3% ± 21%)	37.5%	20.6%	15.4%	11.0%	4.4%	15.0%
3.1 Public University	20	(1.8% ± 19%)	0.0%	0.0%	0.0%	2.0%	1.1%	0.0%
3.2 Government	100	(8.8% ± 13%)	33.3%	17.6%	13.5%	8.0%	2.2%	15.0%
3.3 Other	20	(1.8% ± 03%)	4.2%	2.9%	1.9%	1.0%	1.1%	0.0%
4. Private Organization	144	(12.6% ± 09%)	16.7%	14.7%	15.4%	16.0%	10.0%	5.0%
4.1 Private University	8	(0.7% ± 03%)	0.0%	2.9%	0.0%	4.0%	0.0%	0.0%
4.2 Private Company	119	(10.4% ± 09%)	16.7%	8.8%	15.4%	10.0%	8.9%	5.0%
4.3 Other	17	(1.5% ± 09%)	0.0%	2.9%	0.0%	2.0%	1.1%	0.0%

(*) Average of ALL 26 Brazilian IXPs.

Table 2. Analysis of ASes Profile in PTTMetro

Results: Vertices's Degree

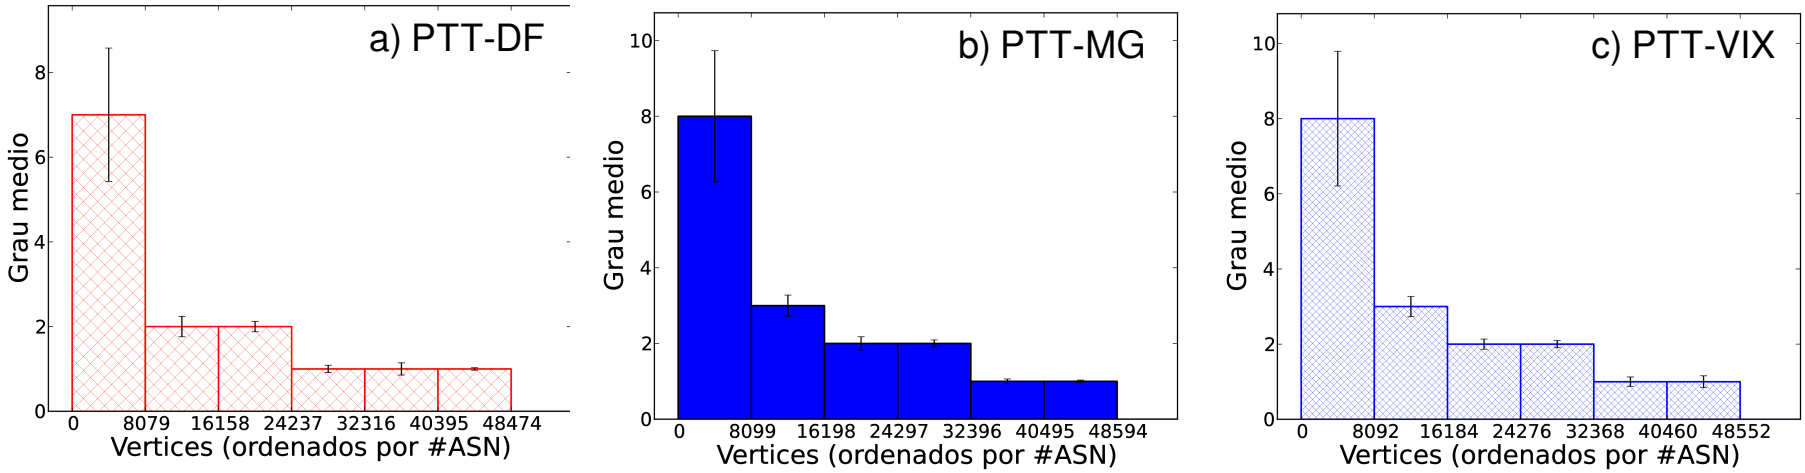

Figure 2. Average Degree of Graphs by ASN

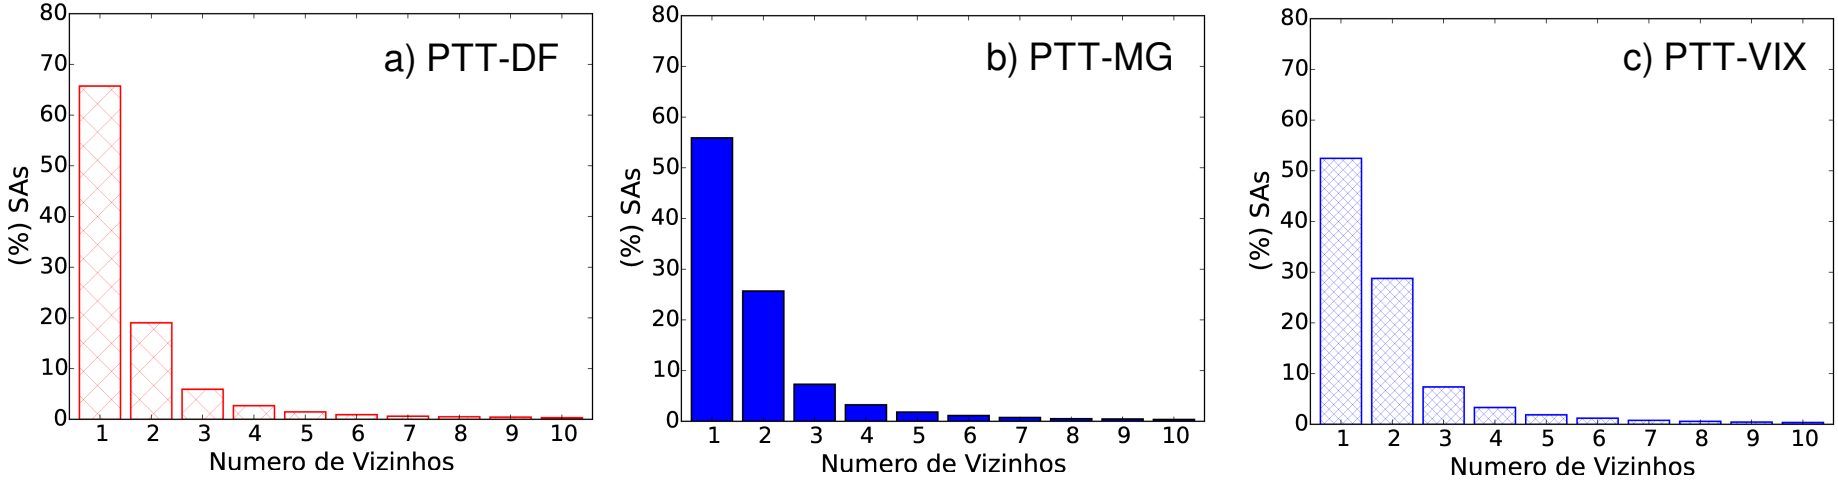

Figure 3. Distribution of Graphs' Degrees

Results: Depth / Diameter

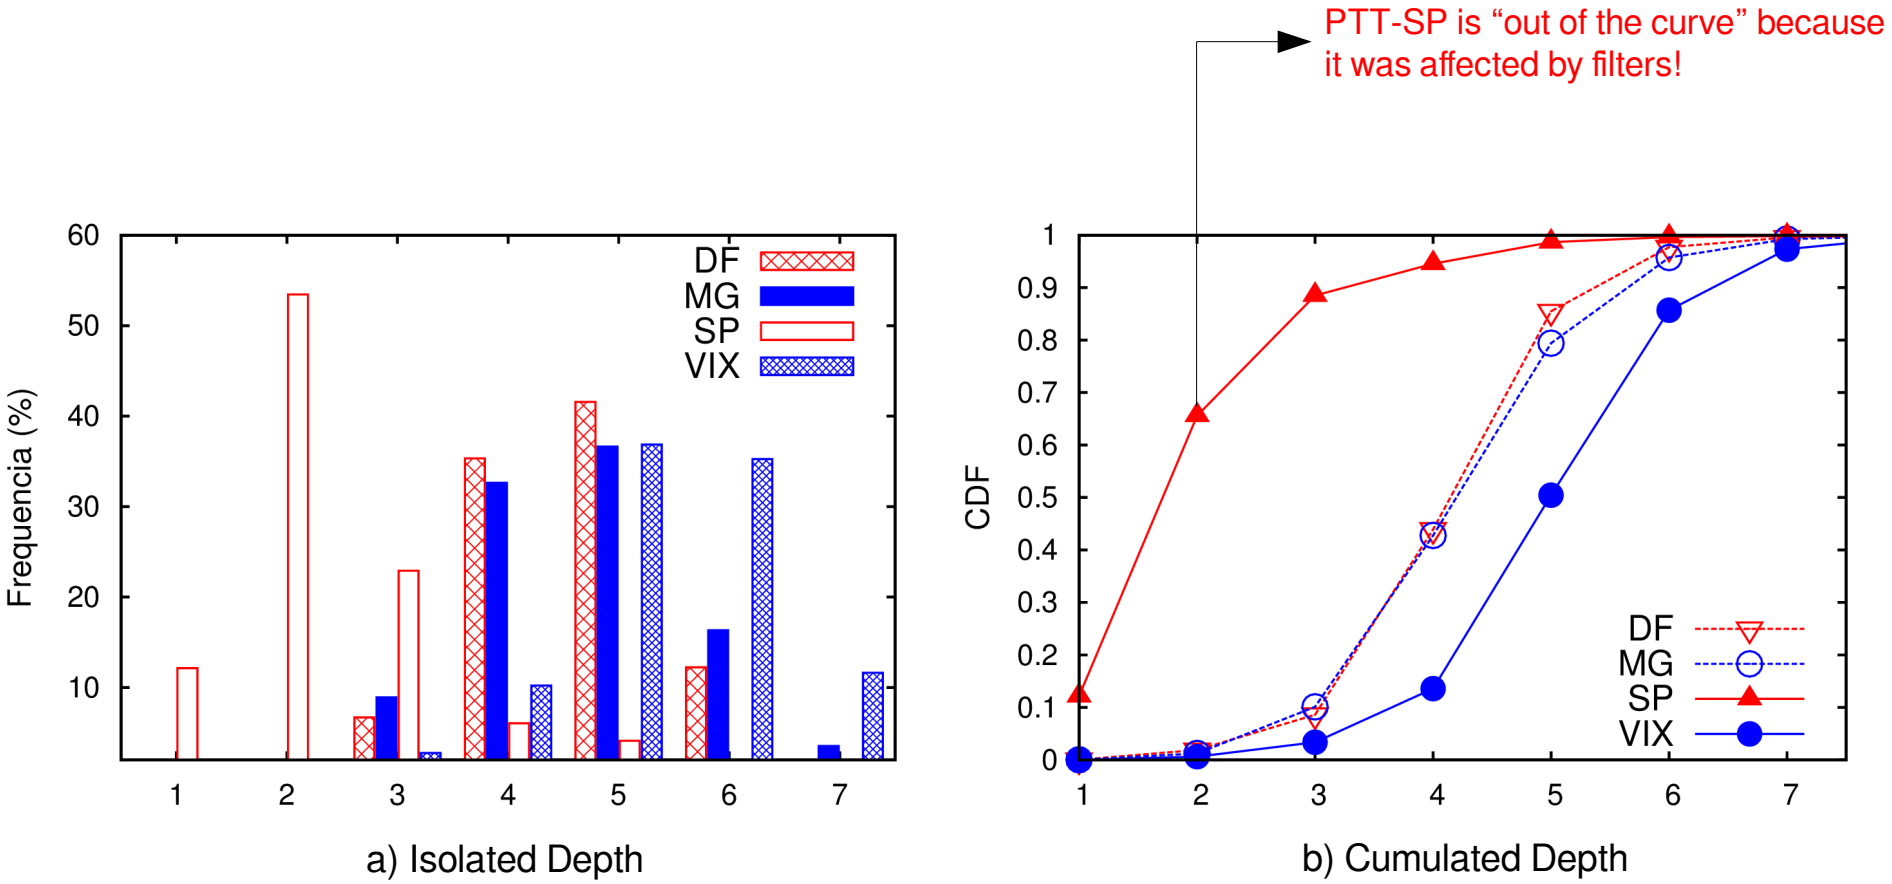

Figure 4. Depth of AS-PATHs

Results: Density of Peering

The density of peering is the ratio between the quantity of members effectively advertising in an IXP and its combination of possible peers.

Metric Description	Brazil (*)	DF	MG	RJ	RS	VIX
Existent Peers (Connections)	126	57	79	271	1,952	71
Combination of Possible Peers	285	276	231	1,275	3,081	153
Density (%)	44.2% ± 23%	20.7%	34.2%	21.3%	63.4%	46.4%

(*) Average of 24 Brazilian IXPs without filters, that is, excluding PTT-PR and PTT-SP.

Table 3. Density of Peers in PTTMetro

Results: AS-Prepend for TE

The AS-PATH attribute is commonly manipulated by an AS to prepend its own ASN for traffic engineering, a technique that makes a route through itself less attractive.

Metric Description	Brazil (*)	DF	MG	RJ	RS	VIX
Routes	832,989	559,159	434,264	1,150,905	1,947,453	2.663.751
Routes with AS-Prepend	295,909	127,184	245,129	294,663	1,710,070	623.965
AS-Prepend X Routes (%)	30.8% ± 22%	22.7%	56.4%	25.6%	87.8%	23,4%
ASes at Graph	43,333	47,176	46,939	47,632	48,351	47.474
ASes with AS-Prepend	7,305	6,206	8,629	8,890	10,803	9.124
AS-Prepend X ASes (%)	16.1% ± 04%	13.2%	18.4%	18.7%	22.3%	19,2%
Members Advertising	18	24	22	51	79	18
Members Advertising with AS-Prepend	6	7	6	19	36	5
AS-Prepend X Members (%)	22.5% ± 19%	29.2%	27.3%	37.3%	45.6%	27.8%

(*) Average of 24 Brazilian IXPs without filters, that is, excluding PTT-PR and PTT-SP.

Table 4. AS-Prepend Analysis in PTTMetro

Results: k-Clique Communities

For lower values of k (mainly 3 and 4) the amount of communities identified are similar (figure 5a) and these communities have great density of connections between its members (figure 5b).

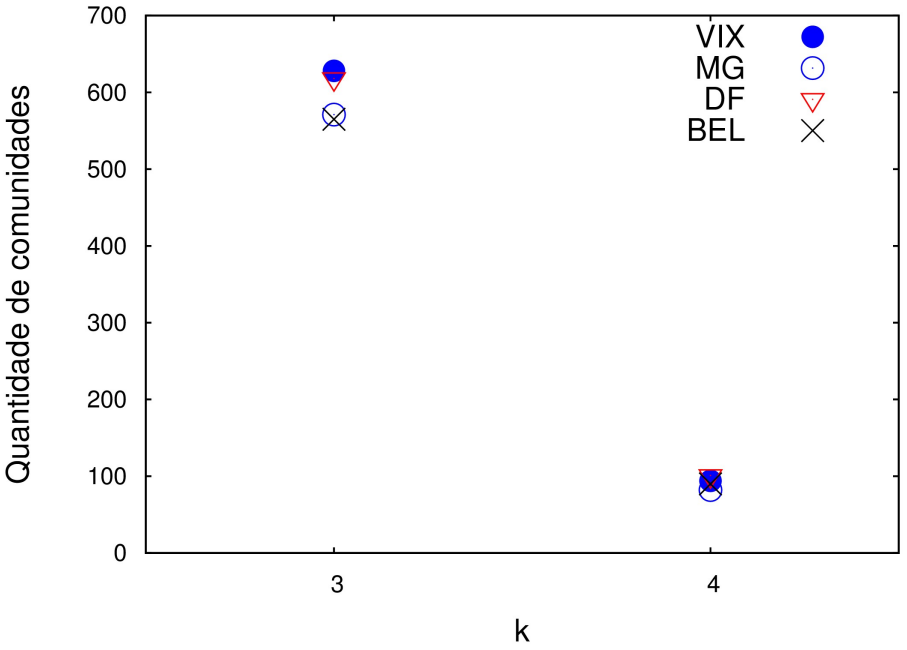

Figure 5a. k3 and k4 Communities

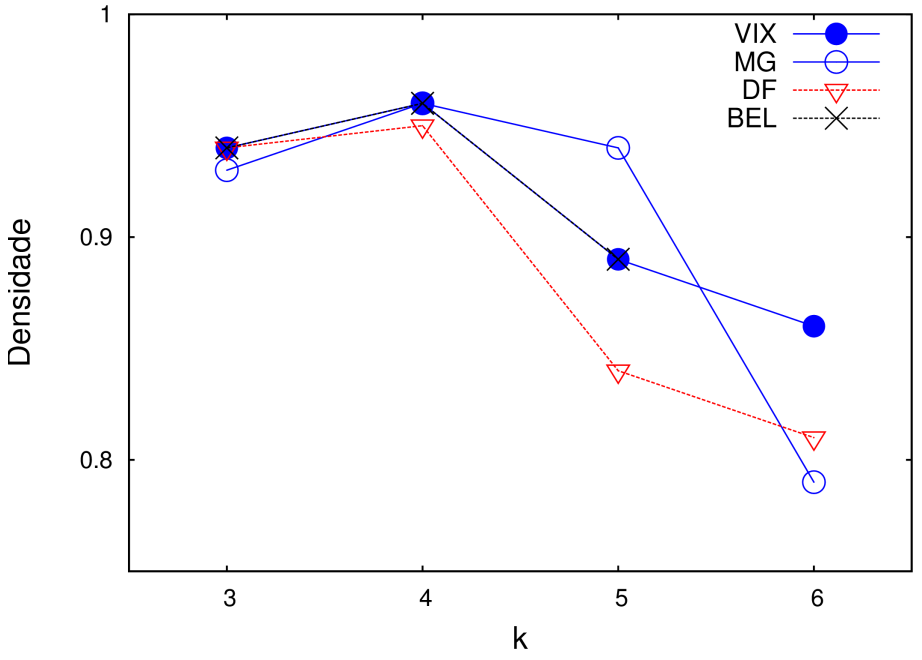

Figure 5b. Density

Results: k-Clique Communities

An interesting observation was the identification of just a few k_{max} communities, most of them tier-1 ISPs with ASN of 3 or 4 digits, while the communities with lower values of k are mostly composed of access providers (tier-2/3 ISPs) and have ASN of 5 digits.

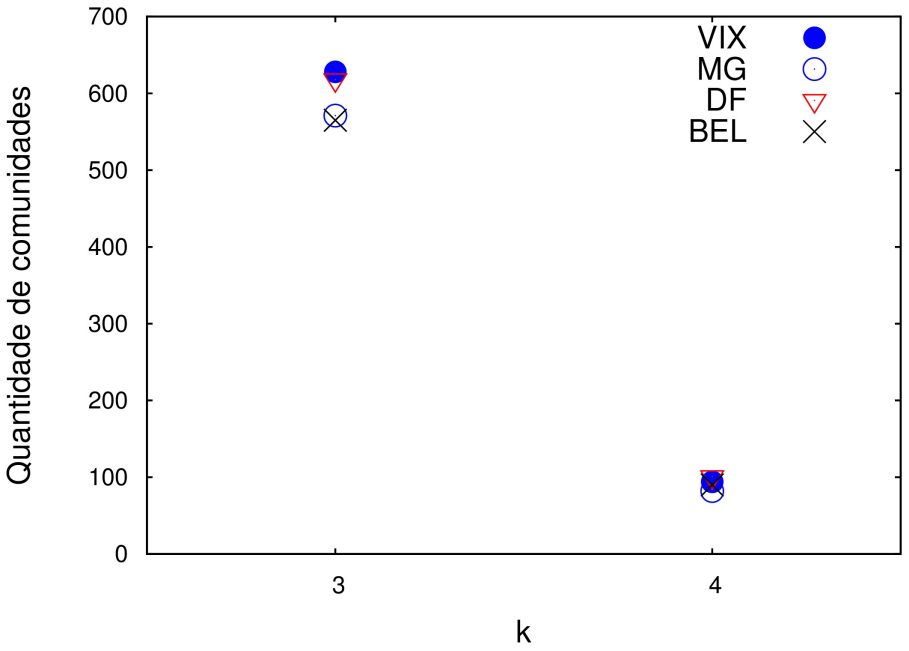

Figure 5a. k3 and k4 Communities

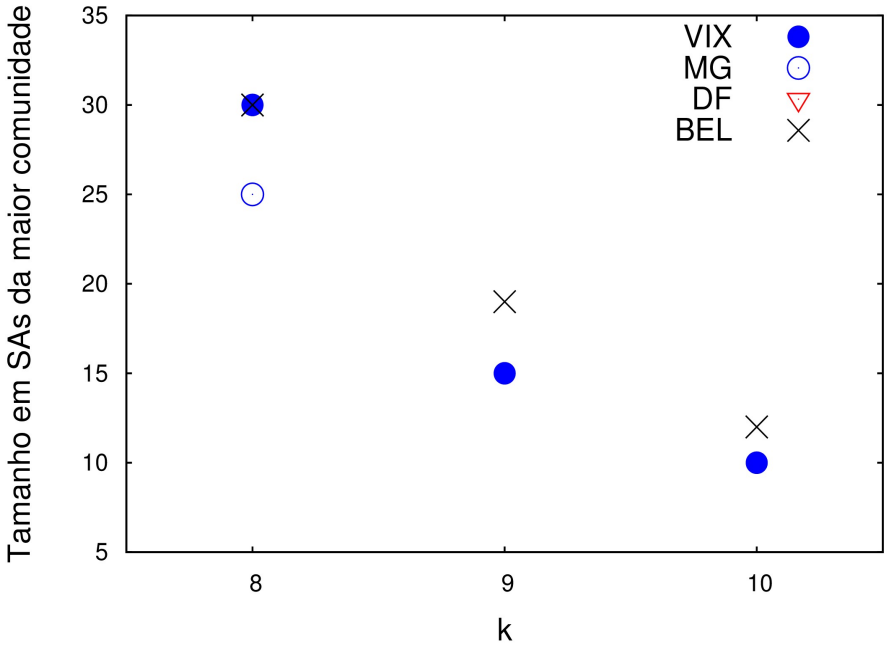

Figure 5c. Larger Communities

Agenda

- Introduction and Background
- Methodology
- Analysis and Results
- Conclusion
- Future Work
- Publications

Conclusion

- By reproducing the connectivity graphs of each Brazilian IXP we could realize the anatomy of this ecosystem in Brazil and we are publicly sharing our 2.5GB dataset with the community.

- For the best of our knowledge, this novel analysis was the first effort in the search for understanding the largest national ecosystem of public IXPs operating in a single country.

Future Work

- Once half of all 26 IXPs of PTTMetro did not allow the extraction of data, we requested NIC.br to make an internal collect to complement our dataset;

  → by the way, we were already attended by them! :-)

- Build a temporal analysis with new samples (snapshots)
- Verify the size of IPv4 and IPv6 prefixes
- Propose a new metric to measure the peering between ASes
- Enhance our study of communities to infer new results
- SDN Peering

Thanks! Obrigado!

shbbrito@dca.fee.unicamp.br

INTRIG
INFORMATION & NETWORKING TECHNOLOGIES RESEARCH & INNOVATION GROUP

http://www.github.com/intrig-unicamp/ixp-ptt-br