POWER-LAWS AND SPECTRAL ANALYSIS OF THE INTERNET TOPOLOGY

Laxmi Subedi

Communication Networks Laboratory
http://www.ensc.sfu.ca/~ljilja/cnl/
School of Engineering Science
Simon Fraser University

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Internet network

- Internet network is a complex network.
- Number of AS has increased approximately ten times over the last ten years.
- It is difficulty to develop representative model of the Internet topology.
- Power-laws and spectral analysis have been used to analyze the Internet topology.
- Properties of the Internet topology are useful:
  - to realistically model the Internet topology for protocols and algorithms evaluation and testing purposes
  - to develop new protocols, algorithms, and new network infrastructure

Project overview

- Analyze the properties of the Internet topology at Autonomous System (AS) level over the period of five years (2003-2008)
- Border Gateway Protocol (BGP) routing datasets collected by:
  - Route Views
  - RIPE (Réseaux IP Européens)
- Method:
  - analysis of power-laws
  - spectral analysis

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Autonomous Systems (ASes)

- ASes:
  - groups of networks sharing the same routing policy
  - identified with Autonomous System Numbers (ASN)
  - ASN assigned by IANA
- Internet topology on AS-level:
  - an arrangement of ASes and their interconnections
- Analyzing the Internet topology and finding properties of associated graphs rely on mining data and capturing information about ASes.

IANA: Internet Assigned Number Authority
http://www.iana.org/assignments/as-numbers

Border Gateway Protocol (BGP)

- Routing table of a BGP router contains AS path information.
- The BGP router uses BGP protocol:
  - inter-AS protocol
  - used to exchange network reachability information among BGP systems
  - reachability information is stored in routing tables

Internet AS-level data

BGP routing tables are collected by:

- Route Views:
  - most participating ASes reside in North America
  - routing data collection process began in 1997
- RIPE:
  - most participating ASes reside in Europe
  - routing data collection process began in 1999
- The BGP routing tables are collected from multiple geographically distributed BGP Cisco routers and Zebra servers.

http://www.routeviews.org
http://www.ripe.net/ris

Internet topology at AS-level

- AS-level datasets from Route Views and RIPE have been extensively used to analyze the Internet topology.

	Route Views	Ripe
Faloutsos, 1999	✓	✗
Chang, 2001	✓	✓
Vukadinovic, 2001	✓	✗
Gkantsidis, 2003	✓	✓

Analyzed datasets

- We analyzed datasets collected at 00:00 am on July 31, 2003 and 00:00 am on July 31, 2008.
- Sample datasets:
  - Route Views:

    TABLE_DUMP| 1050122432| B| 204.42.253.253| 267| 3.0.0.0/8| 267 2914 174 701| IGP| 204.42.253.253| 0| 0| 267:2914 2914:420 2914:2000 2914:3000| NAG| |
  - RIPE:

    TABLE_DUMP| 1041811200| B| 212.20.151.234| 13129| 3.0.0.0/8| 13129 6461 7018 | IGP| 212.20.151.234| 0| 0| 6461:5997 13129:3010| NAG| |

Internet topology at AS level

- Datasets collected from Border Gateway Protocols (BGP) routing tables are used to infer the Internet topology at AS-level.

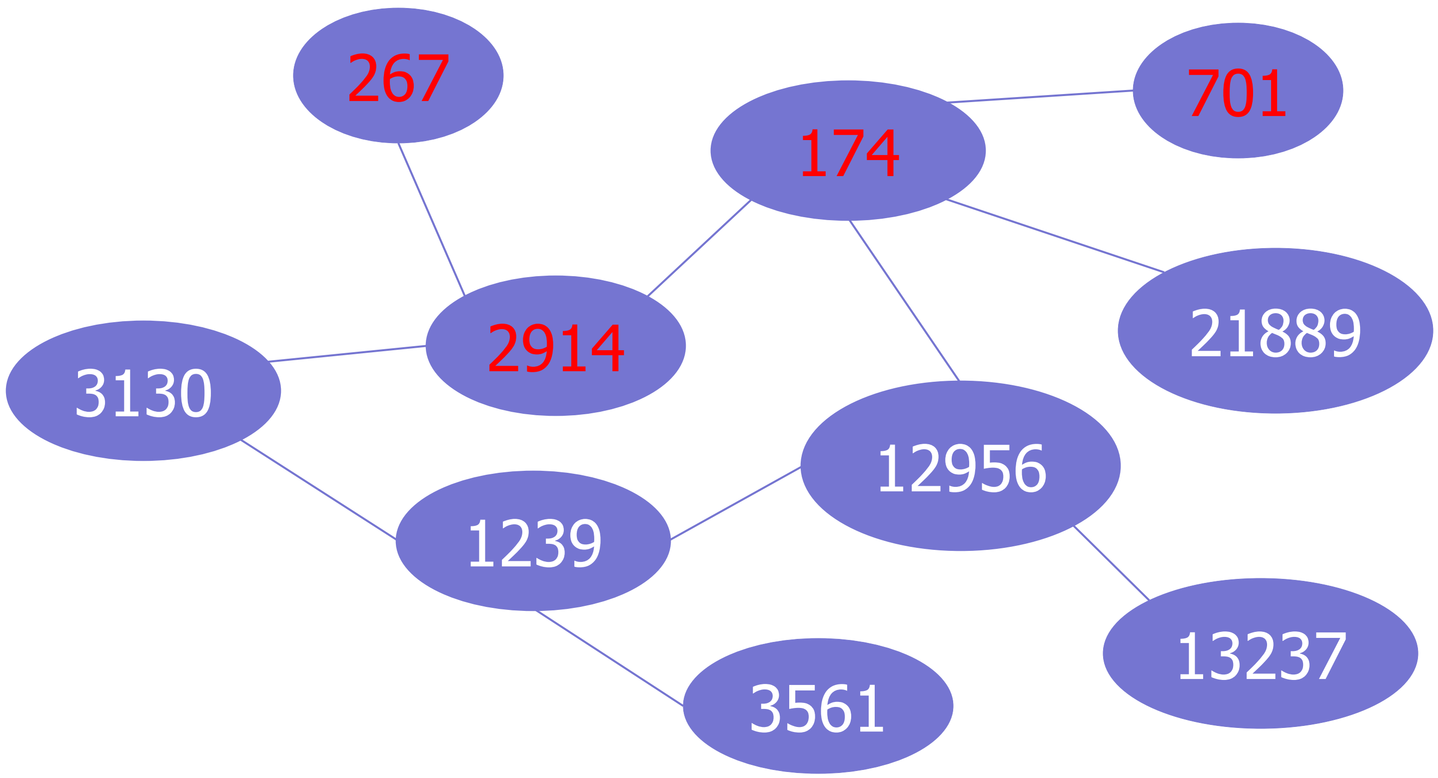

Internet topology and matrices

- Adjacency matrix A(G):

$$A(i,j) = \begin{cases} 1 & \text{if } i \text{ and } j \text{ are adjacent} \\ 0 & \text{otherwise,} \end{cases}$$

where i and j are two nodes.

- Normalized Laplacian matrix NL(G):

$$NL(i,j) = \begin{cases} 1 & \text{if } i = j \text{ and } d_i \neq 0 \\ -\frac{1}{\sqrt{d_i d_j}} & \text{if } i \text{ and } j \text{ are adjacent} \\ 0 & \text{otherwise,} \end{cases}$$

where d_i and d_j are degrees of node i and j, respectively.

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Power-laws

- Power-laws are expressed in the form of:

$$y \propto x^{a},$$

where y and x are the measures of interest and a is a constant.

- The Internet topology is characterized by the presence of various power-laws observed when considering:
  - node degree vs. node rank
  - frequency of node degree vs. node degree
  - CCDF of node degree vs. node degree
  - eigenvalues of the adjacency matrix vs. the order of the eigenvalues

M. Faloutsos, P. Faloutsos, and C. Faloutsos, 1999

Power-laws: node degree vs. rank

- Node degree is the number of edges incident to a node.
- The node degree power-law implies:

$$d_v \propto r_v^R,$$

where d_v is the degree of a node v, r_v is the rank of the node, and R is the exponent of the node degree power-law.

Power-laws: frequency of node degree vs. node degree

- The frequency of a node degree is equal to the number of nodes having the same degree.
- The frequency of node degree power-law implies:

$$f_d \propto d^O,$$

where f_d is the frequency of degree d, d is a node degree, and O is the exponent of the frequency of node degree power-law.

Power-laws: CCDF of node degree vs. node degree

- The complementary cumulative distribution function CCDF is defined as:

$$F_c(x) = P(X > x),$$

  where $P(X>x)$ is the probability that the random variable X has a value greater than x.

- The CCDF of node degree vs. node degree power-law implies:

$$D_d \propto d^D,$$

  where D_d is the CCDF of a node degree d and D is the CCDF power-law exponent.

Power-laws: eigenvalue vs. index

- The power-law for the adjacency matrix implies:

$$\lambda_{ai} \propto i^{\varepsilon},$$

where λ_{ai} is an eigenvalue of the adjacency matrix associated with the increasing sequence of numbers i and ε is the power-law exponent.

- The power-law for the normalized Laplacian matrix implies:

$$\lambda_{Li} \propto i^{L},$$

where λ_{Li} is an eigenvalue of the normalized Laplacian matrix associated with the increasing sequence of numbers i and L is the power-law exponent.

Spectrum of a graph

- Spectrum of a graph is:
  - set of eigenvalues of a matrix
  - closely related to certain graph invariants
  - associated with topological characteristics of the network such as number of edges, connected components, presence of cohesive clusters

- If x is an n-dimensional real vector, then x is called the eigenvector of matrix A with eigenvalue λ if and only if it satisfies:

$$Ax = \lambda x,$$

  where λ is a scalar quantity.

Spectrum of a graph

- The number of times 0 appears as an eigenvalue of the Laplacian matrix is equal to the number of connected components in a graph.
- Algebraic connectivity, the second smallest eigenvalue of a normalized Laplacian matrix is:
  - related to the connectivity characteristic of a graph
- Elements of the eigenvector corresponding to the largest eigenvalue of the normalized Laplacian matrix tend to be positioned close to each other if they correspond to AS nodes with similar connectivity patterns constituting clusters.

M. Fiedler, 1973

D. Vukadinovic, P. Huang, and T. Erlebach, 2001

Spectrum of a graph

- The eigenvectors corresponding to large eigenvalues contain information relevant to clustering.
- Large eigenvalues and the corresponding eigenvectors provide information suggestive to the intracluster traffic patterns of the Internet topology.
- We consider both the adjacency and the normalized Laplacian matrices.

C. Gkantsidis, M. Mihail, and E. Zegura, 2003

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Power-laws and linear regression line

- The power-law exponents are calculated from the linear regression lines:

$$y = 10^b \times x^a,$$

with segment b and slope a and when plotted on a log-log scale.

$$y = 10^b \times x^a$$

$$\log y = \log 10^b + \log x^a$$

$$\log y = b + a \times \log x$$

$$y' = b + ax'$$

Observed power-laws

- Calculated and plotted on a log-log scale are:
  - node degree vs. node rank
  - CCDF of node degree vs. node degree
  - eigenvalues vs. index
- Least square approximation is used to obtain the linear regression line.
- The correlation coefficient is calculated between the regression line and the plotted data.

Power laws: node degree vs. rank

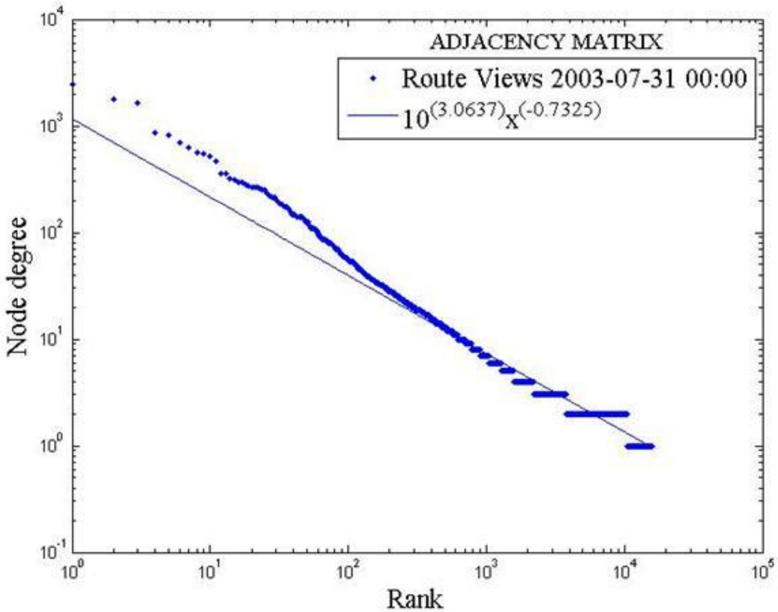

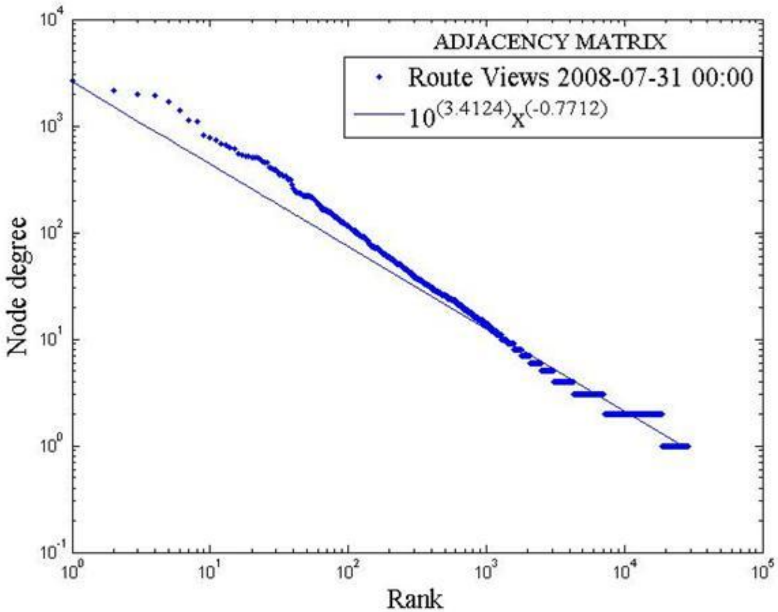

- Route Views 2003 datasets: R= –0.7325 and r= –0.9661
- Route Views 2008 datasets: R= –0.7712 and r= –0.9721

R= power-law exponent; r= correlation coefficient

Power laws: node degree vs. rank

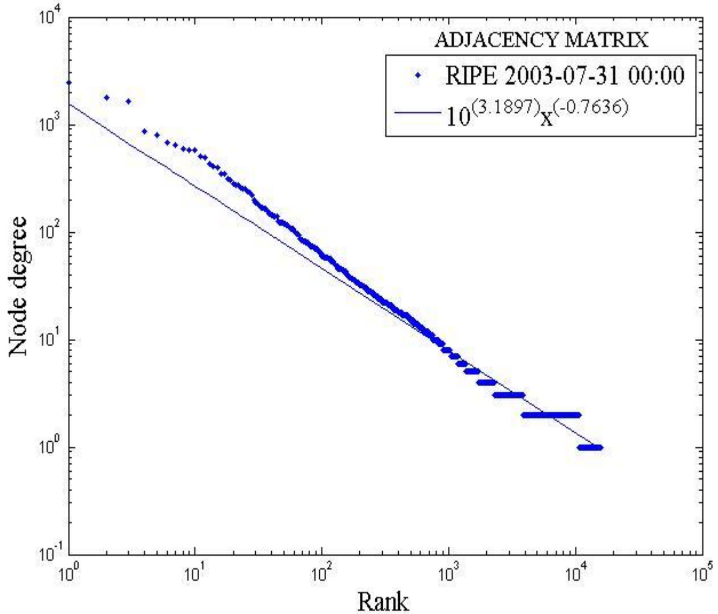

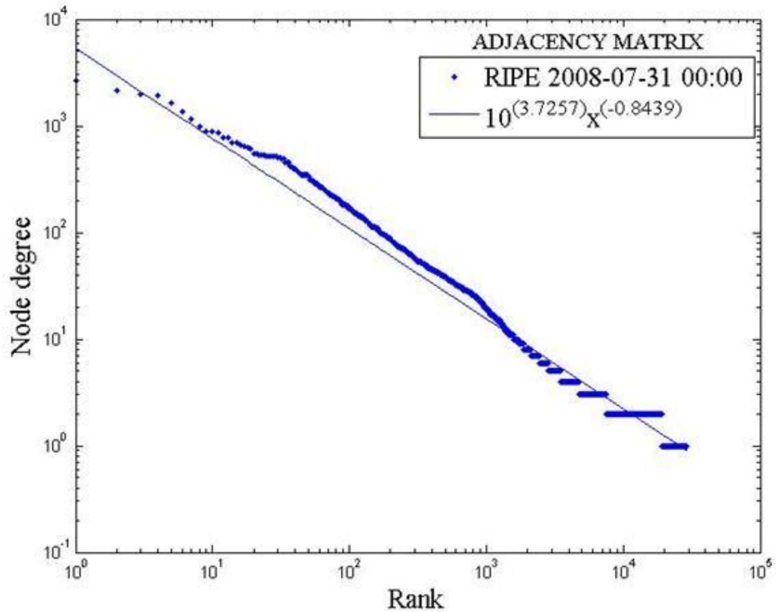

- RIPE 2003 datasets: R= –0.7636 and r= –0.9687
- RIPE 2008 datasets: R= –0.8439 and r= –0.9744

R= power-law exponent; r= correlation coefficient

Confidence intervals

- Six samples is randomly selected from Route Views and RIPE 2003 and 2008 datasets.
- Each dataset is smaller than 30, with unknown standard deviation.
- T-distribution is used to predict the confidence interval at 95% confidence level:

$$\overline{X} - t_{x/2}(s/\sqrt{n}) < \mu < \overline{X} + t_{x/2}(s/\sqrt{n}),$$

X: sample mean

$t_{x/2}$: t-distribution

s: sample standard deviation

n: number of samples

μ: population mean

Confidence interval: node degree vs. rank

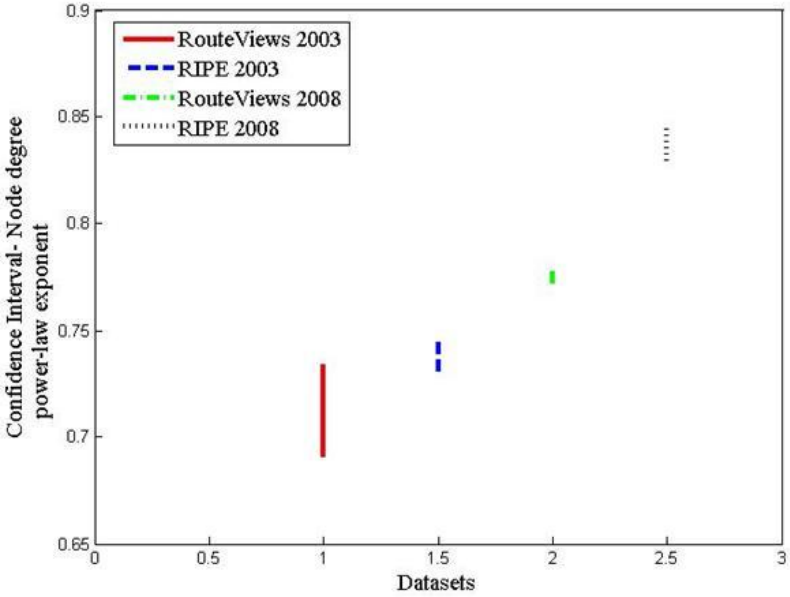

- r> 96% for all datasets

r= correlation coefficient

Power laws: CCDF of node degree vs. node degree

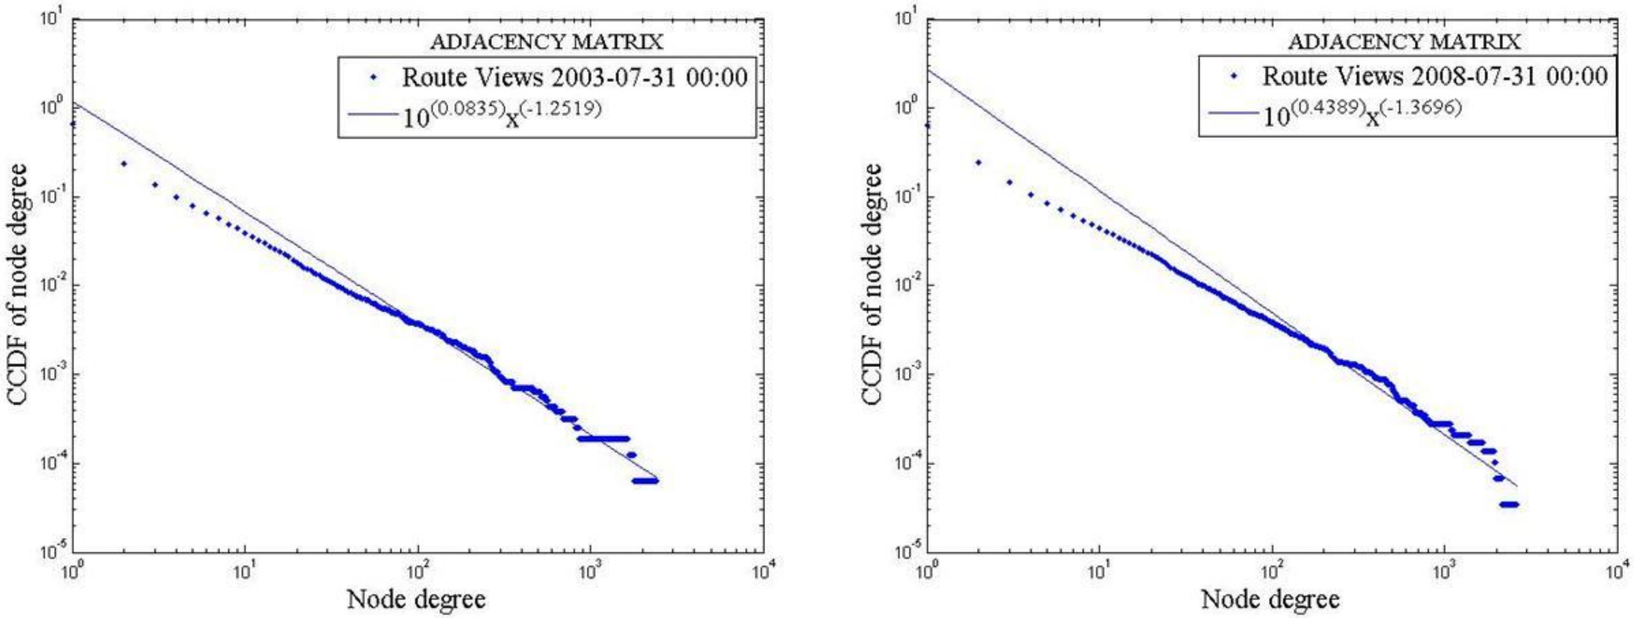

- Route Views 2003 datasets: D= –1.2519 and r= –0.9810
- Route Views 2008 datasets: D= –1.3696 and r= –0.9626

D= power-law exponent; r= correlation coefficient

Power laws: CCDF of node degree vs. node degree

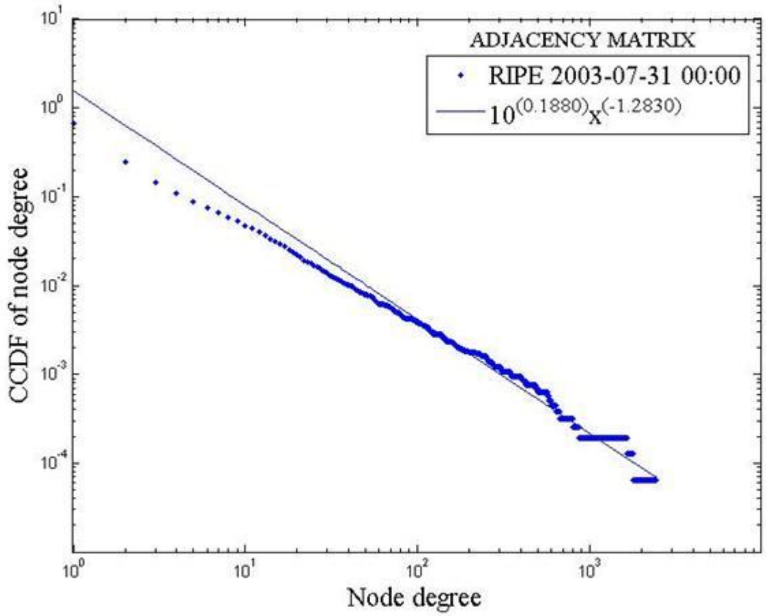

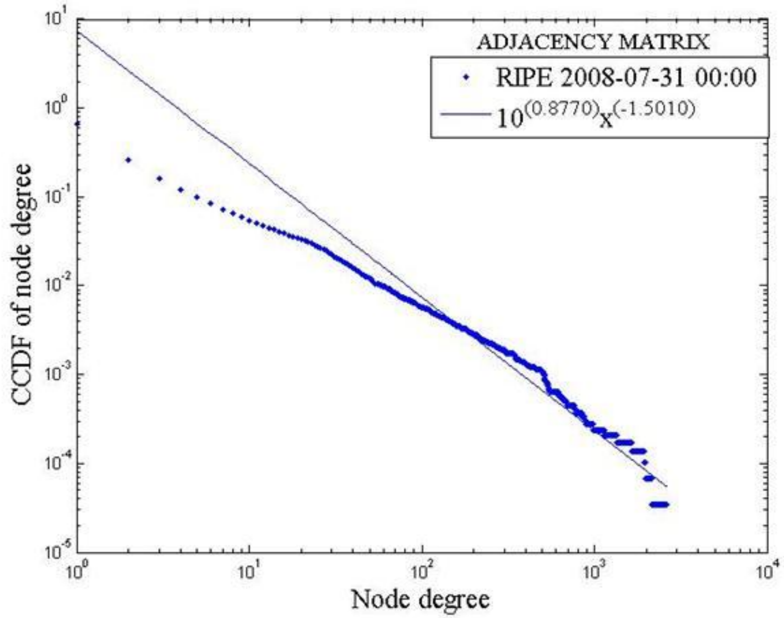

- RIPE 2003 datasets: D= –1.2830 and r= –0.9812
- RIPE 2008 datasets: D= –1.5010 and r= –0.9676

D= power-law exponent; r= correlation coefficient

Confidence interval: CCDF of node degree vs. rank

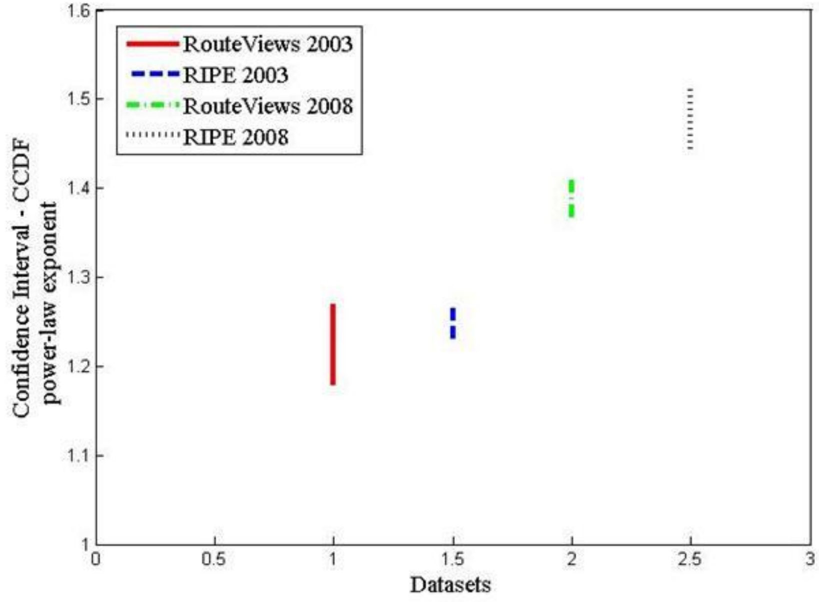

- r> 90% for all datasets

r= correlation coefficient

Eigenvalues of the adjacency matrix

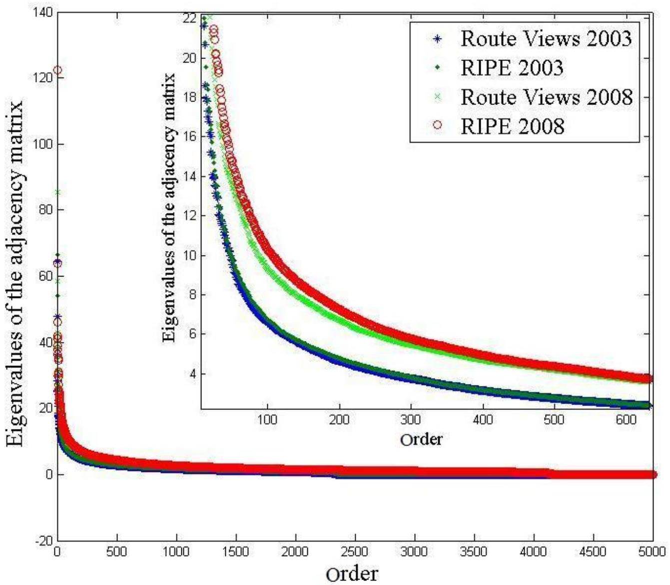

order	Route Views 2003	Route Views 2008	RIPE 2003	RIPE 2008
1	64.30	85.43	66.65	122.28
2	47.75	58.56	54.19	63.94
3	38.15	42.77	38.24	46.14
4	36.23	40.85	36.14	41.98
5	29.88	39.69	31.21	41.08
6	28.50	37.85	27.38	38.93
7	25.47	36.21	26.41	37.94
8	25.06	34.66	25.06	36.47
9	24.13	31.58	23.86	35.08
10	22.51	29.34	23.32	34.47
11	21.61	27.40	22.02	30.97
12	20.69	25.69	21.77	30.54
13	18.58	25.00	20.75	29.68
14	17.94	24.82	19.55	27.03
15	17.78	23.89	18.67	25.74
16	17.31	23.69	18.42	25.35
17	16.99	22.81	17.85	24.83
18	16.75	22.46	17.44	24.30
19	16.22	22.04	17.24	24.06
20	16.01	21.36	16.63	24.00

Power laws: eigenvalues vs. index

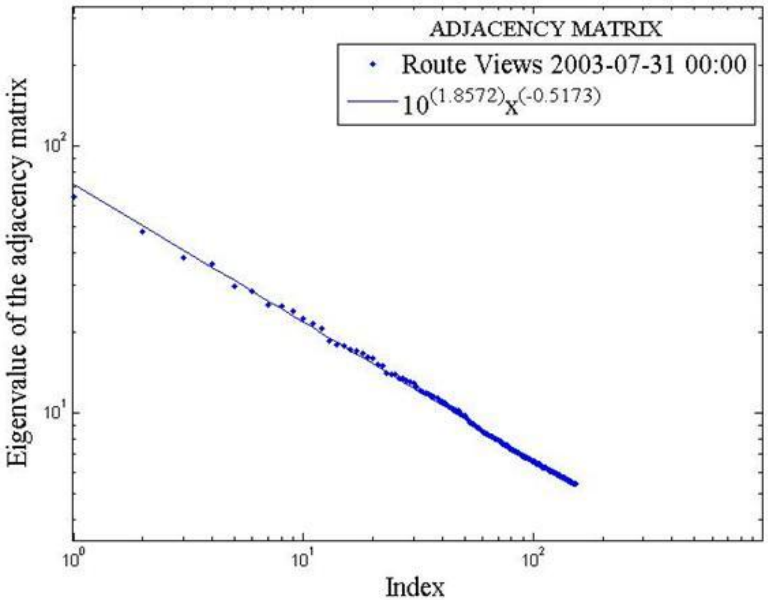

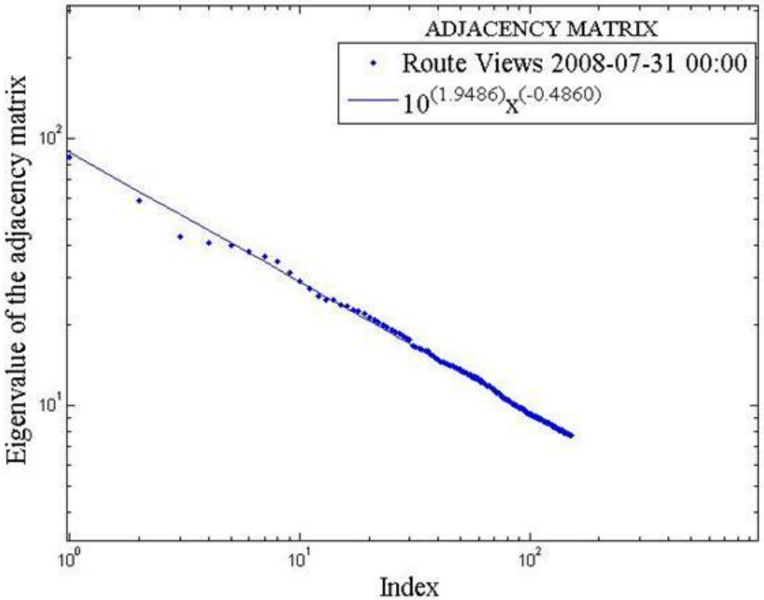

Adjacency matrix:

- Route Views 2003 datasets: ε= –0.5713 and r= –0.9990
- Route Views 2008 datasets: ε= –0.4860 and r= –0.9982

ε= power-law exponent; r= correlation coefficient

Power laws: eigenvalues vs. index

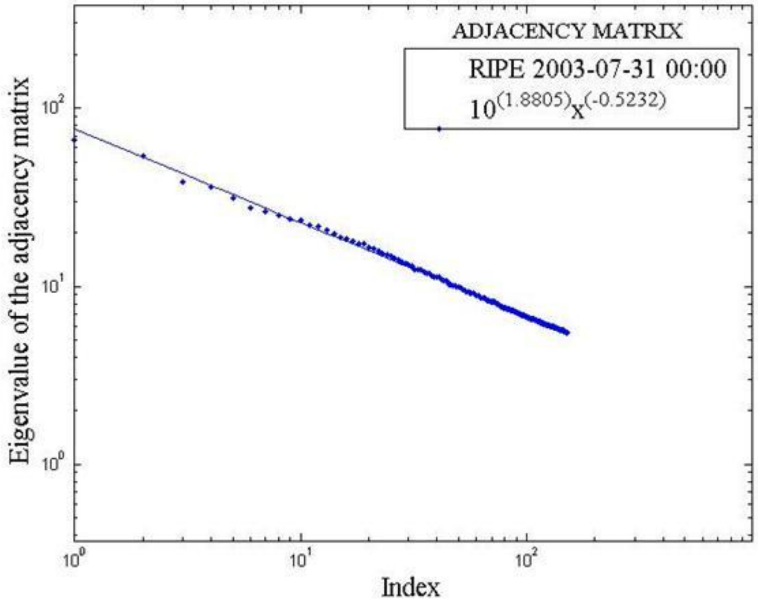

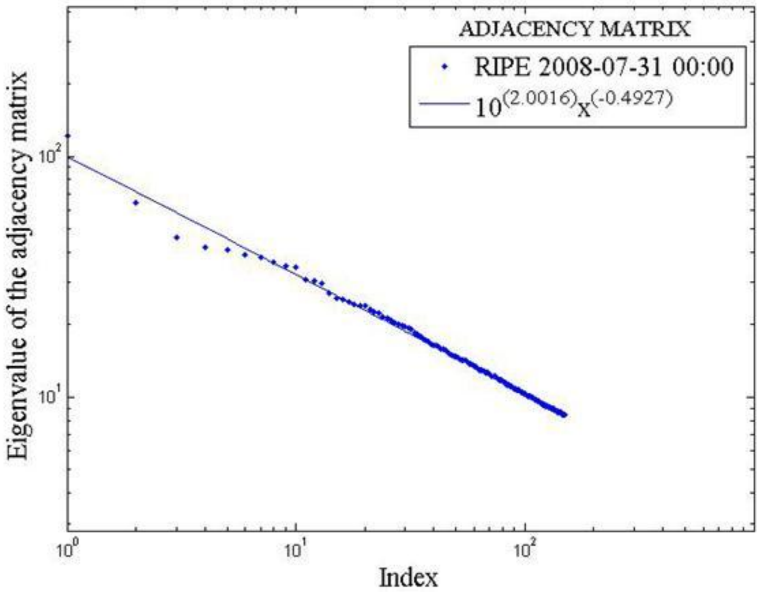

Adjacency matrix:

- RIPE 2003 datasets: ε= –0.5232 and r= –0.9989
- RIPE 2008 datasets: ε= –0.4927 and r= –0.9970

ε= power-law exponent; r= correlation coefficient

Confidence interval: eigenvalues vs. index

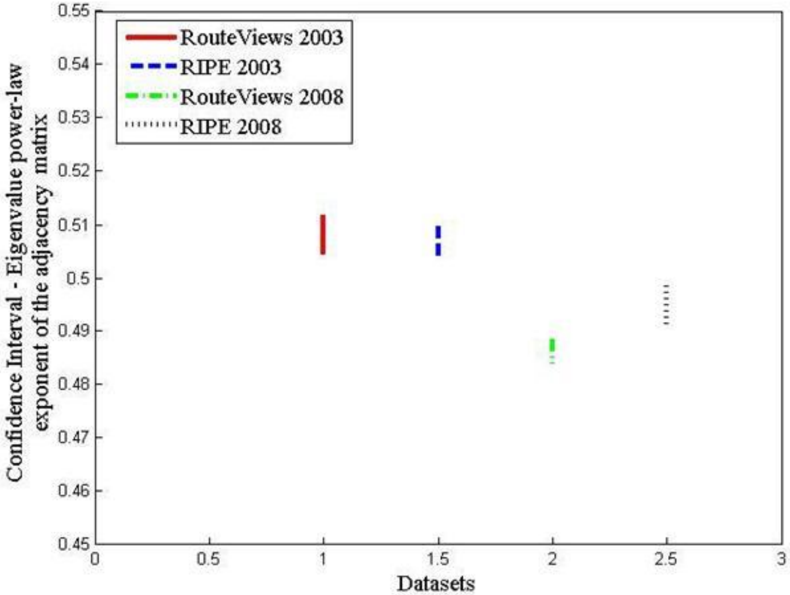

Adjacency matrix:

- r> 99% for all datasets

r= correlation coefficient

Power laws: eigenvalues vs. index

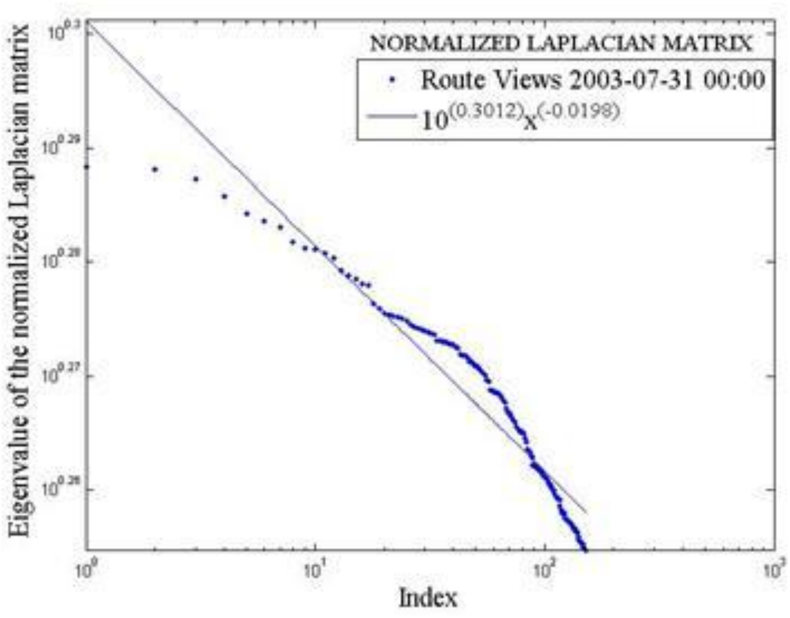

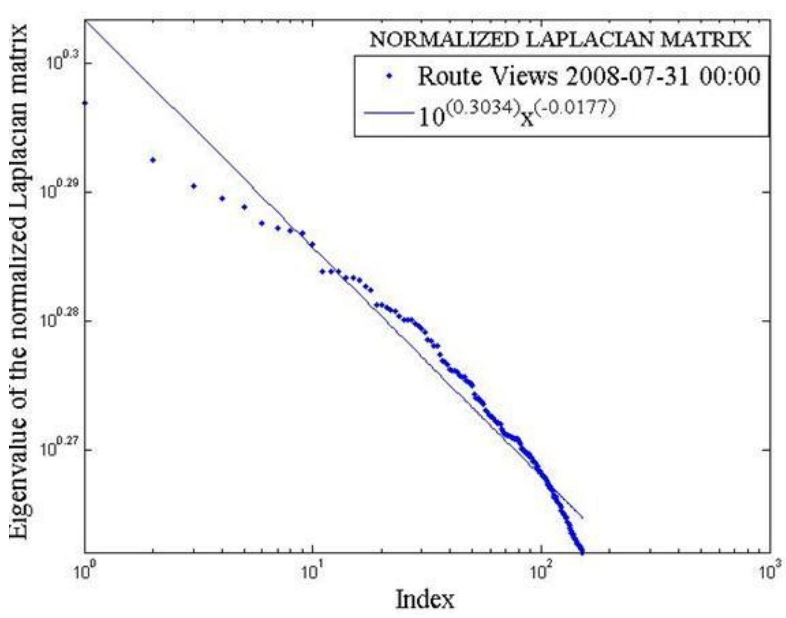

Normalized Laplacian matrix:

- Route Views 2003 datasets: L= –0.0198 and r= –0.9564
- Route Views 2008 datasets: L= –0.0177 and r= –0.9782

L= power-law exponent; r= correlation coefficient

Power laws: eigenvalues vs. rank

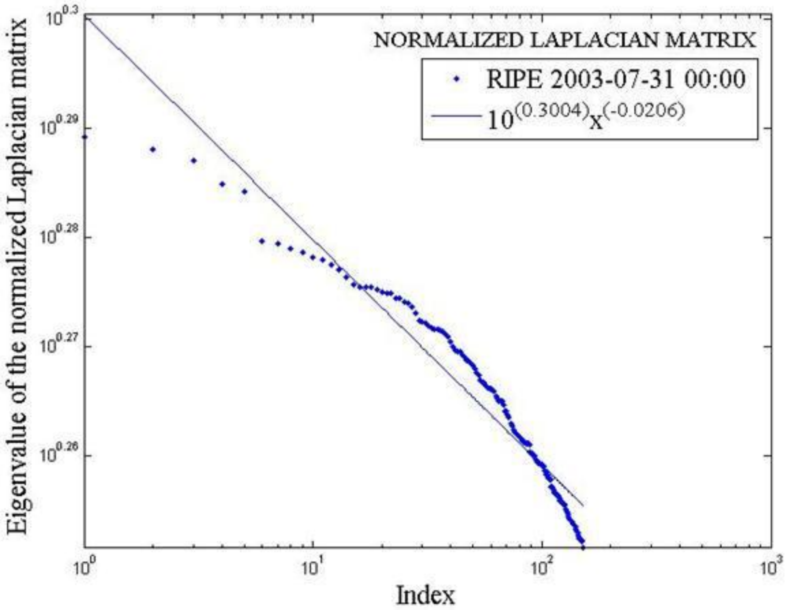

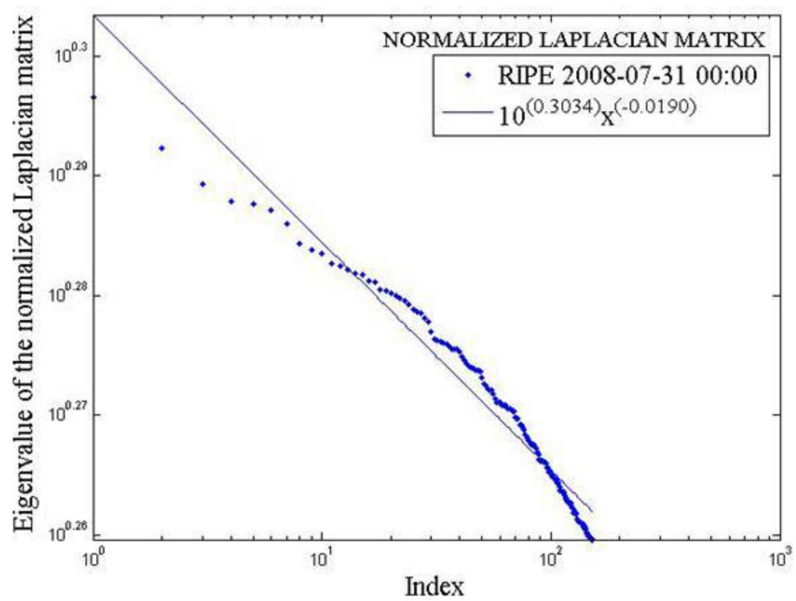

Normalized Laplacian matrix:

- RIPE 2003 datasets: L= –0.5232 and r= –0.9989
- RIPE 2008 datasets: L= –0.4927 and r= –0.9970

L= power-law exponent; r= correlation coefficient

Confidence interval: eigenvalues vs. rank

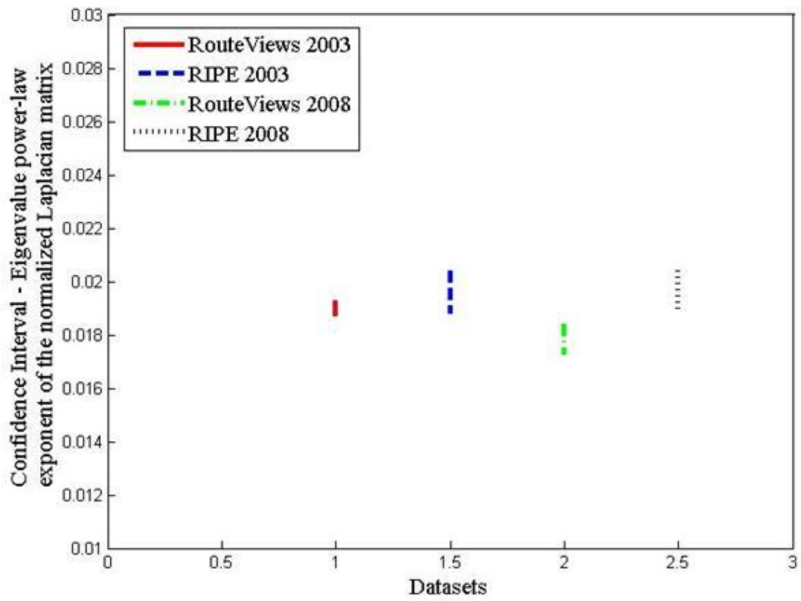

Normalized Laplacian matrix:

- r> 95% for all datasets

r= correlation coefficient

Power-laws: summary

- A high correlation coefficient between the regression line and the plotted data indicates the existence of a power-law.
- Results imply that the node degree, CCDF of node degree, and eigenvalues follow a power-law dependency on the rank, node degree, and index, respectively.
- Power-laws exponents have not substantially changed over the years.

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Clusters of connected ASes: Route Views

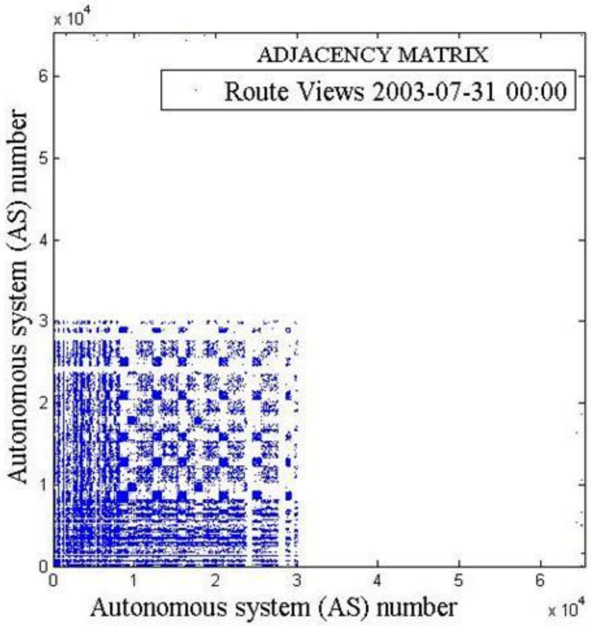

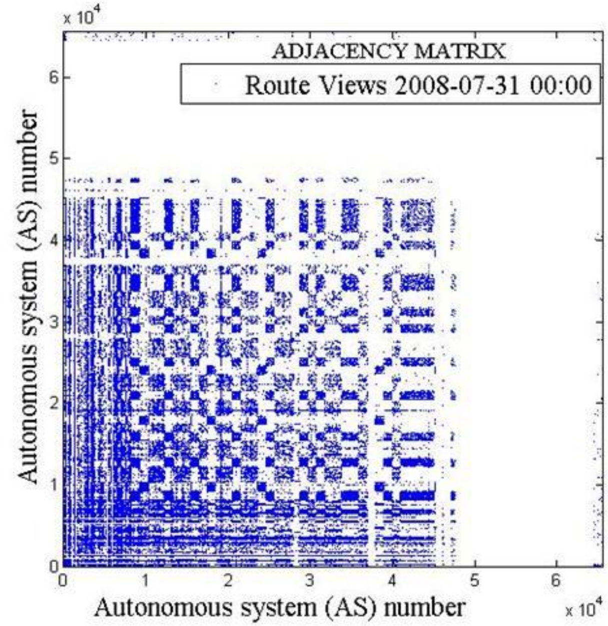

- A dot in the position (x, y) represents the connection patterns between AS nodes.
- Existence of higher connectivity inside a particular cluster and relatively lower connectivity between clusters is visible.

Clusters of connected ASes: Route Views

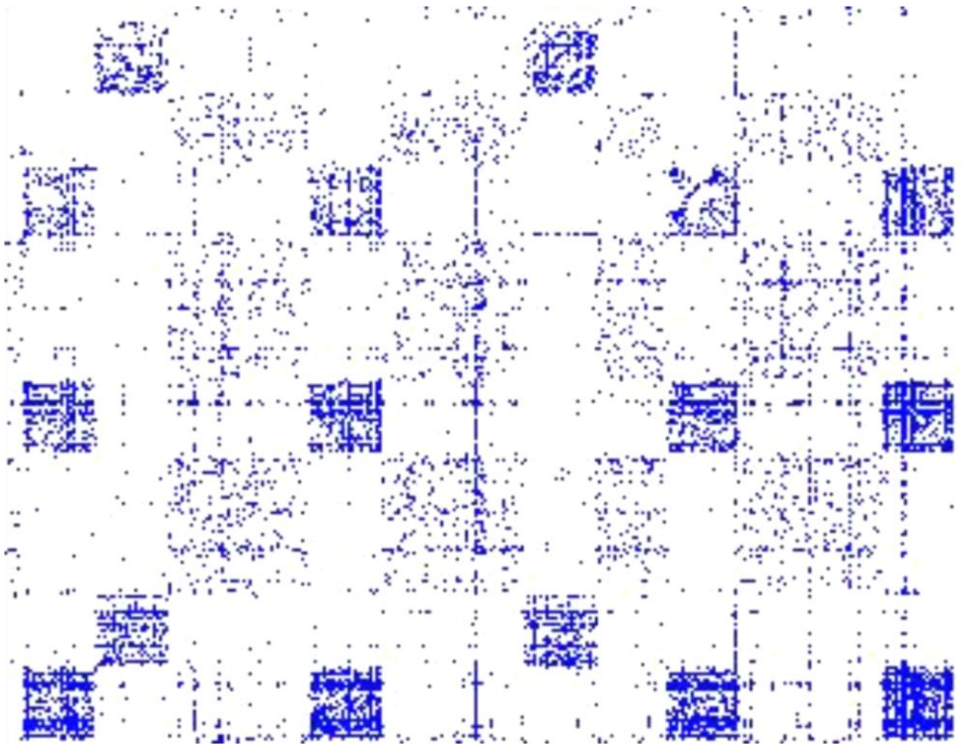

Zoomed view of Route Views 2008 datasets.

Clusters of connected ASes: RIPE

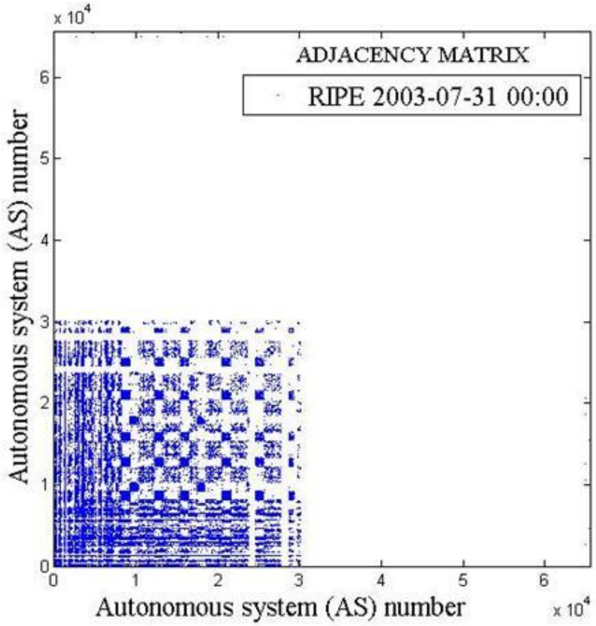

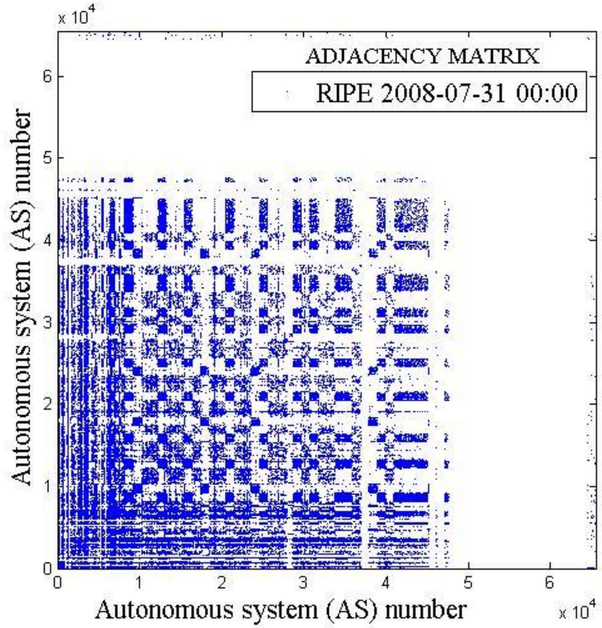

- Similar pattern for Route Views and RIPE 2003 and 2008 datasets

Roadmap

- Introduction
- Internet topology and graph theory
- Power-laws and spectrum of a graph
- Internet routing and BGP datasets
- Power-laws analysis
- Spectral analysis
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Connectivity status: example

- The second smallest eigenvector: 0.35, -0.35, -0.35, 0.41, 0.50, 0.61
- N1(0.35), N2(-0.35), N3(-0.35), N4(0.41), N5(0.50), N6(0.61)
- Sort ASs by element value: N2, N3, N1, N4, N5, N6
- Except N4, all other nodes are connected

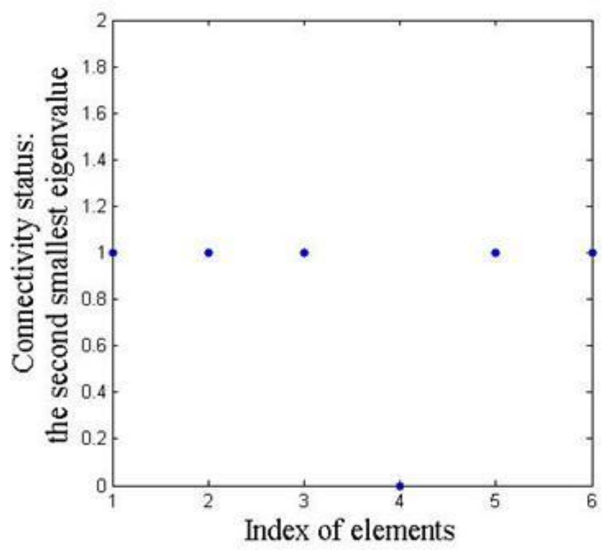

The second smallest eigenvalue: Route Views

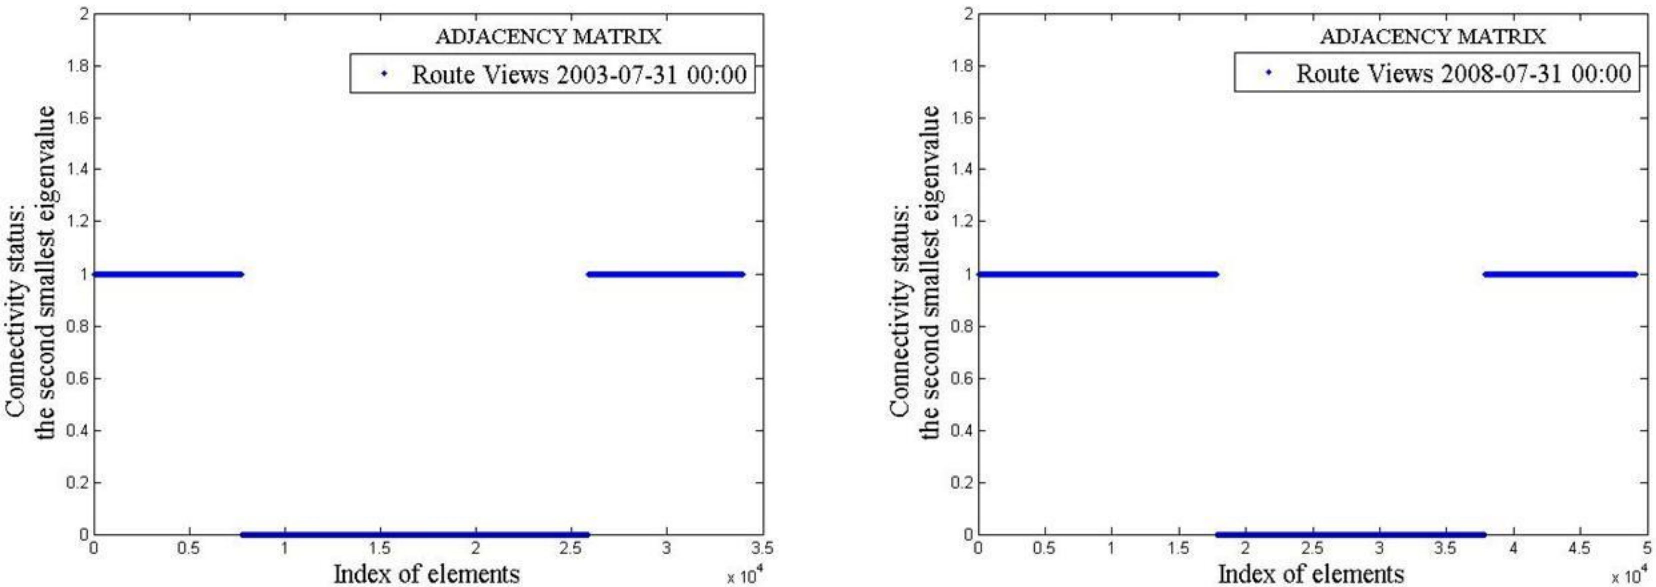

The connectivity status based on the second smallest eigenvalue of the adjacency matrix indicates:

- the connectivity status for Route Views 2003 datasets differs with Route Views 2008 datasets

The second smallest eigenvalue: RIPE

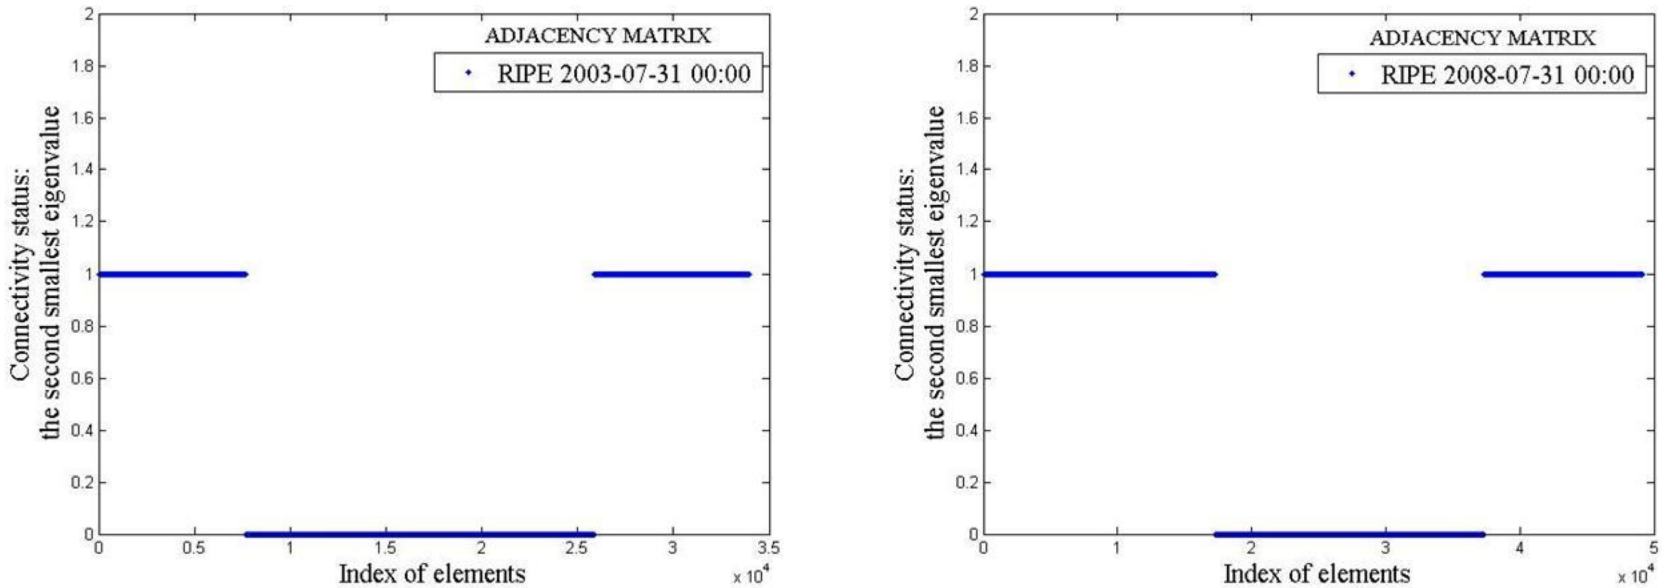

The connectivity status based on the second smallest eigenvalue of the adjacency matrix indicates:

- the connectivity status for RIPE 2003 datasets differs with RIPE 2008 datasets

The largest eigenvalue: Route Views

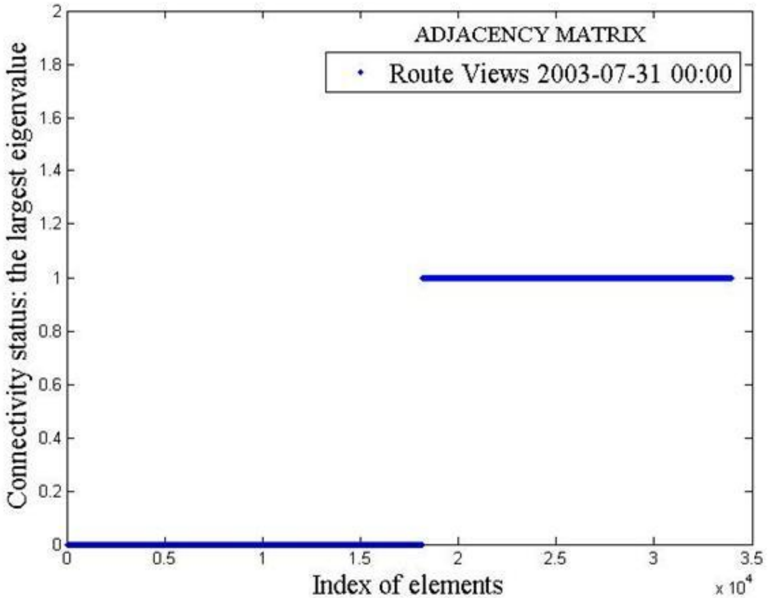

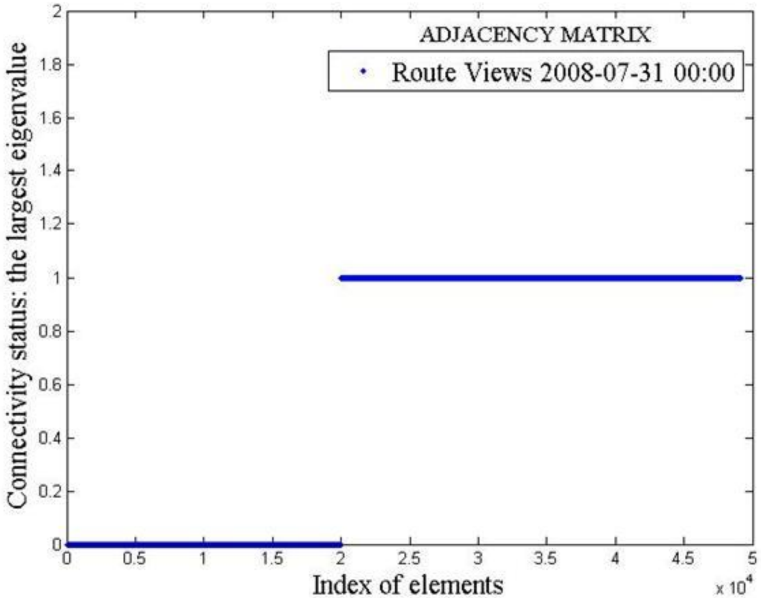

The connectivity status based on the largest eigenvalue of the adjacency matrix indicates:

- the connectivity status for Route Views 2003 differs with Route Views 2008 datasets

The largest eigenvalue: RIPE

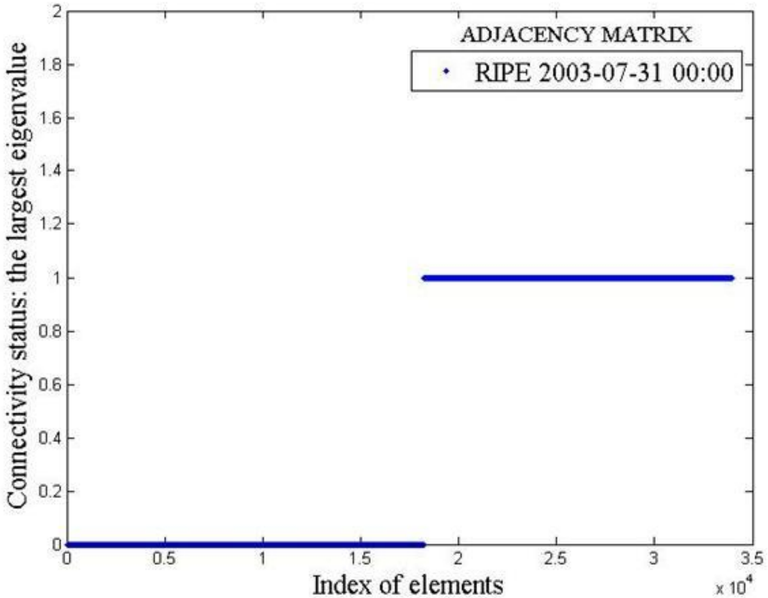

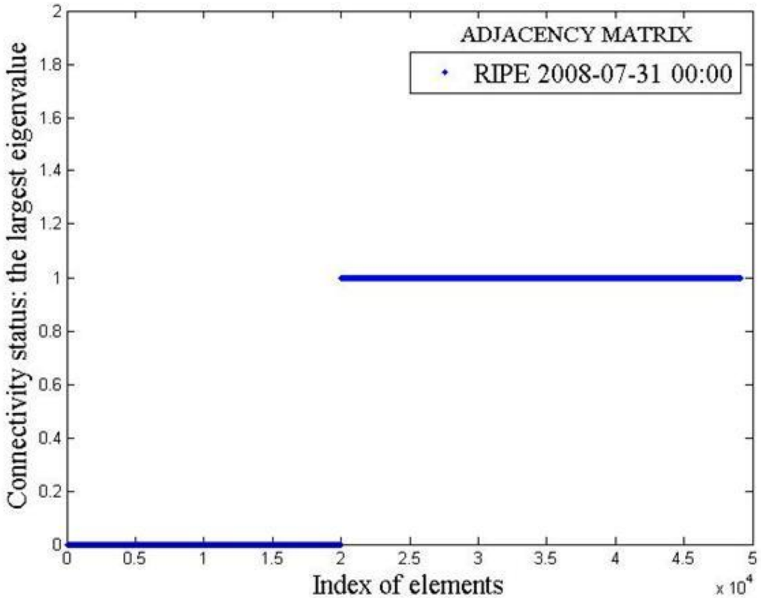

The connectivity status based on the largest eigenvalue of the adjacency matrix indicates:

- the connectivity status for RIPE 2003 differs with RIPE 2008 datasets

The second smallest eigenvalue: Route Views

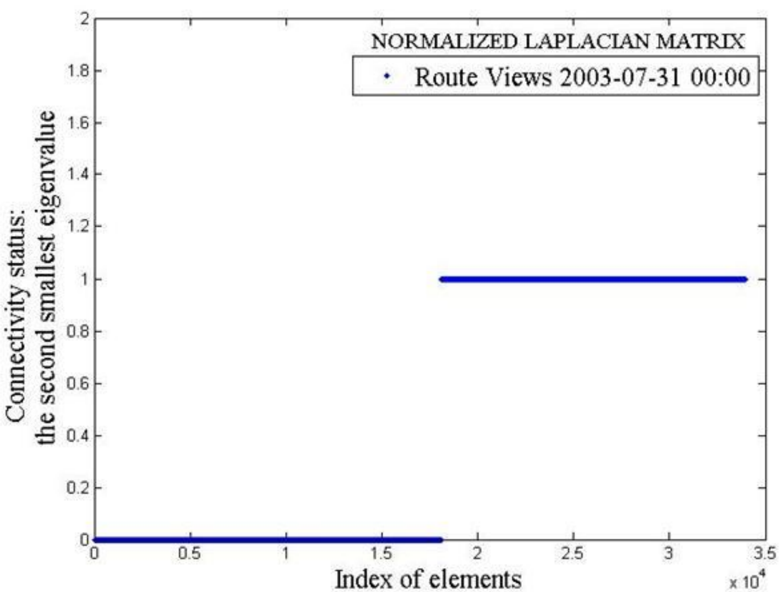

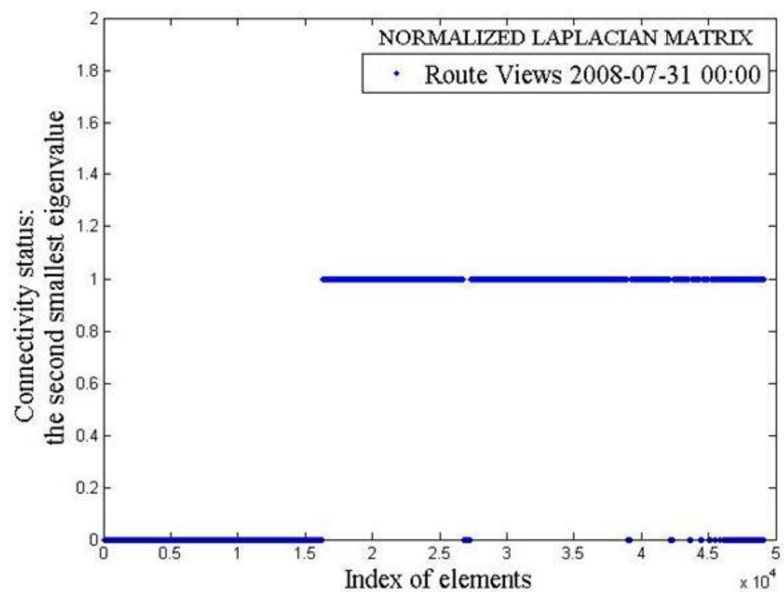

The connectivity status based on the second smallest eigenvalue of the normalized Laplacian matrix indicates:

- the connectivity status for Route Views 2003 differs with Route Views 2008 datasets

The second smallest eigenvalue: RIPE

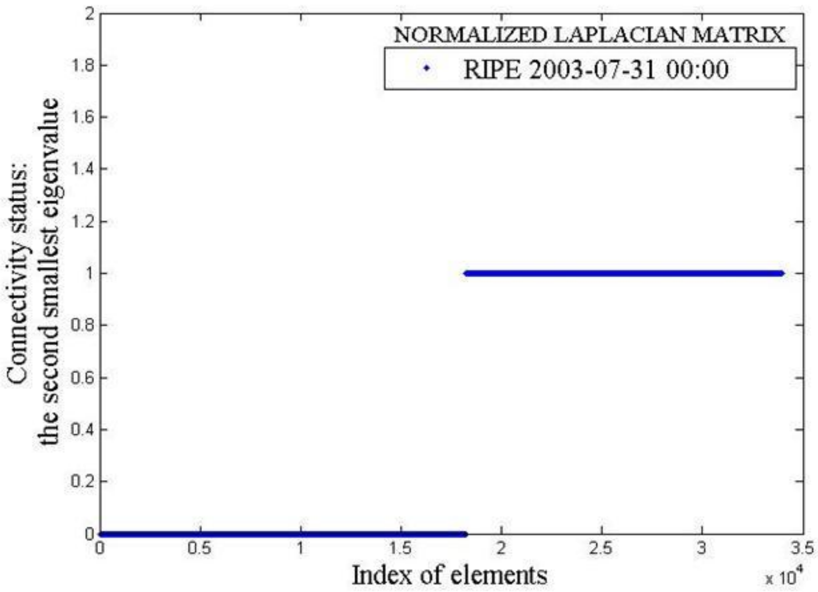

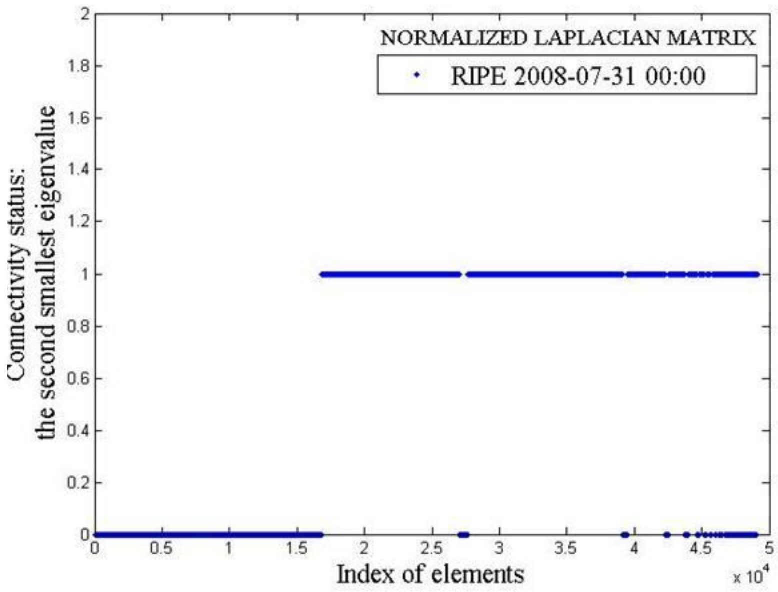

The connectivity status based on the second smallest eigenvalue of the normalized Laplacian matrix indicates:

- the connectivity status for RIPE 2003 differs with RIPE 2008 datasets

The largest eigenvalue: Route Views

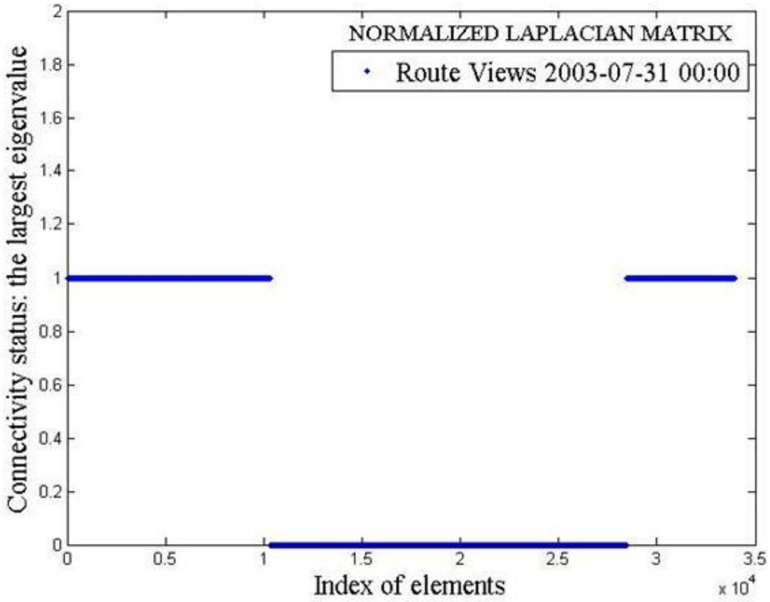

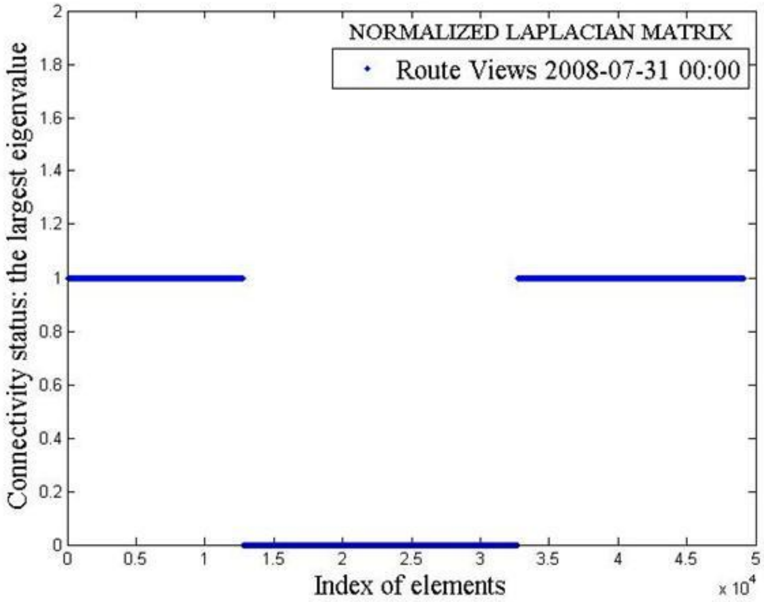

The connectivity status based on the largest eigenvalue of the normalized Laplacian matrix indicates:

- the connectivity status for Route Views 2003 differs with Route Views 2008 datasets

The largest eigenvalue: RIPE

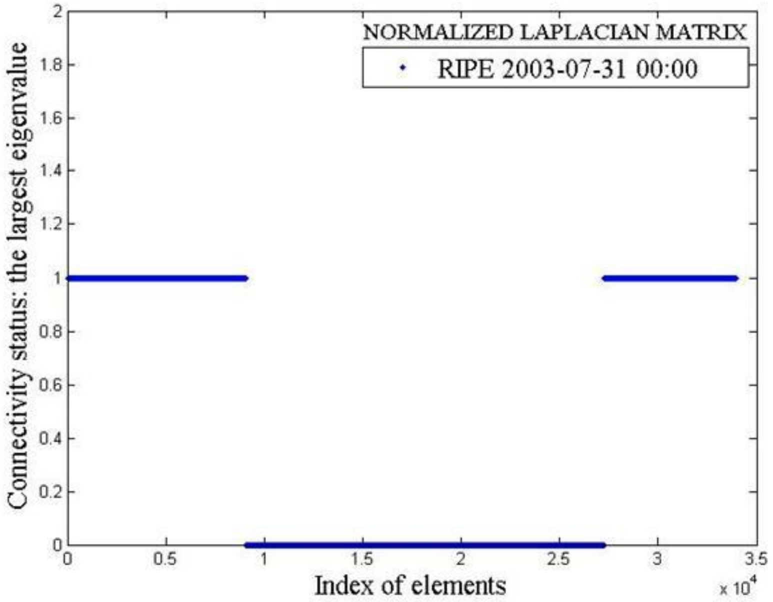

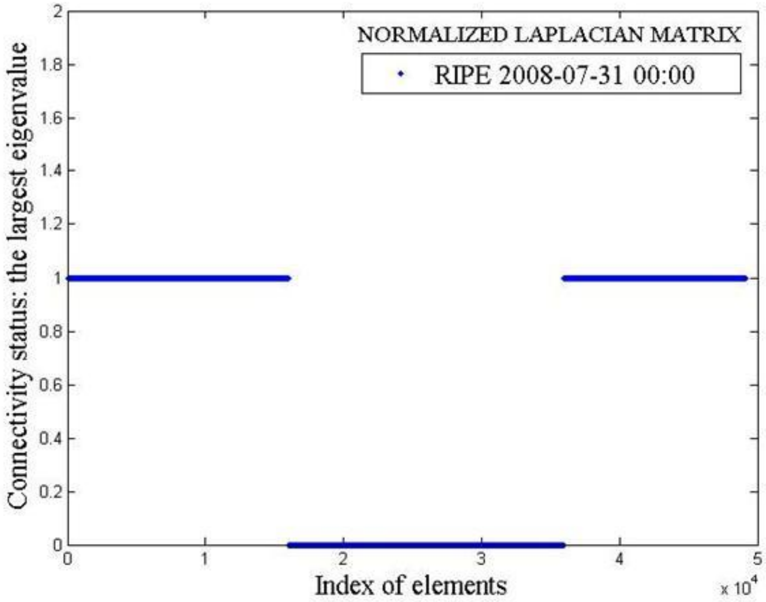

The connectivity status based on the largest eigenvalue of the normalized Laplacian matrix indicates:

- the connectivity status for RIPE 2003 differs RIPE 2008 datasets

Connectivity status: summary

- The second smallest and the largest eigenvalues of both the adjacency and the normalized Laplacian matrix revealed:
  - the connectivity status is different for Route Views 2003 and 2008 datasets
  - the connectivity status is different for RIPE 2003 and 2008 datasets
  - the connectivity status is similar for Route Views and RIPE 2003 and for Route Views and RIPE 2008 datasets
- Connectivity status based on the second smallest eigenvalue of the adjacency matrix is similar to the largest eigenvalue of the normalized Laplacian matrix, and vice versa.
  - this property has its basis in the spectral properties of two matrices since L= D–A

Roadmap

- Introduction
- Internet topology and graph theory
- Power-laws and spectrum of a graph
- Internet routing and BGP datasets
- Power-laws analysis
- Spectral analysis
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Various graphs

- Random graphs:
  - nodes and edges are generated by random process
  - Erdős and Rényi model
- Small world graphs:
  - nodes and edges are generated such that most of the nodes are connected by a small number of nodes inbetween
  - Watts and Strogatz model
- Scale free graphs:
  - graphs whose node degree distribution follow power-law
  - rich get richer
  - Barabási and Albert model

Clusters of AS nodes: small world network

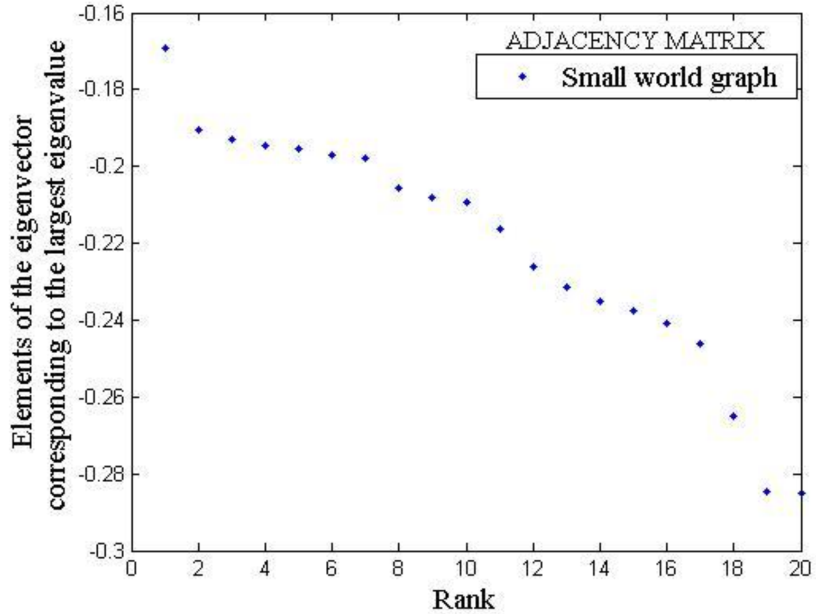

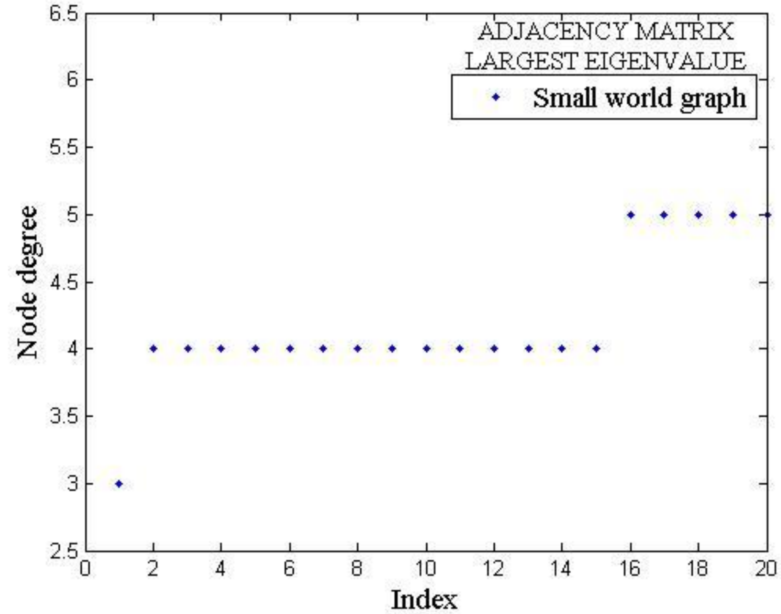

Small world network with 20 nodes:

- nodes having similar degrees are grouped together based on the element values of the eigenvector corresponding to the largest eigenvalue of the adjacency matrix

Clusters of AS nodes: small world network

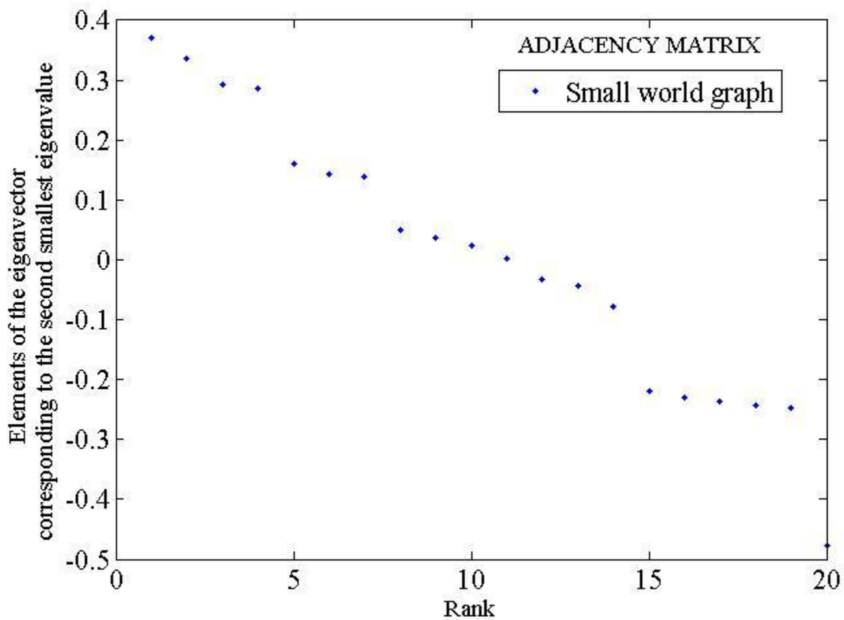

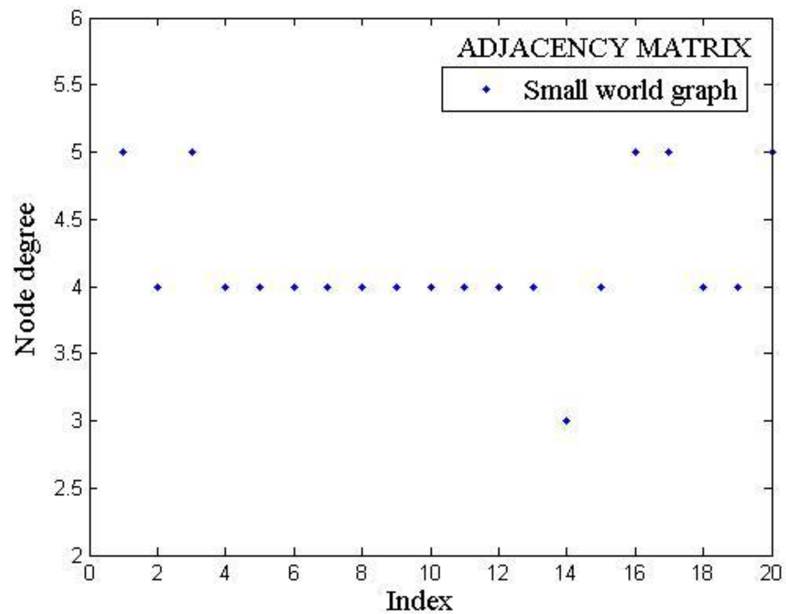

Small world network with 20 nodes:

- nodes having similar degrees are not grouped together based on the element values of the eigenvector corresponding to the second smallest eigenvalue of the adjacency matrix

Eigenvector: the second smallest eigenvalue

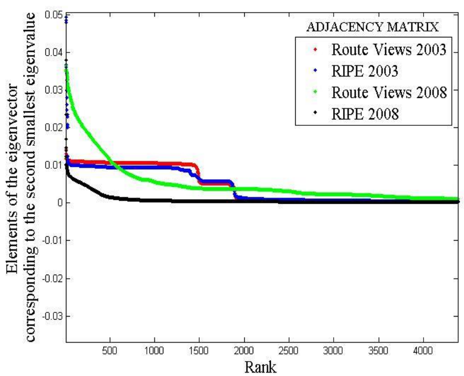

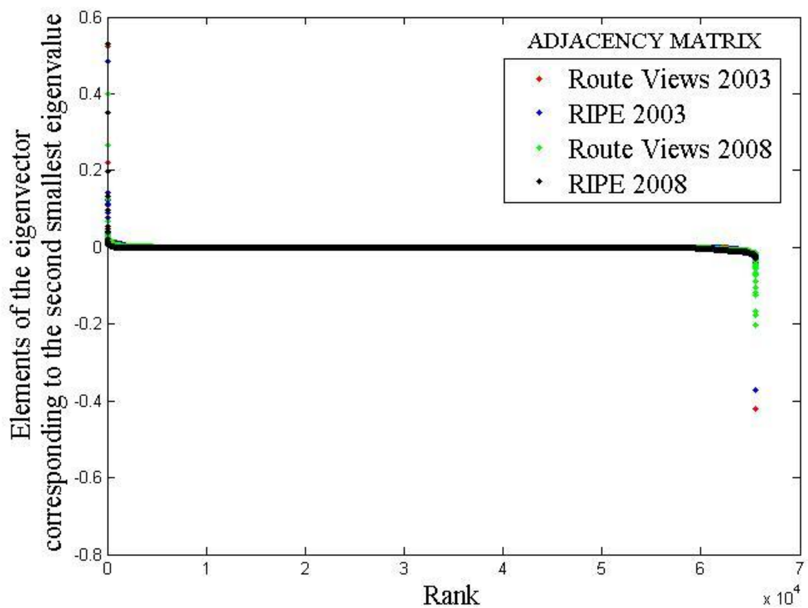

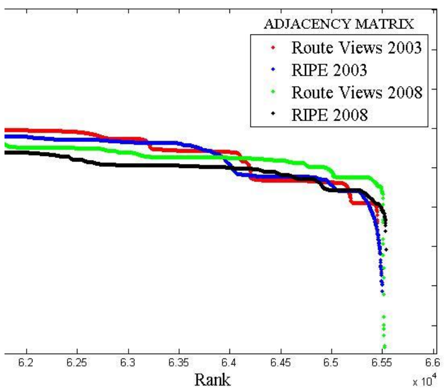

Route Views and RIPE 2003 and 2008 datasets:

- elements of the eigenvectors corresponding to the second smallest eigenvalue of the adjacency matrix

Clusters: Route Views

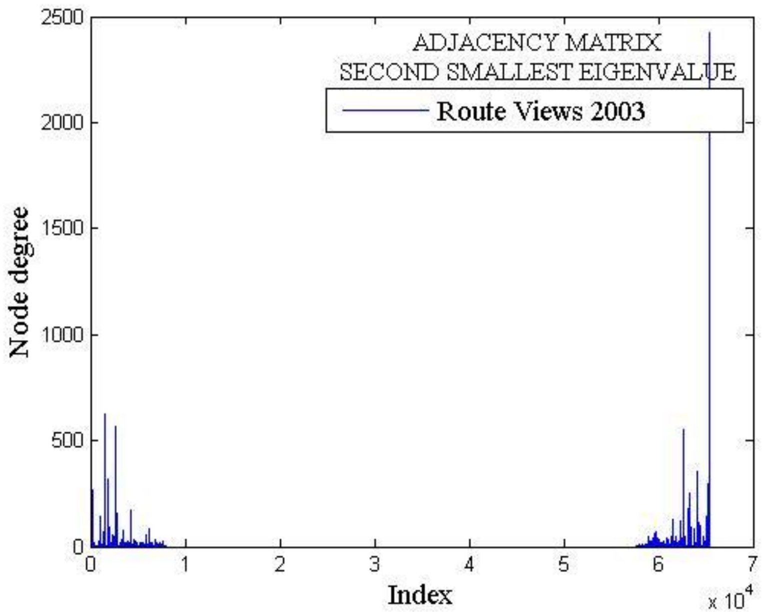

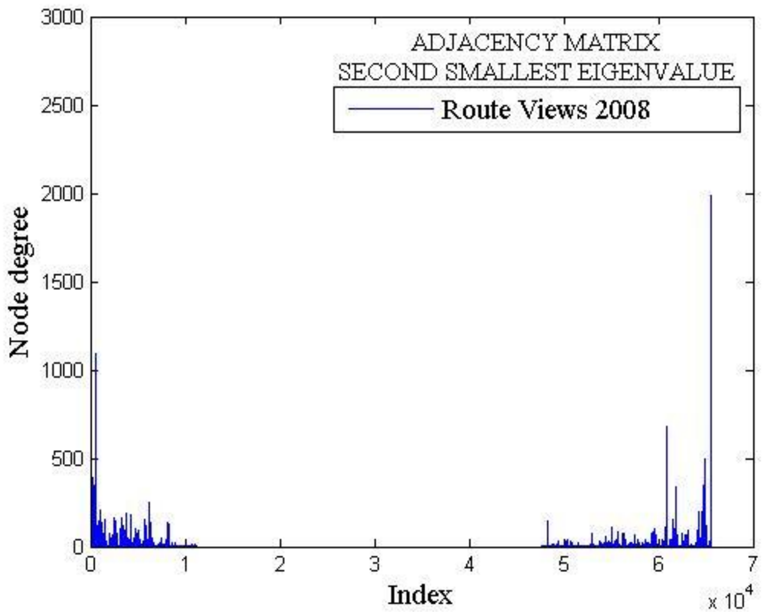

Route Views 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the second smallest eigenvalue of the adjacency matrix divide nodes into two separate clusters of connected nodes

Clusters: RIPE

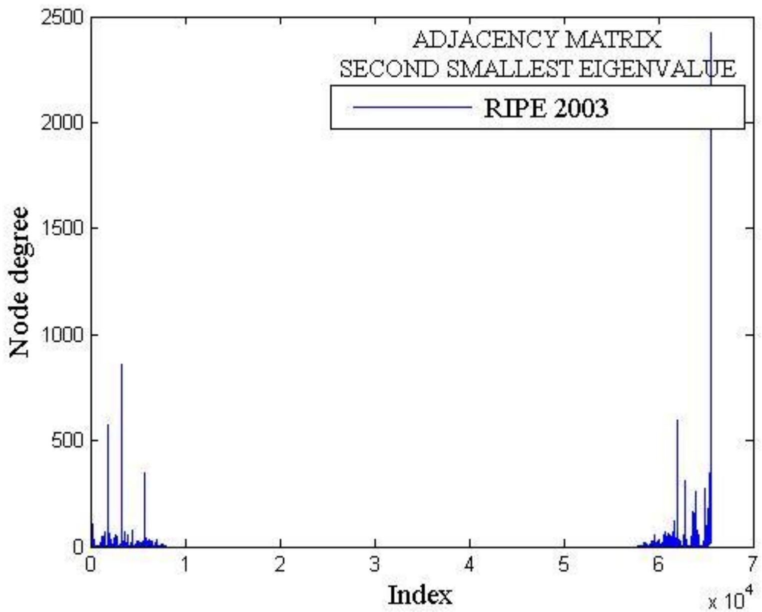

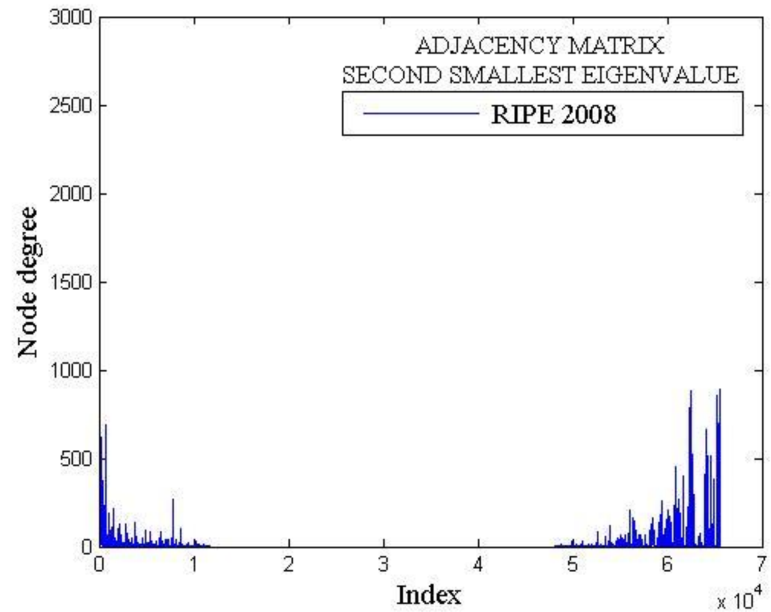

RIPE 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the second smallest eigenvalue of the adjacency matrix divide nodes into two separate clusters of connected nodes

Eigenvector: the largest eigenvalue

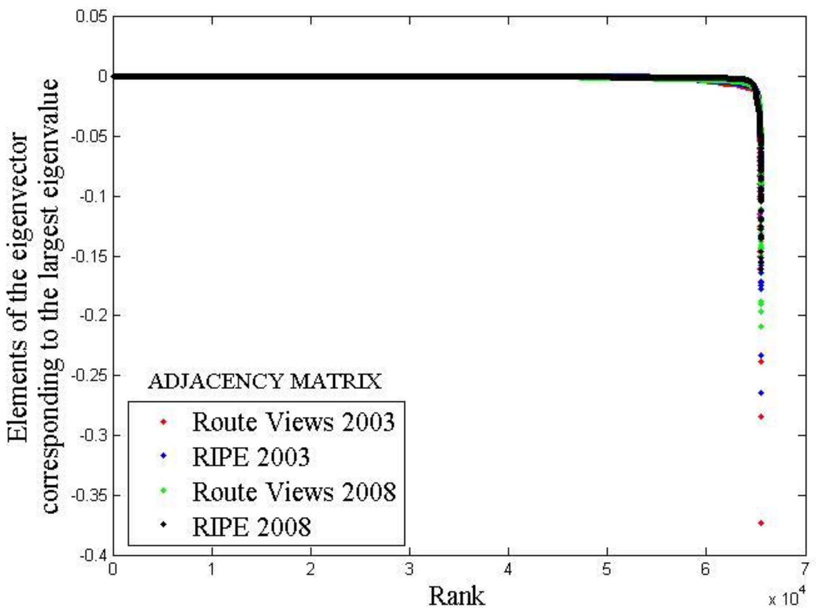

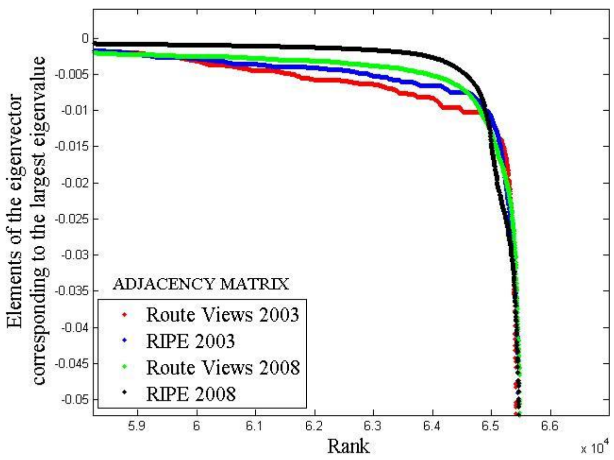

Route Views and RIPE 2003 and 2008 datasets:

- elements of eigenvectors corresponding to the largest eigenvalue of the adjacency matrix

Clusters: Route Views

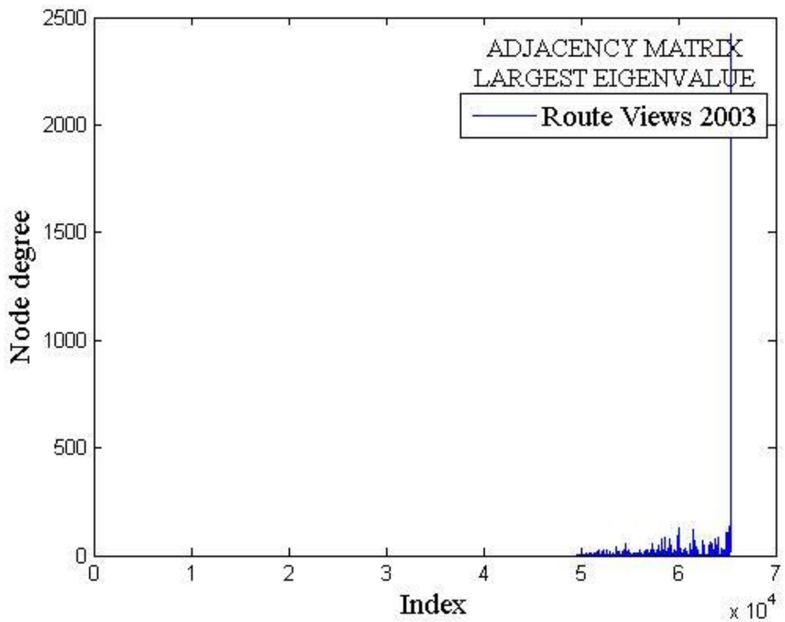

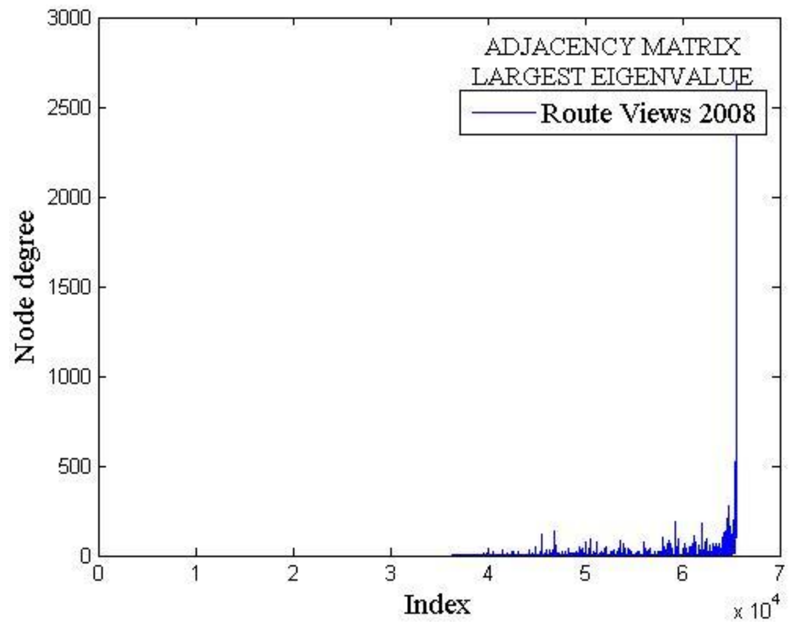

Route Views 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the largest eigenvalue of the adjacency matrix group nodes into a cluster of connected nodes towards the highest end of the rank spectrum

Clusters: RIPE

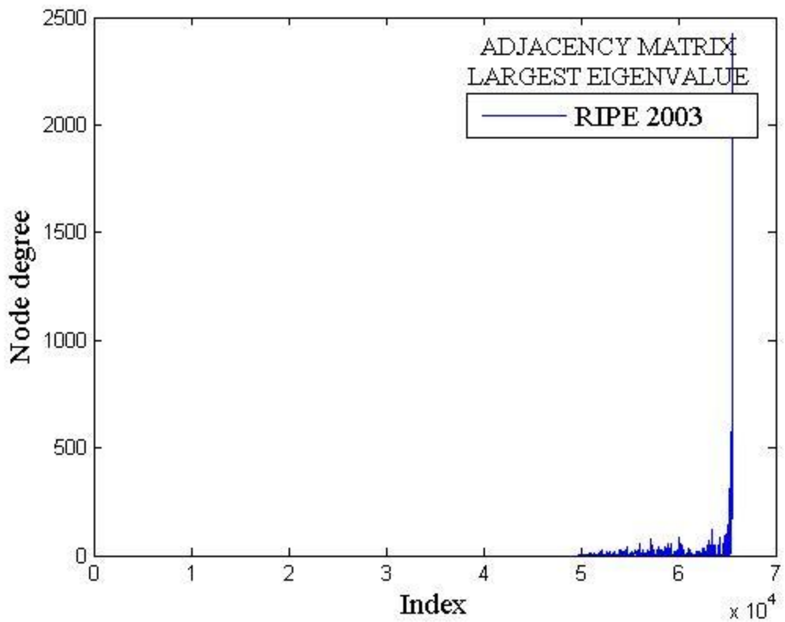

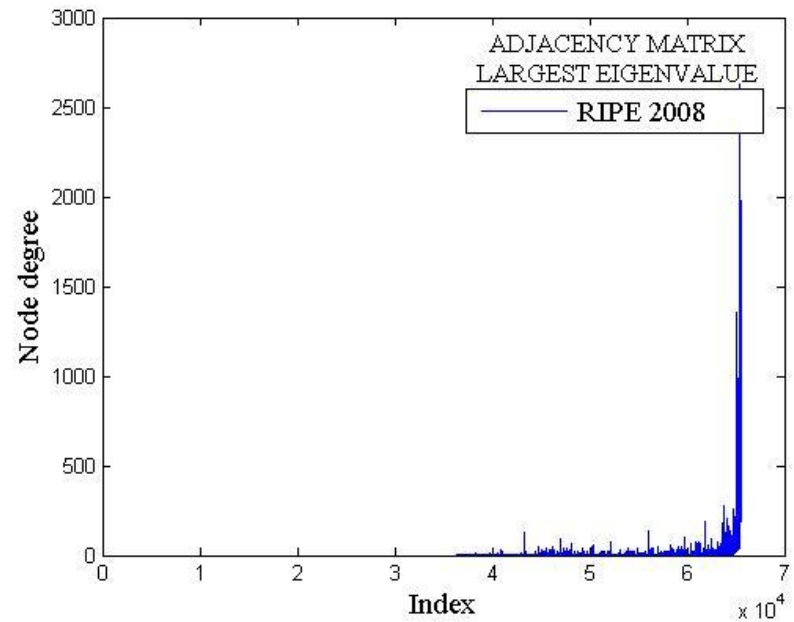

RIPE 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the largest eigenvalue of the adjacency matrix group nodes into a cluster of connected nodes towards the highest end of the rank spectrum

Eigenvector: the second smallest eigenvalue

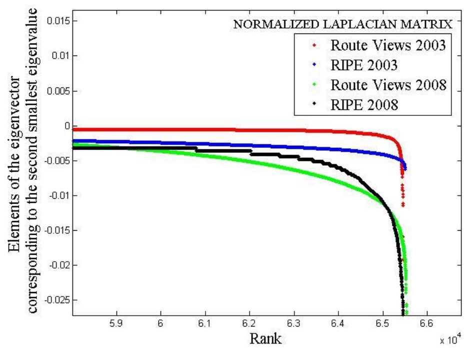

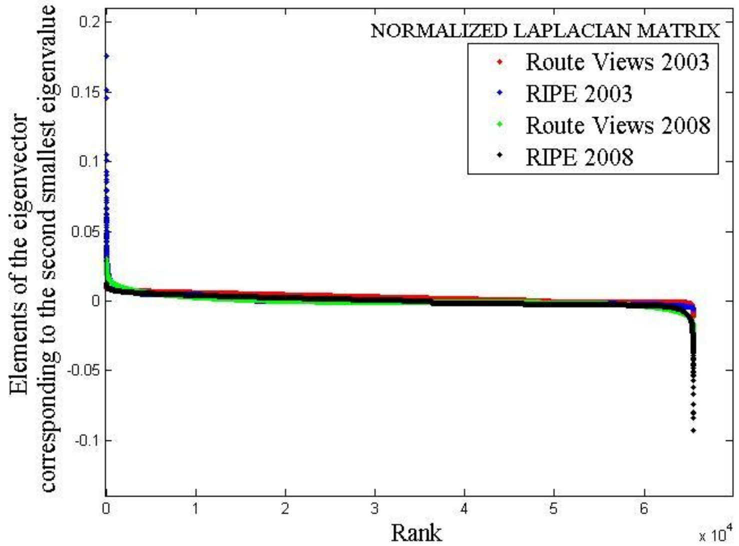

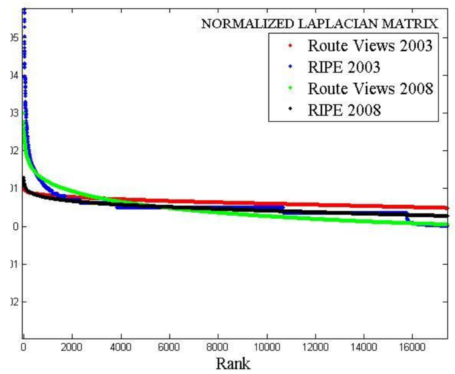

Route Views and RIPE 2003 and 2008 datasets:

- elements of eigenvectors corresponding to the second smallest eigenvalue of the normalized Laplacian matrix

Clusters: Route Views

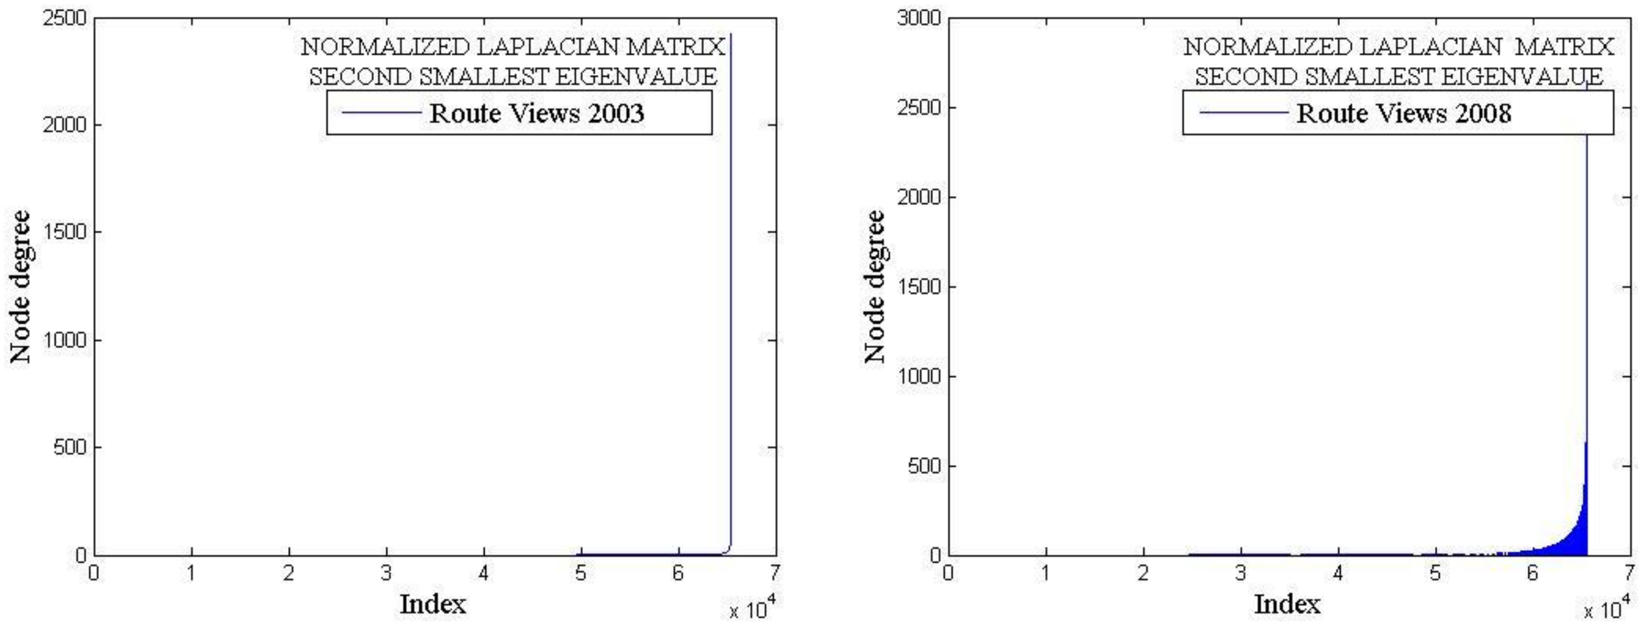

Route Views 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the second smallest eigenvalue of the normalized Laplacian matrix group nodes having similar node degrees

Clusters: RIPE

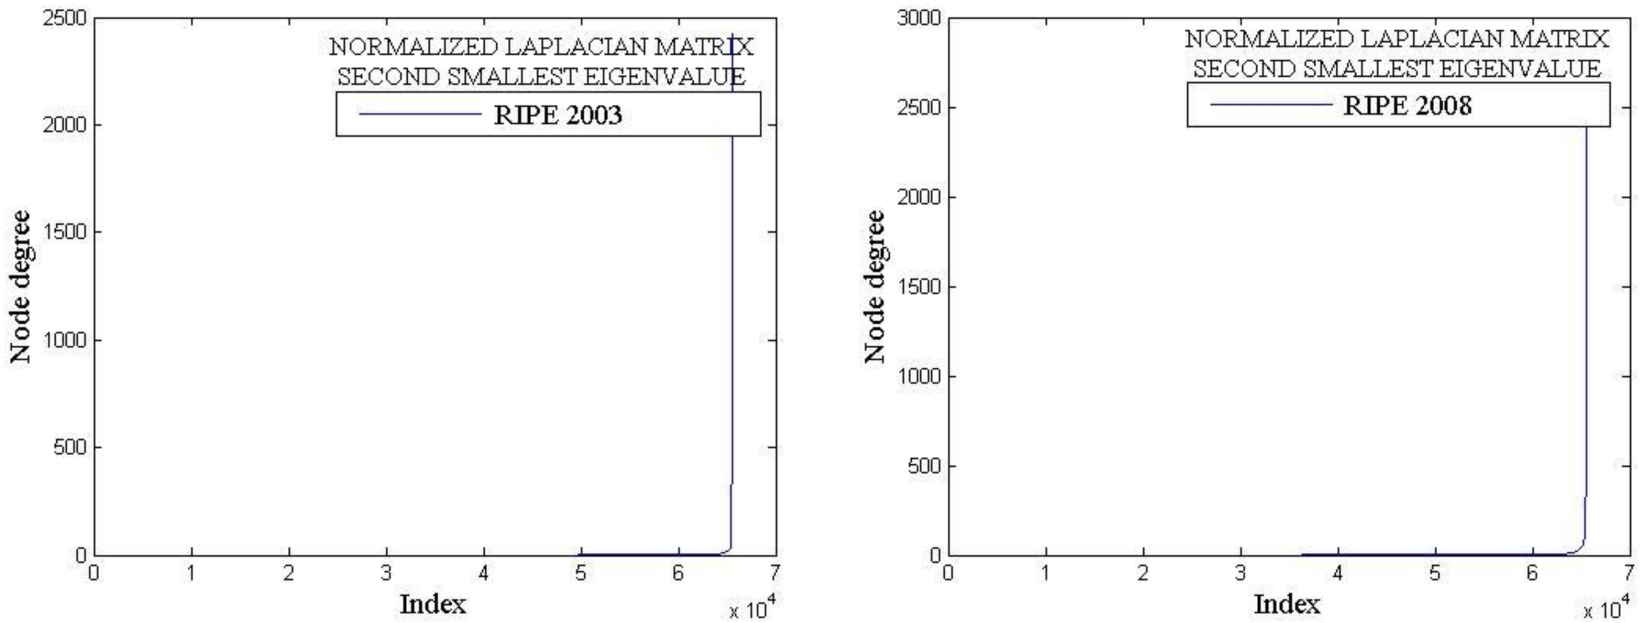

RIPE 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the second smallest eigenvalue of the normalized Laplacian matrix group nodes having similar node degrees

Clusters: RIPE

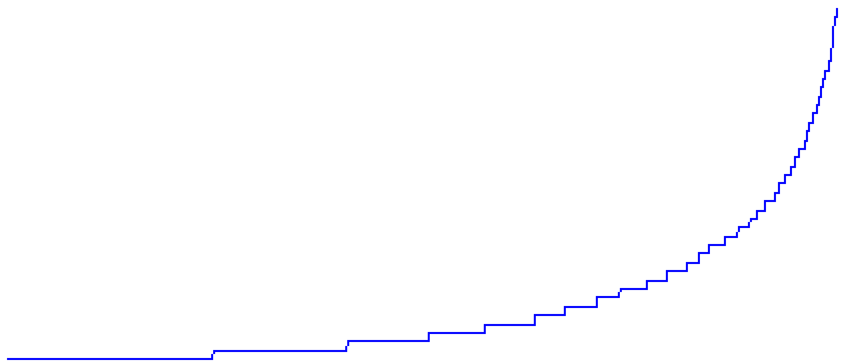

RIPE 2003: zoomed view of node degree vs. rank.

Eigenvector: the largest eigenvalue

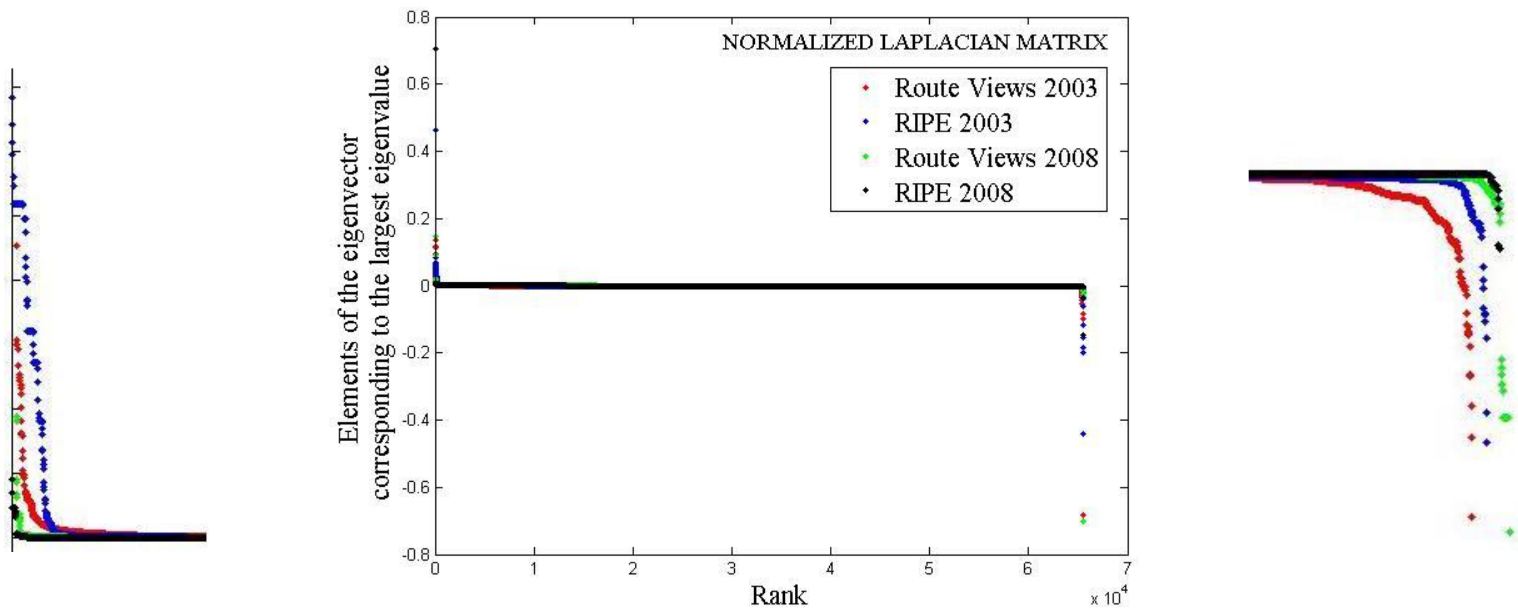

Route Views and RIPE 2003 and 2008 datasets:

- elements of eigenvectors corresponding to the largest eigenvalue of the normalized Laplacian matrix

Clusters: Route Views

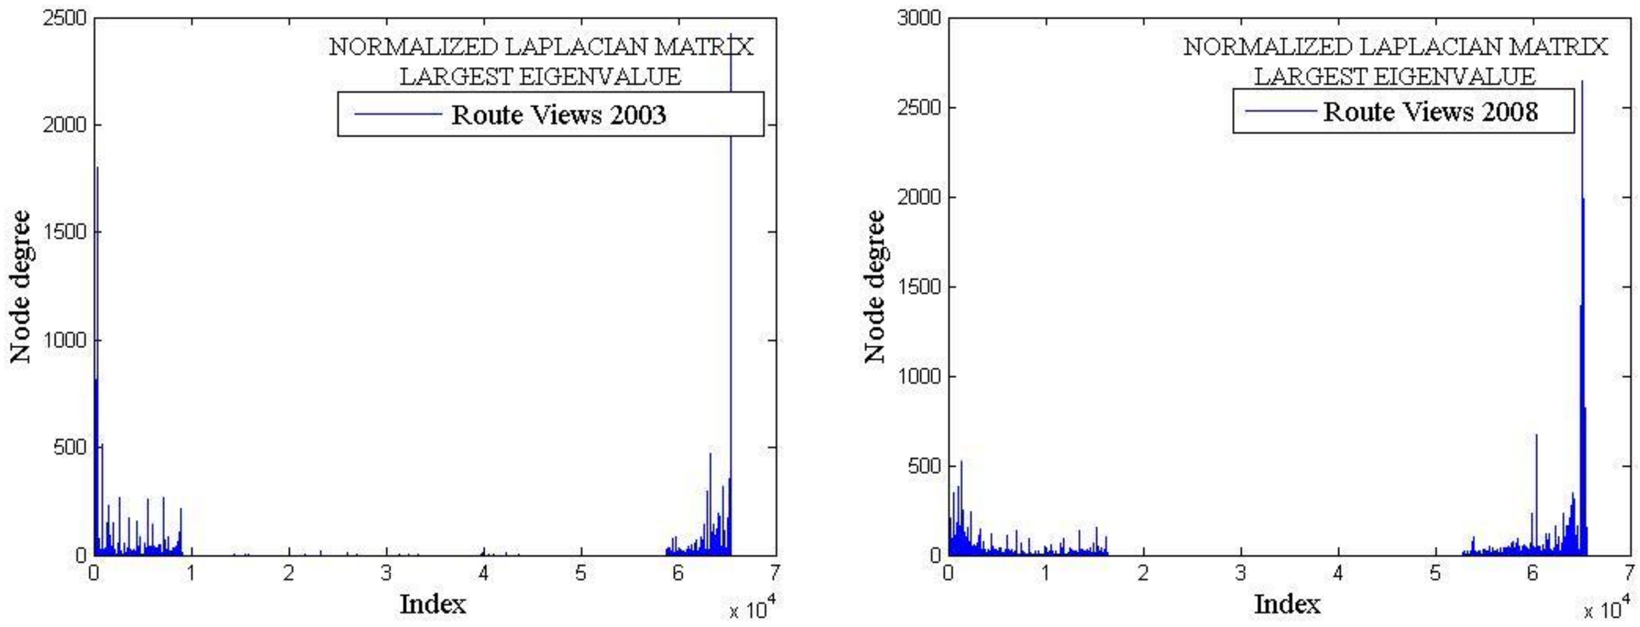

Route Views 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the largest eigenvalue of the normalized Laplacian matrix divide nodes into two clusters of connected nodes

Clusters: RIPE

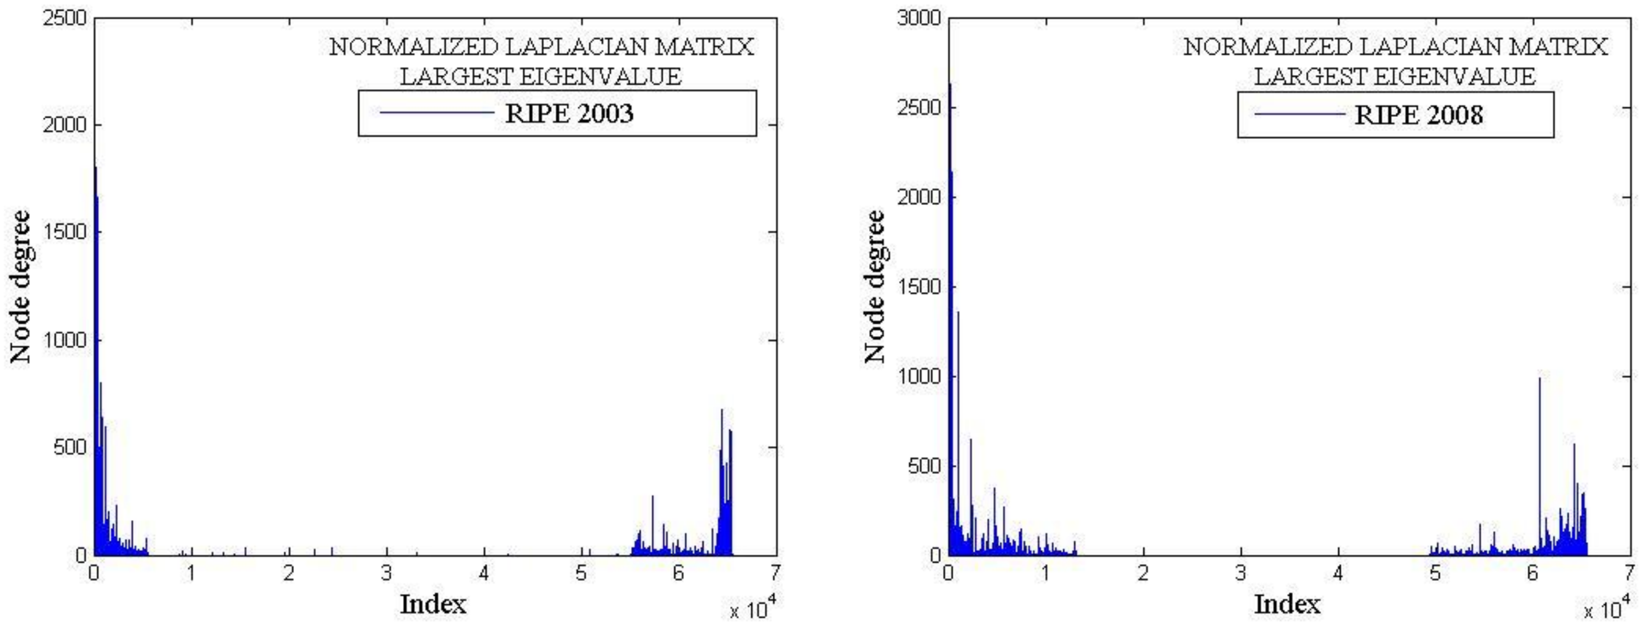

RIPE 2003 and 2008 datasets:

- element values of the eigenvector corresponding to the largest eigenvalue of the normalized Laplacian matrix divide nodes into two clusters of connected nodes

Clusters of AS nodes: summary

- The second smallest eigenvalues of the normalized Laplacian matrix group nodes having similar node degree:
  - group of nodes having larger node degree follows nodes having smaller node degree
- Cluster of nodes based on the elements values of the eigenvector corresponding to the second smallest eigenvalue of the adjacency matrix is similar to the cluster based on the largest eigenvalue of the normalized Laplacian matrix and vice versa.
- Clusters of small world network differ with the Internet graphs.

Roadmap

- Introduction
- Internet topology and BGP datasets
- Power-laws and spectrum of a graph
- Power-laws analysis
- Spectral analysis:
  - connectivity status
  - clusters of ASes
- Conclusions and future work

Conclusions and future work

- Route Views and RIPE datasets reveal similar trends in the development of the Internet topology.
- Power-laws exponents have not significantly changed over the years:
  - indicates they do not capture every property of graph and are only a measure used to characterize the Internet topology
- Spectral analysis reveals new historical trends and notable changes in the connectivity and clustering of AS nodes over the years.
- Element values of the eigenvector corresponding to the second smallest and the largest eigenvalues provide clusters of connected ASes:
  - indicate clusters of connected nodes have changed over time

Conclusions and future work

- Similarity of clusters based on the second smallest eigenvalue of the adjacency matrix to the largest eigenvalue of the normalized Laplacian matrix, and vice versa indicate:
- Clusters based on the second smallest eigenvalues of the normalized Laplacian matrix:
  - group nodes having similar node degree
  - group of nodes having smaller node degree are followed by nodes having larger node degree
  - indicates second smallest eigenvalues of the normalized Laplacian matrix provide node degree information

Conclusions and future work

Future work:

- analysis of the spectral properties of the Internet topology based on matrices such as Laplacian and signless Laplacian
- analysis of the effect of L= D–A in the spectral properties of the adjacency and the normalized Laplacian matrices
- analysis of any significant effect of the observed clusters on:
  - modeling of the Internet topology
  - performance evaluation of protocols and new algorithms

References

- H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, “Towards capturing representative AS-level Internet topologies,” *The International Journal of Computer and Telecommunications Networking*, vol. 44, no. 6, pp. 735-755, Apr. 2004.
- J. Chen and Lj. Trajković, “Analysis of Internet topology data,” in *Proceedings of IEEE International Symposium on Circuits and Systems*, Vancouver, BC, Canada, May 2004, vol. IV, pp. 629–632.
- Q. Chen, H, Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, “The origin of power laws in Internet topologies revisited,” in *Proceedings of INFOCOM*, New York, NY, USA, Apr. 2002, pp. 608–617.C. Gkantsidis, M. Mihail, and E. Zegura, “Spectral analysis of Internet topologies,” in *Proceedings of IEEE INFOCOM*, San Francisco, CA, Mar. 2003, pp. 364–374.
- A. Dhamdhere and C. Dovrolis, "Ten years in the evolution of the Internet ecosystem*,"* in *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, Vouliagmeni, Greece, Oct. 2008, pp.183-196.
- I. Farkas, I. Dernyi, A. Barabasi, and T. Vicsek, “Spectra of “real-world” graphs: Beyond the semicircle law”, *Physical Review E*, vol. 64, no. 2, 026704, pp. 1-12, July 2001.
- L. Gao, “On inferring autonomous system relationships in the Internet,” *IEEE/ACM Trans. Networking*, vol. 9, no. 6, pp. 733–745, Dec. 2001.

References

- S. Jaiswal, A. Rosenberg, and D. Towsley, "Comparing the structure of power-law graphs and the Internet AS graph," in *Proceedings of 12th IEEE International Conference on Network Protocols*, Washington, DC, Oct. 2004, pp. 294-303.
- Y. Li, J. Cui, D. Maggiroini, and M. Faloutsos, "Characterizing and modelling clustering features in AS-level Internet topology," in *Proceedings of IEEE INFOCOM 2008,* Phoenix, Arizona, USA, Apr. 2008, pp. 271-275.
- D. Magoni and J. J. Pansiot, "Analysis of the autonomous system network topology," *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 3, July 2001, pp. 26–37.
- P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. claffy, and A. Vahdat, "The Internet AS-level topology: three data sources and one definitive metric," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, Jan. 2006, pp. 17–26.
- A. Medina, I. Matta, and J. Byers, "On the origin of power laws in Internet topologies," *ACM SIGCOMM Computer Communication Review*, vol. 30, no. 2, pp. 18-28, Apr. 2000.
- H. Haddadi, G. Iannaccone, A. Moore, R. Mortier, M. Rio, "Network topologies: inference, modelling and generation", *IEEE Communications Surveys and Tutorials*, vol. 10, no. 2, pp. 48-67, July 2008.
- C. Gkantsidis, M. Mihail, and E. Zegura, "Spectral analysis of Internet topologies," in *Proceedings of IEEE INFOCOM*, San Francisco, CA, Mar. 2003, pp. 364–374.

References

- M. Najiminaini, L. Subedi, and Lj. Trajković, "Analysis of Internet topologies: A historical view," in Proceedings of IEEE Int. Symp Circuit and Systems, Taipei, Taiwan, May 2009, pp. 1697-1700.
- R. Oliveira, B. Zhang, and L. Zhang, "Observing the evolution of Internet AS topology," ACM SIGCOMM Computer Communication Review, vol. 37, no. 4, pp. 313-324, Jan. 2007.
- A. Pothen, H. Simon, and K. P. Liou, "Partitioning sparse matrices with eigenvalues of graphs," *SIAM Journal of Matrix Analysis*, vol. 11, no. 3, pp. 430-452, July 1990.
- G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos, "Power-laws and the AS-level Internet topology," *IEEE/ACM Trans. Networking*, vol. 11, no. 4, pp. 514-524, Aug. 2003.
- L. Subedi and Lj. Trajković, "Spectral analysis of the Internet topologies," in Proceedings of IEEE Int. Symp Circuit and Systems, Paris, France, May 2010.
- D. Vukadinovic, P. Huang, and T. Erlebach, "On the spectrum and structure of Internet topology graphs," in Proceedings of Innovative Internet Computing Systems, Kühlungsborn, Germany, June 2002, vol. 2346, pp. 83–96.
- D. J. Watts and S. H. Strogatz, "Collective dynamics of small world networks," Nature, vol. 393, pp. 440–442, June 1998.

Acknowledgments

- Committee members:
  - Prof. Mirza Faisal Beg
  - Prof. Ljiljana Trajković
  - Prof. Stephen Hardy
  - Prof. Qianping Gu
- Prof. Kohshi Okumura
- Colleagues:
  - Mohamadreza Najiminaini
  - Shaun Nguyen
  - Eman Elghoneimy
  - Reza Qarehbaghi
  - Sukhchandan Lally

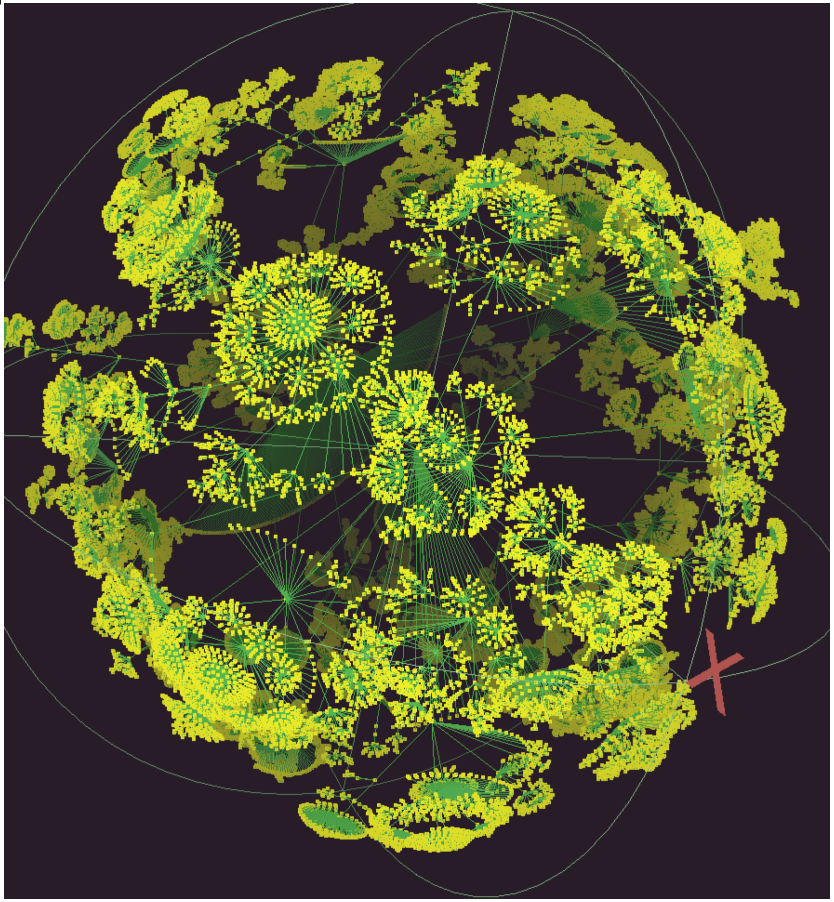

lhr (535,102 nodes and 601,678 links)
http://www.caida.org/tools/visualization/walrus/gallery1/

THANK YOU!