Topological Analysis of the Philippine Internet

Romelo Santos, Author, bit@bitdesigns.net
Rafferty Tan, Author, laffer_077@yahoo.com
William Yu, Adviser, wyu@ateneo.edu
Pierre Tagle, PhD, Adviser, ptagle@ateneo.edu

Department of Information Systems and Computer Science
Ateneo de Manila University
632-4266001

ABSTRACT

This project aims to develop a general methodology of analysis of Internet topology to be immediately applied to the context of the Philippine Internet. This methodology will be applied in the form of an application tool that will provide visualizations of the findings. This system will then allow easier interpretations of analyses of the state of the Philippine Internet, as well as other general subsets of the Internet. This project uses graph theory and linear algebra to build the methods of analysis. It will easily provide relevant information in a meaningful manner. The analysis made will be a significant step in enabling ISPs to clearly visualize and manage the complexity of their networks more effectively and efficiently.

Keywords

Network topology, Internet mapping

1. INTRODUCTION

1.1. Motivation

The Internet is a worldwide network of computers of immense complexity. It was not designed by any single entity but emerged from the *ad hoc* interactions of many entities and continues to grow with no specific global structure.

Developing tools to visualize the Internet has become important to ISPs to gather data that will assist in making strategic decisions concerning network infrastructures, from issues of bandwidth to nodal redundancy. Decisions such as where to connect new ASes or where new connections are needed can be made with the certainty that the outcome of these decisions will be the improvement of the network.

In a network where local Internet connectivity is only beginning to develop and mature, such as that of the Philippines, it is important to determine the strengths and weaknesses of the current network infrastructure. Once a clear idea of the state of the network is acquired, local ISPs can begin to effectively meet the needs of the fast-growing number of Filipinos connecting to the Internet.

1.2. Background

Many projects have already set out to map the Internet but there have not yet been much effort to focus on the Philippine Internet. Ledesma [11] started the effort to study the local Internet topology by developing a system that gathers Border Gateway Protocol information. He used various visualization tools to display the data he had gathered (Figures 1 and 2).

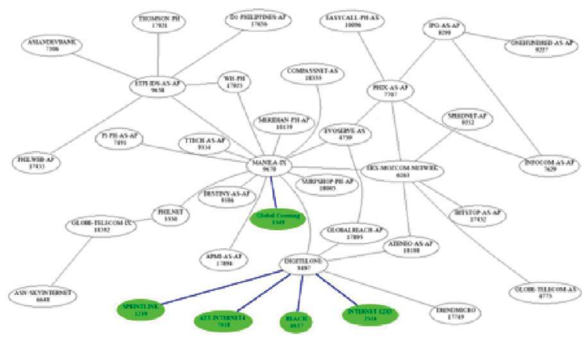

Fig. 1. Philippine Internet Peering (generated by GraphViz)

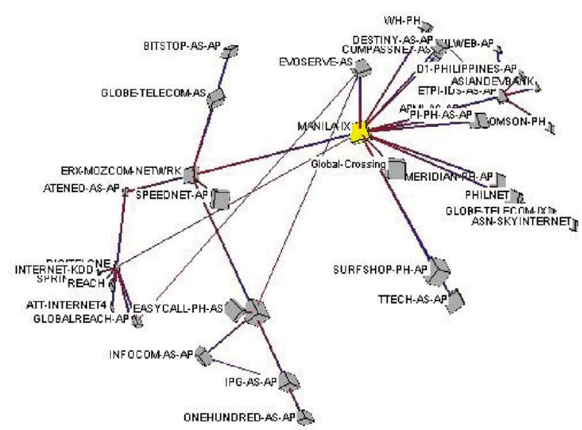

Fig. 2. Philippine Internet Peering (generated by Hypviewer)

Clarito and Viñas [5] further developed the system by storing the BGP information into a Transaction-Time Database.

Current Internet-mapping tools, including the system developed by Ledesma, are able to gather BGP data and/or display a visualization of the information. These tools, however, are not able to analyze the data into a meaningful interpretation for a directly intuitive network analysis.

This study aims to build on Ledesma's project together with data gathered in the Transaction-Time Database and develop a tool to more effectively display the information through visualizations of the graph of the topology and graphical data in a meaningful manner, as well as provide a meaningful interpretation of the data.

One such interpretation is that of inferring Autonomous System relationships. This interpretation aims to classify relationships between pairs of interconnected ASes into customer-provider, peer-to-peer, and sibling-to-sibling relationships according to the contractual agreements between the administrative domains they belong to. Gao [7], was the first to study these relationships forming them into annotated graphs (see Fig. 3.) and enumerated its various applications especially for ISPs. Presently, there is still no publicly available information about inter-AS relationships and ISPs do not register their relationships to Internet registries.

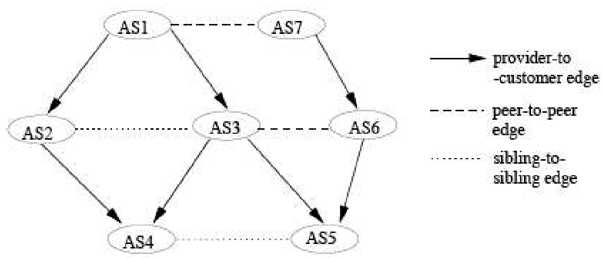

Fig. 3. Annotated AS graph

This project will also produce a study of the current state of the Philippine Internet based on the various metrics gathered, using an analysis method called spectral analysis. Chemists, physicists, and electrical engineers have used this method in the past, but its use in the analysis of network topologies (particularly Internet topologies) is relatively new.

1.3. Objectives and Scope

The earlier studies done concerning Internet topology, AS relationships and spectral analysis were done in the global context. Global properties of the Internet such as behavior of connectivity in different clusters were revealed.

This project aims to derive actual implications of the Internet topology to the operation of the network in terms of performance, connectivity and vulnerability based on the graph and its spectrum. This study will focus on applying spectral analysis on the Philippine subset of the Internet and determine its current state based on certain metrics.

2. RELATED WORKS

2.1. Border Gateway Protocol

The Border Gateway Protocol is an inter-Autonomous System routing protocol. The primary function of a BGP system is to exchange network reachability information with the other BGP systems. This network reachability information includes information on the list of Autonomous Systems (ASes) that reachability information traverses. This information is sufficient to construct a graph of the AS connectivity from which routing loops may be pruned and some policy decisions at the AS level may be enforced. [14]

2.2. Spectral Analysis

Spectral methods (eigendecomposition) have been a part of graph theory for over a century. Network researchers have used spectral methods either implicitly or explicitly since the late 1960's, when computers became generally accessible in most universities.

The eigenvalues of a network are intimately connected to important topological features such as maximum distance across the network (diameter), presence of cohesive clusters, long paths and bottlenecks, and how random the network is. The associated eigenvectors can be used as a natural coordinate system for graph visualization; they also provide methods for discovering clusters and other local features. When combined with other, easily obtained network statistics (e.g., node degree), they can be used to describe a variety of network properties, such as degree of robustness (i.e., tolerance to removal of selected nodes or links), and other structural properties, and the relationship of these properties to node or link attributes in large, complex, multivariate networks. [16]

Seary and Richards in [15] describe spectral methods for partitioning sparse graphs based on three different spectra of a graph. The standard spectrum of a graph is generated from its adjacency matrix. The normal spectrum is generated from the normalized standard adjacency matrix, in which the sum of each column is equal to one. The Laplacian spectrum is based on the Laplacian of the adjacency matrix. The study reveals that the Laplacian and the normal spectra of graphs are much more useful for partitioning graphs.

Fiedler [6] further proposes that the second eigenvalue of Laplacian spectra contains the most relevant information concerning graph properties.

Mohar [12] surveys known results about the spectrum of the Laplacian matrix of graphs with special emphasis on the second smallest Laplacian eigenvalue, popularly known as λ_2 and its relation to numerous graph invariants.

[12] presents upper bounds and lower bounds for the graph diameter, edge and vertex connectivity, and mean distance. All of these properties are important in analyzing communication networks in particular. Edge connectivity is the minimum number of edges that can be removed to disconnect or split the graph. Vertex connectivity is the minimum number of vertices that can be removed to disconnect or split the graph. These two properties correspond to network robustness. Graph diameter is the longest shortest path between all pairs of vertices in the graph. This can describe the worst-case scenario of the transfer of data over the network. Mean distance is the average lengths of all paths between all pairs of vertices in the network. This property is another good indication of nodal connectivity.

Vukadinovic, et. al. [17], first applied spectral analysis to the Internet topology to investigate the topology's properties. Actual data was compared to data gathered from simulated topologies. Their study revealed certain properties that led to a new meaningful structural classification of AS graphs.

Gkantsidis, et. al. [8], continued to study the spectrum of the Internet topology and revealed even more relationships of the eigenvalues/eigenvectors to the actual structure and operation of the Internet.

Gao [7], presented heuristic algorithms that infer the AS graph from BGP routing tables. He formally presented the routing policies implied by AS relationships and derive routing entry patterns as the result of routing policies. He then inferred the AS relationships based on the heuristic that the size of an AS is typically proportional to its degree in the AS graph. The heuristic algorithms he developed then classified the AS pair into having a provider-to-customer, peer-to-peer, or sibling-sibling relationship.

2.1. Otter Visualization Tool

Otter is a CAIDA tool for visualizing arbitrary network data that can be expressed as a set of nodes, links or paths. We developed Otter to handle visualization tasks for a wide variety of Internet data we deal with in our research, including data sets on topology, workload, performance, and routing. We have used Otter to visualize: multicast and unicast topology databases, core BGP routing tables, reachability and delay measurements, SNMP data, and web site directory structures. Otter's strength is in its *data independence*: it can handle any formatted data set consisting of links and nodes. [4]

3. METHODOLOGY

3.1. Project Description

The primary deliverable of this project is the **topology analyzer**, an application that takes BGP data and outputs a data file to be read and displayed by the Otter tool. This tool will do the retrieval of the information from the database, process the data, and finally write the data into a file.

The analysis will result in the classification of the nodes in the graph according to various criteria and in the upper and lower bounds of relevant graph invariants. The main criteria for the node classification will be the function of the node in the graph according to its graph location [8, 9] and its relationships with other nodes [7]. Using spectral analysis, the analysis will compare the numerical values of the graph invariants [1]. These statistics of the AS graph of the Philippines will be compared with the statistics of the AS graphs of other Asian countries.

3.2. Architecture and Design

The application will be subdivided into three components (Fig. 4): the data retrieval component, the data analysis component, and the data output component. The project will be object-oriented and each component will be encapsulated into separate objects.

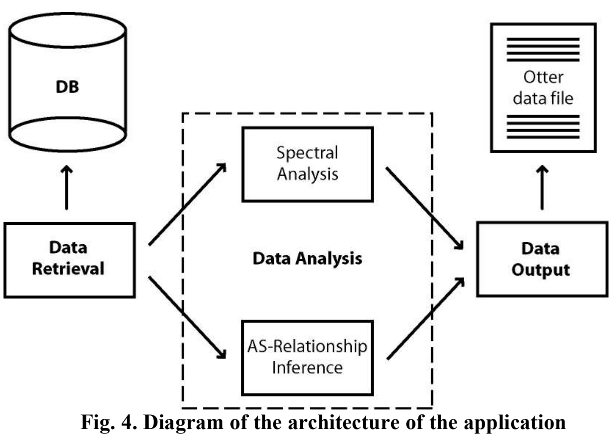

Fig. 4. Diagram of the architecture of the application

3.2.1. Data Retrieval Component

The data retrieval component's role is to connect to the database, retrieve the graph data, and store it in memory to be processed by the Data Analysis Component.

3.2.2. Data Analysis Component

The data analysis component's role is to take the data and perform the spectral analysis and incorporate heuristics that that determine the classifications of the AS pairs.

3.2.3. Data Output Component

The data output component's role is to take the the data from the two former components and output it to a file. This file will be readable by CAIDA's Otter visualization tool.

3.3 Implementation Details

3.3.1. General Implementation Details

The application is implemented using Java Standard Edition, 1.4.1.02. It connects to a PostgreSQL database system.

Our application has been compiled and tested on the following systems:

- Machine 1
  - Windows XP Professional, SP2
  - AMD AthlonXP 1600+, 1.4 GHz
  - 384 MB, PC2100 DDR-RAM
  - Realtek 10/100 (mbps) Ethernet card
- Machine 2
  - MandrakeLinux 10, Kernel version 2.6.3
  - Intel Celeron 1.7 GHz
  - 512 MB, PC2100 DDR-RAM
  - Realtek 10/100 (mbps) Ethernet card

The PostgreSQL database system was hosted on the following system:

- CERSA Web Server
  - Operating system: RedHat Linux 7.2, Kernel version 2.4.2
  - Processor: Intel Pentium III 600 MHz
  - Memory: 256 MB, PC133 SD-RAM
  - Dual IBM 10/100 (mbps) Ethernet cards

The JAMA Java Matrix Package[1], developed by the National Institute of Standards and Technology, is used for the eigendecomposition in the spectral analysis.

3.3.2. Data Retrieval Implementation

Data is retrieved by using the JDBC package. A series of queries with sub-queries needed in this module were made into virtual tables(available in PostgreSQL) to increase readability and ease of use. The implementation gets the connection to the database, queries the database and then parses the data into an otter-readable format. This is done by getting the total number of links and connections. The connections themselves are then created to provide a basis in counting the number of degrees each connected ASes have and a basis also for creating the adjacency matrix both needed in the next module.

3.3.3. Data Analysis Implementation

This module is further subdivided into two submodules, (1) spectral analysis, and (2) AS-relationship inference.

[1]Available at http://math.nist.gov/javanumerics/jama/

3.3.3.1. Spectral Analysis

The spectral analyzer accepts a two-dimensional array that represents the standard adjacency matrix. It creates the normalized and Laplacian adjacency matrices of the standard adjacency matrix. It then creates the eigendecomposition of the matrices and takes the eigenvalues and corresponding eigenvectors. The important eigenvalue of the Laplacian spectrum is the second eigenvalue, λ_2, which is also the second-smallest. These two eigenvalues are used for analysis.

The spectral analysis submodule uses the eigenvalue from the Laplacian spectrum with the inequalities described in [12] to describe edge and vertex connectivity, diameter, and mean distance. We were able to calculate upper and lower bounds for each of these properties by using these inequalities. The bounds concerning vertex connectivity, *$\nu(G)$*, and edge connectivity, *$\eta(G)$*, of a graph of order *n* are:

- $$\lambda_2(G) \leq \nu(G) \leq \eta(G)$$
- $$\lambda_2(G) \geq 2\eta(G)(1-\cos\frac{\pi}{n})$$

The upper and lower bounds of the diameter, *diam(G)*, of the graph of order *n*, are defined as:

- $$diam(G) \geq \frac{4}{n\lambda_2(G)}$$
- $$diam(G) \leq 2\lceil \frac{\Delta+\lambda_2(G)}{4\lambda_2(G)} \ln(n-1) \rceil$$

The upper and lower bounds of the mean distance, ρ, of the *n*-order graph with maximal degree, Δ, are defined as:

- $$\rho \leq \frac{n}{n-1} \lceil \frac{\Delta+\lambda_2(G)}{4\lambda_2(G)} \ln(n-1) \rceil$$
- $$(n-1)\rho \geq \frac{2}{\lambda_2(G)} + \frac{n-2}{2}$$

The lower bound of the number of edges for the approximation of the solution to the max-cut problem, *MC(G)*, is defined as:

- $$MV(G) \leq \frac{n\lambda_n(G)}{4}$$

λ_2 can be used to approximate a solution to the minimum-cut problem [16]. By grouping the nodes according to the sign of its weights, we divide the graph into two subgroups. The links between the nodes of the first group and the nodes of the second group are the links that are to be cut in the min-cut problem.

These results are added to the main data vector and are passed to the data output module.

3.3.3.2. AS-Relationship Inference

The heuristic algorithms for inferring AS relationships are based on the intuition that a provider typically has a larger size than its customer and the size of an AS is typically proportional to its degree in the AS graph. The uphill or downhill top provider of an AS path should be the AS that has the highest degree of all ASes in its maximal uphill or downhill path. Let the *top provider* of an as_path be the AS that has a higher degree between the uphill and downhill top provider. Therefore, the top provider of an AS path is the AS that has the highest degree among the ASes in each AS path.

Knowing the top provider, we can infer that consecutive AS pairs (u1, u2) on the left of the top provider have a customer-to-provider relationship (u2 provides transit services from u1), and consecutive pairs on the right of the top provider have a provider-to-customer relationship (u1 provides transit services to u2). An As pair has a sibling-to-sibling relationship if the pair provides transit services for each other. Further details of this specific is shown in [7].

3.3.4. Data Output Implementation

Results are output in the Otter data file format [5]. Each line of the Otter data file is generated as one string based on the data from the data analysis classes. These strings are then added to a vector that is passed to the file writer class. The application outputs each string in the vector into an Otter data file (.ODF) that can be read and displayed by CAIDA's Otter Visualization tool.

4. RESULTS

4.1. Results of the Spectral Analysis

The upper and lower bounds of three graph properties are taken – edge connectivity, diameter, and mean distance. Edge connectivity is the smallest number edges that can be removed from the graph to disconnect the graph, or in other words, to separate the graph into two disconnected subgraphs. The diameter of a graph is the longest shortest path between all pairs of nodes. The diameter dictates the length of the best worst case for hops between all pairs of nodes. The mean distance is the average path length of all distinct paths between all pairs of nodes. This property is useful for determining the degree of connectivity between nodes.

The edge connectivity of the graph is directly related to the robustness of the graph, therefore a higher lower bound will mean that the network is less vulnerable to becoming disconnected by the removal of links.

In an ideal Internet topology – a complete graph – edge connectivity is *n – 1*, where *n* is the size of the graph. Therefore, a higher value for edge connectivity is more desirable and would imply higher network robustness.

	Philippines	*Japan*	*Singapore*
Lower bound of edge connectivity	0.52052	0.17000	0.17176
Upper bound of edge connectivity	399.01940	1864.41176	479.14000
Lower bound of diameter	0.08838	0.07152	0.14041
Upper bound of diameter	744.47778	20037.17163	4484.46714
Lower bound of mean distance	0.53307	0.53282	0.56760
Upper bound of mean distance	14.78485	4.71969	11.20921

Table 1. Output of the Spectral Analysis. Edge connectivity describes the minimum number of edges that can be removed to disconnect the graph. Diameter is the longest shortest path between all pairs of nodes in the graph. Mean distance is the average path length of all unique paths between all pairs of nodes.

Table 1 indicates that the Philippine AS network has a higher lower bound of edge connectivity, but its upper bound is not that much different from that of Singapore. Japan and Singapore have very similar edge connectivity lower bounds, but Japan has an extremely high edge connectivity upper bound. Considering the number of ASes in the Philippines, compared with the number of ASes of Singapore and Japan, it is not entirely surprising that vulnerability would not differ much. With almost twice as many ASes as Singapore, Japan has a considerably more robust AS network compared to Singapore. Singapore has roughly twice as many ASes as the Philippines, but the robustness of the two AS networks are very similar.

The diameter and mean distance gives a good indication of the number of hops data needs to take to go from one AS to another. This invariant is also connected to vulnerability. When data needs more hops between nodes, there is a greater need for robustness - vulnerability becomes a greater issue.

In terms of diameter lower bounds, once again, the Philippine and Japanese AS networks are quite similar. The Singapore AS network has a diameter lower bound that is twice as high as that of the Philippines and of Japan. Due to the very high numbers of ASes in Japan and Singapore, it is expected that the upper bounds of the diameters of the two networks to be high.

The lower bounds of the mean distances of the networks of the three countries are very similar, so no conclusive observations can be gotten. The upper bound for the mean distance, though, is much more interesting. Japan has the lowest upper bound for distance, despite having the highest number of ASes among the three observed countries. Singapore has the second highest upper bound for mean distance, but is much closer to the upper bound for mean distance of the Philippines. Despite having the smallest AS network in terms of number of nodes, the Philippines has the worst upper bound of mean distance.

By the spectral analysis done, according to the three graph invariants studied, the Philippines is not far behind a well-developed country such as Singapore, but it is also made obvious that the Philippine Internet still is not as robust as that of Japan.

4.2. Results of the AS-Relationship Inference

Philippine Internet peering is still in a limited state according to [13], but our study reveals that the average peering of provider nodes in the Philippines is more similar to the superior network of Japan than that of Singapore.

	Philippines	*Japan*	*Singapore*
Total Number of ASes	87	329	166
Mean Provider Degrees	3.42857	3.45455	2.88

Table 2. Summary of the Output of the AS-Relationship Inference. Mean provider degrees describes the peering of AS providers in the network.

On average, the providers of the Philippines and of Japan have almost one more link to other providers than the providers of Singapore. There is the possibility that Singapore may only lag because of a small number of major transit providers.

The results, though, can be skewed due to the fact that there are less ASes in the Philippines, and thus the ratios will be distorted as most Philippines providers might not appear in the analysis at all.

5. CONCLUSION

The study has resulted in a method of analysis that can be used as a basis of comparison among Internet topologies of subdivisions of the Internet such as the subdivisions among countries.

In terms of robustness, the study did not statistically find any major differences among the three countries, but the minute differences support the assumption that the Philippine Internet, in this aspect, is quit inferior to the Internet of Japan and Singapore.

In terms of connectivity, or the number of hops between nodes, once again, there were no significant differences between the three countries. The minute differences, though, again point to the assumption that the Philippine Internet is closer to the superior Japanese Internet than to the Singaporean Internet.

The AS-relationship inference reveals that the Philippine Internet core of providers is as connected and robust as that the Japanese Internet. Once again, the Internet of Singapore lags in this aspect.

The study has shown that although, quantitatively, the Philippine Internet topology is inferior to those of Japan and Singapore, qualitatively, the Philippines has a decent topological make-up. Of the three countries compared, Japan, unsurprisingly, has the largest and most well-connected Internet map. The Philippines, on the other hand, has shown to be slightly better than Singapore and, in some aspects, even on par with Japan.

6. RECOMMENDATIONS

Although our study has developed an effective method of comparing Internet topologies, the nature of the test data used is not completely reliable. [3] indicates that BGP-based maps only show the best routes and do not necessarily reflect the complete Internet map. In addition, not all ASes in the Philippines use

BGP, and are therefore not included in the analyses. The results may be more refined and accurate if better bases of Internet topology can be used.

The relatively small number of recognizable ASes in a graph, such as that of the Philippines, may also skew the results. Properties such as graph diameter and mean distance will tend to bias towards smaller graphs. Future studies may put the size of the graph into account for a more accurate analysis.

Our developed application for our method of analysis was designed specifically to be used in conjunction with the data warehouse designed by Clarito and Viñas [5]. Our method may be adapted to be used in other implementations so it can be integrated to other Internet topology projects.

7. ACKNOWLEDGEMENTS

We are grateful for the guidance and support of the Department of Information Systems and Computer Science of the Ateneo de Manila University. We would like to thank Dr. Ian Jun L. Garces of the AdMU Mathematics Department for his support regarding the spectra of graphs. We would also like to thank everyone who has made their works regarding BGP data, Internet topologies, and algebraic graph theory available to the research community.

8. REFERENCES

[1] Barabasi, A. and Albert, R. Statistical mechanics of complex networks, *Reviews of Modern Physics, Volume 74.* University of Notre Dame, Notre Dame, Indiana (2002).

[2] Biggs, N. *Algebraic Graph Theory*. Cambridge University Press, London (1974). pp. 1-48.

[3] Broido, A. and Claffy, K. *Internet Topology: Connectivity of IP Graphs*. University of California, San Diego (2001).

[4] Claffy, K., Nemeth, E., and Huffaker, B. *Otter: A general-purpose network visualization tool*. Internet (June 2002).

[5] Clarito, I. and Viñas, R. *Proposed Data Warehousing for Internet Mapping-Related Information Using a Temporal Database*. Undergraduate Computer Science Mini-Thesis. Ateneo de Manila University (2004).

[6] Fiedler, M., Algebraic Connectivity of Graphs, *Czechoslovak Mathematical Journal, 23 (98).* Praha (1973).

[7] Gao, L. *On Inferring Autonomous System Relationships in the Internet*. University of Massachusetts (2000).

[8] Gkantsidis, C., Mihail, M., and Zegura, E. *Spectral Analysis of Internet Topologies.* Research Paper. Georgia Institute of Technology, Atlanta (2002).

[9] Inmon, W. *What is a Data Warehouse?*. Internet. http://www.bilinmon.com/libtrary/whiteprs/earlywp/ttdw.pdf

[10]Kollios, G. *Temporal Databases*. Internet. htp://www.cs.bu.edu/faculty/gkollios/ada02/LectNotes/lect10.ppt

[11]Ledesma, G. Mapping the Philippine Internet using the Border Gateway Protocol. Undergraduate Computer Science Mini-Thesis. Ateneo de Manila University (2003).

[12]Mohar, B. *The Laplacian Spectrum of Graphs*. University of Ljubljana, Yugoslavia (1991).

[13]Paraz, M. and Yu, W. Philippine IX Information, *Asia Pacific Internet Association Newsletter No. 8*. Asia Pacific Internet Association (2002).

[14]Rekhter, Y. and t. Li, Ed. *A Border Gateway Protocol (BGP-4)*. March 1995. Internet. http://www.faqs.org/rfcs/rfc1771.html

[15]Seary, A. and Richards, W. *Partitioning Networks by Eigenvectors*. Simon Fraser University, Burnaby British Columbia, Canada (1995).

[16]Seary, A. and Richards, W. *Spectral methods for analyzing and visualizing networks: an introduction. Research* paper. Simon Fraser University. Burnaby City, Canada (2000).

[17]Vukadinovic, D., Huang, P., and Erlebach, T. *A Spectral Analysis of the Internet Topology.* Technical report. ETH Zurich (2001).